

Essays on the Evaluation of Climate Policies with Quasi-experimental Methods

Piero Basaglia

born in Venice (Italy)

A thesis submitted to Universität Hamburg,
Faculty of Business, Economics, and Social Sciences

for the degree of Doctor rerum politicarum
(Pursuant to the doctoral degree regulations of 18 January 2017)

under the supervision of:
Prof. Dr. Moritz Drupp
Prof. Dr. Andreas Lange

Universität Hamburg
11th October, 2023

Examination commission

Chair: Prof. Dr. Grischa Perino

First assessor: Prof. Dr. Moritz Drupp

Second assessor: Prof. Dr. Andreas Lange

Third assessor: Prof. Dr. Thomas Sterner

Date of the defense (“Disputation”): 08.01.2024

Declaration

I hereby declare that I, Piero Basaglia, have not received any commercial consultation on my doctoral thesis. This thesis has not been accepted as part of any previous doctoral procedure or graded as insufficient.

Hamburg, 08.01.2024

Place, date



Doctoral candidate's signature

Affidavit:

I, Piero Basaglia, hereby declare under oath that I wrote the dissertation titled:

“Essays on the Evaluation of Climate Policies with Quasi-experimental Methods”

myself and in case of cooperation with other researchers pursuant to the enclosed statements in accordance with Section 6 subsection 3 of the Doctoral Degree Regulations of the Faculty of Business, Economics and Social Sciences dated 18 January 2017. I have used no aids other than those indicated.

Hamburg, 08.01.2024

Place, date



Doctoral candidate's signature

Personal declaration for dissertation by chapter (Department of Social Sciences)

Concept/planning: Formulation of the fundamental academic problem, based on previously unanswered theoretical questions, including a summary of general inquiries that can be answered on the basis of analysis or experiments/investigation. Planning of the experiments/analyses and formulation of the methodological approach, including decision about methods and independent methodological development.

Implementation: Degree of integration in the concrete investigations or analyses.

Creation of the manuscript: Presentation, interpretation, and discussion of sought-for findings in the form of an academic article.

Calculation of personal contribution proceeds according to a point system of 1–100 percent.

Candidate has made an independent contribution of 100 percent to at least one of the articles (cf. Chapter 5).

For a second article (cf. Chapter 2), the candidate's personal contribution breaks down thus:

concept/planning	50%
implementation	75%
creation of manuscript	75%

For a third article (cf. Chapter 3), the candidate's personal contribution breaks down thus:

concept/planning	33%
implementation	66%
creation of manuscript	33%

For any further articles (cf. Chapter 4), the candidate's personal contribution breaks down thus:

concept/planning	25%
implementation	50%
creation of manuscript	40%

My respective co-authors have agreed to the above percentages for my own contributions.

Hamburg, 08.01.2024

Place, date



Doctoral candidate's signature

Abstract

English Version *The four chapters of this dissertation combine a set of empirical analyses drawing from several quasi-experimental methods to foster a better understanding of the impacts of climate policies within the realms of environmental, health, innovation, and financial economics. The first chapter provides a comprehensive evaluation of the world's largest implicit carbon tax reform. Leveraging multiple synthetic control methods and a synthetic difference-in-differences estimator, we compare carbon and air pollutant emissions of the actual and counterfactual German transport sector following the 1999 eco-tax reform and find average reductions in external damages of around 80 billion Euros. We further show that environmental taxation induced low-carbon innovation and document much stronger demand responses to environmental tax increases than to market price movements, primarily driven by increased tax salience in newspapers. Our results highlight the key roles of salience and fuel substitution in mediating the effectiveness of carbon taxes to deliver climate and health benefits. The second chapter examines the causal impact of compensation payments for indirect carbon costs embodied in electricity prices for energy-intensive sectors. We use confidential UK plant-level data to exploit firm-level inclusion criteria in both a difference-in-differences with inverse probability weighting and regression discontinuity framework. Our findings suggest that compensated firms increased production and electricity use relative to uncompensated firms, with no significant effect on energy intensity. While compensation lowers leakage risk, it also implies large forgone opportunity costs of public funds and increased mitigation costs of meeting national emission targets. The third chapter analyzes the relationship between climate policy uncertainty and firms' and investors' behavior. Leveraging newspaper data, we develop a set of new indices of climate policy uncertainty, covering the United States with monthly-level variation dating back to 1990, and analyze their impacts on firm-level outcomes such as stock returns, share price volatility, investments in research and development, and employment for all publicly-listed firms in the country. We employ an identification strategy that differentiates sectors by*

their relative exposure to climate policy changes and show that climate policy uncertainty tends to affect all these outcomes considerably, particularly in carbon-intensive sectors, and often more so than existing indices of economic policy uncertainty. The direction of the effect may, however, be driven by the underlying direction of the uncertainty, which we measure explicitly. Finally, the fourth chapter provides novel quasi-experimental evidence on the effects of air pollutants on defensive expenditures and economic productivity to retrieve spatially resolved estimates of the willingness to pay for air quality improvements. To address endogeneity concerns, atmospheric temperature inversions are exploited as a source of quasi-random variation in the spatial concentration of $PM_{2.5}$. Using administrative data from England, I find that a plausibly exogenous $1 \mu g/m^3$ $PM_{2.5}$ shock significantly affects pharmaceutical expenditures and GVA per capita, partly through increased work absenteeism. Leveraging a counterfactual reduction of $1 \mu g/m^3$ of $PM_{2.5}$, I show that health benefits are more pronounced among the elderly and progressively distributed across income levels, while productivity gains are regressive and concentrated in urban areas. These findings imply that incorporating the spatial heterogeneity of pollution-reduction benefits into policy design could enhance the efficiency of environmental regulations and contribute to tackling health inequalities linked to pollution exposure.

Keywords: air pollution; carbon tax; climate policy; compensation schemes; electricity consumption; environmental inequalities; firm decision-making; investments; low-carbon transition; uncertainty.

JEL codes: D22; D81; D83; D84; G10; G18; G32; H23; I12; J14; O32; Q48; Q51; Q52; Q53; Q56; Q58.

Deutsche Übersetzung *Die vier Kapitel dieser Dissertation kombinieren eine Reihe von empirischen Analysen, die sich auf mehrere quasi-experimentelle Methoden stützen, um ein besseres Verständnis der Auswirkungen der Klimapolitik in den Bereichen Umwelt, Gesundheit, Innovation und Finanzwirtschaft zu fördern. Das erste Kapitel bietet eine umfassende Bewertung der weltweit größten impliziten CO₂-Steuerreform. Unter Verwendung mehrerer synthetischer Kontrollmethoden und eines synthetischen Differenz-von-Differenzen-Schätzers vergleichen wir die Kohlenstoff-*

und Luftschadstoffemissionen des tatsächlichen und des kontrafaktischen deutschen Verkehrssektors nach der Ökosteuerreform von 1999 und stellen eine durchschnittliche Verringerung der externen Schäden von rund 80 Milliarden Euro fest. Darüber hinaus zeigen wir, dass die Umweltbesteuerung zu kohlenstoffarmen Innovationen geführt hat, und dokumentieren eine wesentlich stärkere Reaktion der Nachfrage auf Umweltsteuererhöhungen als auf Marktpreisbewegungen, die in erster Linie auf eine erhöhte Bekanntheit der Steuer in den Zeitungen zurückzuführen ist. Unsere Ergebnisse verdeutlichen die Schlüsselrolle der Steuer-Salienz und der Kraftstoffsubstitution bei der Frage, wie wirksam Kohlenstoffsteuern sind, um Klima- und Gesundheitsvorteile zu erzielen. Das zweite Kapitel untersucht die kausalen Auswirkungen von Ausgleichszahlungen für indirekte Kohlenstoffkosten, die in den Strompreisen für energieintensive Sektoren enthalten sind. Wir verwenden vertrauliche Daten aus Großbritannien auf Werksebene, um Einschlusskriterien auf Unternehmensebene sowohl in einem Differenz-von-Differenzen-Rahmen mit inverser Wahrscheinlichkeitsgewichtung als auch in einer Regressions-Diskontinuitätsanalyse auszunutzen. Unsere Ergebnisse deuten darauf hin, dass die entschädigten Unternehmen ihre Produktion und ihren Stromverbrauch im Vergleich zu den nicht entschädigten Unternehmen erhöht haben, ohne dass sich dies signifikant auf die Energieintensität ausgewirkt hätte. Die Entschädigung senkt zwar das Leckagerisiko, bedeutet aber auch hohe entgangene Opportunitätskosten für öffentliche Mittel und erhöhte Minderungskosten für die Erfüllung der nationalen Emissionsziele. Im dritten Kapitel wird die Beziehung zwischen klimapolitischer Unsicherheit und dem Verhalten von Unternehmen und Investoren analysiert. Auf der Grundlage von Zeitungsdaten entwickeln wir eine Reihe neuer Indizes für klimapolitische Unsicherheit, die die Vereinigten Staaten mit monatlichen Schwankungen seit 1990 abdecken, und analysieren ihre Auswirkungen auf Ergebnisse der Unternehmensebene wie Aktienrenditen, Aktienkursvolatilität, Investitionen in Forschung und Entwicklung und Beschäftigung für alle börsennotierten Unternehmen des Landes. Wir verwenden eine Identifizierungsstrategie, die Sektoren nach ihrer relativen Exposition gegenüber klimapolitischen Veränderungen unterscheidet, und zeigen, dass klimapolitische Unsicherheit tendenziell alle diese Ergebnisse erheblich beeinflusst, insbesondere in kohlenstoffintensiven Sektoren, und zwar oft stärker als bestehende Indizes für wirtschaftspolitische Unsicherheit. Die Richtung des

Effekts kann jedoch durch die zugrundeliegende Richtung der Unsicherheit bestimmt werden, die wir explizit messen. Das vierte Kapitel schließlich liefert neuartige quasi-experimentelle Erkenntnisse über die Auswirkungen von Luftschadstoffen auf die Verteidigungsausgaben und die wirtschaftliche Produktivität, um räumlich aufgelöste Schätzungen der Zahlungsbereitschaft für Verbesserungen der Luftqualität zu erhalten. Um Bedenken hinsichtlich der Endogenität auszuräumen, werden atmosphärische Temperaturinversionen als Quelle exogener dynamischer Variationen in der räumlichen Konzentration von Schadstoffen genutzt. Unter Verwendung von Verwaltungsdaten aus England stelle ich fest, dass ein plausibel exogener $1 \mu\text{g}/\text{m}^3$ $\text{PM}_{2.5}$ -Schock die Arzneimittelausgaben und die Bruttowertschöpfung pro Kopf signifikant beeinflusst, zum Teil durch erhöhte Fehlzeiten am Arbeitsplatz. Anhand einer kontrafaktischen Verringerung der $\text{PM}_{2.5}$ um $1 \mu\text{g}/\text{m}^3$ zeige ich, dass die gesundheitlichen Vorteile bei älteren Menschen ausgeprägter sind und sich progressiv über die Einkommensschichten verteilen, während die Produktivitätsgewinne regressiv sind und sich auf städtische Gebiete konzentrieren. Diese Ergebnisse deuten darauf hin, dass die Einbeziehung der räumlichen Heterogenität des Nutzens der Schadstoffreduzierung in die Politikgestaltung die Effizienz von Umweltvorschriften erhöhen und dazu beitragen könnte, gesundheitliche Ungleichheiten im Zusammenhang mit der Schadstoffbelastung zu bekämpfen.

Schlüsselwörter: Luftverschmutzung; Kohlenstoffsteuer; Klimapolitik; Ausgleichsregelungen; Stromverbrauch; ökologische Ungleichheiten; Unternehmensentscheidungen; Investitionen; kohlenstoffarmer Übergang; Unsicherheit.

JEL-Codes: D22; D81; D83; D84; G10; G18; G32; H23; I12; J14; O32; Q48; Q51; Q52; Q53; Q56; Q58.

Acknowledgements

I would like to express my gratitude to my supervisors, Prof Dr Moritz Drupp and Prof Dr Andreas Lange, for their precious guidance in this journey.

Thanks to all my co-authors, Dr. Christopher Hedemann, Dr. Miriam Ferrer González, Dr. Caterina Picuno, Dr. Laura Serra, Rebeca Biancardi Aleu, Dr. Lena Beppi, Lorenzo Romero, Juliane Koch, Dr. Francesco Furini, Dr. Guido Cioni, Michael Turner, Felix Schaumann, Jonas Grunau, Simon Disque, Björn Bos, Aljoscha Minnich, Dr. Shushanik Margaryan, Dr. Silvia Beghelli, Dr Francesca Diluiso and the other Ph.D. colleagues for their daily support, the laughs and tears that truly motivated these years of research. You have been incredibly supportive throughout these years and I will never say thank you enough.

A special thanks go to Dr. Matteo Puglini, my pandemic buddy with whom I spent months locked down at home after having just moved to Germany and navigated through some very uncertain and difficult times.

Another special mention to my Hamburg neighbors, Camilla and Luciana, and my cousin Chiara who made me feel at home in Germany from day 1.

To my parents, Michela and Marco, for always being by my side, ready to motivate me when I need it most.

To my brothers, Paolo and Nanni, and my friends, Laura and Elena, who always knew how to distract me when I needed to disconnect from research and could not do it on my own.

To Beatrice who always pushed for me to pursue a Ph.D. and who has listened to my research ideas for 10 years now.

Finally, to my partner Andrea, who supports me unconditionally in everything I do and who has been by my side until the very last line written in this thesis.

Last but not least, I would like to thank all my friends and the rest of my family, I would not be here without my emotional support network and this thesis is the result of your continued support.

With gratitude, I acknowledge that the four chapters presented in this dissertation have greatly benefited from valuable insights and feedback received at the following conferences and workshops: the 10th Economics of Low-Carbon Markets Workshop, the 28th European Association of Environmental and Resource Economists Conference, the 6th LSE/Imperial/King's Workshop in Environmental Economics, the 2022 Conference of Swiss Economists Abroad, the 2022 Banque de France Conference on Frontiers of Climate and Nature in Macroeconomics and Finance, the 27th European Association of Environmental and Resource Economists Conference, the 13th Toulouse Conference on the Economics of Energy and Climate, the 2022 Conference on Sustainable Resource Use and Economic Dynamics, the 2022 CESifo Area Conference on Energy and Climate Economics, the 96th Western Economic Association International Conference, the 26th European Association of Environmental and Resource Economists Conference, the 2021 Summer Conference of the Association of Environmental and Resource Economists, the 25th European Association of Environmental and Resource Economists Conference, the 2020 Pre-conference workshop of the European Association of Environmental and Resource Economists, the 13th International Workshop on Empirical Methods in Energy Economics, the 5th FSR Annual Conference on the Economic Assessment of European Climate Policies, the 24th European Association of Environmental and Resource Economists Conference, and the 7th IZA Workshop on Environment and Labor Markets.

Contents

Abstract	i
Acknowledgements	v
1 Introduction	1
1.1 Motivation and research focus	1
1.2 This dissertation	8
1.2.1 Chapter 2: Causal effects of fuel taxation and mediating mechanisms for reducing climate and pollution costs	8
1.2.2 Chapter 3: Carbon pricing, compensation, and competitiveness: Lessons from UK manufacturing	9
1.2.3 Chapter 4: Climate policy uncertainty and the behavior of firms and investors	11
1.2.4 Chapter 5: Pollution reduction benefits across space: Quasi-experimental evidence from England	12
1.2.5 Future research	13
2 Causal effects of fuel taxation and mediating mechanisms for reducing climate and pollution costs	17
2.1 Introduction	18
2.2 Methodology	23
2.2.1 Causal Inference Methods	23
2.2.2 Semi-elasticity models	28
2.3 Data	30
2.4 Results from Causal Inference Methods	36

2.4.1	Emissions reductions according to SCMs	36
2.4.2	Emissions reductions according to generalized SCMs	41
2.4.3	Impacts on low-carbon innovation for SCM and GSCM	42
2.4.4	Staggered treatment adoption with an SDID estimator	44
2.5	Results on Fuel and Tax Elasticities	46
2.5.1	Real price semi-elasticities for gasoline and diesel	46
2.5.2	Tax semi-elasticities for gasoline and diesel	47
2.5.3	Emission scenarios and underlying mechanisms	49
2.5.4	Tax salience	53
2.6	Assessing climate and pollution reduction benefits	57
2.7	Conclusion	60
2.A	Background on the eco-tax and data sources	64
2.B	Additional results for the synthetic control methods	69
2.B.1	Placebo in time	72
2.B.2	No sample restrictions	72
2.B.3	Leave-one-out tests	73
2.B.4	Generalized Synthetic Control Method (GSCM)	74
2.C	Elasticities	75
2.C.1	Time-series elasticities with a distributed lag model	78
2.C.2	Fuel substitution due to the eco-tax	79
2.C.3	NO _x emission under different taxation regimes	80
2.C.4	SCMs and the Simulation Approach	81
2.C.5	Impacts of the eco-tax on road casualties	81
2.D	Synthetic difference-in-differences	82
2.E	Salience analysis	86
2.E.1	Search strategies	86
2.E.2	Using information in newspaper articles as an indicator of salience	87

3 Carbon pricing, compensation, and competitiveness: Lessons from UK manufacturing	89
3.1 Introduction	90
3.2 Conceptual framework	96
3.3 Research Design and Data	101
3.3.1 Policy Background	101
3.3.2 Data sources	105
3.3.3 Descriptive statistics	107
3.3.4 Using predicted electricity use to calculate the eligibility criterion	109
3.4 Empirical Strategy	110
3.4.1 Difference-in-differences	110
3.4.2 Fuzzy regression discontinuity design (DiDiD-IV)	114
3.5 Treatment effects of the indirect carbon cost compensation	116
3.5.1 DiD estimates of the average treatment effects	116
3.5.2 Fuzzy RD estimates of the local average treatment effects	121
3.5.3 Comparing the DiD and RD estimates	125
3.6 Discussion on policy implications	126
3.7 Conclusion	129
3.A Research context	132
3.A.1 Electricity prices in the UK and continental Europe	132
3.A.2 Eligible 4-digit industries	133
3.B Additional descriptive material	134
3.B.1 Magnitude of compensation payments	134
3.B.2 Descriptive evidence for outcome variables	135
3.C Predicting electricity consumption	138
3.D Using energy purchases to calculate the running variable	139
3.E Additional results: DiD	140
3.E.1 Complementary results based on the period 2010–2019	140

3.E.2	DiD Results for employment	141
3.E.3	Alternative definition of industry-year fixed effects	142
3.E.4	Alternative sample trimming	143
3.E.5	Alternative proxies for production and electricity intensity	144
3.E.6	Alternative calculations of p-scores	146
3.F	Graphical comparisons of DiD estimates	147
3.G	Dynamic difference-in-differences results	149
3.H	Robustness Checks: Fuzzy RDD	150
3.H.1	Alternative functional specifications	150
3.H.2	RD Results for employment	152
3.H.3	RD Manipulation Test using local polynomial density estimation	152
3.H.4	RD bandwidth selection procedure	153
3.H.5	Alternative bandwidths	155

4 Climate policy uncertainty and the behavior of firms and investors 158

4.1	Introduction	159
4.2	Introducing the CPU index	163
4.2.1	Building the index	163
4.2.2	Validating the CPU index	168
4.2.3	Extending the index	171
4.3	Outcome data and empirical strategy	175
4.3.1	Outcome variables and descriptive statistics	175
4.3.2	Empirical Strategy	178
4.4	Empirical results	180
4.4.1	Firm and investor behavior in response to uncertainty shocks	180
4.4.2	Robustness tests	185
4.4.3	Historical analysis	188
4.4.4	Belief revision	190

4.5	Conclusion	191
4.A	Keyword selection	194
4.A.1	English	194
4.A.2	French	196
4.A.3	German	197
4.A.4	Spanish	198
4.A.5	Italian	198
4.B	Short history of climate policy in the United States	200
4.C	Estimated Outcome Changes	202
4.D	CPU index and other relevant measures	206
4.E	Yearly time series	209
4.F	Additional Descriptive Evidence	210
4.F.1	Employment levels in the estimation sample	210
5	Pollution reduction benefits across space: Quasi-experimental evidence from England	211
5.1	Introduction	212
5.2	Conceptual framework	218
5.3	Data and descriptives	221
5.4	Empirical strategy	230
5.5	Results	233
5.5.1	Nationwide effects on morbidity and productivity	233
5.5.2	Spatial heterogeneity	243
5.6	Policy implications	247
5.7	Conclusion	250
5.A	Descriptive evidence	253
5.A.1	Yearly spatial distribution of inversion events	253
5.A.2	Temporal distribution of inversion episodes	260
5.A.3	Concentrations of PM ₁₀ , NO _x , NO ₂ and C ₆ H ₆ in $\mu\text{g}/\text{m}^3$ across England	261

5.A.4	Prescription data: descriptive statistics	262
5.A.5	Mapping of GP practices across England by postcodes	264
5.A.6	Clinical Commissioning groups in England	265
5.A.7	LAUs in London	266
5.A.8	Pre-existing pollution exposure inequalities along the income distribution	267
5.B	Robustness analyses	268
5.B.1	Effect of air pollution by therapeutic section	268
5.B.2	Additional model specifications by air pollutant	269
5.B.3	Effects on prescriptions items and total costs to the NHS	271
5.B.4	Reduced form (RF) with leads and lags	273
5.B.5	LAUs as alternative unit of observation	273
5.C	Changes in defensive behavior over time	274
5.D	Computing the newspaper-based index	278

Bibliography		278
---------------------	--	------------

List of Tables

2.1	Overview of the specification choices for the SCMs	25
2.2	Effects of the eco-tax with a Generalized Synthetic Control . .	42
2.3	Effects of environmental taxes with a SDID staggered adoption design	45
2.4	Real price semi-elasticities for transport fuels	47
2.5	Eco-tax semi-elasticities for transport fuels	48
2.6	Effects of salience on fuel consumption	56
2.A1	Data Sources	69
2.A2	SCM for CO ₂ : Pre-Treatment Predictor Means for Germany, Base- line Synthetic Germany and the Sample Average	70
2.A3	SCM for PM _{2.5} : Pre-Treatment Predictor Means for Germany, Baseline Synthetic Germany and the Sample Average	70
2.A4	SCM for NO _x : Pre-Treatment Predictor Means for Germany, Base- line Synthetic Germany and the Sample Average	70
2.A5	Effects of the eco-tax with the GSCM (No sample restrictions)	74
2.A6	Gasoline consumption	77
2.A7	Diesel consumption	77
2.A8	Real price elasticities for gasoline after 1991	77
2.A9	Real price elasticities for diesel after 1991	77
2.A10	Fuel substitution: Diesel-to-Gasoline ratio	79
2.A11	Fuel substitution: Share of Diesel	79
2.A12	Effects of environmental taxes with an SDID staggered adoption design (No sample restrictions)	82

2.A13	Effects of environmental taxes with an SDID staggered adoption design (Germany and Sweden)	82
2.A14	Effects of environmental taxes with an SDID staggered adoption design (Additional covariates)	82
2.A15	Effects of salience on gasoline demand (robustness)	88
2.A16	Effects of salience on diesel demand (robustness)	88
3.1	Summary statistics for the period 2005–2011, by compensation status	113
3.2	Average treatment effects of compensation. 2010–2015.	118
3.3	Average treatment effects of compensation. 2010–2015.	118
3.4	ATEs of compensation on sales. 2010–2019.	119
3.5	ATEs of compensation on energy intensity (energy purchases/sales). 2010–2019.	120
3.6	Local average treatment effects of compensation. 2010–2015. Fuzzy RDD.	122
3.7	Costs and benefits	127
3.A1	Eligible industries	133
3.A2	Simple correlation between electricity use and energy purchases. 2005–2011	138
3.A3	Correlation between electricity use and energy purchases, employment, and turnover. 2005–2011	138
3.A4	ATEs of compensation on sales. 2010–2019.	140
3.A5	ATEs of compensation on energy intensity (energy purchases/sales). 2010–2019.	140
3.A6	ATEs of compensation on employment. 2010–2019.	141
3.A7	ATEs of compensation on sales <i>within 2-digit industries</i> . 2010–2019.	142
3.A8	ATEs of compensation on energy intensity (energy purchases/sales) <i>within 2-digit industries</i> . 2010–2019.	142

3.A9	ATEs of compensation on sales <i>with different trimming</i> . 2010–2019.	143
3.A10	ATEs of compensation on energy intensity (energy purchases/sales) <i>with different trimming</i> . 2010–2019.	143
3.A11	ATEs of compensation on QFi variables <i>with different trimming</i> . 2010–2015.	144
3.A12	ATEs of compensation on total output. 2010–2019.	144
3.A13	ATEs of compensation on turnover. 2010–2019.	145
3.A14	ATEs of compensation on energy/output. 2010–2019.	145
3.A15	ATEs of compensation on energy/turnover. 2010–2019.	146
3.A16	ATEs of compensation on sales <i>using different p-scores</i> . 2010–2015.	147
3.A17	LATEs of compensation. Fuzzy RDD controlling for linear distance from the cut-off.	150
3.A18	LATEs of compensation. Fuzzy RDD controlling for quadratic distance from the cut-off.	151
3.A19	LATEs of compensation on employment with different functional forms. 2010–2015. Fuzzy RDD.	152
3.A20	RD Manipulation test following Cattaneo et al. (2020)	152
4.1	Summary statistics	175
4.2	Average carbon intensity by SIC code	178
4.3	Effects of climate policy uncertainty on historical stock price volatility (30-day horizon) and share prices in the US (monthly series).	182
4.4	Effects of climate policy uncertainty on historical stock price volatility (30-day horizon) and share prices in the US (quarterly series).	183
4.5	Effects of climate policy uncertainty on R&D expenses and employment in the US. 1990 - 2018 (quarterly series).	183
4.6	Effects of N-CPU on volatility, share prices, RD expenses and employment in the US. 1990 - 2018 (quarterly series).	186
4.7	Effects of climate policy uncertainty on volatility, share prices, RD expenses and employment in the US. 1996-2018.	186

4.8	Effects of climate policy uncertainty on volatility, share prices, RD expenses and employment in the US. Alternative exposure measure (1).	187
4.9	Effects of climate policy uncertainty on volatility, share prices, RD expenses and employment in the US. Alternative exposure measure (2).	189
4.10	Historical Breakdowns (quarterly series).	190
4.11	Effects of climate policy uncertainty on volatility, share price, RD expenses and employment in the US. 1990 - 2018 (quarterly series). Comparing CPU- and CPU+.	191
4.A1	Estimated changes in volatility associated with CPU changes from 2000 to 2018 for firms in selected industries.	202
4.A2	Estimated changes in share prices associated with CPU changes from 2000 to 2018 for firms in selected industries.	203
4.A3	Estimated changes in RD expenses associated with CPU changes from 2000 to 2018 for firms in selected industries.	204
4.A4	Estimated changes in employment associated with CPU changes from 2000 to 2018 for firms in selected industries.	205
4.A5	Effects of climate policy uncertainty on volatility, share prices, and RD expenses in the US (yearly series).	209
5.1	Effects of a 1 $\mu\text{g}/\text{m}^3$ annual increase in $\text{PM}_{2.5}$ on pharmaceutical expenditures	235
5.2	Effects of a 1 $\mu\text{g}/\text{m}^3$ annual increase in $\text{PM}_{2.5}$ on GVA per capita	235
5.3	Effects of an annual increase in air pollution on pharmaceutical expenditures	241
5.4	Effects of an annual increase in air pollution on GVA per capita	241
5.5	Effects of air pollution on the issuance of sick leaves by diagnosis.	243
5.A1	Exposure to pollution information in the news and demand for pharmaceuticals	277

List of Figures

2.1	Synthetic Control Method results for emissions	38
2.2	Mean annual percentage gap in CO ₂ , PM _{2.5} and NO _x emissions	39
2.3	Inference results for the Synthetic Control Method	40
2.4	Effects of the eco-tax on low-carbon patented technologies . . .	43
2.5	Predicted emissions by fuel under different taxation scenarios .	49
2.6	Share of total emission reductions by fuel due to the eco-tax . .	51
2.7	Substitution towards diesel due to the eco-tax	52
2.8	Underlying mechanisms of reductions in emissions	53
2.9	Evolution of salience in the media over time	54
2.10	Predicted fuel use under different tax and salience scenarios . .	57
2.11	Reductions in external climate and pollution damages due to the eco-tax	59
2.A1	Fuel consumption over time	67
2.A2	Nominal taxes of gasoline and diesel over time	68
2.A3	Real fuel prices and their tax components over time	68
2.A4	Fuel taxes in Germany and the OECD average	68
2.A5	Comparing donor pool weights across SCM specifications . . .	71
2.A6	In-time placebos	72
2.A7	Results with no donor pool restrictions	72
2.A8	Leave-one-out tests	73
2.A9	GSCM with Interactive Fixed Effects Models	74
2.A10	Effects of the eco-tax with the GSCM on low-carbon patents (No sample restrictions)	74

2.A11	Fuel-specific real price and eco tax elasticities with a lead . . .	78
2.A12	Predicted NO _x emissions by fuel under different tax scenarios .	80
2.A13	Gap in per capita emissions: SCMs vs Simulation Approach . .	81
2.A14	Effects of the eco-tax on road casualties	81
2.A15	Dynamic SDID effects in Germany and time-specific weights . .	83
2.A16	Dynamic SDID effects in Sweden and time-specific weights . .	84
2.A17	Dynamic SDID effects in Finland and time-specific weights . .	85
3.1	Carbon prices and compensation payments in the UK	102
3.2	Predicting electricity consumption from energy purchases . . .	109
3.3	Share of compensated plants by electricity cost share. 2005-2011	122
3.4	RD Plot based on quantile spaced number of bins. 2013-2015. .	123
3.5	Comparing ATEs and LATEs across ABS outcome variables. .	126
3.A1	Industrial electricity prices, pre-and post-compensation for selected EU countries in 2016. EUR/MWh.	133
3.A2	Compensation payments as a share of electricity prices. 2013-2017.	134
3.A3	Raw average trends in key outcome variables from ABS over time, by year. 2005-2019.	136
3.A4	Raw average trends in key outcome variables from QFI over time, by year. 2005-2015.	137
3.A5	Comparing predicted and reported electricity use.	138
3.A6	Share of compensated plants, by average energy intensity. . . .	139
3.A7	Comparison of DiD estimates (1)	147
3.A8	Comparison of DiD estimates (2)	148
3.A9	Treatment effects of compensation, by year. 2010-2015.	149
3.A10	RD bandwidth selection procedure	154
3.A11	Comparing LATEs on sales across different bandwidths. 2010- 2015.	155
3.A12	Comparing LATEs on energy intensity (scaled by sales) across different bandwidths. 2010-2015.	156

3.A13	Comparing LATEs on energy purchases across different bandwidths. 2010-2015.	157
4.1	Article counts in the <i>New York Times</i> (US)	166
4.2	Ratio of CPU articles over total articles in the <i>New York Times</i> (US)	167
4.3	CPU index in the United States	168
4.4	CPU index and associated events in the United States	169
4.5	CPU index and associated events in the United States	170
4.6	Comparing the evolution of the CPU and N-CPU indices	172
4.7	Quarterly CPU- index in the United States	174
4.8	Quarterly CPU+ index in the United States	174
4.A1	Comparing the evolution of the CPU to the EPU index developed by Baker et al. (2016)	206
4.A2	Comparing the evolution of the CPU to the VIX index	207
4.A3	Comparing the evolution of the CPU to WTI Crude	207
4.A4	Comparing the evolution of the CPU to Brent Crude	208
4.A5	Frequency distribution of employment levels in the estimation sample	210
5.1	Prescription items per capita across income deprivation quintiles	223
5.2	Thermal inversion episodes as instrument for air pollution	226
5.3	Correlogram of different definitions of inversion episode	227
5.4	Distribution of inversion events (950 hPa) in the UK	228
5.5	Distribution of PM _{2.5} and SO _x concentrations across England	229
5.6	Correlogram of different air pollutants	238

5.7	Alternative model specifications	240
5.8	Distribution of health benefits	246
5.9	Distribution of productivity benefits	247
5.10	Distribution of benefits of a $\mu\text{g}/\text{m}^3$ $\text{PM}_{2.5}$ reduction	248
5.A1	Distribution of inversion events (950 hPa) in 2012	253
5.A2	Distribution of inversion events (950 hPa) in 2013	254
5.A3	Distribution of inversion events (950 hPa) in 2014	255
5.A4	Distribution of inversion events (950 hPa) in 2015	256
5.A5	Distribution of inversion events (950 hPa) in 2016	257
5.A6	Distribution of inversion events (950 hPa) in 2017	258
5.A7	Distribution of inversion events (950 hPa) in 2018	259
5.A8	Frequency monthly counts of inversion episodes	260
5.A9	Distribution of air pollutants concentrations across England	261
5.A10	Spatial distribution of GPs and patients across England	262
5.A11	Spatial distribution of prescription items by therapeutic section	263
5.A12	Spatial distribution of GP practices in English regions	264
5.A13	Spatial distribution of Clinical Commissioning groups (CCGs) in England	265
5.A14	Spatial distribution of LAUs in the London region	266
5.A15	Quadratic predictions of pollution concentrations in 2011 based on income deprivation scores	267
5.A16	Results by therapeutic section	268
5.A17	Results by by air pollutant	269
5.A18	Results by air pollutant	270
5.A19	Effects of air pollutants on prescriptions items	271
5.A20	Effects of air pollutants on total costs to the NHS	272
5.A21	Dynamic RF effects on prescriptions items	273
5.A22	Effects of a $1 \mu\text{g}/\text{m}^3$ annual increase in $\text{PM}_{2.5}$ on pharmaceutical expenditures	273

5.A23	Effects of PM _{2.5} on demand for pharmaceuticals over time. . .	274
5.A24	Evolution of the newspaper index over time. 2012–2018. . . .	276
5.A25	Ratio of matching articles over total articles	279

Chapter 1

Introduction

This chapter serves as an introduction to the doctoral dissertation and it is structured as follows. First, it begins by exploring the overarching motivation and introducing the central environmental issue at the core of its focus: climate change. Second, it discusses different approaches to regulatory interventions to tackle climate change externalities according to economic theory. Third, Section 1.2 outlines the focus of each chapter and suggests directions for future research.

1.1 Motivation and research focus

Tackling climate change. Climate change is a complex and multifaceted global challenge that requires the development of effective policy responses to mitigate and adapt to its impacts. The scientific understanding of climate change and its impacts has evolved over the past few decades, with increasingly robust evidence linking human activities to rising temperatures and other climate-related changes (IPCC, 2022). In the absence of a substantial reduction in greenhouse gas (GHG) emissions within the coming decades, the increase in mean temperature and in the number of extreme weather events will threaten livelihoods and likely exacerbate poverty and accentuate existing inequalities, both on the global stage and within

jurisdictions (Diffenbaugh and Burke, 2019). In response, there has been a growing urgency to develop policy responses capable of not only mitigating and adapting to the consequences of climate change but also shepherding the transition toward a sustainable, low-carbon economy.

The complexity and scale of the climate change challenge necessitate a coordinated global response. Policymakers face the task of devising climate policies that are not only effective in reducing emissions but also equitable, politically feasible, and economically viable. Achieving consensus on emission reduction targets, aligning incentives among countries with varying levels of economic development, and addressing the distributional impacts of climate policies are some of the intricate issues confronting policymakers (Barrett and Stavins, 2003; Lange et al., 2007). Economic costs associated with climate change are substantial. Among others, the disruption of agricultural activities, human loss of life due to increased frequency and severity of weather extreme events, and damage to critical infrastructure impose significant economic burdens (Deschênes and Greenstone, 2007; Newman and Noy, 2023). Effective climate policies are instrumental in mitigating these costs and stimulating innovation and investment in green technologies, positioning countries at the forefront of the emerging low-carbon economy (Ambec et al., 2013).

Notably, climate change is not merely an environmental and economic issue but also a critical determinant of public health (McMichael et al., 2006; Romanello et al., 2021; Baker et al., 2022). Elevated temperatures, altered disease vectors, and compromised air quality are among the pathways through which climate change influences health outcomes. Robust climate policies, by curbing emissions, have the potential to mitigate these adverse health effects and safeguard human well-being, yielding substantial co-benefits (Parry et al., 2015; Vandyck et al., 2020). Furthermore, consideration of the climate-health nexus is of utmost importance when addressing the distributional concerns of climate policies, as it underscores the intricate interplay between environmental impacts and their uneven societal repercussions. Vulnerable populations, including low-income communities and marginalized groups, often bear

a disproportionate burden of the adverse impacts of climate change (Hallegatte and Rozenberg, 2017). Consequently, climate policies must be designed with careful consideration of the distributional implications to ensure that the benefits and costs of mitigation efforts are equitably distributed (Banzhaf et al., 2019; Hernandez-Cortes and Meng, 2023).

In the context of a globalized economy, international cooperation is a prerequisite for effective climate policy. The evolution of climate policy on a global scale has been characterized by a series of setbacks and periods of increased stringency, reflecting the inherent complexities and diverse interests at play. The late 20th century witnessed the emergence of international climate agreements, notably the United Nations Framework Convention on Climate Change (UNFCCC) in 1992 and the Kyoto Protocol in 1997. However, subsequent negotiations faced substantial hurdles, with the withdrawal of the United States from the Kyoto Protocol in 2001 serving as a stark setback. Policy setbacks incur societal costs, as regulatory uncertainty delaying investments in low-carbon technologies may raise GHG levels, exacerbating climate change irreversibly, and ultimately increasing transition costs to a low-carbon economy (Dorsey, 2019).

Nevertheless, the 21st century also bore witness to critical milestones, such as the Paris Agreement in 2015, which marked a resurgence of global climate diplomacy. However, challenges such as differing national interests, enforcement mechanisms, and geopolitical tensions persist in the international arena, complicating the translation of climate policy commitments into concrete action (Nordhaus, 2021). Simultaneously, voluntary pledges announced by individual jurisdictions have been found to fall considerably short of the ambitious targets set by the Paris Agreement to stay below 1.5 °C of global warming (Rogelj et al., 2016), thereby giving rise to legitimate concerns regarding the sufficiency and effectiveness of existing mitigation efforts.

In summary, the exigency of addressing climate change lies at the nexus of environmental, economic, and social imperatives. The efficacy of climate policies is

contingent upon their ability to reconcile the challenges inherent in their design, delivery, and evaluation. Navigating trade-offs between economic growth, social equity, and environmental sustainability necessitates a nuanced understanding of the complexities surrounding climate change and an unwavering commitment to interdisciplinary, evidence-based policy design and assessment. Consequently, our ability to address the detrimental effects of climate change and facilitate the shift toward a low-carbon global economy hinges on a better understanding of the effects of different approaches to climate policymaking.

Economic principles and climate policy. Economists have long identified the emissions of GHGs as a market failure that requires to be addressed (Montgomery, 1972; Arrow et al., 1997). According to economic theory, there are broadly two categories of public policy designs to regulate the emissions of GHGs. First, command-and-control regulations, which involve establishing requirements for the specific technology used (i.e., technology mandates) or setting maximum pollution levels (i.e., performance standards), with subsequent monitoring and enforcement. Second, market-based or incentive-driven approaches involve imposing a price on GHGs, such as carbon dioxide, to incentivize emitters to reduce their emissions. Market-based solutions can be further categorized into those that control the quantity of emissions (e.g., cap-and-trade systems, which set a limit on the quantity of emission permits) and those that regulate the price of emissions (e.g., Pigouvian taxes). Various criteria can be employed to assess the advantages of different regulatory approaches. These criteria may encompass economic efficiency, cost-effectiveness, the distributional patterns of benefits and costs, the capacity to address uncertainties, and the feasibility within the political landscape as potential metrics for evaluation (Goulder and Parry, 2008).

In terms of efficiency, theoretical evidence underscores the superiority of market-based solutions over command-and-control regulations under first-best conditions. This is primarily attributed to the inherent differences among emitters in their

capacity to curb GHG emissions, as measured by their respective marginal abatement costs (Newell and Stavins, 2003). Within this context, incentive-based instruments prevail by aligning marginal costs with marginal benefits across all emitters provided that they possess knowledge of their own abatement costs while policymakers do not. Another theoretical advantage of market-based instruments lies in their ability to harness all available emissions reduction avenues, while command-and-control regulations often overlook some, especially those related to output reduction. Finally, incentive-based instruments typically generate revenues (except in cases where emission permits are allocated across emitters at no cost) that have the potential to mitigate other distortionary taxes and yield efficiency gains. The concept of revenue recycling presents the prospect of a *double dividend* which could enhance environmental quality while concurrently reducing the net welfare cost associated with environmental policy, as discussed by scholars such as Baumol and Oates (1988), Pearce (1991), Carraro et al. (1996), and Chiroleu-Assouline and Fodha (2014).

However, when confronted with the presence of multiple market failures, the unequivocal superiority of market-based instruments may no longer hold on theoretical grounds and a combination of different instruments may be warranted, such as market-based, standards and public subsidies. A key example arises when the administrative costs associated with monitoring emissions soar, rendering command-and-control regulatory approaches a more effective choice (Goulder and Parry, 2008). Additionally, in scenarios characterized by pre-existing distortions within factor markets, clean energy standards tend to outperform price-based instruments due to their lesser implicit taxation on production factors. Without government intervention, competitive markets are also expected to under-incentivize private investment in the development and diffusion of new low-carbon technologies (Rosendahl, 2004; Fischer, 2008). This holds true particularly when the social marginal cost of pollution deviates from the market carbon price and uncertainties remain over the durability of the price signal over the long term (Ambec et al., 2013). The *appropriability* problem due to the public good nature of innovation implies that incentives for clean technology R&D

will be inefficiently low, even if emissions' externalities are appropriately priced. This underscores the need to complement carbon pricing schemes with tailored technology policies, such as government research support projects that can help restore low-carbon invention efforts to an efficient level. Finally, the road transport sector investigated in Chapter 2 serves as an illustrative example of multiple market failures, encompassing not only GHG emissions, but also local air pollution, noise, accident risks, and congestion, and thus necessitating the adoption of second-best policies.

When deciding between price or quantity-based instruments, Weitzman (1974) asserts that the choice hinges on the steepness of the marginal damage function, particularly in situations involving uncertainty regarding aggregate emission reduction costs — a common scenario in practical real-world applications. That is, price-based instruments prove superior when the marginal damage curve exhibits flatness, while quantity-based instruments are preferred when the curve demonstrates pronounced steepness. This reasoning stems from the crucial importance of delivering precise emission targets when even slight changes in emissions yield substantial increases in damage costs. Accordingly, given the relatively flat nature of the marginal damage curve in the context of climate change, economists have traditionally argued in favor of a carbon tax. On a global scale, Weitzman (2015) further advocates for a uniform carbon tax rather than internationally tradable permits, referring to uncertainty surrounding country-specific abatement cost profiles as a key rationale for this preference.

In terms of equity considerations, regulatory intervention can give rise to distributional concerns, particularly when certain individuals bear a disproportionate burden of the regulation. As a case in point, in high-income countries, carbon taxes have often exhibited regressive characteristics if the tax revenues are not effectively redistributed (Callan et al., 2009; Preuss et al., 2021; Köppl and Schratzenstaller, 2023). This regressive nature arises because individuals with lower incomes tend to allocate a more substantial proportion of their consumption expenditures toward carbon-intensive

goods. Existing research has revealed that standards may exhibit a higher degree of regressiveness compared to a carbon tax coupled with lump-sum transfers (Levinson, 2019; Davis and Knittel, 2019). The policy debate, however, often exclusively focuses on the costs of environmental and climate policies for consumers. Applied modeling studies suggest that the overall distributional effects, which consider also source-side impacts on wages and capital incomes, are less regressive or even progressive (Goulder et al., 2019). Nevertheless, this still ignores the distribution of environmental and health co-benefits induced by local air pollution improvements, which are also likely distributed progressively (e.g., Drupp et al., 2018, 2021; Hernandez-Cortes and Meng, 2023). Existing research has indicated that the regressive consumer-side impacts of carbon taxes can be mitigated when the tax proceeds are redistributed through lump-sum transfers to all households, as demonstrated, for instance, by studies by Metcalf (2009), Klenert and Mattauch (2016), and Cronin et al. (2019). However, lump-sum transfers may not entirely address horizontal equity issues due to the significant variability in tax incidence within income deciles (cf., Fischer and Pizer, 2019).

In the current climate policy landscape, there is a broad consensus among economists that putting a price on GHG emissions would be the central element of an efficient policy response to climate change. This was highlighted in 2019 by the “Economists’ Statement on Carbon Dividends” (Wall Street Journal, 2019) and the “Economists’ Statement on Carbon Pricing” (EAERE, 2020), jointly signed by around 5000 economists. Yet, around 77% of global emissions are currently still not facing a carbon price (World Bank, 2022). The lack of more widespread pricing may be a consequence of its limited acceptability, primarily driven by concerns about its efficiency and distributional implications (Klenert et al., 2018). For instance, recent research underscores the significance of how policies are perceived (Maestre-Andrés et al., 2019; Douenne and Fabre, 2022). Contextual factors, such as levels of political trust, have also been recognized as influential determinants of policy acceptance (Rafaty, 2018). These additional dimensions wield substantial influence on the

reception, acceptance, and ultimate effectiveness of regulatory measures. On these premises, a more solid and comprehensive evidence base on the impacts of carbon pricing could play a pivotal role in enhancing its credibility and political support.

1.2 This dissertation

This doctoral dissertation comprises four empirical research papers, each dedicated to exploring distinct facets of climate policymaking. These papers aim to provide novel evidence on the effectiveness of climate policy measures (Chapter 2), examine the trade-offs in policy design concerning mitigation and leakage (Chapter 3), analyze the role of uncertainty in policy implementation (Chapter 4), and investigate the spatial distribution of pollution-reduction benefits (Chapter 5). Below, I provide a summary of the focus of each chapter while introducing the underlying methodological approaches. Finally, I suggest directions for future research.

1.2.1 Chapter 2: Causal effects of fuel taxation and mediating mechanisms for reducing climate and pollution costs

Focus of this chapter. Chapter 2 is a joint collaboration with Sophie Behr (The German Institute for Economic Research) and Moritz Drupp (University of Hamburg). We draw on multiple causal inference methods to conduct an empirical assessment of the world's largest environmental tax reform, the German eco-tax, which increased fuel taxes in Europe's biggest transport sector in yearly steps from 1999 to 2003 up to 15.35 cents per liter.

The analysis begins by estimating the effects of the eco-tax on emissions of CO₂, PM_{2.5}, and NO_x and on low-carbon patenting in the German transport sector.

First, we use the synthetic control method (SCM) (e.g., Abadie, 2021) to build counterfactual Germanies with weighted combinations of donor countries and compare emission paths of the German transport sector and its synthetic counterfactuals. Second, we corroborate our results using the generalized SCM (GSCM) to construct counterfactuals by modeling emissions and low-carbon patenting with interactive fixed effects models (Xu, 2017) and restrict the donor pool to EU countries to rule out that effects are driven by EU-wide regulation, like emission standards (e.g., Reynaert, 2021). Finally, we validate the robustness and external validity of our findings by harnessing the staggered adoption of other environmentally-motivated taxes in Europe in a synthetic difference-in-differences (SDID) design to address potential impacts of concurrent unobserved idiosyncratic shocks (Arkhangelsky et al., 2021).

In addition to our causal findings, we provide complementary evidence on mediating mechanisms by leveraging semi-elasticity models harnessing cross-country panel variation in fuel prices and tax rates. We further explore the role of fuel substitution in navigating the trade-off between attaining climate and pollution targets. Finally, we test whether different demand responses to the eco-tax with respect to market-driven fuel price changes are induced by tax salience, which we measure explicitly based on newspaper data. We thereby provide the first direct empirical evidence for the hypothesis that consumers react more strongly to (environmentally-motivated) fuel taxes the more salient they and their associated price increases are.

1.2.2 Chapter 3: Carbon pricing, compensation, and competitiveness: Lessons from UK manufacturing

Focus of this chapter. Chapter 3 is a joint work with Elisabeth Isaksen (Ragnar Frisch Centre for Economic Research) and Misato Sato (The London School of Economics and Political Science). We empirically examine UK manufacturing firms’

responses to an output-based carbon cost compensation scheme introduced in 2013 following the implementation of a carbon price floor. By shielding firms from the full carbon cost, compensation results in partial carbon cost internalization and likely compromises efficient carbon price incentives to decarbonize industrial production and consumption. In particular, compensation payments based on production volumes (known as “output-based allocation”) essentially reward firms for each unit of production and mitigate the increases in marginal costs of production that result from emissions pricing (Fischer and Fox, 2011) and provide an implicit production subsidy (Fischer and Fox, 2007; Fowlie et al., 2016; Meng, 2017). Dampening incentives to limit supply from energy-intensive sectors means that to achieve the overall ETS cap, the mitigation burden shifts elsewhere (to other sectors or towards greater emissions intensity improvements) which means allowance prices and overall costs rise. This perverse production incentive effect has been highlighted in the theoretical literature (Fischer, 2001; Demailly and Quirion, 2008; Böhringer et al., 2012; Fischer and Fox, 2011) but downplayed in policy debates arguably due to the lack of robust empirical evidence.

To explore the effects of output-based compensation payments, we combine two quasi-experimental research designs; a difference-in-difference (DiD) design with inverse propensity score weighting and a “fuzzy” regression discontinuity design (RDD). In both approaches, we exploit variation caused by the UK eligibility rules for receiving compensation to identify effects. We obtained confidential microdata from the UK secure data lab on economic variables and energy use at the plant level and combined it with a publicly available list of firms that received compensation. While eligibility for compensation is assessed at the firm level, the amount of compensation paid is calculated at the plant level and is linked to the plant’s output. Compared with firm-level analysis, more disaggregated plant-level data is advantageous because firms may operate multiple plants across different sectors. We are comparing similar plants belonging to compensated and non-compensated firms to isolate the effects of the compensation for indirect carbon costs.

1.2.3 Chapter 4: Climate policy uncertainty and the behavior of firms and investors

Focus of this chapter. How does uncertainty in climate policy affect the behavior of firms and investors? Chapter 4 sheds light on this question in a joint collaboration with Stefano Carattini (Georgia State University), Antoine Dechezleprêtre (OECD), and Tobias Kruse (OECD). To address this question, we built a novel index of policy uncertainty specific to climate policy, which allows us to address this question empirically. Our “climate policy uncertainty” index, or CPU, builds on the seminal work of Baker et al. (2016) leveraging textual analysis of newspaper data to proxy economic policy uncertainty and combines their original search strategy with keywords related to climate policy. Our index runs monthly from 1990 onward and covers the main newspapers in the United States. Then, we analyze the relationship between CPU and firm outcomes such as share prices, implied volatility, employment decisions, as well as investments in research and development.

Our approach also takes into account a crucial feature related to climate policy. While in the case of standard economic policy, the economy tends to move along a given trajectory determined by its steady state and uncertainty tends to be detrimental to economic growth, in the case of climate change the economy needs to transition from fossil-fueled activities to a cleaner way of production. Hence, the economy needs to move from one equilibrium, which is carbon intensive, to another equilibrium, which is much cleaner. Since climate change entered the policy arena in the 1980s, both domestic and international climate policymaking have gone through important achievements as well as numerous setbacks. If firms and investors respond to short-term variation in the probability of future policy tightening, rather than adopting long-term goals such as decarbonization, setbacks are likely to benefit them. For this reason, our index is complemented by two sub-indices, aimed at measuring whether the source of uncertainty is an acceleration in the process of decarbonization, or

rather a deceleration.

The primary empirical goal of this chapter is to examine how economic outcomes respond to greater uncertainty about climate policy, also depending on its drivers. To do so, we exploit variations in our CPU index, and its sub-indices, across different months, quarters, or years from 1990 onward. Specifically, we estimate fixed effects models where we interact our news-based indices with the average carbon intensity across 4-digit SIC industries. By doing so, we develop an identification strategy that differentiates firms according to their relative exposure to climate policy risk. Using panel data on publicly listed companies extracted from Compustat, our model tests whether exposure to climate policy risk matters for economic outcomes when greater uncertainty about climate policy materializes as measured by newspaper article coverage.

1.2.4 Chapter 5: Pollution reduction benefits across space: Quasi-experimental evidence from England

Focus of this chapter. How much pollution reduction is socially desirable, and which societal groups would reap the greatest benefits from these reductions? From a social welfare standpoint, answering these questions requires information on how much individuals value pollution control, which can be assessed by retrieving empirical estimates of their willingness to pay for ameliorating air quality.

To delve into these questions of key general economic and policy importance, in Chapter 5 I provide novel quasi-experimental evidence on the causal effects of air pollutants on (i) *defensive expenditures* as proxied by pharmaceutical expenditures and (ii) *economic productivity* measured by GVA per capita to retrieve empirical estimates of the social willingness to pay for air quality improvements that account for heterogeneity in local-scale benefits. To address endogeneity concerns due to residential sorting (cf. Heblich et al. 2021), atmospheric temperature inversions are

exploited as a source of quasi-random dynamic variation in the spatial concentration of air pollutants across England (e.g., Arceo et al., 2016; Dechezleprêtre et al., 2019). I compile a novel and unprecedentedly granular dataset for England that combines high-resolution vertical temperature profiles from the European Centre for Medium-Range Weather Forecasts with gridded pollution maps, administrative practice-level healthcare records covering 54 million residents, and data on economic activities at the district level. The spatial resolution of my data allows me to exploit within-district variation in thermal inversion exposure in a two-stage least squares setting.

To explore spatial heterogeneity, I harness my causal estimation framework and simulate counterfactual reductions of $1 \mu\text{g}/\text{m}^3$ in $\text{PM}_{2.5}$ and investigate the distributions of both health and productivity benefits across the population and their correlation with socio-economic factors. Additionally, I leverage the dynamic nature of the exogenous source of variation in my empirical setting and explore how revealed preferences have evolved over time. I rely on press coverage of the negative health effects of air pollution to quantify the role played by changes in information exposure, which I measure by constructing an index based on textual analysis of British newspapers' articles.

1.2.5 Future research

Saliency and instrument choice. The important role of tax saliency in shaping fuel demand that we document in Chapter 2 can have implications for instrument choice. Due to inherently fluctuating prices, price saliency may likely be less pronounced in the case of emissions trading schemes. As such, increased fuel prices may induce a lower demand response at a given carbon price rate. Investigating the role of saliency for demand responses when policy relies on emission trading schemes is especially important given that the European Union has recently decided to introduce a second emissions trading scheme that encompasses the transport sector

and that may subsequently replace the current approach in many countries to levy taxes on fuels. Such a shift to emission trading schemes may require more targeted communication and information campaigns to yield comparable demand responses as equivalent fuel or carbon tax rates. This is of particular relevance, especially considering that the European Union has recently implemented EU ETS II for the transport sector, signifying a shift away from fuel taxation as the primary policy instrument for curbing greenhouse gas emissions. Future research should investigate the differential salience effects of different policy instruments and compare differential demand responses to different market-based approaches to carbon pricing.

Unilateral policies and mitigation responsibilities in a cap-and-trade. The findings presented in Chapter 3 shed light on how unilateral policy actions taken by individual jurisdictions can intersect with broader inter-jurisdictional climate policies, such as the EU ETS, potentially resulting in the redistribution of mitigation responsibilities among sectors and countries - the so-called *waterbed effect* (Perino, 2018). This has two significant implications. Firstly, compensation schemes can lead to a relocation of mitigation efforts to sectors facing higher abatement costs. Secondly, this sectoral redistribution can impact the distribution of co-benefits arising from carbon mitigation (i.e., reduced emissions of air co-pollutants yielding local health benefits). Investigating these two research directions is essential for gaining a deeper understanding of the economy-wide consequences of carbon pricing policies paired with compensation schemes. Such insights are crucial for policymakers when evaluating these policy design choices in comparison to alternative methods of utilizing auction revenues, including different revenue recycling mechanisms.

Distributional impacts of climate policy. Considering the spatial distribution of benefits associated with various climate policy approaches has the potential to enhance the efficiency of environmental and climate regulations, as discussed in Chapter 5. Currently, there exists limited empirical evidence concerning the distribu-

tional consequences of climate policies, with a predominant focus on consumer-side impacts (Köppl and Schratzenstaller, 2023). Nevertheless, computational modeling studies suggest that the overall distributional effects, when encompassing source-side effects on wages and capital incomes, may be less regressive or even exhibit progressive tendencies (Goulder et al., 2019). However, these analyses still overlook the distribution of environmental and health co-benefits resulting from local air pollution improvements, which, as demonstrated in Chapter 5, are likely to benefit lower-income populations (e.g., Drupp et al., 2018, 2021; Hernandez-Cortes and Meng, 2023). Therefore, future research should prioritize expanding our understanding of the distribution of different types of policy-induced costs and benefits to enable a more comprehensive evaluation of the societal welfare impacts of diverse instruments.

Non-human benefits of climate policy. While much of the existing empirical research on the co-benefits of climate policy primarily focuses on human impacts, it is crucial to account that carbon mitigation and the resulting reductions in co-pollutants can yield significant benefits within the non-human nature (Liang et al., 2020; Lin et al., 2023; Sanderfoot et al., 2022). One example is the reduction in the deposition of acidifying substances like sulfur dioxide (SO_2) and nitrogen oxides (NO_x), largely originating from the combustion of fossil fuels. The deposition of these pollutants can have far-reaching effects on biodiversity and terrestrial ecosystems (Zvereva and Kozlov, 2010), potentially influencing vital ecosystem services, such as the soil carbon cycle. Future research should aim to investigate the extent to which current levels of air pollution deposition may be contributing to shifts in the balance between soil carbon input and output which may ultimately affect the pace of climate change. Future interdisciplinary explorations can shed light on additional non-human benefits resulting from climate mitigation, helping to assess the optimal policy stringency based on a more comprehensive range of benefits that extend beyond the realm of human impacts.

Updating official cost estimates. The official cost estimates to conduct cost-benefit analyses of policies fostering air pollution reduction which have been mentioned in this dissertation are based on health effects driven by selected medical conditions (i.e., affecting the respiratory and cardiovascular systems) and currently do not account for direct productivity effects (Umweltbundesamt, 2012; UK-AIR, 2023). These omissions pose the risk of significantly underestimating the comprehensive net benefits attributable to such regulations. My findings from Chapter 5, along with recent advancements in the literature (e.g., Dechezleprêtre et al., 2019; Leroutier and Ollivier, 2022) show that productivity gains potentially represent a substantial share of the benefits of pollution reduction. Additionally, an emerging body of literature has linked exposure to air pollution to additional health conditions, affecting the central nervous system (Zhang et al., 2018). Accounting for such additional potential health co-benefits, which more immediately benefit those that bear the costs of higher carbon prices, may also be crucial for gathering greater support for climate policies (e.g., Löschel et al., 2021). In forthcoming revisions of governmental guidelines for policy appraisal, it is crucial to incorporate these additional benefits into the assessment framework. Furthermore, my research supports the case for equity weighting criteria in cost-benefit analysis that account for the distribution of pollution reduction benefits among various socio-economic groups: such an approach has the potential to foster policy intervention with a more equitable distribution of net benefits, thereby contributing to enhancing their overall acceptability. The principle of equity weighting is already discussed in the UK government’s guidelines on cost-benefit analysis (Treasury, 2016) while distributional weights are currently employed in Germany in the estimation of climate change damages (UBA, 2019).

Chapter 2

Causal effects of fuel taxation and mediating mechanisms for reducing climate and pollution costs

with Sophie Behr (The German Institute for Economic Research) and Moritz Drupp (University of Hamburg)

SUMMARY. This chapter provides an assessment of how fuel taxation reduces climate and pollution externalities with a quasi-experimental evaluation of the world's largest environmental tax reform. Leveraging multiple causal inference methods, we compare carbon and air pollutant emissions of the actual and counterfactual German transport sector following the 1999 eco-tax reform and demonstrate sizable reductions in carbon, particulate matter, and nitrogen dioxide emissions. Using official cost estimates, the eco-tax saved around 80 billion Euros of external costs, predominantly relating to pollution reduction benefits. We further show that environmental taxation contributed substantially to fostering low-carbon innovation. In complementary analyses, we document much stronger demand responses to increases in environmentally-motivated taxes than to market price movements, which we relate

primarily to increased tax salience in newspapers. Our findings highlight the roles of salience and fuel substitution in mediating the effectiveness of fuel taxes to deliver climate and pollution reduction benefits.

2.1 Introduction

Fuel taxation is a key policy instrument to reduce negative externalities of fossil-fuelled transportation (Parry et al., 2007; Sterner, 2007; Hintermann et al., 2021) and has seen renewed interest due to concerns about climate change, air pollution, and energy security (e.g., Grigolon et al., 2018; Parry et al., 2021). Understanding how fuel taxation affects fuel demand is essential to effectively leverage this tool for policy. Many assessments assume that demand responses to tax changes are equivalent to those of market-driven price variations and estimate limited impacts of carbon taxes (e.g., Green, 2021). In contrast, recent work highlights the considerable role of tax salience effects (e.g., Chetty et al., 2009; Li et al., 2014), which may suggest that more modest taxes may achieve politically targeted fuel reductions. Additionally, carbon abatement represents only part of the economic benefits that can justify fuel taxation. Importantly, transportation causes considerable health damages linked to air pollution (e.g., Schlenker and Walker, 2016; Knittel et al., 2016) and reducing fossil fuel use can thus yield substantial health benefits (e.g., Shaw et al., 2014; Parry et al., 2015). Accounting for such health co-benefits may be important for gathering public support for fuel or carbon pricing.

We investigate the effectiveness of fuel taxation in reducing carbon and air pollutant emissions with a quasi-experimental assessment of the world’s largest environmental tax reform: the German eco-tax. The reform increased fuel taxes in Europe’s biggest transport sector in yearly steps from 1999 to 2003 up to 15.35 cents per liter. In 2003, implicit carbon costs due to the eco-tax amounted to €58 (\$65) per tCO₂ for diesel and €66 (\$74) for gasoline. This was then the second highest effective carbon

price globally—higher alone than federal fuel taxes in the US, where regulation has mainly focused on standards (Jacobsen et al., 2023), and only slightly lower than the Swedish carbon tax on transport fuels that was levied on a much smaller tax base (Andersson, 2019).

Our analysis starts by estimating effects of the eco-tax on emissions of CO₂, PM_{2.5}, and NO_X in the German transport sector drawing on a battery of causal inference methods. First, we use the synthetic control method (SCM) (e.g., Abadie, 2021) to build counterfactual Germanies with weighted combinations of control countries and compare emission paths of the German transport sector and its synthetic counterfactuals.¹ Our SCM results imply that, between 1999 and 2009, the eco-tax led to emission gaps of around 10% for CO₂, 27% for PM_{2.5}, and 13% for NO_X on average across specifications, and to an average reduction in external damages of around 80 billion euros when using official cost estimates.² We, second, corroborate our results using the generalized SCM (GSCM) to construct counterfactuals by modeling emissions with interactive fixed effects models, and restrict the donor pool to EU countries to rule out that effects are driven by EU-wide regulation, like emission standards (e.g., Reynaert, 2021). Finally, we validate the robustness and external validity of our findings by harnessing the staggered adoption of other environmentally-motivated taxes in Europe in a synthetic difference-in-differences (SDID) design to address potential impacts of concurrent unobserved idiosyncratic shocks. While modeling studies consistently indicate considerable health benefits due to lower fossil fuel use (e.g., Shaw et al., 2014; Choma et al., 2021), This chapter is the first observational study to quantify the climate *and* pollution reduction benefits of fuel taxation in a quasi-experimental framework. Our assessment of the world’s largest environmental tax reform complements studies on the role of

¹We draw on a growing literature using SCMs to evaluate policies (e.g., Lindo and Packham, 2017; Cunningham and Shah, 2018), particularly for environmental regulations (e.g., Andersson, 2019; Isaksen, 2020; Bayer and Aklin, 2020; Leroutier, 2022).

²Our SCM results are robust to a host of placebo and sensitivity tests, including in-time placebos, alternative donor pools, sets of predictors, different pre-treatment time frames, the exclusion of one donor country at a time, and permutation tests that apply the SCM to every potential donor country.

emission standards to reduce climate and pollution externalities in the transport sector (e.g., Auffhammer and Kellogg, 2011; Jacobsen et al., 2023; Reynaert, 2021) and substantially extends investigations on the effectiveness of environmental taxes that focused exclusively on carbon abatement.³

We further harness our battery of causal inference methods (SCM, GSCM, SDID) to quantify the impacts of environmental taxation on the development of low-carbon patented technologies, building on Aghion et al. (2016), who use transport fuel prices to proxy carbon prices and link them to an increase in innovation in clean technologies in the automobile sector. In contrast, we investigate low-carbon innovation induced by environmentally-motivated taxation, which may yield a greater response due to the higher salience (Stern, 2012b). By focusing on economy-wide patent data, our empirical strategy captures innovation in response to an implicit carbon price that accounts for unregulated companies, upstream equipment manufacturers (Sanyal and Ghosh, 2013), downstream suppliers (Popp, 2019) and new entrants to the market (Noailly and Smeets, 2015), departing from existing firm-level observational studies exploiting policy inclusion criteria (e.g., Calel and Dechezlepretre, 2016; Calel, 2020). We find that the eco-tax has led to a 6% average yearly increase in patented low-carbon technologies concerning the transport sector. Our results thus indicate considerable potential for fuel or carbon taxes for directing technological innovation (Acemoglu et al., 2012) to increase the fuel efficiency and contribute to reducing abatement costs (e.g., Popp, 2019).

Next, we enrich our causal analyses with explorations of mediating mechanisms, focusing in particular on the roles of fuel substitution and tax salience.⁴ We build on a large literature exploring effects of gasoline and energy prices on fuel demand

³Andersson (2019), Mideksa (2021) and a contemporaneous paper (Runst and Höhle, 2021) examine the effectiveness of carbon or fuel taxes to reduce CO₂ emissions using the SCM. We go beyond in several dimensions by investigating effects on air pollution and low-carbon innovation, by disentangling effects by fuel type to illuminate trade-offs between climate and pollution reduction benefits, and by providing first direct evidence on the key role of tax salience.

⁴Analyzing other mechanisms suggests that the eco-tax has likely contributed to fostering fleet renewal of passenger cars and to fewer passenger-kilometers travelled without reduced overall economic activity.

and emissions (e.g., Dahl and Sterner, 1991; Levin et al., 2017; Linn, 2019; Parry et al., 2021), which often relies on fuel and energy prices as proxies for carbon prices and use price changes over time to estimate impacts on fuel demand. Yet, fuel prices are prone to endogeneity concerns, likely biasing price elasticity estimates downwards (e.g., Kilian, 2009; Davis and Kilian, 2011; Coglianese et al., 2017). We use cross-country panel variations in fuel-specific tax rate changes, coupled with an instrumental variable approach, and a set of distributed lag models to account for potential tax anticipation effects (cf. Kilian and Zhou, 2023). Our focus on fuel-specific demand adjustments departs from previous studies that rely on changes in gasoline consumption as a proxy for aggregate emission reductions (e.g., Davis and Kilian, 2011; Rivers and Schaufele, 2015) and helps to illuminate the role of fuel substitution. Accounting for gasoline-to-diesel substitution is crucial in the European context given its high diesel share (Zimmer and Koch, 2017), and allows quantifying trade-offs between climate and health benefits.

We first estimate price and tax elasticities of demand for gasoline and diesel to disentangle behavioral responses in Germany. Our preferred specifications yield a tax-exclusive price elasticity of demand for gasoline (diesel) of -0.32 (-0.26) and an eco-tax elasticity of demand of -2.7 (-1.1). Fuel-specific eco-tax elasticities are thus 4 to 8.5 times higher than the tax-exclusive price elasticity (a ratio referred to as *tax saliency ratio*), in line with prior findings that changes in taxes are more potent than equivalent market-driven price changes (e.g., Li et al., 2014; Rivers and Schaufele, 2015; Andersson, 2019).⁵ This underscores potentially large biases in policy evaluations that rely on responses to market-driven fuel price changes as a proxy for the effect of environmental taxes.

We then use these fuel-specific tax elasticities to perform simulations and find that around three-quarters of the (simulated) reduction in CO₂ emissions is attributable to

⁵Kilian and Zhou (2023) reconsider the analysis by Li et al. (2014) using a distributed lag model—as in Coglianese et al. (2017)—and find that the tax elasticity is not significantly different from tax-exclusive price elasticity in the US after accounting for anticipation effects. In our setting, even after accounting for anticipatory behavior, we still document sizable and significant *tax saliency ratios*.

lower gasoline use, partly driven by gasoline-to-diesel substitution. Conversely, almost all decreases in $PM_{2.5}$, and more than half of decreases in NO_x emissions, are driven by lowered diesel use due to the eco-tax. This highlights important trade-offs that can arise between climate and air pollution targets, which is particularly relevant for price instruments set on the carbon content of fuels that can foster fuel substitution. Such fuel substitution is—with the exception of Linn (2019)—not accounted for in existing policy evaluations. We complement Linn (2019) by relaxing the assumption that consumers respond similarly to fuel taxes as to other changes in fuel prices. We find that accounting for tax salience effects illuminates a much more sizable trade-off between climate and health benefits. This trade-off, and the associated inefficiency in targeting both climate and pollution targets with one price instrument, is a more general feature of second-best taxation (e.g., Knittel and Sandler, 2018), especially when it is not feasible to tax externalities directly (Jacobsen et al., 2023). Nonetheless, both our causal estimates and simulation results using disentangled elasticities provide evidence that the German eco-tax has led to sizable reductions in these “untaxable” air pollution externalities.

Finally, we advance the literature on the role of salience for environmental policy (e.g., Li et al., 2014; Rivers and Schaufele, 2015; Huse and Koptuyug, 2022) by developing a framework to quantify the role of salience changes in the media in driving the effects of the eco-tax. Similarly to Li et al. (2014), who show that a tax change is associated with a greater increase in media coverage than a comparable change in the tax-exclusive fuel price, we rely on media analysis to explicitly investigate tax salience. Specifically, we construct a newspaper-based index to capture the evolution of eco-tax salience based on textual analysis of German newspaper articles (cf. Gentzkow et al., 2019). Leveraging annual variations in our salience index within our elasticity models, we find that greater tax salience is associated with lower consumption of both gasoline and diesel and that these effects increase with the real eco-tax rate. Our simulations suggest that the salience of the eco-tax is responsible for around 70% (55%) of the observed contraction in gasoline (diesel)

consumption. These results provide first direct evidence for the hypothesis that consumers react more strongly to fuel taxes the more salient they are and imply that targeted measures to increase salience may have considerable potential to enhance the cost-effectiveness of price instruments to internalize externalities.

The chapter proceeds as follows. Section 5.4 details the methodologies employed in our research designs. Section 5.3 discusses the data. Section 2.4 presents results derived from SCMs, while Section 2.5 reports results on fuel and tax elasticities, simulations, and additional mediating mechanisms. Section 2.6 quantifies climate and health benefits, while Section 5.7 concludes. The Appendix contains institutional details on the eco-tax reform and supporting materials for our analyses.

2.2 Methodology

2.2.1 Causal Inference Methods

This section introduces our causal inference methods—the SCM (e.g., Abadie and Gardeazabal, 2003; Abadie, 2021), GSCM (Xu, 2017) and SDID (Arkhangelsky et al., 2021)—and explains how we leverage them to estimate causal effects of environmental taxation on carbon and air pollutant emissions and low-carbon innovation.

The SCM estimator. Suppose there are $J + 1$ countries. Each country is indexed by j , where $j = 1$ denotes the *treated* country (i.e., Germany), while $j = 2, \dots, J + 1$ are *untreated* countries (the *donor pool*), which may be used to construct a control group. The T time periods are divided into pre-treatment and post-treatment (i.e., after the eco-tax reform in 1999) with T_0 as the period prior to the policy ($t = t_0, t_{-1}, \dots, T_0$). Denoting the intervention as I , the SCM considers that the observed outcome, y_{jt} ,

is the effect from the treatment, $\alpha_{jt}I_{jt}$, and the counterfactual outcome, y_{jt}^J :

$$y_{jt} = \alpha_{jt}I_{jt} + y_{jt}^J. \quad (2.1)$$

The idea of the SCM is to construct a vector of weights over J donor countries such that their weighted combination mimics the pre-treatment outcome of the treated country. This weighted combination of donor units is called a synthetic Germany. Defining X_1 as the $k \times 1$ vector of the k characteristics of Germany in the pre-intervention period, and X_0 as the $k \times J$ vector with the same pre-treatment characteristics for donors, the SCM algorithm identifies non-negative donor weights \mathbf{W} , such that $\sum_{j=2}^{J+1} w_j = 1$, to minimize the divergence between pre-treatment characteristics \mathbf{X}_1 and \mathbf{X}_0 of the treated country and the untreated donors. More formally, the vector \mathbf{W}^* is chosen to minimize the mean square prediction error (MSPE) over k pre-treatment characteristics:

$$MSPE = \sum_{m=1}^k v_m (X_{1m} - X_{0m} \mathbf{W})^2, \quad (2.2)$$

where \mathbf{V} is a matrix of non-negative components measuring the relative importance of each predictor, v_m . Given optimal weights w_j^* for each $j = 2, \dots, J + 1$ donor country, the synthetic control at any time t is the weighted combination of the outcome variable (e.g., CO₂ emissions in the transport sector) in the donor countries, $\sum_{j=2}^{J+1} w_j^* y_{jt}$. The treatment effect α_{1t} is then the difference between emissions in the treated country y_{1t} and emissions in the synthetic counterfactual in the post-treatment period, $t > T_0$:

$$\hat{\alpha}_{1t} = y_{1t} - \sum_{j=2}^{J+1} w_j^* y_{jt}. \quad (2.3)$$

Choice of SCM predictors. There are various methods for choosing the relative importance of predictors (v_m) (Abadie and Gardeazabal, 2003; Abadie et al., 2010). The standard approach selects the matrix \mathbf{V} along weights \mathbf{W} to minimize the

⁶The average treatment effect is thus given by: $\hat{\beta}_{1T} = \frac{1}{T} \sum_{t=t_1}^T (y_{1t} - \sum_{j=2}^{J+1} w_j^* y_{jt})$.

pre-treatment difference between actual and synthetic Germany’s emissions, using the *synth* package in STATA by Abadie et al. (2010). Despite being a primarily data-driven approach, there is some discretion in specifying the SCM, which may lead to cherry picking combinations of predictors to influence the result (e.g., Ferman et al., 2020).⁷ Given a lack of consensus on how to choose the best specification, we report results for a range of specifications used in previous SCM evaluations (see Table 2.1).

Table 2.1: Overview of the specification choices for the SCMs

Specification	Lagged outcome variable	Selected literature examples
Baseline	Lagged outcome in 1998 (t_0)	Andersson, 2019; Kaul et al., 2015; Leroutier, 2022
Lags (Mean)	Pre-treatment outcome mean	Abadie and Gardeazabal, 2003; DeAngelo and Hansen, 2014
Lags (All)	Lagged pre-treatment outcome (t_0, t_{-1}, \dots, T_0)	Bohn et al., 2014; Dustmann et al., 2017; Isaksen, 2020
Lags (Selected)	Lagged outcome in 1971, 1980, 1991, 1998	Cavallo et al., 2013; Cunningham and Shah, 2018
Reunification	Lagged outcome in 1991 and 1998	<i>Specific to the German case</i> (cf. Abadie et al., 2015)
Tax anticipation	Lagged outcome in 1999 (t_1)	Abbring and Van den Berg, 2003; Coglianesi et al., 2017
No covariates	Lagged pre-treatment outcome (t_0, t_{-1}, \dots, T_0)	Gobillon and Magnac, 2016; Lindo and Packham, 2017

Notes: Summary of SCM specifications. *Specification* denotes the name that we use for SCM specification henceforth. *Lagged outcome variable* specifies the number and years of the pre-treatment outcome lags. All except *No Covariates* include as predictors (i) GDP per capita (PPP, in mio 2011 USD), (ii) gasoline and (iii) diesel consumption per capita, (iv) the share of the urban population, and (v) the number of vehicles per 1000 people. SCM specifications for NO_x emissions also include (vi) $\text{PM}_{2.5}$ emissions in the transport sector as a general proxy for air pollution to account for the impact of unilateral policies affecting emission levels. We refer to the specification used by Andersson (2019) as the *Baseline* model. We start the post-treatment period in 1999 even if the first fully treated year is 2000 to capture anticipation effects (cf. Section 2.A in the Appendix for details). Our *Tax anticipation* specification provides results when we set t_1 in the year 2000 for comparison.

Statistical inference for the SCM. A key advantage of the SCM is that it offers an approach to causal analysis that does not rely on parallel pre-intervention trends like difference in difference methods. Yet, it does not allow to employ standard (large-sample) inferential methods, primarily because the number of suitable donors and time periods are usually very limited. Abadie et al. (2010, 2015) and Abadie (2021) suggest using placebo experiments using permutation techniques to make inferences. We implement cross-sectional placebo tests by sequentially applying the SCM algorithm to every potential donor country and compare estimated placebo effects with the baseline results for Germany, after accounting for the quality of the pre-treatment match, which we do by scaling effects by the relevant pre-treatment root MSPE (RMSPE). Examining whether potential comparison countries show

⁷While Kaul et al. (2015) point out that including the entire pre-treatment periods of the outcome variable as a predictor causes all other covariates to be obsolete, Ferman et al. (2020) advise using all pre-treatment periods as it is less arbitrary.

larger treatment effects helps assess the robustness of our results. A p-value is then computed as the proportion of control units that have an estimated effect at least as large as Germany's. Suppose that the estimated standardized effect for some post-treatment period is $\hat{\alpha}_{1t}$ and that the distribution of in-place placebo is $\hat{\alpha}_{jt}^{PL} = \{\hat{\alpha}_{jt} : j \neq 1\}$, the one-sided and two-sided p-values are then given by:

$$p = Pr(\hat{\alpha}_{jt}^{PL} \geq \hat{\alpha}_{1t}) \quad \text{and} \quad p = Pr(\hat{\alpha}_{jt}^{PL} \leq \hat{\alpha}_{1t}), \quad (2.4)$$

$$p = Pr(|\hat{\alpha}_{jt}^{PL}| \geq |\hat{\alpha}_{1t}|) = \frac{\sum_{j \neq 1} 1(|\hat{\alpha}_{jt}^{PL}| \geq |\hat{\alpha}_{1t}|)}{J}. \quad (2.5)$$

Following Firpo and Possebom (2018) and Abadie and L'hour (2021), we implement a one-sided test, which allows constructing p-values based on placebo effects, $\hat{\alpha}_{jt}^{PL}$, that yield reductions in post-treatment emissions, as only reductions in emissions due to fuel taxes are of interest for the rank statistics of country-level treatment effects (we also report two-sided p-values). To evaluate how the significance of the effects unfolds over time—as the eco-tax rate increased in yearly steps from 1999 to 2003 (see Figure 2.3)—we apply the permutation-based inference procedure for each post-treatment year.

Generalized SCM with interactive fixed effects models. We additionally draw on GSCMs (Gobillon and Magnac, 2016; Xu, 2017) based on a linear interactive fixed effects (IFE) model (Bai, 2009). The GSCM expands the SCM in several dimensions (see Xu, 2017 for details on the methodology). First, the GSCM allows explicitly absorbing level differences and unobserved time-varying shocks specific to each country with IFE. Second, by including relevant control variables, our IFE model can control for time-varying covariates and explicitly capture heterogeneous influences of other policies across countries, such as the different effects of EU-wide emission standards on European economies and their emissions (cf. Bai, 2009). Third, the GSCM enhances the interpretability of SCM results by providing

uncertainty estimates conditional on observed covariates such as standard errors and confidence intervals to conduct statistical inference. A further advantage of the GSCM estimator is its built-in cross-validation scheme which automatically selects the model specification, limiting arbitrariness and reducing the risks of over-fitting.⁸

Synthetic difference-in-differences (SDID) with a staggered adoption design. To further examine the internal and external validity of our SCM results, we draw to a complementary research design that exploits the staggered introduction of similar environmental fuel taxes within other European transport sectors in our sample, for which causal effects on carbon emissions have been documented before: Finland and Sweden (Andersson, 2019; Mideksa, 2021). To this end, we employ the SDID methodology, which allows combining desirable features of a two-way fixed effects difference-in-differences (TWFE-DID) and the SCM in a staggered adoption setting. Specifically, the SDID estimator incorporates time and unit IFE within the regression function together with *unit-specific* weights to ensure closely matched pre-intervention trends as well as *time-specific* weights that reduce the influence of time periods that significantly differ from post-treatment periods (see Arkhangelsky et al., 2021 for details). This allows the SDID estimator to sidestep some of the typical issues encountered in standard DID and SCM applications, which include the inability to estimate causal relationships when parallel trends are not observed in aggregated data for DID, and the requirement for the treated unit to be located within a *convex hull* of control units in the case of SCM.⁹

In our setting, we harness the earlier introduction of carbon taxes in Finland in 1990 and Sweden in 1991 and leverage treatment status variation across multiple jurisdictions to estimate the causal effects of environmental taxation with a staggered adoption configuration.¹⁰ The key advantage of this design is that the additional cross-

⁸A key difference is that the GSCM employs dimension reduction before re-weighting implying that, unlike the standard SCM, weights cannot be directly interpreted.

⁹The SDID estimator also exhibits greater flexibility than the standard SCM by allowing for level differences between treatment and control groups.

¹⁰According to the World Bank (2022), the Finnish carbon tax was introduced in 1990 with an

country variation further helps mitigate the impact of contemporaneous confounding factors, such as potentially unobserved idiosyncratic external shocks unique to the German context in the post-treatment period which might confound our estimations, thereby enhancing the internal and external validity of our findings. Specifically, the staggered introductions thus offer a suited empirical setting for curbing the potential influence of concurrent alterations in policy, market dynamics, and societal preferences within Germany subsequent to the environmental tax reform.

2.2.2 Semi-elasticity models

We subsequently complement our causal inference methods by estimating price and tax elasticities of gasoline and diesel demand and use these to perform simulations to investigate how tax effectiveness is mediated by salience and fuel substitution using log-linear semi-elasticity models. We estimate fuel-specific elasticities, using two different specifications (cf. Andersson, 2019). First, we calculate real price elasticities and compare them to typical fuel demand elasticities (cf. Equation 2.6: *Real price elasticities*). Second, in line with Li et al. (2014), we split the real price into its three main elements: (i) the eco-tax, (ii) other existing fuel taxes (henceforth the energy tax), and (iii) the remaining tax-exclusive component, here called the raw price (cf. Equation 2.7: *Eco-tax elasticities*).

Real price and tax elasticity in Germany. To compare our findings to Andersson (2019), we first estimate a set of models based on variation in fuel demand within Germany and use the estimated elasticities from Equation 2.7 to simulate predicted pathways of CO₂ and air pollution emissions under different taxation regimes.¹¹ The

initial tax rate of \$1.75 metric ton/CO₂e and has steadily grown to \$27 metric ton/CO₂e by the end of our sample in 2009. In Sweden, the initial rate amounted to around \$41 metric ton/CO₂e in 1991 and rose to \$126 metric ton/CO₂e over the same period.

¹¹We refer to this specification as our *Baseline* model when discussing results in Section 2.5.

static log-linear models for Germany are expressed as:

$$\log(y_t) = \beta_0 + \varphi_1 p_t^{real} + \beta_2 D_t^{eco} + \lambda' \mathbf{X}_t + \epsilon_t \quad (2.6)$$

$$\log(y_t) = \beta_0 + \varphi_2 p_t^{excl} + \varphi_3 p_t^{eco} + \varphi_4 p_t^{energy} + \beta_2 D_t^{eco} + \lambda' \mathbf{X}_t + \epsilon_t \quad (2.7)$$

Elasticity estimates obtained leveraging annual data within a static model typically lie somewhere between short- and long-term elasticities, and are regarded as “intermediate” (Dahl and Sterner, 1991). Outcome y_t refers to log fuel consumption per capita for gasoline or diesel in liters.¹² p_t^{real} is the real retail price, including VAT. p_t^{excl} is the retail price excluding the energy and eco-tax but with VAT, in real terms. p_t^{eco} and p_t^{energy} refer to the eco and energy tax, respectively, including VAT and are included in the models as separate terms (cf. Equation 2.7). D_t^{eco} is a dummy equal to one after the implementation of the eco-tax and zero otherwise. \mathbf{X}_t is a vector of control variables that includes GDP per capita, the unemployment rate, and a time trend. The error terms are denoted by ϵ_t . We estimate the model using an OLS regression. As autocorrelation is detected, we use the Newey-West-estimator, which is robust against autocorrelation and heteroskedasticity.¹³

A standard concern with estimating fuel elasticities is an endogeneity problem, where fuel demand can also affect supply and thus prices (e.g., Kilian, 2009; Coglianese et al., 2017; Kilian and Zhou, 2023). Endogeneity due to reverse causality is arguably a lesser source of concern in a single EU country setting, as crude oil prices are set in a global market and changes in demand in a single country are thus expected to have a relatively marginal impact on overall demand. One possibility to address this issue is to adopt an instrumental variable (IV) approach. In line with Li et al. (2014) and Andersson (2019), we complement our OLS regressions with an IV approach and use the (Brent) crude oil price as an IV to validate the demand elasticities of the real price of gasoline and diesel.

¹²Prior to taking logs, we convert fuel consumption to liters.

¹³Standard errors are calculated using lags chosen following Newey and West (1994).

Fixed effects models with cross-country panel data. We further estimate a set of fuel-specific fixed effects models harnessing cross-country panel variation in fuel prices and tax rates to refine and validate our set of *Real* and *Eco-tax elasticities* for Germany. Crucially, the additional variation across jurisdictions in the estimation sample allows us to include a host of fixed effects to control more precisely for unobserved time-varying confounding factors. The resulting static log-linear fixed effects models are written as

$$\log(y_{it}) = \beta_0 + \varphi_1 p_{it}^{real} + \lambda' \mathbf{X}_{it} + \gamma_i + \psi_t + D_t^{eco} \times \phi_i + \epsilon_{it} \quad (2.8)$$

$$\log(y_{it}) = \beta_0 + \varphi_2 p_{it}^{excl} + \varphi_3 p_{it}^{eco} + \varphi_4 p_{it}^{energy} + \lambda' \mathbf{X}_{it} + \gamma_i + \psi_t + D_t^{eco} \times \phi_i + \epsilon_{it} \quad (2.9)$$

One key difference vis-a-vis Equation 2.6 and 2.7 is the inclusion of γ_i and ψ_t , which refer to country and time fixed effects, respectively. The former absorbs any time-invariant characteristics that might affect fuel demand in each country allowing us to focus on changes within countries over time whereas the latter captures common time trends that affect fuel demand across all countries in the same way (e.g., macroeconomic factors, technological advancements, or global demand changes). The models allow for spatial autocorrelation by clustering standard errors at the country-year level. We also now include our dummy indicator, D_t^{eco} , interacted with country-specific dummies, ϕ_i , to absorb any unobserved jurisdiction-level shocks affecting fuel demand after the implementation of the eco-tax (e.g., nationwide policies affecting fuel demand). Finally, we add an EU-specific time trend among a vector of cross-country control variables, \mathbf{X}_t , to account for common trends in fuel demand among European economies (e.g., EU-wide market and/or policy reforms).

2.3 Data

Our analysis is structured in two parts. In each step, we combine several data. First, we resort to three causal inference methods (SCM, GSCM, and SDID) to

evaluate effects of the eco-tax on CO₂, PM_{2.5} and NO_X emissions, and on low-carbon innovation, building on a panel dataset of OECD countries. Second, we perform complementary (non-causal) analyses on underlying mechanisms. To this end, we estimate price elasticities relying on a time-series dataset constructed specifically for Germany. We then examine the mechanism of tax salience in detail, relying on textual analysis of German newspapers. Table 2.A1 in the Appendix provides a detailed overview of all data sources used.

Emissions in the transport sector. To analyze the effect of the eco-tax reform on CO₂, PM_{2.5}, and NO_X emissions of the transport sector, we construct an annual panel dataset of OECD countries from 1971 to 2009. We obtain CO₂ emissions in metric tons by multiplying total CO₂ emissions from fuel combustion from the International Energy Agency (IEA) by the share of total fuel combustion for transportation. Annual emissions of PM_{2.5} (referring to both exhaust and non-exhaust) and NO_X are extracted from the Emission Database for Global Atmospheric Research (EDGAR) v6.1.^{14,15} GDP data refers to expenditure-side real GDP at current purchasing power parities (in million 2011 USD) from the Penn World Table. Data for population, the share of urban population and diesel and gasoline consumption per capita in kg of oil equivalent are from the World Bank, and the number of vehicles from Dargay et al. (2007).

We limit our dataset to OECD countries, as these share more structural similarities with Germany in terms of their economic situation, emissions, and form of government,

¹⁴We use EDGAR as this computes sectoral-level emissions based on the IPCC recommendations, which assured consistency in time, harmonized sector definitions, and thus comparability across countries. This has clear benefits over national emission inventories, which have two key caveats: (i) a much shorter pre-treatment period (only 1990 onwards) and (ii) methodological inconsistencies in officially-reported pollution data across time and countries, which may hinder direct comparability and increase measurement error. Weights computed with shorter pre-treatment periods (T_0), and outcome variables with substantial noise, may increase biases in SCM (Ferman and Pinto, 2021) and GSCM (Xu, 2017).

¹⁵Key benefits of looking at transport emissions directly include capturing changes in country-specific emission factors for technologies in the transport sector, as well as fuel demand adjustments made on the extensive margin in response to a fuel tax, such as substitution between modes of transport.

which is desirable for the SCM (Abadie, 2021). To build a suitable synthetic control for Germany, we exclude a number of countries. First, data for the Baltic countries, Costa Rica, Slovakia, Czech Republic, and Slovenia is very sparse (especially prior to 1989), which is why we cannot consistently use them for the SCM starting from 1971. Second, we exclude countries that have implemented an explicit CO₂ price in the transport sector. This concerns Finland, Sweden, Norway, and the Netherlands (Kossoy et al., 2015).¹⁶ As a number of countries implemented carbon taxes in the transport sector in 2009 or shortly thereafter, our analysis focuses on the time frame up to 2009.¹⁷ Third, we exclude countries that implemented fuel taxes in the transport sector that are not labeled as carbon taxes—similar to the eco-tax in Germany. This includes Italy, the UK (OECD, 2001), and Spain (Bosch, 2001). Fourth, we exclude Japan due to its very successful top runner program implemented in 1998 that set requirements for the fuel efficiency of vehicles (Osamu, 2012). Fifth, we exclude Ireland due to its exceptional economic growth in the 1990s. Finally, we exclude Austria and Luxembourg due to non-negligible fuel tourism.¹⁸ These restrictions, mostly due to carbon and fuel taxation, leave us with a main sample of 20 countries for the time frame from 1971 to 2009 that we use in our battery of causal inference approaches (SCM, GSCM, SDID).¹⁹

¹⁶Although Denmark implemented a carbon tax around the same time, it did not include the transport sector, which is why Denmark remains in the sample (Andersson, 2019). The same holds for Poland, as its carbon tax in cost per ton of CO₂ was negligible at a few cents (Kossoy et al., 2015).

¹⁷Other key rationales to end our sample in 2009 include the introduction of a vehicle circulation tax based on carbon emission rates implemented in Germany from July 2009 (Klier and Linn, 2015), the roll-out of low-emissions zones from 2008 onward (Wolff and Perry, 2010), stricter car emission performance standards for new car types introduced in 2009 (Reynaert, 2021), as well as the considerable collusion on emission control technologies by German carmakers thereafter (Ale-Chilet et al., 2021).

¹⁸Luxembourg’s fuel sales are 5 to 8 times higher per capita than those of the neighboring countries (Dings, 2004). Austria, too, has very low taxes with a tax minimum in 2005 and a downward trend from 1997 onwards. This is a contrast to tax increases in Germany and Italy in 1999. As a result, more fuel tourism has likely taken place and emission data is not reliable (Dings, 2004).

¹⁹An exception is that the GSCM also runs with unbalanced samples which allow us to include countries excluded due to sparse data in earlier periods (7 in total). For each method, we show that our findings are not significantly altered by these sample restrictions in the Appendix.

Low-carbon innovation: Patent data. To measure innovation, we use patent data from the OECD Patent Database. Patent documents are categorized into climate change mitigation patents in accordance with the Y02 tagging scheme of the Cooperative Patent Classification. We extract a panel dataset of climate change mitigation patents related to transportation (Y02T category) filed by inventors in OECD countries spanning from 1985 (earliest availability) to 2009. We focus on triadic patent families to improve the quality and the international comparability of patent counts.²⁰ Triadic patents are a sub-set of patents taken at the European Patent Office, the Japan Patent Office and the US Patent and Trademark Office that protect the same invention. Since only patents applied for in all three are included, we address concerns related to home advantage and the influence of geographical location. Moreover, triadic patents are generally of higher value: patentees only take on the additional costs and delays of extending protection to other countries if they deem it worthwhile (Aghion et al., 2016).²¹ Patents in our data are counted according to the earliest priority date, which corresponds to the first patent application worldwide and is, thus, closest to the invention date.

Consumption and real price of transport fuels. To estimate price and tax elasticities and disentangle the different taxation changes, we first construct an annual time-series dataset for Germany, spanning from 1971 to 2009.²² The data for the gasoline and diesel prices reflect yearly consumer prices for both fuels including VAT. We convert all nominal prices to real prices, including the energy and eco-tax rates and the strategic reserve component (the Appendix details data sources). As VAT is not only imposed on the tax-free price p but also on the eco and energy taxes,

²⁰We treat multiple application filings of an invention (i.e., a patent family) as one innovation. We focus on patent families to capture the number of low-carbon technologies that are developed in Germany rather than the count of underlying patent applications.

²¹Considering the number of jurisdictions in which a patent application is filed is a common approach to capture patent quality (e.g., Cabel and Dechezlepretre, 2016).

²²A peculiarity of Germany is its division until the year 1990. As there was no market economy in East Germany, there were no market prices and no taxes in the same sense as in West Germany. All prices that will be discussed in this chapter thus relate only to West Germany prior to 1991, while price data from 1991 onwards, and all fuel consumption data, reflects the entirety of Germany.

τ^{eco} and τ^{energy} , and the strategic reserve, τ^{sr} , in the same way as on the price, the retail price p^r can be defined as:

$$p^r = (p + \tau^{eco} + \tau^{energy} + \tau^{sr}) * (1 + VAT) \quad (2.10)$$

To account for this, VAT is already included in each retail price element.²³ All prices given in Deutsche Mark (DM) are converted to Euro, and all nominal prices and absolute tax rates into real 1995 values. We chose 1995 as a convenient base year close to the implementation of the eco-tax. Whenever a tax rate changed within a year, we weighted rates according to the date at which the change took place and used these average tax rates. The (Brent) crude oil price used for the IV regressions is from the IEA, converted from USD per barrel to €/l using the Eurostat (2020) €/USD exchange rate.

Second, we construct an annual panel dataset with country-level diesel and gasoline prices to expand our country-level time series. Data on gasoline (diesel) prices and taxes is consistently available from the IEA for 24 (19) major countries starting from 1978 onwards.²⁴ We harness the additional cross-country variation to estimate a set of fuel-specific price and tax semi-elasticity models which employ a host of fixed effects to control more precisely for unobserved time-varying confounding factors.

Salience: Newspaper data. We further examine the role of salience as a mediating mechanism of consumers' responses to the eco-tax. To this end, we rely on newspaper data as a reasonably representative proxy of tax salience within the media, as newspapers still reached the majority of the adult population in Germany in the mid-2000s. We extract information from the Factiva database, which stores all articles published by major newspapers either in their print or online format, and

²³If the eco-tax was raised by 10 cents, the fuel price would increase by 11.90 cents with a VAT rate of 19%. Thus, the eco and energy tax rates include VAT. In our calculations, the price increase is attributed to a change in the eco-tax rate.

²⁴The difference in the number of jurisdictions covered by the IEA data for gasoline and diesel prices stems from the lack of a sizable market for diesel in a number of jurisdictions.

use this to develop a newspaper-based index to capture the evolution of salience of the eco-tax-inclusive price based on textual analysis of newspaper articles (Gentzkow et al., 2019). We restrict our analysis to the four largest newspapers retrievable from Factiva, as relying on a single source provides consistent, comparable, and thus more robust counts. While, as a consequence, our text analysis does not capture Germany’s largest tabloid (*BILD*), it captures a representative account of newspaper salience in four of the largest nationwide newspapers from both the center-left (*Der Spiegel* and *Die Zeit*) and center-right spectrum (*Die Welt* and *Focus*).

We designed a text-based search strategy to identify newspaper articles that specifically discuss the repercussions of the eco-tax on the price increase of transport fuels. We prefer this approach vis-a-vis focusing on articles that more broadly discuss the ecological tax reform to ensure that our index more accurately proxies an indicator of eco-tax price salience. Our main salience index is thus constructed using the number of articles published in leading German national newspapers after 1991—as there was no unified press prior to German reunification—that discuss the effects of the eco-tax on fuel prices scaled by newspaper-specific publishing trends specific to the topic of environmental taxation. We scale our frequency counts to ensure that spikes in our index are not driven by newspaper-specific trends in reporting of environmental issues, which has experienced steadily growing attention in the German public media (Schmidt et al., 2013), and may reflect changes in editorial priorities over time. To obtain newspaper article counts, we rely on a set of text-based search strategies that identify around 5,700 unique articles. After scaling the raw counts, we standardize each newspaper’s series, average across all papers, and normalize the resulting index to 100 over the period. We follow the same standardization and normalization procedure by Baker et al. (2016) to leverage newspaper data in an empirical setting. A description of our search strategies and the steps to construct newspaper-based indices is detailed in Section 2.E.2 of the Appendix.

2.4 Results from Causal Inference Methods

This section presents how we leverage our battery of causal inference methods (SCM, GSCM, SDID) described in Section 2.2.2 to estimate the impact of the eco-tax on carbon and local air pollution emissions within the German transport sector. We, first, focus on examining emission reductions according to the SCM, performing inference using permutation tests and assessing their robustness using standard sensitivity and placebo tests (cf. Section 2.4.1) before turning to the GSCM (cf. Section 2.4.2). Subsequently, we use SCM and GSCM to examine impacts on low-carbon innovation (cf. Section 2.4.3). Finally, we use SDID to examine the internal and external validity of the SCM and GSCM findings on both emissions reductions and low-carbon innovation (cf. Section 2.4.4). Additional supporting evidence is available in Section 2.B of the Appendix.

2.4.1 Emissions reductions according to SCMs

Panels (a), (c) and (e) in Figure 2.1 plot the path of CO_2 , $\text{PM}_{2.5}$ and NO_X emissions in the German transport sector (solid line) and in synthetic Germanies (dashed lines) across specifications (cf. Table 2.1) from 1971 to 2009. The overlap between the solid and dashed line before 1999 captures the quality of the pre-treatment fit achieved by the SCM; the same comparison after 1999 plots the dynamic treatment effects for the eleven years that followed. All panels reveal sizable emission reductions following the eco-tax reform.

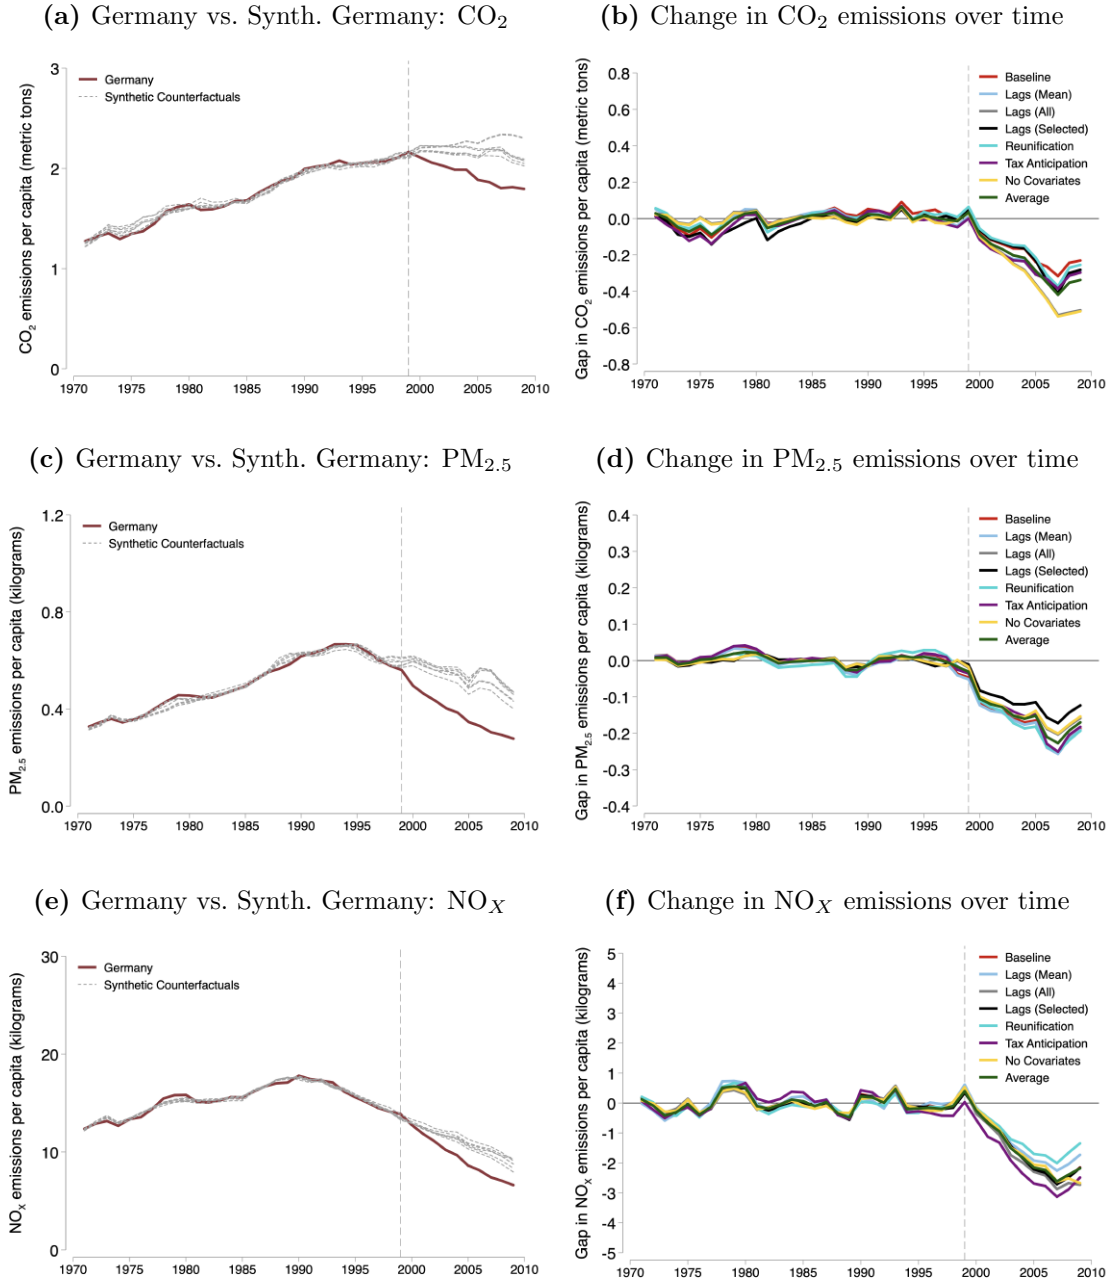
The validity of SCM effects depends on synthetic Germany's ability to replicate emissions from the German transportation sector prior to the eco-tax introduction. Panels (a) and (b) show that prior to the treatment, emissions from transportation in Germany and its synthetic counterpart exhibit a high degree of similarity, with an average absolute difference of slightly more than 0.02 metric tons of CO_2 , less

than 0.01 kg of PM_{2.5} and around 0.22 kg of NO_X. Figure 2.A5 in the Appendix plots the distribution of country-specific weights across all specifications and shows that the composition of our synthetic Germanies varies considerably across outcomes and specifications. Tables 2.A2 - 2.A4 in the Appendix compares the values of key predictors for Germany prior to 1999 with those for our baseline synthetic Germany (cf. Section 2.3). Overall, synthetic Germany exhibits a much more refined fit compared with the donor pool average.

Panels (b), (d) and (e) of Figure 2.1 report the estimated gap in metric tons of CO₂ and kg of PM_{2.5} and NO_X emissions across the seven SCM specifications (colored lines), where *Average* refers to the average estimated emission gap (green line). All specifications point to sizable decreases in CO₂, PM_{2.5} and NO_X emissions in the transport sector following the eco-tax reform. Panel (b) shows that the distance between Germany and the synthetic Germanies is steadily growing between 1999 and 2007.²⁵ In 2007, this distance was on average -0.42 metric tons of CO₂ per capita, equivalent to a 19 percent reduction. Between 1999 and 2009, annual emission reduction amounted to 0.23 metric tons of CO₂ per capita on average, which cumulatively sums up to 208,216,572 tons of CO₂. Panel (d) presents the emission gap over time for PM_{2.5}. On average, 0.15 kg of per capita PM_{2.5} less were emitted each year in comparison to a scenario with no eco-tax, which amounts to total PM_{2.5} savings of around 135,632 tons. Finally, Panel (f) displays emission gaps for NO_X. Following the eco-tax reform, per capita NO_X emissions were lower by 1.5 kg, on average, with a cumulative reduction in NO_X of 1,347,190 tons.


²⁵There are different possible explanations for the convergence in emissions after 2007. An obvious one is the financial crisis, which evolved into an economic crisis across the EU in 2008, which likely affected German transport differently than that of donor countries, implying that synthetic Germany may not describe the counterfactual after 2007/2008 as accurately as before. Another explanation is decreasing fuel taxes in real terms. As the last increase of the eco-tax took place in 2003, the real fuel tax on gasoline and diesel has been decreasing ever since then due to inflation.

Figure 2.1: Synthetic Control Method results for emissions



Notes: The figure plots the estimated reductions in CO₂, PM_{2.5} and NO_x emissions relative to (synthetic) counterfactuals. Panels (a) and (b) refer to reductions in CO₂ emissions per capita in metric tons or percentage terms (as indicated on the respective y-axis). Panels (c) - (f) refer to reductions in PM_{2.5} and NO_x emissions per capita expressed either in kg. Panels (a), (c) and (e) plot the absolute paths of emissions in Germany and Synthetic Germanies for our specifications (see Table 2.1). Panels (b), (d) and (f) report gaps in emissions over time relative to synthetic Germanies, estimated by our seven different SCM specifications and their average.

Figure 2.2: Mean annual percentage gap in CO₂, PM_{2.5} and NO_X emissions




Notes: The figure plots the average annual percentage gap for each specification in CO₂, PM_{2.5} and NO_X emissions between Germany and a synthetic counterfactual development reported in Figure 2.1.

Figure 2.2 plots mean annual changes in emissions in percentage terms to put into perspective the distribution of the effect magnitudes from different specifications. CO₂ per capita emissions decrease, on average, from 8.1% to 13.4% between 1999 and 2009, conditional on the specification used, while PM_{2.5} and NO_X per capita emission reductions range between 22.4% - 30.3% and 10% - 16.5%, respectively. Our finding that emission reductions due to the eco-tax are sizable is thus robust across a range of specifications.

Inference from permutation tests for the SCM

We rely on permutation tests to gauge the significance of our treatment effects. Figure 2.3 plots estimated one-sided p-values in each post-intervention year. We report yearly permutations for a number of SCM specifications: (i) *Baseline*, (ii) *Baseline* restricting the pre-intervention period after German reunification in 1991, (iii) *Tax Anticipation*, and (iv) *No covariates* following Ferman et al. (2020). Overall, the distribution of the estimated p-values is centered well below a 10% threshold level, and generally below a 5% threshold, particularly after the last eco-tax rate increase in 2003. The mean joint two-sided p-values are below 5% for CO₂ and NO_X and below 1% for PM_{2.5} (see Figure 2.3).

Figure 2.3: Inference results for the Synthetic Control Method



Notes: The figure plots estimated one-sided p-values (primary left-hand side y-axis) computed as the proportion of effects from control units as great as the treated unit in each post-intervention period, after scaling it by the relevant pre-treatment RMSPE (Abadie, 2021). Joint two-sided p-values represent the proportion of placebos that have a ratio of post-treatment RMSPE over pre-treatment RMSPE at least as large as the average ratio for Germany. The gray bars plot the annual real eco-tax rate in 1995 cents (secondary right-hand side y-axis). The darker gray bars indicate the post-treatment periods where the full nominal eco-tax rate increase fuel was in place.

Additional sensitivity and placebo tests for the SCM

Our findings are robust to a host of standard sensitivity and placebo tests, including in-time placebos and the use of alternative donor pools.

In-time tests. For the in-time placebos, the year of treatment is shifted to a number of years prior to the actual eco-tax reform. Any sizable and enduring placebo effect would cast doubt on the validity of the results from Figure 2.1. Figure 2.A6 in the Appendix shows that the synthetic control closely resembles the actual emission trajectories in Germany after the placebo treatment and that no significant divergence is detected.

Alternative donor pools. To investigate the sensitivity of our emission results to the composition of the donor pool, we perform the following tests: (i) implementing the SCM without any sample restriction either with the inclusion of covariates as predictors or solely based pre-treatment lags, (ii) excluding only countries that

implemented carbon taxes, and (iii) “leave-one-out” tests (cf., Abadie et al., 2015), where we sequentially exclude from the restricted donor pool all control countries with a weight larger than 0.001 (0.1%). The results (see Figures 2.A7 and 2.A8 in the Appendix) show that none of the possible alternative donor pool compositions yield a consistent non-negative post-intervention gap.

2.4.2 Emissions reductions according to generalized SCMs

We next construct GSCM counterfactuals by modeling emissions of countries with interactive fixed effects (IFE) models. First, we include controls to explicitly account for the impacts of EU membership, namely a binary EU member indicator and a dummy identifying EU member countries after 2005 (denoted *IFE only*).²⁶ Second, we additionally model each country’s emissions as a function of their economic activity (*Economic activity*), proxied by GDP per capita (Bayer and Aklin, 2020). Finally, we restrict the donor pool to EU countries (*EU only*) to further address concerns that effects may be partly driven by EU-wide regulation, such as emission standards (e.g., Reynaert, 2021). The inclusion of IFE crucially curbs the heterogeneous influence of other nationwide policies affecting emissions in our setting. Wald tests for pre-treatment fitting checks show that all the different models capture the variability in the data well prior to the eco-tax reform, validating the main identification assumption. Table 2.2 summarizes our GSCM results. We report mean reductions of emissions due to the eco-tax with bootstrapped 95% confidence intervals. Our GSCM results are comparable in magnitude to the average SCM results reported in Figure 2.1, pointing towards slightly larger magnitudes in carbon reductions and

²⁶We include this dummy to control for potential spillovers due to the EU Emissions Trading Scheme (EU ETS), introduced in 2005, and the EU-wide PM₁₀ limits in cities, also introduced in 2005. These spillovers are likely not substantial, as transport emissions were not covered by the EU ETS and have not decreased due to the scheme (Bayer and Aklin, 2020). Further, Germany failed to meet the 2005 PM₁₀ limits, triggering infringement proceedings in 2009, and EU-wide PM₁₀ limits on were not very effective initially, with 70% of all cities with larger populations than 250,000 having exceeded the limits at some point as of 2007 (Wolff and Perry, 2010). Some German municipalities responded by implementing low-emission zones from 2008 onward, limiting access for highly-polluting vehicles within city centers.

Table 2.2: Effects of the eco-tax with a Generalized Synthetic Control

	IFE only	Economic activity	EU only
Panel A: CO₂ (t)			
Mean [95% CI]	-0.43 [-0.53; -0.34]	-0.39 [-0.50; -0.25]	-0.44 [-0.57; -0.29]
Panel B: PM_{2.5} (kg)			
Mean [95% CI]	-0.15 [-0.26; -0.04]	-0.14 [-0.25; -0.07]	-0.21 [-0.27; -0.13]
Panel C: NO_x (kg)			
Mean [95% CI]	-1.98 [-3.32; -0.24]	-1.65 [-3.09; -0.14]	-3.34 [-5.33; -0.26]
Observations	1053	939	451
Countries	27	27	14
Wald test p-value	<0.001	<0.001	<0.001

Notes: Summary of average treatment effects and 95% confidence intervals for different GSCM specifications. Wald test p-values refer to pre-treatment fitting tests (cf. Xu, 2017): for each specification, we report the highest p-values across panels. All models include IFE and a binary indicator for German reunification and post-intervention. *IFE only* includes a dummy for EU membership and for EU member countries after 2005 *Economic activity* also controls for GDP per capita, while *EU only* restricts the donor pool to EU countries.

almost identical average reductions in air pollution.²⁷

2.4.3 Impacts on low-carbon innovation for SCM and GSCM

In this section, we provide complementary evidence on the role of the eco-tax in spurring low-carbon innovation leveraging again the SCM and GSCM. Figure 2.4 plots the estimated gaps in low-carbon patents per million of the population over time for both methods.²⁸ In this context, the GSCM conveniently enables absorbing level differences in the determinants of innovative behavior across time with IFE, such as differences in government policies and support measures and human capital, but also other key unobservable determinants of innovation such as cultural attitudes toward risk-taking and entrepreneurship, making it a suited approach to quantify innovation effects.²⁹ Here, we can see that incorporating IFE translates into a superior pre-treatment fit yielded by the GSCM vis-a-vis the standard SCM, which

²⁷Table 2.A5 in the Appendix shows that our results are not affected by sample trimming. Additionally, Figure 2.A13 in the Appendix compares the dynamic treatment effects across all our different empirical strategies.

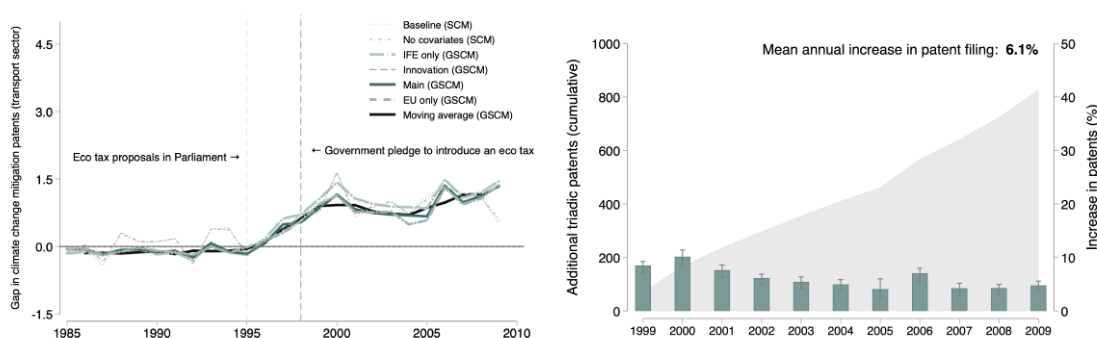
²⁸We further report a smoothed specification using a three-year moving average to account for the fluctuating nature of patent data (Griliches, 1990).

²⁹This differs from the case of air emissions which are linearly associated with fuel consumption.

Figure 2.4: Effects of the eco-tax on low-carbon patented technologies

(a) Germany vs. Synthetic Germany

(b) Change in patents over time



Notes: Our *Baseline* SCM includes total triadic patents per capita, the share of climate change mitigation patents related to transportation and a pre-treatment lag as predictors whereas *No covariates* is only based on pre-treatment lags. Our different GSCM specifications with IFE include different sets of controls. *IFE only:* (i) a binary variable identifying EU countries and (ii) a binary variable indicating whether a country was regulated by EU-wide regulations after 2005. *Innovation* further accounts for: (iii) total triadic patents per capita and (iv) share of climate change mitigation patents related to transportation. Our *Main* specification additionally controls for: (v) GDP per capita. *EU only:* estimates our *Main* specification restricting the sample to countries in the EU. *Moving average:* estimates our *Main* specification relying on a 3-year moving average instead of annual patent counts. Panel (b) refers to our *Main* specification. Percentage increases are computed as the estimated increase of triadic patents induced by the eco-tax scaled by the annual number of climate change mitigation patents related to transportation in Germany.

serves to derive unbiased and robust estimates. The extended pre-intervention period allows us to account for anticipatory behavior in the years leading up to the eco-tax reform (e.g., Lemoine, 2017). Overall, we observe that all our different specifications point to a sizable increase in low-carbon innovation following the eco-tax reform with an additional 0.91 patents per million population each year.³⁰ We also detect some anticipatory innovation responses starting after the parliamentary debate on the eco-tax reform first gained momentum in 1995 (Beuermann and Santarius, 2006).

Figure 2.4 presents effects in percentage and cumulative terms. The gray area plots the cumulative number of additional patents throughout the post-intervention period, while the green bar charts show the annual percentage increase in patents induced by the eco-tax with bootstrapped 95% confidence intervals. On average, between 1999 and 2009, the eco-tax induced an annual increase of 6.1% in carbon mitigation patents related to transportation, which cumulatively resulted in 826 additional patented technologies vis-a-vis a scenario without the eco-tax. This finding complements previous studies on the innovation response of regulated companies to carbon pricing

³⁰Our results remain unchanged when we do not trim the sample (cf. Figure 2.A10 in the Appendix).

schemes (e.g., Calel and Dechezlepretre, 2016), which generally find limited aggregate effects.³¹

2.4.4 Staggered treatment adoption with an SDID estimator

Finally, we present our complementary results from the staggered treatment adoption design described in Section 2.2.1 to address the concern that contemporaneous post-treatment trends could confound our SCM estimations for Germany, and to investigate the external validity of our German case study. Table 2.3 displays the average treatment effects of environmental fuel taxes on air emissions and low-carbon patenting harnessing their gradual rollout in Finland (in 1990), Sweden (in 1991), and Germany (in 1999). We also report average effects by country to assess heterogeneity. Additional graphical evidence on the dynamic unfolding of effects by jurisdiction and the pre-treatment time-weights employed in the SDID estimations is available in Figures 2.A15-2.A17 in the Appendix.

Overall, the average SDID estimates for emission reductions are comparable in magnitude to the SCM (cf. Figure 2.1) and GSCM results (cf. Table 2.2) for Germany, with larger effects for nitrogen dioxides driven by Finland.³² The estimated average effects on low-carbon innovation for the three countries point to a yearly increase of 0.6 patents per million population. This figure is around half in size compared to the point estimate for Germany of 1.1, which is instead more comparable to the SCM and GSCM estimate ranging between 0.9 and 1 (cf. Figure 2.4). Given Germany's

³¹Two key differences may explain our larger magnitudes. First, employing an economy-wide approach can additionally capture innovation occurring along the supply chain and across unregulated agents, due to pass-through of regulatory costs or knowledge spillovers (Popp, 2019). Second, innovation in the automotive industry is arguably of larger importance in Germany than in other countries that do not feature comparatively large automobile industries.

³²The larger NO_x effects for Finland despite its lower tax rate may be, partly, explained by two factors: (i) a period of economic recession during the 1990s which had deep and persistent repercussions on the Finnish economy and labor market, plausibly affecting fuel demand (cf. Mideksa, 2021) and (ii) a concurrent nationwide tax relief scheme on the registration of cars with catalytic converters which might have contributed to shape and accelerate the fleet renewal rate towards less-polluting vehicles (Ministry of the Environment, 1995). By contrast, Germany and Sweden did not levy registration taxes on new vehicles (aside from VAT) but imposed considerably higher and comparable fuel taxes (ACEA, 2022). Results based on Germany and Sweden only are presented in Table 2.A13 in the Appendix.

record of innovative culture and prominent market share in the automobile sector, it is unsurprising that the average SDID estimate for the three countries indicates a more modest impact of environmental fuel taxes on innovation relative to Germany-specific estimates. Accordingly, we attribute greater external validity to our average estimate, as it is better poised to mitigate the influence of the idiosyncratic features of Germany’s innovative behavior within the transportation sector. Our calculations suggest that the early introduction of implicit carbon pricing in Finland and Sweden cumulatively translated into 134 additional low-carbon patented technologies vis-a-vis a counterfactual scenario.

On the whole, our SDID cross-country findings corroborate the internal validity of our synthetic control results of the German environmental tax reform on air emissions and low-carbon patenting. Furthermore, while individual magnitudes differ across countries in plausible ways, the SDID analysis also provides evidence of the external validity of our synthetic control results on broadly comparable effects in terms of emissions reductions, and with qualitatively robust but smaller effects on low-carbon innovation.

Table 2.3: Effects of environmental taxes with a SDID staggered adoption design

	CO ₂ emissions (t)	PM _{2.5} emissions (kg)	NO _x emissions (kg)	Low-carbon patents
Environmental fuel taxation	-0.24*** (0.05)	-0.10*** (0.03)	-2.77*** (0.97)	0.64** (0.29)
Observations	858	858	858	550
Countries	22	22	22	22
Average point estimates by country				
Finland	-0.24	-0.10	-4.19	0.15
Sweden	-0.21	-0.07	-1.89	0.71
Germany	-0.32	-0.14	-1.75	1.14
German SCM effects for comparison				
Germany (SCM)	-0.23	-0.15	-1.50	0.99
Germany (GSCM)	-0.39	-0.14	-1.65	0.91

Notes: All outcome variables are expressed in per capita terms. Patents are expressed in per million population terms. Standard errors are computed using the bootstrap variance estimation algorithm outlined in Arkhangelsky et al. (2021) based on multiple treated units. All regressions include unit-specific and time-specific fixed effects and control for GDP per capita and a binary variable indicating whether a country was regulated by EU-wide regulations after 2005. Results with additional covariates can be found in Table 2.A14 of the Appendix. Germany (SCM) is based on our Average SCM results (cf. Figure 2.1) whereas Germany (GSCM) refers to our *Economic Activity* GSCM specification (cf. Table 2.2).

2.5 Results on Fuel and Tax Elasticities

This section leverages the semi-elasticity models described in Section 2.2.2 to disentangle the effects of the eco-tax, the energy tax, and VAT in order to compare behavioral responses from changes to the eco-tax rate and equivalent fuel real price changes. We report elasticities leveraging both national time series and cross-country panel variation.

2.5.1 Real price semi-elasticities for gasoline and diesel

Table 2.4 reports estimates from *Real price elasticities* specifications (cf. Section 2.2.2) for gasoline (left panel) and diesel (right panel) consumption.³³ Using our estimate from column (1) in the left panel (cf., Table 2.4), we derive a real price elasticity of gasoline of -0.54.³⁴ The IV regression yields a very similar price elasticity of -0.50 (column (2) of Table 2.4), indicating that the endogeneity of gasoline prices is likely not a major concern in our setting. To test the instrument's relevance condition, we use an F-test for that single instrument. For the price of gasoline, the F-statistic is 69.47 suggesting that the relevance condition is fulfilled and that Brent crude oil price can be considered a suitable instrument for gasoline prices. The cross-country elasticity based on OECD data presented in column (3) also yields a very similar elasticity of 0.50. The right panel of Table 2.4 displays results for diesel consumption from the real price elasticity specification (cf. Section 2.2.2). The real price elasticity of demand for diesel shown in column (1) of Table 2.4 is somewhat lower than for gasoline at -0.34. The IV regression in column (2) yields an estimate of -0.28, which deviates slightly more than the IV and OLS regressions for gasoline, but is still

³³See Tables 2.A8 and 2.A9 in the Appendix for robustness results using a shorter time frame (1991–2009).

³⁴To calculate elasticities from our log-level model estimates ($\log(Y) = a + bX$), the coefficient for each tax is multiplied by the average sample mean of the real fuel price (90 cents for gasoline and 76 cents for diesel), as the elasticity of demand is given by $\epsilon = \frac{dY}{dX} * \frac{X}{Y}$. This implies that $\frac{dY}{dX} = be^a e^{bX}$. Plugging in, we obtain $\epsilon = \frac{be^a e^{bX}}{e^a e^{bX}} * X = bX$.

Table 2.4: Real price semi-elasticities for transport fuels

	(1) OLS: Baseline	(2) IV: Brent crude	(3) OLS: Fixed effects		(1) OLS: Baseline	(2) IV: Brent crude	(3) OLS: Fixed effects
Gasoline Price	-0.00603** (0.00278)	-0.00553* (0.00305)	-0.00553*** (0.00103)	Diesel Price	-0.00440*** (0.00103)	-0.00361*** (0.000856)	-0.00454*** (0.00147)
Instrument F-statistic		69.47		Instrument F-statistic		168.86	
Price elasticity	0.54	0.50	0.50	Price elasticity	0.34	0.28	0.34
Sample	Germany	Germany	OECD	Sample	Germany	Germany	OECD
Controls	✓	✓	✓	Controls	✓	✓	✓
Observations	38	38	765	Observations	39	39	574

Notes: The dependent variable is the log of fuel consumption in liters per capita, which refers to total fuel consumption or either gasoline or diesel consumption (as indicated by the column heading). Columns (2) use the Brent crude oil price as an instrumental variable for the real fuel price. Results for gasoline consumption refer to 1972-2009 in column (1) due to missing price data prior to 1972. Newey-West standard errors in parentheses are heteroskedasticity and autocorrelation robust. Standard errors in columns (1) - (2) are calculated relying on the automatic bandwidth selection procedure following Newey and West (1994). Standard errors in column (3) are clustered at the country-year level * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

sufficiently close to corroborate the magnitude of the real price elasticity for diesel. Once again, our real price elasticity for Germany is almost identical to what we yield in column (3) harnessing cross-country variation across OECD jurisdictions, indicating that Germany does not display distinctive fuel price elasticity patterns compared to other OECD countries, reinforcing the broader applicability of our results. Overall, our estimates fall into the range of price elasticities of demand in the literature (e.g., Frondel and Vance, 2014).

2.5.2 Tax semi-elasticities for gasoline and diesel

The left panel of Table 2.5 displays results for gasoline consumption from the *Eco-tax elasticities* specifications (cf. Section 2.2.2).³⁵ The OLS results in column (1) in Table 2.5 indicate that the price elasticity of demand for the price excluding the energy and the eco-tax (but including the VAT) is -0.32. The energy tax elasticity of demand, instead, amounts to -0.22. Both elasticities are computed relying on coefficients that exhibit a considerably lower significance. This contrasts the eco-tax elasticity of demand, which is estimated at -2.7 and is thus around 8.5 times larger than the tax-exclusive price elasticity. The eco-tax elasticity of diesel demand is also significantly higher than that for the real price. The right panel of Table 2.5 displays the results for the different tax rates for diesel. Using column (1) in Table 2.5, the elasticity for the real price, excluding the energy and eco-tax, is -0.26. The energy

³⁵We cannot reject the hypothesis of full pass-through, see Section 2.C in the Appendix.

Table 2.5: Eco-tax semi-elasticities for transport fuels

	(1)	(2)	(3)		(1)	(2)	(3)
	OLS: Baseline	OLS: Fixed effects	OLS: Fixed effects		OLS: Baseline	OLS: Fixed effects	OLS: Fixed effects
Raw price of Gasoline	-0.00357* (0.00179)	-0.00256 (0.00165)	-0.00427** (0.00163)	Raw price of Diesel	-0.00346*** (0.00104)	-0.00525*** (0.00164)	-0.00506** (0.00187)
Energy Tax on Gasoline	-0.00242 (0.00476)	-0.00485*** (0.00128)	-0.00413*** (0.000569)	Energy Tax on Diesel	-0.00729** (0.00292)	0.000388 (0.00152)	0.000937 (0.00155)
Eco-tax on Gasoline	-0.0306*** (0.00700)	-0.0296*** (0.00479)	-0.0247*** (0.00350)	Eco-tax on Diesel	-0.0143*** (0.00359)	-0.0232*** (0.00351)	-0.0197*** (0.00328)
Raw price = Eco-tax (p-value)	<0.001	<0.001	<0.001	Raw price = Eco-tax (p-value)	<0.001	<0.001	<0.001
Eco-tax elasticity	2.7	2.7	2.2	Eco-tax elasticity	1.1	1.7	1.5
Sample	Germany	OECD	EU	Sample	Germany	OECD	EU
Controls	✓	✓	✓	Controls	✓	✓	✓
Observations	38	765	509	Observations	39	574	415

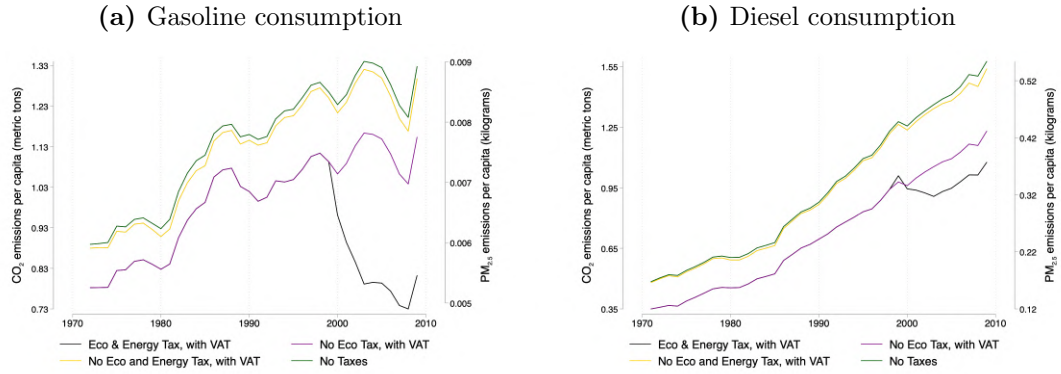
Notes: The dependent variable is the log of fuel consumption in liters per capita of either gasoline or diesel consumption (as indicated by the column heading). Results for gasoline consumption in column (1) refer to 1972-2009 due to missing price data prior to 1972. Newey-West standard errors in parentheses are heteroskedasticity and autocorrelation robust. Standard errors in column (1) are calculated relying on the automatic bandwidth selection procedure following Newey and West (1994). Standard errors in columns (2) - (3) are clustered at the country-year level * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

tax elasticity of demand is -0.56, slightly higher than the price elasticity. The eco-tax elasticity is again the highest level at -1.1, about 4 times larger than the tax-exclusive price elasticity. These magnitudes and the corresponding *tax saliency ratios* are corroborated by our set of fixed effects models both when leveraging panel variation across OECD jurisdictions in column (2) and restricting the sample to EU members only in column (3). It follows that an increase in the eco-tax predicts a stronger response in demand than that of a market-driven price change for both gasoline and diesel.³⁶

Li et al. (2014) discuss two underlying reasons that would reconcile our findings and explain the estimated stronger response to the eco-tax. The first one is *persistence*, meaning that consumers rely on tax changes to build expectations for the future price of gasoline. A tax increase may thus be perceived as more enduring than market-driven price fluctuations, which, in turn, would stimulate a stronger consumer response. The second is *saliency*, meaning that consumers are more aware of the price increase due to media coverage. We investigate the role of greater media tax salience in driving behavioral responses to changes in the eco-tax in Section 2.5.4.

³⁶We additionally amend our semi-elasticity models with a lead to test whether consumers increased their purchases of transport fuel in anticipation of tax increases, which could potentially bias estimated price and eco-tax coefficients (Coglianese et al., 2017). We do not find evidence of a potential anticipatory effect, and the estimated real price and eco-tax elasticities are very similar to the main result reported in Tables 2.4 - 2.5 (see Figure 2.A11 in the Appendix). One explanation is that anticipatory behavior is a lesser source of concern when dealing with yearly data as compared to relying on monthly variation.

Figure 2.5: Predicted emissions by fuel under different taxation scenarios



Notes: The figures plot predicted emissions from the eco-tax specification of our log-level semi-elasticity models (cf. Section 2.2.2) under different taxation scenarios. We rely on the estimated fuel-specific price and tax elasticities computed from our estimates from column (1) in Table 2.5. Panel (a) refers to predicted emissions from gasoline consumption, while Panel (b) covers diesel. In each panel, the left-hand side primary y-axis refers to per capita CO₂ in metric tons, while the right-hand side secondary y-axis refers to per capita PM_{2.5} in kg. The top green line displays predicted emissions when the eco and energy tax elasticities are set to zero, and VAT is deducted from the fuel price. For the yellow line, the eco and energy tax elasticities are set to zero but VAT is included. The purple line shows how predicted emissions change when the eco-tax is set to zero, but we include the energy tax and VAT. The black line provides predicted emissions using the full model with differentiated tax and price elasticities. The corresponding simulations for NO_x emissions can be found in Figure 2.A12 in the Appendix.

2.5.3 Emission scenarios and underlying mechanisms

We next rely on fuel-specific price and tax elasticities estimates from columns (1) in Table 2.5 to predict CO₂ and PM_{2.5} (and NO_x) emissions for different taxation scenarios, namely a scenario where no VAT and no taxes are introduced, a scenario where either VAT or VAT and the energy tax is added to the price of fuels, and, finally, a scenario where all are implemented.³⁷ We refer to this as the *Simulation Approach*.

Predicted emissions in the Simulation Approach. Panels (a) and (b) in Figure 2.5 summarize the estimated evolution of CO₂ (left-hand side primary y-axis) and PM_{2.5} (right-hand side secondary y-axis) emissions by fuel in the German transport sector under different tax regimes. The black line represents predicted

³⁷The combustion of one liter gasoline (diesel) emits 2.235kg (2.66kg) of CO₂ (US EPA, 2005). Using this factor, the predicted log gasoline (diesel) consumption values can first be turned into liters and then CO₂ emissions. To estimate PM_{2.5} exhaust emissions from fuel consumption, we rely on average emission factors by the European Environment Agency (EEA) for gasoline (diesel) vehicles in Germany (Ntziachristos and Samaras, 2019) of 0.02 grams (1.12 grams) of PM_{2.5} per kg of gasoline (diesel). Although EEA only reports emission factors for PM without specifying the size range, it clarifies that PM mass emissions in vehicle exhaust mainly fall in the PM_{2.5} category.

emissions accounting for all existing tax measures, including the eco-tax, energy tax, and VAT. The purple line plots the estimated evolution of emissions in the absence of the eco-tax, while the yellow line depicts the expected path of emissions with neither the eco-tax nor the energy tax, solely incorporating VAT. The green line shows predicted emissions without any tax policies. The gap between the black and purple line highlights the estimated emission gap attributable to the eco-tax, while the other lines represent alternative counterfactuals.

Panel (a) in Figure 2.5 shows that, between the years 1999 and 2009, the decrease in emissions of CO₂ (PM_{2.5}) from gasoline induced by the eco-tax was around 0.27 tons (0.002 kg) per capita on average per year. Similarly, Panel (b) provides the estimated emission reductions for diesel. The corresponding mean decline in annual emissions of CO₂ (PM_{2.5}) from diesel induced by the eco-tax was around 0.11 tons (0.04 kg) per capita, i.e. less marked than for gasoline due to the lower eco-tax elasticity for diesel.³⁸


Panels (a) and (b) in Figure 2.6 contrast the estimated share of aggregate reductions in emissions attributable to contractions in either gasoline or diesel use for CO₂ and PM_{2.5}, additionally including reductions in NO_x emissions. On average across our time frame, contractions in gasoline (diesel) use were responsible for around 72% (28%) of overall reductions in CO₂ emissions. Conversely, reduced diesel use is responsible for almost the entirety (95%) of the reduction of PM_{2.5} emissions. In other words, on average, reductions in diesel consumption have contributed around 21 (0.4) times more to the decline in PM_{2.5} (CO₂) emissions relative to gasoline.

Fuel substitution and abatement trade-offs. Diesel fuel vehicles contribute considerably more to emissions of fine particulates, such as PM_{2.5}, than gasoline vehicles.³⁹ However, diesel vehicles have lower CO₂ emissions rates per traveled

³⁸Note that our simulations for PM_{2.5} are not directly comparable to our SCM results as the former only accounts for exhaust emissions, thus missing a share of total PM_{2.5} emission reductions.

³⁹Relying on emission factors provided by the EEA for Germany, the average PM_{2.5} emission factor for diesel vehicles is around 56 times larger than that for gasoline (Ntziachristos and Samaras, 2019).

Figure 2.6: Share of total emission reductions by fuel due to the eco-tax




Notes: The figures above plot the share of total predicted emissions reductions by fuel type from our log-level semi-elasticity models (cf. Section 2.2.2). The share of total emission reductions for each fuel type is computed from the estimated post-treatment gap in emissions from gasoline (diesel) consumption due to the eco-tax, which refers to the distance between the bottom black line and the purple line in Figure 2.5.

kilometer compared to gasoline vehicles, by around 20% for otherwise virtually identical vehicles (Linn, 2019), as diesel engines are typically much more fuel-efficient. It follows that policy measures that foster a switch from gasoline vehicles to diesel vehicles (e.g., taxes based on the carbon content of fuels), could, in turn, lead to a decrease in CO₂ emissions but also an increase in PM_{2.5} emissions. Previous research on fuel and carbon taxation has not explicitly considered this trade-off in policy evaluations, with the exception of Linn (2019).

Figure 2.7 plots the estimated gasoline-to-diesel substitution induced by the eco-tax (cf., Table 2.A11 in the Appendix), implying that part of the contraction in CO₂ linked to reduced gasoline use came at the expense of greater PM_{2.5} emissions due to fuel substitution. We estimate that the share of diesel consumption is predicted to have increased by around 4% more than it would have had in the absence of the eco-tax throughout the post-treatment period. Our calculations suggest that gasoline-to-diesel substitution due to the eco-tax led to a cumulative increase in PM_{2.5} exhaust emissions of around 25,000 tons.

Fleet renewal and passenger-kilometers. An important argument for regulating emissions in the transport sector is that it can prompt a more rapid adoption of more efficient vehicles (e.g., Jacobsen et al., 2023). Panel (a) in Figure 2.8 provides descriptive evidence of the change in fleet renewal rate by plotting the share of new

Figure 2.7: Substitution towards diesel due to the eco-tax



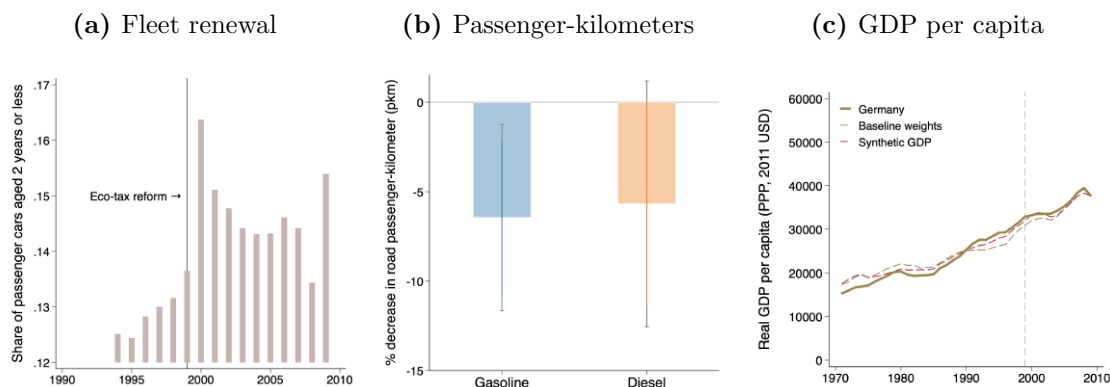
Notes: The figures plots the annual predicted substitution towards diesel from our semi-elasticity models (cf. Table 2.A11 in the Appendix).

passenger car registrations in the German fleet over time. We observe a discontinuity following 1999: after the eco-tax reform, the share of new passenger cars increased on average by 2%. Drawing a connection between this trend and our findings on low-carbon innovation (cf. Section 2.4.3), it seems plausible that the eco-tax has played a role in accelerating the adoption of cleaner vehicles, which could, at least partly, explain the contraction in emissions. We then resort to our semi-elasticity models to investigate how changes in the eco-tax rate affected the volume of road passenger transport, proxied by passenger-kilometers (pkm). Panel (b) in Figure 2.8 shows that, on average, the eco-tax is associated with a decrease in pkm by around 6.5% (5.7%).⁴⁰ These results offer suggestive evidence indicating that a share of the estimated emission reductions can be attributed to both an accelerated fleet renewal and a reduction in the volume of road passenger transport.

Decoupling. A common contention against the implementation of carbon taxation revolves around potential detrimental effects on economic growth. We thus investigate

⁴⁰We provide complementary, suggestive evidence that the eco-tax, and the consequent estimated reduction in pkm, led to fewer road casualties (fatalities and injuries), which represent a considerable externality of road transport (e.g., Anderson and Auffhammer, 2014). Again leveraging our semi-elasticity models, we find that the introduction of the eco-tax is, on average, associated with decreased road casualties by approximately 11% (cf. Figure 2.A14 in the Appendix). This underscores that the externality reductions we capture here—focused solely on climate and health benefits linked to air pollution—likely represent a conservative estimate of the benefits generated by the eco-tax.

Figure 2.8: Underlying mechanisms of reductions in emissions



Notes: Panel (a) plots the share of new passenger cars in the German fleet (aged 2 years or less) using data from the UNECE Statistical Database. Panel (b) plots the estimated percentage reductions in passenger-kilometers (pkm) by fuel for the average eco-tax rate of 13 cents. Data on pkm was retrieved from OECD Statistics. Panel (c) plots the evolution of GDP per capita in Germany and compares it with synthetic counterfactual developments.

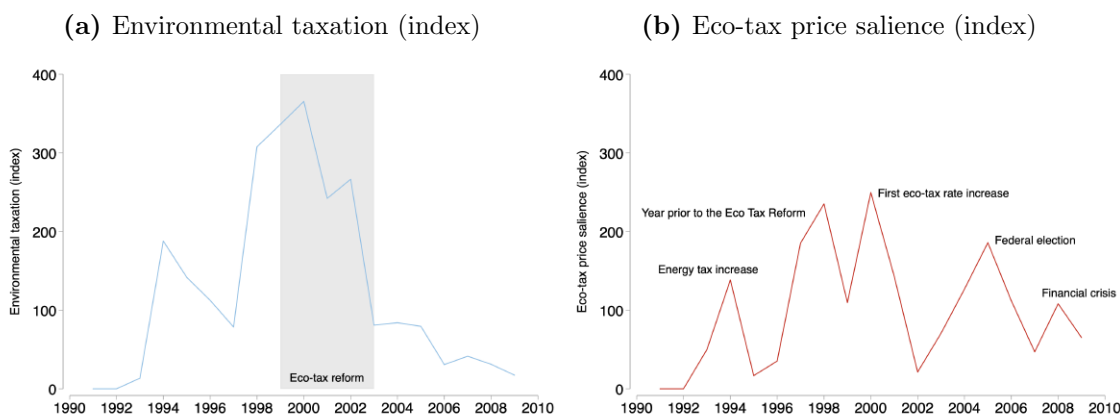
whether the observed reduction in emissions may have occurred alongside a reduction in economic activity. Figure 2.8 plots the evolution of GDP per capita in Germany relative to synthetic counterfactuals. Specifically, we rely on (i) the *Baseline* weights to construct a *no eco-tax* synthetic GDP development and (ii) an additional SCM specification where we further include lagged GDP in 1989 and 1991 as special predictors to account for the effect of German reunification. In both cases, we do not document any observable long-term negative effects on GDP from the eco-tax reform.

2.5.4 Tax salience

Our analysis continues by quantifying the role of eco-tax price salience in the media in driving the estimated effects of the eco-tax, drawing on a growing number of economic studies leveraging newspaper data as source of variation in the salience of events (e.g., Li et al., 2014; Baker et al., 2016; Beach and Hanlon, 2023).

Figure 2.9 illustrates the evolution of our newspaper-based indices. Panel (a) plots general publishing trends related to environmental taxation scaled by the total number of published articles. We observe a clear surge in news coverage of environmentally-motivated taxation, particularly during the years leading up to and during the

Figure 2.9: Evolution of salience in the media over time



Notes: Based on yearly series from 1991 to 2009. Authors' own calculations based on newspaper articles from Factiva (cf. Section 2.3). The two indices have a Pearson's correlation of 0.56. A detailed description of the steps undertaken to construct the newspaper indices can be found in Section 2.E of the Appendix.

implementation of the eco-tax, indicating that the policy sparked substantial media attention. However, this time series provides no insights into the specific focus of these article counts and whether the price effects of fuel taxes were being discussed. Thus, Panel (b) displays our *eco-tax price salience index*, introduced in Section 2.3, which specifically traces variations in the salience of fuel price increases induced by the eco-tax. We leverage changes in the eco-tax price salience index to investigate how variations in tax salience affect fuel-specific consumption responses. Specifically, we amend our semi-elasticity models (cf. Section 2.2.2) by additionally interacting our eco-tax price salience index with the annual real rate of the eco-tax.⁴¹ This allows us to empirically isolate how salience affects fuel use in accordance with the evolution of the eco-tax. Our identification strategy captures the additional effect on fuel demand decrease (at a given tax rate) due to to greater salience of eco-tax induced price increases in the media.⁴²

First, Columns (1) in Table 2.6 reports our coefficients of the amended semi-elasticity model based on time-series fuel demand variation in Germany. The significant negative interaction term indicates that greater eco-tax price salience is associated

⁴¹The interaction term will thus equal 0 prior to the eco-tax reform by design.

⁴²Our regressions focus on salience in the previous year, as print and digital news coverage tend to peak prior to actual or proposed changes to the eco-tax rate (cf. Li et al., 2014). Another key rationale for this approach is that it may help lessen the scope for reverse causality.

with lower consumption of both gasoline and diesel and that these effects increase with the eco-tax rate. Furthermore, the eco-tax elasticities tend to converge to the tax-exclusive real price elasticities after explicitly accounting for tax salience (henceforth, the *non-salient eco-tax elasticity*), suggesting that much of the divergence in the response for the increase in the eco-tax—relative to market-driven price changes—can be explained by tax salience in our model. Gasoline demand tends to be more responsive to variation in salience than diesel. This may be due for two reasons. Despite nearly one-third of all passenger cars in Germany having diesel engines at the time, diesel is widely used in freight transportation, which tends to be less price-responsive. Additionally, corporate fleets often use diesel, where companies cover part of the fuel cost, reducing consumers' price sensitivity.

Second, to address potential measurement error and endogeneity concerns, Columns (2) in Table 2.6 employs our fixed effects estimator to fit the amended semi-elasticity model described above, harnessing cross-country panel data to obtain improved within-country estimates.⁴³ Notably, the inclusion of our set of time fixed effects here allows us to control for any common developments across OECD and EU member states as well as shifts in cultural norms and societal concerns specific to each country, which can influence fuel demand based on attitudes toward car ownership, public transportation, and environmental consciousness after the eco-tax reform. Looking at the interaction term, our more refined fixed effects estimates yield comparable magnitudes, which corroborate our initial findings based on the time-series regressions based on Germany only. Furthermore, after controlling for the impact of tax salience, the two approaches also notably yield very similar *non-salient eco-tax elasticities*.

Finally, to better disentangle eco-tax price salience from time-varying factors driving media attention on the topic of environmental policy, Column (3) includes our

⁴³National newspapers are plausibly a significant source of information in our setting, but they may not be the sole determinants of tax salience for consumers. It is also plausible that our index of eco-tax price salience is correlated with other factors that create awareness, such as public discourse, government announcements, or advocacy campaigns which may introduce measurement errors in our index. Nevertheless, while the presence of such measurement error may influence precision in the magnitude of the estimated salience effects mediated solely by newspapers, it is less likely to alter the direction and significance of our findings.

Table 2.6: Effects of salience on fuel consumption

	(1)	(2)	(3)		(1)	(2)	(3)
	OLS: Baseline	OLS: Fixed effects	OLS: Fixed effects		OLS: Baseline	OLS: Fixed effects	OLS: Fixed effects
Raw price of Gasoline	-0.00280 (0.00176)	-0.00262 (0.00164)	-0.00257 (0.00164)	Raw price of Diesel	-0.00326*** (0.000900)	-0.00533*** (0.00163)	-0.00558*** (0.00159)
Energy Tax	-0.00338 (0.00489)	-0.00489*** (0.00128)	-0.00488*** (0.00128)	Energy Tax	-0.00723** (0.00348)	0.000340 (0.00152)	0.0000739 (0.00137)
Eco-tax	-0.00773 (0.00557)	-0.00683*** (0.00113)	-0.00594*** (0.00145)	Eco-tax	-0.00818*** (0.00275)	-0.00840*** (0.00264)	-0.00647** (0.00304)
Eco-tax x Eco-tax price salience (lag)	-0.00441** (0.00190)	-0.00414*** (0.000968)	-0.00467*** (0.00101)	Eco-tax x Eco-tax price salience (lag)	-0.00120* (0.000689)	-0.00272*** (0.000925)	-0.00359*** (0.00106)
Environmental taxation (lag)			-0.0121** (0.00559)	Environmental taxation (lag)			-0.0184** (0.00645)
Raw price = Eco-tax (p-value)	0.319	0.069	0.189	Raw price = Eco-tax (p-value)	0.081	0.473	0.845
Non-salient eco-tax elasticity	0.70	0.62	0.54	Non-salient eco-tax elasticity	0.62	0.63	0.49
Sample	Germany	OECD	OECD	Sample	Germany	OECD	OECD
Controls	✓	✓	✓	Controls	✓	✓	✓
Observations	38	765	765	Observations	39	574	574

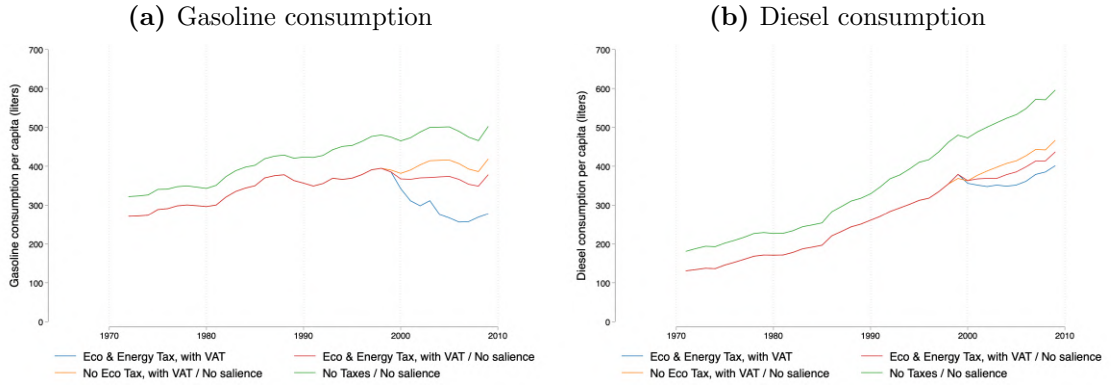
Notes: The dependent variable is log total fuel consumption in liters per capita for either gasoline or diesel (see column headings). Results for gasoline in column (1) refer to 1972-2009 due to missing price data prior to 1972. Standard errors in column (1) are calculated relying on the automatic bandwidth selection procedure following Newey and West (1994). Standard errors in columns (2) - (3) are clustered at the country-year level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

environmental taxation index (cf., Figure 2.9a) as an additional control. By doing so, we can account for the relative impacts of underlying peaks in the coverage of environmental taxes in the German news on fuel demand. Even in this case, the effects of eco-tax salience still hold while the environmental taxation index exhibits a significant and negative coefficient suggesting that deviations in news coverage of environmentally-motivated fuel tax schemes are associated with an overall decrease in fuel demand (a 10% increase reduces demand by around 0.1 - 0.2%).

To provide some perspective on the role of salience, let's consider the average eco-tax rate for gasoline (diesel) in real terms of 13 cents per liter. Our most conservative estimates from column (1) suggest that when our salience index exhibits an increase of a standard deviation relative to the mean, the additional reduction of gasoline (diesel) consumption induced by salience amounts to 4.3% (1.2%).⁴⁴ Leveraging our results from column (1), Figure 2.10 plots predicted gasoline and diesel consumption in the German transport sector under different taxation regimes and compares their evolution with and without salience. According to our model, eco-tax driven price

⁴⁴Our salience index exhibits a mean of 100 with a standard deviation of 76. A standard deviation increase thus represents a 76% increase relative to the mean. Both fuel consumption and the salience index are expressed in log terms in our model. Denoting the coefficient of the interaction term as φ_6 , we can interpret the estimated coefficients, $\hat{\varphi}_6$, as follows: For the average eco-tax rate of 13 cents, a standard deviation increase (or 76% increase relative to the mean) in our salience index will lead to an additional percentage reduction in fuel consumption amounting to $13 \times [(1.01^{\hat{\varphi}_6} - 1) \times 100] \times 0.76$.

Figure 2.10: Predicted fuel use under different tax and salience scenarios



Notes: The figures plot predicted fuel consumption from our amended log-level semi-elasticity models (cf. Section 2.2.2 and 2.5.4) under different taxation scenarios. We rely on the estimated fuel-specific price and tax elasticities computed from our estimates from columns (1) in Table 2.6. Specifically, Panel (a) refers to predicted per capita gasoline consumption (in liters), while Panel (b) is based on predicted per capita diesel consumption (in liters). The top green line displays predicted emissions in the absence of taxes, which means both the eco and energy tax elasticities are set to zero, and the VAT is deducted from the fuel price. For the orange line, the eco-tax elasticity is set to zero but the VAT-inclusive energy tax is now included. The red line shows how predicted emissions change when we include both the eco and energy taxes with the VAT but we set salience (as proxied by our newspaper-based index) equal to zero. The bottom blue line provides predicted emissions using the full model described in Section 2.5.4 with the differentiated tax and price elasticities which additionally includes the salience interactive term.

salience is responsible for around 70% (55%) of the contraction in gasoline (diesel) consumption in our simulation. Overall, these results corroborate the hypothesis that consumers react more strongly—relative to market prices—to environmental taxes that are salient.⁴⁵

2.6 Assessing climate and pollution reduction benefits

While previous reports suggested that environmental improvements due to the German eco-tax have been limited (Steiner and Cludius, 2010), we document substantial reductions in emissions. To quantify benefits from reduced climate and pollution costs, we apply official cost estimates from the first comprehensive guidelines by the

⁴⁵Salience can interact with other mechanisms that may lead to larger *tax salience ratios* but are hard to isolate, including the expected persistence of the price increase (e.g., Li et al., 2014) or the moral desirability of demand reductions (e.g., Mideksa, 2021). Our tax salience effects may thus also capture increased persistence expectations or a stronger signal that demand reductions are socially desirable.

Umweltbundesamt (2012). We, first, apply these to a prior evaluation of carbon emission reductions and, subsequently, illustrate results for our simulations and causal inference methods.


The Umweltbundesamt (2012) recommended using a social costs of carbon (SCC) per ton of CO₂ emitted in 2010 by 80 Euros (in 2010 Euros), and provided disaggregated cost estimates for PM_{2.5} in the transport sector, distinguishing costs of PM_{2.5} exhaust emissions released within (364,100 €/t) and outside of cities (122,800 €/t), recognizing that within city emissions contribute more to human health costs. Using their reported breakdown of the share of PM_{2.5} within and outside of cities for different transport modes, we compute a weighted average of PM_{2.5} damages. External costs of NO_x are not distinguished across locations, with an average cost estimate of 15,400 €/t.⁴⁶ We transform all cost estimates from a base year 2010 to 2022 values using official inflation statistics.

Steiner and Cludius (2010) estimate a price elasticity of fuel demand of -0.18 based on household survey data and attribute -0.1 to the tax elasticity component, with which they quantify reductions of CO₂ emissions due to the eco-tax, amounting to 120 kg CO₂ per household per year. Multiplying with the yearly number of households in Germany from 1999 to 2009, this sums up to 50.73 million tons of CO₂ emissions. Evaluating these emission reductions with the 2010 SCC in 2022 Euros yields a climate benefit of 4.9 billion Euros (first bar of Panel (a) in Figure 2.11).

We, first, contrast this with results from our Simulation approach (Panel (a) in Figure 2.11). Simulating emission reductions of CO₂ relative to the counterfactual without the eco-tax yields 344 million fewer tons, and an aggregate climate benefit of around 35 billion Euros. We further simulate reductions of PM_{2.5} and NO_x emissions

⁴⁶The cost estimates provided by Umweltbundesamt (2012) were derived from the EU NEEDS project which—in the time frame relevant for our historical analysis—provided the most comprehensive cost estimates available. In terms of damage sources, the full PM_{2.5} damages were related to (human) health damages, which compares to a health damages share of 82% for NO_x, where the remaining share of damages derives from biodiversity loss (14%), crop yield damages (3%) and material damages (1%). See Umweltbundesamt (2012) for further details.

Figure 2.11: Reductions in external climate and pollution damages due to the eco-tax



Notes: The figure plots the estimated reductions in external climate and pollution damages based on estimates from (a) the Simulation Approach and (b) the Synthetic Control Methods on CO₂ (green), PM_{2.5} (orange), and NO_x (blue) reductions and compares their magnitudes with the implied estimates from Steiner and Cludius (2010). GSCM refers to our *Economic Activity* specification (cf. Table 2.2). Simulation results for PM_{2.5} emissions do not account for non-exhaust emissions. Aggregate cost reductions are computed relying on pollutant-specific official cost estimates provided by the Umweltbundesamt (2012) and expressed in 2022 Euros.

of 36,368 tons and 1.08 million tons,⁴⁷ translating into pollution reduction benefits of 31 billion Euros. In sum, the Simulation Approach suggests that the eco-tax has reduced external damages by 66 billion Euros, 13 times as much as the previous estimate.

We further consider alternative scenarios, first starting with a *non-salient eco-tax* scenario. We estimate that external damage reduction would have been around two-thirds smaller at 23 billion Euros in the absence of a salient eco-tax. Second, we consider a scenario with *no fuel substitution* from gasoline to diesel induced by the eco-tax.⁴⁸ External damage reductions would have amounted to 55.5 billion Euros with *no fuel substitution*, with a very different composition: While not switching to diesel would have led to much lower climate benefits (34.9 vs. 18.7 billion Euros), benefits to due reducing PM_{2.5} would have been higher (30.9 vs. 36.7 billion Euros).

⁴⁷To estimate NO_x emissions from fuel consumption, we rely on estimates from the EEA on average emission factors for gasoline (diesel) vehicles in Germany (Ntziachristos and Samaras, 2019) of 5.61 (20.1) grams of NO_x per kg of gasoline (diesel).

⁴⁸We compute the *no fuel substitution* scenario by holding annual traveled km per capita fixed. As gasoline vehicles are less fuel efficient than comparable diesel vehicles, this assumption implies that the foregone increase in diesel use due to fuel substitution translates into a 1.2 times increase in gasoline use to account for lower fuel efficiency (Linn, 2019). Foregone gasoline-to-diesel substitution is computed using column (3) in Table 2.A11 in the Appendix. We then add (subtract) the estimated foregone substitution towards diesel to predicted gasoline (diesel) use from column (1) in Table 2.5.

We now move to quantifying externality reductions using our causal inference estimates, primarily relying on the SCM.⁴⁹ The first three bars in Panel (b) of Figure 2.11 show results of our SCM for specifications yielding minimal, average and maximal emission reductions whereas the last bar refers to our GSCM results. The average across all seven SCM specifications suggests benefits from reduced carbon and pollution costs due to the eco-tax of 80.7 billion euros, more than 16 times as much as the estimate by Steiner and Cludius (2010).⁵⁰ The GSCM yields a slightly higher benefit estimate at around 95 billion Euros (final bar in Figure 2.11). For comparison, the average effect for Germany using the SDID methodology yields a reduction in external costs of 91 billion Euros.

Overall, our results suggest that the eco-tax was orders of magnitude more effective in reducing external damages than previously suggested. Crucially, evaluations of fuel or carbon taxes that focus solely on climate benefits (e.g., Andersson, 2019; Mideksa, 2021; Runst and Höhle, 2021) miss a substantial share of benefits. For the case of the German eco-tax, we estimate that neglecting health co-benefits due to reduced air pollution would miss the majority share—between 63% (GSCM), 69% (SDID) and 75% (average SCM)—of the reductions in external damages.

2.7 Conclusion

This chapter provides the most comprehensive assessment thus far of the effectiveness of fuel taxation to reduce climate and local pollution externalities with a quasi-experimental evaluation of the world’s largest environmental tax reform. Our battery of causal inference designs demonstrates that the German eco-tax introduced in 1999

⁴⁹The SCMs and simulation results for PM_{2.5} emissions are not directly comparable, as the latter relies on conversion factors that do not account for non-exhaust emissions (Ntziachristos and Samaras, 2019).

⁵⁰Note that the EDGAR data and the emission factors used in the simulation approach are based on laboratory emission rates which tend to underestimate actual on-road nitrogen dioxides and particulate matter emissions (Crippa et al., 2018), also partly due to the recent *Dieselgate* scandal (Grange et al., 2020). Our estimated impacts on on-road emissions may thus represent lower-bound estimates.

has led to sizable reductions in CO₂, PM_{2.5} and NO_X emissions. Using official cost estimates, we show that the eco-tax has saved more than 80 billion Euros of external damages between 1999 and 2009. The majority of reductions in external costs relate to reduced air pollution and associated health benefits. We further document that the eco-tax has induced low-carbon innovation, leading to more than 800 additional low-carbon patented technologies in Germany that may have contributed to lowering abatement costs. The external validity of the estimated air emissions reductions in Germany is corroborated by a synthetic difference-in-difference approach, leveraging the staggered early implementation of environmental fuel taxes in Sweden and Finland. By contrast, we only detect similarly substantial effects on low-carbon innovation in Sweden, likely due to the importance of innovation in the automobile industry in both Sweden and Germany.

In addition to our causal findings, we provide complementary evidence on mediating mechanisms, suggesting that the eco-tax has likely contributed to fostering fleet renewal of passenger cars and to reduced passenger-kilometers traveled, without having reduced economic activity. We further highlight the key role of fuel substitution for navigating the trade-off between attaining climate and pollution targets. Finally, we show that the much higher demand response to the eco-tax is primarily due to increased tax salience, which we measure explicitly based on newspaper data. We thereby provide the first direct empirical evidence for the hypothesis that consumers react more strongly to (environmentally-motivated) fuel taxes the more salient they and their associated price increases are.

Overall, our results underscore the pivotal roles of co-pollution, innovation, fuel substitution and tax salience for the effectiveness of fuel taxes to reduce external damages and carry important policy implications. First, a sole focus on carbon abatement—as is common in the literature (e.g., Andersson, 2019; Leroutier, 2022; Runst and Höhle, 2021)—substantially underestimates the potential of taxes on fossil fuels to reduce externalities.⁵¹ Thus, accounting for reductions in pollution costs and

⁵¹Our results likely still provide a lower-bound of eco-tax induced externality reductions, as the

associated health co-benefits is crucial when evaluating the benefits of carbon pricing. Accounting for such health co-benefits, which more immediately benefit those that bear the costs of higher fuel prices, may also be crucial for gathering support for fuel and climate policies (e.g., Löschel et al., 2021). Our finding is also relevant for evaluating distributional effects. While the consumer costs of fuel taxation tend to burden lower-income households disproportionately (e.g., Sterner, 2012a; Känzig, 2023), poorer households may also benefit disproportionately from better air quality (e.g., Banzhaf et al., 2019; Hernandez-Cortes and Meng, 2023). Consequently, the true incidence of fuel taxation is likely less regressive as often suggested based on consumer costs only (e.g., Drupp et al., 2021).

Second, and relatedly, it is important for evaluations of fuel and carbon pricing to consider the trade-offs that can arise between climate and air pollution targets (e.g., Linn, 2019; Parry et al., 2021). We show that this is particularly relevant in the context price instruments set on the carbon content of fuels that can foster gasoline-to-diesel substitution. While this general feature of second-best taxation (Knittel and Sandler, 2018) is less important in the US, due to a predominant share of gasoline-fuelled cars, it is key when evaluating fuel pricing schemes in Europe (Zimmer and Koch, 2017; Linn, 2019). We show that relaxing the assumption that consumers respond similarly to fuel taxes as to other sources of fuel price variation (Linn, 2019) suggests that policymakers have to navigate a much larger trade-off between climate and pollution-reduction benefits.

Third, we shed light on the potential of environmentally-motivated taxation to spur low-carbon innovation by capturing economy-wide responses to an implicit carbon tax. Our approach thus complements previous studies that focused on the innovation response of regulated companies (e.g., Calel and Dechezlepretre, 2016), which generally find limited aggregate effects, and provides indirect evidence on the potential magnitude of the additional innovation occurring along the supply

eco-tax may also have contributed to reducing congestion (e.g., Hintermann et al., 2021), fatality risk (e.g., Anderson and Auffhammer, 2014) or the reliance on fossil imports and related security concerns.

chain and across unregulated agents, for instance, due to pass-through of regulatory costs or knowledge spillovers. By permanently reducing abatement costs, induced innovation is a key dimension to capture when conducting a comprehensive cost-benefit assessment of the climate benefits of fuel and carbon taxation measures. Our results document that regulatory-induced innovation responses to a carbon price can be sizable when considering economy-wide effects.

Finally, our results underscore the crucial role of tax salience for fostering the effectiveness of fuel taxation and carbon pricing (cf. Chetty et al., 2009; Li et al., 2014; Rivers and Schaufele, 2015). This implies that complementary informational measures may have considerable potential to foster climate, health and security benefits through a greater demand response at a given tax rate, and hence enhance the cost-effectiveness of price instruments to internalize externalities. This important role of salience, however, is a double-edged sword for policy design. While this is good news for policies aimed at reducing external costs or attaining specific climate or pollution targets, as fuel or carbon taxes may yield larger demand responses than is routinely considered in policy analysis using price elasticities estimated solely on market-price movements (e.g., Edenhofer et al., 2019). Tax salience may, however, also impede more stringent future policies due to stronger public resistance, such as in the case of the French “Yellow vests” (Douenne and Fabre, 2022). Indeed, while there were plans to further increase the eco-tax rate over time, the yearly increases were discontinued in 2003 due to public resistance, and only picked up again in 2021, then under the explicit label of carbon pricing.

2.A Background on the eco-tax and data sources

Taxing oils and fuels has a long history in Germany; the first mineral oil tax was established in 1939 for several fuels, including fuel oil, and other mineral oils such as gasoline and petroleum (Bundesministerium der Finanzen, 2014). In the 1980s, Binswanger (1992) suggested an ecological tax to internalize the externalities from the transport sector by implementing a tax at a low level and raising it until emissions have decreased to an environmentally sustainable level (Knigge and Görlach, 2005). The ecological fiscal reform (henceforth eco-tax reform) then came into effect in April 1999 taxing fuels, gas, electricity, and heating oil (Bundesgesetzblatt I, S.378, 1999; Steiner and Cludius, 2010). Note that this means that most of the first half of the year 1999 is not treated, which we consider in our analysis. In each year between 1999 and 2003, the fuel tax on gasoline and diesel was increased by 3.07 cents (6 Pfennig) per liter. This led to a total tax increase of 15.35 cents per liter for gasoline and diesel and is hereafter referred to as the eco-tax.

The law was updated in 2002, when some tax rates were increased and special rules implemented (Bundesgesetzblatt I, S. 2432., 1999; Bundesgesetzblatt I, S. 4602, 2002). Due to economic and social concerns, the eco-tax was exempted in other sectors was subject to substantial exceptions; thus it only affected the price of fuels and the use of electricity for less energy-intensive industries (Knigge and Görlach, 2005; Bach, 2009). For this reason, we focus our analysis on the German transport sector only instead of total economy-wide emissions. Since 2003, the eco-tax rate has not been changed, implying that nominal taxes on transport fuels have remained the same since 2003 up until the introduction of an explicitly labeled CO₂-price in January 2021. The revenue generated by the eco-tax overwhelmingly goes toward the German pension fund as reducing the statutory payments toward the pension fund was one of the key goals of the tax reform (Beuermann and Santarius, 2006). Out of the 18.7 billion euros that were raised by the eco-tax in 2003, 16.1 billion

euros went to the pension fund (Kohlhaas et al., 2005).

Below, we report some descriptive statistics to provide context on the German transport sector before and after the implementation of the eco-tax. Figure 2.A1 plots total fuel consumption by fuel type over time whereas Figure 2.A2 shows the nominal mineral oil tax from 1939 to 2009 for gasoline and diesel. For real values and other tax rates, please refer to Figure 2.A3. Over time, this law was changed frequently until its name was eventually changed to energy taxation law in 2006 Bundesministerium der Finanzen (2014). This is why we refer to the mineral oil tax as “energy tax” henceforth.

As mentioned in the main text, the German eco-tax is not a direct carbon tax, however, it can be interpreted as one. As of 2020, the total energy tax per liter of gasoline is 65.45 cents (Bundesministerium der Finanzen, 2014). The combustion of one liter of gasoline emits 2.325 kg of CO₂ (US EPA, 2005). If this is taken as a base, the energy tax on gasoline indirectly amounts to 281.51€ per ton of CO₂. The numbers are slightly different for diesel with 2.660 kg of CO₂ emitted as a result of the combustion of one liter and an energy tax of 47.04 cents per liter (US EPA, 2005; Bundesministerium der Finanzen, 2014). Still, this amounts to a price of 176.84€ per ton of CO₂. Prior to the eco-tax reform, the energy tax resulted in an indirect carbon tax of 215.53€ per ton of CO₂ for gasoline and 119.17€ for diesel. This means, that the eco-tax increased the effective carbon price by 57.67€ (\$65.17) for diesel and 65.98€ (\$74.56) for gasoline between 1999 and 2003. Thereby the eco-tax effectively represented the second highest tax on CO₂ in the world at that time.⁵² Figure 2.A4 compares the evolution of fuel-specific tax rates in Germany to the OECD average to put magnitudes into perspective in relation to the donor pool of countries employed for the synthetic control methods (SCMs).

The eco-tax reform was the major policy approach to curbing transport-related

⁵²The World Bank (2020) counts seven CO₂ taxes in 2003, with the highest in Sweden (\$89.65), followed by Norway (\$44.53). The German eco-tax is not classified.

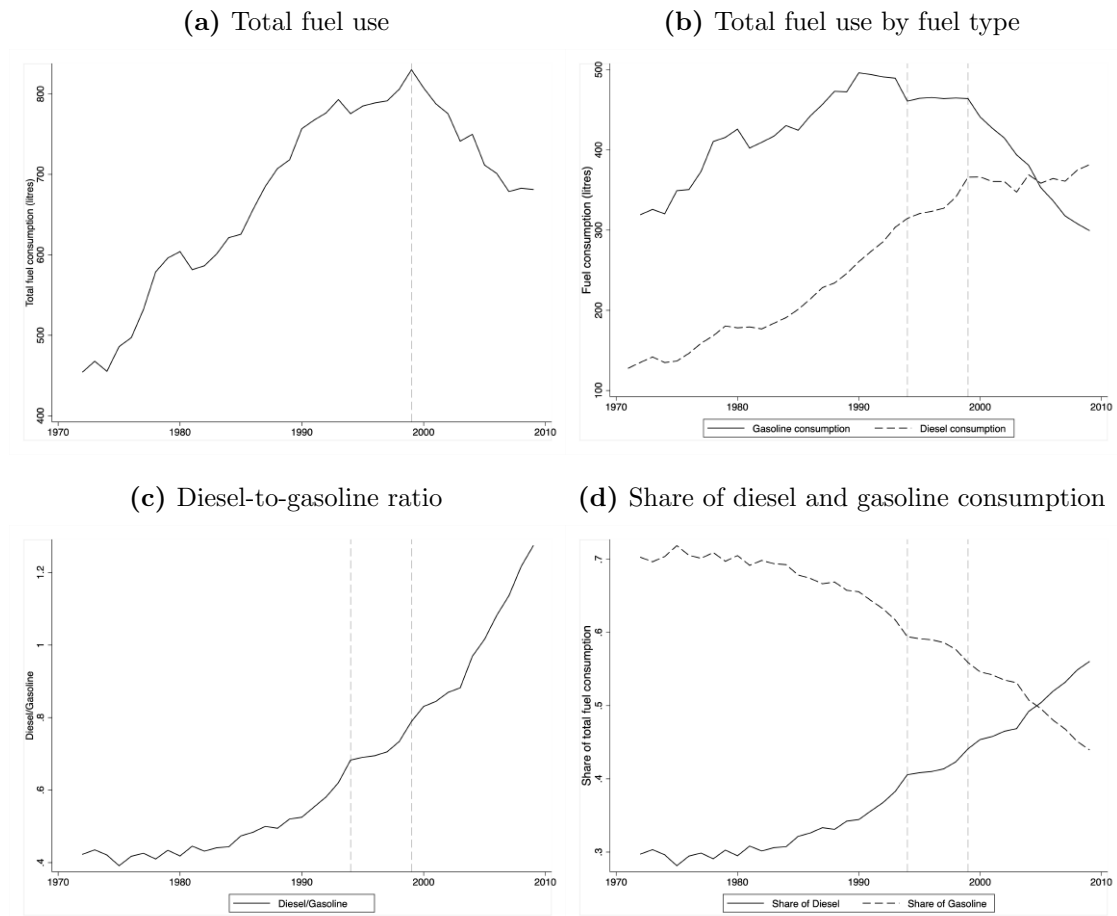
emissions in Germany until the final years of our time horizon.⁵³ Yet, towards the end of our considered time horizon—constituting one of the reasons why we ended our analysis in 2009—a few cities introduced environmental zones that restricted the entry of certain vehicles to the city center starting in the year 2008. Subsequently, other cities implemented and strengthened the standard of environmental zones. While Holman et al. (2015) only show small improvements in local air quality in German low emission zones, Wolff and Perry (2010) find that this policy was successful in reducing negative health outcomes in the regulated areas, without examining potential spillovers to other locations within Germany. Thus, while we cannot fully exclude that low emission zones may have contributed to emission reductions in the final two years of our analysis (where the magnitude of the estimated treatment effects starts decreasing), they are very unlikely to be a considerable driver, as they only affected a small number of cities at the end of our time frame.

Furthermore, as a response to the financial crisis in 2008, the German government decided to boost new car sales by temporarily paying a scrappage subsidy. This program was available from January 2009 until September 2009. A subsidy of 2500 Euro was available when disposing of one's car under the condition that it was at least nine years old and that a new car that fulfills certain vehicle emission standards (i.e., the Euro Norm 4 criteria) is bought. Helm et al. (2023) show local air pollution improvements following the policy. The next permanent nationwide taxation reform, after the introduction of the eco-tax, concerned vehicle circulation taxation. The reform was agreed on in 2009 and came into force on July 1st, 2009. With it, the taxation method to calculate vehicle circulation taxes was adjusted to include CO₂ emissions for cars that were first registered from July 2009 onward, which decreased car registrations (Klier and Linn, 2015). Nevertheless, Flintz et al. (2022) conclude that it did not have the necessary drive that is needed to abate emissions considerably.

⁵³Solely some changes to tax breaks for commuters in Germany were revised in the core time frame, with limited expected effects on emissions and/or innovation. Specifically, the maximum tax break was increased between 2001 and 2004 to 0.36 Euro per kilometer for commuting distances up to ten kilometers, for anything above up to 0.4 Euros per kilometer. In 2004 this law (EStG §9(1)4) was changed to a universal tax break of 0.30 Euro per kilometer (Weiss, 2009).


Both of these subsidies might have changed the composition of the newly registered vehicles at the end of our time horizon, and indeed we observe a spike in fleet renewal in the final year (see Panel (a) of Figure 2.8) which was still smaller than the spike in fleet renewal following the introduction of the eco-tax. Our sample thus ends at the onset of the next major nationwide policy reform concerning the transportation sector following the eco-tax reform.

Figure 2.A1: Fuel consumption over time




Notes: Data on fuel consumption is expressed in liters per capita or percentage terms, as denoted on the y-axis.

Figure 2.A2: Nominal taxes of gasoline and diesel over time



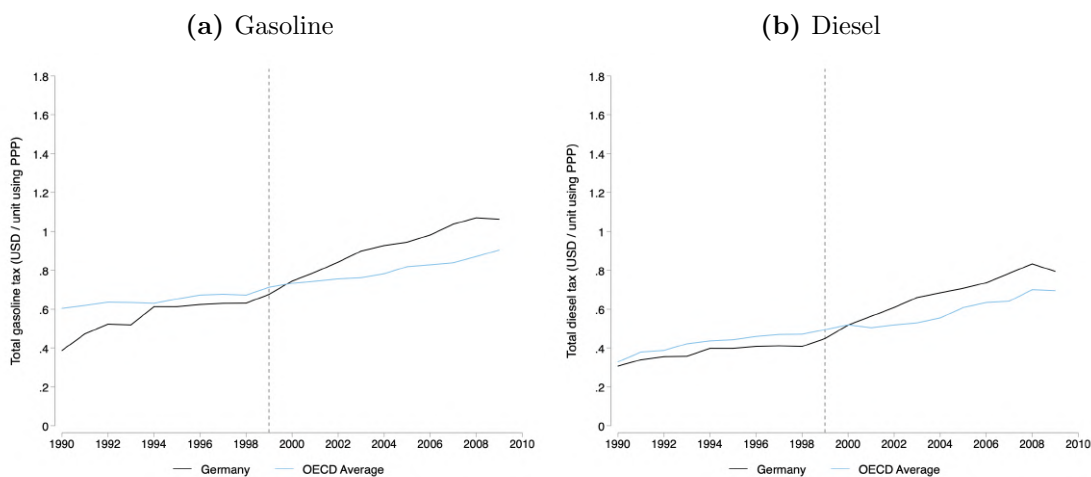
Notes: The figure above plots nominal taxes of gasoline and diesel from 1939 to 2009 as reported by the Bundesministerium der Finanzen (2014). Note that whenever a tax changes throughout a year, the average tax is calculated and shown here. Numbers are in cents.

Figure 2.A3: Real fuel prices and their tax components over time



Notes: Prices are in 1995€. Own calculations.

Figure 2.A4: Fuel taxes in Germany and the OECD average



Notes: Prices are in USD using PPP. Source: IEA Energy Prices and Taxes Statistics.

Table 2.A1: Data Sources

Variable	Source
Share of CO ₂ emissions from transport	Data downloaded from World Bank
CO ₂ emissions from fuel combustion	IEA
PM _{2.5} and NO _x emissions from EDGAR	EDGAR
Population	World Bank
Expenditure-side real GDP at current PPPs (in mil. 2011 US\$)	Penn World Tables
Urban population (% of total population)	World Bank
Road sector diesel (1) and gasoline (2) fuel consumption per capita (kg of oil equivalent)	World Bank (1), World Bank (2)
Road sector gasoline fuel consumption per capita (kg of oil equivalent)	Mineralwirtschaftsverband
Consumer price index for Germany (1995=100)	Statistisches Bundesamt (Destatis)
Strategic Reserve for Gasoline and Diesel in DM/t	Erdölbevorratungsverband
Energy Tax for diesel and gasoline in cents per litre	Bundesministerium für Finanzen
Eco Tax for diesel and gasoline in cents per litre	Bundesministerium für Finanzen
Value-added tax rate	Statista
Fuel prices and taxes for OECD countries	IEA Energy Prices and Taxes Statistics (Commercial data)
Unemployment Rate	Bundesagentur für Arbeit
U.S. Crude Oil First Purchase Price (Dollars/Barrel)	EIA
Euro/ECU exchange rates - annual data	Eurostat
Vehicles ownership per 1,000 people	Received from Professor Gately (Dargay et al., 2007).
Low-carbon patents related to transportation: triadic patent families (1) and total (2)	OECD (1), OECD (2)
Newspaper-specific article frequency counts	Factiva (Commercial data)
Road passenger transport (pkm)	OECD
Vehicle registrations by age	UNECE
Road casualties	OECD

2.B Additional results for the synthetic control methods

This section of the Appendix provides additional supporting material and results related to the synthetic control method (SCM) and its generalized version (GSCM). Specifically, this section contains the following material: Tables 2.A2 - 2.A4 report country-specific weights used for the construction of our synthetic counterfactuals in Figure 2.1. The three panels in Figure 2.A6 plot in-time placebo tests when we assign a fake treatment to Germany in 1995. Figure 2.A7 reports our results leveraging the standard SCM when we do not impose any of the sample restrictions discussed in Section 2.3. Figure 2.A8 reports leave-one-out tests (cf. Abadie et al., 2015) for

our *Baseline* (i.e., Panels a, c and e) and *No covariates* specifications (i.e., Panels b, d and f). The former is in line with the recommendations in Kaul et al. (2015), while the latter follows Ferman et al. (2020). Finally, Figure 2.B.4 plots the dynamic treatment effects estimated for each of our GSCM specifications presented in Section 2.4.2.

Table 2.A2: SCM for CO₂: Pre-Treatment Predictor Means for Germany, Baseline Synthetic Germany and the Sample Average

Variables	Germany	Synthetic	Sample Mean
GDP per capita	22,197.42	23,615.94	17,972.24
Diesel consumption per capita	185.23	185.27	130.29
Gasoline consumption per capita	332.55	332.77	343.23
Share of urban population	0.73	0.73	0.73
Number of vehicles per 1,000 people	410.34	410.48	290.14
CO2 from transport in 1998	2.10	2.10	2.12

All variables except lagged CO₂ per capita are averaged from 1971-1998. GDP per capita is measured at current PPPs in million 2011 USD. Gasoline and diesel consumption is measured in kg of oil equivalent. Share of urban population is measured as a percentage of total population. CO₂ emissions are measured in metric tons per capita and are retrieved from the IEA.

Table 2.A3: SCM for PM_{2.5}: Pre-Treatment Predictor Means for Germany, Baseline Synthetic Germany and the Sample Average

Variables	Germany	Synthetic	Sample Mean
GDP per capita	22,197.42	22,346.93	17,972.24
Diesel consumption per capita	185.23	170.25	130.29
Gasoline consumption per capita	332.55	367.82	343.23
Share of urban population	0.73	0.75	0.73
Number of vehicles per 1,000 people	410.34	410.39	290.14
PM _{2.5} from transport in 1998	0.58	0.61	0.58

All variables except lagged PM_{2.5} per capita are averaged from 1971-1998. GDP per capita is measured at current PPPs in million 2011 USD. Gasoline and diesel consumption is measured in kg of oil equivalent. Share of urban population is measured as a percentage of total population. PM_{2.5} emissions are measured in kg per capita and are retrieved from the EDGAR v6.1 database.


Table 2.A4: SCM for NO_x: Pre-Treatment Predictor Means for Germany, Baseline Synthetic Germany and the Sample Average

Variables	Germany	Synthetic	Sample Mean
GDP per capita	22,197.42	22,199.20	17,972.24
Diesel consumption per capita	185.23	179.35	130.29
Gasoline consumption per capita	332.55	303.51	343.23
Share of urban population	0.73	0.76	0.73
Number of vehicles per 1,000 people	410.34	360.88	290.14
PM _{2.5} from transport	0.50	0.50	0.42
NO _x from transport in 1998	14.13	14.26	16.72


All variables except lagged NO_x per capita are averaged from 1971-1998. GDP per capita is measured at current PPPs in million 2011 USD. Gasoline and diesel consumption is measured in kg of oil equivalent. Share of urban population is measured as a percentage of total population. NO_x emissions are measured in kg per capita and are retrieved from the EDGAR v6.1 database.

Figure 2.A5: Comparing donor pool weights across SCM specifications


(a) Synthetic Germany: CO₂



(b) Synthetic Germany: PM_{2.5}




(c) Synthetic Germany: NO_x



Notes: The figure plots the estimated country-specific weights assigned by the synthetic control algorithms across our set of SCM specifications (cf. Table 2.1).

2.B.1 Placebo in time


Figure 2.A6: In-time placebos



Notes: The figure plots the in-time placebo for our results on (a) CO_2 , (b) $\text{PM}_{2.5}$, and (c) NO_x emissions where a placebo treatment is assigned in 1995.

2.B.2 No sample restrictions

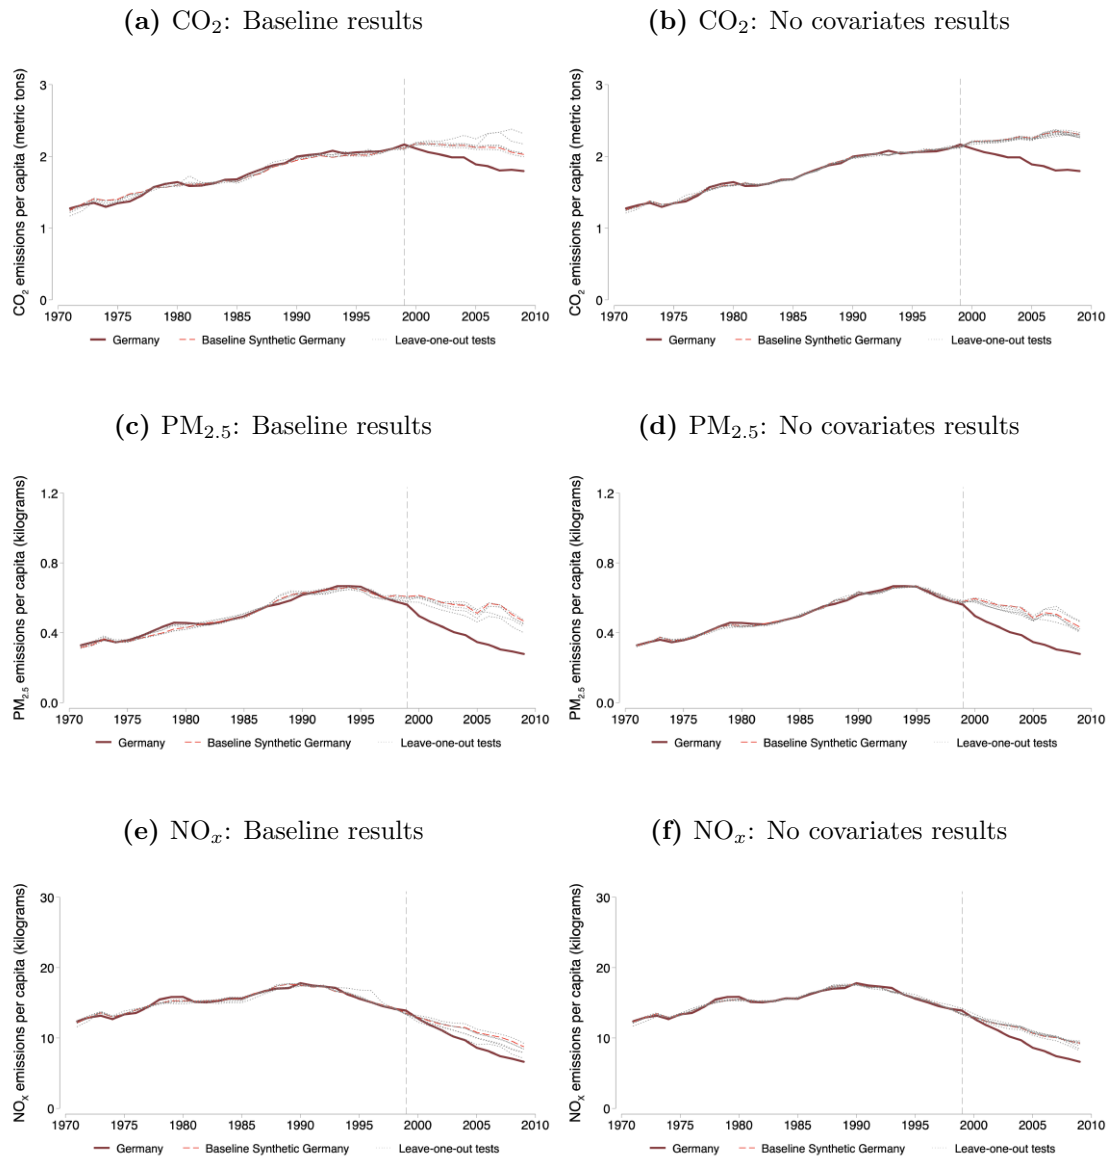
Figure 2.A7: Results with no donor pool restrictions



Notes: The figure plots our Baseline SCM results without applying the sample description described in Section 2.3.

2.B.3 Leave-one-out tests

Figure 2.A8: Leave-one-out tests




Notes: The figure plots leave-one-out tests following Abadie et al. (2015) where we iteratively exclude countries that receive at least a 1% in the construction of the synthetic counterfactual. More details can be found in Section 2.4.1.

2.B.4 Generalized Synthetic Control Method (GSCM)

Figure 2.A9: GSCM with Interactive Fixed Effects Models

(a) Change in CO₂ over time (b) Change in PM_{2.5} over time (c) Change in NO_x over time




Notes: The figure plots the estimated gaps in emissions relative to a synthetic counterfactual development based on a Generalized Synthetic Control Method with interactive fixed effects models Xu (2017). More details on the GSCM specifications can be found in Section 2.4.2.

Table 2.A5: Effects of the eco-tax with the GSCM (No sample restrictions)

	CO ₂ emissions (t)	PM _{2.5} emissions (kg)	NO _x emissions (kg)
Mean [95% CI]	-0.57 [-0.74; -0.30]	-0.14 [-0.27; -0.07]	-1.39 [-2.78; -0.21]
Observations	1482	1482	1482
Countries	38	38	38

Notes: All outcome variables are expressed in per capita terms. The table displays the estimated average gaps in emissions relative to a synthetic counterfactual development based on a Generalized Synthetic Control Method with interactive fixed effects models Xu (2017). More details on the GSCM specifications can be found in Section 2.4.2.

Figure 2.A10: Effects of the eco-tax with the GSCM on low-carbon patents (No sample restrictions)



Notes: The figure plots the estimated gaps in emissions relative to a synthetic counterfactual development based on a Generalized Synthetic Control Method with interactive fixed effects models Xu (2017). More details on the GSCM specifications can be found in Section 2.4.2.

2.C Elasticities

This Section is structured as follows. First, we provide evidence of tax pass-through to prove that taxes are noticeable to consumers in our setting. Second, Tables 2.A7 - 2.A7 provide a host of robustness tests for our real and eco-tax elasticity results presented in Section 2.5. Figure 2.A11 plots our elasticity results when employing a distributed lag model with one lead to account for anticipatory behavior (Coglianese et al., 2017; Kilian and Zhou, 2023).⁵⁴ Table 2.A10 and 2.A11 provides evidence of gasoline-to-diesel substitution in our setting again leveraging the semi-elasticity models presented in Section 2.2.2. Figure 2.A12 displays predicted NO_x emissions under different taxation scenarios complementing Figure 2.5 in the main text. Figure 2.A13 compares the dynamic treatment effects across all the different empirical strategies employed in our study, namely the (a) SCM, (b) the generalized SCM and (c) the simulation approach.⁵⁵ Finally, Figure 2.A14 leverages again the semi-elasticity models to provide some complementary suggestive evidence on the average effects of the eco-tax on road casualties (i.e., considering fatalities and injuries).

Tax pass-through. Before computing fuel-specific price and tax elasticities, we check if the tax increases get effectively passed through to the retail price of fuel to ensure that changes in taxation are noticeable to consumers (c.f. Andersson, 2019). We use first-differencing to regress the crude oil price i and the combined nominal energy and eco tax $\tau^{eco,energy}$ on the retail fuel price p^* of gasoline and diesel, respectively:

$$\Delta p_t^* = \alpha_0 + \alpha_1 \Delta \sigma_t + \alpha_2 \Delta \tau_t^{eco,energy} + \epsilon_t \quad (2.11)$$

⁵⁴We additionally run first-differences models including different sets of leads and lags of the normalized tax change, as in Kilian and Zhou (2023). We produce a distribution of p-values for testing the null of equal effects between tax-exclusive and eco-tax price changes: Across all the different specifications, we reject the null hypothesis of equal effects between tax-exclusive and tax-only price changes in our setting. Results are available upon request.

⁵⁵Note that simulated PM_{2.5} emissions are not directly comparable to our SCMs results as the former do not account for non-exhaust emissions.

The p-values of a linear Wald test show that for both regressions, the tax coefficient α_2 is not significantly different from unity.⁵⁶ For gasoline, α_2 equals 0.94 (with a 95% confidence interval of [0.79; 1.08]). The result is comparable for diesel, where the coefficient is 0.86 [0.54; 1.17]. We repeat the estimation with the tax rates being formally separated into energy and eco-tax in the model:

$$\Delta p_t^* = \alpha_0 + \alpha_1 \Delta \sigma_t + \alpha_3 \Delta \tau_t^{energy, VAT} + \alpha_4 \Delta \tau_t^{eco, VAT} + \epsilon_t \quad (2.12)$$

Again, we are not able to reject the hypothesis that there is full pass-through.⁵⁷ This indicates that fuel taxes have been noticeable for consumers and that we can interpret our estimates of fuel-specific tax elasticities as price elasticities of demand.

⁵⁶The p-value of the linear Wald test for $\Delta \alpha_2 = 1$ is equal to 0.38 for gasoline and 0.34 for diesel.

⁵⁷For gasoline, α_3 equals 0.92 [0.75; 1.09] and α_4 1.02 [0.83; 1.20]. While the eco-tax coefficient for diesel is similar at 0.96 [0.49; 1.43], the one for the energy tax is slightly lower at 0.64 [0.02; 1.25]. The p-values of the linear Wald tests for $\Delta \alpha_3 = 1$ are 0.34 for gasoline and 0.24 for diesel, and 0.84 and 0.87 for $\Delta \alpha_4 = 1$, respectively.

Table 2.A6: Gasoline consumption

	Real price	Aggregate tax	Eco-tax
Real price of Gasoline	-0.00603** (0.00278)		
Raw price of Gasoline (only VAT)		-0.00584* (0.00331)	-0.00357* (0.00204)
Energy + Eco Tax		-0.00798** (0.00375)	
Energy Tax on Gasoline			-0.00242 (0.00497)
Eco Tax on Gasoline			-0.0306*** (0.00773)
Dummy Eco Tax	-0.154 (0.131)	-0.144 (0.126)	0.104** (0.0393)
Trend	0.00158 (0.0138)	-0.00328 (0.0118)	0.0240 (0.0210)
GDP per capita	0.000174 (0.0116)	0.00893 (0.0168)	-0.0245 (0.0318)
Unemployment rate	0.0292 (0.0176)	0.0311* (0.0177)	0.00902 (0.0239)
Observations	38	38	38

Table 2.A7: Diesel consumption

	Real price	Aggregate tax	Eco-tax
Real price of Diesel	-0.00440*** (0.00103)		
Raw price of Diesel (only VAT)		-0.00384*** (0.000908)	-0.00346*** (0.00104)
Energy + Eco Tax		-0.0111*** (0.00141)	
Energy Tax on Diesel			-0.00729** (0.00292)
Eco Tax on Diesel			-0.0143*** (0.00359)
Dummy Eco Tax	-0.0205 (0.0564)	0.0574* (0.0315)	0.0794*** (0.0174)
Trend	0.0189*** (0.00587)	0.0104** (0.00456)	0.0187** (0.00774)
GDP per capita	0.0177*** (0.00528)	0.0287*** (0.00702)	0.0201** (0.00753)
Unemployment rate	0.0107* (0.00558)	0.0104* (0.00538)	0.00651 (0.00816)
Observations	39	39	39

Notes: The dependent variable is the log of fuel consumption in liters per capita, which refers to total fuel consumption or either gasoline or diesel consumption (as indicated by the column heading). Prices are in 1995€. Results for gasoline consumption refer to 1972-2009 due to missing price data prior to 1972. Unemployment is measured as percentage of total labor force. Newey-West standard errors in parentheses are heteroskedasticity and autocorrelation robust. Standard errors are calculated relying on the automatic bandwidth selection procedure following Newey and West (1994). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 2.A8: Real price elasticities for gasoline after 1991

	(1) OLS	(2) OLS	(3) OLS	(4) IV: Brent Crude
Real price of Gasoline	-0.00698*** (0.00142)	-0.00693*** (0.00150)	-0.00510*** (0.000592)	-0.00531*** (0.000640)
Dummy Eco Tax	0.105** (0.0371)	0.106*** (0.0354)	0.106*** (0.0164)	0.106*** (0.0135)
Trend	-0.0237*** (0.00703)	-0.0217* (0.0111)	-0.0336*** (0.00544)	-0.0332*** (0.00505)
GDP per capita		-0.00311 (0.00686)	0.00795 (0.00636)	0.00793 (0.00575)
Unemployment rate			0.0181*** (0.00309)	0.0178*** (0.00268)
N	19	19	19	19

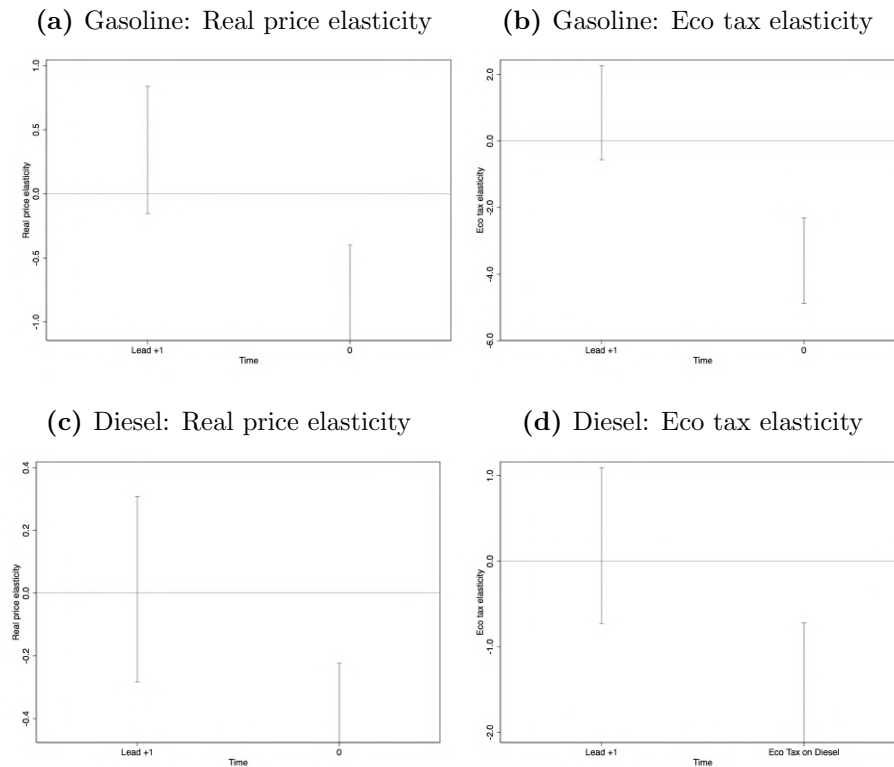
Table 2.A9: Real price elasticities for diesel after 1991

	(1) OLS	(2) OLS	(3) OLS	(4) IV: Brent Crude
Real price of Diesel	-0.00404** (0.00161)	-0.00456*** (0.00112)	-0.00358*** (0.000318)	-0.00317*** (0.000315)
Dummy Eco Tax	0.0687*** (0.0151)	0.0635*** (0.0167)	0.0634*** (0.0111)	0.0632*** (0.00961)
Trend	0.0206*** (0.00596)	0.0108 (0.00670)	0.00457 (0.00264)	0.00384* (0.00227)
GDP per capita		0.0172*** (0.00388)	0.0217*** (0.00355)	0.0211*** (0.00235)
Unemployment rate			0.0104** (0.00388)	0.0113*** (0.00378)
N	19	19	19	19

Notes: The dependent variable is the log of fuel consumption in liters per capita, which refers to total fuel consumption or either gasoline or diesel consumption (as indicated by the column heading). Columns (4) use the Brent crude oil price as an instrumental variable for the real fuel price. Prices are in 1995€. Unemployment is measured as percentage of total labor force. Newey-West standard errors in parentheses are heteroskedasticity and autocorrelation robust. Standard errors are calculated relying on the automatic bandwidth selection procedure following Newey and West (1994). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

2.C.1 Time-series elasticities with a distributed lag model

Figure 2.A11: Fuel-specific real price and eco tax elasticities with a lead



Notes: The figure plots the estimated fuel-specific elasticities of gasoline and diesel demand by amending our log-level semi-elasticity models with the introduction of a lead (c.f. Section 2.2.2). Specifically, Panel (a) and (c) show the real price elasticity of gasoline and diesel demand respectively (c.f. Table 2.4 and 2.4). Panel (b) and (d) display the gasoline and diesel eco tax elasticities (c.f. Table 2.5 and 2.5). Prices are in 1995€. Results for gasoline consumption refer to 1972-2009 due to missing price data prior to 1972. Unemployment is measured as percentage of total labor force. Confidence intervals are based on Newey-West standard errors are heteroskedasticity and autocorrelation robust. Standard errors are calculated relying on the automatic bandwidth selection procedure following Newey and West (1994). * $p < 0.05$, ** $p < 0.01$, ***

2.C.2 Fuel substitution due to the eco-tax

Table 2.A10: Fuel substitution: Diesel-to-Gasoline ratio

	(1) Diesel/Gasoline	(2) Diesel/Gasoline	(3) Diesel/Gasoline
Raw price of Gasoline (only VAT)	0.00187 (0.00241)	0.00185 (0.00124)	0.00184 (0.00126)
Energy Tax on Gasoline	0.00471*** (0.00123)	-0.000316 (0.00263)	0.000991 (0.00237)
Eco Tax on Gasoline	0.0175*** (0.00634)	0.0157*** (0.00465)	0.0126** (0.00482)
Dummy Eco Tax	-0.0108 (0.0276)	-0.0619** (0.0296)	-0.0377 (0.0242)
Trend	0.0126*** (0.00306)	-0.00700 (0.00700)	0.00671 (0.0152)
GDP per capita		0.0372** (0.0149)	0.0214 (0.0187)
Unemployment rate			-0.0142
Observations	38	38	38

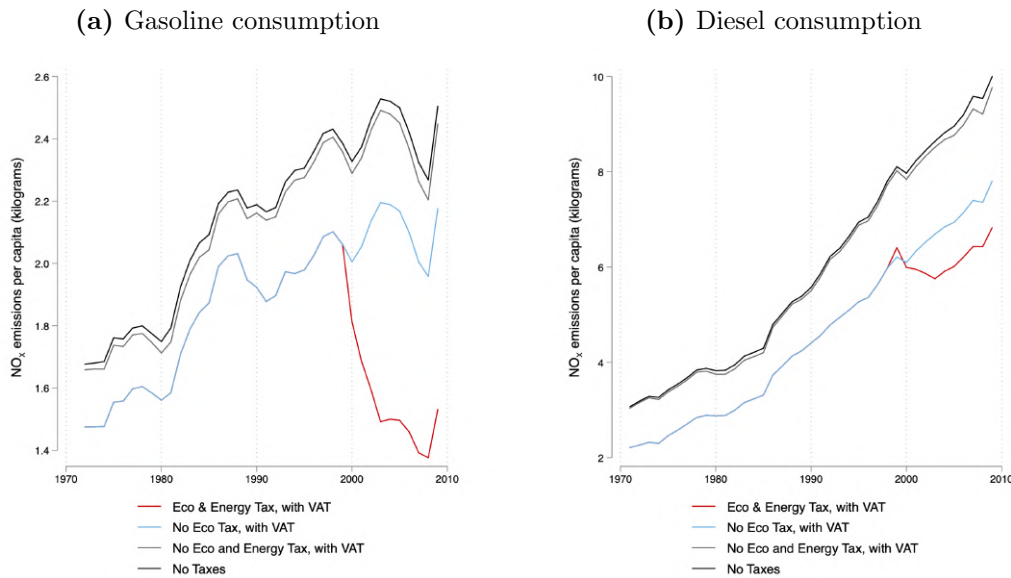
Table 2.A11: Fuel substitution: Share of Diesel

	(1) Share of Diesel	(2) Share of Diesel	(3) Share of Diesel
Raw price of Gasoline (only VAT)	0.000255 (0.000565)	0.000250 (0.000314)	0.000248 (0.000317)
Energy Tax on Gasoline	0.00179*** (0.000396)	0.000697 (0.000658)	0.000917 (0.000721)
Eco Tax on Gasoline	0.00415*** (0.00144)	0.00376*** (0.00110)	0.00325*** (0.00108)
Dummy Eco Tax	0.00367 (0.00645)	-0.00746 (0.00747)	-0.00339 (0.00630)
Trend	0.00482*** (0.000731)	0.000554 (0.00168)	0.00286 (0.00378)
GDP per capita		0.00810** (0.00352)	0.00546 (0.00491)
Unemployment rate			-0.00239 (0.00420)
Observations	38	38	38

Notes: The dependent variable is either (a) the ratio of diesel-to-gasoline consumption in litres per capita or (b) the share of diesel of total fuel consumption in percentage terms (as indicated by the column heading). Prices are in 1995€. Results for gasoline consumption refer to 1972-2009 due to missing price data prior to 1972. Unemployment is measured as percentage of total labor force. Newey-West standard errors in parentheses are heteroskedasticity and autocorrelation robust. Standard errors are calculated relying on the automatic bandwidth selection procedure following Newey and West (1994). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

2.C.3 NO_X emission under different taxation regimes

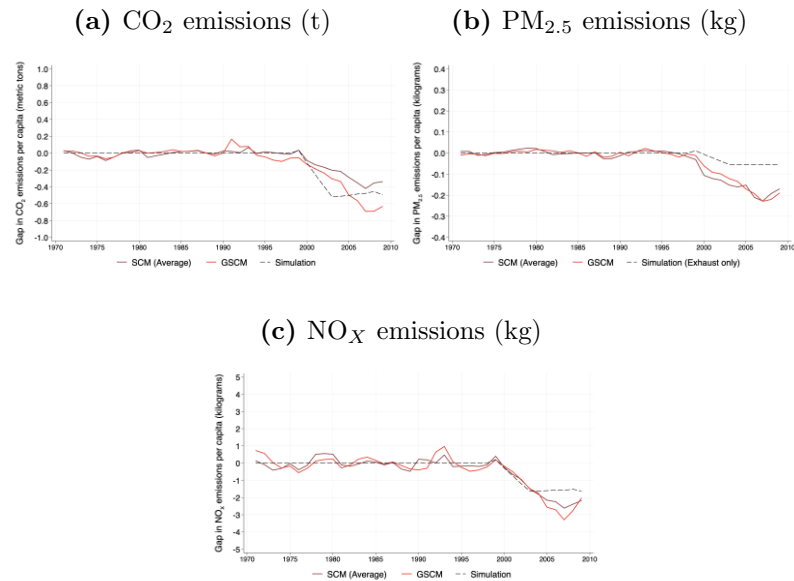
Figure 2.A12: Predicted NO_X emissions by fuel under different tax scenarios



Notes: The figures above plot predicted emissions from the eco-tax specification of our log-level semi-elasticity models (c.f. Section 2.2.2) under different taxation scenarios. We rely the estimated fuel-specific price and tax elasticities computed from our estimates from column (3) in Tables 2.5 and 2.5. Panel (a) refers to predicted emissions from gasoline consumption, while Panel (b) covers diesel consumption. In each panel the y-axis refers to per capita NO_X in kilograms. The top black line displays predicted emissions when the eco and energy tax elasticities are set to zero, and VAT is deducted from the fuel price. For the gray line, the eco and energy tax elasticities are set to zero, but VAT is included. The light blue line shows how predicted emissions change when the eco tax is set to zero, but we include the energy tax and VAT. The red line provides predicted emissions using the full model with differentiated tax and price elasticities.

2.C.4 SCMs and the Simulation Approach


Figure 2.A13: Gap in per capita emissions: SCMs vs Simulation Approach



Notes: The figures above plot the estimated average gap in per capita emissions from our synthetic control experiments (c.f. Section 2.2.1) and the simulation approach based on our log-level semi-elasticity models (c.f. Section 2.2.2). Nationwide reductions in emissions in the simulation approach have been computed by accounting for predicted emission reductions from both gasoline and diesel. Note that simulated PM_{2.5} emissions are not directly comparable to our SCMs results as the former do not account for non-exhaust emissions.

2.C.5 Impacts of the eco-tax on road casualties

Figure 2.A14: Effects of the eco-tax on road casualties



Notes: The dependent variable is the number of road casualties (i.e., including fatalities and injuries) in logarithmic terms. The estimated effects refer to the average eco-tax rate of 13 cents. All regressions control for the fuel raw price, the energy tax rate, GDP per capita (in 1995€), the unemployment rate, and include a time trend as well as a dummy for the post-treatment period (i.e., equal to 1 after 1999). We use Newey-West standard errors that are heteroskedasticity and autocorrelation robust following Newey and West (1994).

2.D Synthetic difference-in-differences

Table 2.A12: Effects of environmental taxes with an SDID staggered adoption design (No sample restrictions)

	CO ₂ emissions (t)	PM _{2.5} emissions (kg)	NO _x emissions (kg)	Low-carbon patents
Environmental fuel taxation	-0.228*** (0.077)	-0.103*** (0.025)	-3.423** (1.369)	0.349 (0.312)
Observations	1209	1209	1209	775
Countries	31	31	31	31

Notes: All outcome variables are expressed in per capita terms. Patents are expressed in per million population terms. Standard errors are computed using the bootstrap variance estimation algorithm outlined in Arkhangelsky et al. (2021), which requires multiple treated units. All regressions include unit-specific and time-specific fixed effects and control for GDP per capita and include a binary variable indicating whether a country was regulated by EU-wide regulations after 2005.

Table 2.A13: Effects of environmental taxes with an SDID staggered adoption design (Germany and Sweden)

	CO ₂ emissions (t)	PM _{2.5} emissions (kg)	NO _x emissions (kg)	Low-carbon patents
Environmental fuel taxation	-0.250*** (0.064)	-0.098*** (0.037)	-1.835* (1.112)	0.929*** (0.137)
Observations	819	819	819	525
Countries	21	21	21	21

Notes: All outcome variables are expressed in per capita terms. Patents are expressed in per million population terms. Standard errors are computed using the bootstrap variance estimation algorithm outlined in Arkhangelsky et al. (2021), which requires multiple treated units. All regressions include unit-specific and time-specific fixed effects and control for GDP per capita and include a binary variable indicating whether a country was regulated by EU-wide regulations after 2005.

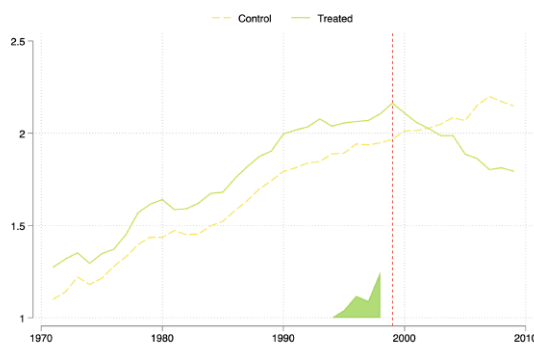
Table 2.A14: Effects of environmental taxes with an SDID staggered adoption design (Additional covariates)

	CO ₂ emissions (t)	PM _{2.5} emissions (kg)	NO _x emissions (kg)	Low-carbon patents
Environmental fuel taxation	-0.245*** (0.057)	-0.099** (0.039)	-2.771** (1.104)	0.643*** (0.246)
Observations	858	858	858	550
Countries	22	22	22	22


Notes: All outcome variables are expressed in per capita terms. Patents are expressed in per million population terms. Standard errors are computed using the bootstrap variance estimation algorithm outlined in Arkhangelsky et al. (2021), which requires multiple treated units. All regressions control for GDP per capita and include unit-specific and time-specific fixed effects as well as a binary variable indicating whether a country was regulated by EU-wide regulations after 2005. Emissions reductions are estimated by additionally controlling for pre-treatment diesel and gasoline consumption. For low-carbon innovation, we control for pre-treatment triadic patents per capita.

Figure 2.A15: Dynamic SDID effects in Germany and time-specific weights


(a) Change in CO₂ over time




(b) Change in PM_{2.5} over time



(c) Change in NO_x over time




(d) Change in patents over time




Notes: The figure plots the graphical results from our SDID staggered adoption design presented in Section 2.4.4. Time weights are represented in light green at the bottom of the pre-intervention period.

Figure 2.A16: Dynamic SDID effects in Sweden and time-specific weights


(a) Change in CO₂ over time




(b) Change in PM_{2.5} over time



(c) Change in NO_x over time




(d) Change in patents over time




Notes: The figure plots the graphical results from our SDID staggered adoption design presented in Section 2.4.4. Time weights are represented in light green at the bottom of the pre-intervention period.

Figure 2.A17: Dynamic SDID effects in Finland and time-specific weights


(a) Change in CO₂ over time




(b) Change in PM_{2.5} over time



(c) Change in NO_x over time



(d) Change in patents over time



Notes: The figure plots the graphical results from our SDID staggered adoption design presented in Section 2.4.4. Time weights are represented in light green at the bottom of the pre-intervention period.

2.E Salience analysis

The following section provides additional information on the salience analysis conducted in Section 2.5.4. This section of the Appendix is structured in three parts. First, we report the different search strategies that were used to extract frequency counts of newspapers' articles from Factiva. Second, we provide a detailed description of the construction of our set of newspaper-based indices that were employed in the empirical analysis. Finally, we present a set of robustness checks for our empirical analysis of salience effects presented in Section 2.5.4.

2.E.1 Search strategies

Here below, we report the three different search strategies that were developed to download articles' count used in the construction of our indices. A brief description of each strategy will follow. *Strategy # 1* restricts our search to articles talking about environmental/ecological taxation. This provides us with a clearer idea of publishing trends directly related to environmental taxation and will be used to scale frequency counts of a more targeted search strategy that specifically captures price salience. Finally, *Strategy # 2* is employed to identify articles talking about environmental/ecological taxation *and* resulting in increases in fuel prices. Here, we use a double AND operator to impose that at least one keyword from each of the brackets that come after the operator must appear in the article.

Strategy #1: Environmental taxation trends. (Ökosteur* or "Ökologische Steuerreform" or Umweltsteuer* or "Ökologische Finanzreform" or Umweltabgabe*)

Strategy #2: Eco tax price salience. (Ökosteur* or "Ökologische Steuerreform" or Umweltsteuer* or "Ökologische Finanzreform" or Umweltabgabe*) AND (Dieselpreis or Benzinpreis) AND (Preissteigerung or Preisanstieg or Preiserhöhung or Anstieg or ansteigen or steigen or zunehmen or Zunahme or Erhöhung or erhöhen

or anheben or aufschlagen or Aufschlag or angestiegen or zugenommen or erhöht* or angehoben or aufgeschlagen)

2.E.2 Using information in newspaper articles as an indicator of salience

For each newspaper, we separately downloaded the annual count of articles that are picked up by our search strategies. To account for publishing trends specific to the topic of environmental taxation, we begin by computing a simple newspaper-specific ratio of articles matching Strategy #2 over the frequency counts from Strategy #1. A challenge with these raw article ratios is that the number of articles varies a lot across newspapers and time, making it difficult to simply average the ratios across several newspapers. We, therefore, apply the standardization approach of Baker et al. (2016) to obtain our salience index.

We begin with the simple ratio of articles matching Strategy #2 divided by the total article counts for Strategy #1 for each newspaper, and then divide this ratio by the newspaper-specific standard deviation across all years. This creates a newspaper-specific time series with a unit standard deviation across the entire time interval, which ensures that the volatility of the index is not driven by the higher volatility of a particular newspaper. We then average these standardized series across all newspapers within each year. Lastly, we normalize the yearly series to a mean of 100 over the entire time interval to develop our main salience index. This procedure allows us to explicitly capture variation over time in the price salience of the eco-tax while accounting for newspaper-specific publishing trends concerning the topic of environmental taxation.

Table 2.A15: Effects of salience on gasoline demand (robustness)

	(1)	(2)	(3)	(4)
Raw price of Gasoline (only VAT)	-0.00266 (0.00242)	-0.00282 (0.00179)	-0.00280 (0.00176)	-0.000497 (0.00130)
Energy Tax	-0.00610** (0.00234)	-0.00243 (0.00505)	-0.00338 (0.00489)	-0.00717 (0.00427)
Eco Tax	-0.00656 (0.00492)	-0.0103 (0.00632)	-0.00773 (0.00557)	0.00947 (0.0105)
Eco Tax x Salience Index	-0.00531*** (0.00132)	-0.00433** (0.00203)	-0.00441** (0.00190)	-0.00199** (0.000739)
L.Eco Tax x Salience Index				-0.000459 (0.00197)
L2.Eco Tax x Salience Index				-0.00622*** (0.00208)
Dummy Eco Tax	-0.0227 (0.0399)	0.0323 (0.0868)	0.0135 (0.0738)	-0.195 (0.116)
Trend	0.0153*** (0.00391)	0.0296** (0.0143)	0.0198 (0.0221)	0.0135 (0.0206)
GDP per capita		-0.00274 (0.00310)	-0.00161 (0.00316)	0.000913 (0.00298)
Unemployment rate			0.0101 (0.0256)	-0.00325 (0.0213)
N	38	38	38	37

Table 2.A16: Effects of salience on diesel demand (robustness)

	(1)	(2)	(3)	(4)
Raw price of Diesel (only VAT)	-0.00306*** (0.000766)	-0.00318*** (0.00103)	-0.00326*** (0.000900)	-0.00197*** (0.000620)
Energy Tax	-0.00103 (0.00293)	-0.00537 (0.00337)	-0.00723** (0.00348)	-0.00773** (0.00342)
Eco Tax	-0.0119*** (0.00129)	-0.00998*** (0.00232)	-0.00818*** (0.00275)	0.000528 (0.00367)
Eco Tax x Salience Index	-0.000999* (0.000497)	-0.00123 (0.000732)	-0.00120* (0.000689)	0.000337 (0.000440)
L.Eco Tax x Salience Index				-0.00166* (0.000891)
L2.Eco Tax x Salience Index				-0.00187* (0.000971)
Dummy Eco Tax	0.0821*** (0.0209)	0.0647** (0.0309)	0.0558* (0.0306)	-0.0411 (0.0343)
Trend	0.0356*** (0.00176)	0.0262*** (0.00611)	0.0184* (0.00955)	0.0177 (0.0105)
GDP per capita		0.00131 (0.00101)	0.00210** (0.000995)	0.00272** (0.00106)
Unemployment rate			0.00636 (0.00894)	0.00131 (0.00931)
N	39	39	39	37

Notes: The dependent variable is the log of fuel consumption in liters per capita, which refers to total fuel consumption or either gasoline or diesel consumption (as indicated by the column heading). Prices are in 1995€. Results for gasoline consumption refer to 1972-2009 due to missing price data prior to 1972. Unemployment is measured as percentage of total labor force. Newey-West standard errors in parentheses are heteroskedasticity and autocorrelation robust. Standard errors are calculated relying on the automatic bandwidth selection procedure following Newey and West (1994). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Chapter 3

Carbon pricing, compensation, and competitiveness: Lessons from UK manufacturing

with Elisabeth Isaksen (Ragnar Frisch Centre for Economic Research) and Misato Sato (The London School of Economics and Political Science)

SUMMARY. Carbon pricing is often paired with compensation to carbon-intensive firms, with the aim of mitigating carbon leakage risk. This paper examines the causal impact of compensation payments for indirect carbon costs embodied in electricity prices. We use confidential UK administrative microdata to exploit firm-level inclusion criteria in both a difference-in-differences and regression discontinuity framework. Findings suggest that compensated firms increased production and electricity use relative to uncompensated firms, with no significant effect on energy intensity. While compensation lowers leakage risk, it also implies large forgone opportunity costs of public funds and increased mitigation costs of meeting national emission targets.

3.1 Introduction

Policies to establish a carbon price have proliferated in recent years. Currently, 73 of such initiatives collectively cover 23% of global emissions (The World Bank, 2023). While carbon pricing is considered an essential part of the solution to achieving a cost-effective decarbonization of the economy, there is a long-standing concern that carbon price incentives are being compromised by the concessions offered to industry (Fischer and Fox, 2007; Sterner and Muller, 2008; Rosendahl, 2008). For example, the EU Emissions Trading System (EU ETS) gives energy-intensive sectors free allocation of allowances for the direct costs of carbon emissions. Additionally, some countries also compensate energy-intensive firms for the indirect carbon costs embodied in electricity prices. These cost containment measures make the policy more politically acceptable and are increasingly justified on grounds of alleviating carbon leakage risk¹ (Sato et al., 2022), thus target energy-intensive manufacturing firms operating in regional or global markets with limited ability to pass through carbon costs to consumers (Ganapati et al., 2020).

By shielding firms from the full carbon cost, however, such compensation may compromise efficient carbon price incentives to decarbonize industrial production and consumption. Studies have shown that adjusting free allocation volumes over time can create incentives for polluters to emit more in the present to obtain more free allocations in the future (Rosendahl, 2008)² contrary to earlier claims that market outcomes and efficiency are independent of how allowances were allocated (Montgomery, 1972).³ Compensation linked to current production volumes essentially provides an implicit production subsidy and dampens the carbon price signal (Fischer and Fox, 2007; Fowlie et al., 2016; Meng, 2017), also limiting carbon cost pass through

¹Carbon leakage is often defined as a policy-induced relocation of emissions to countries with more lenient carbon policies.

²This is known as “output-based” allocation and is in contrast to allocation based on historic output or emissions known as “grandfathering” or “ex-ante” allocation.

³Free allocation does not alter the emissions cap and therefore the aggregate effectiveness of a carbon market. However, it is associated with inefficiency losses, as is explained below.

to consumers and foregoing demand side substitution. This means that to achieve the overall emission reduction targets, the mitigation burden shifts elsewhere (to other sectors or towards greater emissions intensity improvements), which means carbon prices and overall costs rise. This perverse production incentive effect has been highlighted in the literature (Fischer, 2001; Demailly and Quirion, 2008; Böhringer et al., 2012; Fischer and Fox, 2011) but downplayed in policy debates arguably due to the lack of robust empirical evidence.

This paper contributes to the literature by empirically examining UK manufacturing firms' responses to an indirect carbon cost compensation scheme. Starting in 2013, the EU ETS allows participating states to partially shield electro-intensive firms from the indirect carbon cost induced by emissions trading, due to carbon cost pass-through in the power sector (European Commission, 2020b). This is expected to continue, for example, Germany, France and Poland have committed to compensating in total an estimated €27.5 billion, €13.5 billion and €10 billion, respectively, between 2021 and 2030 (European Commission DG Competition, 2022). Given the large fiscal implications involved and number of countries compensating indirect carbon costs, there is surprisingly little empirical evidence on their impacts.⁴ Of the Member States providing compensation the UK's compensation was relatively generous because electricity prices reflect relatively high carbon costs induced not only by the EU ETS but also the carbon price floor implemented in 2013 that more than tripled the cost of power sector emissions.

We combine two quasi-experimental research designs, namely a difference-in-difference (DiD) design with inverse propensity score weighting and a "fuzzy" regression discontinuity (RD) design. The two methods complement each other by addressing different types of potential selection biases, and by providing different types of treatment estimates. In both approaches, we exploit the variation caused by the UK

⁴The UK, Germany, Belgium, the Netherlands, Greece, Lithuania, Slovakia, France, Finland, Luxembourg, Poland, Romania Spain, and Norway all provide monetary compensation to electro-intensive firms for higher indirect carbon costs induced by the EU ETS in 2020 (European Commission, 2020b). The total compensation distributed in 2017 by EU countries (for indirect costs incurred in 2016) amounted to €694 million (European Commission, 2018).

eligibility rules for receiving the compensation to identify effects. To be eligible for the program, a firm first needs to operate in a 4-digit NACE industry that is deemed eligible for compensation. Second, eligible firms need to document that the *firm's* overall electricity costs as a share of gross value added (GVA) amounts to at least 5%, where calculations are based on historical values.⁵ Third, the firm needs to apply to the compensation scheme, documenting that it meets the two eligibility criteria. The second and third requirements imply that there are likely both compensated and uncompensated firms operating plants in the same narrowly defined industries, which we can exploit to identify how plants respond to higher indirect carbon costs with and without compensation in place.

To examine how plants respond to indirect carbon cost compensation, we combine confidential microdata from the UK secure data lab on economic variables and energy use at the plant-level with a publicly available list of firms that received compensation. While eligibility for compensation is assessed at the firm level, the amount of compensation paid is calculated at the plant level and is linked to the plant's output. Compared with firm-level analysis, more disaggregated plant-level data are advantageous because firms may operate multiple plants across different sectors. In the analysis, we are comparing similar plants belonging to compensated and non-compensated firms to isolate the effects of compensation for indirect carbon costs, going well beyond previous analysis relying on cross-sectoral or cross-country variation (Ferrara and Giua, 2022).

As a first step, we develop a static conceptual framework to elucidate how compensation payments affect firms' adaptation to indirect carbon costs. The compensation payout is based on historical output multiplied by an electricity intensity benchmark, but if an installation significantly extends (reduces) its production, then baseline output can be increased (reduced) to reflect the capacity or production changes. Our

⁵Note that this criteria is calculated at the firm level, i.e., the legal entity, and not at the plant level. This means that for firms operating multiple plants where some are very electro-intensive while others are not, there may be electro-intensive plants belonging to firms that do not pass the 5% eligibility test.

framework illustrates how, analogous to output-based free allocation in emissions trading, firms receiving compensation for the indirect carbon costs embodied in electricity prices face weaker incentives to contract output, while the incentives to improve electricity intensity of production remain intact. As a consequence, the overall electricity use is expected to increase for compensated firms compared to uncompensated firms.⁶

Our empirical analysis delivers three key results. First, in line with our theoretical prediction, we find that compensated plants increased production relative to non-compensated plants. Results from the DiD estimation show that compensation led to an increase in the sales of own goods by around 16% in the post-treatment period (2013–2015). These results are supported by the fuzzy RD design, where we find a 30% increase in own sales for the compensated plants, with an estimated lower bound of 26% (reduced form estimate). Second, our results point to an increase in electricity use (measured in physical units) of around 22% as a result of compensation accompanied by no significant changes in electricity intensity. Relatedly, we additionally document an increase in carbon emissions of approximately 22% for compensated plants vis-à-vis their uncompensated counterparts.⁷ Finally, we find that energy intensity (scaled by sales) did not experience any significant changes in both the DiD and RD designs. Overall, we find robust evidence in line with our theoretical predictions that incomplete carbon price internalization created by output-based compensation provisions for carbon and energy-intensive industries weakens incentives to reduce output and hence overall energy consumption. Our DiD findings exhibit robustness across a range of tests, including variations in the time frames used to compute p-scores, industry-specific effects defined at different digit levels, sample trimming based on electricity intensity to mitigate the influence of outliers, considering different time horizons in the estimations, extended post-treatment periods, and the utilization

⁶Analytical models on this topic tend to compare one allocation approach over another e.g. Hagem et al. (2020); Fowlie et al. (2016). Our model instead compares the effect of treatment on compensated firms with that on non-compensated firms.

⁷Due to the small sample size, the effects of compensation on electricity use in physical units, carbon emissions and the associated measures of electricity intensity are only produced in the DiD design.

of diverse proxies for production and energy usage. Additionally, our results from the RD design are robust to multiple bandwidth selections and alternative functional forms (Lee and Lemieux, 2010).

Our findings provide several important policy implications for carbon pricing in the UK and elsewhere where free allocation, compensation and exemptions remain commonplace (European Commission, 2020a; The World Bank, 2023).⁸ While carbon leakage may have been limited,⁹ industry compensation represents a substantial forgone carbon tax revenue that could be employed towards driving forward the transition to net zero. We find robust evidence that compensation encourages firms to increase production and thereby pollute more, shifting the mitigation burden elsewhere in the economy where emissions abatement may be costlier (Martin et al., 2014b). Moreover, output-based compensation to industry also limits cost pass through, thus also hindering mitigation through demand-side response (Quirion, 2009). Our results hence underscore the need for complementary measures to encourage consumers to substitute away from energy- and carbon-intensive goods.

Our paper contributes to a broader literature on the incentives effects of industry compensation in climate policy. Free allocation in emissions trading and the distortions that can arise from specific designs of free allocation rules have been extensively studied. For example, *ex-ante* free allocation based on historic activity can generate large windfall profits (Laing et al., 2014) and over-allocation (Martin et al., 2014b), and lead to early action problems, distorting investment decisions or reducing incentives to phase out inefficient technologies (Stern and Muller, 2008; Venmans, 2016) but can be rectified through benchmarking (Neuhoff et al., 2006; Zetterberg, 2014); closure provisions create incentives to delay exit (Verde et al., 2019); combining free allocation with activity thresholds create incentives to artificially inflate output in low-activity installations (Branger et al., 2015). Our

⁸Even in the EU where the Carbon Border Adjustment Mechanism (CBAM) will be introduced in 2026 to reduce the risk of leakage, free allocation is scheduled to continue until 2035 (Morgado Simões, 2023).

⁹Given large volumes of free allocation, it is not surprising that studies on the EU ETS find limited evidence to support leakage (Naegele and Zaklan, 2019; Verde, 2020)

empirical analysis particularly complements literature on output-based free allocation that primarily uses theoretical and modelling approaches and highlights perverse production incentives while improving leakage outcomes (Fischer, 2001; Fischer and Fox, 2007; Demailly and Quirion, 2008; Böhringer et al., 2014). Rosendahl and Storrøsten (2015) show that output-based allocation (OBA) in general gives stronger incentives to improve abatement technology due to a higher permit price but the effects of OBA is heterogeneous across types of firms and sectors. Finally, research has shown that opportunity costs of compensation are high in part because they are coarsely or ill-targeted (Martin et al., 2014b; Fowlie and Reguant, 2022).

Some papers have explored other carbon cost compensation measures including refunding of emission payments (Martin et al., 2014b; Hagem et al., 2020), and relatedly, exemptions and rebates for energy taxes (Ito, 2015; Gerster and Lamp, 2023).¹⁰ On the compensation scheme for indirect carbon costs, to our knowledge, there is only one other empirical analysis (Ferrara and Giua, 2022), but their empirical approach using firms in other countries or sectors without compensation as controls is problematic.¹¹ We are the first paper to rigorously examine the effects of indirect carbon cost compensation.

Our study also complements and expands the knowledge base on how carbon pricing affects carbon and energy-intensive firms (Martin et al., 2014a; Petrick and Wagner, 2014; Aldy and Pizer, 2015; Klemetsen et al., 2020; Marin and Vona, 2021; Dechezleprêtre et al., 2023; Colmer et al., 2023)¹², including the specific papers on the UK Carbon Price Floor (Abrell et al., 2022; Leroutier, 2022). The latter studies examine the direct impact of the policy on decarbonizing the UK electricity sector, while we

¹⁰In contrast to exemptions for energy- and carbon-related taxes, the CO₂ price compensation scheme is designed in a way that aims to restore some of the incentives created by the initial carbon pricing policy. We would therefore expect the mechanism and impacts to differ from an exemption scheme.

¹¹To distinguish the compensation scheme's causal effect from other factors unrelated to the program is difficult under this choice of control group. The countries self-selecting into giving out compensation are likely to be different from other countries in terms of observable and unobservable factors. The sectors selected for compensation have been assessed as energy intensive and at high risk of relocation, hence likely to be different from non-eligible sectors.

¹²See Laing et al. (2014), Martin et al. (2016) and Dechezleprêtre et al. (2023) for EU ETS reviews and Green (2021) for a review of the empirical carbon pricing literature.

study the indirect effects of carbon pricing via higher electricity prices, as well as how these indirect costs are mediated through a compensation scheme.

The remainder of the paper is structured as follows. We first lay out a simple conceptual framework to characterize the compensation scheme's impact on firms in Section 3.2. We then give some essential policy background on the UK carbon pricing and compensation scheme, introduce the data, and provide descriptive statistics in Sections 3.3. Section 3.4 details our two empirical strategies. Section 3.5 presents our main results and compares the estimates from both strategies. Section 3.6 presents some back-of-the-envelope calculations to provide perspective on the trade-offs between preventing leakage and fostering carbon abatement, before we conclude in Section 3.7.

3.2 Conceptual framework

Here we use a simple framework to characterize the theoretical predictions of manufacturing plants' behavior in the presence of indirect carbon costs with and without compensation, drawing inspiration from Hagem et al. (2020) and Fowlie et al. (2016).

Suppose that production causes direct carbon emissions from the combustion of fossil fuels (e_i is the emission intensity or emissions per unit of output q_i for firm i) as well as indirect carbon emissions through the use of electricity (el_i is the firm-specific electricity intensity). Each firm can reduce its overall emissions ($e_i \cdot q_i$) and electricity use ($el_i \cdot q_i$) by reducing production (q_i) and/or by lowering the respective intensities – by installing abatement equipment to lower e_i or electricity saving technology to lower el_i . Suppose that firms face two types of carbon costs. First, firms pay a *direct carbon cost* that depends on the output, q_i , the emission intensity, e_i , and an equilibrium emission permit price, τ , (or more generally, the monetized damages associated with an additional tonne of carbon emissions). Second, firms face an *indirect carbon costs* via carbon embodied in electricity prices that is a function

of output, q_i , the electricity intensity, el_i , and the electricity price, $p_{el}(\tau_{el})$. Note that the electricity price is a function of the carbon tax levied on the power sector: $p_{el}(\tau_{el})$.¹³

We consider a sector that consists of firms indexed by $i=1, \dots, n$ with each firm producing quantity q_i of a homogeneous good, operating in perfectly competitive global markets where all firms are price takers.¹⁴ We apply the standard assumptions that marginal costs of production is positive and increasing: $c'_i > 0$, $c''_i > 0$ and abstract from exit and entry decisions. The profit of a single plant is given by:

$$\pi_i = pq_i - c_i(q_i) - my_i - nz_i - \underbrace{\varphi_i(q_i, e_i(y_i), \tau)}_{\text{Direct carbon costs}} - \underbrace{\psi_i(q_i, el_i(z_i), p_{el}(\tau_{el}))}_{\text{Electricity costs}} \quad (3.1)$$

where p is the product price, $c_i(q_i)$ is the cost of output, q_i – excluding electricity use –, m is the (annuity) price per unit of abatement equipment y_i , and n is the (annuity) price per unit of electricity saving equipment z_i . The parameter $\varphi_i(q_i, e_i(y_i), \tau)$ indicates the direct carbon costs and the parameter $\psi_i(q_i, el_i(z_i), p_{el}(\tau_{el}))$ indicates the electricity costs. The electricity costs include an *indirect* carbon cost component, represented by τ_{el} , which is the carbon price in the electricity sector.

For direct carbon costs, we assume that permits are allocated based on units of production multiplied by an industry-specific emission intensity benchmark, \bar{e}_j i.e. output-based allocation. Thus, direct carbon costs (φ) to the firm will be:

$$\varphi(q_i, e_i(y_i), \tau) = q_i \cdot \tau (e_i(y_i) - \bar{e}_j) \quad (3.2)$$

¹³Assuming 100 % pass-through of carbon taxes in the power sector, the full carbon cost associated with electricity generation is born by the users of electricity.

¹⁴This assumption is in line with the UK Government's underlying assumption of UK firms being unable to pass through domestic carbon taxes to product prices.

(i) No compensation for indirect carbon costs

Under no compensation for indirect carbon costs, the cost of electricity consumption (ψ) to the firm will be:

$$\psi(q_i, el_i(z_i), p_{el}) = q_i \cdot el_i(z_i) \cdot p_{el}(\tau_{el}) \quad (3.3)$$

where τ_{el} is the carbon price faced by electricity generators.¹⁵ Intuitively, any increase (decrease) in τ_{el} or electricity intensity el_i would translate into higher (lower) ψ .

Maximizing the profit function with respect to output q_i and electricity saving investments z_i yields the following first-order conditions:

$$\frac{p - \overbrace{c'_i(q_i) - \tau \cdot [e_i(y_i) - \bar{e}_j]}^{\Delta \text{direct carbon costs}}}{el_i(z_i)} = p_{el}(\tau_{el}) \quad (3.4)$$

$$-\frac{n}{q_i \cdot el'_i(z_i)} = p_{el}(\tau_{el}) \quad (3.5)$$

The left-hand side of Equation (3.4) expresses the marginal cost of reducing electricity use through output reductions, and the left-hand side of Equation (3.5) expresses the marginal cost of reducing electricity use through technology investments.

(ii) Compensation for indirect carbon costs

If compensation is introduced to offset the indirect carbon cost component of electricity prices, based on industry-specific electricity intensity benchmarks (denoted by \bar{e}_j) and baseline output subject to dynamic updating,¹⁶ then it follows that the

¹⁵Assuming complete cost pass-through in the electricity sector, the carbon price faced by UK power plants will be equal to: $\tau_{el} \equiv \tau + \text{Carbon Price Support}$.

¹⁶Both assumptions match how compensation payments for higher electricity costs induced by the EU ETS are calculated across Member States, where baseline output is updated on a quarterly basis and electricity consumption efficiency benchmarks (in MWh/tonne of output and defined at

cost of electricity consumption (ψ) to the firm will be:

$$\psi(q_i, el_i(z_i), p_{el}) = q_i \cdot \left[\underbrace{el_i(z_i) \cdot p_{el}(\tau_{el})}_{\text{Electricity cost per tonne}} - \underbrace{e\bar{l}_j \cdot \tau_{el} \cdot A_i}_{\text{Compensation per tonne}} \right] \quad (3.6)$$

where $e\bar{l}_j$ is the electricity intensity benchmark for industry j (tCO₂/tonne) and A is the aid share.^{17,18}

Conditional on being compensated, maximizing the profit function with respect to output q_i and electricity saving investments z_i yields the following first-order conditions:

$$\frac{p - c'_i(q_i) - \overbrace{\tau \cdot [e_i(y_i) - e_j]}^{\Delta \text{direct carbon costs}}}{el_i(z_i)} = p_{el}(\tau_{el}) - \overbrace{e\bar{l}_j \cdot \tau_{el} \cdot A_i}^{\text{compensation}} \quad (3.7)$$

$$-\frac{n}{q_i \cdot el'_i(z_i)} = p_{el}(\tau_{el}) \quad (3.8)$$

As (3.5) = (3.8), the first-order condition w.r.t. the electricity-saving investment z_i is the same regardless of the compensation payments for the indirect carbon costs. Put differently, the social marginal cost of electricity reduction through technology investments is equal to the level of the electricity price for all firms. However, we see that the first-order condition w.r.t. output, q_i , has changed relative to no compensation: (3.4) \neq (3.7). From (3.7) we see that the marginal cost of lower electricity use through output reductions is no longer equal to the electricity price $p_{el}(\tau_{el})$, but equal to $p_{el}(\tau_{el})$ minus the compensation payments per unit of output ($e\bar{l}_j \cdot \tau_{el} \cdot A_i$).

Prodcom 8 level) are defined as the product-specific electricity consumption per tonne of output achieved by the most electricity-efficient methods of production for the product considered (EU 2012/C 158/04).

¹⁷Over the time frame considered in this paper, the EU Commission recommendations state that aid intensity should not exceed 85% of the eligible costs incurred in 2013, 2014 and 2015 and 80% of the eligible costs incurred in 2016 (EU 2012/C 158/04).

¹⁸In the case of complete pass-through of the power sector carbon price τ_{el} to electricity prices, $A_i = 1$, and $el_i = e\bar{l}_j$, the compensation per tonne received by the firm would equal the increased electricity cost per tonne due to the higher τ_{el} . If instead $el_i < e\bar{l}_j$, compensation payments per unit of output will be larger than the carbon price-induced increase in the electricity price.

Introducing compensation payments for indirect carbon costs increases the cost of reducing electricity use through output reductions, as reduced production leads to lower compensation payments – this marginal loss of compensation via reduced output equals $\bar{e}_{lj} \cdot \tau_{el} \cdot A_i$. Hence, the firm’s marginal cost of reduced output exceeds the social cost of reduced output. While higher electricity prices induced by carbon pricing in the power sector make production more costly, compensation payments make production less costly.

Testable predictions of firms’ production behavior

By comparing models (i) and (ii), we formalize the following hypothesis of how plants respond to an increase in the indirect carbon cost τ_{el} :

- **Prediction 1** *Compensated plants’ production will contract less vis-à-vis uncompensated plants.*
- **Prediction 2** *Compensated and uncompensated plants have the same incentives to invest in electricity-saving technology. Therefore, a similar effect of an increase in τ_{el} on the electricity intensity is expected for compensated and uncompensated plants.*
- **Prediction 3** *Based on predictions 1 and 2, we expect that compensated plants’ overall electricity use will contract less vis-à-vis uncompensated plants*

These predictions compare the effects of output-based compensation on treated and non-treated firms, in contrast with the predictions in previous papers, which compare the effects of output-based compensation on treated firms vis-à-vis other allocation methods such as auctioning or grandfathering (e.g. Fowlie et al., 2016; Rosendahl, 2008).

In the following sections, we empirically test these theoretical predictions by applying a difference-in-difference and regression discontinuity design to the UK indirect

carbon cost compensation scheme. In the next section, we describe the research design and data used in our empirical analysis.

3.3 Research Design and Data

3.3.1 Policy Background

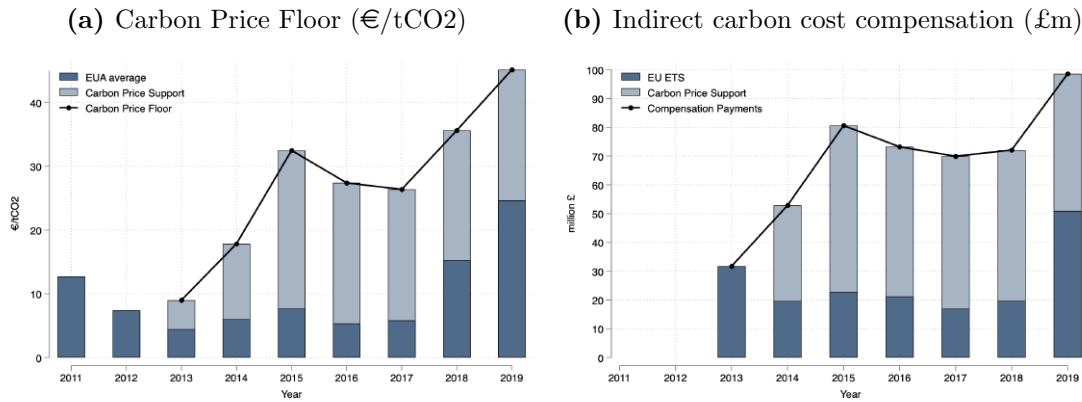
The design of the carbon pricing and compensation schemes plays a central role in our empirical strategy, so it is essential to understand how the relevant policies were rolled out.

In 2005, a EU-wide carbon price was introduced for the manufacturing and power sectors with the introduction of the EU ETS. The carbon price can affect manufacturing firms in two ways. First, regulated firms have to purchase and surrender EU Allowances (EUAs) for each tonne of CO₂ emitted in the previous year (*direct* ETS costs). Second, firms also pay for the carbon price reflected in higher electricity prices (*indirect* ETS costs) due to electricity producers passing forward the carbon price on to consumers (Sijm et al., 2006; Fabra and Reguant, 2014; Hintermann, 2016). To prevent carbon leakage, the ETS Directive gives free allocation to leakage-exposed sectors to limit their exposure to direct carbon costs. Since 2013, “the 2012 Guidelines” also allowed EU ETS countries to grant State aid to compensate selected electro-intensive industries for indirect carbon costs (European Commission, 2020b).

In the UK, in addition to the EU ETS, a Carbon Price Floor was unilaterally introduced on April 1 2013, applying only to electricity generation and immediately raising the carbon price faced by UK power plants. The initial idea of the policy was to first set the desired carbon price floor (path) and then stipulate the tax needed to top up the EUA price with the Carbon Price Support (CPS). From 2016, however, the UK Government decided to freeze the CPS at £18/tCO₂, which meant that the policy effectively functioned as an additional tax on carbon emissions that came on

top of the EUA price. As seen from Panel (a) in Figure 3.1, CO₂ prices faced by power plants were 2-5 times larger than the EUA price.

Figure 3.1: Carbon prices and compensation payments in the UK



Notes: Panel (a) illustrates the two elements of the carbon price faced by UK power plants. For the period 2013 to 2015, the Carbon Price Support, i.e., the tax, was set to 4.94, 9.55, and 18.08 £/tCO₂. From 2016, Carbon Price Support (tax) was frozen at £18/tCO₂. Approximate calculations using the yearly average of EUA prices in €/tCO₂ from sandbag.org.uk, the Carbon Price Support rates in £/tCO₂ from Hirst (2018), and GBP/EUR exchange rates. Panel (b) summarizes the annual compensation payments made by the UK government for EU ETS and CPS indirect carbon costs communicated directly by the Department for Business and Trade through a freedom of information request.

The UK CPS was expected to accelerate the decarbonization of the UK power sector and came in response to the general concern in the years leading up to phase III of the scheme that the EUA price was too low (UK BEIS, 2019); in 2012, the average allowance price was around €7/tCO₂. But simultaneously, it spiked substantial concerns about leakage and loss of competitiveness of UK electro-intensive manufacturing firms vis-à-vis competitors abroad.^{19,20} To mitigate the potential adverse effects on domestic firms and win political support, the CPS was accompanied by a compensation scheme for the additional costs it entailed. This was meant to start in 2013, but was only approved by the EU Commission in March 2014, when it came into effect. This was combined with another compensation introduced

¹⁹Even before the Carbon Price Floor, UK industrial sectors voiced strong concerns about electricity prices for several reasons. Over the past decade, UK manufacturing companies have paid relatively high electricity prices compared to their counterparts in neighboring countries such as France, Germany, and Italy, but the differences are mitigated by compensation for policy costs; see Figure 3.A1 in Appendix 3.A. Electricity has been the main source of energy in the UK manufacturing sector as a whole since 2006 (UK BEIS, 2018).

²⁰Grubb and Drummond (2018) quantify the relative contribution of various components to UK industrial prices. They stipulate that costs induced by the CPS and the EU ETS accounted for approximately 25% of the industrial electricity price in 2016. Cambridge Econometrics (2017) report a lower number: As a proportion of the industry electricity price in 2016, the indirect EU ETS carbon cost and the CPS amounted to around 9%.

in January 1, 2013 for the indirect carbon costs induced by the EU ETS. Since 2013, carbon prices have been higher in the UK due to the CPS (see Figure 3.A1 in Appendix 3.A), but so were the compensations. Panel (b) in Figure 3.1 summarizes the annual payments made by the UK government for compensation for EU ETS and CPS indirect carbon costs, demonstrating an upward trend in correlation with the rise in the price of EUA allowances in more recent years.

Eligibility

We exploit a discontinuity in the eligibility rules for indirect carbon cost compensation to test the effect of the compensation on firms' economic and environmental outcomes. Eligibility for compensation for the indirect costs of both the EU ETS and the Carbon Price Support was based on two criteria. First, the firm needs to manufacture a product in the UK within an eligible sector defined by the 4-digit NACE code. The European Commission selected a list of eligible sectors with a high risk of carbon leakage.²¹ Appendix Table 3.A1 lists the 15 eligible industries according to the 4-digit NACE code (European Commission, 2012).²²

With the aim of a more targeted compensation scheme, the UK Government also imposed a second eligibility criteria: a firm needs to show that its indirect carbon costs (the combined costs of EU ETS and the Carbon Price Support) would amount to 5% or more of its gross value added. Specifically, this so-called 5 % filter test was calculated in the following way:

$$\frac{\text{electricity consumption (MWh)} \times \text{price impact (£/MWh)}}{\text{Gross Value Added (£)}} \geq 5\%, \quad (3.9)$$

²¹Generally, the compensation schemes need to comply with the principles set out in the Environmental and Energy Aid Guidelines and the ETS State and Guidelines adopted by the European Commission. The first set of guidelines states that Member States are allowed to partially compensate large electricity users for the indirect costs of taxes on energy products, when those taxes have the same aim and effect as the ETS carbon allowance price. The criteria for choosing eligible firms and calculating compensation levels need to be the same as those in the ETS State aid Guidelines.

²²This list was subsequently revised down from 15 to 10 in 2020.

where electricity consumption and gross value added (GVA) are average values for the period 2005-2011, and the price impact was set to £19/MWh in real 2007 prices. As calculations were based on historical values, there was a limited ability for firms to adjust consumption or production to ensure that they were eligible for compensation. Both electricity costs and GVA had to be calculated at the aggregate legal entity level, i.e., the firm. For multi-plant firms, this implied that parts of the electricity use and GVA might stem from activity unrelated to the manufacture of the eligible product(s). If these activities were less energy intensive, it would lower the firm's average electricity intensity, and hence make it harder to meet the eligibility criteria.

Even if a firm meets the two criteria, it also needs to submit an application to receive the compensation. Crucially for identification, the multiple criteria implies that we might have three types of firms within a narrowly defined eligible industry: (i) firms that passed the 5% filter test and received compensation, (ii) firms that *would* pass the 5% filter test, but did not apply, (iii) firms that did not pass the 5% filter test. This makes it possible to exploit within-industry variation to estimate impacts of the compensation scheme.²³

Compensation calculation

While the 5% test requires a calculation at the aggregate *firm* level, the amount of compensation is calculated based on *installation* level data.^{24,25} Compensation

²³In addition to the criteria listed, a firm was also eligible for compensation if it could document that a close competitor received compensation. A close competitor is defined as a firm producing the same product, as defined by the 8-digit Prodcom classification. Additionally, a firm is also granted compensation if it can demonstrate that it failed the 5% test because of the inclusion of business activity that did not relate to the manufacture of the eligible product(s).

²⁴In the compensation scheme an installation was defined as a stationary technical unit where one or more activities associated with the manufacture of the eligible product are carried out.

²⁵It is then possible that two plants have the exact same electricity intensity (el/GVA) but only one of the plants are eligible for compensation because the plant's owner firm passes the eligibility test. The ineligible plant might be part of a multi-plant firm, where the other plants are less energy intensive. Generally, we would expect that firms with a secondary industry code that makes them eligible are less likely to receive compensation compared to firms with a primary industry code that is eligible.

payments based on installation-level data are calculated using the following formula:

$$\begin{aligned} & \text{Baseline output of product X (tonne)} \times \\ & \text{Electricity consumption efficiency benchmark (MWh/tonne)} \times \\ & \text{Emission factor (tCO}_2\text{/MWh)} \times \\ & [\text{Carbon Price Support (£/tCO}_2\text{)} + \text{EUA forward price at year t-1 (£/tCO}_2\text{)}] \times \\ & \text{Aid share (e.g. 80\%).} \end{aligned} \tag{3.10}$$

The baseline output corresponds to the average production of the eligible product in tonnes per year at the installation over the reference period 2005–2011. However, if an installation significantly extended its production, the baseline output could be increased in proportion to the production extension. Also, if an installation significantly reduced its production, the aid would be reduced according to a stepwise function.²⁶ Payments to firms are made quarterly, and firms were required to inform the UK Government quarterly of any significant increases or reductions in their production. There is hence a degree of dynamic updating of the baseline, which means that compensation payments can potentially be affected by a firm’s recent production.

3.3.2 Data sources

To examine the indirect effect of carbon pricing on manufacturing, we combine several data sources at the firm and plant levels, primarily confidential microdata from the UK secure data lab. While the disaggregated data offers rich detail, it also poses challenges for analysis due to the relatively small sample size because some

²⁶If production was reduced by less than 50%, there would be no reduction in the aid amount. If reduced between 50% and 75%, an installation would only receive 50% of the aid amount. If reduced by 90% or more, and installation would not receive any compensation. Conditional on eligibility, there may be perverse incentives around the thresholds to artificially inflate production especially during economic downturns in order to receive full compensation as documented in the case of ETS free allocation by Branger et al. (2015).

data sources are surveys.

Compensation schemes: A list of firms that received compensation for indirect carbon costs between 2016 and 2019 is publicly available from the Department for Energy Security and Net Zero (DESNZ)²⁷ website. We assume that the same firms also received compensation for the years 2013 to 2015.²⁸ There were in total 59 firms that received compensation in 2016 for the indirect costs induced by the EU ETS and the Carbon Price Support.

Economic data: We use plant-level data²⁹ on employment and economic outcomes from restricted microdata maintained by the Office for National Statistics (ONS). Our core dataset is the Annual Business Survey (ABS), which is an annual survey of businesses covering production, construction, distribution, and service industries. ABS is the largest business survey conducted by the ONS and covers around 62,000 plants. The sample design is a stratified random sample using three stratification variables: employment, geography, and the 4-digit Standard Industrial Classification (SIC) code. From the ABS, we collect information on SIC codes, employment, sales of own goods, production value, turnover, gross value added (GVA), and energy expenditures for the period 2005 to 2019.^{30,31} Monetary values are adjusted for inflation, with 2010 serving as the base year, based on official inflation statistics.

²⁷Formerly Department for Business, Energy & Industrial Strategy (BEIS)

²⁸While information on which firms received compensation before 2016 is not publicly available, we were told in conversations with the former Department for Business, Energy & Industrial Strategy (BEIS) that it is safe to assume that the list of firms are approximately the same as for 2013-2015.

²⁹A “plant” corresponds to a “reporting unit”, which holds the mailing address for the business and is the unit for which businesses report their survey data to the UK Office for National Statistics. A reporting unit represents an aggregation level that is more granular than an “enterprise unit” (which may be subdivided into several reporting units) and more aggregated than a “local unit” (which may be combined to form one reporting unit to reduce compliance costs). It is the lowest aggregation level for which most business data are available. Within our sample, around 16% of compensated enterprise units represent multi-plant firms. For more details see Criscuolo et al. (2003).

³⁰This includes a period of economic turmoil following the 2016 EU Referendum in the UK. We show in robustness tests that the results are consistent including 2016-2019.

³¹The ABS was merged with the names of compensated firms by first manually matching the compensated firms names with Bureau van Dijk’s Orbis data to obtain the Company Registration Number (CRN), which can then be linked to the company IDs in the confidential data (Enterprise Reference Number).

Energy and Electricity use: To examine how electricity use is impacted by carbon pricing and the compensation scheme, we collect detailed information from the Quarterly Fuels Inquiry (QFI). The QFI provides quarterly information on the value and the quantity of fuels used by a small sub-sample of UK manufacturing plants. Before 2008 the survey covered around 1200 plants, while after 2008, the survey only covered around 600 plants. The survey is maintained by the ONS on behalf of the DESNZ. Unfortunately, this data is not available beyond 2015. Observations are aggregated to the annual level and then linked to the ABS. Because the QFI covers a smaller sample than the data on economic variables and is not available beyond 2015, to have sufficient power to test some of our hypotheses, we rely on reported energy costs from the ABS as a proxy (see section 3.3.4).

Electricity related indirect emissions: To calculate indirect carbon emissions embodied in electricity, we combine detailed electricity use in physical units from the QFI with emission factors provided by the UK DESNZ.³²

3.3.3 Descriptive statistics

Table 3.1 is based on plant-level microdata from the ABS and the QFI and shows summary statistics by compensation status. The sample is restricted to manufacturing industries (SIC 7-33).

To test our first hypothesis on the effect of compensation on production, we use sales of own goods as our main dependent variable and proxy for production volumes, and other proxies including total output, GVA, and total turnover in robustness checks (see Panel A). Given that protecting jobs is a frequently used argument to justify compensation, we also examine the effects of the compensation scheme on employment but regard this outcome as less tightly linked to production volumes. Comparing compensated and non-compensated plants, we see from Panel A in Table

³²Government conversion factors for company reporting of greenhouse gas emissions can be found here.

3.1 that compensated plants are larger than the non-compensated manufacturing plants in terms of both production, employment, and gross value added. We also see that there is a limited number of compensated plants in our sample, ranging from 70 to 119 depending on the variable of interest. By contrast, the number of non-compensated manufacturing plants is between 8,976 and 16,180 plants.

To test our second hypothesis on electricity intensity impacts, we focus on electricity use in kWh (from the QFI) as a share of sales of own goods (Panel B) and energy purchases as a share of sales (Panel A), we also provide results for a wide range of intensity measures in robustness checks.

To test the third hypothesis on the effects on electricity consumption, we focus on electricity use in kWh (Panel B) as the variable that is closest to what we would like to test. However, due to the smaller sample size in the QFI (34 compensated plants and 739 non-compensated plants), we also examine the effects on the energy purchases variable from the larger ABS (Panel A) even if this variable is a proxy for electricity use.

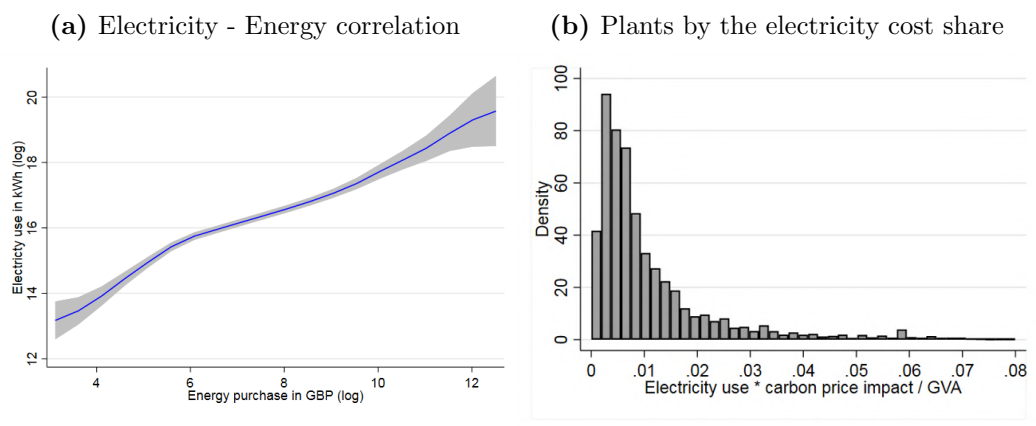
Comparing compensated and noncompensated plants, we see that compensated plants are larger, use more energy, and are more energy-intensive than noncompensated plants. Clearly, we need to account for this selection bias in our estimation in order to recover causal estimates of the compensation scheme.³³ Table 3.1 also highlights the challenge we face in terms of sample size, with the limited number of compensated plants in our sample relative to the number of noncompensated manufacturing plants particularly for the QFI sample.

³³Additional descriptive evidence on our key outcome variables, including plots showing the development in variables over time, are provided in Appendix 3.B.

3.3.4 Using predicted electricity use to calculate the eligibility criterion

One key data challenge we face is the limited availability of plant-level data on electricity consumption. The QFI is a relatively small sample and data on electricity use in kWh is only available for a small subset of plants (Table 3.1, Panel B) up to 2015. To circumvent this problem, we use the relationship between energy purchases (in £) from the ABS and electricity use (in kWh) from the QFI sample to predict electricity consumption for the larger ABS sample up to 2019. Panel (a) in Figure 3.2 shows the strong and positive relationship between electricity use and total energy purchases. The raw correlation ranges from 0.91 to 0.93, depending on sample restriction (see Table 3.A2 in Appendix 3.C where we detail the procedure used to make out-of-sample predictions of electricity consumption).³⁴

Figure 3.2: Predicting electricity consumption from energy purchases



Notes: Panel (a) plots the correlation between log electricity use and log energy purchase in 2011, with 95% confidence interval and local smoothing. Panel (b) shows the distribution of plants by the electricity cost share, using the formula outlined in Equation 3.9 and predicted electricity use. Data source: the Annual Business Survey (ABS) and the Quarterly Fuels Inquiry (QFI). The population is restricted to plants in SIC 7-33 industries.

Predicted electricity consumption is then used to calculate the electricity cost intensity for all plants in the sample, to evaluate the eligibility criterion described in Equation 3.9. As we will see in Section 3.4, having a measure of the electricity cost intensity is important in the empirical strategies we use (as a matching variable in the DiD

³⁴In robustness tests, we use energy purchases directly to calculate an energy cost intensity, and instead infer the likely cut-off value; see Appendix 3.C.

estimation and as the running variable c_i in the RD design). Note that we do not use predicted electricity use, or any variable derived from predicted electricity use, as an outcome variable in the main analysis presented in Section 3.5.

Panel (b) in Figure 3.2 plots the distribution of the calculated electricity intensity criteria based on Equation 3.9. We see that most firms' intensity is much lower than 5%. There is also no detectable bunching right above the 5% criterion, which suggests that plants are not able to manipulate the running variable c_i (see Section 3.4.2 for more details). A McCarty test also gives no indication of bunching at the 5% eligibility cut-off; see Table 3.A20.

3.4 Empirical Strategy

Faced with challenges around selection bias and sample size, our approach to examining the indirect impacts of carbon pricing via electricity prices on manufacturing firms with and without compensation schemes in place is the following. Acknowledging that no single approach can adequately overcome all threats to identification, we pursue two empirical strategies: i) a difference-in-differences (DiD) strategy with inverse propensity score weighting and industry-specific time trends, and ii) a “fuzzy” regression discontinuity (RD) design, where we exploit the discontinuous jump in the probability of receiving compensation at the eligibility thresholds. We then compare the results from the two strategies.

3.4.1 Difference-in-differences

Our first strategy is to exploit variation within narrowly defined industries in a difference-in-differences (DiD) framework. When $Comp_{ijt}$ is a dummy that indicates if firm i in industry j receives compensation payments at time t , the DiD estimator

is written as:

$$y_{ijt} = \beta_1 Comp_{ijt} + X'_{ijt} + \gamma_i + \delta_{jt} + \epsilon_{ijt}, \quad (3.11)$$

where y_{ijt} is a placeholder for a relevant plant-level outcome (e.g., production, electricity use, or electricity intensity). X'_{ijt} is a vector of plausibly exogenous covariates, γ_i are firm-specific fixed effects, and ϵ_{ijt} is the idiosyncratic error term. The main identifying assumption is that, in the absence of compensation payments, the compensated and uncompensated firms would have followed parallel trends in the outcome variable. One potential threat to identification is industry-specific shocks. By including industry-specific time dummies, δ_{jt} , we absorb time-varying shocks at the 3-digit industry level, which means that identification is based on variation within narrowly defined industries.³⁵

However, there is still the possibility of selection bias within industries across treated and non-treated groups, such as systematic differences in electricity intensity. To account for such within-industry differences in observables, we combine the DiD design with inverse propensity score weighting. Specifically, we use a propensity score estimator to reweight plants in Equation 3.11 to reflect the differences in the probability of getting compensation. We estimate the propensity score (\hat{p}) based on a proxy of the pre-treatment electricity intensity, and lagged values of the outcome variable. On the former, we estimate the propensity score based on an electricity intensity measure that is as similar as possible to the eligibility criteria (see Equation 3.9 in Section 3.3.1) where electricity intensity is defined relative to firm-level GVA. As mentioned, due to the small sample size of the QFI, where electricity use is reported, we instead use predicted electricity use to calculate the eligibility criteria; (Section 3.3.4). The propensity score is calculated separately for each 3-digit SIC industry, based on the period 2005-2011. These years correspond to the period used by the UK Government to calculate the electricity cost share, which again determines whether a plant passes the 5% filter test. The propensity score estimates

³⁵We also show effects for 2 digit industries in robustness checks; see Section 3.5.1. Due to the small sample size, there is a trade-off between accounting for detailed industry-specific trends and ensuring that we have sufficient observations to recover precise estimates.

are then transformed into weights and used in panel regressions. Specifically, we weight each compensated plant by $1/\hat{p}$, and weight each uncompensated plant by $1/(1 - \hat{p})$. This allows us to recover an estimate of the average treatment effect (ATE) of compensation on the outcome of interest (Imbens, 2004).³⁶

To verify if pre-treatment trends are parallel and to examine how the treatment unfolds over time, we also estimate a dynamic version of the DiD with leads and lags. Specifically, we interact the treatment variable, $Comp_{ijt}$, with time dummies, where we use the year before the first treatment year as the reference category. If we denote M as the number of leads and K as the number of lags, we can estimate the unfolding of the treatment with the following regression:

$$y_{ijt} = \sum_{m=0}^M \beta_{-m} Comp_{ijt-m} + \sum_{k=1}^K \beta_{+k} Comp_{ijt+k} + X'_{ijt} \beta_2 + \gamma_i + \delta_{jt} + \epsilon_{ijt}, \quad (3.12)$$

where lead m captures potential deviations in the pre-treatment m years before treatment and lag k captures the effect of the policy k years after the start of the treatment.

Even if pre-treatment trends are parallel, and we ensure that any differences in initial electricity intensity are accounted for, there might still be a component of non-random self-selection into the compensation scheme that influences the development in production, energy use, and financial performance in the post-intervention period. For example, as firms applying to the compensation scheme will likely incur fixed costs in preparing the necessary accounting and administrative work, firms with lower levels of electricity use (but still above the eligibility threshold) might find it too costly to apply. While in principle selection effects can be addressed by adding additional (time-varying) control variables and matching on additional pre-treatment observables, selection might in part be driven by unobserved factors. It is therefore difficult to fully account for potential self-selection effects.

³⁶This approach avoids discarding non-matching observations, retaining a larger estimation sample and hence greater statistical power for inference. See e.g., Guadalupe et al. (2012) for a similar approach.

Table 3.1: Summary statistics for the period 2005–2011, by compensation status

	Compensated	N	Other	N	Difference
Panel A: Variables from the Annual Business Survey (ABS)					
Sales of own goods	10.35 (1.506)	111	7.086 (2.223)	14770	3.264*** (0.211)
Total output	10.36 (1.457)	112	7.208 (2.182)	15503	3.149*** (0.207)
Total turnover	10.38 (1.436)	112	7.303 (2.167)	15713	3.073*** (0.205)
Production value	10.98 (1.219)	70	7.157 (2.483)	8976	3.819*** (0.297)
GVA (Market Prices)	16.04 (1.447)	118	13.37 (2.074)	15463	2.662*** (0.191)
Employment	5.157 (1.161)	119	2.925 (1.694)	16180	2.233*** (0.156)
Productivity (turnover / employment)	5.432 (0.806)	118	4.345 (0.900)	15708	1.087*** (0.0831)
Energy purchases (£)	6.583 (1.791)	99	3.275 (2.232)	15272	3.308*** (0.225)
Energy purchases /Sales	-3.124 (0.879)	118	-3.898 (0.893)	14284	0.773*** (0.0825)
Energy purchases /Output	-3.192 (0.868)	120	-4.035 (0.931)	15039	0.843*** (0.0853)
Energy purchases /Turnover	-3.288 (0.989)	121	-4.118 (0.917)	15278	0.830*** (0.0837)
Energy purchases /Production	-3.052 (1.094)	77	-3.989 (1.014)	8590	0.937*** (0.116)
Energy purchases /GVA	-8.854 (1.175)	119	-10.15 (1.169)	14934	1.299*** (0.108)
Energy purchases /Employment	1.789 (1.324)	109	0.241 (1.184)	15310	1.547*** (0.114)
Panel B: Variables from the Quarterly Fuels Inquiry (QFI)					
Electricity use (kWh)	17.44 (1.796)	33	14.99 (1.768)	729	2.451*** (0.315)
Electricity use / Sales	6.148 (1.044)	32	4.974 (1.213)	706	1.174*** (0.218)
Electricity use / Output	6.005 (0.942)	30	4.891 (1.197)	707	1.114*** (0.221)
Electricity use / Turnover	5.839 (1.196)	33	4.785 (1.230)	726	1.054*** (0.219)
Electricity use / Production	6.145 (0.967)	25	5.224 (1.125)	321	0.921*** (0.232)
Electricity use / GVA	0.530 (1.306)	32	-0.992 (1.331)	720	1.522*** (0.240)
Electricity use / Employment	11.53 (1.321)	31	9.712 (1.353)	728	1.819*** (0.248)
Electricity emissions	24.31 (1.288)	24	22.30 (1.356)	287	2.019*** (0.287)
Panel C: Variables that are calculated based on the ABS and QFI					
Predicted electricity use (kWh)*	15.05 (1.563)	93	12.01 (2.141)	15248	3.037*** (0.222)
Electricity intensity based on Eq. (3.9)	-3.889 (0.942)	116	-5.110 (0.980)	7130	1.220*** (0.0917)

Notes: The table shows summary statistics at the plant level for the period 2005–2011, which is the baseline period used to determine eligibility for the compensation scheme. All variables are in logs. The sample is restricted to manufacturing industries (SIC 7-33). N refers to the number of plants. Source: ABS and QFI. *See Section 3.3.4 for details.

3.4.2 Fuzzy regression discontinuity design (DiDiD-IV)

In an alternative empirical approach, we take advantage of thresholds that influence the eligibility for treatment to identify causal effects.³⁷ In our setting, we can exploit that there is a change in the probability of treatment at two eligibility thresholds: (i) the industry code, and (ii) electricity costs are at least 5% of GVA over a baseline period. While these two thresholds may not perfectly determine whether a firm gets compensation, they still create a discontinuity in the *probability* of treatment.

The intuition behind a fuzzy RD is related to the instrumental variable strategy, and the fuzzy RD can be estimated using two-stage least squares. When $comp_{ijt}$ is a dummy that indicates if firm i in industry j receives compensation payments in year t , then the first stage, reduced form, and the second stage are:

First stage:

$$comp_{ijt} = \pi_1 \underbrace{post_t \times \mathbb{1}\{c_i \geq c_0\} \times \mathbb{1}\{elig_j = 1\}}_{Instrument} + X'_{ijt}\beta + \gamma_i + \mu_{ijt} \quad (3.13)$$

Reduced form:

$$y_{ijt} = \pi_2 \underbrace{post_t \times \mathbb{1}\{c_i \geq c_0\} \times \mathbb{1}\{elig_j = 1\}}_{Instrument} + X'_{ijt}\beta + \gamma_i + e_{ijt} \quad (3.14)$$

Second stage:

$$y_{ijt} = \beta_1 \widehat{comp}_{ijt} + X'_{ijt}\beta_2 + \gamma_i + \epsilon_{ijt}, \quad (3.15)$$

where $post_t$ is equal to 1 for the year 2013 and onwards and 0 otherwise, $\mathbb{1}\{elig_j = 1\}$ indicates if a plant operates in a 4-digit industry eligible for compensation, and $\mathbb{1}\{c_i \geq c_0\}$ indicates if a plant's electricity intensity is above the eligibility cut-off

³⁷In general, regression discontinuity designs (RD) can be either sharp or fuzzy. A sharp RD exploits the fact that passing a specific cut-off value deterministically leads to treatment. By contrast, a fuzzy RD allows for a smaller jump in the probability of assignment to treatment at the threshold (Imbens and Lemieux, 2008).

c_0 (Equation 3.9 in Section 3.3.1). c_i is often referred to as the "assignment" or "running" variable. When the running variable exceeds the cut-off value, c_0 , it induces a change in the probability of a plant receiving compensation. In our context, higher electricity intensity increases, by definition, the likelihood that a plant i will be closer to the cut-off. If the compensation scheme matters, this will induce a change in the outcome variable, y_{ijt} at the cutoff. As not all plants that are eligible receive compensation payments, the change in the outcome variable at the cut-off needs to be rescaled by the jump in the probability of treatment, i.e.,: $\beta_1 = \frac{\pi_2}{\pi_1}$. The estimate corresponds to β_1 in the second stage estimation (Eq. 3.15). Using a 2SLS framework, we can estimate a weighted local average treatment effect (LATE) for the compensated firms, where the weights reflect the *ex-ante* likelihood that plant i is near the threshold (Imbens and Lemieux, 2008). This represents a LATE for a small subgroup of the sample composed of highly electricity-intensive firms close to the 5% cut-off and is therefore not directly comparable to the ATE, which is evaluated based on the entire population of plants. As shown in Figure 3.2b, the 5% threshold is in the right tail of the electricity intensity distribution. Therefore, the subsample of observations used for estimating the LATE represents a small group of highly electricity intensive plants.

Note that the increased probability of receiving compensation as the electricity intensity crosses the eligibility cut-off ($\mathbb{1}\{c_i \geq c_0\}$) only applies to plants operating in eligible industries ($\mathbb{1}\{elig_j = 1\}$). The instrumental variable (IV) is hence the interaction between these two indicator variables. By including $post_t \times \mathbb{1}\{elig_j = 1\}$ and $post_t \times \mathbb{1}\{c_i \geq c_0\}$ in the vector of covariates X'_{ijt} , we allow for eligible industries and plants with an electricity intensity above the cut-off c_0 to develop differently over time.³⁸ By exploiting variation along three dimensions (pre and post, eligible and non-eligible industries, above and below the electricity cut-off), the empirical

³⁸We allow for several different functional forms of $post_t \times \mathbb{1}\{c_i \geq c_0\}$ in our regressions; linear and 2nd degree polynomial distance from the cut-off and equal-sized bins on each side of the cut-off. To control for $post_t \times \mathbb{1}\{elig_j = 1\}$, we combine $post_t$ with a dummy variable indicating if the 2-digit SIC industry is eligible for compensation. We include the control at the 2-digit level, as including industry-specific trends at the 3 or 4 digit level is too demanding and leaves us with very little identifying variation.

strategy could also be interpreted as a difference-in-difference-in-difference (DiDiD) combined with instrumental variables (IV).

A causal interpretation of β_1 relies on several identifying assumptions. First, the probability of treatment has to jump at the cut-off, c_0 . This assumption is usually evaluated by looking at the first stage (see Section 3.5.2). The second identification assumption is that plants cannot manipulate the running variable, c_i , which in our case is the industry code and the electricity cost share. The latter is based on historical electricity consumption and gross value added and is therefore difficult to manipulate. A McCarty test also shows no sign of bunching around the threshold value (see Section 3.3.4). Industry codes are assigned to plants and should in principle not be manipulable. Third, we must assume monotonicity, i.e., that crossing the threshold cannot simultaneously cause some units to get compensation and others to move out of the compensation scheme.

Beyond these identifying assumptions, one obvious threat to identification is the small sample size, especially the small number of compensated plants included in the QFI. Given the limited number of observations close to the threshold in our data, we are forced to increase the bandwidth. This introduces the possibility of increased bias, given that a wider bandwidth increases the likelihood of systematic differences between firms positioned above and below the cut-off.

3.5 Treatment effects of the indirect carbon cost compensation

3.5.1 DiD estimates of the average treatment effects

Tables 3.2 and 3.3 present the main results from the DiD estimation (Equation 3.11) using data from the ABS and QFI, respectively. We additionally report p-values

from a mean comparison test of lagged outcomes categorized by treatment status to present corroborative evidence on the robustness of the parallel trend assumption after IPW. To recall, the ABS sample is larger than the QFI but energy purchase is used as a proxy for electricity consumption. The treatment group is defined as plants belonging to a firm that received compensation for the indirect carbon costs induced by the EU ETS and the UK Carbon Price Support. In all regressions, the sample is restricted to manufacturing industries (SIC 7-33) and plants with at least one observation in the post-treatment period.

First in terms of production, in line with Prediction 1, our results indicate that compensation led to an increase in our main proxy indicator “sales of own goods” by around 16% in the post-treatment period. This estimated effect is based on a comparison of compensated and non-compensated plants with similar electricity intensity and sales figures in the pre-treatment period; see column (1) in Table 3.2. The estimated treatment effect is robust across a number of tests which are presented in Section 3.5.1. In other words, our results suggest that compensation is doing its job in combating the displacement of production and carbon leakage that could arise from climate policy induced electricity price differentials. Interestingly, we do not find any significant effect on employment (cf. Table 3.A6), productivity or GVA (cf. Figure 3.A7). In other words, our results fail to support claims that carbon pricing or higher energy costs lead to job losses.

In terms of electricity intensity, both our results using the QFI (electricity use/sales, Table 3.3 column 2) and ABS (Energy purchase/sales, Table 3.2 column 3) that the difference between compensated plants and non-compensated plants is not statistically significant. This is in line with Prediction 2.

As production is higher, we find broadly that overall electricity consumption is also higher for compensated firms, broadly in line with Prediction 3. In other words, the compensation is dampening the effect of the carbon price signal on discouraging energy use and therefore emissions. Estimates using actual electricity use data from

the QFI (Table 3.3 column 1 and 3) indicate that compensation increased electricity use by 22%, and electricity-related carbon emissions by 23%. Instead when using energy purchases data from the ABS as a proxy, we find a positive effect that is not statistically significant (Table 3.2, column 2).

Table 3.2: Average treatment effects of compensation. 2010–2015.

	Source: ABS		
	Sales of own goods	Energy purchases	Energy intensity
	(1)	(2)	(3)
Compensation	0.156** (0.0638)	0.300 (0.182)	-0.123 (0.102)
Observations	532	303	688
N Compensated	27	14	27
N Other	97	65	157
Plant FE	✓	✓	✓
Year×Industry FE (3-digit SIC code level)	✓	✓	✓
Mean electricity intensity 05-11: compensated	0.035	0.036	0.031
Mean electricity intensity 05-11: other	0.035	0.033	0.036
P-value: mean-comparison test	0.725	0.524	0.107
Mean outcome pre-treatment: compensated	11.135	7.346	-3.147
Mean outcome pre-treatment: other	11.147	7.327	-2.980
P-value: mean-comparison test	0.944	0.961	0.248

Notes: Table shows the coefficient β_1 estimated from Equation 3.11. Dependent variables are in logs. Standard errors are clustered at the firm level. All regressions include year \times industry fixed effects at the 3-digit SIC code level, and are weighted by the inverse propensity score. We drop plants with an electricity intensity based on Eq. (3.9) below 0.01. Data sources: Annual Business Survey (ABS) and Quarterly Fuels Inquiry (QFI). See reference list for full citation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.3: Average treatment effects of compensation. 2010–2015.

	Source: QFI		
	Electricity use	Electricity intensity	Indirect CO ₂ emissions
	(1)	(2)	(3)
Compensation	0.220** (0.0900)	0.140 (0.189)	0.225** (0.0884)
Observations	413	598	426
N Compensated	15	16	14
N Other	65	106	68
Plant FE	✓	✓	✓
Year×Industry FE (1-digit SIC code level)	✓	✓	✓
Mean electricity intensity 05-11: compensated	0.036	0.034	0.037
Mean electricity intensity 05-11: other	0.037	0.034	0.037
P-value: mean-comparison test	0.795	0.958	0.583
Mean outcome pre-treatment: compensated	17.305	5.936	23.491
Mean outcome pre-treatment: other	17.392	5.995	23.541
P-value: mean-comparison test	0.607	0.706	0.795

Notes: Table shows the coefficient β_1 estimated from Equation 3.11. Dependent variables are in logs. Standard errors are clustered at the firm level. All regressions include year \times industry fixed effects at the 1-digit SIC code level, and are weighted by the inverse propensity score. We drop plants with an electricity intensity based on Eq. (3.9) below 0.01. Data sources: Annual Business Survey (ABS) and Quarterly Fuels Inquiry (QFI). See reference list for full citation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

ATEs over time

We also present the dynamic version of the DiD (Figure 3.A9, which plots the annual DiD coefficients estimated from Equation 3.12 and shows how treatment effects unfold over time. It also shows the validity of parallel pre-treatment trends leading up to 2013 when compensation was first paid out (for indirect costs incurred in 2012) – the same year as the introduction of the UK Carbon Price Support in the UK power sector. Figure 3.A9, Panel (a) shows that difference in production levels between compensated and non-compensated firms emerged already in 2013, but grew more in 2014. Figure 3.A9, Panel (b) instead shows that for electricity intensity (proxied by energy purchases over sales), the gap widened in 2013 but closed in subsequent years.

Our main estimates are based on a post-treatment period that ranges from 2013 to 2015 as this is the only estimation window where information both from the ABS and the QFI is available. Nevertheless, ensuring comparability across results for different variables comes at the expense of shrinking the estimation sample size. Tables 3.4 and 3.5 provide additional results for outcome variables that are available beyond that period to corroborate our findings from Table 3.2. The corresponding results for employment are presented in Table 3.A6 in the Appendix.

Table 3.4: ATEs of compensation on sales. 2010–2019.

Sales of own goods					
	2015	2016	2017	2018	2019
Compensation	0.156** (0.0638)	0.164** (0.0763)	0.126* (0.0705)	0.147** (0.0701)	0.144** (0.0693)
Obs	532	717	851	1069	1186
N compensated	27	36	39	40	40
N other	97	127	132	156	158
Energy intensity 05-11 (Treat)	0.035	0.035	0.035	0.034	0.034
Energy intensity 05-11 (Control)	0.036	0.037	0.036	0.036	0.036
p-value (mean-comparison test)	0.725	0.424	0.628	0.596	0.597
outcome pre-treatment (Treat)	11.135	10.867	10.831	10.744	10.743
outcome pre-treatment (Control)	11.147	10.905	10.906	10.844	10.844
p-value (mean-comparison test)	0.944	0.822	0.664	0.558	0.553

Notes: Table shows the coefficient β_1 estimated from the DiD equation. Dependent variables are in logs. Standard errors are clustered at the firm level. All regressions include year x industry fixed effects at the 3-digit SIC code level and are weighted by the inverse propensity score. We drop plants with an electricity intensity based on Eq. (3.9) below 0.01. Data sources: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.5: ATEs of compensation on energy intensity (energy purchases/sales). 2010–2019.

	Energy intensity				
	2015	2016	2017	2018	2019
Compensation	-0.123 (0.102)	-0.0421 (0.112)	0.0519 (0.163)	0.0158 (0.111)	0.0177 (0.108)
Obs	688	989	1222	1445	1611
N compensated	27	42	45	45	45
N other	157	202	218	233	239
Energy intensity 05-11 (Treat)	0.031	0.030	0.037	0.031	0.031
Energy intensity 05-11 (Control)	0.036	0.034	0.037	0.036	0.035
p-value (mean-comparison test)	0.107	0.056	0.907	0.059	0.107
outcome pre-treatment (Treat)	-3.147	-3.087	-2.750	-3.089	-3.090
outcome pre-treatment (Control)	-2.980	-2.999	-2.990	-3.005	-3.013
p-value (mean-comparison test)	0.248	0.416	0.029	0.415	0.456

Notes: Table shows the coefficient β_1 estimated from the DiD equation. Dependent variables are in logs. Standard errors are clustered at the firm level. All regressions include year \times industry fixed effects at the 3-digit SIC code level and are weighted by the inverse propensity score. We drop plants with an electricity intensity based on Eq. (3.9) below 0.01. Data sources: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Robustness checks for DiD estimation

Our DiD results are robust to a number of tests. To mitigate concerns about how the global financial crisis might affect the computation of our p-scores, and our estimates accordingly, we show that our coefficients are robust to the use of an alternative time horizon to compute our p-scores ranging from 2010 to 2012 (see Appendix 3.E.6). We also show how our results change when we trim the sample by dropping plants with an electricity intensity based on Eq. (3.9) below different thresholds to ensure that our results are not driven by sample trimming decisions (see Appendix 3.E.4). Additionally, Appendix 3.E.3 shows how our results change when incorporating industry-specific effects at a broader sectoral level (2-digit level), thereby trading off some precision in the identification strategy to expand our estimation sample. Finally, Tables 3.A12 - 3.A14 in the Appendix provide a set of alternative estimations relying on different proxies for production and energy intensity from the ABS sample. These findings are summarized in Figures 3.A7 - 3.A8 in the Appendix which provides a graphical comparison of the estimated effects across the array of robustness tests across all outcome variables.

3.5.2 Fuzzy RD estimates of the local average treatment effects

Turning now to the RD estimation, we start by presenting the graphical evidence and estimated coefficients of the first stage and the reduced form, before turning to the instrumental variable estimates (second stage). Note that we present estimates from the first stage, reduced form and second stage using different functional forms as controls.


First stage and reduced form

Figure 3.3 illustrates the first stage, showing the share of compensated plants for different intervals of the electricity cost intensity. The sample is restricted to eligible industries, and averages within each bin are based on data from the period 2005–2011. Predicted electricity use is used to calculate the electricity cost intensity. As expected, we observe a sharp discontinuous jump in the share of compensated plants as we cross the eligibility cut-off; for plants with an electricity cost intensity between 5-6%, over half of the plants receive compensation payments. The exact height of the bars located to the left of the threshold is suppressed due to confidentiality reasons, but the share of compensated plants is below 10% for those bins.

Table 3.6, Panel A, reports the estimated first stage based on Equation 3.13, where we include both eligible and non-eligible industries as well as firm- and sector-year³⁹-specific fixed effects. The estimated coefficients reflect the probability of receiving compensation payments if the plant is above the 5% eligibility cut-off *and* operates in an eligible industry. The estimated probability of receiving compensation is 0.88 and the F-statistic of the excluded instrument is around 75. Thus, our first-stage results show that our instrument is a strong predictor of receiving compensation.

³⁹Sector-year fixed effects are included at the 2-digit level. Including this at the 3-digit level of disaggregation was not possible due to issues of sample size.

Figure 3.3: Share of compensated plants by electricity cost share. 2005-2011



Notes: Figure shows the share of compensated plants by the electricity cost share (electricity use*carbon price impact/GVA), using predicted electricity use. The height of the bars reflect mean values for plants located within the indicated electricity cost share bins. The precise height of the bars located to the left of the indicated threshold is censored due to disclosure concerns. The sample is restricted to eligible 4-digit industries. Data source: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI). See reference list for full citation.

Table 3.6: Local average treatment effects of compensation. 2010–2015. Fuzzy RDD.


Source: ABS			
	Sales of own goods (1)	Energy purchases (2)	Energy intensity (3)
Panel A: First stage	0.879*** (0.101)	0.879*** (0.101)	0.879*** (0.101)
Panel B: Reduced form	0.264** (0.125)	0.209 (0.187)	-0.0562 (0.134)
Panel C: Second stage	0.301** (0.131)	0.238 (0.199)	-0.0639 (0.156)
Panel D: OLS	0.164 (0.102)	0.263** (0.130)	0.103 (0.105)
Observations	253	249	335
N Compensated	20	20	20
N Other	49	48	27
F statistics	75.47	75.44	75.39
Functional form	Bins	Bins	Bins

Notes: Tables show the coefficients estimated from the first stage, reduced form, and second stage of the fuzzy regression discontinuity design. Dependent variables are given by the table headings and are measured in logs. Standard errors are clustered at the firm level. The sample estimation period is 2010-2015. Cutoff value: 0.05. Bandwidth: cutoff value +/-0.007. Bandwidth refers to the range of electricity intensity values (electricity use * carbon cost / GVA) used to restrict the sample. Each stage of the estimation includes firm-level and 2-digit sector-specific year fixed effects. Data source: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 3.4 shows graphical evidence of the discontinuous jump in our outcome variables at the threshold value, i.e., the reduced form effect. The RD plots are based on polynomial regressions over quantile-spaced bins, where we follow Calonico et al. (2015) to determine the optimal number of bins. Each dot represents a local mean for each bin. The figure shows a jump in sales of own goods (cf. Panel (a)) and electricity consumption (proxied by energy purchases, (cf. Panel (c))) at the 5% eligibility cut-off, indicating that the compensation had an effect on these outcomes.

Table 3.6, Panel B, reports the reduced form coefficients estimated based on Equation 3.14. The coefficients represent a lower bound of the effect of the compensation scheme (in the RDD sample) as not all plants that meet the eligibility criteria receive compensation. The reduced form estimates could be interpreted as “intention to treat”, which has the advantage that they do not rely on the exclusion restriction for unbiasedness. A statistically significant jump in outcome is observed for sales of own goods (0.26) but not for the other outcome variables. Additional results based on alternative specifications, different samples, and different outcome variables are presented in Appendix 3.H.

Figure 3.4: RD Plot based on quantile spaced number of bins. 2013-2015.



Notes: Figure shows data-driven regression discontinuity plots using polynomial regression based on quantile-spaced numbers of bins. Optimal number of bins has been selected following Calonico et al. (2015). Cutoff: 0.05. Data sources: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI). See reference list for full citation.

Main RD estimates

We now rescale the jump in (reduced form) outcomes by the jump in the (first stage) treatment probability to obtain the second stage estimates around the cut-off. A

causal interpretation of the findings relies on the assumption that crossing the 5% eligibility threshold only impacts plants via the probability of receiving compensation and reflects a LATE. Due to the smaller sample size around the threshold, our RD estimates are only based on the ABS sample.⁴⁰ For our main results, we report RD estimates with a ± 0.007 bandwidth (which restricts the sample to companies whose electricity cost share amounts to an interval between 4.3% and 5.7%) following the data-driven procedure to identify optimal estimation windows in RD settings by Calonico et al. (2020). More details on this procedure can be found in Section 3.H.4 in Appendix 3.H.

Table 3.6 Panel C reports the second stage RD estimates. We find evidence of a causal effect of compensation on production, proxied by sales, which increased by 30% for compensated plants relative to similar noncompensated plants. The effect on electricity consumption, proxied by energy purchases, is positive and large (24%) but not statistically significant, while we find a negative and non-significant effect on energy intensity. Overall, these findings are in line with our three predictions and DiD results and provide additional evidence that compensation for higher electricity prices particularly boosts production volumes for the compensated. Overall, while pointing towards the same general conclusions, compared to our ATEs, the RD estimates are larger in magnitude, suggesting that as expected, the effects of compensation tend to be larger for more electricity-intensive plants.

Robustness checks

We additionally perform a number of robustness tests to further investigate the validity of our baseline RD findings. Specifically, we produce RD estimates with different assumptions on the functional form where we amend Equations 3.13, 3.14, and 3.15 by additionally accounting for the linear (see Table 3.A17 in the Appendix)

⁴⁰The RD estimates for QFI variables yield statistically inconclusive results due to the very limited sample size within the bandwidth considered for the estimation and cannot be reported due to disclosure concerns.

and quadratic (see Table 3.A18 in the Appendix) distance of each observation from the threshold (cf., Section 3.4.2). We also examine the robustness of our main estimations with different bandwidth choices and generate a distribution of estimated effects across different estimation window sizes (see Section 3.H.5 in Appendix 3.H).

3.5.3 Comparing the DiD and RD estimates

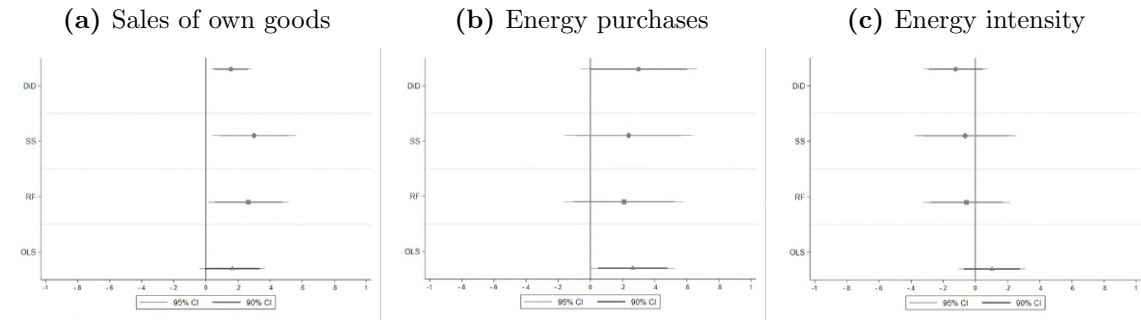
The balance of evidence from DiD and RD approaches is summarized in Figure 3.5. On the whole, results from both strategies indicate that the compensation scheme had a positive impact on sales and energy consumption with no detectable significant improvements in energy intensity. However, the magnitude of the treatment effect estimates differs between the two approaches, with the local average treatment effects (LATEs) estimated by the RD approach being larger than the average treatment effects (ATEs) estimated by the DiD approach.

One first reason for this difference in magnitude is that the two strategies focus on different populations with the RD approach focusing on a few plants around the discontinuity threshold in the electricity intensity distribution (see Panel (b) in Figure 3.2). This means that the RD approach may be interpreted as the treatment effect of the compensation scheme for plants that are most likely to be affected by the policy. In contrast, the DiD approach estimates an average effect of the compensation scheme that is representative for the broader population of manufacturing plants, regardless of their relative position in the electricity intensity distribution.

Another reason for the difference in magnitude may be linked to the identification strategy used in each approach. Our DiD approach combined with IPSW assumes that the weighted treatment and control groups are comparable in all other respects except for the treatment. However, this assumption may not hold if there are unobservable differences between the treatment and control groups that affect the outcomes of interest. The RD approach, on the other hand, relies on a discontinuity in the policy rule to identify the treatment effect. This means that the RD approach

is better able to control for unobservable factors that may affect the outcomes of interest.

Figure 3.5: Comparing ATEs and LATEs across ABS outcome variables.



Notes: Figure compares estimated coefficients across different empirical strategies and estimation samples. DD refers to the Difference-in-difference (DiD) estimates presented in Section 3.5.1. SS, RF, and OLS refer to the second stage, reduced form, and OLS estimates, respectively, presented in Section 3.5.2. All outcome variables are in log terms. Data sources: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI). See reference list for full citation.

3.6 Discussion on policy implications

Assessments of the effectiveness of anti-leakage policies typically focus on whether there is evidence of leakage occurring, without explicitly considering the costs of measures. This section aims to shed light on the trade-offs between preventing leakage and forgoing abatement.

Table 3.7 reports back-of-the-envelope calculations on the costs and benefits of the compensation scheme based on our DiD estimates. The estimated value associated with maintaining higher sales, calculated based on our DiD estimates of the average treatment effect (Table 3.2), is in the ballpark of £2 billion per year. The increase in production led to an increase in electricity use, cumulatively amounting to 2.35 TWh (or 2.5% of total annual industrial electricity consumption). The associated annual increase in indirect CO₂ emissions due to greater electricity use amounts to approximately 1.5 million tonnes CO₂.

The foregone reductions in indirect carbon emissions are valued at 36 to 377 million £ per year, depending on the CO₂ price assumption used. The upper bound estimate

Table 3.7: Costs and benefits

	Total
Number of compensated firms	59
Estimated value of increased production	£2,000 million / year
Estimated value of increased GVA	£232 million / year
Estimated forgone reduction in electricity use	2.35TWh / year
Increased indirect emissions	1.56 million tonnes / year
Increased indirect emissions - lower bound	£36 million / year
Increased indirect emissions - upper bound	£377 million / year
Compensation for CO ₂ costs	£72.4 million / year
Increase in production per £ of compensation	£27.6
Increase in GVA per £ of compensation	£3.2
Value of increased indirect emissions per £ compensation - lower bound	£0.5
Value of increased indirect emissions per £ compensation - upper bound	£5.2

Note: £ are reported in 2020-values. Compensation payments are computed by averaging the values reported between 2013 and 2019 (cf. Section 3.3.1). We calculate increases in production and indirect emission for the average compensated firm in our sample by leveraging our DiD estimates of the average treatment effect presented in Table 3.2 and 3.3. Specifically, we calculate firm-specific mean increases in sales (as a proxy for production) and indirect emissions by multiplying the corresponding estimated ATE from Eq. 3.11 with mean pre-treatment outcome levels of sales (with a mean value of 173,749 thousand £) and indirect emissions (with a mean value of 117,904 tonnes) in each compensated firm. We additionally compute the implied increase in GVA leveraging our additional estimates summarized in Figure 3.A7. We obtain cumulative values by multiplying the estimated mean firm-level increases by the total number of compensated firms. *Lower bound* increased indirect emissions (£) are calculated based on the average EUA price in 2020 (which amounted to 22.83 £). Upper bound increased indirect emissions (£) are estimated using UK official guidelines on the social costs of carbon (SCC) of £241 £ / tonne of carbon dioxide emitted.

uses current official recommendations on the social cost of carbon (SCC) from the UK government⁴¹ while the lower bound estimate uses the average EUA clearing prices as an alternative market-based proxy for the cost of a tonne of CO₂. The upper bound estimate is less informative here because given the ETS cap, the SCC reflects abatement costs elsewhere in the economy.

The substantial increase in production indicates that the compensation scheme has contributed to shielding energy-intensive firms from higher electricity costs by acting as an implicit production subsidy. When comparing the magnitudes to the direct annual cost of the scheme of around 72 million £ (cf. Section 3.3.1), each pound of compensation on average has yielded more than one pound in production value (proxied by sales) and GVA. Yet the collateral increase in indirect emissions among compensated energy-intensive firms is sizable, corresponding to around 4.3% (1.3%) of annual industrial (nationwide) emissions from electricity use.

Therefore, in line with empirical studies that find limited evidence of carbon leakage

⁴¹Under current guidelines, the UK government recommends using a social costs of carbon (SCC) per tonne of carbon dioxide emitted of £241 (in 2020 £) for policy appraisal and evaluations. See here for further details.

from the EU ETS due to generous free allocation (e.g. Naegele and Zaklan, 2019), our results indicate that the indirect carbon compensation scheme is working, insofar as production displacement and carbon leakage is being discouraged. However, the known downsides of preventing leakage through an output-based compensation have also materialized. Compensation dampens the carbon price signal which is intended to reduce emissions by discouraging the production of CO₂ intensive goods. It creates perverse incentives on the supply side to artificially inflate output, resulting in higher emissions compared to a scenario without compensation.

Interestingly, indirect cost compensation had no statistically significant effect on employment (cf. Table 3.A6), suggesting that increased electricity prices due to carbon pricing have not led to the displacement of workers in electro-intensive sectors. We also do not find any significant effect on GVA, which is a proxy for value added, or on productivity (cf. Figure 3.A7). The scheme also did not hamper technological improvements in terms of increased energy efficiency in compensated firms vis-à-vis non-compensated firms.

These results have implications for both the economic efficiency and distributional outcomes from carbon pricing. It is likely that compensation to electro-intensive sectors increases the overall compliance cost for meeting mitigation goals. As compensation targets energy-intensive sectors, abatement responsibilities in the EU ETS would shift toward sectors with relatively lower energy intensity – due to the so-called *waterbed effect* (cf. Perino, 2018). This shift represents an adjustment in the distribution of the compliance costs associated with reducing CO₂ emissions. Sectors with lower energy intensity may find it costlier to implement emissions reduction measures compared to energy-intensive sectors. As the abatement burden shifts to these sectors, the cost-effectiveness of the emissions reduction program may diminish and increase overall compliance costs for the cap-and-trade system (Martin et al., 2014b). In the context of an inter-jurisdictional cap-and-trade system, this additionally implies that unilateral compensation schemes have the potential to shift the distribution of abatement responsibilities across countries, effectively

redistributing not only carbon abatement costs but also the local health co-benefits associated with reduced emissions of air pollutants from CO₂ combustion (e.g., Cushing et al., 2018; Banzhaf et al., 2019; Hernandez-Cortes and Meng, 2023).

Additionally, under output-based compensation, theory predicts that producers will not pass on the full CO₂ cost to product prices (Quirion, 2009). Without the full CO₂ cost pass-through, incentives along the production and consumption chain to substitution away from energy-intensive goods are dampened. This suggests the need for supplementary consumption-based measures to encourage mitigation through demand-side substitution. For example, embodied carbon standards, green procurement, and climate excise contribution are discussed in the literature (Grubb et al., 2022).

Finally, using ETS auction revenue to compensate energy-intensive companies for higher carbon costs comes at the trade-off of other climate-related investments or redistributing climate policy costs to the public through alternative revenue recycling schemes (such as lump sum transfers), which could contribute to enhancing the public acceptability of carbon pricing schemes (Baranzini et al., 2017; Douenne and Fabre, 2022). The need to consider the opportunity costs of public funds devoted to compensation schemes becomes more salient with the anticipated substantial future payments driven by the recent surge in carbon prices within the EU ETS. These forthcoming payments are expected to lead to substantial transfers that disproportionately benefit a select few energy-intensive firms and their capital owners, underscoring pivotal equity implications in the distribution of climate policy costs.

3.7 Conclusion

Governments pursuing ambitious climate policies encounter a complex challenge characterized by a delicate balancing act. On one hand, they must incentivize emission reduction efforts, and on the other, they must mitigate the risk of carbon

leakage and competitive disadvantage for domestic industries. This conundrum necessitates the deployment of comprehensive strategies. One approach is to pair carbon pricing with schemes that compensate energy-intensive firms for higher carbon costs or electricity prices. Such policies may help obtain political buy-in from industry and alleviate adverse economic effects. At the same time, a carbon cost containment measure by its nature is likely to delay industrial decarbonization.

While the downsides of output-based free allocation or compensation have been known, perhaps they have been downplayed due to the lack of empirical evidence. We use UK microdata and idiosyncrasies in the eligibility criteria to examine the impact of indirect carbon cost compensation on firms output, electricity use, electricity intensity, and emissions.

We find robust evidence that as intended, compensation limits carbon leakage. It does so by attenuating the carbon price signal and discouraging energy-intensive firms from reducing production, electricity use and emissions. Our back-of-the-envelope calculations suggest that each pound of compensation yields more than a pound in production value and GVA, but the increase in indirect emissions among compensated energy-intensive firms is also sizable.

In the context of a cap-and-trade scheme, these findings carry important implications for the distribution of mitigation burdens across sectors. Dampening incentives to limit supply from energy-intensive sectors means that to achieve the overall ETS cap, mitigation shifts elsewhere (to other sectors or towards greater emissions intensity improvements) which implies allowance prices and overall costs would rise (Martin et al., 2014b).

Compensation for indirect carbon costs as well as free allocation is, however, likely to prevail for some time.⁴² Free allocation in the EU ETS is also set to continue until

⁴²The UK has committed to continued compensation to 2025 (Department for Business, Energy & Industrial Strategy, UK, 2022) while international CO₂ price differences prevail, and industrial carbon neutral technologies are not yet widely available. In Europe, several governments have already committed compensation payments until 2030.


2028 (European Parliament , 2021) even after the introduction of the Carbon Border Adjustment Mechanism (CBAM) to reduce leakage risk for EU exporters because the proposed CBAM targets imports only. Indeed, free allocation continues to be the default anti-leakage policy across emission trading schemes worldwide, not least because it is hugely advantageous for obtaining political buy-in for carbon pricing from industry (Sato et al., 2022). Our results help make these difficult trade-offs faced by policy makers more explicit, by quantifying the increased production by energy and emission intensive firms due to compensation payments.

3.A Research context

3.A.1 Electricity prices in the UK and continental Europe

Electricity prices are kept low in continental Europe, often through discounts or exemptions for industrial users. For example, in Germany, the regulatory approach taken to recover network and policy costs protects electro-intensive industries by recovering costs primarily from domestic and commercial users. By contrast, in the UK these costs are spread relatively evenly across all electricity consumers. In France, the industry has been able to collectively negotiate long-term contracts for lower electricity prices, whereas the UK market has no collectively negotiated contracts and few contracts with a duration beyond a couple of years ahead. Higher levels of interconnection on the continent also allow policy choices to lower industrial electricity prices. For example in Italy, the government facilitated large energy-intensive companies to purchase cheap electricity from neighboring countries in exchange for investments in expanding interconnection capacity. Furthermore, wholesale price differences between the UK and continental Europe are driven by differences in fossil fuel prices, renewable penetration, and the merit order effect (Grubb and Drummond, 2018).

Figure 3.A1: Industrial electricity prices, pre-and post-compensation for selected EU countries in 2016. EUR/MWh.



Source: Adapted from *Scenario S2* in Grubb and Drummond (2018). Carbon price compensation covers EU ETS costs and costs induced by the UK carbon price floor/tax.

3.A.2 Eligible 4-digit industries

Table 3.A1: Eligible industries


Industry	NACE Rev. 1
Mining of Iron Ore	1310
Mining of chemical and fertiliser minerals	1430
Preparation and spinning of cotton-type fibres	1711
Manufacture of leather clothes	1810
Manufacture of pulp*	2111
Manufacture of paper and paperboard	2112
Manufacture of other inorganic basic chemicals	2413
Manufacture of other organic basic chemicals	2414
Manufacture of fertilisers and nitrogen compounds	2415
Manufacture of plastics in primary forms*	2416
Manufacture of man-made fibres	2470
Manufacture of basic iron and steel and of ferro-alloys	2710
Aluminium production	2742
Lead, zinc and tin production	2743
Copper production	2744

Note: For industries noted by *, only a subset of products are eligible for compensation. *Source:* European Commission (2012)

3.B Additional descriptive material

3.B.1 Magnitude of compensation payments

Figure 3.A2: Compensation payments as a share of electricity prices. 2013-2017.

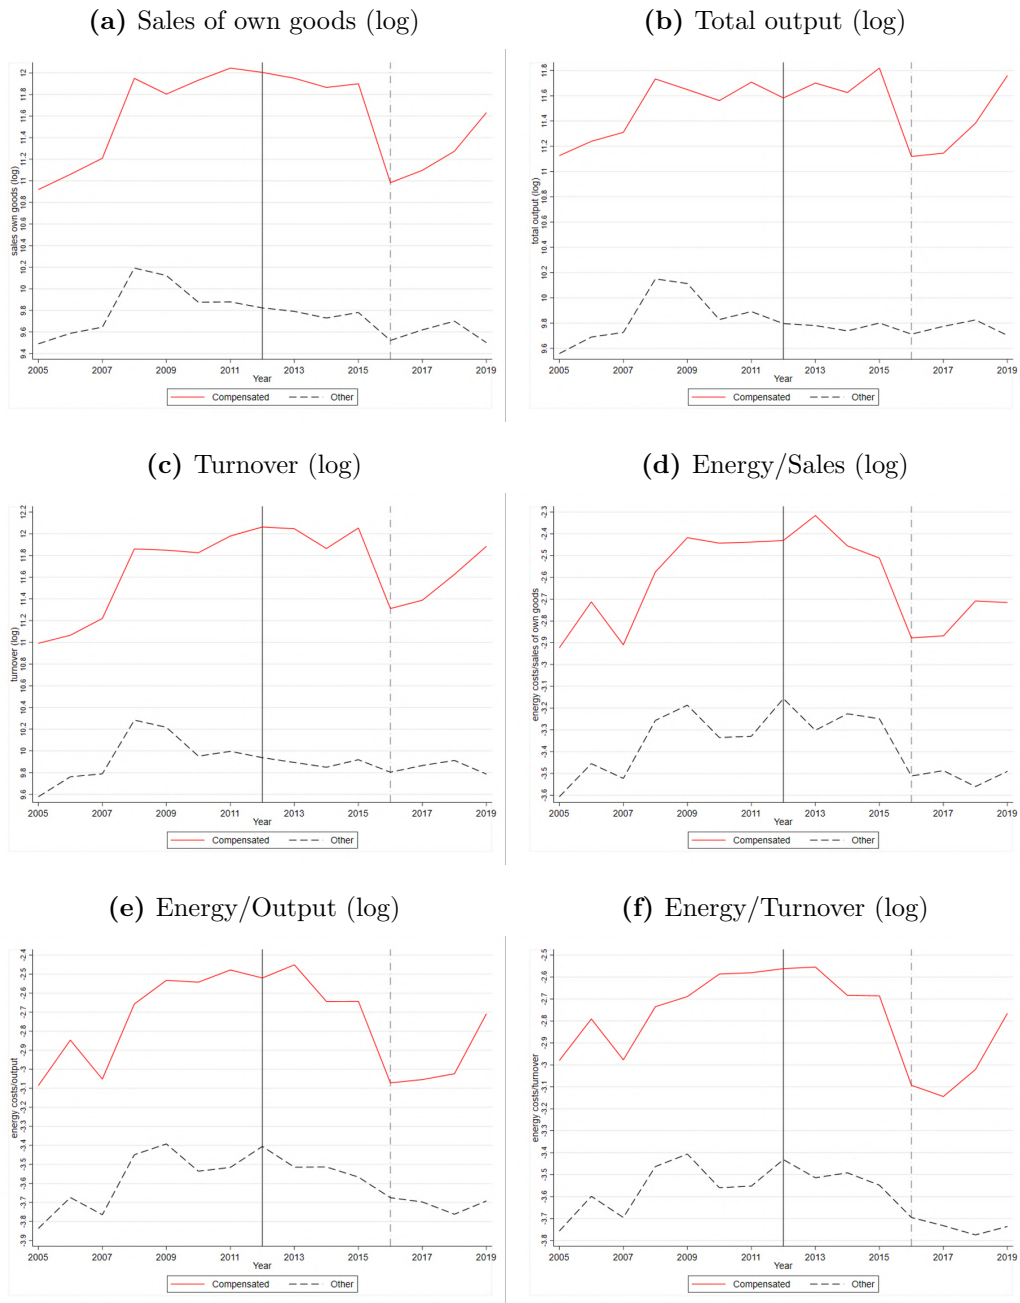


Notes: Own calculations based on compensation formula and average electricity prices from the UK Department for Business, Energy & Industrial Strategy (2018), Table 3.1.4: Prices of fuels purchased by manufacturing industry in Great Britain.

3.B.2 Descriptive evidence for outcome variables

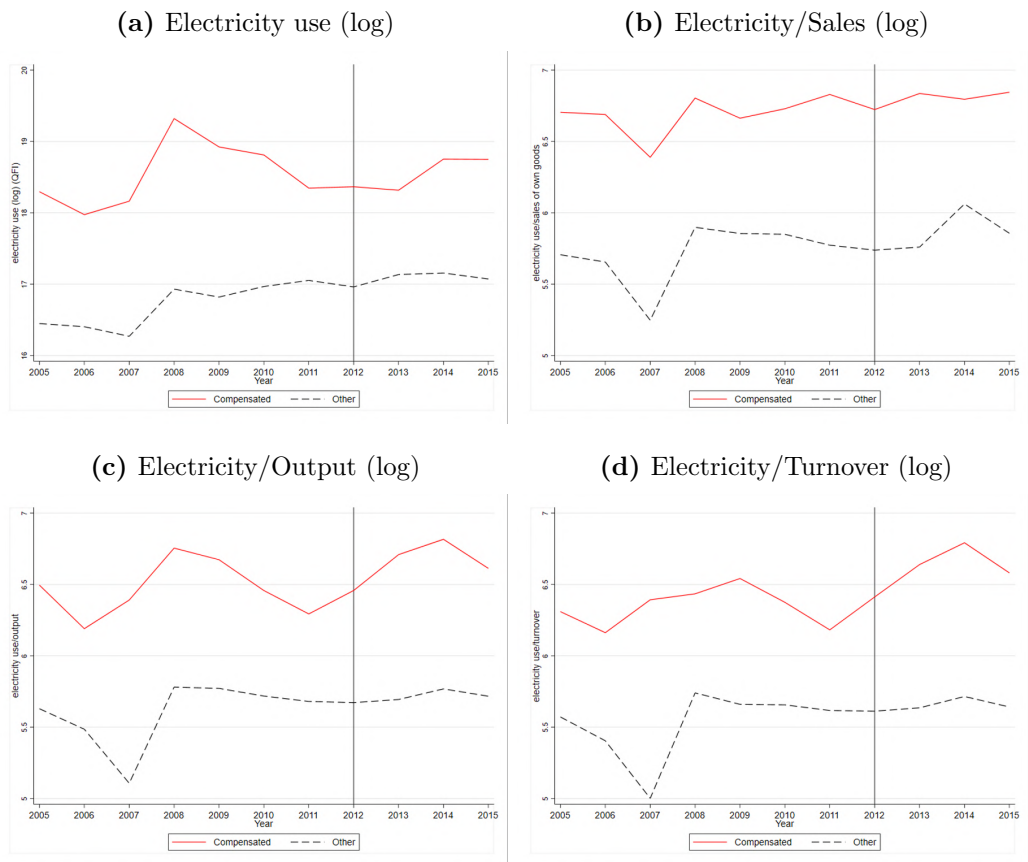
Figures 3.A3 and 3.A4 provide additional descriptive evidence for our key outcome variables. The raw mean trends exhibit a steady significant decrease in energy intensity for the average plant in the sample, following the introduction of the UK Carbon Price Floor in 2013 as can be seen in Panel (d), Panel (e), and Panel (f) in Figure 3.A3. Although the raw pre-post mean comparison already provides some exploratory evidence, it does not necessarily capture the causal effect of the regulation, as there are many possible channels that could plausibly explain the observed drop in energy intensity. Additionally, there has been a remarkable decrease in production levels as shown by Panel (a), Panel (b), and Panel (c) in Figure 3.A3 both in the late 2000s and in 2016. These drops coincide respectively with the global financial crisis in 2008-09 and the EU membership Referendum, that took place in the UK in 2016 (Brexit). Descriptive evidence from Figure 3.A4 indicates that there has been a tendency to increase electricity consumption and electricity intensity among compensated plants vis-a-vis uncompensated plants (that do not exhibit any trend deviation) following 2013.

Figure 3.A3: Raw average trends in key outcome variables from ABS over time, by year. 2005-2019.



Notes: Figures plot the average values of key outcomes variable over time by treatment status. Data sources: Annual Business Survey (ABS). The vertical line indicates the year before the carbon price floor was introduced.

Figure 3.A4: Raw average trends in key outcome variables from QFI over time, by year. 2005-2015.



Notes: Figures plot the average values of key outcomes variable over time over time by treatment status. Data sources: Quarterly Fuels Inquiry (QFI). The vertical line indicates the year before the carbon price floor was introduced.

3.C Predicting electricity consumption

Table 3.A2: Simple correlation between electricity use and energy purchases. 2005–2011

	Electricity (log)	Electricity (log)	Electricity (log)
Energy purch. (log)	0.905*** (0.0112)	0.926*** (0.0135)	0.920*** (0.0228)
R ²	0.645	0.670	0.669
Obs	3583	2324	805
Sample	QFI, all SIC	eligible SIC, 2 digit	eligible SIC, 4 digit

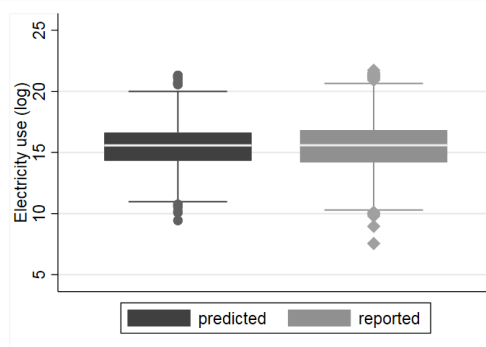
Notes: Table shows correlations between electricity consumption and energy purchases. Both variables are in logs. Electricity use is only available for plants part of the QFI. Data source: ARDx and QFI. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A3: Correlation between electricity use and energy purchases, employment, and turnover. 2005–2011

	Electricity (log)
Energy purch. (log)	0.567*** (0.0237)
Employment (log)	0.316*** (0.0400)
Turnover (log)	0.148*** (0.0323)
Constant	8.030*** (0.0740)
R ²	0.740
Obs	3582
Sample	QFI, all SIC

Notes: Table shows the correlation between predicted and electricity use and observed energy purchases. Regression include industry dummies at the 4-digit level. Data source: ARDx and QFI. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 3.A5: Comparing predicted and reported electricity use.




Notes: Box plot shows the distribution of predicted electricity use and reported electricity use. Sample is restricted to units with reported electricity use. Data source: ABS and QFI.

3.D Using energy purchases to calculate the running variable

As an alternative to using predicted electricity consumption to calculate the second eligibility criteria, we have also tried to use infer the electricity intensity cut-off value, X_0 , using total energy purchase in GBP. Total energy purchases and GVA are available from the Annual Business Survey (ABS), hence avoiding the problem om limited coverage of the Quarterly Fuels Inquiry (QFI). Figure 3.A6 plots the share of plants that receive compensation, by pre-treatment energy intensity. Energy intensity is defined as energy purchases divided by GVA, and the sample is restricted to the eligible industries listed in Table 3.A1. From the figure, there is a clear jump in the probability of a plant receiving compensation when the energy intensity is above 12 %.

Figure 3.A6: Share of compensated plants, by average energy intensity.



Notes: Energy intensity is measured as plant-level energy costs divided by gross value added (GVA). Values are averaged over the years 2009 to 2012. Sample is restricted to the eligible industries listed in Table 3.A1, and to plants with an energy intensity less than 0.30 to account for outliers. The sample is split into 20 bins containing an equal number of plants. Each data point in the graph reflects the average share of compensated plants within each of these 20 bins. Data source: Annual Business Survey (ABS).

3.E Additional results: DiD

3.E.1 Complementary results based on the period 2010–2019

Table 3.A4: ATEs of compensation on sales. 2010–2019.

Sales of own goods					
	2015	2016	2017	2018	2019
Compensation	0.156** (0.0638)	0.164** (0.0763)	0.126* (0.0705)	0.147** (0.0701)	0.144** (0.0693)
Obs	532	717	851	1069	1186
N compensated	27	36	39	40	40
N other	97	127	132	156	158
Energy intensity 05-11 (Treat)	0.035	0.035	0.035	0.034	0.034
Energy intensity 05-11 (Control)	0.036	0.037	0.036	0.036	0.036
p-value (mean-comparison test)	0.725	0.424	0.628	0.596	0.597
outcome pre-treatment (Treat)	11.135	10.867	10.831	10.744	10.743
outcome pre-treatment (Control)	11.147	10.905	10.906	10.844	10.844
p-value (mean-comparison test)	0.944	0.822	0.664	0.558	0.553

Notes: Table shows the coefficient β_1 estimated from the DiD equation. Dependent variables are in logs. Standard errors are clustered at the firm level. All regressions include year x industry fixed effects at the 3-digit SIC code level and are weighted by the inverse propensity score. We drop plants with an electricity intensity (i.e., predicted electricity consumption/GVA) below 0.1. Data sources: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A5: ATEs of compensation on energy intensity (energy purchases/sales). 2010–2019.

Energy intensity					
	2015	2016	2017	2018	2019
Compensation	-0.123 (0.102)	-0.0421 (0.112)	0.0519 (0.163)	0.0158 (0.111)	0.0177 (0.108)
Obs	688	989	1222	1445	1611
N compensated	27	42	45	45	45
N other	157	202	218	233	239
Energy intensity 05-11 (Treat)	0.031	0.030	0.037	0.031	0.031
Energy intensity 05-11 (Control)	0.036	0.034	0.037	0.036	0.035
p-value (mean-comparison test)	0.107	0.056	0.907	0.059	0.107
outcome pre-treatment (Treat)	-3.147	-3.087	-2.750	-3.089	-3.090
outcome pre-treatment (Control)	-2.980	-2.999	-2.990	-3.005	-3.013
p-value (mean-comparison test)	0.248	0.416	0.029	0.415	0.456

Notes: Table shows the coefficient β_1 estimated from the DiD equation. Dependent variables are in logs. Standard errors are clustered at the firm level. All regressions include year x industry fixed effects at the 3-digit SIC code level and are weighted by the inverse propensity score. We drop plants with an electricity intensity (i.e., predicted electricity consumption/GVA) below 0.1. Data sources: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.E.2 DiD Results for employment

Table 3.A6: ATEs of compensation on employment. 2010–2019.

	Employment				
	2015	2016	2017	2018	2019
Compensation	0.0355 (0.0437)	0.0292 (0.0451)	0.0122 (0.0482)	0.0230 (0.0492)	0.0239 (0.0503)
Obs	669	923	1106	1337	1492
N compensated	28	40	42	43	43
N other	139	176	184	205	208
Energy intensity 05-11 (Treat)	0.035	0.035	0.035	0.035	0.035
Energy intensity 05-11 (Control)	0.035	0.037	0.037	0.036	0.036
p-value (mean-comparison test)	0.920	0.637	0.660	0.699	0.698
outcome pre-treatment (Treat)	5.401	5.306	5.295	5.246	5.245
outcome pre-treatment (Control)	5.363	5.108	5.130	5.116	5.117
p-value (mean-comparison test)	0.744	0.086	0.151	0.255	0.260
p-value (mean-comparison test)	0.907	0.555	0.478	0.312	0.220

Notes: Table shows the coefficient β_1 estimated from the DiD equation. Dependent variables are in logs. Standard errors are clustered at the firm level. All regressions include year x industry fixed effects at the 3-digit SIC code level and are weighted by the inverse propensity score. We drop plants with an electricity intensity (i.e., predicted electricity consumption/GVA) below 0.1. Data sources: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.E.3 Alternative definition of industry-year fixed effects

Table 3.A7: ATEs of compensation on sales *within 2-digit industries*. 2010–2019.

Sales of own goods (effects at 2-digit level)					
	2015	2016	2017	2018	2019
Compensation	0.146*** (0.0536)	0.185** (0.0720)	0.193*** (0.0741)	0.195*** (0.0707)	0.190*** (0.0724)
Obs	925	1189	1409	1709	1881
N compensated	31	41	44	45	45
N other	185	216	225	250	252
Energy intensity 05-11 (Treat)	0.045	0.039	0.040	0.038	0.038
Energy intensity 05-11 (Control)	0.035	0.035	0.035	0.034	0.034
p-value (mean-comparison test)	0.001	0.170	0.155	0.233	0.239
outcome pre-treatment (Treat)	9.538	10.729	10.718	10.669	10.667
outcome pre-treatment (Control)	11.003	10.876	10.862	10.810	10.813
p-value (mean-comparison test)	0.000	0.266	0.273	0.269	0.254

Notes: Table shows the coefficient β_1 estimated from the DiD equation. Dependent variables are in logs. Standard errors are clustered at the firm level. All regressions include year x industry fixed effects at the 2-digit SIC code level and are weighted by the inverse propensity score. We drop plants with an electricity intensity (i.e., predicted electricity consumption/GVA) below 0.1. Data sources: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A8: ATEs of compensation on energy intensity (energy purchases/sales) *within 2-digit industries*. 2010–2019.

Energy intensity (effects at 2-digit level)					
	2015	2016	2017	2018	2019
Compensation	-0.0906 (0.0622)	0.0808 (0.0961)	0.130 (0.104)	0.114 (0.0806)	0.108 (0.0814)
Obs	1097	1447	1778	2080	2304
N compensated	39	49	54	56	56
N other	261	319	342	356	365
Energy intensity 05-11 (Treat)	0.035	0.039	0.039	0.034	0.034
Energy intensity 05-11 (Control)	0.037	0.037	0.037	0.037	0.037
p-value (mean-comparison test)	0.576	0.596	0.528	0.260	0.278
outcome pre-treatment (Treat)	-2.999	-2.953	-2.937	-3.133	-3.133
outcome pre-treatment (Control)	-2.973	-3.026	-3.013	-3.029	-3.038
p-value (mean-comparison test)	0.814	0.437	0.410	0.236	0.274

Notes: Table shows the coefficient β_1 estimated from the DiD equation. Dependent variables are in logs. Standard errors are clustered at the firm level. All regressions include year x industry fixed effects at the 2-digit SIC code level and are weighted by the inverse propensity score. We drop plants with an electricity intensity (i.e., predicted electricity consumption/GVA) below 0.1. Data sources: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.E.4 Alternative sample trimming

Table 3.A9: ATEs of compensation on sales *with different trimming*. 2010–2019.

Sales of own goods (Intensity >0.05)					
	2015	2016	2017	2018	2019
Compensation	0.136** (0.0653)	0.158** (0.0665)	0.130** (0.0651)	0.114* (0.0643)	0.115* (0.0644)
Obs	914	1252	1552	1875	2081
N compensated	33	50	62	63	63
N other	199	249	277	309	315
Energy intensity 05-11 (Treat)	0.031	0.030	0.030	0.031	0.031
Energy intensity 05-11 (Control)	0.026	0.027	0.027	0.026	0.026
p-value (mean-comparison test)	0.123	0.255	0.189	0.070	0.062
outcome pre-treatment (Treat)	10.441	10.276	10.227	10.208	10.194
outcome pre-treatment (Control)	10.673	10.474	10.449	10.396	10.386
p-value (mean-comparison test)	0.233	0.254	0.204	0.277	0.270

Notes: Table shows the coefficient β_1 estimated from the DiD equation. Dependent variables are in logs. Standard errors are clustered at the firm level. All regressions include year x industry fixed effects at the 3-digit SIC code level and are weighted by the inverse propensity score. We drop plants with an electricity intensity (i.e., predicted electricity consumption/GVA) below 0.05. Data sources: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A10: ATEs of compensation on energy intensity (energy purchases/sales) *with different trimming*. 2010–2019.

Energy intensity (Intensity >0.05)					
	2015	2016	2017	2018	2019
Compensation	-0.144 (0.0908)	-0.0380 (0.0974)	0.0105 (0.102)	0.0401 (0.0962)	0.0322 (0.0933)
Obs	997	1427	1789	2140	2371
N compensated	34	56	71	72	72
N other	234	299	331	363	371
Energy intensity 05-11 (Treat)	0.028	0.028	0.028	0.028	0.028
Energy intensity 05-11 (Control)	0.029	0.028	0.030	0.029	0.028
p-value (mean-comparison test)	0.785	0.920	0.432	0.665	0.757
outcome pre-treatment (Treat)	-3.249	-3.094	-3.108	-3.108	-3.118
outcome pre-treatment (Control)	-3.242	-3.254	-3.238	-3.252	-3.275
p-value (mean-comparison test)	0.969	0.124	0.198	0.152	0.117

Notes: Table shows the coefficient β_1 estimated from the DiD equation. Dependent variables are in logs. Standard errors are clustered at the firm level. All regressions include year x industry fixed effects at the 3-digit SIC code level and are weighted by the inverse propensity score. We drop plants with an electricity intensity (i.e., predicted electricity consumption/GVA) below 0.05. Data sources: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A11: ATEs of compensation on QFi variables *with different trimming*. 2010–2015.

	Electricity use	Electricity intensity	Carbon Emissions
Compensation	0.240*** (0.0831)	0.185 (0.186)	0.243*** (0.0818)
Obs	472	712	490
N compensated	15	16	15
N other	76	128	79
Industry effects digit	1	1	1
Industry effects-year FE	Yes	Yes	Yes
EUTL-year-Industry effects	No	No	No
Energy intensity 05-11 (Treat)	0.036	0.033	0.036
Energy intensity 05-11 (Control)	0.034	0.031	0.032
p-value (mean-comparison test)	0.700	0.546	0.329
outcome pre-treatment (Treat)	17.299	5.901	23.512
outcome pre-treatment (Control)	17.360	5.960	23.562
p-value (mean-comparison test)	0.715	0.692	0.796

Notes: Table shows the coefficient β_1 estimated from the DiD equation. Dependent variables are in logs. Standard errors are clustered at the firm level. All regressions include year x industry fixed effects at the 1-digit level and are weighted by the inverse propensity score. We drop plants with an electricity intensity (i.e., predicted electricity consumption/GVA) below 0.05. Data sources: Quarterly Fuels Inquiry (QFI).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.E.5 Alternative proxies for production and electricity intensity

Table 3.A12: ATEs of compensation on total output. 2010–2019.

	Total output				
	2015	2016	2017	2018	2019
Compensation	0.147* (0.0798)	0.185** (0.0868)	0.156* (0.0895)	0.150* (0.0884)	0.157* (0.0902)
Obs	528	710	862	1105	1230
N compensated	26	37	40	41	41
N other	99	125	134	162	165
Energy intensity 05-11 (Treat)	0.035	0.034	0.035	0.034	0.034
Energy intensity 05-11 (Control)	0.036	0.037	0.037	0.037	0.037
p-value (mean-comparison test)	0.586	0.338	0.403	0.350	0.351
outcome pre-treatment (Treat)	10.960	10.752	10.721	10.651	10.650
outcome pre-treatment (Control)	11.116	10.873	10.839	10.778	10.778
p-value (mean-comparison test)	0.307	0.454	0.472	0.422	0.419

Notes: Table shows the coefficient β_1 estimated from the DiD equation. Dependent variables are in logs. Standard errors are clustered at the firm level. All regressions include year x industry fixed effects at the 3-digit SIC code level and are weighted by the inverse propensity score. We drop plants with an electricity intensity (i.e., predicted electricity consumption/GVA) below 0.1. Data sources: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A13: ATEs of compensation on turnover. 2010–2019.

	Turnover				
	2015	2016	2017	2018	2019
Compensation	0.141*	0.158**	0.159**	0.192**	0.195**
	(0.0721)	(0.0761)	(0.0778)	(0.0796)	(0.0803)
Obs	550	755	907	1151	1270
N compensated	25	37	39	41	41
N other	105	133	141	168	170
Energy intensity 05-11 (Treat)	0.035	0.035	0.035	0.034	0.034
Energy intensity 05-11 (Control)	0.036	0.036	0.036	0.035	0.035
p-value (mean-comparison test)	0.725	0.603	0.681	0.722	0.722
outcome pre-treatment (Treat)	11.120	10.936	10.906	10.831	10.830
outcome pre-treatment (Control)	11.130	10.946	10.915	10.850	10.859
p-value (mean-comparison test)	0.954	0.954	0.957	0.906	0.858

Notes: Table shows the coefficient β_1 estimated from the DiD equation. Dependent variables are in logs. Standard errors are clustered at the firm level. All regressions include year x industry fixed effects at the 3-digit SIC code level and are weighted by the inverse propensity score. We drop plants with an electricity intensity (i.e., predicted electricity consumption/GVA) below 0.1. Data sources: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A14: ATEs of compensation on energy/output. 2010–2019.

	Energy/Output				
	2015	2016	2017	2018	2019
Compensation	-0.0615	0.0728	0.0601	0.00245	0.00502
	(0.128)	(0.141)	(0.141)	(0.132)	(0.125)
Obs	648	896	1145	1371	1524
N compensated	24	39	43	45	45
N other	150	179	201	217	223
Energy intensity 05-11 (Treat)	0.036	0.035	0.035	0.035	0.035
Energy intensity 05-11 (Control)	0.035	0.036	0.038	0.037	0.036
p-value (mean-comparison test)	0.982	0.858	0.341	0.463	0.754
outcome pre-treatment (Treat)	-3.062	-2.966	-2.976	-2.952	-2.949
outcome pre-treatment (Control)	-3.042	-3.033	-3.051	-3.058	-3.078
p-value (mean-comparison test)	0.907	0.555	0.478	0.312	0.220

Notes: Table shows the coefficient β_1 estimated from the DiD equation. Dependent variables are in logs. Standard errors are clustered at the firm level. All regressions include year x industry fixed effects at the 3-digit SIC code level and are weighted by the inverse propensity score. We drop plants with an electricity intensity (i.e., predicted electricity consumption/GVA) below 0.1. Data sources: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A15: ATEs of compensation on energy/turnover. 2010–2019.

	Energy/Turnover				
	2015	2016	2017	2018	2019
Compensation	-0.0445 (0.101)	0.0134 (0.116)	0.0476 (0.146)	0.0319 (0.138)	0.00431 (0.135)
Obs	705	983	1232	1455	1609
N compensated	27	39	44	45	46
N other	162	206	224	238	241
Industry effects digit	3	3	3	3	3
Energy intensity 05-11 (Treat)	0.030	0.029	0.036	0.036	0.036
Energy intensity 05-11 (Control)	0.035	0.036	0.033	0.032	0.032
p-value (mean-comparison test)	0.075	0.009	0.145	0.113	0.042
outcome pre-treatment (Treat)	-3.300	-3.254	-3.019	-3.029	-2.998
outcome pre-treatment (Control)	-3.123	-3.128	-3.135	-3.170	-3.177
p-value (mean-comparison test)	0.181	0.221	0.268	0.180	0.090
outcome pre-treatment (Control)	-3.042	-3.033	-3.051	-3.058	-3.078
p-value (mean-comparison test)	0.907	0.555	0.478	0.312	0.220

Notes: Table shows the coefficient β_1 estimated from the DiD equation. Dependent variables are in logs. Standard errors are clustered at the firm level. All regressions include year x industry fixed effects at the 3-digit SIC code level and are weighted by the inverse propensity score. We drop plants with an electricity intensity (i.e., predicted electricity consumption/GVA) below 0.1. Data sources: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.E.6 Alternative calculations of p-scores

In our main results, we estimate propensity score weights based on the period 2005–2011. These years correspond to the period used by the Government to calculate the electricity cost share, which again determines whether a plant passes the 5% filter test.

In Table 3.A16, we show that our main results are robust to the use of an alternative time horizon to compute our p-scores ranging from 2010 to 2012, which mitigate potential concerns about how the global financial crisis might affect the p-score estimation.

Table 3.A16: ATEs of compensation on sales *using different p-scores*. 2010–2015.

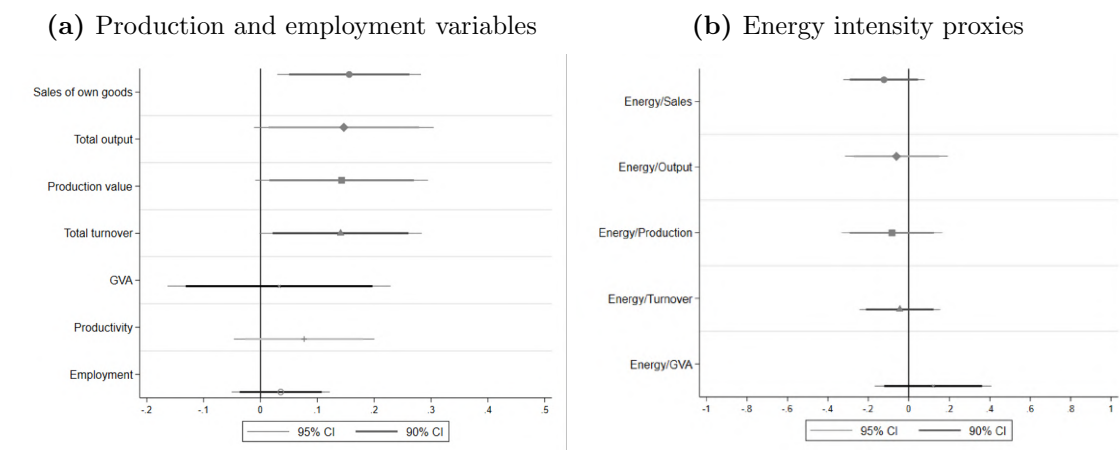
	ABS			QFI	
	Sales	Energy intensity	Electricity use	Indirect Emissions	Electricity intensity
Compensation	0.121* (0.0646)	0.0762 (0.0839)	0.166* (0.1000)	0.177* (0.0981)	0.118 (0.173)
Obs	474	629	489	464	630
N compensated	20	24	13	13	16
N other	84	137	84	78	113
Industry effects digit	3	3	1	1	1
Industry effects-year FE	Yes	Yes	Yes	Yes	Yes
Energy intensity 05-11 (Treat)	0.034	0.031	0.043	0.042	0.035
Energy intensity 05-11 (Control)	0.038	0.036	0.034	0.034	0.033
p-value (mean-comparison test)	0.311	0.100	0.091	0.092	0.535
outcome pre-treatment (Treat)	11.213	-3.099	17.200	23.389	5.999
outcome pre-treatment (Control)	11.218	-3.076	17.329	23.598	6.032
p-value (mean-comparison test)	0.977	0.844	0.468	0.240	0.834

Notes: Table shows the coefficient β_1 estimated from the DiD equation. Dependent variables are in logs. Standard errors are clustered at the firm level. All regressions include year x industry fixed effects at the 3-digit SIC code level (ABS variables) or 1-digit level (QFI variables) and are weighted by the inverse propensity score. We drop plants with an electricity intensity (i.e., predicted electricity consumption/GVA) below 0.1. Data sources: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

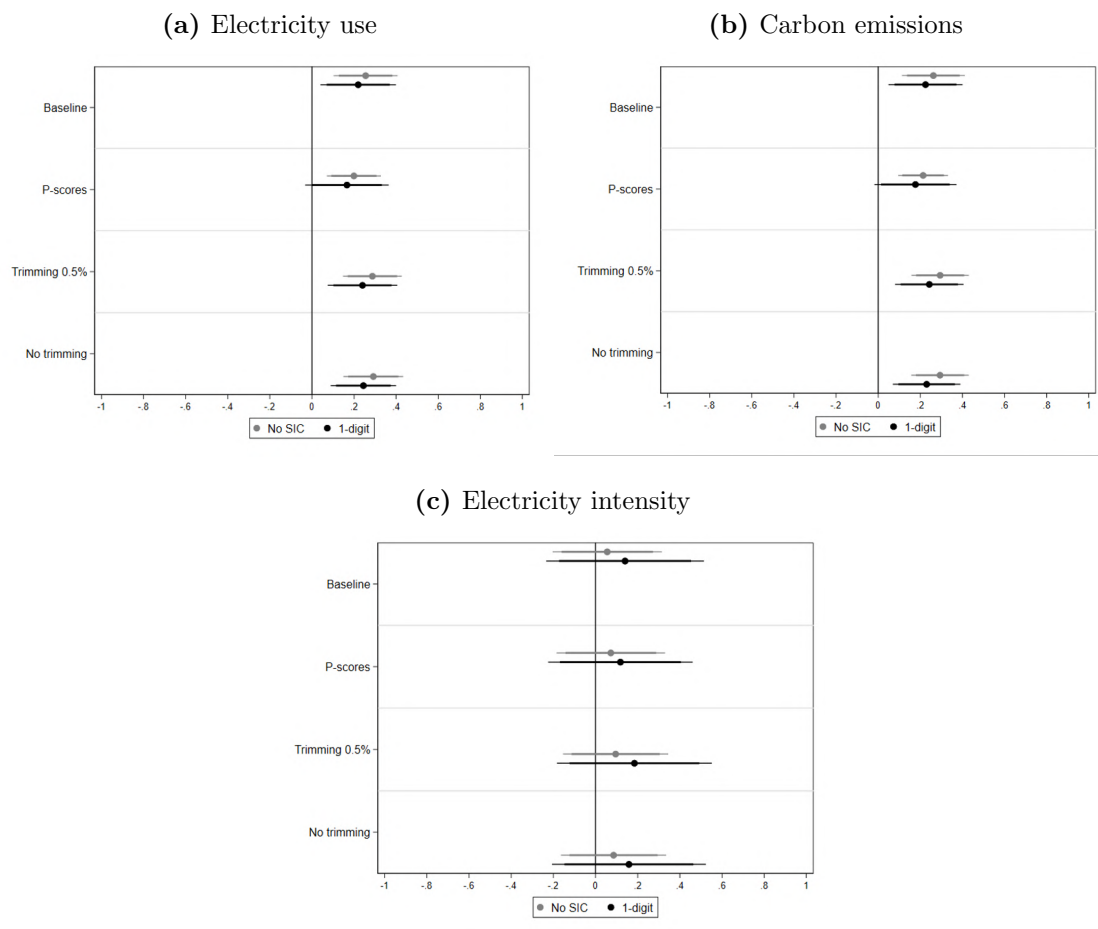
3.F Graphical comparisons of DiD estimates

Figure 3.A7: Comparison of DiD estimates (1)



Notes: All regressions include year x industry fixed effects at the 3-digit SIC code level (ABS variables) or 1-digit level (QFI variables) and are weighted by the inverse propensity score. The post-treatment period is 2013–2015. We drop plants with an electricity intensity (i.e., predicted electricity consumption/GVA) below 0.1. Data sources: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI).

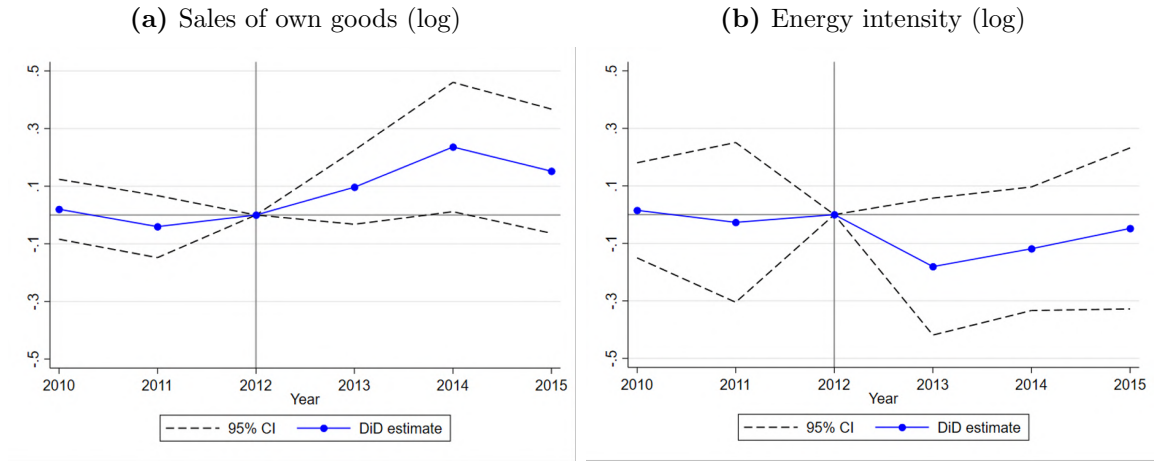
Figure 3.A8: Comparison of DiD estimates (2)



Notes: All regressions include year \times industry fixed effects at the 3-digit SIC code level (ABS variables) or 1-digit level (QFI variables) and are weighted by the inverse propensity score. The post-treatment period is 2013–2015. We drop plants with an electricity intensity (i.e., predicted electricity consumption/GVA) below 0.1. Data sources: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI).

3.G Dynamic difference-in-differences results

Figure 3.A9: Treatment effects of compensation, by year. 2010-2015.



Notes: Figures plot the coefficients $\sum_{m=0}^M \beta_{-m}$ and $\sum_{k=1}^K \beta_{+k}$ estimated from equation 3.12. The dependent variable is given by the subfigure headings. All dependent variables are in logs. The connected lines depict the estimated yearly treatment effect, while the dashed lines indicate 95% confidence intervals. We drop plants with an electricity intensity (i.e., predicted electricity consumption/GVA) below 0.01. All regressions include plant fixed effects and industry specific year dummies at the 3 digit level. Standard errors are clustered at the firm level. Data sources: Annual Business Survey (ABS) and Quarterly Fuels Inquiry (QFI).

3.H Robustness Checks: Fuzzy RDD

3.H.1 Alternative functional specifications

Table 3.A17: LATEs of compensation. Fuzzy RDD controlling for linear distance from the cut-off.

	Sales of own goods	Energy purchases	Energy Intensity
Panel A: First Stage	0.899*** (0.0891)	0.924*** (0.0848)	0.924*** (0.0848)
Panel B: Reduced Form	0.254** (0.125)	0.170 (0.220)	-0.0725 (0.150)
Panel C: Second Stage	0.282** (0.129)	0.184 (0.227)	-0.0784 (0.167)
Panel D: OLS	0.160 (0.107)	0.274* (0.139)	0.124 (0.115)
Observations	253	252	249
N Compensated	20	20	20
N Other	49	48	47
F statistics	101.69	118.73	118.63
Functional form	Linear	Linear	Linear

Notes: Tables show the coefficients estimated from the first stage, reduced form, and second stage of the fuzzy regression discontinuity design. Dependent variables are given by the table headings and are measured in logs. Each stage of the estimation includes include $post_t \times \mathbb{1}\{elig_j = 1\}$ and $post_t \times \mathbb{1}\{c_i \geq c_0\}$ and firm-level fixed effects (see Section 3.4.2 for more details). Standard errors are clustered at the firm level. The sample estimation period is 2010-2015. Cutoff value: 0.05. Bandwidth: cutoff value +/-0.007. Bandwidth refers to the range of electricity intensity values (electricity use * carbon cost / GVA) used to restrict the sample. Data source: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.A18: LATEs of compensation. Fuzzy RDD controlling for quadratic distance from the cut-off.

	Sales of own goods	Energy purchases	Energy Intensity
Panel A: First Stage	0.923*** (0.103)	0.917*** (0.104)	0.918*** (0.104)
Panel B: Reduced Form	0.273* (0.140)	0.0591 (0.174)	-0.211* (0.124)
Panel C: Second Stage	0.296** (0.138)	0.0644 (0.186)	-0.230 (0.144)
Panel D: OLS	0.150 (0.110)	0.254* (0.137)	0.107 (0.131)
Observations	253	252	249
N Compensated	20	20	20
N Other	49	48	47
F statistics	79.91	77.79	77.87
Functional form	Polynomial	Polynomial	Polynomial

Notes: Tables show the coefficients estimated from the first stage, reduced form, and second stage of the fuzzy regression discontinuity design. Dependent variables are given by the table headings and are measured in logs. Each stage of the estimation includes include $post_t \times \mathbb{1}\{elig_j = 1\}$ and $post_t \times \mathbb{1}\{c_i \geq c_0\}$ and firm-level fixed effects (see Section 3.4.2 for more details). Standard errors are clustered at the firm level. The sample estimation period is 2010-2015. Cutoff value: 0.05. Bandwidth: cutoff value +/-0.007. Bandwidth refers to the range of electricity intensity values (electricity use * carbon cost / GVA) used to restrict the sample. Data source: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.H.2 RD Results for employment

Table 3.A19: LATEs of compensation on employment with different functional forms. 2010–2015. Fuzzy RDD.

	Employment		
Panel A: First stage	0.879*** (0.101)	0.899*** (0.0891)	0.923*** (0.103)
Panel B: Reduced form	0.0612 (0.125)	0.0693 (0.0837)	0.0245 (0.0858)
Panel C: Second stage	0.0696 (0.0890)	0.0771 (0.0893)	0.0265 (0.0916)
Panel D: OLS	0.0371 (0.0557)	0.0290 (0.0569)	0.0228 (0.0522)
Observations	256	256	256
N Compensated	20	20	20
N Other	50	50	50
F statistics	75.51	101.77	79.82
Functional form	Bins	Linear	Quadratic

Notes: Tables show the coefficients estimated from the first stage, reduced form, and second stage of the fuzzy regression discontinuity design. Dependent variables are given by the table headings and are measured in logs. Each stage of the estimation includes include $post_t \times \mathbb{1}\{elig_j = 1\}$ and $post_t \times \mathbb{1}\{c_i \geq c_0\}$ and firm-level fixed effects (see Section 3.4.2 for more details). Standard errors are clustered at the firm level. The sample estimation period is 2010-2015. Cutoff value: 0.05. Bandwidth: cutoff value \pm 0.007. Bandwidth refers to the range of electricity intensity values (electricity use * carbon cost / GVA) used to restrict the sample. Data source: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.H.3 RD Manipulation Test using local polynomial density estimation

Table 3.A20 reports the results of our density discontinuity tests (or manipulation testing) following Cattaneo et al. (2020).


Table 3.A20: RD Manipulation test following Cattaneo et al. (2020)

	2010	2011	2012
p-value (robust bias-corrected)	0.2194	0.3674	0.2361
Window	\pm 0.01	\pm 0.01	\pm 0.01

3.H.4 RD bandwidth selection procedure

We implement the window-selection procedure based on balance tests for RD designs under local randomization introduced by Calonico et al. (2015, 2017). Specifically, this procedure involves constructing a sequence of nested windows around the RD cutoff and undertaking binomial tests for the running variable and hypothesis tests for a set of covariates. Then, the selected window is the largest window around the cutoff such that the minimum p-value of the balance test is larger than 0.10. To produce Figure 3.A10, we select proxies for production levels and energy intensity (i.e., sales of own goods and electricity scaled by sales as a measure of intensity, respectively) as covariates and focus on the pre-treatment period to select the largest inference window where local randomisation is assumed to hold where we can empirically show that the distribution of observed covariates does not change discontinuously at the threshold to a significant extent. Here, we report the selection of covariates that produced the most conservative (or lowest) p-values in our runs and opt for an optimal window of ± 0.007 from the cutoff. As the choice of covariates bears an arbitrary component, we run the same procedure outlined above with a different selection of covariates, and the resulting p-value for the window length that we selected (± 0.007) ranges from around 0.14 (as shown below) to around 0.5 (when we include other production values such as total output, turnover, and production value). We then test the extent to which our results are affected by different bandwidth choices in the following section. Nevertheless, due to the limited sample size around the threshold, we face a trade-off between moving closer to the threshold where the assumption of local randomisation becomes increasingly more plausible and model precision. See Section 3.4.2 for more details on our RD setting.


Figure 3.A10: RD bandwidth selection procedure



Notes: Figure plots the minimum p-value of a balance test following Calonico et al. (2020).


3.H.5 Alternative bandwidths

Figure 3.A11: Comparing LATEs on sales across different bandwidths. 2010-2015.




Notes: Figure plots the coefficients estimated from the second stage of the fuzzy regression discontinuity design. Dependent variables are indicated in the caption and are measured in logs. Each stage of the estimation includes $post_t \times \mathbb{1}\{elig_j = 1\}$ and $post_t \times \mathbb{1}\{c_i \geq c_0\}$ and firm-level fixed effects (see Section 3.4.2 for more details). Standard errors are clustered at the firm level. The sample estimation period is 2010-2015. Cutoff value: 0.05. Bandwidth: cutoff values range from +/-0.006 to +/-0.0105 with a 0.0005 step-wise increase in the estimation window from left to right. Bandwidth refers to the range of electricity intensity values (electricity use * carbon cost / GVA) used to restrict the sample. Data source: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI).

Figure 3.A12: Comparing LATEs on energy intensity (scaled by sales) across different bandwidths. 2010-2015.



Notes: Figure plots the coefficients estimated from the second stage of the fuzzy regression discontinuity design. Dependent variables are indicated in the caption and are measured in logs. Each stage of the estimation includes $post_t \times \mathbb{1}\{elig_j = 1\}$ and $post_t \times \mathbb{1}\{c_i \geq c_0\}$ and firm-level fixed effects (see Section 3.4.2 for more details). Standard errors are clustered at the firm level. The sample estimation period is 2010-2015. Cutoff value: 0.05. Bandwidth: cutoff values range from +/-0.006 to +/-0.0105 with a 0.0005 step-wise increase in the estimation window from left to right. Bandwidth refers to the range of electricity intensity values (electricity use * carbon cost / GVA) used to restrict the sample. Data source: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI).

Figure 3.A13: Comparing LATEs on energy purchases across different bandwidths. 2010-2015.



Notes: Figure plots the coefficients estimated from the second stage of the fuzzy regression discontinuity design. Dependent variables are indicated in the caption and are measured in logs. Each stage of the estimation includes $post_t \times \mathbb{1}\{elig_j = 1\}$ and $post_t \times \mathbb{1}\{c_i \geq c_0\}$ and firm-level fixed effects (see Section 3.4.2 for more details). Standard errors are clustered at the firm level. The sample estimation period is 2010-2015. Cutoff value: 0.05. Bandwidth: cutoff values range from +/-0.006 to +/-0.0105 with a 0.0005 step-wise increase in the estimation window from left to right. Bandwidth refers to the range of electricity intensity values (electricity use * carbon cost / GVA) used to restrict the sample. Data source: Annual Business Survey (ABS), and Quarterly Fuels Inquiry (QFI).

Chapter 4

Climate policy uncertainty and the behavior of firms and investors

with Stefano Carattini (Georgia State University), Antoine Dechezleprêtre (OECD), and Tobias Kruse (OECD)

SUMMARY. Whether and how firms are affected by uncertainty revolving around the implementation of climate policy is pivotal to fostering a low-carbon transition and has implications for the potential systemic risk related to the coordinated implementation of ambitious climate policy. We develop a new index of climate policy uncertainty, covering the United States with monthly-level variation between 1990 and 2018. We leverage the variations in our index to analyze the relationship between climate policy uncertainty and firm-level outcomes such as stock volatility, share price, investments in research and development, and employment for all publicly listed firms in the country. We find that climate policy uncertainty tends to considerably affect these outcomes, and often more so than existing indices of economic policy uncertainty. The direction of the uncertainty matters as well, as measured by sub-indices capturing whether the uncertainty reflects potential acceleration or deceleration in climate policymaking.

4.1 Introduction

Understanding the behavior of agents such as firms and investors is a crucial component of economics, with important implications for society at large. Often, firms' and investors' decisions are analyzed with respect to a change in policy or another aspect that determines the environment in which they operate. However, at least as often firms and investors need to take decisions in a context of substantial uncertainty. Within this context, economists have long recognized the role of uncertainty (Bernanke, 1983; McDonald and Siegel, 1986), but only relatively recently started measuring it in a systematic way (Baker et al., 2016; Hassan et al., 2019).

Climate change is one of the most pressing issues of this century. The need to mitigate climate change has been known among scientists for some four decades and in policy circles at least since the early 1990s. However, recent human history shows that there is much more uncertainty on the implementation of climate policy, domestically and internationally, than there is on climate change itself (Moore et al., 2022). Hence, climate change, with its all-encompassing need for change, offers a suitable context to examine the behavior of firms and investors under uncertainty. How firms and investors respond to uncertainty related to climate policy also has implications for climate mitigation. In particular, firms' expectations about future climate action may influence their decisions concerning innovation as well as the choice of inputs, in particular labor and capital, which are crucial dimensions for the transition to a cleaner economy. Yet, these dimensions have been examined mostly in response to actual policy changes (Martin et al., 2014a; Aghion et al., 2016; Calel and Dechezlepretre, 2016; Yamazaki, 2017).

Further, analyzing firms' and investors' responses to uncertainty in climate policy-making is also informative for the analysis of *transition risk*, potential systemic risk driven by the relatively abrupt implementation of ambitious climate policy after decades of delay. With the Paris Agreement, countries committed to reducing green-

house gas emissions to keep temperature increases within 1.5-2°C above pre-industrial levels. Each signatory pledged to reduce emissions, in absolute terms or relative to a business-as-usual scenario (Tobin et al., 2018). With the quantity-based approach behind the Paris Agreement, voluntary pledges set the ambition. Then, policymakers need to identify ways to make sure that a set of instruments is implemented to meet the pledges. Hence, policies systematically trail ambition (Harstad, 2022). One implication that follows from this approach is that firms may be misaligned with long-run climate goals.

This misalignment could have, in turn, two main implications: first, firms may continue to invest in “dirty” technologies, leading to continued emissions as well as potential asset stranding and investors’ losses once policy risk is materialized (see van der Ploeg and Rezai, 2020b for a review). As the former governor of the Bank of England Mark Carney made clear in his well-known 2015 speech (Carney, 2015), private losses should not be a concern for the regulator. Asset values should reflect fundamentals and investors have known about the need to tackle climate change for decades. However, adjustments in the stock market due to the abovementioned asset stranding could potentially lead to a systemic shock, especially considering that carbon-intensive sectors can represent up to half of an advanced economy’s standard portfolio (Battiston et al., 2017; ECB, 2021). Many influential voices have raised concerns about systemic risk related to a potentially abrupt transition to a low-carbon economy, including central banks and financial regulators in some of the world’s major economies (Carney, 2015; Vermeulen et al., 2018; Banque de France, 2019; Rudebusch, 2021). Hence, it is of fundamental importance to examine the behavior of firms and investors in presence of uncertainty on the likelihood and timing of future climate policy developments.

How does uncertainty in climate policy affect the behavior of firms and investors? To address this question, we built the first index of policy uncertainty specific to climate policy, which allows us to address this question empirically. Our “climate policy uncertainty” index, or CPU, combines the original search strategy in Baker

et al. (2016) with keywords related to climate policy. Our index runs monthly from 1990 to 2018 and covers the main newspapers in the United States. Then, we analyze the relationship between CPU and firm outcomes such as share prices, volatility, employment decisions, as well as investments in research and development.

Our approach also takes into account a crucial feature related to climate policy. While in the case of standard economic policy, the economy tends to move along a given trajectory determined by its steady state and uncertainty tends to be detrimental to economic growth, in the case of climate change the economy needs to transition from fossil-fueled activities to a cleaner way of production. Hence, the economy needs to move from one equilibrium, which is carbon intensive, to another equilibrium, which is much cleaner. Since climate change entered the policy arena in the 1980s, both domestic and international climate policymaking have gone through important achievements as well as numerous setbacks. If firms and investors respond to short-term variation in the probability of future policy tightening, rather than adopting long-term goals such as decarbonization, setbacks are likely to benefit them. For this reason, our index is complemented by two sub-indices, aimed at measuring whether the source of uncertainty is an acceleration in the process of decarbonization, or rather a deceleration.

The primary empirical goal of this chapter is to examine how economic outcomes respond to greater uncertainty about climate policy, also depending on its drivers. To do so, we exploit variations in our Climate Policy Uncertainty (CPU) index, and its sub-indices, across different months, quarters, or years from 1990 to 2018. Specifically, we estimate fixed effects models where we interact our news-based indices with the average carbon intensity across narrowly-defined industries. By doing so, we develop an identification strategy that differentiates firms according to their relative exposure to climate policy risk. Using panel data on publicly-listed companies, our model tests whether exposure to climate policy risk matters for economic outcomes when greater uncertainty about climate policy materializes as measured by newspaper article coverage.

Overall, we find that an increase in our index is associated with greater stock price volatility and lower share prices, as well as reductions in R&D efforts and annual employment levels. Our back-of-the-envelope calculations suggest that climate policy uncertainty over the last two decades is responsible for an average upward shift in volatility of approximately 3%, a decrease of around 13% in share prices, and decreases in R&D investments of around 12%. Additionally, we detect negative but limited effects on employment. Yet, our results exhibit considerable heterogeneity exist across industries, with more pronounced effects observed in carbon-intensive sectors. Our analyses also suggest that the direction of the estimated effects matters, so the impact of climate policy uncertainty on the outcomes of interest depends on the underlying drivers of climate policy uncertainty. Finally, further empirical investigations reveal that firm-level economic outcomes are more sensitive to uncertainty about climate policy when changes in expectations point towards more stringent regulation in the future. Our results are robust to a host of sensitivity tests.

We contribute to four strands of literature. First, a growing literature examining the role of policy uncertainty on a wide range of outcomes (Bernanke, 1983; McDonald and Siegel, 1986; Hassett and Metcalf, 1999; Handley and Limão, 2015; Baker et al., 2016; Hassan et al., 2019), including investments in green technologies in a set of specific contexts (Fabrizio, 2013; Dorsey, 2019). We contribute to this literature by introducing an index of climate policy uncertainty, which allows us to examine firms' and investors' responses to changes in the probability of climate policy tightening for the largest firms in the United States over about four decades.

Second, a recent theoretical literature on firms' and investors' decisions under the specter of future climate policy (Rozenberg et al., 2018; van der Ploeg and Rezai, 2020a; van Benthem et al., 2022), including implication in terms of systemic risk (Carattini et al., 2021; Diluiso et al., 2021), and a recent set of empirical applications testing the theory (Carattini and Sen, 2019; Sen and von Schickfus, 2020; Engle et al., 2020; Krueger et al., 2020). We contribute to this literature by providing additional empirical evidence on a range of firm-level outcomes from shocks in climate policy

uncertainty.

Third, a stream of research examining the role of innovation in response to environmental regulation, analyzing, theoretically and empirically, the role of directed technical change (e.g., Bovenberg and Smulders, 1995; Porter and van der Linde, 1995; Popp, 2002; Acemoglu et al., 2012; Aghion et al., 2016; Caelal and Dechezlepretre, 2016; see also Ambec et al. 2013 for a review). Unlike the existing literature, which infers mostly from existing policies, leveraging changes in stringency, our study focuses on variations in uncertainty, leading to adjustments in firms' beliefs about the likelihood of future policy tightening or weakening.

Fourth, empirical literature shows relatively muted changes in employment following the tightening of environmental regulation (Martin et al., 2014a; Yamazaki, 2017). Also in this case, we contribute to the literature by covering changes in uncertainty about potential regulatory changes, rather than only realized policy shocks, and uncovering their effects on employment levels over a long period for a large number of firms.

The chapter proceeds as follows. Section 4.2 introduces our CPU index as well as its sub-indices. Section 4.3 describes the data and empirical approach. Section 4.4 presents our empirical results. Section 4.5 concludes.

4.2 Introducing the CPU index

4.2.1 Building the index

This study builds upon the work of Baker et al. (2016) in order to develop an indicator of climate policy uncertainty using a comparable methodological approach, which we detail in what follows. To build their index of Economic Policy Uncertainty in the U.S., Baker et al. (2016) count the frequency of newspaper articles that

contain the following trio of terms: (1) “economic” or “economy”; (2) “uncertain” or “uncertainty”; and (3) “Congress”, “deficit”, “Federal Reserve”, “legislation”, “regulation” or “White House”. To build our index of Climate Policy Uncertainty (CPU), we similarly created a lexicon of words and combinations of words as our search strategy. To ensure that we capture the right concept, we create a separate lexicon of words for each of the three components (Climate, Policy, and Uncertainty). The first category includes terms such as “pollution”, “CO2”, or “climate change” which refer to a specific concern related to climate change. It also includes terms referring to technologies addressing these concerns such as “solar PV” or “renewable”. The second category includes terms related to policy-making such as “regulation”, “legislation”, or “tax”, but also terms more specific to environmental policies such as “emissions trading scheme” or “cap and trade”. The full list of keywords used in these two components is listed in Appendix 4.A. The third category includes the words “uncertain” or “uncertainty”. Selected articles have to include at least one term from each category.

We initially created the lexicon in English in order to capture articles in English-speaking countries. All keywords were then translated, by native speakers, in several other languages. Appendix 4.A provides the keyword selection for all languages.¹ In this chapter, we use the English version, applied to the United States. The main challenge in creating an indicator of policy uncertainty based on counts of newspaper articles is the possible inclusion of “false positives”, which are articles that are not relevant but are still selected based on the search strategy. Such false positive results would inflate our index and incorrectly indicate higher levels of uncertainty. In turn, they would introduce a downward bias in the empirical analyses. To reduce as much as possible the likelihood of including such false positives, we read several hundreds of

¹To ensure that our index is consistently observed across countries, we avoid using country-specific terms. For example, we do not include the exact name of environmental ministries, departments, or environmental protection agencies. The names of ministries or departments dealing with environmental and climate change topics tend to change with governments, which makes them difficult to track consistently across countries and time. In the United Kingdom, for instance, the Department for Energy and Climate Change became part of the Department for Business, Energy Industrial Strategy in July 2016 following a change in government.


randomly selected articles and recursively adjusted the search strategy. We manually coded the randomly selected articles as relevant and irrelevant. We were thereby able to adjust the search strategy systematically to increase the ratio of relevant articles to above 80%, which is considered a reasonable compromise between including as many relevant articles as possible and limiting the extent of false positives.

To ensure that the selected articles talk about climate policy and not about the climate in one part and about unrelated policies in another, we imposed the restriction that terms from the policy category have to be located within the same paragraph from the respective word in the climate category. We thereby contribute to ensuring that the two terms are related to each other in the newspaper article. A difficulty when using terms such as “environment” or “climate” is that they can also be used to describe other concepts such as “business climate”, “business environment” or “policy environment”. We therefore explicitly excluded all articles that used one of these expressions. An additional challenge in creating topic-specific policy uncertainty indices is that they tend to require many more search terms compared to general economic policy uncertainty indicators. This is necessary to ensure that as many topic-related events as possible are picked up. Baker et al. (2016) are able to obtain comprehensive coverage of economic policy uncertainty with ten search terms for the United States. For our climate policy uncertainty index, we apply more than 60 search terms. Since newspaper coverage of climate-related policy uncertainty is typically smaller than coverage of economic policy uncertainty, our search strategy needs to be sufficiently sensitive in order to observe as many topic-specific events as possible. For the United States, the CPU index covers the years from 1990 until 2018. The main reason to start the index in 1990 is that, prior to that date, the number of available newspaper articles is smaller and potentially too small.

To construct their indicator of economic policy uncertainty for the United States, Baker et al. (2016) use data from 10 leading newspapers. Limiting the search to leading newspapers ensures the quality of the underlying articles and avoids including newspapers that only exceptionally report on the topic, spuriously creating huge

volatility over time. For each newspaper, we separately downloaded the annual count of articles that are picked up by our search strategy as well as the total number of articles published by the outlet. Two online newspaper databases were used to download the article counts, Factiva and Nexis, covering different sets of newspapers. As an illustration, Figure 4.1 shows the annual article counts for the *New York Times* (United States). These time series show the trends in overall articles (left axis) and in articles on climate policy uncertainty (right axis). The number of annual articles related to climate policy uncertainty varies between 0 and 300, with a significant year-on-year variation. Overall, the frequency of articles on climate policy uncertainty appears to have increased in the recent period, but the total number of articles published has increased as well.

Figure 4.1: Article counts in the *New York Times* (US)




Note: Yearly series from 1990 to 2018.
Source: Factiva.

In order to account for this rising trend in total articles published, we first compute a simple newspaper-specific ratio of articles on climate policy uncertainty over the total article count by newspaper. This ratio is displayed in Figure 4.2 for the same

newspaper, the New York Times. Over time, less than 2 in 1000 articles deal with climate policy uncertainty in the New York Times, further justifying our choice to use multiple keywords to cast as wide a net as possible given the specificity of the topic of interest in the general press.

Figure 4.2: Ratio of CPU articles over total articles in the *New York Times* (US)




Note: Based on yearly series from 1990 to 2018.
Source: Factiva.

A challenge with these raw article ratios is that the number of articles varies a lot across newspapers and time, making it difficult to simply average the ratios across several newspapers in a given country. We, therefore, apply the standardization approach of Baker et al. (2016) to obtain our CPU index. We begin with the simple ratio of articles on climate policy uncertainty divided by the total article counts for each newspaper, as illustrated in Figure 4.2. For each newspaper, we then divide this ratio by the newspaper-specific standard deviation across all years. This creates a newspaper-specific time series with unit standard deviation across the entire time interval, which ensures that the volatility of the overall country-level index is not driven by the higher volatility of a particular newspaper. We then average these

standardized series across all newspapers within each country by year. Lastly, we normalize the country-specific series to a mean of 100 over the time interval.

Figure 4.3: CPU index in the United States




Note: Based on yearly series from 1990 to 2018.
Source: Authors' own calculations based on newspaper articles from Factiva.

4.2.2 Validating the CPU index

As a first approach to validate our index, and following Baker et al. (2016), we link the country-specific peaks to relevant events such as the discussion or implementation of major climate policies. To verify that our index varies in conjunction with the respective events and the corresponding realization of uncertainty, we read the headlines of the first hundred articles that were downloaded for the peak years. In the United States, the index has pronounced peaks in 2001, 2010, and 2017, as shown in Figure 4.3. The first peak in 2001 is linked to the Energy Plan published by the George W. Bush administration that included environmental deregulation, in particular with

respect to oil and gas explorations. While the event itself created climate policy uncertainty by lowering environmental standards, the lengthy discussion around the publication of the plan also contributed to the spike in the index. The spike in 2010 is driven by the Democratic party withdrawing a major bill on climate change due to insufficient support in Congress. Moreover, the prior discussion on whether the bill might achieve sufficient support in Congress and whether the Democratic party might be willing to amend the bill contributed to the uncertainty. The third spike in 2017 is in turn related to uncertainty arising from President Trump’s withdrawal from the Paris Agreement and efforts to revoke clean energy and climate policies. Appendix 4.B provides an extended list of major events in the United States relevant to climate policy uncertainty, which we leverage later in this section as well as in the remainder of the chapter. Visibly, such major events related to both instances of progress as well as setbacks in dealing with climate change, supporting the generation of sub-indices, as described in the next section, capturing these two forces, respectively.

Figure 4.4: CPU index and associated events in the United States




Note: Based on yearly series from 1990 to 2018.
Source: Authors’ own calculations based on newspaper articles from Factiva.

In addition to the annual time series of the CPU index, we are also able to establish

a monthly index for the United States, where the annual number of articles related to climate policy uncertainty is high enough to be further disaggregated. This more granular data allows us to examine the variation in the index in more detail, which we do in Figure 4.5, as well as to analyze responses by high-frequency variables such as share prices and volatility to climate policy uncertainty shocks, which we do in the following sections, among other outcome variables, using either monthly or quarterly series.

Figure 4.5: CPU index and associated events in the United States



Note: Based on monthly series from 1990 to 2018.
Source: Authors' own calculations based on newspaper articles from Factiva.

Based on a careful analysis of the newspaper article headlines and abstracts we are again able to link the peaks to particular policy events and their corresponding realization of uncertainty. Interestingly, we observe a trend in the topics of climate policy uncertainty moving from energy-related issues largely concerned with energy security and energy prices in the 1990s and early 2000s to increasing attention explicitly given to air pollution and climate change issues from the late 2000s onwards.² The early events include the uncertainty around energy prices following

²Since our baseline search strategy includes keywords potentially relating to local air pollution,

Iraq's invasion of Kuwait as well as the above-mentioned discussion around the Energy Plan of the George W. Bush administration in 2001 that included deregulation in particular for oil and gas exploration. The later events include in particular uncertainty arising around the discussion and abrupt withdrawal of a bill to regulate ozone emissions in September 2011 under the Obama administration as well as the election of President Trump, who then announced a planned withdrawal of the United States from the agreement.

It is, however, important to note that, by design, annual (Figure 4.3) and monthly (Figure 4.5) time series can identify different peaks. Such difference can arise if, for instance, the discussion of a policy change spreads across many months within a single year. The frequency per month may be relatively low, but if all the articles are aggregated within a year, they can lead to a peak in the annual time series. In the United States, this occurred for instance with the 2010 withdrawal of the climate change bill under the Obama administration. While it appears as a spike in the yearly chart, the spike in the monthly series is less marked. Figure 4.3 shows elevated levels of climate policy uncertainty throughout 2010. The withdrawal of the bill was not a major surprise as it had already appeared that the administration did not have sufficient support in Congress to see it pass. Therefore, the combination of both annual and monthly time series provides unique insights as it allows us to examine all policy events from both perspectives. Section 4.D further compares the index to other relevant measures, such as the EPU from Baker et al. (2016), the Chicago Board Options Exchange's CBOE Volatility Index (hereafter referred to as VIX), and oil price volatility.


4.2.3 Extending the index

We provide two extensions to the index, which are documented in more detail, including the keyword searches, in Appendix 4.A. First of all, since our baseline

we run an additional newspaper article search that excludes them. More details are provided in Section 4.2.3.

search strategy includes keywords potentially relating to local air pollution, we run a new newspaper article search that excludes them. Figure 4.6 plots the evolution of the two indices since 1990. Overall, they exhibit a correlation of 0.9923. The index resulting from the narrower search, which we denote as N-CPU for Narrow Climate Policy Uncertainty, is used for robustness tests in Section 4.4.2.

Figure 4.6: Comparing the evolution of the CPU and N-CPU indices




Note: Based on quarterly series from 1990 to 2018.
Source: Authors' own calculations based on newspaper articles from Factiva.

The second extension follows from an important observation about the difference between our CPU index and the EPU index developed by Baker et al. (2016). While in the case of standard economic uncertainty, any increase in the index, and thus in the underlying uncertainty, is overall detrimental to economic output, the case of climate policy uncertainty is very different. The EPU index largely measures the effect of uncertainty as a destabilizing factor from a trajectory of economic growth. In contrast, the CPU index measures the uncertainty surrounding the pace at which the economy is expected to move from business as usual to carbon neutrality. In this context, there is a trade-off between current economic output and climate

change mitigation, so an increase in climate policy uncertainty has two effects: first, a negative shock on economic output due to the direct effect of uncertainty, as analyzed by Baker et al. (2016); second, an effect that depends on how beliefs on the pace of the transition towards a cleaner economy are adjusted. Indeed, the process of implementing climate policy, both domestically and internationally, has had many instances of acceleration and deceleration. While when climate change entered the political arena in the '80s and '90s expectations might have been that of a relatively quick transition to fewer fossil fuels, as recommended by scientists, it later became apparent that (international) climate change mitigation would have been harder to achieve than coordination in banning products responsible for ozone depletion as done with the Montreal Protocol. In more recent times, however, unilateral initiatives, followed by the Paris Agreement, and the emergence of a new generation of environmental leaders, have pointed to an acceleration in climate change mitigation. Over only a few years, carbon pricing went from covering 15% of global emissions to about 22.5% (The World Bank, 2023). Unless investors are totally aligned with long-term climate goals as provided by climate scientists and unmoved by present political developments, which do not seem to be the case (see e.g. Carattini and Sen, 2019), we would expect stock markets to make gains when new developments point to additional delays in climate action and to make losses, everything else equal, when new developments point to an acceleration in climate action.


Hence, it is important not only to analyze variation in the CPU index but also to try to disentangle its drivers, whether an increase in uncertainty suggests that the transition is slowing or accelerating. To this end, we performed two additional separate searches, adding keywords related to progress and failure, respectively, to the standard keyword search. We denote the resulting sub-indices as CPU+ when belief revision goes towards more climate action (hence “plus” for more action) and CPU- when belief revision goes towards less climate action (hence “minus” for less action). Figures 4.7 and 4.8 plot the evolution of the sub-indices over time, linking respective index-specific peaks to policy-relevant events.

Figure 4.7: Quarterly CPU- index in the United States



Note: Based on quarterly series from 1990 to 2018.
Source: Authors' own calculations based on newspaper articles from Factiva.

Figure 4.8: Quarterly CPU+ index in the United States



Note: Based on quarterly series from 1990 to 2018.
Source: Authors' own calculations based on newspaper articles from Factiva.

4.3 Outcome data and empirical strategy

4.3.1 Outcome variables and descriptive statistics

To examine whether climate policy uncertainty has an impact on economic outcomes, we combine several data sources on publicly listed firms. In particular, we investigate firms' and investors' responses to uncertainty about climate policy by focusing on share prices and volatility, research and development (R&D) expenses, and employment. All variables except volatility are obtained from Standard Poor's Compustat, specifically from Compustat North America, which includes information for companies listed in the United States and Canada. Additionally, we combine information retrieved from Options Metrics, which provides firm-level historical volatility over different time horizons since the mid-1990s in our main estimations, as well as longer time horizons in alternative specifications. Table 4.1 is based on firm-level information for publicly listed companies in the US between 1990 and 2018.

Table 4.1: Summary statistics

Variable	Obs	Mean	Std. Dev.	Min	Max	Median
Volatility (30 days)	298103	-1.042	.637	-6.398	2.47	-1.043
Share Price	1433855	.87	3.247	-13.816	13.423	2.053
R&D	288375	.768	2.324	-6.908	9.299	.761
Employment (Annual)	220356	8.692	38.375	0	2545.209	.574

Notes: Table shows summary statistics for publicly-listed companies in the US between the years 1990-2018 without sample restrictions. Variables are expressed in log terms.

We use daily data on volatility from Option Metrics' volatility which provides historical information on firm-level 30-day volatility. We include in our dataset all options that have been traded on the Chicago Board of Options and Exchange since 1996.³ In line with Baker et al. (2016), we consider historical share price volatility as a proxy for firm-level uncertainty. We then combine information on share prices,

³As the information on stock-price volatilities is only available from 1996, we check whether our results change when running our estimations for other outcome variables from 1996 onward

research and development (R&D) expenses, and employment for the entire universe of publicly traded firms in the United States since 1990. Our analysis with share price as an outcome variable covers around 10,000 listed companies in the United States. Share prices refer to a stock's closing price, which is the standard benchmark used by investors to track its performance in time. Compustat North America provides information on share prices since 1962, which allows us to analyze the relationship between this outcome variable and the index since 1990. R&D expenses are included in the sample to proxy firm-level innovative behavior. These expenses are defined as the costs incurred throughout a given quarter that cover the development of new products or services. Information on R&D expenses is only available since 1989. Finally, employment refers to the annual level of employees in a given company since 1990.

Further, we are interested in analyzing whether the above-mentioned economic outcomes are differently affected based on the exposure to climate policy risk of each firm, which we proxy by emission intensity. To this end, we combine information on emissions from the US Environmental Protection Agency (EPA). The EPA tracks facility-level emissions of air pollutants, through the Greenhouse Gas Reporting Program (GHGRP). The GHGRP collects annual information on the emissions of different greenhouse gases, primarily on carbon dioxide (CO₂).⁴

The reporting program provides data on individual facilities, thus offering an opportunity to disaggregate nationwide emissions estimates to narrowly defined industries or specific companies. The database covers approximately 85% to 90% of total greenhouse gas emissions in the United States from 2010 to 2018.⁵ This includes data

⁴Carbon dioxide is the greenhouse gas (GHG) emitted in the largest quantities: carbon dioxide emissions reported in 2018 represented 90.9% of the total emissions of GHGs reported during the year. Other greenhouse gases covered include methane (CH₄), nitrous oxide (N₂O), and fluorinated GHGs (HFCs, PFCs, SF₆). In 2018, methane emissions represent 7.6% of total GHG emissions, N₂O represented around 1.0%, and fluorinated gases accounted for around 0.5%.

⁵There are specific thresholds above which reporting is required within a given industry. In general, the threshold is set at 25,000 metric tons CO₂-e per year. However, all facilities in the following industry categories must report regardless of annual emissions: Electricity Generation, Petroleum Refineries, Adipic Acid Production, Ammonia Manufacturing, HCFC-22 Production from HFC-23 Destruction, Nitric Acid Production, Petrochemical Production, Phosphoric Acid Production, Silicon Carbide Production, Titanium Dioxide Production, Aluminum Production,

on direct emissions reported by stationary sources, covering nearly all direct emissions from electricity generation and most emissions from industry, which account for approximately 50% of total nationwide emissions. In addition, this also includes GHG data reported by suppliers of fossil fuels and industrial gases, which account for the vast majority of emissions from transportation, commercial, and residential sources, representing roughly 40% of total US emissions. The GHGRP does not include emissions from the agriculture and land use sectors or other small sources of emissions.

Our main model specifications differentiate firms by their relative exposure to climate policy changes. The underlying intuition is that more pollution-intensive firms would be more exposed to the possibility of more stringent climate regulation in the future. To compute this exposure, we draw on facility-level information on air emissions from the GHGRP. As a first step, we match Compustat firms to reporting facilities using the names of their parent companies, which is provided by the EPA.

We do so using Standard Poor Capital IQ's Identifier Converter which allows identifying company identifiers of all public firms using company names. The GHGRP database includes detailed ownership percentages of facilities by multiple parents, and we rely on these values to assign each facility's pollution to its parent companies. Through this match, we yield parent firms' annual levels of carbon dioxide emissions, which we use to obtain firm-level intensities as the ratio of total air emissions to total revenue. We then aggregate emission intensity levels to obtain the ratio of air emissions to revenues in each four-digit industry by year.

Finally, we average these ratios to compute our exposure measure for each four-digit SIC industry. Table 4.2 displays intensity by 4-digit SIC code averaged across main industry group classifications. Nevertheless, there is substantial variation in average carbon intensity across 4-digit SIC codes within industry groups. For instance, carbon intensity in manufacturing ranges from relatively low values in the food

Cement Production, Lime Manufacturing, Soda Ash Production. More information on reporting requirements by industry can be found [here](#).

industry to much higher levels for the manufacturing of cement and metal products. Specifically, Cookies Crackers (SIC 2052) exhibits an average intensity of around 0.6 metric tons of carbon emissions per million of revenue generated compared to 3310 and almost 4300 metric tons per million in Cement, Hydraulic (SIC 3241) and Fabricated Metal Products (SIC 3490) respectively. Similarly, intensity values in Services range from 0.09 metric tons/million in Life Insurance (SIC 6311) to 4000 metric tons/million in Oil Royalty Traders (SIC 6792). Tables 4.A1 - 4.A4 in the Appendix provide the corresponding intensity figures for a number of other selected industries.

Table 4.2: Average carbon intensity by SIC code

Industry Description	Range of 4-digit SIC Codes	Average Intensity
Mining	1000-1499	9.92
Construction	1500-1799	9.27
Manufacturing	2000-3999	8.52
Transport, Communications, Electric, Gas and Sanitary Service	4000-4999	9.11
Wholesale and Retail Trade	5000-5999	8.79
Finance, Insurance and Real Estate	6000-6799	8.97
Services	7000-8999	7.61
Median Sample Intensity	9.34	

Notes: Industry-level intensities are expressed as natural logs and averaged across SIC codes. CO₂ emissions are retrieved from the EPA Greenhouse Gas Reporting Program (GHGRP). Carbon intensities are measured in myriagrams CO₂-e to total revenue (in millions of dollars).

4.3.2 Empirical Strategy

The main empirical goal of this chapter is to examine how selected economic outcomes respond to greater uncertainty about climate policy, also depending on its drivers. To do so, we exploit variations in the Climate Policy Uncertainty (CPU) index across different months, quarters, or years from 1990 to 2018. Our empirical strategy consists in estimating fixed effects models where we interact our news-based index with the measure of exposure to climate policy risk described in Section 4.3.1. This additional source of variation allows controlling for unobserved time-varying confounders. These specifications test whether effects for firms with greater exposure to climate policy shocks covary more strongly with our index. In other words, the model tests whether exposure to climate policy risk matters for economic outcomes when changes in

uncertainty about climate policy materialize in the news. We estimate the following equation:

$$\ln y_{it} = \beta_1 CPU_t \times Exp_j + \beta_2 X'_{jt} + \delta_t + \gamma_i + e_{ijt}, \quad (4.1)$$

where y_{it} represents one of the outcome variables presented in Section 4.3.1. CPU_t refers to our Climate Policy Uncertainty index in a given time period t , whereas Exp_j refers to our intensity measures computed for each 4-digit SIC industry, j . e_{ijt} is the idiosyncratic error term. The main identifying assumption in the model is that companies operating in high-emitting sectors tend to be more exposed to climate policy uncertainty. One potential threat to identification is firm and time-specific shocks. By including firm-specific fixed effects, γ_i , and time fixed effects, δ_t , we are able to capture time-constant firm-specific factors as well as absorb unobserved time-varying shocks. Without the interaction term, CPU_t is collinear with the time fixed effects and drops out from the equation.

Furthermore, we include a vector of controls, X' , to evaluate to what extent our CPU measure tells us anything different from other measures of uncertainty and policy uncertainty. First, the most obvious choice is to control for variation in the Economic Policy Uncertainty (EPU) index developed by Baker et al. (2016). By doing so, we can assess whether our climate policy uncertainty index can be a significant predictor of firm-level economic outcomes after controlling for the impact of general economic policy uncertainty. Both indices are constructed using scaled frequency counts of newspaper articles, but they differ conceptually. While the EPU index is designed to measure policy-related uncertainty for the economy as a whole, our CPU index quantifies uncertainty specifically related to climate policy. Drawing on Baker et al. (2016), we interact $\ln(\text{EPU})$ with SIC-specific contract intensity. The latter reflects the average ratio of federal purchases to revenue in each four-digit industry and captures exposure to uncertainty about government purchases. The intuition is that the effects of economic policy uncertainty tend to be more sizable in industries disproportionately relying on direct sales of goods and

services to the federal government. Secondly, we test whether our results change when controlling for overall economic uncertainty, approximated by the VIX index - a common measure of expectations of further stock market volatility computed as the 30-day volatility in the S&P500 index. Even in this case, we interact $\ln(\text{VIX})$ with SIC-specific contract intensity to differentiate firms by their exposure to overall uncertainty. Finally, we additionally control for fluctuations in the price of oil proxied by the West Texas Intermediate (WTI). This is because persistent spikes in oil prices may also ultimately affect the performance and thus the valuation of companies largely relying on carbon-intensive production processes.

4.4 Empirical results

4.4.1 Firm and investor behavior in response to uncertainty shocks

We are interested in firms' and investors' responses to uncertainty along the following outcome variables: share prices and volatility, research and development expenses, and employment. In our empirical analyses, we focus first on our main index and analyze its relationship with our main outcomes of interest. Then, we test the robustness of our main findings with a host of sensitivity tests. Lastly, we analyze belief revision, leveraging the sub-indices defined as CPU+ and CPU-.

We now describe the analyses using our main index. We start with share prices and share price volatility. Table 4.3 displays results from regressing firms' 30-day historical stock price volatility and share prices on climate policy uncertainty. Our estimates of interest are reported in the first row, with robust standard errors clustered at the firm level. Taking advantage of the high-frequency nature of stock market variables, we provide estimates using both monthly and quarterly time series of our index. Both levels of aggregation provide unique perspectives into the evolution of climate policy

uncertainty and contribute to providing a more complete picture for our empirical analysis (see section 4.2). While the quarterly time series provides insights into the effects of enduring uncertainty across months, the monthly series allows investigating prompt responses to uncertainty shocks occurring within shorter time frames. In our regressions, we use 30-day volatility implied by firm-level equity options. We calculate the average volatility over all trading days in a given month or quarter to match stock market data. Our sample extends from 1990 to 2018, as most of our outcome variables are consistently available from 1990 onward only (as described in Section 4.3.1). However, recall that information on stock price volatility is available from 1996 onward only. Hence, in Table 4.7 we analyze all outcome variables using 1996 to 2018 as an estimation window.

Tables 4.3 and 4.4 report results from our monthly and quarterly-level specifications respectively. Overall, our monthly-level estimates indicate that an increase in our index is associated with greater historical stock price volatility and lower share prices. Specifically, we find that for a firm with median exposure, a 1% increase in CPU over a given month leads to an increase of around 0.05% in volatility ($0.0055 \times 9.34 = 0.047$) and a reduction of 0.09% in share price ($0.01 \times 9.34 = 0.09$). In line with our expectations, we observe that firms operating in more carbon-intensive 4-digit SIC industries tend to respond more strongly to variations in climate policy uncertainty. Table 4.4 shows how these estimates change when turning to our quarterly-level specifications. Overall, both specifications yield similar results, but the magnitude of the estimated relationships is larger with quarterly series. To assess these magnitudes, our quarterly-level coefficients now predict for a firm with median exposure that a 1% increase in CPU would lead to an increase of 0.08% in volatility and a reduction of 0.3% in share price. These results reveal that stock market performances tend to be more sensitive to spikes in climate policy uncertainty when the latter persists over multiple months. Table 4.A5 in the Appendix extends our approach to the annual series. Even in this case, results suggest that the more persistent the shock, the larger the effect. Furthermore, to put our coefficients into perspective, the quarterly

CPU index rose on average by 40.5 log points from 2000 to 2018. Assuming a median exposure, this implies an estimated upward shift in volatility of approximately 3% ($0.405 \times 0.00816 \times 9.34 \times 100$) and an overall decrease of around 13% ($0.405 \times 0.0334 \times 9.34 \times 100$) in share prices attributed to variation in the CPU. Nevertheless, the estimated relationships between fluctuations in the CPU index and stock market variables vary considerably in relation to industry-level carbon intensity. Tables 4.A1 and 4.A2 in the Appendix compute the implied changes in volatility and share prices from 2000 to 2018 across different industries to explore heterogeneity across firms more in detail.

Table 4.3: Effects of climate policy uncertainty on historical stock price volatility (30-day horizon) and share prices in the US (monthly series).

	Volatility (30)			Share Price		
	(1)	(2)	(3)	(1)	(2)	(3)
CPU x CO2 intensity	0.00571*** (0.00159)	0.00565*** (0.00159)	0.00509*** (0.00141)	-0.0124** (0.00484)	-0.0121** (0.00483)	-0.00998*** (0.00352)
VIX x Contract Intensity	0.124 (0.0986)			-0.511* (0.269)		
EPU x Contract Intensity		-0.0540 (0.137)	-0.0407 (0.137)		0.512 (0.385)	0.473 (0.384)
WTI x CO2 intensity			0.00338 (0.00280)			-0.00996 (0.00972)
Search Strategy	Original	Original	Original	Original	Original	Original
N	273367	273367	273367	956480	956480	956480
R-squared	0.642	0.642	0.642	0.815	0.815	0.815
Number of firms	3237	3237	3237	8775	8775	8775
Firm effects	Yes	Yes	Yes	Yes	Yes	Yes
Time effects	Month	Month	Month	Month	Month	Month
First Year	1996	1996	1996	1990	1990	1990
Last Year	2018	2018	2018	2018	2018	2018

Notes: Variables are averaged across months and expressed as natural logs. Standard errors in parentheses are based on clustering at the firm level. Sample is restricted to companies with consecutive observations over the whole period of time where they have been listed. CO₂ emissions are retrieved from the EPA Greenhouse Gas Reporting Program (GHGRP). Intensity measures are computed following the procedure described in Section 4.3.1.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Next, Table 4.5 examines the extent to which variations in climate policy uncertainty affect firm-level expenses in research and development and employment. As before, our specifications differentiate firms by their relative exposure to climate policy changes. Table 4.5 relies on the quarterly series, as these variables are not available at the monthly level on Compustat.

Table 4.4: Effects of climate policy uncertainty on historical stock price volatility (30-day horizon) and share prices in the US (quarterly series).

	Volatility (30)			Share Price		
	(1)	(2)	(3)	(1)	(2)	(3)
CPU x CO2 intensity	0.00892*** (0.00252)	0.00879*** (0.00252)	0.00816*** (0.00233)	-0.0391*** (0.00854)	-0.0385*** (0.00853)	-0.0334*** (0.00640)
VIX x Contract Intensity	0.0942 (0.0995)			-0.593** (0.269)		
EPU x Contract Intensity		-0.116 (0.148)	-0.105 (0.149)		0.585 (0.459)	0.517 (0.459)
WTI x CO2 intensity			0.00237 (0.00274)			-0.0134 (0.00948)
Search Strategy	Original	Original	Original	Original	Original	Original
N	97863	97863	97863	440903	440903	440903
R-squared	0.689	0.689	0.689	0.786	0.786	0.787
Number of firms	3374	3374	3374	11033	11033	11033
Firm effects	Yes	Yes	Yes	Yes	Yes	Yes
Time effects	Quarter	Quarter	Quarter	Quarter	Quarter	Quarter
First Year	1996	1996	1996	1990	1990	1990
Last Year	2018	2018	2018	2018	2018	2018

Notes: Variables are averaged across quarters and expressed as natural logs. Standard errors in parentheses are based on clustering at the firm level. Sample is restricted to companies with consecutive observations over the whole period of time where they have been listed. CO₂ emissions are retrieved from the EPA Greenhouse Gas Reporting Program (GHGRP). Intensity measures are computed following the procedure described in Section 4.3.1.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4.5: Effects of climate policy uncertainty on R&D expenses and employment in the US. 1990 - 2018 (quarterly series).

	R&D			Employment		
	(1)	(2)	(3)	(1)	(2)	(3)
CPU x CO2 intensity	-0.0296** (0.0128)	-0.0295** (0.0128)	-0.0312*** (0.0109)	-0.0627*** (0.00809)	-0.0630*** (0.00810)	-0.0364*** (0.00560)
VIX x Contract Intensity	-0.416 (0.457)			-0.382 (0.244)		
EPU x Contract Intensity	0.0430	0.0596 (0.831)	-0.330 (0.828)	-0.482*	0.585 (0.288)	0.517 (0.288)
WTI x CO2 intensity			0.00577 (0.0147)			-0.0384*** (0.00624)
Search Strategy	Original	Original	Original	Original	Original	Original
N	94915	94915	94915	79465	79465	79465
R-squared	0.889	0.889	0.889	0.943	0.943	0.943
Number of firms	3038	3038	3038	8273	8273	8273
Firm effects	Yes	Yes	Yes	Yes	Yes	Yes
Time effects	Quarter	Quarter	Quarter	Year	Year	Year
First Year	1990	1990	1990	1990	1990	1990
Last Year	2018	2018	2018	2018	2018	2018

Notes: Variables are averaged across quarters or years and expressed as natural logs. Standard errors in parentheses are based on clustering at the firm level. Sample is restricted to companies with consecutive observations over the whole period of time where they have been listed. CO₂ emissions are retrieved from the EPA Greenhouse Gas Reporting Program (GHGRP). Intensity measures are computed following the procedure described in Section 4.3.1.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

One of the main challenges to achieving a successful transition towards a low-carbon economy is to create incentives to trigger firms' investment in low-carbon technologies. As anticipated, policy uncertainty introduces an element of risk for private companies which may affect their investment behavior. Sustained additional risk due to uncertainty about climate policy developments may therefore constitute a potential barrier to delivering the necessary low-carbon investments. At the same time, the direction of uncertainty is especially important in this context, as analyzed in Section 4.4.4. In Table 4.5, we focus on the aggregate effect of a change in the CPU index.

First, we consider the impact of CPU on research and development. As with share prices, we find a moderately large and statistically significant negative coefficient on climate policy uncertainty changes for R&D expenses, particularly for firms with high exposure to climate policy changes. These results are in line with predictions from the real options theory. High levels of uncertainty may depress firm-level investment by prompting preventive delays due to investment irreversibility (Dixit, 1989; Pindyck, 1988; Bloom et al., 2007), which is an especially important source of concern in the case of R&D investments (Dixit et al., 1994). Consider again the climate policy uncertainty changes from 2000 to 2018. Assuming a median exposure, the implied quarterly decreases in R&D expenses amount to almost 12%. In other words, in the absence of climate policy uncertainty, our estimates predict that research and development efforts since 2000 may have been greater by as much as one-tenth. Even in this case, the implied changes in R&D investments vary substantially across industries, ranging from an estimated modest decrease of around 2% for firms in Life Insurance (SIC 6311) to reductions of more than 16% for those operating in Fabricated Metal Products (SIC 3490).

Finally, we explore the relationship between climate policy uncertainty and employment. These analyses rely on yearly data, as company-level employment data are available only at the annual level on Compustat. Our coefficients in Table 4.5 suggest that uncertainty about climate policy is associated with negative effects on annual

employment levels, particularly for firms in high-emitting sectors. Working again with the changes in climate policy uncertainty from 2000 to 2018, we estimate that for a firm with median exposure, the implied changes in annual employment is around 13%. The implied effects at the firm level are relatively moderate if we consider that more than 80% of the companies in our estimation sample employ less than 10 workers (see Figure 4.A5 in the Appendix). Hence, the estimated relationship between CPU and aggregate employment levels is expected to be modest. Tables 4.A3 and 4.A4 in the Appendix further investigate the relationship of climate policy uncertainty changes to the cross-sectional structure of R&D investment rates and employment levels across different industries.

4.4.2 Robustness tests

This section presents a number of additional results for robustness purposes. Our main robustness tests, as presented in what follows, include (1) the use of an alternative version of our index (N-CPU) introduced in Section 4.2.3; (2) a different estimation window that ensures comparability across all our outcome variables; (3) a number of other industry-level policy exposure measures. Table 4.7 assesses the sensitivity of our results to an alternative version of our CPU index computed with a search strategy restricted to climate policy keywords (see Appendix 4.A). By doing so, we investigate whether the differences in topical scope between the original and the restricted version of the index alter our estimations to a considerable degree. The key rationale is to verify whether our estimated relationships might be driven by uncertainty about policy developments targeting other environmental concerns, such as local air pollution, rather than climate regulation. Results in Table 4.6 are all comparable in terms of size and significance to those presented in Tables 4.4 - 4.5, suggesting that our estimations are fundamentally driven by uncertainty related to policies addressing climate change.

Next, Table 4.7 explores whether our results change when running our estimations

Table 4.6: Effects of N-CPU on volatility, share prices, RD expenses and employment in the US. 1990 - 2018 (quarterly series).

	(3)	(3)	(3)	(3)
	Volatility (30)	Share Price	R&D	Employment
CPU x Industry CO2 intensity	0.00860*** (0.00237)	-0.0352*** (0.00654)	-0.0312*** (0.0112)	-0.0367*** (0.00548)
EPU x Contract Intensity	-0.104 (0.149)	0.517 (0.459)	0.0613 (0.828)	-0.482* (0.288)
WTI x Industry CO2 intensity	0.00218 (0.00273)	-0.0123 (0.00940)	0.00645 (0.0146)	-0.0374*** (0.00620)
Search Strategy	Restricted	Restricted	Restricted	Restricted
N	97863	441044	94915	79512
R-squared	0.689	0.786	0.889	0.943
Number of firms	3374	11033	3038	8276
Firm effects	Yes	Yes	Yes	Yes
Time effects	Quarter	Quarter	Quarter	Year
First Year	1996	1990	1990	1990
Last Year	2018	2018	2018	2018

Notes: Variables are averaged across quarters and expressed as natural logs. Standard errors in parentheses are based on clustering at the firm level. Sample is restricted to companies with consecutive observations over the whole period of time where they have been listed. CO₂ emissions are retrieved from the EPA Greenhouse Gas Reporting Program (GHGRP). Intensity measures are computed following the procedure described in Section 4.3.1.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4.7: Effects of climate policy uncertainty on volatility, share prices, RD expenses and employment in the US. 1996-2018.

	(3)	(3)	(3)	(3)
	Volatility (30)	Share Price	R&D	Employment
CPU x Industry CO2 intensity	0.00816*** (0.00233)	-0.0354*** (0.00667)	-0.0346*** (0.0113)	-0.0371*** (0.00574)
EPU x Contract Intensity	-0.105 (0.149)	0.652 (0.504)	-0.407 (0.917)	-0.443 (0.278)
WTI x Industry CO2 intensity	0.00237 (0.00274)	-0.0125 (0.00918)	0.00509 (0.0134)	-0.0329*** (0.00596)
Search Strategy	Original	Original	Original	Original
N	97863	384403	82217	66711
R-squared	0.689	0.796	0.894	0.951
Number of firms	3374	10442	2865	7646
Firm effects	Yes	Yes	Yes	Yes
Time effects	Quarter	Quarter	Quarter	Year
First Year	1996	1996	1996	1996
Last Year	2018	2018	2018	2018

Notes: Variables are averaged across quarters and expressed as natural logs. Standard errors in parentheses are based on clustering at the firm level. Sample is restricted to companies with consecutive observations over the whole period of time where they have been listed. CO₂ emissions are retrieved from the EPA Greenhouse Gas Reporting Program (GHGRP). Intensity measures are computed following the procedure described in Section 4.3.1.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4.8: Effects of climate policy uncertainty on volatility, share prices, RD expenses and employment in the US. Alternative exposure measure (1).

	(2)	(2)	(2)	(2)
	Volatility (30)	Share Price	R&D	Employment
CPU x Industry CO2 intensity	0.00671*** (0.00180)	-0.0318*** (0.00694)	-0.0323*** (0.0105)	-0.0590*** (0.00720)
EPU x Industry CO2 intensity	-0.00494** (0.00214)	-0.0139 (0.00992)	0.0111 (0.0123)	0.0242*** (0.00724)
Search Strategy	Original	Original	Original	Original
N	123871	520061	95708	87534
R-squared	0.711	0.794	0.889	0.941
Number of firms	4366	13341	3079	9155
Firm effects	Yes	Yes	Yes	Yes
Time effects	Quarter	Quarter	Quarter	Year
First Year	1996	1990	1990	1990
Last Year	2018	2018	2018	2018

Notes: Variables are averaged across quarters and expressed as natural logs. Standard errors in parentheses are based on clustering at the firm level. Sample is restricted to companies with consecutive observations over the whole period of time where they have been listed. CO₂ emissions are retrieved from the EPA Greenhouse Gas Reporting Program (GHGRP). Intensity measures are computed following the procedure described in Section 4.3.1.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

starting from 1996. Our volatility measure retrieved from Options Metrics is only available from 1996. Therefore, in order to test to what extent differences in terms of significance and magnitude between our estimates could be potentially explained by distinct estimation windows, we additionally estimate our main specifications using a common time frame, i.e. 1996 to 2018. Even in this case, we yield comparable results to Tables 4.4 - 4.5. Finally, we consider additional industry-level economic policy exposure measures that we interact with the EPU index to investigate whether different approaches to measuring exposure to government policy risks affect our results.

First, we interact the EPU index with the measure of exposure to climate policy risk described in Section 4.3.1. Results are presented in Table 4.8. The aim is to provide a direct comparison between our CPU and the EPU index. Overall, the estimated relationships between climate policy uncertainty and economic outcomes are still comparable to our main specifications. The coefficient on EPU, however, bears limited direct economic interpretation. Nevertheless, we see these results

as additional supporting evidence that our index can be a significant predictor of firm-level economic outcomes even after controlling for the impacts of economic policy uncertainty.

Second, we compute Herfindahl-Hirschman (HHI) concentration indices using Compustat information on sales and industry definitions. Within every SIC 4-digit industry, we sum up the squared ratios of firm sales to the total industry sales in the year prior to our estimation period. Then, we assign the estimated pre-sample industry-level HHI to each firm and interact it with the EPU index. Companies may exhibit different responses to changes in economic policy uncertainty depending on the amount of competition among them. On the one hand, firms in sectors where market power is more concentrated may be less sensitive to changes in EPU because they have more monopolistic positions. On the other hand, companies operating in more concentrated industries tend to be larger and more actively traded in the stock market, making them more exposed to regulatory risk changes. Results are presented in Table 4.9. Overall, both alternative measures and specifications yield significant results similar to the results displayed in Tables 4.4 - 4.5 under specification (2).

4.4.3 Historical analysis

As discussed in Section 4.2.3, a pivotal aspect to investigate in the context of climate policy developments is the direction of the uncertainty. Throughout the years, climate action has experienced many instances of acceleration and deceleration. Greater uncertainty may arise either from expectations of additional delays in climate action or anticipated greater stringency in future climate regulation. In the following section, we provide historical breakdowns for our main specifications to investigate whether the direction of our estimated relationships changes in accordance with the underlying drivers of climate policy uncertainty.

Table 4.10 reports our estimated coefficients when running specification (2) from Tables 4.4 - 4.5 across consecutive shorter time frames in our sample. In line with

Table 4.9: Effects of climate policy uncertainty on volatility, share prices, RD expenses and employment in the US. Alternative exposure measure (2).

	(2)	(2)	(2)	(2)
	Volatility (30)	Share Price	R&D	Employment
CPU x Industry CO2 intensity	0.00595** (0.00248)	-0.0373*** (0.00787)	-0.0331*** (0.0128)	-0.0571*** (0.00768)
EPU x HHI	0.0173 (0.0108)	-0.196*** (0.0305)	0.0144 (0.0390)	0.000169 (0.0192)
Search Strategy	Original	Original	Original	Original
N	108472	518800	95586	87434
R-squared	0.689	0.794	0.889	0.941
Number of firms	3706	13258	3074	9149
Firm effects	Yes	Yes	Yes	Yes
Time effects	Quarter	Quarter	Quarter	Year
First Year	1996	1990	1990	1990
Last Year	2018	2018	2018	2018

Notes: Variables are averaged across quarters and expressed as natural logs. Standard errors in parentheses are based on clustering at the firm level. Sample is restricted to companies with consecutive observations over the whole period of time where they have been listed. CO₂ emissions are retrieved from the EPA Greenhouse Gas Reporting Program (GHGRP). Intensity measures are computed following the procedure described in Section 4.3.1.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

our expectations, the direction of the estimated effects seems to change depending on the estimation window. For instance, although we primarily observe a significant and negative effect on share prices throughout the entire estimation period, this effect appears to turn positive between 2010 and 2014. This coincides with the 2010 withdrawal of the US climate change bill under the Obama administration as well as the President's retreat on stricter ozone standards over the following year. Such developments may have signaled additional delays in climate action at a national level, plausibly leading investors to revise their expectations of regulatory risk downwards. Intuitively, we would expect stock markets to make gains when new developments point to a setback in climate action unless investors are thinking in terms of long-term climate goals. Similarly, the estimated effect for volatility turns negative from 2015 onward in stark contrast to the trend that characterized the preceding decade. This occurred in conjunction with the election of President Trump, which represented a clear shift from the policy priorities and goals of the preceding administration's climate agenda. In addition, our results suggest that the change in the direction of environmental policy in the United States under the Trump

administration has been accompanied by significant reductions in R&D efforts.

Table 4.10: Historical Breakdowns (quarterly series).

Dependent Variable	1990 - 1994	1995 - 1999	2000 - 2004	2005 - 2009	2010 - 2014	2015 - 2018
volatility						
CPU x CO2 Intensity		0.00231 (0.00395)	-0.00458 (0.00565)	0.0170*** (0.00657)	0.0132** (0.00555)	-0.0146*** (0.00376)
Share Price						
CPU x CO2 Intensity	-0.00851* (0.00503)	-0.00723** (0.00332)	0.000502 (0.0116)	-0.0370** (0.0159)	0.0357*** (0.0135)	-0.00457 (0.00674)
R&D Expenses						
CPU x CO2 Intensity	-0.0218 (0.0134)	0.0103 (0.00888)	0.0581*** (0.0157)	-0.0271 (0.0325)	-0.0402* (0.0232)	-0.0462*** (0.0131)
Employment (Annual)						
CPU x CO2 Intensity	-0.0566*** (0.0122)	0.0146 (0.00964)	0.0352** (0.0138)	-0.0361** (0.0166)	0.00714 (0.0110)	-0.0217*** (0.00581)

Notes: Variables are averaged across quarters and expressed as natural logs. Standard errors in parentheses are based on clustering at the firm level. Sample is restricted to companies with consecutive observations over the whole period of time where they have been listed. CO₂ emissions are retrieved from the EPA Greenhouse Gas Reporting Program (GHGRP). Intensity measures are computed following the procedure described in Section 4.3.1.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4.4.4 Belief revision

A key question that arises at this point of the analysis is how economic outcomes respond to the different underlying drivers of uncertainty about climate policy. To this end, we turn to the sub-indices introduced in Section 4.2.3, namely CPU+ and CPU- (see Appendix 4.A). Making use of the sub-indices allows us to systematically disentangle the effects of climate policy uncertainty when belief revision goes towards more or less climate action. The estimated coefficients using both indices are presented in Table 4.11.

Overall, these results suggest that economic outcomes are more sensitive to uncertainty about climate policy when expectations point towards more stringent regulation. Comparing the results for the two sub-indices, the coefficients on CPU+ consistently exhibit larger coefficients, in absolute value. This holds particularly true in the case of share prices. Stock market reactions appear to be remarkably more sensitive to uncertainty related to potential policy developments increasing

climate ambition. This implies that the relationship between share prices and CPU estimated by our main specifications is likely to be primarily driven by belief revision towards more regulatory stringency. This also appears to be the case for the effect on employment. In fact, when regressing annual employment levels on each sub-index respectively, we only yield a significant coefficient for CPU+. The difference in the estimated effects is less marked for the effects on volatility and R&D expenses.

Table 4.11: Effects of climate policy uncertainty on volatility, share price, RD expenses and employment in the US. 1990 - 2018 (quarterly series). Comparing CPU- and CPU+.

	CPU-				CPU+			
	Volatility (30)	Share Price	R&D	Employment	Volatility (30)	Share Price	R&D	Employment
CPU x Industry CO2 intensity	0.00631*** (0.00137)	-0.0174*** (0.00267)	-0.00730* (0.00424)	-0.00554 (0.00338)	0.00850*** (0.00216)	-0.0439*** (0.00636)	-0.0174* (0.0104)	-0.0422*** (0.00552)
EPU x Contract Intensity	-0.124 (0.146)	0.548 (0.459)	0.0950 (0.836)	-0.394 (0.291)	-0.108 (0.149)	0.550 (0.457)	0.121 (0.827)	-0.426 (0.289)
WTI x Industry CO2 intensity	0.00244 (0.00275)	-0.0161* (0.00972)	0.00243 (0.0150)	-0.0441*** (0.00650)	0.00214 (0.00271)	-0.00979 (0.00934)	0.00423 (0.0144)	-0.0353*** (0.00612)
Search Strategy	Negative	Negative	Negative	Negative	Positive	Positive	Positive	Positive
N	97196	437888	94104	79509	97863	441044	94915	79512
R-squared	0.689	0.787	0.889	0.943	0.689	0.787	0.889	0.943
Clustered S.E.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of firms	3372	11032	3038	8276	3374	11033	3038	8276
Firm effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time effects	Quarter	Quarter	Quarter	Year	Quarter	Quarter	Quarter	Year
First Year	1996	1990	1990	1990	1996	1990	1990	1990
Last Year	2018	2018	2018	2018	2018	2018	2018	2018

Notes: Variables are averaged across quarters and expressed as natural logs. Standard errors in parentheses are based on clustering at the firm level. Sample is restricted to companies with consecutive observations over the whole period of time where they have been listed. CO₂ emissions are retrieved from the EPA Greenhouse Gas Reporting Program (GHGRP). Intensity measures are computed following the procedure described in Section 4.3.1.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4.5 Conclusion

Since private sector investments in low-carbon technologies are fundamentally dependent upon expectations over future climate policy stringency, an important barrier to private sector investment in such technologies may be policy uncertainty. Analyzing firms' and investors' responses to climate policy shocks is also crucial to devising the best possible approach to transition to a low-carbon economy, which may include a wide array of environmental, fiscal, innovation, and macroprudential policies. However, firms and investors may not only react to realized climate policy shocks, but also to changes in the probability of these shocks happening, which is what we define

as climate policy uncertainty. Since climate change became a policy issue in the early 1980s, domestic and international climate policy has attempted, through periods of progress and others of setbacks, to move the economy from a carbon-intensive to a low-carbon equilibrium.

To capture firms' and investors' responses to climate policy uncertainty, we develop a novel newspaper-based index capturing climate policy uncertainty in several major economies and then study its relationship with a set of key firm-level outcomes for the United States, covering publicly-listed firms from 1990 onward. We analyze outcomes such as share price volatility and share price, employment decisions, and investments in research and development. Our approach also accounts for the fact that uncertainty may sometimes reflect a slowdown in the transition to a cleaner economy, and sometimes to a breakthrough or acceleration. As a result, we developed two sub-indices, capturing either source of uncertainty.

Overall, we find that an increase in climate policy uncertainty is linked with larger historical stock price volatility as well as lower share prices. Similarly, climate policy uncertainty is negatively associated with R&D investments and annual employment. Our calculations indicate that climate policy uncertainty in the past two decades has led to an average increase in volatility of approximately 3%, a decrease of around 13% in share prices, and a reduction of about 12% in R&D investments. We also find some negative, albeit modest, impacts on employment. Notably, these effects vary significantly across industries, with more substantial repercussions observed in carbon-intensive sectors. The variation in R&D confirms previous research considering actual policy changes as a source of variation, suggesting that firms base their decisions on whether to innovate not only based on regulatory changes, but also on expectations thereof. The negative, but rather small changes in employment are also consistent with the existing literature, which points to relatively small changes in employment following climate policy tightening. In all our results, the source of the uncertainty matters, though. In periods in which climate policy was stalling, several outcomes reacted positively to higher uncertainty, as it might have pointed to further divisions

among legislators. Consistently, our sub-indices indicate stronger reactions to climate policy uncertainty when the latter is driven more by policy tightening than inaction.

Our results suggest that the recent increase in climate policy uncertainty has significantly slowed down investments in R&D in the most carbon-intensive sectors, which are major contributors to greenhouse gas emissions and local air pollutants. These findings offer micro-level evidence supporting the notion that stability in climate policies is crucial for facilitating the transition towards a low-carbon economy. Notably, deferred investments in low-carbon technologies will consequently result in higher carbon concentrations, further exacerbating climate change, and might significantly increase the cost of transitioning to a low-carbon economy irreversibly (Dorsey, 2019). Incorporating mechanisms within the initial policy design to restrict arbitrary adjustments may help mitigate uncertainty and minimize its adverse effects and reduce the overall transition costs (Annicchiarico et al., 2022).

4.A Keyword selection

The following subsections report the keyword selection for our free-text search strategies in Factiva in all languages.

4.A.1 English

Original Search Strategy: (energy or "the environment" or environmental* or "climate change" or "global warming" or climate not ("business climate" or "political climate" or "economic climate" or "regulatory climate" or "legal climate")) or carbon or emission* or "greenhouse gas" or GHG or "carbon dioxide" or CO2 or methane or CH4 or pollut* or "sulphur oxide" or "sulfur oxide" or SOx or "sulphur dioxide" or "sulfur dioxide" or SO2 or "nitrogen oxide" or NOx or "nitrogen dioxide" or NO2 or "particulate matter" or "fine particulates" or "fine particle" or "PM2.5" or "PM10" or ozone or renewable or hydro or "wind power" or "wind energy" or "wind farm" or "wind farms" or "wind turbine" or "wind turbines" or photovoltaic or PV or solar or biomass or "electric vehicle" or "electric vehicles" or "electric car" or "electric cars" or "hybrid vehicle" or "hybrid vehicles" or EV) same ((policy not "monetary policy") or policies or regulation* or legislation* or law or laws or fee or fees or tax or taxes or standard or standards or certificate* or subsidy or subsidies or pricing or ETS or feed-in-tariff* or "trading scheme" or "trading system" or "cap and trade" or "emissions trading" or label or "eco-label") and (unclear or vague or uncertain or uncertainty)

Restricted Search Strategy (N-CPU): (energy or "the environment" or environmental* or "climate change" or "global warming" or (climate not ("business climate" or "political climate" or "economic climate" or "regulatory climate" or "legal climate"))) or carbon or emission* or "greenhouse gas" or GHG or "carbon dioxide" or CO2 or methane or CH4 or renewable or hydro or "wind power" or

"wind energy" or "wind farm" or "wind farms" or "wind turbine" or "wind turbines"
or photovoltaic or PV or solar or biomass or "electric vehicle" or "electric vehicles"
or "electric car" or "electric cars" or "hybrid vehicle" or "hybrid vehicles" or EV
) same ((policy not "monetary policy") or policies or regulation* or legislation* or
law or laws or fee or fees or tax or taxes or standard or standards or certificate*
or subsidy or subsidies or pricing or ETS or feed-in-tariff* or "trading scheme" or
"trading system" or "cap and trade" or "emissions trading" or label or "eco-label")
and (unclear or vague or uncertain or uncertainty)

Search Strategy with additional keywords related to progress (**CPU+**): (energy
or "the environment" or environmental* or "climate change" or "global warming"
or climate not ("business climate" or "political climate" or "economic climate" or
"regulatory climate" or "legal climate") or carbon or emission* or "greenhouse gas"
or GHG or "carbon dioxide" or CO2 or methane or CH4 or pollut* or "sulphur
oxide" or "sulfur oxide" or SOx or "sulphur dioxide" or "sulfur dioxide" or SO2 or
"nitrogen oxide" or NOx or "nitrogen dioxide" or NO2 or "particulate matter" or
"fine particulates" or "fine particle" or "PM2.5" or "PM10" or ozone or renewable or
hydro or "wind power" or "wind energy" or "wind farm" or "wind farms" or "wind
turbine" or "wind turbines" or photovoltaic or PV or solar or biomass or "electric
vehicle" or "electric vehicles" or "electric car" or "electric cars" or "hybrid vehicle"
or "hybrid vehicles" or EV) same ((policy not "monetary policy") or policies or
regulation* or legislation* or law or laws or fee or fees or tax or taxes or standard or
standards or certificate* or subsidy or subsidies or pricing or ETS or feed-in-tariff*
or "trading scheme" or "trading system" or "cap and trade" or "emissions trading"
or label or "eco-label") and (unclear or vague or uncertain or uncertainty) **and**
(progress or implementation or adoption or consensus or action or success
or achievement)

Search Strategy with additional keywords related to failure (**CPU-**): (energy or "the
environment" or environmental* or "climate change" or "global warming" or climate
not ("business climate" or "political climate" or "economic climate" or "regulatory

climate” or ”legal climate”) or carbon or emission* or ”greenhouse gas” or GHG or ”carbon dioxide” or CO2 or methane or CH4 or pollut* or ”sulphur oxide” or ”sulfur oxide” or SOx or ”sulphur dioxide” or ”sulfur dioxide” or SO2 or ”nitrogen oxide” or NOx or ”nitrogen dioxide” or NO2 or ”particulate matter” or ”fine particulates” or ”fine particle” or ”PM2.5” or ”PM10” or ozone or renewable or hydro or ”wind power” or ”wind energy” or ”wind farm” or ”wind farms” or ”wind turbine” or ”wind turbines” or photovoltaic or PV or solar or biomass or ”electric vehicle” or ”electric vehicles” or ”electric car” or ”electric cars” or ”hybrid vehicle” or ”hybrid vehicles” or EV) same ((policy not “monetary policy”) or policies or regulation* or legislation* or law or laws or fee or fees or tax or taxes or standard or standards or certificate* or subsidy or subsidies or pricing or ETS or feed-in-tariff* or ”trading scheme” or ”trading system” or ”cap and trade” or ”emissions trading” or label or ”eco-label”) and (unclear or vague or uncertain or uncertainty) **and (slowdown or delay or disagreement or failure or rejection or postponement or setback)**

4.A.2 French

(”l’énergie” or énergétiqu* or environnementa* or écologique* or “changement climatique” or “réchauffement climatique” or climatique* or pollution or polluant* or carbone or ”gaz à effet de serre” or ”dioxyde de carbone” or CO2 or méthane or CH4 or ”oxyde de soufre” or SO2 or ”dioxyde de soufre” or SOx or ”oxyde d’azote” or NOx or ”dioxyde d’azote” or ”particules fines” or PM2,5 or PM10 or ozone or éolien* or (solaire* not ”système solaire”) or photovoltaïque* or hydraulique* or biomasse or ”énergies renouvelables” or ”énergie renouvelable” or ”voitures électriques” or ”voiture électrique” or ”voiture hybride” or ”voitures hybrides”) same ((politiqu* not ”politique monétaire”) or réglementation* or lois or loi or redevance* or tax* or impôt* or norme* or tarification* or ”tarif de rachat” or certificat* or subvention* or ETS or ”marché d’émissions” or ”droits à polluer” or ”système d’échanges” or ”SEQE”) and (incertitude* or incertain or incertaine or incertains or incertaines or

”peu clair” or ”pas clair”)

4.A.3 German

(Energiewende or ”Erneuerbare*Energien*Gesetz” or ”EEG-Einspeisevergütung” or ”EEG-Umlage” or Klimapolitik or Energiepolitik or Umweltpolitik or Luftreinhaltepolitik or Luftreinhalteplan or (”die Umwelt” or ökologisch or Klimawandel or Erderwärmung or ”globale Erwärmung” or ”Klimaerwärmung” or ”das Klima” or ”dem Klima” or ”des Klimas” or Klima* or ”die Umwelt” or ”der Umwelt” or Umwelt* or ”die Energie” or ”der Energie” or Energie* not (Geschäftsklima or ”politisches Klima” or ”wirtschaftliches Klima” or ”Wirtschaftsklima” or ”Regulierungsklima” or ”regulatorisches Klima” or ”Rechtsklima” or ”rechtliches Klima” or ”gesellschaftliches Klima” or ”Gesellschaftsklima”)) or Kohlenstoff* or Treibhausgas* or THG* or Kohlendioxid* or Kohlenstoffdioxid* or CO₂* or Methan* or CH₄* or Schadstoff* or Umweltverschmutzung* or Luftverschmutzung* verschmutz* or Schwefeloxid* or SO_x* or Schwefeldioxid* or SO₂* or Stickoxid* or NO_x* or Stickstoffdioxid* or NO₂* or Partikel* or Feinpartikel* or Feinstaub* or PM_{2,5} or PM₁₀* or Ozon* or erneuerbar* or Hydro* or Windenergie* or Windpark* or Windkraftanlage* or Photovoltaik* or PV or Solar* or Biomasse* or Elektrofahrzeug* or Elektroauto* or ”E-Auto*” or Hybridfahrzeug* or Hybridauto*) same ((Politik nicht Geldpolitik) or Richtlinie or Richtlinien or Reform or Reformen or Regulierung or Regulierungen or Vorschrift or Vorschriften or Gesetz or Gesetze or Gebühr or Gebühren or Abgabe or Abgaben or Maßnahme or Maßnahmen or Steuer or Steuern or Standard or Standards or Zertifikat or Zertifikate or Subvention or Subventionen or Preisgestaltung or Emissionshandel or ETS or Einspeisetarif or Einspeisetarife or Einspeisevergütung or Einspeisevergütungen or Handelssystem or Handelssysteme or ”Cap and Trade” or Emissionshandel or Label or Kennzeichen or ”Umweltzeichen” or ”Umweltabzeichen” or Umlage)) and (unklar or vage or unsicher or Unsicherheit)

4.A.4 Spanish

(”la energía” or *energético** or ”medio ambiente*” or *ecológico** or ”cambio climático” or ”calentamiento global” or *climatico* or *contaminación* or *contaminante** or *polución* or *carbono* or ”gases de efecto invernadero” or ”dióxido de carbono” or CO₂ or *metano* or CH₄ or ”óxido de azufre” or SO₂ or ”dióxido de azufre” or SO_x or ”óxido de nitrógeno” or NO_x or ”dióxido de nitrógeno” or ”partículas finas” or ”partículas en suspensión” or PM_{2.5} or PM₁₀ or *ozono* or *eólico** or ”tecnología* solar*” or ”panel* solar*” or ”placa* solar*” or ”central* solar*” or *fotovoltaico** or ”energía hidráulica” or *hidroeléctric** or *biomasa* or ”energías renovables” or ”energías verdes” or ”energías alternativas” or ”energías limpias” or ”renovables” or ”auto* eléctrico*” or ”coche* eléctrico*” or ”auto* híbrido*” or ”coche* híbrido*”) same ((*política** not ”política monetaria”) or *regulación** or *ley* or *leyes* or *impuesto** or *estándar** or ”tarifa de alimentación” or *certificado** or *subsidio** or ETS or ”mercado* de emisión*” or ”derecho* a contaminar” or ”sistema de comercio” or ”ETS”) and (*incertidumbre** or *incierto** or ”no es claro*” or “no está claro*” or ”no son claros*” or ”no están claros*”)

4.A.5 Italian

(*energia* or *energetic** or ”l’ambiente” or *ambiental** or *ecologic** or “riscaldamento globale” or *climatic** or *carbonio* or (*emissioni* not(”emissioni obbligatorie” or ”emissioni del Tesoro”)) or “gas a effetto serra” or “gas ad effetto serra” or “gas serra” or “anidride carbonica” or CO₂ or *metano* or CH₄ or *inquinamento** or *inquinante* or “ossid* di zolfo” or SO_x or “diossido di zolfo” or “biossido di zolfo” or “anidride solforosa” or “SO₂” or “ossido di azoto” or “monossido di azoto” or NO_x or “diossido di azoto” or “biossido di azoto” or NO₂ or “particelle fini” or “particolato atmosferico” or “particelle solide” or “particelle piccole” or “polveri sottili” or “particolato grossolano” or “particolato” or “materiale particolato” or

“PM10” or “PM2,5” or ozono or rinnovabil* or idroelettric* or idraulic* or eolic* or (solare not(“sistema solare” or “anno solare” or “eritema solare” or “ustione solare” or “trattamento solare”)) or fotovoltaic* or biomass* or “auto elettric*” or “vehicol* elettric*” or “auto ibrid*”) same ((politica not(“politica monetaria”)) or regolament* or regolamentazione or legislazione or legge or tasse or canone or standard not(“Standard Poor’s”) or certificat* or * certificazion* or sussidi or sussidio or sovvenzion* or ETS or “Sistema ES” or “feed in tariff*” or “conto energia” or “scambio di quote” or ”regime di scambio” or ”sistema di scambio ” or ”decarbonizzazione” or “effetto serra” or ”cap and trade” or “mercato dei diritti per l’emissione” or “etichett* ambiental*” or norma or norme or “marchio ambientale” or eco-etichett* or “etichett* ecologic*” or “eco-label” or normative or normativa) and (incerto or incerti or incertezza or incertezze) not (spread or bond)

4.B Short history of climate policy in the United States

1970. National Environmental Policy Act (NEPA) signed by President Nixon - "The Environmental Decade".

1980. Carter signed into law a bill that established Superfund.

1980. Congress appointed the National Academy of Sciences to carry out a comprehensive study on the impacts of rising CO₂ emissions.

1981. For the first time, a federal agency (EPA) declared that global warming was "not a theoretical problem but a threat whose effects will be felt within a few years", with potentially "catastrophic" consequences.

1988. The IPCC was established by the World Meteorological Organization (WMO) and the United Nations Environment Programme (UNEP).

1990. Finland is the first country in the world to introduce a carbon tax (followed the year after by Sweden and Norway).

1990. Amendments to the Clean Air Act of 1970: substantially increased the authority and responsibility of the federal government (i.e. introduction of an SO₂ cap-and-trade program).

1992. Bush opposed international efforts at the Earth Summit in Rio de Janeiro, Brazil - "new rules to limit carbon dioxide emissions would hurt economic growth".

1997. The US Senate voted unanimously under the Byrd-Hagel Resolution that the United States would not be ratifying the Kyoto Protocol.

2001. President Bush released his National Energy Policy ("NEP").

2003. The Clear Skies Act fails to become federal law of the United States.

2005. The European Union Emissions Trading System (EU ETS) was launched.

2009. President Barack Obama in his inaugural address called for the expanded use of renewable energy to meet the challenges of energy security and climate change.

2011. Obama Administration abandons plans for stricter ozone standards proposed

by the Environmental Protection Agency that would have significantly reduced emissions of smog-causing chemicals.

2015-2016. The United States became a signatory to the Paris Agreement.

2017. President-elected Donald Trump announced that the U.S. would cease all participation in the 2015 Paris Agreement.

4.C Estimated Outcome Changes

In order to investigate the relationship between climate policy uncertainty changes to the cross-sectional structure of stock market variables, R&D investment rates, and employment levels we compute the implied changes in our outcome variables from 2000 to 2018, relying on the estimation presented in Section 4.4.

Table 4.A1: Estimated changes in volatility associated with CPU changes from 2000 to 2018 for firms in selected industries.

Outcome Measure And Industry	(1) Carbon Intensity (log)	(2) Δ CPU (log points)	(3) Coeff. on ln(CPU) *Intensity	Estimated Change (1x2x3) in %
volatility				
Mining				
Mining And Quarrying Of Nonmetallic Minerals	11.97	40.5	0.00816	4.0
Crude Petroleum Natural Gas	11.24	40.5	0.00816	3.7
Metal Mining	10.53	40.5	0.00816	3.5
Bituminous Coal Lignite Mining	10.07	40.5	0.00816	3.3
Manufacturing				
Miscellaneous Fabricated Metal Products	12.97	40.5	0.00816	4.3
Cement, Hydraulic	12.71	40.5	0.00816	4.2
Miscellaneous Products of Petroleum Coal	11.97	40.5	0.00816	4.0
Agricultural Chemicals	11.91	40.5	0.00816	3.9
Pulp Mills	10.1	40.5	0.00816	3.3
Broadwoven Fabric Mills, Cotton	8.54	40.5	0.00816	2.8
Beverages	6.42	40.5	0.00816	2.1
Canned, Fruits, Veg, Preserves, Jams Jellies	5.93	40.5	0.00816	2.0
Cookies Crackers	4.12	40.5	0.00816	1.4
Transport, Communications, Electric, Gas and Sanitary Service				
Electric Services	12.69	40.5	0.00816	4.2
Deep Sea Foreign Transportation of Freight	9.21	40.5	0.00816	3.0
Natural Gas Distribution	8.45	40.5	0.00816	2.8
Refuse Systems	8.4	40.5	0.00816	2.8
Finance, Insurance and Real Estate				
Oil Royalty Traders	12.92	40.5	0.00816	4.3
Real Estate Investment Trusts	9.07	40.5	0.00816	3.0
Miscellaneous Business Credit Institution	3.31	40.5	0.00816	1.1
Life Insurance	2.22	40.5	0.00816	0.7
Services				
Engineering Services	7.46	40.5	0.00816	2.5
Personal Services	4.31	40.5	0.00816	1.4

Notes: Variables are averaged across quarters and expressed as natural logs. Standard errors in parentheses are based on clustering at the firm level. Sample is restricted to companies with consecutive observations over the whole period of time where they have been listed. CO₂ emissions are retrieved from the EPA Greenhouse Gas Reporting Program (GHGRP). Intensity measures are computed following the procedure described in Section 4.3.1.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4.A2: Estimated changes in share prices associated with CPU changes from 2000 to 2018 for firms in selected industries.

Outcome Measure And Industry	(1) Carbon Intensity (log)	(2) Δ CPU (log points)	(3) Coeff. on ln(CPU) *Intensity	Estimated Change (1x2x3) in %
Share Prices				
<i>Mining</i>				
Mining And Quarrying Of Nonmetallic Minerals	11.97	40.5	0.0334	16.2
Crude Petroleum Natural Gas	11.24	40.5	0.0334	15.2
Metal Mining	10.53	40.5	0.0334	14.2
Bituminous Coal Lignite Mining	10.07	40.5	0.0334	13.6
<i>Manufacturing</i>				
Miscellaneous Fabricated Metal Products	12.97	40.5	0.0334	17.5
Cement, Hydraulic	12.71	40.5	0.0334	17.2
Miscellaneous Products of Petroleum Coal	11.97	40.5	0.0334	16.2
Agricultural Chemicals	11.91	40.5	0.0334	16.1
Pulp Mills	10.1	40.5	0.0334	13.7
Broadwoven Fabric Mills, Cotton	8.54	40.5	0.0334	11.6
Beverages	6.42	40.5	0.0334	8.7
Canned, Fruits, Veg, Preserves, Jams Jellies	5.93	40.5	0.0334	8.0
Cookies Crackers	4.12	40.5	0.0334	5.6
<i>Transport, Communications, Electric, Gas and Sanitary Service</i>				
Electric Services	12.69	40.5	0.0334	17.2
Deep Sea Foreign Transportation of Freight	9.21	40.5	0.0334	12.5
Natural Gas Distribution	8.45	40.5	0.0334	11.4
Refuse Systems	8.4	40.5	0.0334	11.4
<i>Finance, Insurance and Real Estate</i>				
Oil Royalty Traders	12.92	40.5	0.0334	17.5
Real Estate Investment Trusts	9.07	40.5	0.0334	12.3
Miscellaneous Business Credit Institution	3.31	40.5	0.0334	4.5
Life Insurance	2.22	40.5	0.0334	3.0
<i>Services</i>				
Engineering Services	7.46	40.5	0.0334	10.1
Personal Services	4.31	40.5	0.0334	5.8

Notes: Variables are averaged across quarters and expressed as natural logs. Standard errors in parentheses are based on clustering at the firm level. Sample is restricted to companies with consecutive observations over the whole period of time where they have been listed. CO₂ emissions are retrieved from the EPA Greenhouse Gas Reporting Program (GHGRP). Intensity measures are computed following the procedure described in Section 4.3.1.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4.A3: Estimated changes in RD expenses associated with CPU changes from 2000 to 2018 for firms in selected industries.

Outcome Measure And Industry	(1) Carbon Intensity (log)	(2) Δ CPU (log points)	(3) Coeff. on ln(CPU) *Intensity	Estimated Change (1x2x3) in %
R&D expenses				
<i>Mining</i>				
Mining And Quarrying Of Nonmetallic Minerals	11.97	40.5	0.0312	15.1
Crude Petroleum Natural Gas	11.24	40.5	0.0312	14.2
Metal Mining	10.53	40.5	0.0312	13.3
Bituminous Coal Lignite Mining	10.07	40.5	0.0312	12.7
<i>Manufacturing</i>				
Miscellaneous Fabricated Metal Products	12.97	40.5	0.0312	16.4
Cement, Hydraulic	12.71	40.5	0.0312	16.1
Miscellaneous Products of Petroleum Coal	11.97	40.5	0.0312	15.1
Agricultural Chemicals	11.91	40.5	0.0312	15.0
Pulp Mills	10.1	40.5	0.0312	12.8
Broadwoven Fabric Mills, Cotton	8.54	40.5	0.0312	10.8
Beverages	6.42	40.5	0.0312	8.1
Canned, Fruits, Veg, Preserves, Jams Jellies	5.93	40.5	0.0312	7.5
Cookies Crackers	4.12	40.5	0.0312	5.2
<i>Transport, Communications, Electric, Gas and Sanitary Service</i>				
Electric Services	12.69	40.5	0.0312	16.0
Deep Sea Foreign Transportation of Freight	9.21	40.5	0.0312	11.6
Natural Gas Distribution	8.45	40.5	0.0312	10.7
Refuse Systems	8.4	40.5	0.0312	10.6
<i>Finance, Insurance and Real Estate</i>				
Oil Royalty Traders	12.92	40.5	0.0312	16.3
Real Estate Investment Trusts	9.07	40.5	0.0312	11.5
Miscellaneous Business Credit Institution	3.31	40.5	0.0312	4.2
Life Insurance	2.22	40.5	0.0312	2.8
<i>Services</i>				
Engineering Services	7.46	40.5	0.0312	9.4
Personal Services	4.31	40.5	0.0312	5.4

Notes: Variables are averaged across quarters and expressed as natural logs. Standard errors in parentheses are based on clustering at the firm level. Sample is restricted to companies with consecutive observations over the whole period of time where they have been listed. CO₂ emissions are retrieved from the EPA Greenhouse Gas Reporting Program (GHGRP). Intensity measures are computed following the procedure described in Section 4.3.1.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4.A4: Estimated changes in employment associated with CPU changes from 2000 to 2018 for firms in selected industries.

Outcome Measure And Industry	(1) Carbon Intensity (log)	(2) Δ CPU (log points)	(3) Coeff. on Log(CPU) *Intensity	Estimated Change (1x2x3) in %
Employment				
<i>Mining</i>				
Mining And Quarrying Of Nonmetallic Minerals	11.97	40.5	0.0364	17.6
Crude Petroleum Natural Gas	11.24	40.5	0.0364	16.6
Metal Mining	10.53	40.5	0.0364	15.5
Bituminous Coal Lignite Mining	10.07	40.5	0.0364	14.8
<i>Manufacturing</i>				
Miscellaneous Fabricated Metal Products	12.97	40.5	0.0364	19.1
Cement, Hydraulic	12.71	40.5	0.0364	18.7
Miscellaneous Products of Petroleum Coal	11.97	40.5	0.0364	17.6
Agricultural Chemicals	11.91	40.5	0.0364	17.6
Pulp Mills	10.1	40.5	0.0364	14.9
Broadwoven Fabric Mills, Cotton	8.54	40.5	0.0364	12.6
Beverages	6.42	40.5	0.0364	9.5
Canned, Fruits, Veg, Preserves, Jams Jellies	5.93	40.5	0.0364	8.7
Cookies Crackers	4.12	40.5	0.0364	6.1
<i>Transport, Communications, Electric, Gas and Sanitary Service</i>				
Electric Services	12.69	40.5	0.0364	18.7
Deep Sea Foreign Transportation of Freight	9.21	40.5	0.0364	13.6
Natural Gas Distribution	8.45	40.5	0.0364	12.5
Refuse Systems	8.4	40.5	0.0364	12.4
<i>Finance, Insurance and Real Estate</i>				
Oil Royalty Traders	12.92	40.5	0.0364	19.0
Real Estate Investment Trusts	9.07	40.5	0.0364	13.4
Miscellaneous Business Credit Institution	3.31	40.5	0.0364	4.9
Life Insurance	2.22	40.5	0.0364	3.3
<i>Services</i>				
Engineering Services	7.46	40.5	0.0364	11.0
Personal Services	4.31	40.5	0.0364	6.4


Notes: Variables are averaged across quarters and expressed as natural logs. Standard errors in parentheses are based on clustering at the firm level. Sample is restricted to companies with consecutive observations over the whole period of time where they have been listed. CO₂ emissions are retrieved from the EPA Greenhouse Gas Reporting Program (GHGRP). Intensity measures are computed following the procedure described in Section 4.3.1.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4.D CPU index and other relevant measures


The following section compares the evolution of the CPU index to other relevant uncertainty measures, such as the EPU from Baker et al. (2016), the Chicago Board Options Exchange's CBOE Volatility Index, and oil price volatility as proxied by changes in West Texas Intermediate and Brent Crude.

Figure 4.A1: Comparing the evolution of the CPU to the EPU index developed by Baker et al. (2016)




Note: Based on quarterly series from 1990 to 2018.
Source: Authors' own calculations based on newspaper articles from Factiva.

Figure 4.A2: Comparing the evolution of the CPU to the VIX index




Note: Based on quarterly series from 1990 to 2018.
Source: Authors' own calculations based on newspaper articles from Factiva.

Figure 4.A3: Comparing the evolution of the CPU to WTI Crude



Note: Based on quarterly series from 1990 to 2018.
Source: Authors' own calculations based on newspaper articles from Factiva.

Figure 4.A4: Comparing the evolution of the CPU to Brent Crude



Note: Based on quarterly series from 1990 to 2018.
Source: Authors' own calculations based on newspaper articles from Factiva.

4.E Yearly time series

Table 4.A5: Effects of climate policy uncertainty on volatility, share prices, and RD expenses in the US (yearly series).

	(2)	(2)	(2)
	Volatility (30)	Share Price	R&D
CPU x CO2 intensity	0.0128*** (0.00428)	-0.0485*** (0.00944)	-0.0346*** (0.0119)
Search Strategy	Original	Original	Original
N	28367	130194	43151
R-squared	0.730	0.779	0.898
Number of firms	3358	11640	4587
Firm effects	Yes	Yes	Yes
Time effects	Year	Year	Year
First Year	1996	1990	1990
Last Year	2018	2018	2018


Notes: Variables are averaged across quarters and expressed as natural logs. Standard errors in parentheses are based on clustering at the firm level. Sample is restricted to companies with consecutive observations over the whole period of time where they have been listed. CO₂ emissions are retrieved from the EPA Greenhouse Gas Reporting Program (GHGRP). Intensity measures are computed following the procedure described in Section 4.3.1.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4.F Additional Descriptive Evidence

4.F.1 Employment levels in the estimation sample

Figure 4.A5: Frequency distribution of employment levels in the estimation sample



Chapter 5

Pollution reduction benefits across space: Quasi-experimental evidence from England

SUMMARY. This study provides novel quasi-experimental evidence on the effects of air pollutants on defensive expenditures and economic productivity to retrieve spatially resolved estimates of the willingness to pay for air quality improvements. To address endogeneity concerns, atmospheric temperature inversions are exploited as a source of quasi-random variation in the spatial concentration of $\text{PM}_{2.5}$. Using administrative data from England, I find that a plausibly exogenous $1 \mu\text{g}/\text{m}^3$ $\text{PM}_{2.5}$ shock significantly affects pharmaceutical expenditures and GVA per capita, partly through increased work absenteeism. Leveraging a counterfactual reduction of $1 \mu\text{g}/\text{m}^3$ of $\text{PM}_{2.5}$, I show that health benefits are more pronounced among the elderly and progressively distributed across income levels, while productivity gains are regressive and concentrated in urban areas. These findings imply that incorporating the spatial heterogeneity of pollution-reduction benefits into policy design could enhance the efficiency of environmental regulations and contribute to tackling health inequalities linked to pollution exposure.

5.1 Introduction

There is growing evidence that shows that modest concentrations of air pollutants affect various economic outcomes, including human health¹ and productivity², through changes in respiratory, cardiovascular, and cognitive functions, raising questions about the efficiency of current ambient pollution standards. From a social welfare perspective, optimal air pollution regulation requires information on the extent to which individuals value the control of air pollutants, or - in other words - their willingness to pay (WTP) for air quality improvements (Greenstone and Jack, 2013). Defensive behaviors offer one viable channel for estimating part of the demand for air quality improvements, as compensatory adaptation has an opportunity cost (e.g., Deschenes et al., 2017; Ito and Zhang, 2020). Nevertheless, an accurate measurement of the WTP requires estimating both *opportunity* and *direct* costs associated with air pollution (Becker, 1965; Grossman, 1972).

Yet, empirical estimates of the WTP for clean air are still scarce, primarily due to the paucity of exogenous shocks in air quality for empirical applications. Furthermore, the scarcity of suitable empirical settings typically limits the scope to carry out heterogeneity analyses to account for how the WTP varies across space as this would ideally require exploiting extended variation in air pollution for a broad representative sample of the population. As a result, the lack of context-specific welfare estimates prevents opportunities to enhance public welfare through the design of more efficient air pollution regulation and more precise benefit-cost analyses that can accurately account for the distribution of air pollution costs across the population (Muller and Mendelsohn, 2009).

This paper estimates the causal effects of PM_{2.5} concentrations on nationwide (i) *defensive expenditures*, as measured by expenditures on pharmaceuticals, and (ii)

¹See Pope III and Dockery, 2006; Brook et al., 2010; Deryugina et al., 2019; Wu et al., 2020.

²These include Graff Zivin and Neidell, 2012; Hanna and Oliva, 2015; Dechezleprêtre et al., 2019; He et al., 2019; Fu et al., 2021; Sarmiento, 2022.

economic productivity, proxied by local gross value added (GVA) per capita. To this end, I compile a novel dataset for England that merges granular vertical temperature profiles from the European Centre for Medium-Range Weather Forecasts (ECMWF) with high-resolution gridded pollution maps, public health care records, and data on economic activities at the district level. England’s centralized, publicly-funded National Health System (NHS) provides access to healthcare records for its 54 million automatically registered citizens, offering an ideal setting for this analysis. I find that a plausibly exogenous $1 \mu\text{g}/\text{m}^3$ pollution shock causes significant increases in pharmaceutical expenditures as well as a reduction in GVA per capita over the same year. Specifically, I estimate that a 1 microgram per cubic meter ($\mu\text{g}/\text{m}^3$) increase in the annual average concentration of $\text{PM}_{2.5}$ leads to an increment in expenditures on pharmaceuticals of 32.7% (around £1.2 billion or £22 per capita annually) and a reduction in gross value added (GVA) per capita of 1.6% (around £13 billion or £425 per capita annually).

To circumvent concerns of endogeneity due to residential sorting (e.g., Chay and Greenstone, 2005; Lee and Lin, 2018; Hebllich et al., 2021), atmospheric temperature inversions occurring at different pressure levels are exploited as a source of quasi-random dynamic variation in the spatial concentration of pollutants across England.³ Specifically, my empirical results are estimated by leveraging inversions as an instrumental variable (IV) in a two-stage least squares (2SLS) framework. In my baseline regressions, the instrument reflects the annual frequency of thermal inversion detected within a given grid, defined as a positive upward temperature gradient between the two pressure levels closest to the surface (defined as 1000 hPa and 950 hPa). The higher level of granularity of my data compared to previous studies allows exploiting within-district variation in thermal inversion exposure and detect inversions on a high periodicity (i.e., 6 hours) to address concerns that harnessing low-frequency events as instruments may lead to inflated estimates due to

³See Hicks et al. (2016); Arceo et al. (2016); Chen et al. (2017); Jans et al. (2018); Dechezleprêtre et al. (2019); Sager (2019); Molina (2021); Cui et al. (2023) for examples of similar approaches.

low statistical power (Bagilet and Zabrocki-Hallak, 2022).⁴

On aggregate, my estimates are approximately 9 times higher than the current damage cost per change in concentration used by the Department for Environment Food and Rural Affairs (DEFRA) to conduct cost-benefit analyses of pollution control policies in the UK. For comparison, the assumed damage cost of a 1 $\mu\text{g}/\text{m}^3$ increase in $\text{PM}_{2.5}$ equal £50.12 per capita annually (UK-AIR, 2023).⁵ The sizable gap relative to my quasi-experimental estimates can be primarily ascribed to two crucial differences. First, while current nationwide official cost estimates focus exclusively on health costs induced by cardiovascular and respiratory conditions, which have been traditionally associated with air pollution exposure (e.g., Ward 2015; Deschenes et al. 2017; Deryugina et al. 2019), this paper advances from the existing pollution-health literature by building on recent epidemiological evidence and additionally accounting for the effects of pollution on the nervous system (e.g., Wang et al., 2017; Zhang et al., 2018; Peeples, 2020; Aguilar-Gomez et al., 2022; Cook et al., 2023; Krebs and Luechinger, 2024). Second, despite a growing body of country-level evidence regarding the adverse economic consequences of pollution (e.g., Dechezleprêtre et al., 2019), the current damage cost per change in concentration in the UK have not yet integrated the direct economic costs of pollution linked to non-health impacts. This omission could significantly underestimate the overall societal benefits of such policies (e.g., Hunt et al. 2016; Leroutier and Ollivier 2022; Borgschulte et al. 2022). In contrast, I further quantify productivity losses induced by air pollution concentration to provide more comprehensive damage estimates and inform the calibration of environmental regulations that can help optimize both public health and economic growth.

The IV estimates are robust to several robustness checks, including a battery of different instrument definitions based on different pressure levels and nocturnal

⁴In contrast to existing studies that typically rely on NASA’s MERRA-2 database, which offers data at a coarser 60km x 60km resolution, the vertical temperature profiles used in this study are accessible at a much finer 10km x 10km resolution which substantially reduces measurement error.

⁵This value refers to the *Central* estimate. The upper-end annual value (the *High* case) of £156.52 per capita is still almost 3 times lower compared to my quasi-experimental estimations.

inversions, alternative proxies of air pollution concentrations, the inclusion of flexible linear and quadratic weather controls, absorbing detailed heterogeneous local trends, controlling for other ambient pollutants, and different clustering choices to account for spatial autocorrelation. Moreover, they are significantly larger (by approximately fourfold) than my ordinary least squares (OLS) estimates, indicating the potential for sizable bias in observational studies of air pollution exposure that do not account for its endogenous nature. Finally, I estimate a reduced form (RF) placebo specification that includes leads and lags of my instrument to rule out any anticipatory behavior.

I additionally provide complementary evidence of a causal nexus between plausible quasi-random pollution shocks and the number of sick leave certificates issued related to pollution-exposure conditions to shed light on the absenteeism channel that links health effects and foregone productivity (e.g., Holub et al., 2020). This finding carries distributional implications as a reduced capacity to work due to pollution-driven morbidity may trap low-income individuals in an illness-poverty cycle (cf. Ketcham et al. 2019), due to their often disproportionate exposure (e.g., Colmer et al. 2020). Reducing air pollution may thus additionally act as a channel to alleviate income inequalities that could persist through succeeding generations (e.g., Isen et al. 2017; Chetty and Hendren 2018).

In the second part of the empirical analysis, I leverage my quasi-experimental framework to simulate counterfactual reductions of $1 \mu\text{g}/\text{m}^3$ in $\text{PM}_{2.5}$ and investigate the spatial heterogeneity of health and productivity benefits across demographics and socio-economic groups. My findings suggest that health benefits tend to be larger for the elderly population and progressively distributed along the income distribution. This gap is consistent with previous literature pointing to existing inequalities along the income distribution in the adaptive capacity to environmental damages (see Drupp et al. 2021 for a recent review). Furthermore, By contrast, productivity gains tend to be regressive and concentrated in large urban areas, with the largest effects found in high-GVA districts of the capital. Precisely, my spatially-resolved estimates of the implied WTP range from around £170 to more than £1500 per

capita, in comparison to a mean WTP of around £445 per capita that approaches the conservative side of the distribution of the estimates. This extensive variability demonstrates the potential for substantial bias in cost-benefit analyses based on WTP estimates that do not account for spatial heterogeneity.

This paper's main contributions to the literature are twofold. First, I contribute to the quasi-experimental literature that evaluates the costs of air pollution. The existing literature has focused primarily on health costs to individuals (e.g., Chay and Greenstone, 2003; Chen et al., 2013; Schlenker and Walker, 2016; Anderson, 2020; Barreca et al., 2021), but often focuses exclusively on localized effects of Low Emission Zones (e.g., Rohlf et al. 2020; Margaryan 2021), a specific demographic (e.g., Currie and Neidell 2005; Knittel et al. 2016; Klauber et al. 2021) or investigates the effects of pollution on a specific array of health conditions (e.g., Neidell 2004). Additionally, while the existing economics literature on the health costs of pollution primarily investigates impacts on *mortality* - which largely concentrate among the elderly population (Deryugina et al., 2019) - this study focuses on *morbidity* costs. This allows to provide a more representative read-out of both acute and chronic conditions induced by air pollution exposure.⁶ Although potential mortality effects call for regulatory attention, my results show that overlooking the share of economic costs attributable to morbidity impacts would lead to a severe underestimation of the total economic costs of pollution.

From this strand of the literature, the most closely related study to my analysis is Pimpin et al. (2018) which estimates the morbidity costs of air pollution to the NHS in England by relying on microsimulations from prescription data.⁷ Simulations, however, are prone to endogeneity concerns, potentially impeding the accurate

⁶Focusing on morbidity has at least two additional key advantages: (a) morbidity tends to be a more sensitive indicator of pollution effects due to the relative immediacy of the impact; and (b) indicators of morbidity effects - such as foregone economic productivity per capita - allow for direct quantification of economic benefits rather than relying on existing estimates of the value of statistical life (VSL).

⁷A key difference is that Pimpin et al. (2018) focus on a limited number of health conditions, including asthma, chronic obstructive pulmonary disease, coronary heart disease, stroke, type 2 diabetes, dementia, and lung cancer.

identification of causal effects and leading to a downward bias in the estimation. The authors estimate that a $1 \mu\text{g}/\text{m}^3$ reduction in population exposure to $\text{PM}_{2.5}$ and NO_2 would result in savings of approximately 100 million per year in NHS and social care costs. This is considerably lower than my quasi-experimental estimates of the health costs of $\text{PM}_{2.5}$ pollution even before accounting for productivity losses. My results suggest that neglecting the direct economic costs of pollution leads to a large underestimation of its total societal costs.

Within this literature, a growing number of studies has also examined how pollution affects productivity through its impacts on workers (e.g., Graff Zivin and Neidell, 2012; Chang et al., 2016; Meyer and Pagel, 2017; He et al., 2019; Chang et al., 2019; Sarmiento, 2022; Adhvaryu et al., 2022), labor supply (e.g., Hanna and Oliva 2015) or firm's productivity (Fu et al., 2021). However, most of these studies are based on specific settings (e.g., single or selected production sites operating in a single sector) raising questions on their external validity as empirical inputs to compute representative economy-wide estimates and conduct cost-benefit analyses to inform policymaking. An exception is Dechezleprêtre et al. (2019) who provides causal estimates on the effects of air pollution on economy-wide reductions in European economic activity. In comparison, my study quantifies productivity effects relying on more granular novel data which allows to provide more precise estimations and disentangle heterogeneous effects across demographics and socio-economic strata. This paper departs from the literature by providing the first nationwide quasi-experimental estimates of the costs of air pollution that jointly account for both *health* and *productivity* costs as well as their spatial heterogeneity.

Second, this paper relates to the literature on private adaptations to environmental conditions. Previous studies have shown that individuals engage in a range of options available to them for adapting to changes in environmental conditions, such as defensive expenditures (e.g., Deschenes et al., 2017; Sun et al., 2017; Zhang and Mu, 2018; Williams and Phaneuf, 2019; Ito and Zhang, 2020) and avoidance behaviors (e.g., Moretti and Neidell, 2011; Zivin et al., 2011; Chen et al., 2020). Nevertheless,

only a handful of studies attempted to account for how private adaptation and the implied WTPs vary across different demographics, but this is generally hindered by data coarseness (cf. Drupp et al. 2021). My empirical analysis complements the existing literature by providing novel quasi-experimental evidence on how accounting for demand heterogeneity along the income distribution affects the computation of revealed preference estimates of the societal benefits of air pollution reductions.

The remainder of this paper is structured as follows. Section 5.2 introduces a conceptual framework to elucidate the focus of the empirical investigation and detail how the effects of pollution on (i) defensive expenditures and (ii) economic productivity relate to the WTP for clean air. Sections 5.3 and 5.4 describe the data and the 2SLS strategy. Section 5.5 presents the empirical results. Section 5.6 discusses policy implications. Section 5.7 concludes.

5.2 Conceptual framework

Becker-Grossman health production function. This section lays out a conceptual framework drawing on the Becker-Grossman health production function to elucidate how the effects of pollution on (i) defensive expenditures and (ii) economic productivity relate to the WTP for clean air (Becker, 1965; Grossman, 1972). The model shows that an accurate measurement of the WTP requires knowledge of both (i) and (ii), as well as an understanding of how air pollution affects health outcomes - including morbidity and mortality. In this setting, the health production function takes the following form:

$$H = H(D(\psi), \psi(P)) \quad (5.1)$$

As shown in Eq. 5.1, pollution-driven sickness episodes, ψ , depend on the concentration level ($\mu\text{g}/\text{m}^3$) of the ambient pollutant (P). Defensive expenditures (D) are in turn determined by these episodes. Similarly to Graff Zivin and Neidell (2013) and Deschenes et al. (2017), this model assumes that defensive expenditures can be made before or after the exposure to pollutants and refers to both avoidance and

mitigating behavior that reduces the negative health consequences from air pollution exposure. Individuals gain utility from the consumption of non-health-related goods (C), leisure (L), and health and are assumed to receive nonlabor income (I) from either capital or transfer payments and to work for a given wage rate (w). The individual utility maximization problem is:

$$\max_{C, L, H} \mathcal{L} = U(C, L, H) + \lambda[I + w(H)[T - L] - p_C C - p_D D] \quad (5.2)$$

Given that defensive expenditures and the use of health care services ultimately depend on ambient pollution levels, the relationship between health and pollution levels can be expressed as the following total derivative of Eq. 5.1:

$$\frac{dH}{dP} = \underbrace{\left(\frac{\partial H}{\partial D} \frac{\partial D}{\partial \psi} + \frac{\partial H}{\partial \psi} \right)}_{\frac{dH}{d\psi}} \left(\frac{\partial \psi}{\partial P} \right) \quad (5.3)$$

In this setting, the effect of air pollution on population health therefore depends on two distinct components: the relationship between pollution and sickness episodes ($\partial\psi/\partial P$) and the extent to which these episodes translate into lowered health status ($\partial H/\partial\psi$), which is mitigated by individual defensive behavior. Denoting the costs of regulation (R) as p_R , the marginal WTP for air quality improvements (w_R) associated with environmental regulation can be expressed in monetary terms with the following decomposition:

$$\frac{\partial P}{\partial R} p_R = w_R = \underbrace{\frac{\partial D}{\partial P} p_D}_{(I')} + \underbrace{\frac{\partial w}{\partial H} \frac{dH}{dP}}_{(II')} + \frac{\partial U}{\partial H} \frac{dH}{dP} \frac{1}{\lambda} \quad (5.4)$$

Equation 5.4 shows that the marginal WTP for clean air w_R is composed of three terms. The first term (**I'**) reflects the cost of defensive expenditures, valued at their market price, p_D . The second (**II'**) captures the effect of pollution on productive working time, valued at the wage rate. Finally, the third component represents the

disutility of pollution-induced sickness, valued in monetary terms. Optimal regulation is defined at the point where the marginal costs of environmental regulation are equal to the reduced costs associated with that marginal reduction in pollution. The primary empirical goal of this paper is to develop a measure of marginal WTP that is based on I' and II' .

Divergence from the Neoclassical framework. The remainder of this section discusses how the empirical setting of this paper deviates from this neoclassical framework.

First, England relies on a publicly-funded health system, thus the cost burden of increased healthcare use would generally not affect individuals directly or, at least, to a lesser extent as opposed to health systems based on private medical insurance. That said, rather than affecting the total economic benefits of pollution control, the funding structure of the healthcare system is expected to alter their distribution. How costs are ultimately split is not expected to affect the estimation of the social WTP (Deschenes et al., 2017).

Second, in the English public NHS, the marginal cost of pharmaceuticals to the consumer may be smaller than their market price. Prescription medications face a fixed charge that is updated every year and functions as a copayment for pharmaceuticals. As the remainder of the pharmaceutical cost is covered by the NHS, the empirical estimation of the social WTP will account for the full transacted price of pharmaceuticals (as listed in the national Drug Tariff)⁸ to capture the increase in both private *and* public spending linked to their consumption.

Third, British legislation mandates that every employee has a legal right to a paid sick leave should they be too ill to work, implying that temporary pollution-driven sickness episodes are likely not going to materialize as individual wage reductions. Lost work time due to pollution-driven sickness will still be accounted for in the

⁸The NHS Prescription Services produces the Drug Tariff monthly on behalf of the Department of Health and Social Care and can be accessed here.

empirical investigation as a cost to the economy, but one borne by employers in the form of foregone productivity.

Fourth, the health outcome of interest in my empirical analysis is pollution-driven respiratory, cardiovascular, and cognitive morbidity episodes treated by medications.

Finally, my WTP estimates provide a conservative lower-bound estimate of pollution reduction's total economic benefits: for instance, they do not account for additional defensive investments (e.g., Ito and Zhang 2020), effects on mortality (e.g., Deryugina et al. 2019) as well as additional benefits including material damage (e.g., Brimblecombe 2003) and other minor discomforts such as decreased visibility (Hyslop, 2009).

5.3 Data and descriptives

To carry out the empirical analysis, I compile a novel annual dataset spanning from 2012 to 2018 that merges nationwide gridded reanalysis pollution and weather data with a comprehensive panel dataset of administrative healthcare records and official data on local economic activity. This section describes the different data sources, cleaning procedures, and key variables employed in this study.

Pharmaceutical prescriptions. Monthly pharmaceutical prescriptions are drawn from NHS Digital for the period from 2012 to 2019. Practice-level national prescription data is provided by the NHS in England and released under the terms of the Open Government Licence. This study aims to capture sickness episodes induced and exacerbated by pollution exposure that is treated by prescription medications to account for both acute *and* chronic effects in the assessment of the economic burden of morbidity.⁹ England offers an ideal setting to carry out this analysis as the country relies on a publicly-funded universal NHS where residents are automatically

⁹Many existing studies look at emergency rooms or hospital visits, which tend to be representative of more acute episodes rather than chronic morbidity. Acute effects only account for a limited share of the overall burden of health conditions attributable to air pollution (Chanel et al., 2016).

registered. The concentration of healthcare services into a single provider allows me to exploit detailed comparable information on practice-level healthcare records for an average of approximately 54 million registered patients across the country. The dataset typically contains over 10 million records per month covering each practice in England, providing information on the consumption of over 20 thousand different prescription items.¹⁰

The key variables that I extract for my analysis are the practice code and its postcode, the medication identifier, the number of prescription items, and the associated expenditures as measured by the Net Ingredient Cost (NIC).¹¹ Each medication is classified under a specific therapeutic section within the British National Formulary (BNF) chapter.¹² Informed by the epidemiological literature, the categories selected for this study are cardiovascular (BNF section 2), respiratory (BNF section 3), and central nervous (BNF section 4) systems. General practitioner (GP) practices' location has been geo-coded (cf. Figure 5.A12 in the Appendix) and assigned to 5km x 5km grids under the British National Grid (BNG) reference system, which serves as the unit of observation in the analysis.¹³ As I only observe the practice address, I assume that people tend to register with practices closer to their residential address, within the same grid. This assumption is supported by a recent study indicating that, in England, individuals tend to select a practice located around 2 kilometers away from where they live (Santos et al., 2017).

To account for potential reporting errors, I exclude extreme outliers identified as practices reporting prescription items per capita above the 99th percentile. I ad-


¹⁰The data includes prescriptions written by GPs and other non-medical prescribers (such as nurses and pharmacists) who are attached to practices. Where prescribing cannot be linked to a practice, the data is excluded. This accounts for less than 0.1% of all prescribed items. It does not include items that have been dispensed in England but prescribed in another country. Furthermore, the prescription data does not cover voluntary private health insurance. However, the latter only accounts for less than 0.03% of annual expenditures on medical goods in UK accounts (ONS, 2018).

¹¹This refers to the net ingredient costs of prescribed items following the price listed in the national Drug Tariff. I additionally observe total costs for the NHS linked to pharmaceutical prescriptions in each practice which I leverage for robustness exercises.

¹²This can be accessed from [here](#).

¹³Utilizing LAUs as an alternative level of aggregation produces results within the 95% confidence intervals of my baseline findings (cf. Figure 5.A22).

Figure 5.1: Prescription items per capita across income deprivation quintiles



Note: The boxplots above display the the distribution of prescription items per capita across income deprivation quintiles, where 1 refers to the least deprived areas of the country and 5 represent the most deprived ones.

ditionally gather yearly information on practice characteristics from the General Practice Workforce data available from 2012. It contains information on registered patients, GPs headcount, its breakdown by age and gender, and country of qualification. Descriptive evidence on the distribution of GPs, total patient counts, and prescription items by therapeutic section throughout the country are summarized in Figures 5.A10 and 5.A11.

Census-level socio-economic indicators. I match all practice postcodes with the respective census output areas and their associated Index of Multiple Deprivation (IMD). This is a multidimensional composite index including dimensions related to income, employment, health, education, and crime. The four constituent nations of the UK have each developed their own IMD. The IMD for England is published by the UK Ministry of Housing, Communities and Local Government (MHCLG). These have been built to identify small area concentrations of deprivation, and are based on a methodology developed at the University of Oxford Social Disadvantage Research Centre (Noble et al., 2006). I consider IMD scores at the Middle Layer Super Output Area (MSOAs) level.¹⁴

¹⁴MSOAs are a statistical geography created for the Census of England and Wales with a typical population between 7000 and 10000 people.

Harnessing the IMD, the boxplots in Figure 5.1 provide descriptive evidence on how demand for prescription items differs along the income distribution in the sample. Specifically, I plot prescription items per capita across income deprivation score quintiles. The somewhat lower median in wealthier areas (or, in other words, in the first income deprivation quintile) implies that, on average, residents in these regions require or use fewer prescription items compared to those in lower-income areas. Interquartile ranges, instead, are comparable across income deprivation quintiles indicating that the variability in prescription item usage is overall consistent among income groups. This suggests that the majority of individuals in both wealthier and poorer areas have prescription item counts that fall within a similar range, implying a level of uniformity in pharmaceutical consumption across income brackets. This preliminary descriptive understanding sets the stage for my subsequent causal analysis of the impacts of pollution on health outcomes.

Sick leaves. I additionally retrieve information on the annual number of fit notes issued by NHS local health authorities and diagnoses (identified by an ICD10 chapter).¹⁵ Digital fit notes, also known as the *Med3 form*, were introduced in April 2010 in England, Wales, and Scotland. Nevertheless, data is publicly available from April 2015 onward only. A fit note is given after a patient has been sick for 7 days, at which point they can no longer self-certify. Differently from the prescription data, information on sick leaves is provided at the clinical commissioning group level (CCG).¹⁶ CCGs were NHS organizations set up by the Health and Social Care Act 2012 to oversee the delivery of healthcare services in local areas across England but were abolished in 2022. One limitation pertains to the unavailability of precise information on the exact count of lost workdays; instead, CCG data from NHS Digital only captures the issuance of fit notes. I rely on this data to provide complementary insights into the effect of air pollution on work absenteeism.

¹⁵This is the 10th revision of the International Statistical Classification of Diseases and Related Health Problems (ICD). More information can be found here.

¹⁶Figure 5.A13 in the Appendix 5.A maps their distribution across the country.


Economic activity. Information on annual estimates of balanced UK gross value added (GVA) at the LAU level is obtained from the Office of National Statistics (ONS).¹⁷ At the time of the analysis, England comprised a total of 317 local authorities (or districts for simplification), covering the entire country (details can be found here on the government website). In this study, GVA per capita serves as a proxy for economic productivity. The primary rationale for this choice is that GVA stands as among the most accessible metrics for gauging productivity, which facilitates international comparisons to contextualize magnitudes. I combine information on local economic activity with other official LAU-level statistics from the ONS such as population and median gross weekly salary.

Thermal inversions as IV. Obtaining empirical estimates of the effects of air pollution exposure is challenging, primarily due to the scarcity of exogenous shocks in air quality driven by sorting dynamics (e.g., Heblich et al. 2021). To address endogeneity, I rely on atmospheric temperature inversions at different pressure levels as a source of exogenous variation in the spatial concentration of air pollutants (e.g., Dechezleprêtre et al., 2019). Figure 5.2 illustrates the concept of using inversions as an IV for pollution.

To leverage inversions in a two-stage least squares (2SLS) framework, I rely on vertical temperature profiles with a 10 km x 10 km resolution from the ECMWF extracted from the UERRA dataset (Copernicus Climate Change Service, 2019). The most recent complete year included in the dataset is 2018, which therefore determines the last year in the estimation sample. The spatial granularity of this novel dataset allows computing inversion episodes with greater precision and accuracy compared to previous studies that have relied on this instrument. The interval between pressure levels is 25 hPa, starting from the lowest 1000 hPa pressure level (which approximately corresponds to 30m above sea level). I extract data on a 6-hour frequency on surface-

¹⁷LAU stands for local administrative units which is a classification of spatial units used for statistical production across the European Union, aligning closely with district delineations in many regions.

Figure 5.2: Thermal inversion episodes as instrument for air pollution




Note: The figure above illustrates the concept of thermal inversions. Under normal conditions, air temperature typically decreases with increasing altitude. However, under specific atmospheric conditions known as thermal inversions, this natural relationship undergoes a reversal. During a thermal inversion, warmer air lies above cooler air, hindering the vertical movement of air masses. This condition traps pollutants closer to the ground, leading to the buildup of air emissions and reduced air quality.

level temperature (1000 hPa) and for the closest pressure levels, namely 975 hPa, 950 hPa, and 925 hPa.

I define inversions as a positive upward temperature gradient between the pressure level considered (either 975 hPa, 950 hPa or 925 hPa) and the surface (1000 hPa) calculated on a 6-hour frequency. The inversion instrument is defined as the annual frequency of thermal inversion events detected within a given geographical unit. I rely on the highest time frequency available to capture as many inversion events as possible to address concerns that low-frequency events as instruments may lead to inflated estimates due to low statistical power when estimating acute health effects (Bagilet and Zabrocki-Hallak, 2022). In the definition of the instrumental variable, there is a trade-off between relying on inversion episodes occurring closer to the surface, which are more likely to induce a pollution exposure shock but exhibit lower frequencies, and considering additional pressure levels. Such variation in frequency counts can be seen in Figure 5.A8 which plots the matrix of the monthly frequency of inversions at different pressure levels.

My baseline regressions rely on inversion episodes between the two pressure levels closest to the surface (which corresponds to 1000 hPa and 950 hPa). This is the first pressure level that does not exclude considerable areas of the country

Figure 5.3: Correlogram of different definitions of inversion episode




Note: The figure plots the correlation matrix across different definitions of the instrumental variable. *hPa925* refers to inversion events occurring below a pressure level of 925 hPa; *hPa950* refers to inversion events occurring below a pressure level of 950 hPa; *hPa975* refers to inversion events occurring below a pressure level of 975 hPa; *Nights* refers to inversion events occurring at night and below a pressure level of 950 hPa.

at a higher altitude, which may otherwise result in unrepresentative results. As a robustness check, I additionally consider inversion episodes occurring at night to alleviate concerns regarding potential daytime inversions being noticeable for individuals (cf. Sager 2019). Figure 5.3 plots the correlation matrix among alternative inversion definitions utilized in the empirical analysis, illustrating consistently high correlations between my baseline and alternative definitions. A graphical example of the computation of inversion episodes below 950 hPa averaged across years is provided in Figure 5.4. The corresponding figures for each year can be found in the Appendix (cf. Figures 5.A1 - 5.A7).

Meteorological conditions. Gridded weather controls including mean ground-level temperature ($^{\circ}\text{C}$), precipitation (mm), wind speed (knots), and relative humidity (%) are sourced from the UK Meteorological (Met) Office. While weather information is accessible in the ECMWF dataset, I opt for ground measurements by the UK Met Office due to their higher resolution (i.e., 5km x 5km grids), which reduces

Figure 5.4: Distribution of inversion events (950 hPa) in the UK



Notes: The figure plots the distribution of annual inversion events below 950 hPa computed on a 6-hour frequency averaged between 2012 and 2018. The spatial resolution corresponds to 5km x 5km grids following the BNG reference system.

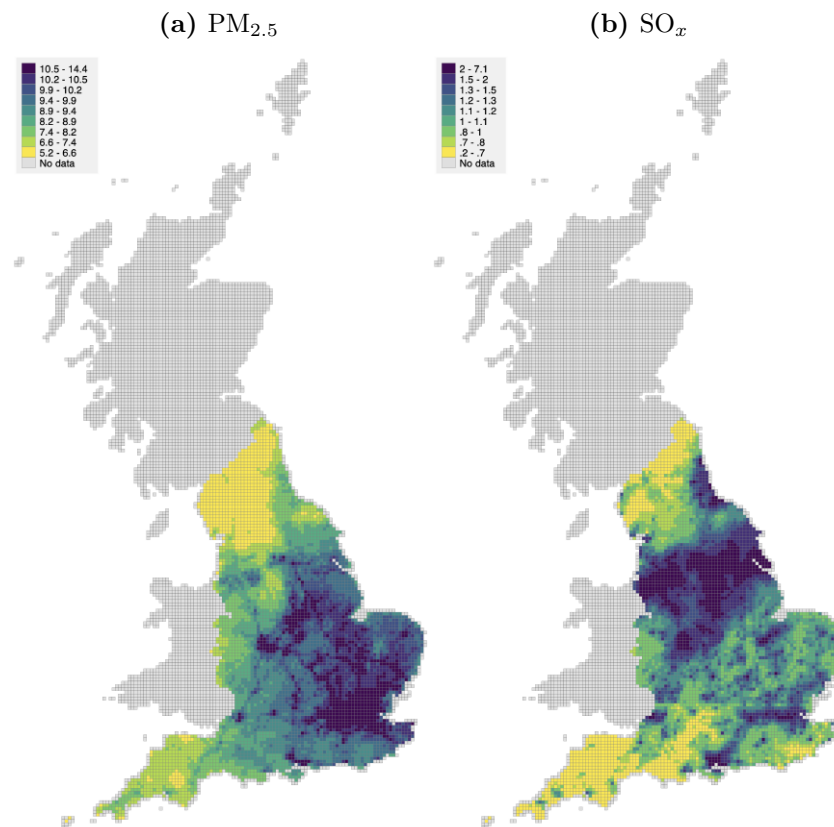
measurement errors.

Air pollutants. I rely on background pollution maps at 1km x 1km resolution that are modelled each year under the UK Department for Environment, Food and Rural Affairs (DEFRA) Modelling of Ambient Air Quality (MAAQ) contract. These maps are used to provide policy support for DEFRA and, formerly, to fulfill the UK's reporting obligations to Europe and cover key major air pollutants, namely concentrations of $PM_{2.5}$, PM_{10} , NO_x , NO_2 , SO_x , and C_6H_6 . They are made publicly available for public health research applications. Following the World Health Organization (2016), I focus on $PM_{2.5}$ as a general indicator of air pollution and harness information on other air pollutants for further robustness analyses.

The use of gridded reanalysis pollution data provides a key advantage for measuring the average pollution exposure of people living in a specific area compared to

pollution-monitor readings. Pollution monitors are often limited in number and placed strategically to measure emissions from specific sources, such as highways. Therefore, monitor readings may not provide an accurate reflection of the average pollution exposure in a particular location (e.g., Zou 2021). In contrast, reanalysis data synthesizes information derived from pollution monitors with the assistance of a chemistry-transport model that takes into account all potential pollution sources, thereby yielding a more comprehensive and holistic measurement of average exposure to pollution. Figure 5.5 plots the average concentration distributions of two key pollutants in my sample, namely $PM_{2.5}$ and SO_x , across England.

Figure 5.5: Distribution of $PM_{2.5}$ and SO_x concentrations across England



Notes: The figures above plot the distribution of $PM_{2.5}$ and SO_x concentrations in England across 5km x 5km grids following the BNG reference system. The legend displays concentration defined in $\mu g/m^3$. The corresponding Figures for the other pollutants in my sample can be found in Figure 5.A9 in the Appendix.

Inequality in exposure to air pollutants. Figure 5.A15 in the Appendix provides some additional descriptive evidence on the distribution of pollution concentration across socio-economic characteristics. Specifically, they plot the pre-sample

(values are for 2011) quadratic prediction of local pollutant concentration based on income deprivation scores to investigate pollution exposure inequality along the income distribution. Results across all pollutants provide descriptive evidence of the degree of inequality in the distribution of pollution across different socio-economic groups in the sample, allowing for nonlinearities. In line with existing studies, these findings underscore a consistent regressive distribution of pollution exposure across income brackets, displaying varying degrees among different air pollutants (cf. Colmer et al. 2020; Jbaily et al. 2022). This initial descriptive exploration not only lays the foundation for understanding the distribution of pollution across different income groups before the period of analysis but also serves as a critical step in justifying the adoption of an instrumental variable approach. By shedding light on the disparities in pollution exposure, it underscores the need to address estimation biases from existing pre-sample sorting dynamics and provides empirical support for the use of instrumental variables in the empirical analysis (cf. Heblich et al. 2021).

5.4 Empirical strategy

2SLS estimation. The main empirical goal of this paper is to estimate the effects of short-term exposure to air pollutants on morbidity and productivity, accounting for potentially confounding factors. I model this relationship relying on a two-stage estimation. In the first stage, I predict the concentration of different air pollutants based on the observed frequency of thermal inversions occurring below a given altitude h , as measured by atmospheric pressure levels (in hPa). In the second stage, I estimate the effect of the predicted concentration estimate on (i) pharmaceutical expenditures and (ii) GVA per capita. The main identifying assumption of my IV approach is that, after flexibly controlling for a set of fixed effects and weather variables, changes in a geographical unit's annual frequency of inversion episodes are unrelated to changes in any of the outcome variables except through their impact on

air pollution.¹⁸

For each air pollutant p , the reduced form (Eq. 5.5), the first stage (Eq. 5.6) and the second stage (Eq. 5.7) are written as:

$$y_{i \subset j, t}^p = \alpha_1^p \underbrace{\text{Frequency of Inversions}_{i \subset j, t, h}}_{\text{Instrument}} + \gamma_{i \subset j, t}^p + \delta_t^p + \phi_j^p + \epsilon_{i \subset j, t, h}^p \quad (5.5)$$

$$C_{i \subset j, t}^p = \alpha_1^p \underbrace{\text{Frequency of Inversions}_{i \subset j, t, h}}_{\text{Instrument}} + \gamma_{i \subset j, t}^p + \delta_t^p + \phi_j^p + \epsilon_{i \subset j, t, h}^p \quad (5.6)$$

$$y_{i \subset j, t}^p = \beta_1^p \hat{C}_{i \subset j, t}^p + \gamma_{i \subset j, t}^p + \delta_t^p + \phi_j^p + \epsilon_{i \subset j, t, h}^p \quad (5.7)$$

where $y_{i \subset j, t}^p$ is the natural logarithm of the outcome variable, which is either (i) the expenditure on prescriptions for cardiovascular, respiratory, and nervous conditions or (ii) GVA per capita in each geographical unit, i , within a LAU, j , and year, t . $C_{i \subset j, t}^p$ reflects the annual average concentration of pollutant p measured in $\mu\text{g}/\text{m}^3$. The model absorbs LAUs (ϕ_j^p) and year fixed effects (δ_t^p), meaning that my identification strategy exploits within-district variation in thermal inversion exposure (Correia 2016, 2019). The former absorbs, for instance, spatial variation in healthcare quality and access, diagnostic standards, and environmental quality, whereas the latter controls flexibly for common time-varying shocks, such as those induced by any NHS or environmental policy changes in the sample period. The vector $\gamma_{i \subset j, t}^p$ includes covariates such as yearly average wind speed, ground-level temperature, humidity, and rainfall to account for other environmental factors that may affect the outcomes. $\epsilon_{i \subset j, t, h}^p$ is the idiosyncratic error term.

Since pollution observed in a given grid is likely driven by emissions elsewhere that also affect nearby grids, all grid-level inferences allow for correlations in errors across

¹⁸An underlying assumption in my empirical strategy is that the assumed residence corresponds to the location of exposure to air pollution. However, people are also exposed to air pollution at their place of work, place of leisure, or while commuting. Should this not be the case, the measurement error in pollution exposure would be inflated and my estimates could be biased toward zero (attenuation bias). I additionally leverage night inversions to investigate the sensitivity of this assumption.

neighboring grids by relying on a more aggregate cluster dimension: my baseline models rely on 50km grids from the BNG reference system but results are robust to a range of clustering choices (cf. Appendix 5.B). All estimates are weighted by either the number of patients (when modeling pharmaceutical consumption) or the total population (when modeling GVA per capita) within a given geographical unit. This is to account for the fact that the relevant population is unevenly distributed across districts in the country and exposed to different levels of pollution depending on their location. I additionally report OLS estimates of Eq. 5.7 to test whether these are prone to bias as exposure to pollution is not randomly assigned and is likely measured with error.

Identifying assumptions. The estimation of an unbiased estimate of the causal effect of air pollution with an IV approach rests on meeting a set of identifying assumptions. Precisely, these assumptions include (i) instrument relevance, (ii) the exclusion restriction, and (iii) monotonicity - as outlined in Angrist and Imbens (1995).

Instrument relevance requires a significant impact of inversion episodes on air pollution concentrations, which can be directly tested through the first stage (cf. Eq. 5.6). The exclusion restriction implies that inversion episodes need to be randomly assigned, meaning that inversions are expected to affect any outcome of interest only through their impact on pollution concentration. Given that inversions originate from continental-scale air movements, they are plausibly unlikely to be affected by local-scale socio-economic factors that are being modeled in this study, supporting their credibility as a source of quasi-random variation. That is, thermal inversions increase pollution levels without being correlated with either the causes of polluting emissions (e.g., industry or transport) or its effects (e.g., on health and productivity). To further bolster the validity of the exclusion restriction, the 2SLS estimation employed in this setting additionally controls for weather conditions that may correlate with inversion frequency (e.g., temperature) while also impacting socio-economic outcomes.

This ensures that the estimation is specifically isolating the effect of increased air pollution, excluding the influence of other co-varying weather conditions. Finally, the monotonicity assumption rules out scenarios where inversions consistently produce an opposite impact on pollution compared to the overall trend (i.e., reductions rather than increases in pollution concentrations). If these assumptions hold and treatment effects remain constant conditional on a set of covariates, the 2SLS approach will yield an unbiased estimate of the local average treatment effect (LATE).

5.5 Results

The following section presents results from the two-stage estimation strategy described in Section 5.4. As a first step, I present the causal effects of instrumented pollution shocks on nationwide (i) pharmaceutical expenditures and (ii) GVA per capita in Section 5.5.1. Subsequently, in Section 5.5.2, I leverage these causal estimations to simulate counterfactual reductions of $1 \mu\text{g}/\text{m}^3$ in $\text{PM}_{2.5}$ and investigate the respective distributions of predicted health and productivity benefits and their correlation with socio-economic factors.

5.5.1 Nationwide effects on morbidity and productivity

Tables 5.1 and 5.2 report estimates of equation 5.7, where I regress (i) pharmaceutical expenditures and (ii) GVA per capita (cf. Section 5.3) on instrumented pollution and controls. One key difference between the estimation of (i) and (ii) is the unit of observation. While the former estimation relies on data at the 5km x 5km grid level, the latter is based on LAUs as it represents the highest available spatial resolution (cf. Section 5.3).¹⁹ I begin by presenting first-stage and second-stage

¹⁹Another difference relates to the time fixed effects. The higher granularity of the data in the estimation of effects on (i) pharmaceutical expenditures allows me to include region-specific year (instead of nationwide year effects) effects to absorb variation in regional effects in my baseline regression while leaving enough identifying variation to estimate the model. Figure 5.7 shows that the inclusion of more detailed trends in the estimation of (ii) GVA per capita does not affect my

results obtained from different definitions of inversion episodes, which are derived considering different atmospheric pressure levels, h . Specifically, column (1) presents results for my baseline regression model based on inversions between the two pressure levels closest to the surface, namely 1000 hPa and 950 hPa (cf. Section 5.3). Column (2) and column (3) report results when considering 925 hPa and 975 hPa, respectively, as alternative upper pressure levels.

First-stage estimates. Panel A of Tables 5.1 and 5.2 display the estimates of the first stage estimation from Eq. 5.6. The large first-stage F-statistics confirm that the frequency of inversion episodes is a strong relevant predictor of air pollution levels, implying that *weak instrument bias* is not a source of concern in this setting.²⁰ The validity of the identification strategy is further corroborated by the results of the Lagrange Multiplier (LM) test for under-identification. Taking column (1) from Table 5.1 as an example, mean annual inversion episodes in our sample (which amounts to 121 episodes) are estimated to cause an increase of around $0.9 \mu\text{g}/\text{m}^3$ in annual average $\text{PM}_{2.5}$ concentration (i.e., 0.00732×121). Column (2) and column (3) show that the coefficients are not sensitive to harnessing different instrument definitions. Similarly, Table 5.2 exhibits the corresponding results for the first-stage estimation in the LAU-level sample, employed to model GVA per capita. Even in this case, each column exhibits considerable statistical significance, affirming their validity as instruments.

Second-stage estimates. Panel B of Tables 5.1 and 5.2 reports estimates of Eq. 5.7, where I regress pharmaceutical expenditures and GVA per capita on instrumented pollution and controls. Across the three specifications, the coefficients on instrumented pollution show that a $1 \mu\text{g}/\text{m}^3$ increase in the concentration of any of the pollutants significantly affects defensive investment behavior (cf. Table 5.1)

estimations (cf. *Local trends*).

²⁰Following Deryugina et al. (2019), Tables 5.1 and 5.2 present first-stage F-statistics computed assuming that errors are homoskedastic. This allows for comparison to the Stock and Yogo (2005) critical values, whose validity relies on the homoskedasticity assumption.

Table 5.1: Effects of a 1 $\mu\text{g}/\text{m}^3$ annual increase in $\text{PM}_{2.5}$ on pharmaceutical expenditures

Instrument: Inversion frequency below pressure level (h)	(1) $h = 950$ hPa	(2) $h = 925$ hPa	(3) $h = 975$ hPa
Panel A: First Stage (Eq. 5.6)	0.00732*** (0.00113)	0.00757*** (0.00111)	0.00661*** (0.00113)
Panel B: Second Stage (Eq. 5.7)	0.283** (0.114)	0.243** (0.107)	0.296** (0.124)
Cragg-Donald Wald F-statistic	888.3	803.3	767.2
Kleibergen-Paap rk LM statistic	$P < 0.000$	$P < 0.000$	$P < 0.000$
N	41830	41830	41830
Weather Controls	✓	✓	✓
LAU Effects	✓	✓	✓
Time Effects	Region x Year	Region x Year	Region x Year
Clustered Std Error	50km	50km	50km

Notes: Table shows the coefficients estimated from the first (Panel A) and second stage (Panel B) of the IV approach, where the frequency of thermal inversions is used as an instrument for pollution concentrations (cf. Section 5.4). GP practices' location has been geo-coded using GIS tools and assigned to 5km x 5km grids following the BNG reference system. Pharmaceutical expenditures are reported by each registered practice and have been aggregated at a 5km x 5km grid level. $\text{PM}_{2.5}$ concentration refers to average background annual average concentrations in $\mu\text{g}/\text{m}^3$ across 5km x 5km grids based on data reported by DEFRA. Thermal inversions are defined as a positive upward temperature gradient from the surface, and calculated on a 6-hour frequency using data from the ECMWF. Weather controls include mean ground-level temperature ($^{\circ}\text{C}$), precipitation (mm), wind speed (knots), and relative humidity (%). These are retrieved on a 3-hour frequency from the UK Met Office and aggregated at the yearly level. See Section 5.3 for more details. All regressions control for the total number of GPs in the area and are weighted by the total number of patients in each grid. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5.2: Effects of a 1 $\mu\text{g}/\text{m}^3$ annual increase in $\text{PM}_{2.5}$ on GVA per capita

Instrument: Inversion frequency below pressure level (h)	(1) $h = 950$ hPa	(2) $h = 925$ hPa	(3) $h = 975$ hPa
Panel A: First Stage (Eq. 5.6)	0.0170*** (0.00147)	0.00931*** (0.00141)	0.0153*** (0.00146)
Panel B: Second Stage (Eq. 5.7)	-0.0160** (0.00784)	-0.0254* (0.0135)	-0.0142 (0.00896)
Cragg-Donald Wald F-statistic	181.9	47.79	135.2
Kleibergen-Paap rk LM statistic	$P < 0.000$	$P < 0.000$	$P < 0.000$
N	2219	2219	2219
Weather Controls	✓	✓	✓
LAU Effects	✓	✓	✓
Time Effects	Year	Year	Year
Clustered Std Error	LAU	LAU	LAU

Notes: Table shows the coefficients estimated from the first (Panel A) and second stage (Panel B) of the IV approach, where the frequency of thermal inversions is used as an instrument for pollution concentrations (cf. Section 5.4). GVA per capita represents the ratio of local GVA, sourced from the UK ONS, divided by the total population in the LAU. $\text{PM}_{2.5}$ concentration refers to average background annual average concentrations in $\mu\text{g}/\text{m}^3$ across LAUs based on data reported by DEFRA. Thermal inversions are defined as a positive upward temperature gradient from the surface, and calculated on a 6-hour frequency using data from the ECMWF. Weather controls include mean ground-level temperature ($^{\circ}\text{C}$), precipitation (mm), wind speed (knots), and relative humidity (%). These are retrieved on a 3-hour frequency from the UK Met Office and aggregated at the yearly level. See Section 5.3 for more details. All regressions are weighted by the total population in each LAU. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

and economic productivity (cf. Table 5.2). More precisely, taking again the baseline estimation from column (1), a 1 $\mu\text{g}/\text{m}^3$ increase in $\text{PM}_{2.5}$ causes a 32.7% increase in pharmaceutical expenditure and 1.6% decrease in GVA per capita (following exponential transformation).²¹

To provide perspective to the magnitude of the cost estimates, we can compute the implied estimated annual nationwide health and productivity costs based on the coefficients presented in Tables 5.1 and 5.2. Specifically, a 1 $\mu\text{g}/\text{m}^3$ in the annual concentration of $\text{PM}_{2.5}$ translates into approximately £1.24 billion/year of additional health costs and around £13 billion/year of productivity losses, as proxied by GVA per capita.²² These findings highlight how both healthcare and productivity costs individually represent a sizable share of the economic costs of air pollution exposure. It follows that cost-benefit analyses based on empirical estimates that focus solely on either one of the two cost elements are likely to underestimate the potential welfare benefits associated with air pollution reduction. Taking $\text{PM}_{2.5}$, considering both health and productivity effects translates into a mean WTP of around £445 per capita which is around 9 times higher than the current damage costs per change in concentration reported by the UK government (UK-AIR, 2023).

Alternative model specifications. To corroborate the robustness of the findings discussed above, I examine various alternative model specifications deviating from the 2SLS approach described in Section 5.4, which are summarized in Figure 5.7. Overall, the series of robustness tests provides statistically comparable estimates, with my baseline mean point estimates typically leaning towards the more conservative end of their distribution.

First, I report results from a simple OLS estimation, where I directly regress pharmaceutical expenditures and GVA per capita on $\text{PM}_{2.5}$ pollution concentration and

²¹The estimated productivity losses attributable to $\text{PM}_{2.5}$ are close to the upper bounds of the estimated effects documented on a European Union-wide scale by Dechezleprêtre et al. (2019).

²²I rely on statistics from the Office of National Statistics to extract the average number of employed people in the country within the period under investigation, which is around 31 million.

control variables (cf. *OLS* specification). Notably, in line with previous quasi-experimental pollution studies (e.g., Deryugina et al. 2019), OLS results are consistently lower in magnitude (around 4 times) compared to IV estimates, highlighting the potential for bias in observational studies that do not account for the endogenous nature of air pollution.

Second, to address potential concerns of bias stemming from avoidance behavior associated with the visibility of daytime inversions, I introduce the use of night inversions (cf. Sager 2019) as a simultaneous additional instrument (cf. *Nights* specification).

Third, other potential concerns involve residual (i) unobserved local heterogeneity over time which could introduce omitted variable bias (Gormley and Matsa, 2014) and (ii) interactive weather conditions that might be (non-linearly) correlated with inversions and affect pharmaceutical consumption or productivity (e.g., extreme heat with high humidity). To this end, I introduce two additional specifications where (i) detailed heterogeneous trends at either the grid or LAU level are absorbed²³ in the estimation (cf. *Local trends* specification), and (ii) I additionally control for both the linear and quadratic interactions among the set of weather controls (cf. *Interactions* specification).


Fourth, the estimation of causal effects resulting from PM_{2.5} pollution faces a challenge due to the correlation between multiple air pollutants such as NO_x or SO_x, often originating from the same emission sources (cf. Figure 5.6). Consequently, detected adverse effects associated with air pollution may be attributable to any single or a combination of these pollutants. To tackle this concern, I present coefficient estimates of the impact of PM_{2.5} while incorporating NO_x and SO_x as covariates within the 2SLS estimation procedure (cf. *Other pollutants* specification). To additionally support my findings in Tables 5.1 and 5.2, I additionally replicate my baseline results

²³The process involves absorbing heterogeneous slopes with the estimator outlined in Correia (2016, 2019), which accommodates distinct coefficients for individual regressors across each fixed effect category.

in column (1) using different measures of air pollution, namely PM_{10} , NO_x , NO_2 , SO_x , and C_6H_6 . These results consistently provide qualitative support for my baseline findings, as summarized in Tables 5.3 to 5.4.

Finally, I report results where I allow for serial correlation in errors by clustering standard errors in two dimensions at the level of grids (or LAUs) and years (cf. *Year Cluster* specification).

Figure 5.6: Correlogram of different air pollutants




Note: The figure plots the correlation matrix across different concentrations of air pollutants (in $\mu g/m^3$). *PM* refers to $PM_{2.5}$.

Additional robustness tests. In the Appendix, I further conduct a placebo exercise where I estimate a RF of my baseline specification (cf. Eq. 5.5) that includes a set of leads and lags of my instrument to test whether the estimated effects are driven by contemporaneous inversion episodes in a given year and mitigate concerns that individuals may anticipate inversion episodes to a significant extent thus biasing any estimation on instrumented pollution (see Figure 5.A21 in the Appendix). Notably, the outcome of this placebo exercise rules out the prospect of anticipation acting as a confounding factor in my estimations. The Appendix also contains supplementary analyses displaying additional results on the impact of air pollution on item demand and overall estimated cost increases to the NHS, which are summarized in Figures 5.A19 - 5.A20. Finally, Figure 5.A16 in the Appendix


presents separate estimations for each BNF therapeutic section (cf. Section 5.3). The largest effects are detected for health conditions affecting the central nervous system, indicating that the health effects of air pollution extend significantly beyond the respiratory and cardiovascular diseases typically highlighted in prior research. (e.g., Deschenes et al. 2017).

Figure 5.7: Alternative model specifications

(a) Pharmaceutical expenditures



(b) GVA per capita



Notes: The figures above plot the estimated effect from the second stage of the alternative specifications of Eq. 5.7 described in Section 5.5.1. Panel (a) refers to results for pharmaceutical expenditures, whereas Panel (b) is based on results for GVA per capita. The bar charts display mean point estimates with overlaid lines representing confidence intervals, respectively delineating 95% confidence intervals in gray and 90% confidence intervals in color. Additional estimates based PM_{10} , NO_x , NO_2 , SO_x , and C_6H_6 are summarized in Figures 5.A17 and 5.A18 and can be found in the Appendix.

Table 5.3: Effects of an annual increase in air pollution on pharmaceutical expenditures

	(1)	(2)	(3)	(4)	(5)	(6)
	PM _{2.5}	PM ₁₀	NO _x	NO ₂	SO _x	C ₆ H ₆
	($\mu\text{g}/\text{m}^3$)	($\mu\text{g}/\text{m}^3$)	($\mu\text{g}/\text{m}^3$)	($\mu\text{g}/\text{m}^3$)	(0.1 $\mu\text{g}/\text{m}^3$)	(0.1 $\mu\text{g}/\text{m}^3$)
Panel A: OLS estimates	0.0679*** (0.0168)	-0.0134 (0.0171)	0.00551 (0.00616)	0.0227** (0.00636)	0.0100*** (0.00155)	0.109*** (0.0198)
Panel B: Second Stage (Eq. 5.7)	0.283** (0.114)	0.150** (0.0639)	0.0868*** (0.0326)	0.147*** (0.0523)	0.0600** (0.0231)	0.387** (0.147)
Cragg-Donald Wald F-statistic	888.3	1558.8	192.4	223.5	210.2	289.6
Kleibergen-Paap rk LM statistic	$P < 0.000$	$P < 0.000$	$P < 0.000$	$P < 0.000$	$P < 0.000$	$P < 0.000$
N	41830	41830	41830	41830	41830	41830
Weather Controls	✓	✓	✓	✓	✓	✓
LAU Effects	✓	✓	✓	✓	✓	✓
Region x Year Effects	✓	✓	✓	✓	✓	✓
Clustered Std Error	50km	50km	50km	50km	50km	50km

Notes: Table shows the coefficients estimated from the first (Panel A) and second stage (Panel B) of the IV approach, where the frequency of thermal inversions is used as an instrument for pollution concentrations (cf. Section 5.4). GP practices' location has been geo-coded using GIS tools and assigned to 5km x 5km grids following the BNG reference system. Pharmaceutical expenditures are reported by each registered practice and have been aggregated at a 5km x 5km grid level. Pollution concentrations refer to average background annual average concentrations in $\mu\text{g}/\text{m}^3$ across 5km x 5km grids based on data reported by DEFRA. Thermal inversions are defined as a positive upward temperature gradient from the surface below 950 hPa, and calculated on a 6-hour frequency using data from the ECMWF. Weather controls include mean ground-level temperature ($^{\circ}\text{C}$), precipitation (mm), wind speed (knots), and relative humidity (%). These are retrieved on a 3-hour frequency from the UK Met Office and aggregated at the yearly level. See Section 5.3 for more details. All regressions control for the total number of GPs in the area and are weighted by the total number of patients in each grid. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5.4: Effects of an annual increase in air pollution on GVA per capita

	(1)	(2)	(3)	(4)	(5)	(6)
	PM _{2.5}	PM ₁₀	NO _x	NO ₂	SO _x	C ₆ H ₆
	($\mu\text{g}/\text{m}^3$)	($\mu\text{g}/\text{m}^3$)	($\mu\text{g}/\text{m}^3$)	($\mu\text{g}/\text{m}^3$)	(0.1 $\mu\text{g}/\text{m}^3$)	(0.1 $\mu\text{g}/\text{m}^3$)
Panel A: OLS estimates	-0.00383 (0.00309)	-0.00196 (0.00257)	-0.00110 (0.000831)	-0.00170 (0.00151)	-0.000332 (0.000378)	-0.00417* (0.00236)
Panel B: Second Stage (Eq. 5.7)	-0.0160** (0.00784)	-0.0111** (0.00547)	-0.00489* (0.00250)	-0.00781** (0.00393)	-0.00211** (0.00106)	-0.0118** (0.00589)
Cragg-Donald Wald F-statistic	181.9	265.0	117.7	147.6	146.3	239.2
Kleibergen-Paap rk LM statistic	$P < 0.000$	$P < 0.000$	$P < 0.000$	$P < 0.000$	$P < 0.000$	$P < 0.000$
N	2219	2219	2219	2219	2219	2219
Weather Controls	✓	✓	✓	✓	✓	✓
LAU Effects	✓	✓	✓	✓	✓	✓
Year Effects	✓	✓	✓	✓	✓	✓
Clustered Std Error	LAU	LAU	LAU	LAU	LAU	LAU

Notes: Table shows the coefficients estimated from the first (Panel A) and second stage (Panel B) of the IV approach, where the frequency of thermal inversions is used as an instrument for pollution concentrations (cf. Section 5.4). GVA per capita represents the ratio of local GVA, sourced from the UK ONS, divided by the total population in the LAU. Pollution concentrations refer to average background annual average concentrations in $\mu\text{g}/\text{m}^3$ across LAUs based on data reported by DEFRA. Thermal inversions are defined as a positive upward temperature gradient from the surface below 950 hPa, and calculated on a 6-hour frequency using data from the ECMWF. Weather controls include mean ground-level temperature ($^{\circ}\text{C}$), precipitation (mm), wind speed (knots), and relative humidity (%). These are retrieved on a 3-hour frequency from the UK Met Office and aggregated at the yearly level. See Section 5.3 for more details. All regressions are weighted by the total population in each LAU. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Health-productivity-poverty nexus. The estimated sizable declines in economic productivity may be attributed to various underlying mechanisms. For instance, previous literature points out that air pollution shocks may manifest in increased absenteeism (e.g., Holub et al. 2020), diminished worker productivity (e.g., Leroutier and Ollivier 2022), direct impact through agricultural productivity (e.g., Avnery et al. 2011), or by hampering the accumulation of human capital (e.g., Ebenstein et al. 2016). All of these channels offer plausible explanations for the observed reductions in economic activity. In this section, I provide additional causal evidence to link pollution-driven morbidity to productivity losses by investigating work absenteeism as a driving factor. Specifically, I estimate the causal effects of pollution shocks on countrywide work absenteeism leveraging information on the local issuance of sick leaves.

To carry out this analysis, I amend the empirical strategy described in Section 5.4 in three ways. First, the analysis is carried out by exploiting variations within clinical commissioning groups (CCG) instead of leveraging within-district variation.²⁴ Second, the estimation sample begins in 2016, which reflects the more restricted temporal coverage of the sick leaves data. Third, I estimate the reduced-form effect of thermal inversions (cf. Eq. 5.5) on the number of sick leaves issued in a given period with a Poisson regression to accommodate the need to model count data. The results of the RF approach should be interpreted as the effect of a general pollution shock induced by a thermal inversion episode rather than the effect of a specific pollutant.²⁵

Table 5.5 reports results across the different health domains under investigation. The estimated Poisson regression coefficients report that pollution shocks are associated with a general increase in the issuance of sick leaves across all categories. The largest effect is detected for health conditions related to the central nervous system (column 3 in Table 5.5). Taking the estimated effect for all diagnoses from specification (1),

²⁴This is the highest level of resolution at which sick leaves data by diagnosis is available.

²⁵The key rationale behind this approach is the more limited power in the model due to the reduced number of observations which may be too demanding for a 2SLS estimation.

we can interpret the Poisson regression coefficients as follows: a standard deviation increase in pollution shocks driven by inversions²⁶ would lead to an overall increase of around 6% in sick leaves (i.e., $33 \times [e^{0.00184} - 1] \times 100\%$). This roughly translates into an annual increase of 23 thousand sick leaves being issued nationwide.

Although the limited granularity of sick leave data prevents the precise quantification of total lost workdays, these causal insights imply that increased absenteeism represents a relevant mechanism contributing to productivity losses due to air pollution. This finding also implies that a reduced capacity to work caused by pollution might trap low-income individuals in a cycle of mounting morbidity and poverty (cf. Isen et al. 2017; Chetty and Hendren 2018; Ketcham et al. 2019), due to their higher level of exposure.

Table 5.5: Effects of air pollution on the issuance of sick leaves by diagnosis.

	(1) All	(2) Cardiovascular	(3) Nervous	(4) Respiratory
Inversion episodes	0.00184* (0.000954)	0.00219* (0.00126)	0.00257** (0.00115)	0.00149 (0.00100)
N	1701	567	567	567
Controls	✓	✓	✓	✓
Time Effects	NHS Region x Year	NHS Region x Year	NHS Region x Year	NHS Region x Year
CCG Effects	✓	✓	✓	✓
Clustered Std Error	CCG	CCG	CCG	CCG

Notes: Table shows the coefficients estimated from the RF of the IV approach, where the number of sick leaves is regressed on the frequency of thermal inversions (cf. Section 5.4). Sick leaves refer to the annual count of fit notes issued in each CCG. Year effects are specific to each NHS region in the country, namely London and Central, North, and South of England. Thermal inversions are defined as a positive upward temperature gradient from the surface below 950 hPa, and calculated on a 6-hour frequency using data from the ECMWF. Weather controls include linear and quadratic mean ground-level temperature ($^{\circ}\text{C}$), precipitation (mm), wind speed (knots), and relative humidity (%) as well as their interactions. These are retrieved on a 3-hour frequency from the UK Met Office and aggregated at the yearly level. See Section 5.3 for more details. All models additionally control for the total number of GPs and patient counts within older age brackets (i.e., 45-64, 65-74, and 75-84). Regressions are weighted by the total number of registered patients in each CCG. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.5.2 Spatial heterogeneity

The empirical analysis continues by examining spatial heterogeneity. To this end, I leverage the second-stage predicted estimates from Equation 5.7 to compute a counterfactual pollution reduction scenarios as follows:

²⁶This equals an increase of 33 episodes in the annual frequency of inversions when considering inversion episodes occurring below 950 hPa.

$$\Delta \widehat{y_{i \subset j, t}^{PM_{2.5}}} = \widehat{\beta}_1^{PM_{2.5}} \Delta \widehat{C_{i \subset j, t}^{PM_{2.5}}} \quad (5.8)$$

Specifically, I focus on a $1 \mu\text{g}/\text{m}^3$ nationwide reduction in the annual concentration of $\text{PM}_{2.5}$ (i.e., $\Delta \widehat{C_{i \subset j, t}^{PM_{2.5}}} = 1 \mu\text{g}/\text{m}^3$) to map the distribution of the corresponding predicted (i) health and (ii) productivity benefits per capita across the country.

Health benefits. Panel (a) in Figure 5.10 plots predicted health benefits per capita induced by the simulated $1 \mu\text{g}/\text{m}^3$ reduction shock in $\text{PM}_{2.5}$ estimated from Eq. 5.8. Here, health benefits are defined as the reduction in pharmaceutical expenditures per registered patient due to a $1 \mu\text{g}/\text{m}^3$ decrease in $\text{PM}_{2.5}$. Overall, we observe significant heterogeneity across different areas, with values ranging from less than £1 to around £140 per capita (cf. Panel (a) in Figure 5.9). Leveraging census-level variation in the IMD (cf. Section 5.3), I systematically explore how the distribution of the estimated benefits varies across different socio-economic strata and how this affects the computation of the implied WTP. Specifically, to delve deeper into spatial disparities, I compute quadratic predictions ($\Delta \widehat{y_{i \subset j, t}^{PM_{2.5}}} = \widehat{\beta}_0 + \widehat{\beta}_1 x + \widehat{\beta}_2 x^2$) of health benefits in my sample based on income deprivation scores, accommodating nonlinearities. These correlational results are presented in Panel (b) of Figure 5.9 and reveal that health benefits from $\text{PM}_{2.5}$ reduction tend to be distributed pro-poor, exhibiting a non-linear increase with income deprivation.


Specifically, predicted health benefits increase by around threefold along the spectrum of income deprivation, from roughly £10 to £30 per capita. This result is consistent with prior findings in the literature that highlighted how households' adaptive capacity to environmental shocks tends to correlate with disposable income (cf. Drupp et al., 2021), which can in turn exacerbate environmental inequalities along the income distribution. Case in point, other than sorting themselves into less polluted areas, households within the upper end of the income distribution are also more likely to

afford preemptive defensive expenses, such as purchasing an air purifier (e.g., Ito and Zhang 2020), which would, in turn, explain the reduced health impacts following a pollution shock.²⁷ Contingent on fulfilling its identifying assumptions, the 2SLS design addresses potential sorting bias by leveraging on quasi-random shocks in pollution. Hence, this implies that the observed differential effects across income groups are plausibly driven by residual differences in adaptive capacity rather than self-selection based on pollution concentrations.

Finally, I harness the same approach to investigate the distribution of health benefits across the age structure of registered patients across the sample. As illustrated in Panel (c) of Figure 5.9, when the age structure of the patient cohort skews towards a larger proportion of individuals aged 65 years or older, the predicted health benefits exhibit a discernible non-linear increase. Precisely, predicted health benefits exhibit a fourfold range (£6 - £25 per capita) across the spectrum of elderly patient shares. This insight has significant policy implications, suggesting that pollution-reduction interventions can have disproportionately positive effects on the health status of older demographics, corroborating previous findings on mortality effects (e.g., Deryugina et al. 2019).

²⁷While England's public NHS is expected to mitigate income inequalities in healthcare access, my dataset lacks information on individual households' adaptive capacity, preventing the inclusion of additional preemptive defensive behavior investments (and their distribution) in my analysis. Hence, my estimates present a conservative estimate of health-related air pollution costs (cf. Section 5.2).

Figure 5.8: Distribution of health benefits




Notes: Panel (a) displays the frequency distribution of predicted health benefits from a $\mu\text{g}/\text{m}^3$ $\text{PM}_{2.5}$ reduction in the second-stage estimation (cf. Eq. 5.8). Panel (b) reports the quadratic prediction of predicted health benefits based on income deprivation scores with the corresponding 95% confidence intervals. Panel (c) plots the quadratic prediction of predicted health benefits based on the share of patients aged 65 or more in a given grid with the corresponding 95% confidence intervals.

Productivity benefits. Panel (b) in Figure 5.10 builds upon the same simulated reduction shock in $\text{PM}_{2.5}$ from Eq. 5.8 and presents findings on the distribution of productivity benefits per capita. These benefits reflect variations in GVA per capita attributed to the simulated reduction in $\text{PM}_{2.5}$ concentration. This graph highlights the presence of substantial spatial heterogeneity, with estimated economic advantages stemming from productivity gains spanning a wide spectrum, ranging from roughly £170 to over £1500 per capita. Panel (a) in Figure 5.9 illustrates the frequency distribution of the estimated productivity benefits. Overall, these findings show how the higher data granularity uncovers substantial spatial heterogeneity in pollution-driven productivity effects that would be overlooked in nationwide or regional studies (Dechezleprêtre et al., 2019).

Leveraging a quadratic prediction again ($\Delta \widehat{y}_{iC}^{PM_{2.5}} = \hat{\beta}_0 + \hat{\beta}_1 x + \hat{\beta}_2 x^2$), Panel (b) shows that predicted economic benefits arising from productivity gains display a non-linear increase alongside the rise in median weekly income in the area, which implies that the direct economic advantages associated with pollution reduction are regressively concentrated in more affluent areas of the country.

To gain further insights into this pattern, Panel (c) in Figure 5.9 provides a ranking of LAUs in England, listing those with the highest estimated productivity benefits

Figure 5.9: Distribution of productivity benefits



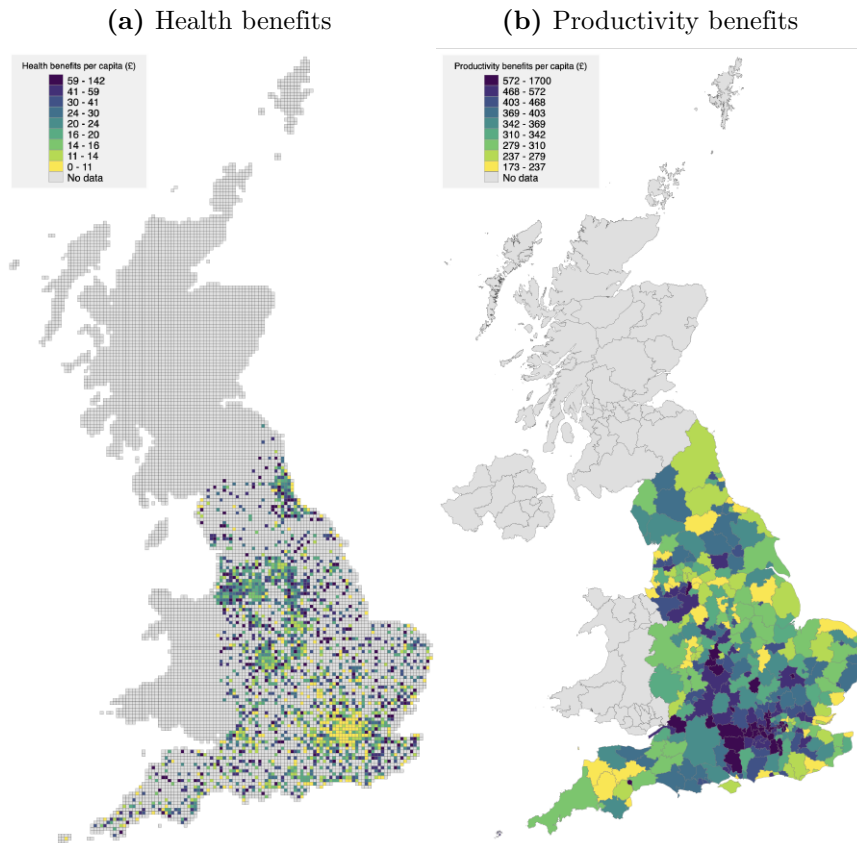
Notes: Panel (a) displays the frequency distribution of predicted productivity benefits from a $\mu\text{g}/\text{m}^3$ $\text{PM}_{2.5}$ reduction in the second-stage estimation (cf. Eq. 5.8). Panel (b) reports the quadratic prediction of predicted health benefits based on median weekly incomes in a given LAU with the corresponding 95% confidence intervals. Panel (c) plots a ranking of the 10 highest estimated productivity effects across areas in England, excluding the financial district (i.e., the City of London) due to its outlier status. A map plotting the distribution of LAUs in the London region can be found in the Appendix (cf. Figure 5.A14).

in descending order. Notably, the most significant effects are primarily driven by the London area, with Westminster (home to the United Kingdom’s Houses of Parliament) occupying the top position in the ranking. This ranking, however, crucially excludes the financial district of the City of London, where the estimated productivity benefits translate into a sizable annual figure of more than 100 thousand pounds per capita, due to the concentrated presence of numerous high-GVA financial service activities in this area. Collectively, these findings suggest that when considering productivity effects alone, prioritizing pollution reduction efforts in urban areas would likely prove more efficient.

5.6 Policy implications

Spatial disparities and distributional trade-offs. Overall, the findings presented in Section 5.5.2 provide important insights into the complex trade-offs that policymakers must navigate when designing pollution reduction policies. The observed spatial disparities in productivity benefits, as highlighted in Section 5.5.2, underscore the pivotal role of local economic factors in shaping the outcomes of these policies. High-earning urban centers, exemplified by London, are poised to

Figure 5.10: Distribution of benefits of a $\mu\text{g}/\text{m}^3$ $\text{PM}_{2.5}$ reduction



Notes: The figures above plot the distribution of benefits across England induced by a $1 \mu\text{g}/\text{m}^3$ nationwide reduction in the annual concentration of $\text{PM}_{2.5}$ based on predicted values from Eq. 5.8. The resolution is contingent on the highest available spatial granularity. Panel (a) refers to health benefits due to reduced pharmaceutical spending, whereas Panel (b) refers to productivity gains as proxied by GVA per capita. In the legends, displayed values are capped at the 99th percentile of the value distribution to ensure a clearer representation of the data.

benefit significantly from pollution reduction initiatives due to their concentration of economic activities and high productivity levels. This urban advantage implies a disproportionate positive direct economic impact of nationwide air quality improvements in such areas.

It follows that failure to acknowledge spatial heterogeneity may result in overestimating or underestimating the economic gains associated with local pollution reduction measures, which can subsequently misguide spatial resource allocation and policy prioritization. Nevertheless, how such productivity gains are distributed among socio-economic strata within different regions is ultimately mediated by many factors, such as national fiscal redistributive policies or the underlying distribution of capital ownership and wages, which I do not directly observe in my sample.

While urban regions may experience more substantial economic gains, the health benefits stemming from an overall pollution reduction exhibit a distinct distribution. These health advantages, as discussed, tend to favor individuals and communities that are economically disadvantaged, particularly among the elderly population. This pro-poor characteristic underscores the potential for air quality improvements to alleviate existing health inequalities (cf. Section 5.3) and improve the health status of vulnerable segments of the population with less adaptive capacity to environmental damage. Neglecting to consider this distributional dimension has the potential to introduce bias in WTP estimates by not accurately reflecting the preferences and health improvements experienced by various socioeconomic and demographic groups (cf. Bento et al. 2015).

It follows that accounting for spatial heterogeneity in pollution-reduction benefits has the potential to enhance public welfare through the design of more efficient regulation and more precise benefit-cost analyses that can accurately account for the distribution of different types of air pollution benefits across the population (Muller and Mendelsohn, 2009). Therefore, the challenge for policymakers lies in finding a balance between the distribution of different types of societal benefits to design effective and equitable local pollution reduction strategies. An approach that has been discussed in the literature entails distributional weights, which can be chosen to reflect societal inequality aversion in line with ethical considerations (Adler, 2016; Drupp et al., 2021).²⁸

Co-benefits of climate policy. Finally, my findings also carry important policy implications concerning the quantification and distribution of co-benefits of climate change mitigation, which may foster air pollution reductions as a byproduct of decreased reliance on fossil fuels (e.g., Wagner and De Preux 2016; Vandyck et al.

²⁸The principle of equity weighting has already been enshrined in the UK official guidelines on cost-benefit analysis (Treasury, 2016). Distributional weights are also integrated into Germany's approach for estimating climate change damages (UBA, 2019). These two real-world instances exemplify the use of equity and distributional weights in economic and environmental policy assessment.

2020; Basaglia et al. 2023). Currently, the public discourse typically centers around the costs of climate policies for current consumers, which tend to be regressive in developed countries (Sterner, 2012a; Klenert and Mattauch, 2016). Applied modeling studies indicate that overall distributional outcomes, additionally incorporating source-side impacts on wages and capital incomes (Fullerton and Muehlegger 2019), exhibit a reduced regressive nature or sometimes even demonstrate a progressive trend (Goulder et al., 2019). Yet, this still ignores the distribution of health benefits, which I demonstrate to be potentially sizable and distributed pro-poor. Hence, the overall distributional burden of climate policy might be less prominent than previously assumed on poorer households, who typically endure higher exposure to air pollutants. Relatedly, this suggests that policy-driven air quality improvements may offer a potential avenue for mitigating environmental inequities (e.g., Cushing et al., 2018; Hernandez-Cortes and Meng, 2023). Considering and effectively communicating these health co-benefits, which provide immediate benefits to individuals impacted by policy costs, could significantly contribute to gaining support for stricter climate policies (Löschel et al., 2021), which face considerable political and public resistance (Carattini et al., 2019; Douenne and Fabre, 2022).

5.7 Conclusion

There is currently limited empirical evidence to inform policymaking on which nationwide environmental policies would be socially desirable. This paper provides the first quasi-experimental estimate of the nationwide cost of air pollution that jointly accounts for both *health* and *productivity* costs and examines their spatial heterogeneity.

To address spurious correlation concerns in the distribution of air pollution, I exploit atmospheric temperature inversions as a source of exogenous dynamic variation in the spatial concentration of air pollution across England. I find that a plausibly

exogenous $1 \mu\text{g}/\text{m}^3$ pollution shock causes significant increases in pharmaceutical expenditures as well as a reduction in GVA per capita over the same year.

Specifically, I find that a $1 \mu\text{g}/\text{m}^3$ increase in the annual average concentration of $\text{PM}_{2.5}$ leads to an increment in expenditures on pharmaceuticals of 32.7% (or £1.2 billion annually) and a reduction in gross value added (GVA) per capita of 1.6% (or around £13 billion annually). Taken together, these estimates amount to a total per capita damage of approximately £445. In comparison, this value significantly exceeds current cost estimates used by the UK government to assess nationwide optimal policy stringency and previous findings in the literature that focused exclusively on either health or productivity costs (cf. Pimpin et al., 2018; Dechezleprêtre et al., 2019; UK-AIR, 2023).

Furthermore, I leverage my causal estimates to simulate counterfactual reductions of $1 \mu\text{g}/\text{m}^3$ in $\text{PM}_{2.5}$ and investigate the individual distributions of both health and productivity benefits and their correlation with demographic and socio-economic factors. While health benefits tend to be larger for the elderly and progressively distributed along the income distribution, productivity gains tend to be regressive and concentrated in large urban areas.

These findings carry four key policy implications. First, accurately measuring both the health and productivity impacts of pollution is pivotal to providing comprehensive damage estimates that inform the calibration of environmental regulations that optimize both public health and economic growth. Failure to capture either cost dimension likely results in underestimating the societal benefits associated with pollution reduction measures, which can subsequently misguide spatial resource allocation and policy prioritization.

Second, the empirical evidence presented in this study suggests that neglecting local variations in the demand for air quality can introduce sizable bias into WTP estimates, potentially compromising the accuracy and validity of cost-benefit assessments of optimal environmental policy stringency. Policymakers face the challenge of balancing

different types of societal benefits to design efficient and equitable local pollution reduction strategies. This may involve considering distributional weights for economic and environmental policy assessment, as proposed in the literature (e.g., Adler 2016).


Third, I provide complementary causal insights linking pollution-driven productivity losses to work absenteeism. This finding carries further distributional implications as a reduced capacity to work due to pollution-driven morbidity may trap low-income individuals in an illness-poverty cycle (Ketcham et al., 2019), owing to their typically disproportionate exposure (e.g., Samoli et al. 2019; Colmer et al. 2020). Reducing air pollution could thus further serve as a means to alleviate income disparities that might endure across successive generations (e.g., Isen et al. 2017; Durlauf and Seshadri 2018; Chetty and Hendren 2018).

Finally, while prior economics literature on the health costs of pollution primarily focuses on its effects on mortality (e.g., Deryugina et al. 2019), my findings show that disregarding the economic burden of morbidity impacts would significantly underestimate the overall economic costs of pollution. Additionally, I show that these impacts extend beyond respiratory and cardiovascular diseases typically considered in previous studies (e.g., Deschenes et al. 2017).

5.A Descriptive evidence


5.A.1 Yearly spatial distribution of inversion events

Figure 5.A1: Distribution of inversion events (950 hPa) in 2012




Notes: The figure plots the distribution of inversion events below 950 hPa computed on a 6-hour frequency in 2012. The spatial resolution corresponds to 5km x 5km grids following the BNG reference system.

Figure 5.A2: Distribution of inversion events (950 hPa) in 2013




Notes: The figure plots the distribution of inversion events below 950 hPa computed on a 6-hour frequency in 2013. The spatial resolution corresponds to 5km x 5km grids following the BNG reference system.

Figure 5.A3: Distribution of inversion events (950 hPa) in 2014




Notes: The figure plots the distribution of inversion events below 950 hPa computed on a 6-hour frequency in 2014. The spatial resolution corresponds to 5km x 5km grids following the BNG reference system.

Figure 5.A4: Distribution of inversion events (950 hPa) in 2015




Notes: The figure plots the distribution of inversion events below 950 hPa computed on a 6-hour frequency in 2015. The spatial resolution corresponds to 5km x 5km grids following the BNG reference system.

Figure 5.A5: Distribution of inversion events (950 hPa) in 2016




Notes: The figure plots the distribution of inversion events below 950 hPa computed on a 6-hour frequency in 2016. The spatial resolution corresponds to 5km x 5km grids following the BNG reference system.

Figure 5.A6: Distribution of inversion events (950 hPa) in 2017



Notes: The figure plots the distribution of inversion events below 950 hPa computed on a 6-hour frequency in 2017. The spatial resolution corresponds to 5km x 5km grids following the BNG reference system.


Figure 5.A7: Distribution of inversion events (950 hPa) in 2018



Notes: The figure plots the distribution of inversion events below 950 hPa computed on a 6-hour frequency in 2018. The spatial resolution corresponds to 5km x 5km grids following the BNG reference system.

5.A.2 Temporal distribution of inversion episodes

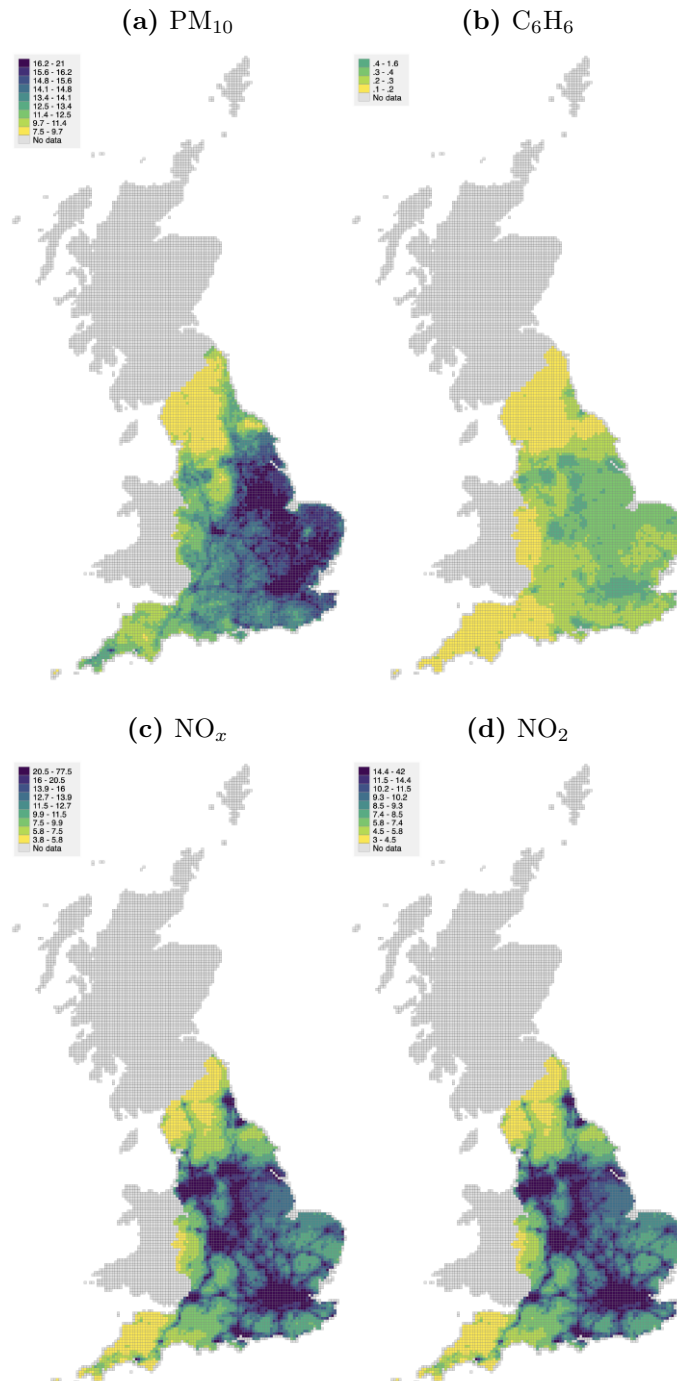
Figure 5.A8: Frequency monthly counts of inversion episodes



Notes: Figures plot a matrix of frequency counts of inversion events occurring each month below a given pressure level (hPa) indicated by the figure heading.

5.A.3 Concentrations of PM_{10} , NO_x , NO_2 and C_6H_6 in $\mu\text{g}/\text{m}^3$ across England

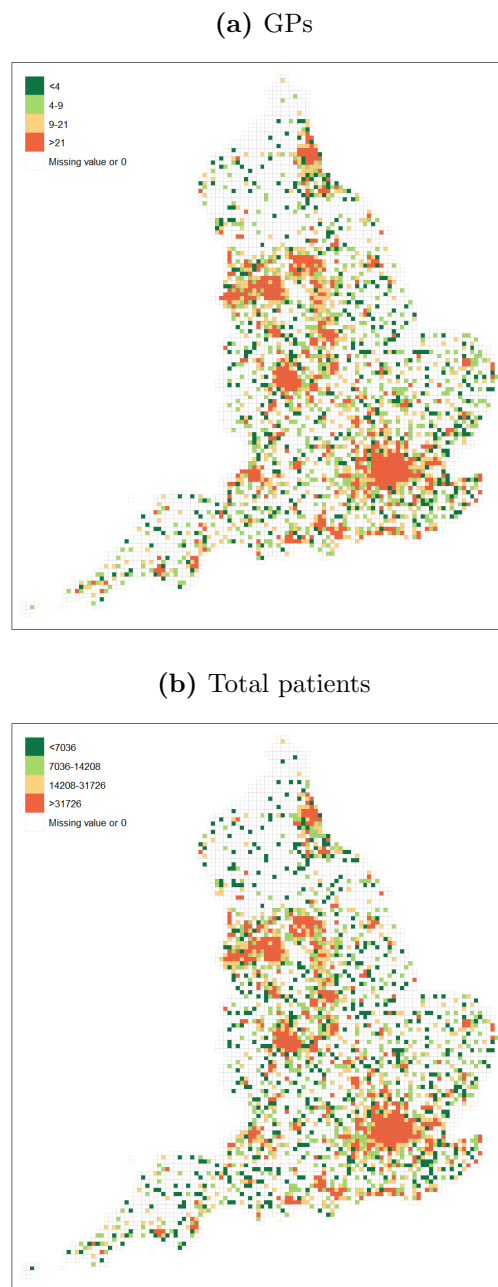
Figure 5.A9: Distribution of air pollutants concentrations across England



Notes: The figures above plot the distribution of pollution concentrations in England averaged from 2012 and 2018 across 5km x 5km grids following the BNG reference system. Pollution concentration refers to average background annual average concentrations in $\mu\text{g}/\text{m}^3$ across 5km x 5km grids based on data reported by DEFRA. The legend displays concentration defined in $\mu\text{g}/\text{m}^3$.

5.A.4 Prescription data: descriptive statistics

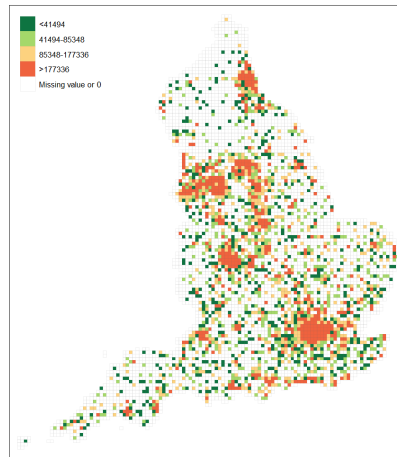
Figure 5.A10: Spatial distribution of GPs and patients across England



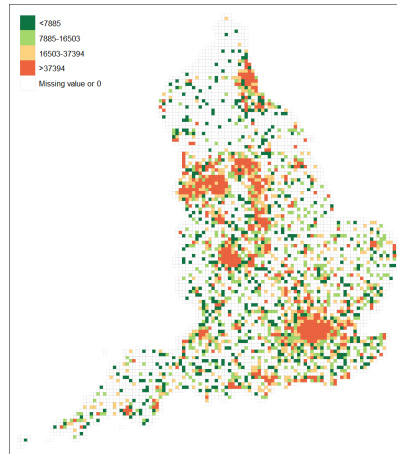
Notes: Figures plot the distribution of the average number of general practitioners (GPs) and the total number of patients reported by each registered practice to NHS Digital aggregated at a 5km x 5km grid level.

Figure 5.A11: Spatial distribution of prescription items by therapeutic section

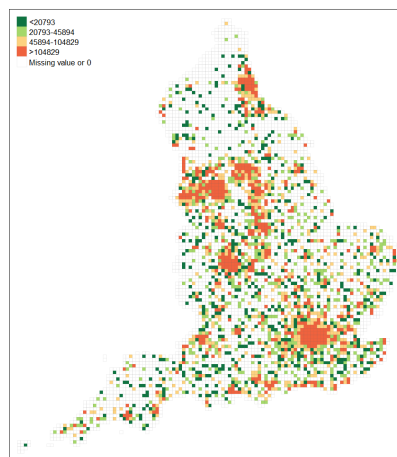
(a) Cardiovascular



(b) Respiratory



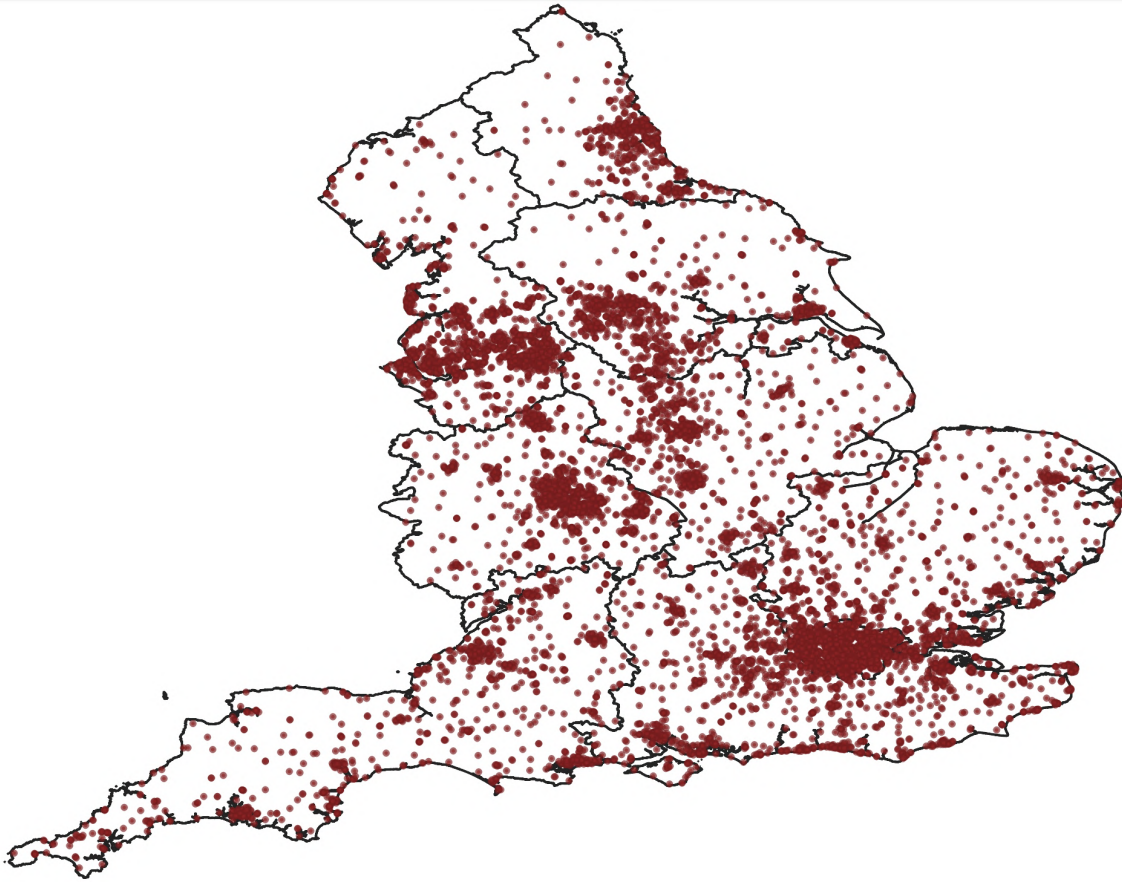
(c) Central nervous



Notes: Figures plot the distribution of the average number of yearly prescription items by therapeutic section issued by each registered practice to NHS Digital aggregated at a 5km x 5km grid level.

5.A.5 Mapping of GP practices across England by postcodes

Figure 5.A12: Spatial distribution of GP practices in English regions



Notes: The figure above displays the geographic distribution of GP practices in England, utilizing Geographic Information System (GIS) tools.

5.A.6 Clinical Commissioning groups in England

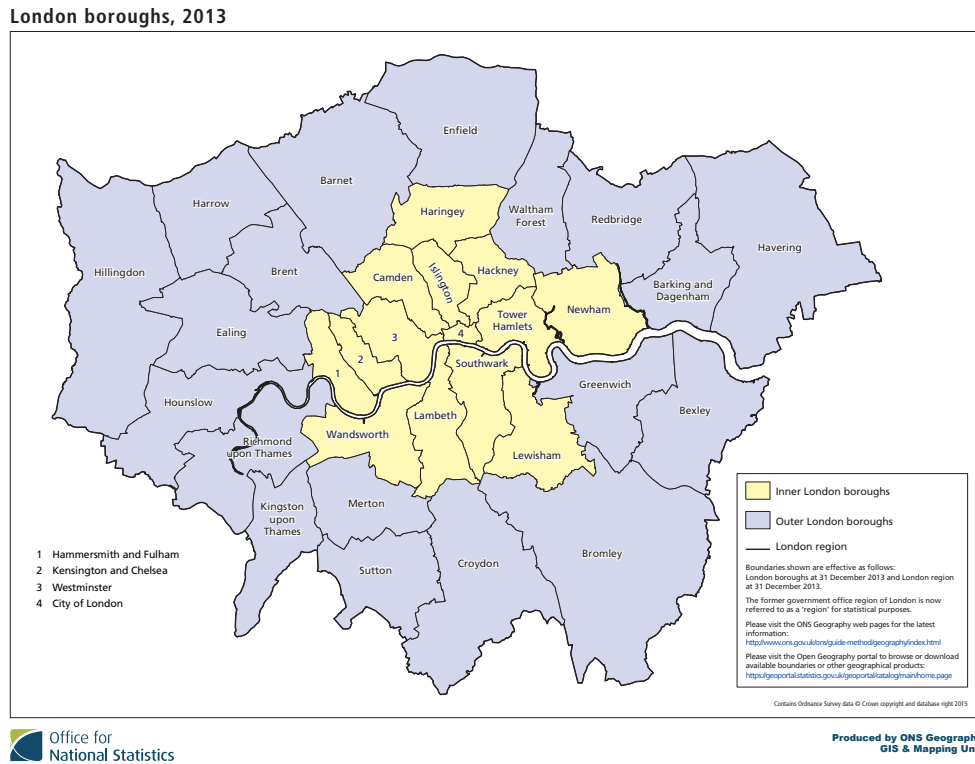
Figure 5.A13: Spatial distribution of Clinical Commissioning groups (CCGs) in England



Notes: The figure above plots the distribution of Clinical Commissioning groups (CCGs) in England.

5.A.7 LAUs in London

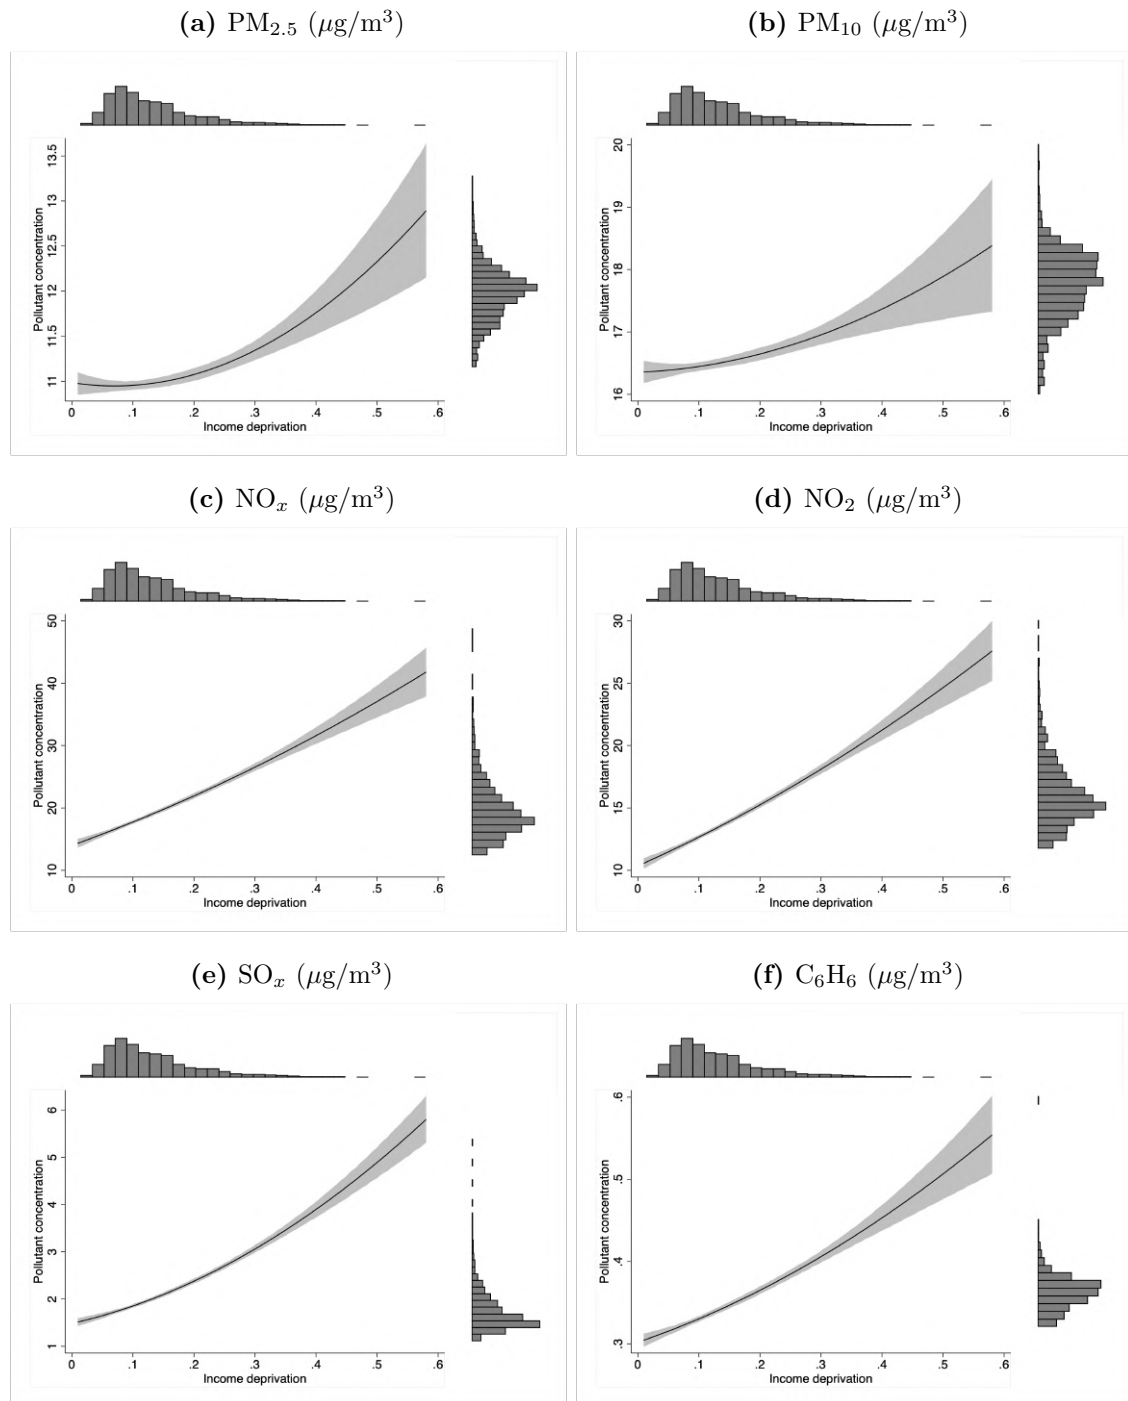
Figure 5.A14: Spatial distribution of LAUs in the London region



Notes: The figure is sourced from the Office for National Statistics (ONS) website and plots the distribution of LAUs (which correspond to boroughs) in the London region.

5.A.8 Pre-existing pollution exposure inequalities along the income distribution

Figure 5.A15: Quadratic predictions of pollution concentrations in 2011 based on income deprivation scores




Notes: Quadratic predictions of the pollutant concentration (as indicated in the heading) with the 95% confidence interval based on income deprivation scores with histograms plotting the frequency distributions of each variable indicated in the *x*- and *y*-axis. See Section 5.3 for more details.

5.B Robustness analyses


5.B.1 Effect of air pollution by therapeutic section

Figure 5.A16: Results by therapeutic section


(a) Cardiovascular



(b) Respiratory



(c) Central nervous

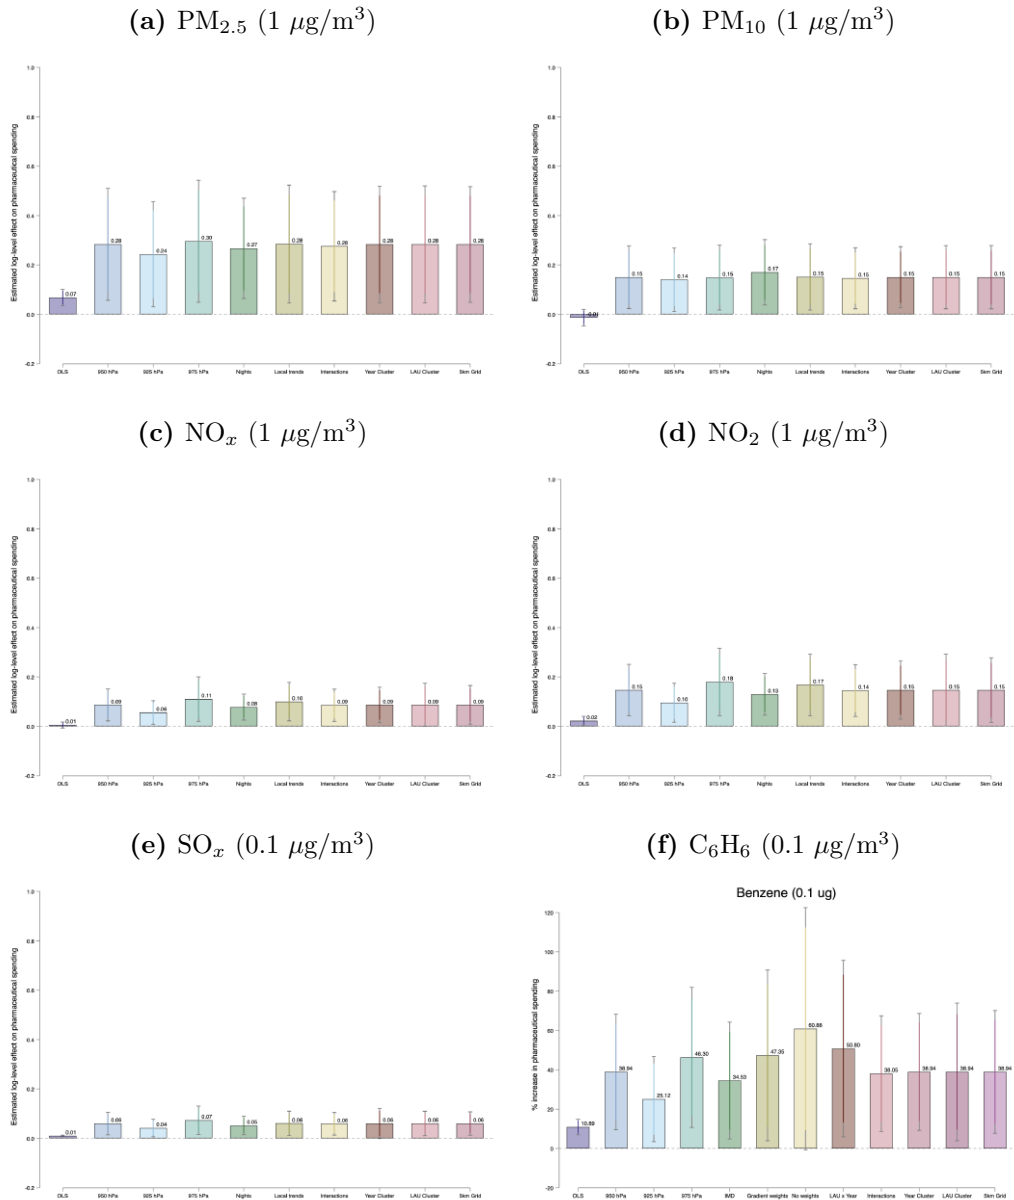


Notes: The figures above plot the estimated effect from the second stage (Eq. 5.7) of the 2SLS estimation described in Section 5.4 by BNF therapeutic section. The bar charts display mean point estimates with overlaid lines representing confidence intervals, respectively delineating 95% confidence intervals in gray and 90% confidence intervals in color. Pharmaceutical expenditures are reported by each registered practice and have been aggregated at a 5km x 5km grid level. Pollution concentration refers to average background annual average concentrations in $\mu\text{g}/\text{m}^3$ across 5km x 5km grids based on data reported by DEFRA. Weather controls include mean ground-level temperature ($^{\circ}\text{C}$), precipitation (mm), wind speed (knots), and relative humidity (%). All regressions additionally control for the total number of GPs and are weighted by the total number of patients in each grid.

5.B.2 Additional model specifications by air pollutant

Effects on pharmaceutical spending

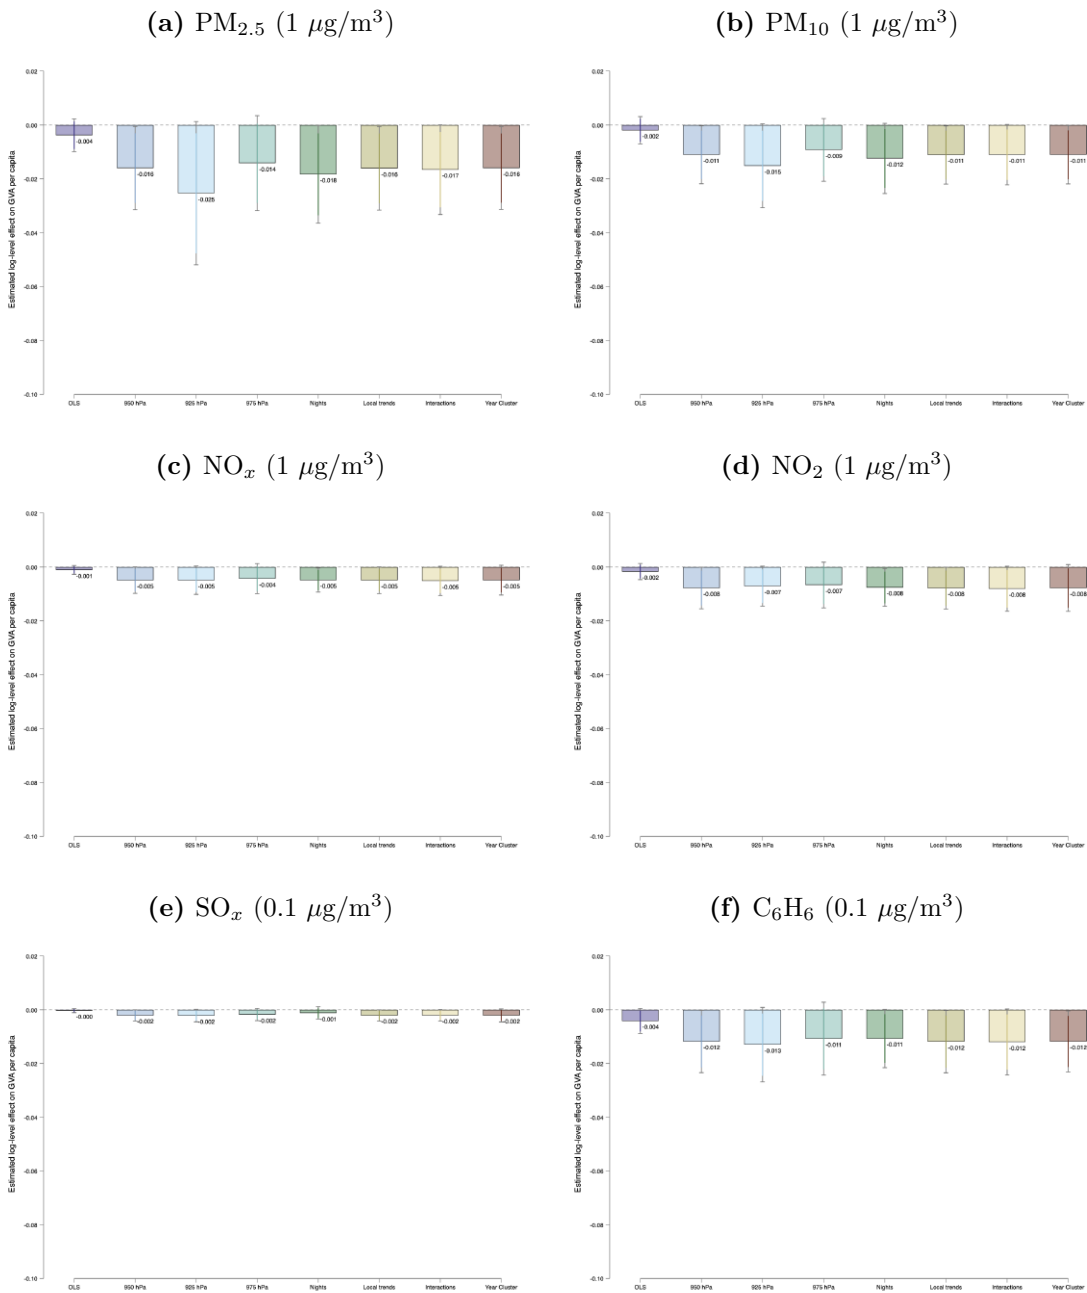
Figure 5.A17: Results by by air pollutant



Notes: The figures above plot the estimated effect from the second stage (Eq. 5.7) of the 2SLS estimation described in Section 5.4. The bar charts display mean point estimates with overlaid lines representing confidence intervals, respectively delineating 95% confidence intervals in gray and 90% confidence intervals in color. Pharmaceutical expenditures are reported by each registered practice and have been aggregated at a 5km x 5km grid level. Pollution concentration refers to average background annual average concentrations in $\mu\text{g}/\text{m}^3$ across 5km x 5km grids based on data reported by DEFRA. Weather controls include mean ground-level temperature ($^{\circ}\text{C}$), precipitation (mm), wind speed (knots), and relative humidity (%). All regressions additionally control for the total number of GPs and are weighted by the total number of patients in each grid. 950 hPa represents the baseline specification presented in the paper. Details on alternative specifications can be found in Section 5.5.1. *LAU cluster* and *5km Grid* report results for the baseline model with clustered standard errors at the LAU and 5km grid level, respectively.

Effects on GVA per capita

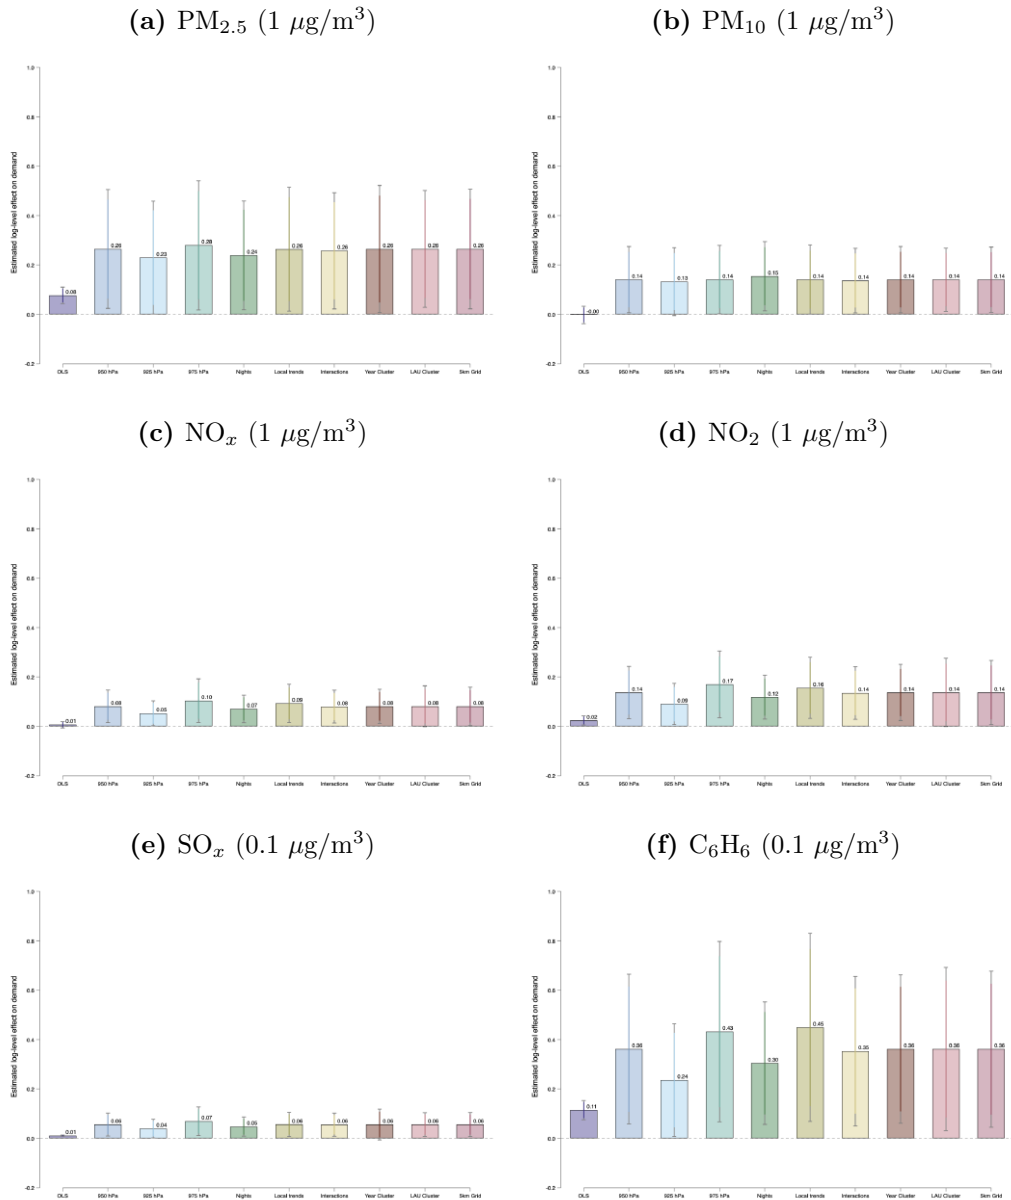
Figure 5.A18: Results by air pollutant



Notes: The figures above plot the estimated effect from the second stage (Eq. 5.7) of the 2SLS estimation described in Section 5.4. The bar charts display mean point estimates with overlaid lines representing confidence intervals, respectively delineating 95% confidence intervals in gray and 90% confidence intervals in color. The outcome is GVA per capita which represents the ratio of local GVA, sourced from the UK ONS, divided by the total population in the LAU. Pollution concentration refers to average background annual average concentrations in $\mu\text{g}/\text{m}^3$ across LAUs based on data reported by DEFRA. Thermal inversions are defined as a positive upward temperature gradient from the surface, and calculated on a 6-hour frequency using data from the ECMWF. Weather controls include mean ground-level temperature ($^{\circ}\text{C}$), precipitation (mm), wind speed (knots), and relative humidity (%). All regressions are weighted by the total population in each LAU. 950 hPa represents the baseline specification presented in the paper. Details on alternative specifications can be found in Section 5.5.1.

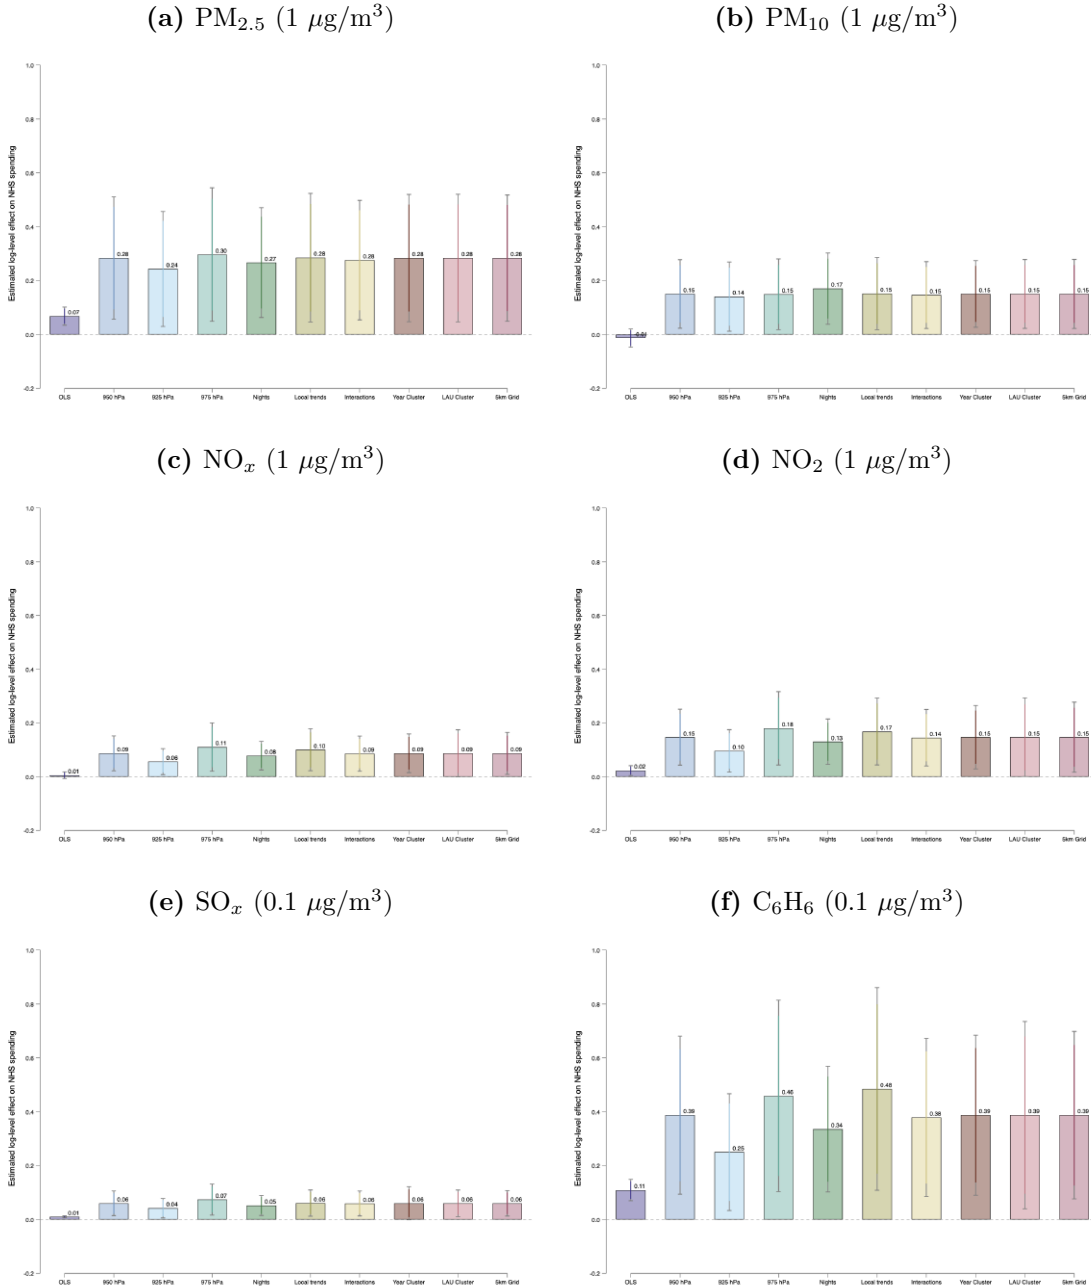
5.B.3 Effects on prescriptions items and total costs to the NHS

Figure 5.A19: Effects of air pollutants on prescriptions items



Notes: The figures above plot the estimated effect from the second stage (Eq. 5.7) of the 2SLS estimation described in Section 5.4. The bar charts display mean point estimates with overlaid lines representing confidence intervals, respectively delineating 95% confidence intervals in gray and 90% confidence intervals in color. The outcome is the number of prescription items reported by each registered practice and has been aggregated at a 5km x 5km grid level. Pollution concentration refers to average background annual average concentrations in $\mu\text{g}/\text{m}^3$ across 5km x 5km grids based on data reported by DEFRA. Weather controls include mean ground-level temperature ($^{\circ}\text{C}$), precipitation (mm), wind speed (knots), and relative humidity (%). All regressions additionally control for the total number of GPs and are weighted by the total number of patients in each grid. 950 hPa represents the baseline specification presented in the paper. Details on alternative specifications can be found in Section 5.5.1. *LAU cluster* and *5km Grid* report results for the baseline model with clustered standard errors at the LAU and 5km grid level, respectively.


Figure 5.A20: Effects of air pollutants on total costs to the NHS



Notes: The figures above plot the estimated effect from the second stage (Eq. 5.7) of the 2SLS estimation described in Section 5.4. The bar charts display mean point estimates with overlaid lines representing confidence intervals, respectively delineating 95% confidence intervals in gray and 90% confidence intervals in color. The outcome is the estimated overall cost to the NHS (e.g., accounting for discounts from suppliers and deducting patients' out-of-pocket expenses) reported by each registered practice and has been aggregated at a 5km x 5km grid level. Pollution concentration refers to average background annual average concentrations in $\mu g/m^3$ across 5km x 5km grids based on data reported by DEFRA. Weather controls include mean ground-level temperature ($^{\circ}C$), precipitation (mm), wind speed (knots), and relative humidity (%). All regressions additionally control for the total number of GPs and are weighted by the total number of patients in each grid. 950 hPa represents the baseline specification presented in the paper. Details on alternative specifications can be found in Section 5.5.1. *LAU cluster* and *5km Grid* report results for the baseline model with clustered standard errors at the LAU and 5km grid level, respectively.

5.B.4 Reduced form (RF) with leads and lags


Figure 5.A21: Dynamic RF effects on prescriptions items



Notes: The figures above plot the estimated 99% confidence intervals from the reduced form (Eq. 5.5) estimation described in Section 5.4. Specifically, the model regresses pharmaceutical expenditures on the frequency of inversion episodes and includes a set of leads and lags of inversion events (as indicated by the figure heading). Pharmaceutical expenditures are reported by each registered practice and have been aggregated at a 5km x 5km grid level. Weather controls include mean ground-level temperature ($^{\circ}\text{C}$), precipitation (mm), wind speed (knots), and relative humidity (%). All regressions additionally control for the total number of GPs and are weighted by the total number of patients in each grid.

5.B.5 LAUs as alternative unit of observation

Figure 5.A22: Effects of a $1 \mu\text{g}/\text{m}^3$ annual increase in $\text{PM}_{2.5}$ on pharmaceutical expenditures




Notes: The figures above plot the estimated effect from the second stage (Eq. 5.7) of the 2SLS estimation described in Section 5.4. The bar charts display mean point estimates with overlaid lines representing confidence intervals, respectively delineating 95% confidence intervals in gray and 90% confidence intervals in color. Pharmaceutical expenditures are reported by each registered practice and have been aggregated at the LAU level. $\text{PM}_{2.5}$ concentration refers to average background annual average concentrations in $\mu\text{g}/\text{m}^3$ across LAUs based on data reported by DEFRA. Weather controls include mean ground-level temperature ($^{\circ}\text{C}$), precipitation (mm), wind speed (knots), and relative humidity (%). All regressions additionally control for the total number of GPs and are weighted by the total number of patients in each grid. *Year effects* controls for year fixed effects and represents the baseline specification. *Regional trends* additionally absorb heterogeneous trends across English regions whereas *50km trends* account for trends at the 50km x 50km grid level.

5.C Changes in defensive behavior over time

The concluding section of the empirical analysis delves into the temporal evolution of defensive behavior and investigates the influence of shifts in pollution information on these changes. I begin by assessing whether revealed preference estimates vary over time: Figure 5.A23 plots the estimated second-stage coefficients on the effects of a $PM_{2.5}$ shock on the demand for prescription items across different time windows in the sample. The figure illustrates an upward trend in the estimated demand effects over time, indicating an average increase in defensive behavior responses to pollution increases.

Figure 5.A23: Effects of $PM_{2.5}$ on demand for pharmaceuticals over time.



Note: The figure plots the results of the 2SLS estimation framework outlined in Section 5.4 when splitting the sample before and after 2015. All estimations include yearly effects, fixed effects at the municipality level, and control for surface-level temperature, rainfall, wind speed, and relative humidity in a given grid. The time frame considered for each estimation is indicated in the x-axis. The two bars plot 90% confidence intervals.

Drawing on a growing number of economic studies leveraging newspaper data as source of variation in the salience of events (e.g., Li et al., 2014; Baker et al., 2016; Beach and Hanlon, 2023), I empirically evaluate the extent to which greater information provision in the media could explain this trend leveraging textual analysis of newspaper articles. Specifically, I construct a newspaper-based index to capture variation across time in the public discourse within England on the health implications of exposure to air pollution. The index reflects the frequency of print and online

articles in leading newspapers, namely *The Guardian*, *The Telegraph*, and *The Independent* that contain terms related to a discussion on the health effects of air pollution.²⁹

Exposure to pollution information through newspapers. First, I developed a search strategy to identify newspaper articles that discuss air pollution which contain keywords such as "particulate matter" or "aerosol particles" or "PM2.5" or "air pollution" and then narrow down the focus to articles that explicitly discuss the health effects of pollution exposure by additionally including terms such as "health" or "morbidity" or "mortality" or "sick*" "ill*" as well as terms related to the health domains more commonly linked to pollution, for instance, "respiratory" or "cardiovascular" or "nervous". scaled by newspaper-specific publishing trends to ensure that spikes in our index are not driven by newspaper-specific shocks. After scaling the raw counts, I standardize each newspaper's series, average across all papers, and normalize the resulting index to 100 over the period, following the same standardization and normalization procedure by Baker et al. (2016) to leverage newspaper data in an empirical setting. The complete search strategy and a detailed description of the steps undertaken to construct my newspaper index is discussed later in the Appendix.


Figure 5.A24 plots the evolution of the newspaper-based index over the period of time under investigation.³⁰ In general, there has been a consistent and substantial rise in newspaper coverage of air pollution, reflecting greater awareness of environmental topics (Schmidt et al., 2013) and mirroring the expanding scientific literature connecting air pollution exposure to an increasingly diverse range of health conditions (Romanello et al., 2021). In the remainder of this section, I will leverage variations

²⁹I restrict the analysis to the three largest newspapers retrievable from Factiva, as relying on a single source provides consistent, comparable, and thus more robust counts.

³⁰Particulate matter pollution has dominated the mainstream public discussion on the health effects of air pollution over the last decades. It is therefore difficult to disentangle articles that deal with the effects of specific air pollutants other than fine particles. For this reason, variations in the newspaper-based index should be interpreted bearing in mind that most of the increase in pollution information provision is disproportionately linked to fine particles.

in the newspaper-based index to empirically investigate how changes in pollution information available to the wider public affect defensive investment responses.

Figure 5.A24: Evolution of the newspaper index over time. 2012–2018.



Note: Based on yearly series from 2012 to 2018.

Source: Authors' own calculations based on newspaper articles from Factiva. The following search strategy was applied to identify newspaper articles used to compute this index: (particulate or "particulate matter" or "suspended particulate matter" or "aerosol particles" or "PM2.5" or "particulate emission*" or "air pollution") and (health or morbidity or mortality or sick* ill* or death or harmful or respiratory or cardiovascular or nervous)

The role of pollution information in the news. Table 5.A1 presents the estimated effects attributed to changes in the newspaper index. I amend my reduced form equation (see Eq. 5.5) by including two additional coefficients: (1) an interaction between the newspaper index and the concentration of $PM_{2.5}$, and (2) an interaction between the squared newspaper index and $PM_{2.5}$. The quadratic form facilitates the identification of any non-linear relationships between information provision and defensive behavior. By having an interaction term, my identification strategy captures the additional effect on demand for pharmaceuticals - at a given $PM_{2.5}$ concentration - due to additional pollution information in the media. Crucially, the interactive term allows for incorporating fixed effects without dropping the newspaper index due to collinearity with time effects.

Comparing columns (1) and (2), the estimated coefficients reveal that an increment in pollution exposure as proxied by the newspaper-based index correlates with an increase in pharmaceutical demand, yet at a diminishing rate. To gauge the

Table 5.A1: Exposure to pollution information in the news and demand for pharmaceuticals

	Items (log)	Items (log)
PM _{2.5} ($\mu\text{g}/\text{m}^3$) \times News Index (log)	0.0163*** (0.00338)	0.0638*** (0.00755)
PM _{2.5} ($\mu\text{g}/\text{m}^3$) \times News Index (log) ²		-0.00969*** (0.00120)
Controls	✓	✓
N	41830	41830

Notes: Tables show the coefficients estimated from the reduced form of the IV approach, where the frequency of thermal inversions is used as an instrument for pollution concentrations. GP practices' location has been geocoded using GIS tools and assigned to 5km x 5km grids in accordance with the Ordnance Survey National Grid reference system. Total items prescribed are reported by each registered practice to the NHS Business Services Authority (BSA) and have been aggregated at a 5km x 5km grid level. Pollutants' concentration refers to background annual average concentrations on a 5km x 5km grid as reported by the British Department for Environment, Food and Rural Affairs. Thermal inversions are defined as a positive upward temperature gradient from the surface, and calculated on a 6-hours frequency using data from the ECMWF. Climate controls are retrieved on a 3-hours frequency from the UK Met Office and aggregated at the yearly level. Coefficients are weighted by the population of each LAU so as to be representative of the average citizen in England rather than the average LAU.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

magnitude of these coefficients, let's consider the median PM_{2.5} concentration in the estimation sample (9.74) and the coefficients from column (2), a 5% rise in the index is linked to an approximate 0.75% increase in demand for prescription items. Taken together, these findings indicate that the provision of information regarding the effects of pollution plays a significant role in shaping revealed preferences in line with recent findings by Ito and Zhang (2020). This implies that when access to information is limited, consumers' actual willingness to pay for air quality enhancements may be underestimated. Unequal access to information across socioeconomic strata is therefore likely to play an important role in generating spatial disparities in defensive behavior. The extent of unequal access to information across socioeconomic strata is therefore likely to play an important role in generating spatial disparities in defensive behavior (cf., Ramírez et al., 2019; Hausman and Stolper, 2021).

5.D Computing the newspaper-based index

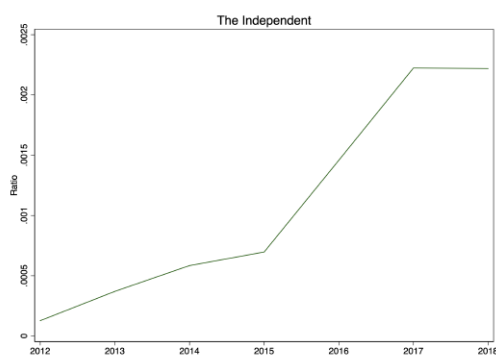
Search Strategy: (particulate or "particulate matter" or "suspended particulate matter" or "aerosol particles" or "PM2.5" or "particulate emission*" or "air pollution") and (health or morbidity or mortality or sick* ill* or death or harmful or respiratory or cardiovascular or nervous)

To construct my newspaper-based index I limited the search to leading newspapers to ensure the quality of the underlying articles and avoid including newspapers that only exceptionally report on the topic, spuriously creating huge volatility over time. Here, I focus on the Independent, the Telegraph, and the Guardian, which are the largest newspapers in the country covered by Factiva. For each newspaper, I separately downloaded the annual count of articles that are picked up by my search strategy as well as the total number of articles published by the outlet. To account for potential trends in publishing over the years, I start by computing a simple newspaper-specific ratio of articles matching my search strategy over the total article count by newspaper. A challenge with these raw article ratios is that the number of articles varies a lot across newspapers and time, making it difficult to simply average the ratios across several newspapers in a given country. I, therefore, apply the standardization approach of Baker et al. (2016) to obtain my newspaper index. I begin with the simple ratio of articles on climate policy uncertainty divided by the total article counts for each newspaper (see Figure 5.A25), and then divide this ratio by the newspaper-specific standard deviation across all years.

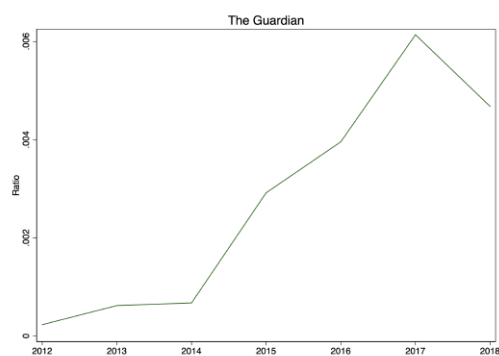
This creates a newspaper-specific time series with a unit standard deviation across the entire time interval, which ensures that the volatility of the index is not driven by the higher volatility of a particular newspaper. I then average these standardized series across all newspapers within each country by year. Lastly, I normalize the yearly series to a mean of 100 over the time interval.

Figure 5.A25: Ratio of matching articles over total articles


(a) The Independent



(b) The Guardian



(c) The Telegraph



Notes: Based on yearly series from 2012 to 2018.

Source: Authors' own calculations based on newspaper articles from Factiva.

Bibliography

- Abadie, A. (2021). Using synthetic controls: Feasibility, data requirements, and methodological aspects. *Journal of Economic Literature*, 59(2):391–425.
- Abadie, A., Diamond, A., and Hainmueller, J. (2010). Synthetic control methods for comparative case studies: Estimating the effect of california’s tobacco control program. *Journal of the American statistical Association*, 105(490):493–505.
- Abadie, A., Diamond, A., and Hainmueller, J. (2015). Comparative politics and the synthetic control method. *American Journal of Political Science*, 59(2):495–510.
- Abadie, A. and Gardeazabal, J. (2003). The economic costs of conflict: A case study of the basque country. *American economic review*, 93(1):113–132.
- Abadie, A. and L’hour, J. (2021). A penalized synthetic control estimator for disaggregated data. *Journal of the American Statistical Association*, 116(536):1817–1834.
- Abbring, J. H. and Van den Berg, G. J. (2003). The nonparametric identification of treatment effects in duration models. *Econometrica*, 71(5):1491–1517.
- Abrell, J., Kosch, M., and Rausch, S. (2022). How effective is carbon pricing? a machine learning approach to policy evaluation. *Journal of Environmental Economics and Management*, 112:102589.
- ACEA (2022). European automobile manufacturers’ association tax guide.

- Acemoglu, D., Aghion, P., Bursztyn, L., and Hemous, D. (2012). The environment and directed technical change. *American Economic Review*, 102(1):131–166.
- Adhvaryu, A., Kala, N., and Nyshadham, A. (2022). Management and shocks to worker productivity. *Journal of Political Economy*, 130(1):1–47.
- Adler, M. D. (2016). Benefit–cost analysis and distributional weights: An overview. *Review of Environmental Economics and Policy*.
- Aghion, P., Dechezleprêtre, A., Hemous, D., Martin, R., and Van Reenen, J. (2016). Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy*, 124(1):1–51.
- Aguilar-Gomez, S., Dwyer, H., Graff Zivin, J., and Neidell, M. (2022). This is air: The “nonhealth” effects of air pollution. *Annual Review of Resource Economics*, 14:403–425.
- Aldy, J. E. and Pizer, W. A. (2015). The competitiveness impacts of climate change mitigation policies. *Journal of the Association of Environmental and Resource Economists*, 2(4):565–595.
- Ale-Chilet, J., Chen, C., and Li, J. (2021). Colluding against environmental regulation. Technical report, CEPR Discussion Papers.
- Ambec, S., Cohen, M. A., Elgie, S., and Lanoie, P. (2013). The Porter hypothesis at 20: Can environmental regulation enhance innovation and competitiveness? *Review of Environmental Economics and Policy*, 7(1):2–22.
- Anderson, M. L. (2020). As the wind blows: The effects of long-term exposure to air pollution on mortality. *Journal of the European Economic Association*, 18(4):1886–1927.
- Anderson, M. L. and Auffhammer, M. (2014). Pounds that kill: The external costs of vehicle weight. *Review of Economic Studies*, 81(2):535–571.

- Andersson, J. J. (2019). Carbon taxes and CO₂ emissions: Sweden as a case study. *American Economic Journal: Economic Policy*, 11(4):1–30.
- Angrist, J. and Imbens, G. (1995). Identification and estimation of local average treatment effects.
- Annicchiarico, B., Carattini, S., Fischer, C., and Heutel, G. (2022). Business cycles and environmental policy: A primer. *Environmental and energy policy and the economy*, 3(1):221–253.
- Arceo, E., Hanna, R., and Oliva, P. (2016). Does the effect of pollution on infant mortality differ between developing and developed countries? evidence from Mexico City. *The Economic Journal*, 126(591):257–280.
- Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W., and Wager, S. (2021). Synthetic difference-in-differences. *American Economic Review*, 111(12):4088–4118.
- Arrow, K., Jorgenson, D., Krugman, P., Nordhaus, W., Solow, R., et al. (1997). The economists' statement on climate change.
- Auffhammer, M. and Kellogg, R. (2011). Clearing the air? The effects of gasoline content regulation on air quality. *American Economic Review*, 101(6):2687–2722.
- Avnery, S., Mauzerall, D. L., Liu, J., and Horowitz, L. W. (2011). Global crop yield reductions due to surface ozone exposure: 1. year 2000 crop production losses and economic damage. *Atmospheric Environment*, 45(13):2284–2296.
- Bach, S. (2009). Zehn Jahre ökologische Steuerreform: Finanzpolitisch erfolgreich, klimapolitisch halbherzig. *DIW Wochenbericht*, 76(14):218–227.
- Bagilet, V. and Zabrocki-Hallak, L. (2022). Why some acute health effects of air pollution could be inflated. Technical report, I4R Discussion Paper Series.
- Bai, J. (2009). Panel data models with interactive fixed effects. *Econometrica*, 77(4):1229–1279.

- Baker, R. E., Mahmud, A. S., Miller, I. F., Rajeev, M., Rasambainarivo, F., Rice, B. L., Takahashi, S., Tatem, A. J., Wagner, C. E., Wang, L.-F., et al. (2022). Infectious disease in an era of global change. *Nature Reviews Microbiology*, 20(4):193–205.
- Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring economic policy uncertainty. *The quarterly journal of economics*, 131(4):1593–1636.
- Banque de France (2019). Greening the financial system: The new frontier. Technical report, Banque de France.
- Banzhaf, S., Ma, L., and Timmins, C. (2019). Environmental justice: The economics of race, place, and pollution. *Journal of Economic Perspectives*, 33(1):185–208.
- Baranzini, A., Van den Bergh, J. C., Carattini, S., Howarth, R. B., Padilla, E., and Roca, J. (2017). Carbon pricing in climate policy: seven reasons, complementary instruments, and political economy considerations. *Wiley Interdisciplinary Reviews: Climate Change*, 8(4):e462.
- Barreca, A. I., Neidell, M., and Sanders, N. J. (2021). Long-run pollution exposure and mortality: Evidence from the acid rain program. *Journal of Public Economics*, 200:104440.
- Barrett, S. and Stavins, R. (2003). Increasing participation and compliance in international climate change agreements. *International Environmental Agreements*, 3:349–376.
- Basaglia, P., Behr, S. M., and Drupp, M. A. (2023). De-fueling externalities: Causal effects of fuel taxation and mediating mechanisms for delivering climate and health benefits. *CESifo Working Paper*, (No. 10508).
- Battiston, S., Mandel, A., Monasterolo, I., Schütze, F., and Visentin, G. (2017). A climate stress-test of the financial system. *Nature Climate Change*, 7(4):283–288.

- Baumol, W. J. and Oates, W. E. (1988). *The theory of environmental policy*. Cambridge university press.
- Bayer, P. and Aklin, M. (2020). The european union emissions trading system reduced co2 emissions despite low prices. *Proceedings of the National Academy of Sciences*, 117(16):8804–8812.
- Beach, B. and Hanlon, W. W. (2023). Culture and the historical fertility transition. *The Review of Economic Studies*, 90(4):1669–1700.
- Becker, G. S. (1965). A theory of the allocation of time. *The economic journal*, 75(299):493–517.
- Bento, A., Freedman, M., and Lang, C. (2015). Who benefits from environmental regulation? evidence from the clean air act amendments. *Review of Economics and Statistics*, 97(3):610–622.
- Bernanke, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *The Quarterly Journal of Economics*, 98(1):85–106.
- Beuermann, C. and Santarius, T. (2006). Ecological tax reform in Germany: Handling two hot potatoes at the same time. *Energy Policy*, 34(8):917–929.
- Binswanger, H. C. (1992). Von der umweltpolitik zur ökologischen wirtschaftspolitik. In *Verkehr und Umwelt*, pages 9–25. Springer.
- Bloom, N., Bond, S., and Van Reenen, J. (2007). Uncertainty and investment dynamics. *The review of economic studies*, 74(2):391–415.
- Bohn, S., Lofstrom, M., and Raphael, S. (2014). Did the 2007 legal Arizona workers act reduce the state’s unauthorized immigrant population? *Review of Economics and Statistics*, 96(2):258–269.
- Böhringer, C., Carbone, J. C., and Rutherford, T. F. (2012). Unilateral climate policy design: Efficiency and equity implications of alternative instruments to reduce carbon leakage. *Energy Economics*, 34, Supplement 2:S208–S217.

- Böhringer, C., Fischer, C., and Rosendahl, K. E. (2014). Cost-effective unilateral climate policy design: size matters. *Journal of Environmental Economics and Management*, 67(3):318–339.
- Borgschulte, M., Molitor, D., and Zou, E. Y. (2022). Air pollution and the labor market: Evidence from wildfire smoke. *Review of Economics and Statistics*, pages 1–46.
- Bosch, X. (2001). Spain opts to fuel health funds with petrol sales. *BMJ*, 323(7324):1271.
- Bovenberg, A. L. and Smulders, S. (1995). Environmental quality and pollution-augmenting technological change in a two-sector endogenous growth model. *Journal of Public Economics*, 57(3):369–391.
- Branger, F., Ponsard, J.-P., Sartor, O., and Sato, M. (2015). EU ETS, free allocations, and activity level thresholds: the devil lies in the details. *Journal of the Association of Environmental and Resource Economists*, 2(3):401–437.
- Brimblecombe, P. (2003). *The effects of air pollution on the built environment*, volume 2. World Scientific.
- Brook, R. D., Rajagopalan, S., Pope III, C. A., Brook, J. R., Bhatnagar, A., Diez-Roux, A. V., Holguin, F., Hong, Y., Luepker, R. V., Mittleman, M. A., et al. (2010). Particulate matter air pollution and cardiovascular disease: an update to the scientific statement from the american heart association. *Circulation*, 121(21):2331–2378.
- Bundesgesetzblatt I, S. 2432. (1999). Gesetz zur fortführung der ökologischen steuerreform.
- Bundesgesetzblatt I, S. 4602 (2002). Gesetz zur fortentwicklung der ökologischen steuerreform.
- Bundesgesetzblatt I, S.378 (1999). Gesetz zum einstieg in die ökologische steuerreform.

- Bundesministerium der Finanzen (2014). Entwicklung der Energie-(vormals Mineralöl-) und Stromsteuersätze in der Bundesrepublik Deutschland. Bonn.
- Calel, R. (2020). Adopt or innovate: Understanding technological responses to cap-and-trade. *American Economic Journal: Economic Policy*, 12(3):170–201.
- Calel, R. and Dechezlepretre, A. (2016). Environmental policy and directed technological change: evidence from the european carbon market. *Review of economics and statistics*, 98(1):173–191.
- Callan, T., Lyons, S., Scott, S., Tol, R. S., and Verde, S. (2009). The distributional implications of a carbon tax in ireland. *Energy Policy*, 37(2):407–412.
- Calonico, S., Cattaneo, M. D., and Farrell, M. H. (2020). Optimal bandwidth choice for robust bias-corrected inference in regression discontinuity designs. *The Econometrics Journal*, 23(2):192–210.
- Calonico, S., Cattaneo, M. D., Farrell, M. H., and Titiunik, R. (2017). rdrobust: Software for regression-discontinuity designs. *The Stata Journal*, 17(2):372–404.
- Calonico, S., Cattaneo, M. D., and Titiunik, R. (2015). Optimal data-driven regression discontinuity plots. *Journal of the American Statistical Association*, 110(512):1753–1769.
- Cambridge Econometrics (2017). Competitiveness impacts of carbon policies on uk energy-intensive industrial sectors to 2030. *Report*.
- Carattini, S., Heutel, G., and Melkadze, G. (2021). Climate policy, financial frictions, and transition risk. Technical Report w28525, National Bureau of Economic Research.
- Carattini, S., Kallbekken, S., and Orlov, A. (2019). How to win public support for a global carbon tax. *Nature*, 565(7739):289–291.
- Carattini, S. and Sen, S. (2019). Carbon taxes and stranded assets: Evidence from Washington state. Technical Report 7785, CESifo Group Munich.

- Carney, M. (2015). Breaking the tragedy of the horizon—climate change and financial stability. *Speech given at Lloyd's of London*, 29:220–230.
- Carraro, C., Galeotti, M., and Gallo, M. (1996). Environmental taxation and unemployment: some evidence on the ‘double dividend hypothesis’ in europe. *Journal of Public Economics*, 62(1-2):141–181.
- Cattaneo, M. D., Jansson, M., and Ma, X. (2020). Simple local polynomial density estimators. *Journal of the American Statistical Association*, 115(531):1449–1455.
- Cavallo, E., Galiani, S., Noy, I., and Pantano, J. (2013). Catastrophic natural disasters and economic growth. *Review of Economics and Statistics*, 95(5):1549–1561.
- Chanel, O., Perez, L., Künzli, N., and Medina, S. (2016). The hidden economic burden of air pollution-related morbidity: evidence from the aphekom project. *The European journal of health economics*, 17(9):1101–1115.
- Chang, T., Graff Zivin, J., Gross, T., and Neidell, M. (2016). Particulate pollution and the productivity of pear packers. *American Economic Journal: Economic Policy*, 8(3):141–169.
- Chang, T. Y., Graff Zivin, J., Gross, T., and Neidell, M. (2019). The effect of pollution on worker productivity: evidence from call center workers in china. *American Economic Journal: Applied Economics*, 11(1):151–72.
- Chay, K. Y. and Greenstone, M. (2003). The impact of air pollution on infant mortality: evidence from geographic variation in pollution shocks induced by a recession. *The Quarterly Journal of Economics*, 118(3):1121–1167.
- Chay, K. Y. and Greenstone, M. (2005). Does air quality matter? evidence from the housing market. *Journal of political Economy*, 113(2):376–424.
- Chen, S., Chen, Y., Lei, Z., and Tan-Soo, J.-S. (2020). Impact of air pollution on short-term movements: evidence from air travels in china. *Journal of Economic Geography*, 20(4):939–968.

- Chen, S., Oliva, P., and Zhang, P. (2017). The effect of air pollution on migration: evidence from china. Technical report, National Bureau of Economic Research.
- Chen, Y., Ebenstein, A., Greenstone, M., and Li, H. (2013). Evidence on the impact of sustained exposure to air pollution on life expectancy from china’s huai river policy. *Proceedings of the National Academy of Sciences*, 110(32):12936–12941.
- Chetty, R. and Hendren, N. (2018). The impacts of neighborhoods on intergenerational mobility i: Childhood exposure effects. *The Quarterly Journal of Economics*, 133(3):1107–1162.
- Chetty, R., Looney, A., and Kroft, K. (2009). Saliency and taxation: Theory and evidence. *American Economic Review*, 99(4):1145–1177.
- Chiroleu-Assouline, M. and Fodha, M. (2014). From regressive pollution taxes to progressive environmental tax reforms. *European Economic Review*, 69:126–142.
- Choma, E. F., Evans, J. S., Gómez-Ibañez, J. A., Di, Q., Schwartz, J. D., Hammitt, J. K., and Spengler, J. D. (2021). Health benefits of decreases in on-road transportation emissions in the United States from 2008 to 2017. *Proceedings of the National Academy of Sciences*, 118(51):e2107402118.
- Coglianesi, J., Davis, L. W., Kilian, L., and Stock, J. H. (2017). Anticipation, tax avoidance, and the price elasticity of gasoline demand. *Journal of Applied Econometrics*, 32(1):1–15.
- Colmer, J., Hardman, I., Shimshack, J., and Voorheis, J. (2020). Disparities in pm_{2.5} air pollution in the United States. *Science*, 369(6503):575–578.
- Colmer, J., Martin, R., Muûls, M., and Wagner, U. J. (2023). Does pricing carbon mitigate climate change? firm-level evidence from the european union emissions trading scheme. *CRC TR 224, Discussion Paper No. 232. Forthcoming in The Review of Economic Studies*.

- Cook, N., Heyes, A., and Rivers, N. (2023). Clean air and cognitive productivity: Effect and adaptation. *Journal of the Association of Environmental and Resource Economists*, 10(5):1265–1308.
- Copernicus Climate Change Service (2019). Copernicus Climate Change Service (C3S), Climate Data Store (2019): UERRA regional reanalysis for Europe on single levels from 1961 to 2019. <https://doi.org/10.24381/cds.32b04ec5>. Accessed on 21-01-2020.
- Correia, S. (2016). A feasible estimator for linear models with multi-way fixed effects. *Preprint at <http://scorreia.com/research/hdfe.pdf>*.
- Correia, S. (2019). Reghdfe: Stata module to perform linear or instrumental-variable regression absorbing any number of high-dimensional fixed effects.
- Crippa, M., others Gridded emissions of air pollutants for the period 1970–2012 within edgar v4. 3.2. *Earth Syst. Sci. Data*, 10(4):1987–2013.
- Criscuolo, C., Haskel, J., and Martin, R. (2003). Building the evidence base for productivity policy using business data linking. *Economic Trends*, 600:39–51.
- Cronin, J. A., Fullerton, D., and Sexton, S. (2019). Vertical and horizontal redistributions from a carbon tax and rebate. *Journal of the Association of Environmental and Resource Economists*, 6(S1):S169–S208.
- Cui, J., Huang, S., and Wang, C. (2023). The impact of air quality on innovation activities in china. *Journal of Environmental Economics and Management*, 122:102893.
- Cunningham, S. and Shah, M. (2018). Decriminalizing indoor prostitution: Implications for sexual violence and public health. *The Review of Economic Studies*, 85(3):1683–1715.
- Currie, J. and Neidell, M. (2005). Air pollution and infant health: what can we

- learn from california's recent experience? *The Quarterly Journal of Economics*, 120(3):1003–1030.
- Cushing, L., Blaustein-Rejto, D., Wander, M., Pastor, M., Sadd, J., Zhu, A., and Morello-Frosch, R. (2018). Carbon trading, co-pollutants, and environmental equity: Evidence from california's cap-and-trade program (2011–2015). *PLoS medicine*, 15(7):e1002604.
- Dahl, C. and Sterner, T. (1991). Analysing gasoline demand elasticities: a survey. *Energy economics*, 13(3):203–210.
- Dargay, J., Gately, D., and Sommer, M. (2007). Vehicle ownership and income growth, worldwide: 1960-2030. *The energy journal*, 28(4).
- Davis, L. W. and Kilian, L. (2011). Estimating the effect of a gasoline tax on carbon emissions. *Journal of Applied Econometrics*, 26(7):1187–1214.
- Davis, L. W. and Knittel, C. R. (2019). Are fuel economy standards regressive? *Journal of the Association of Environmental and Resource Economists*, 6(S1):S37–S63.
- DeAngelo, G. and Hansen, B. (2014). Life and death in the fast lane: Police enforcement and traffic fatalities. *American Economic Journal: Economic Policy*, 6(2):231–257.
- Dechezleprêtre, A., Nachtigall, D., and Venmans, F. (2023). The joint impact of the european union emissions trading system on carbon emissions and economic performance. *Journal of Environmental Economics and Management*, 118:102758.
- Dechezleprêtre, A., Rivers, N., and Stadler, B. (2019). The economic cost of air pollution: Evidence from europe.
- Demailly, D. and Quirion, P. (2008). European emission trading scheme and competitiveness: A case study on the iron and steel industry. *Energy Economics*, 30(4):2009–2027.

- Department for Business, Energy & Industrial Strategy, UK (2022). Review of the schemes to compensate certain energy intensive industries for indirect emission costs in electricity prices: Government response. Online.
- Deryugina, T., Heutel, G., Miller, N. H., Molitor, D., and Reif, J. (2019). The mortality and medical costs of air pollution: Evidence from changes in wind direction. *American Economic Review*, 109(12):4178–4219.
- Deschênes, O. and Greenstone, M. (2007). The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather. *American economic review*, 97(1):354–385.
- Deschenes, O., Greenstone, M., and Shapiro, J. S. (2017). Defensive investments and the demand for air quality: Evidence from the nox budget program. *American Economic Review*, 107(10):2958–89.
- Diffenbaugh, N. S. and Burke, M. (2019). Global warming has increased global economic inequality. *Proceedings of the National Academy of Sciences*, 116(20):9808–9813.
- Diluiso, F., Annicchiarico, B., Kalkuhl, M., and Minx, J. C. (2021). Climate actions and macro-financial stability: The role of central banks. *Journal of Environmental Economics and Management*, 110:102548.
- Dings, J. (2004). Fuelling oil demand: What happened to fuel taxation in europe? *European Federation for Transport and Environment*. <http://www.transportenvironment.org/sites/default/files/media/2011>, 2013.
- Dixit, A. (1989). Entry and exit decisions under uncertainty. *Journal of political Economy*, 97(3):620–638.
- Dixit, R. K., Dixit, A. K., and Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton university press.

- Dorsey, J. (2019). Waiting for the courts: Effects of policy uncertainty on pollution and investment. *Environmental and Resource Economics*, 74(4):1453–1496.
- Douenne, T. and Fabre, A. (2022). Yellow vests, pessimistic beliefs, and carbon tax aversion. *American Economic Journal: Economic Policy*, 14(1):81–110.
- Drupp, M. A., Kornek, U., Meya, J., and Sager, L. (2021). Inequality and the environment: The economics of a two-headed hydra.
- Drupp, M. A., Meya, J. N., Baumgärtner, S., and Quaas, M. F. (2018). Economic inequality and the value of nature. *Ecological Economics*, 150:340–345.
- Durlauf, S. N. and Seshadri, A. (2018). Understanding the great gatsby curve. *NBER Macroeconomics Annual*, 32(1):333–393.
- Dustmann, C., Schönberg, U., and Stuhler, J. (2017). Labor supply shocks, native wages, and the adjustment of local employment. *Quarterly Journal of Economics*, 132(1):435–483.
- EAERE (2020). Economists’ statement on carbon pricing.
- Ebenstein, A., Lavy, V., and Roth, S. (2016). The long-run economic consequences of high-stakes examinations: Evidence from transitory variation in pollution. *American Economic Journal: Applied Economics*, 8(4):36–65.
- ECB (2021). Financial stability review. Technical Report May 2021.
- Edenhofer, O., Flachsland, C., Kalkuhl, M., Knopf, B., and Pahle, M. (2019). Options for a co2 price reform. mcc-pik expertise for the expert ”a ndigenrat for the assessment of macroeconomic development.
- Engle, R. F., Giglio, S., Kelly, B., Lee, H., and Stroebel, J. (2020). Hedging climate change news. *The Review of Financial Studies*, 33(3):1184–1216.
- European Commission (2012). Guidelines on certain state aid measures in the context of the greenhouse gas emission allowance trading scheme post-2012. *Official Journal of the European Union*.

- European Commission (2018). Report on the functioning of the european carbon marke. Report from the COMmission to the European Parliament and the Council COM(2018) 842 final, European Commission.
- European Commission (2020a). Commission staff working document evaluation accompanying the document: Impact assessment on Guidelines on certain State aid measures in the context of the system for greenhouse gas emission allowance trading post 2021. Technical Report SWD(2020) 194 final, European Commission.
- European Commission (2020b). Impact assessment accompanying the document communication from the commission on guidelines on certain state aid measures in the context of the system for greenhouse gas emission allowance trading post 2021. COMMISSION STAFF WORKING DOCUMENT.
- European Commission DG Competition (2022). State Aid SA.100559 (2022/N) – Germany Compensation for indirect ETS costs in Germany for 2021-2030. Decision text.
- European Parliament (2021). Draft report on the proposal for a regulation of the european parliament and of the council establishing a carbon border adjustment mechanism (COM(2021)0564 – C9-0328/2021 – 2021/0214(COD)). Online publication.
- Eurostat (2020). Euro/ECU exchange rates - annual data. Accessed: 2020-09-21.
- Fabra, N. and Reguant, M. (2014). Pass-through of emissions costs in electricity markets. *American Economic Review*, 104(9):2872–99.
- Fabrizio, K. R. (2013). The effect of regulatory uncertainty on investment: Evidence from renewable energy generation. *The Journal of Law, Economics, and Organization*, 29(4):765–798.
- Ferman, B. and Pinto, C. (2021). Synthetic controls with imperfect pretreatment fit. *Quantitative Economics*, 12(4):1197–1221.

- Ferman, B., Pinto, C., and Possebom, V. (2020). Cherry picking with synthetic controls. *Journal of Policy Analysis and Management*, 39(2):510–532.
- Ferrara, A. R. and Giua, L. (2022). Indirect cost compensation under the eu ets: A firm-level analysis. *Energy Policy*, 165:112989.
- Firpo, S. and Possebom, V. (2018). Synthetic control method: Inference, sensitivity analysis and confidence sets. *Journal of Causal Inference*, 6(2).
- Fischer, C. (2001). Rebating environmental policy revenues: Output-based allocations and tradable performance standards. RFF Discussion Paper.
- Fischer, C. (2008). Emissions pricing, spillovers, and public investment in environmentally friendly technologies. *Energy Economics*, 30(2):487–502.
- Fischer, C. and Fox, A. K. (2007). Output-based allocation of emissions permits for mitigating tax and trade interactions. *Land economics*, 83(4):575–599.
- Fischer, C. and Fox, A. K. (2011). The role of trade and competitiveness measures in US climate policy. *American Economic Review*, 101(3):258–262.
- Fischer, C. and Pizer, W. A. (2019). Horizontal equity effects in energy regulation. *Journal of the Association of Environmental and Resource Economists*, 6(S1):S209–S237.
- Flintz, J., Frondel, M., Horvath, M., and Vance, C. (2022). *Emission effects of Germany’s vehicle taxation: Recent empirical evidence*. Number 978. Ruhr Economic Papers.
- Fowlie, M., Reguant, M., and Ryan, S. P. (2016). Market-based emissions regulation and industry dynamics. *Journal of Political Economy*, 124(1):249–302.
- Fowlie, M. L. and Reguant, M. (2022). Mitigating emissions leakage in incomplete carbon markets. *Journal of the Association of Environmental and Resource Economists*, 9(2):307–343.

- Frondel, M. and Vance, C. (2014). More pain at the diesel pump? An econometric comparison of diesel and petrol price elasticities. *Journal of Transport Economics and Policy (JTEP)*, 48(3):449–463.
- Fu, S., Viard, V. B., and Zhang, P. (2021). Air pollution and manufacturing firm productivity: Nationwide estimates for china. *The Economic Journal*, 131(640):3241–3273.
- Fullerton, D. and Muehlegger, E. (2019). Who bears the economic burdens of environmental regulations? *Review of Environmental Economics and Policy*.
- Ganapati, S., Shapiro, J. S., and Walker, R. (2020). Energy cost pass-through in us manufacturing: Estimates and implications for carbon taxes. *American Economic Journal: Applied Economics*, 12(2):303–42.
- Gentzkow, M., Kelly, B., and Taddy, M. (2019). Text as data. *Journal of Economic Literature*, 57(3):535–574.
- Gerster, A. and Lamp, S. (2023). Energy tax exemptions and industrial production. *CESifo Working Paper No. 10232. (Conditionally accepted at the Economic Journal)*.
- Gobillon, L. and Magnac, T. (2016). Regional policy evaluation: Interactive fixed effects and synthetic controls. *Review of Economics and Statistics*, 98(3):535–551.
- Gormley, T. A. and Matsa, D. A. (2014). Common errors: How to (and not to) control for unobserved heterogeneity. *The Review of Financial Studies*, 27(2):617–661.
- Goulder, L. H., Hafstead, M. A., Kim, G., and Long, X. (2019). Impacts of a carbon tax across us household income groups: What are the equity-efficiency trade-offs? *Journal of Public Economics*, 175:44–64.
- Goulder, L. H. and Parry, I. W. (2008). Instrument choice in environmental policy. *Review of environmental economics and policy*.

- Graff Zivin, J. and Neidell, M. (2012). The impact of pollution on worker productivity. *American Economic Review*, 102(7):3652–73.
- Graff Zivin, J. and Neidell, M. (2013). Environment, health, and human capital. *Journal of Economic Literature*, 51(3):689–730.
- Grange, S. K., Farren, N. J., Vaughan, A. R., Davison, J., and Carslaw, D. C. (2020). Post-dieselgate: Evidence of nox emission reductions using on-road remote sensing. *Environmental science & technology letters*, 7(6):382–387.
- Green, J. F. (2021). Does carbon pricing reduce emissions? a review of ex-post analyses. *Environmental Research Letters*.
- Greenstone, M. and Jack, B. K. (2013). Envirodevonomics: A research agenda for a young field. Technical report, National Bureau of Economic Research.
- Grigolon, L., Reynaert, M., and Verboven, F. (2018). Consumer valuation of fuel costs and tax policy: Evidence from the european car market. *American Economic Journal: Economic Policy*, 10(3):193–225.
- Griliches, Z. (1990). Patent statistics as economic indicators: A survey. *Journal of Economic Literature*, 28(4):1661–1707.
- Grossman, M. (1972). On the concept of health capital and the demand for health. *Journal of Political economy*, 80(2):223–255.
- Grubb, M. and Drummond, P. (2018). Uk industrial electricity prices: competitiveness in a low carbon world. *Report*.
- Grubb, M., Jordan, N. D., Hertwich, E., Neuhoff, K., Das, K., Bandyopadhyay, K. R., van Asselt, H., Sato, M., Wang, R., Pizer, B., and Oh, H. (2022). Carbon leakage, consumption, and trade. *Annual Review of Environment and Resources*, 47(1):null.
- Guadalupe, M., Kuzmina, O., and Thomas, C. (2012). Innovation and foreign ownership. *American Economic Review*, 102(7):3594–3627.

- Hagem, C., Hoel, M., and Sterner, T. (2020). Refunding emission payments: Output-based versus expenditure-based refunding. *Environmental and Resource Economics*, 77(3):641–667.
- Hallegatte, S. and Rozenberg, J. (2017). Climate change through a poverty lens. *Nature Climate Change*, 7(4):250–256.
- Handley, K. and Limão, N. (2015). Trade and investment under policy uncertainty: Theory and firm evidence. *American Economic Journal: Economic Policy*, 7(4):189–222.
- Hanna, R. and Oliva, P. (2015). The effect of pollution on labor supply: Evidence from a natural experiment in Mexico City. *Journal of Public Economics*, 122:68–79.
- Harstad, B. (2022). Pledge-and-review bargaining: From Kyoto to Paris. *Economic Journal*.
- Hassan, T. A., Hollander, S., van Lent, L., and Tahoun, A. (2019). Firm-level political risk: Measurement and effects. *The Quarterly Journal of Economics*, 134(4):2135–2202. Publisher: Oxford Academic.
- Hassett, K. A. and Metcalf, G. E. (1999). Investment with uncertain tax policy: Does random tax policy discourage investment. *The Economic Journal*, 109(457):372–393.
- Hausman, C. and Stolper, S. (2021). Inequality, information failures, and air pollution. *Journal of Environmental Economics and Management*, 110:102552.
- He, J., Liu, H., and Salvo, A. (2019). Severe air pollution and labor productivity: Evidence from industrial towns in China. *American Economic Journal: Applied Economics*, 11(1):173–201.
- Heblich, S., Trew, A., and Zylberberg, Y. (2021). East-side story: Historical pollution and persistent neighborhood sorting. *Journal of Political Economy*, 129(5):1508–1552.

- Helm, I., Koch, N., and Rohlf, A. (2023). The effects of cash for clunkers on local air quality. *Journal of Urban Economics*, 138:103576.
- Hernandez-Cortes, D. and Meng, K. C. (2023). Do environmental markets cause environmental injustice? evidence from california’s carbon market. *Journal of Public Economics*, 217:104786.
- Hicks, D., Marsh, P., and Oliva, P. (2016). Air pollution and procyclical mortality: Causal evidence from thermal inversions. *NBER working paper*.
- Hintermann, B. (2016). Pass-through of CO2 emission costs to hourly electricity prices in Germany. *Journal of the Association of Environmental and Resource Economists*, 3(4):857–891.
- Hintermann, B., Schoeman, B. M., Molloy, J., Götschli, T., Castro, A., Tchervenkov, C., Tomic, U., and Axhausen, K. W. (2021). Pigovian transport pricing in practice. Technical report, WWZ Working Paper.
- Hirst, D. (2018). Carbon price floor (CPF) and the price support mechanism. *Report*.
- Holman, C., Harrison, R., and Querol, X. (2015). Review of the efficacy of low emission zones to improve urban air quality in european cities. *Atmospheric Environment*, 111:161–169.
- Holub, F., Hospido, L., and Wagner, U. J. (2020). Urban air pollution and sick leaves: Evidence from social security data.
- Hunt, A., Ferguson, J., Hurley, F., and Searl, A. (2016). Social costs of morbidity impacts of air pollution. *OECD Environment Working Papers*, 99.
- Huse, C. and Koptuyug, N. (2022). Salience and policy instruments: Evidence from the auto market. *Journal of the Association of Environmental and Resource Economists*, 9(2):345–382.
- Hyslop, N. P. (2009). Impaired visibility: the air pollution people see. *Atmospheric Environment*, 43(1):182–195.

- Imbens, G. W. (2004). Nonparametric estimation of average treatment effects under exogeneity: A review. *Review of Economics and Statistics*, 86(1):4–29.
- Imbens, G. W. and Lemieux, T. (2008). Regression discontinuity designs: A guide to practice. *Journal of Econometrics*, 142(2):615–635.
- IPCC (2022). *Climate Change 2022: Impacts, Adaptation and Vulnerability*. Summary for Policymakers. Cambridge University Press, Cambridge, UK and New York, USA.
- Isaksen, E. T. (2020). Have international pollution protocols made a difference? *Journal of Environmental Economics and Management*, 103:102358.
- Isen, A., Rossin-Slater, M., and Walker, W. R. (2017). Every breath you take—every dollar you’ll make: The long-term consequences of the clean air act of 1970. *Journal of Political Economy*, 125(3):848–902.
- Ito, K. (2015). Asymmetric incentives in subsidies: Evidence from a large-scale electricity rebate program. *American Economic Journal: Economic Policy*, 7(3):209–37.
- Ito, K. and Zhang, S. (2020). Willingness to pay for clean air: Evidence from air purifier markets in china. *Journal of Political Economy*, 128(5):1627–1672.
- Jacobsen, M. R., Sallee, J. M., Shapiro, J. S., and Van Benthem, A. A. (2023). Regulating untaxable externalities: Are vehicle air pollution standards effective and efficient? *Quarterly Journal of Economics*, 138(3):1907–1976.
- Jans, J., Johansson, P., and Nilsson, J. P. (2018). Economic status, air quality, and child health: Evidence from inversion episodes. *Journal of Health Economics*, 61:220–232.
- Jbaily, A., Zhou, X., Liu, J., Lee, T.-H., Kamareddine, L., Verguet, S., and Dominici, F. (2022). Air pollution exposure disparities across us population and income groups. *Nature*, 601(7892):228–233.

- Känzig, D. R. (2023). The unequal economic consequences of carbon pricing. *NBER, Working Paper*(31221).
- Kaul, A., Klößner, S., Pfeifer, G., and Schieler, M. (2015). Synthetic control methods: Never use all pre-intervention outcomes together with covariates.
- Ketcham, J. D., Kuminoff, N. V., and Mathes, S. (2019). The illness-poverty-amenity trap. Unpublished working paper.
- Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–1069.
- Kilian, L. and Zhou, X. (2023). Heterogeneity in the pass-through from oil to gasoline prices: A new instrument for estimating the price elasticity of gasoline demand. *FRB of Dallas Working Paper*.
- Klauber, H., Holub, F., Koch, N., Pestel, N., Ritter, N., and Rohlf, A. (2021). Killing prescriptions softly: Low emission zones and child health from birth to school.
- Klemetsen, M. E., Rosendahl, K. E., and Jakobsen, A. L. (2020). The impacts of the EU ETS on Norwegian plants' environmental and economic performance. *Climate Change Economics*, 11(1).
- Klenert, D. and Mattauch, L. (2016). How to make a carbon tax reform progressive: The role of subsistence consumption. *Economics letters*, 138:100–103.
- Klenert, D., Mattauch, L., Combet, E., Edenhofer, O., Hepburn, C., Rafaty, R., and Stern, N. (2018). Making carbon pricing work. *Nature Climate Change*, 8(8):669–677.
- Klier, T. and Linn, J. (2015). Using taxes to reduce carbon dioxide emissions rates of new passenger vehicles: evidence from france, germany, and sweden. *American Economic Journal: Economic Policy*, 7(1):212–242.
- Knigge, M. and Görlach, B. (2005). *Auswirkungen der ökologischen Steuerreform auf private Haushalte*. Ecologic.

- Knittel, C. R., Miller, D. L., and Sanders, N. J. (2016). Caution, drivers! children present: Traffic, pollution, and infant health. *Review of Economics and Statistics*, 98(2):350–366.
- Knittel, C. R. and Sandler, R. (2018). The welfare impact of second-best uniform-pigouvian taxation: Evidence from transportation. *American Economic Journal: Economic Policy*, 10(4):211–242.
- Kohlhaas, M. et al. (2005). *Gesamtwirtschaftliche Effekte der ökologischen Steuerreform*. DIW Berlin.
- Köppl, A. and Schratzenstaller, M. (2023). Carbon taxation: A review of the empirical literature. *Journal of Economic Surveys*, 37(4):1353–1388.
- Kossoy, A., Peszko, G., Oppermann, K., Prytz, N., Klein, N., Blok, K., Lam, L., Wong, L., and Borkent, B. (2015). *State and trends of carbon pricing 2015*. Washington DC: World Bank Publications.
- Krebs, B. and Luechinger, S. (2024). Air pollution, cognitive performance, and the role of task proficiency. *Journal of the Association of Environmental and Resource Economists*. Forthcoming.
- Krueger, P., Sautner, Z., and Starks, L. T. (2020). The importance of climate risks for institutional investors. *The Review of Financial Studies*, 33(3):1067–1111.
- Laing, T., Sato, M., Grubb, M., and Comberti, C. (2014). The effects and side-effects of the eu emissions trading scheme. *WIREs Clim Change*, 5(4):509–519.
- Lange, A., Vogt, C., and Ziegler, A. (2007). On the importance of equity in international climate policy: An empirical analysis. *Energy Economics*, 29(3):545–562.
- Lee, D. S. and Lemieux, T. (2010). Regression discontinuity designs in economics. *Journal of economic literature*, 48(2):281–355.

- Lee, S. and Lin, J. (2018). Natural amenities, neighbourhood dynamics, and persistence in the spatial distribution of income. *The Review of Economic Studies*, 85(1):663–694.
- Lemoine, D. (2017). Green expectations: Current effects of anticipated carbon pricing. *Review of Economics and Statistics*, 99(3):499–513.
- Leroutier, M. (2022). Carbon pricing and power sector decarbonization: Evidence from the uk. *Journal of Environmental Economics and Management*, 111:102580.
- Leroutier, M. and Ollivier, H. (2022). The cost of air pollution for workers and firms. mimeo. 13th Toulouse Conference on The Economics of Energy and Climate.
- Levin, L., Lewis, M. S., and Wolak, F. A. (2017). High frequency evidence on the demand for gasoline. *American Economic Journal: Economic Policy*, 9(3):314–47.
- Levinson, A. (2019). Energy efficiency standards are more regressive than energy taxes: Theory and evidence. *Journal of the Association of Environmental and Resource Economists*, 6(S1):S7–S36.
- Li, S., Linn, J., and Muehlegger, E. (2014). Gasoline taxes and consumer behavior. *American Economic Journal: Economic Policy*, 6(4):302–42.
- Liang, Y., Rudik, I., Zou, E. Y., Johnston, A., Rodewald, A. D., and Kling, C. L. (2020). Conservation cobenefits from air pollution regulation: Evidence from birds. *Proceedings of the National Academy of Sciences*, 117(49):30900–30906.
- Lin, J., Hui, D., Kumar, A., Yu, Z.-G., and Huang, Y. (2023). Climate change and/or pollution on the carbon cycle in terrestrial ecosystems. *Frontiers in Environmental Science*, 11:1253172.
- Lindo, J. M. and Packham, A. (2017). How much can expanding access to long-acting reversible contraceptives reduce teen birth rates? *American Economic Journal: Economic Policy*, 9(3):348–76.

- Linn, J. (2019). Interactions between climate and local air pollution policies: The case of european passenger cars. *Journal of the Association of Environmental and Resource Economists*, 6(4):709–740.
- Löschel, A., Pei, J., Wang, R., Sturm, B., Buchholz, W., and Zhao, Z. (2021). The demand for global and local environmental protection: Experimental evidence from climate change mitigation in beijing. *Land Economics*, 97(1):137–154.
- Maestre-Andrés, S., Drews, S., and van den Bergh, J. (2019). Perceived fairness and public acceptability of carbon pricing: a review of the literature. *Climate policy*, 19(9):1186–1204.
- Margaryan, S. (2021). Low emission zones and population health. *Journal of health economics*, 76:102402.
- Marin, G. and Vona, F. (2021). The impact of energy prices on socioeconomic and environmental performance: Evidence from french manufacturing establishments, 1997–2015. *European Economic Review*, 135:103739.
- Martin, R., De Preux, L. B., and Wagner, U. J. (2014a). The impact of a carbon tax on manufacturing: Evidence from microdata. *Journal of Public Economics*, 117:1–14.
- Martin, R., Muûls, M., Preux, L. B. d., and Wagner, U. J. (2014b). Industry compensation under relocation risk: A firm-level analysis of the EU emissions trading scheme. *American Economic Review*, 104:2482–2508.
- Martin, R., Muûls, M., and Wagner, U. J. (2016). The Impact of the European Union Emissions Trading Scheme on Regulated Firms: What Is the Evidence after Ten Years? *Review of Environmental Economics and Policy*, 10(1):129–148.
- McDonald, R. and Siegel, D. (1986). The value of waiting to invest. *The Quarterly Journal of Economics*, 101(4):707–728.

- McMichael, A. J., Woodruff, R. E., and Hales, S. (2006). Climate change and human health: present and future risks. *The Lancet*, 367(9513):859–869.
- Meng, K. C. (2017). Using a free permit rule to forecast the marginal abatement cost of proposed climate policy. *American Economic Review*, 107(3):748–84.
- Metcalf, G. E. (2009). Designing a carbon tax to reduce us greenhouse gas emissions. *Review of Environmental Economics and Policy*.
- Meyer, S. and Pagel, M. (2017). Fresh air eases work—the effect of air quality on individual investor activity. Technical report, National Bureau of Economic Research.
- Mideksa, T. K. (2021). Pricing for a cooler planet: An empirical analysis of the effect of taxing carbon. *CESifo Working Paper*, No. 9172.
- Ministry of the Environment (1995). Finland’s national report under the united nation’s framework convention on climate change.
- Molina, T. (2021). Pollution, ability, and gender-specific investment responses to shocks. *Journal of the European Economic Association*, 19(1):580–619.
- Montgomery, W. D. (1972). Markets in licenses and efficient pollution control programs. *Journal of economic theory*, 5(3):395–418.
- Moore, F. C., Lacasse, K., Mach, K. J., Shin, Y. A., Gross, L. J., and Beckage, B. (2022). Determinants of emissions pathways in the coupled climate–social system. *Nature*, 603(7899):103–111.
- Moretti, E. and Neidell, M. (2011). Pollution, health, and avoidance behavior evidence from the ports of los angeles. *Journal of human Resources*, 46(1):154–175.
- Morgado Simões, H. (2023). EU carbon border adjustment mechanism: Implications for climate and competitiveness. Briefing - EU Legislation in Progress PE 698.889, European Parliament.

- Muller, N. Z. and Mendelsohn, R. (2009). Efficient pollution regulation: getting the prices right. *American Economic Review*, 99(5):1714–39.
- Naegele, H. and Zaklan, A. (2019). Does the eu ets cause carbon leakage in european manufacturing? *Journal of Environmental Economics and Management*, 93:125–147.
- Neidell, M. J. (2004). Air pollution, health, and socio-economic status: the effect of outdoor air quality on childhood asthma. *Journal of health economics*, 23(6):1209–1236.
- Neuhoff, K., Martinez, K. K., and Sato, M. (2006). Allocation, incentives and distortions: the impact of EU ETS emissions allowance allocations to the electricity sector. *Climate Policy*, 6(1):73–91.
- Newell, R. G. and Stavins, R. N. (2003). Cost heterogeneity and the potential savings from market-based policies. *Journal of Regulatory Economics*, 23:43–59.
- Newey, W. K. and West, K. D. (1994). Automatic lag selection in covariance matrix estimation. *The Review of Economic Studies*, 61(4):631–653.
- Newman, R. and Noy, I. (2023). The global costs of extreme weather that are attributable to climate change. *Nature Communications*, 14(1):6103.
- Noailly, J. and Smeets, R. (2015). Directing technical change from fossil-fuel to renewable energy innovation: An application using firm-level patent data. *Journal of Environmental Economics and Management*, 72:15–37.
- Noble, M., Wright, G., Smith, G., and Dibben, C. (2006). Measuring multiple deprivation at the small-area level. *Environment and planning A*, 38(1):169–185.
- Nordhaus, W. (2021). Dynamic climate clubs: On the effectiveness of incentives in global climate agreements. *Proceedings of the National Academy of Sciences*, 118(45):e2109988118.

- Ntziachristos, L. and Samaras, Z. (2019). Passenger cars, light commercial trucks, heavy-duty vehicles including buses and motor cycles category. EMEP/EEA emission inventory guidebook 2019.
- OECD (2001). Environmentally related taxes in oecd countries: Issues and strategies.
- ONS (2018). Healthcare expenditure: Uk health accounts 2018.
- Osamu, K. (2012). The role of standards: The japanese top runner program for end-use efficiency., chapter 24: Historical case studies of energy technology innovation. *The global energy assessment*.
- Parry, I., Mylonas, V., and Vernon, N. (2021). Mitigation policies for the paris agreement: An assessment for g20 countries. *Journal of the Association of Environmental and Resource Economists*, 8(4):797–823.
- Parry, I., Veung, C., and Heine, D. (2015). How much carbon pricing is in countries' own interests? The critical role of co-benefits. *Climate Change Economics*, 6(04):1550019.
- Parry, I. W. H., Walls, M., and Harrington, W. (2007). Automobile externalities and policies. *Journal of economic literature*, 45(2):373–399.
- Pearce, D. (1991). The role of carbon taxes in adjusting to global warming. *The economic journal*, 101(407):938–948.
- Peeples, L. (2020). How air pollution threatens brain health. *Proceedings of the National Academy of Sciences*, 117(25):13856–13860.
- Perino, G. (2018). New eu ets phase 4 rules temporarily puncture waterbed. *Nature Climate Change*, 8(4):262–264.
- Petrick, S. and Wagner, U. (2014). The impact of carbon trading on industry: Evidence from german manufacturing firms. *Working paper*.

- Pimpin, L., Retat, L., Fecht, D., de Preux, L., Sassi, F., Gulliver, J., Belloni, A., Ferguson, B., Corbould, E., Jaccard, A., et al. (2018). Estimating the costs of air pollution to the national health service and social care: An assessment and forecast up to 2035. *PLoS medicine*, 15(7):e1002602.
- Pindyck, R. S. (1988). *Capital risk and models of investment behaviour*. Springer.
- Pope III, C. A. and Dockery, D. W. (2006). Health effects of fine particulate air pollution: lines that connect. *Journal of the air & waste management association*, 56(6):709–742.
- Popp, D. (2002). Induced innovation and energy prices. *The American Economic Review*, 92(1):160–180.
- Popp, D. (2019). Environmental policy and innovation: A decade of research. *NBER Working Paper*, (w25631).
- Porter, M. E. and van der Linde, C. (1995). Toward a new conception of the environment-competitiveness relationship. *The Journal of Economic Perspectives*, 9(4):97–118.
- Preuss, M., Reuter, W. H., Schmidt, C. M., et al. (2021). Distributional effects of carbon pricing in germany. *FinanzArchiv: Public Finance Analysis*, 77(3):287–316.
- Quirion, P. (2009). Historic versus output-based allocation of GHG tradable allowances: a comparison. *Climate Policy*, 9(6):575–592.
- Rafaty, R. (2018). Perceptions of corruption, political distrust, and the weakening of climate policy. *Global Environmental Politics*, 18(3):106–129.
- Ramírez, A. S., Ramondt, S., Van Bogart, K., and Perez-Zuniga, R. (2019). Public awareness of air pollution and health threats: challenges and opportunities for communication strategies to improve environmental health literacy. *Journal of Health Communication*, 24(1):75–83.

- Reynaert, M. (2021). Abatement strategies and the cost of environmental regulation: Emission standards on the European car market. *Review of Economic Studies*, 88(1):454–488.
- Rivers, N. and Schaufele, B. (2015). Salience of carbon taxes in the gasoline market. *Journal of Environmental Economics and Management*, 74:23–36.
- Rogelj, J., Den Elzen, M., Höhne, N., Fransen, T., Fekete, H., Winkler, H., Schaeffer, R., Sha, F., Riahi, K., and Meinshausen, M. (2016). Paris agreement climate proposals need a boost to keep warming well below 2 c. *Nature*, 534(7609):631–639.
- Rohlf, A., Holub, F., Koch, N., and Ritter, N. (2020). The effect of clean air on pharmaceutical expenditures. *Economics Letters*, 192:109221.
- Romanello, M., McGushin, A., Di Napoli, C., Drummond, P., Hughes, N., Jamart, L., Kennard, H., Lampard, P., Rodriguez, B. S., Arnell, N., et al. (2021). The 2021 report of the Lancet Countdown on health and climate change: code red for a healthy future. *The Lancet*, 398(10311):1619–1662.
- Rosendahl, K. E. (2004). Cost-effective environmental policy: implications of induced technological change. *Journal of Environmental Economics and Management*, 48(3):1099–1121.
- Rosendahl, K. E. (2008). Incentives and prices in an emissions trading scheme with updating. *Journal of Environmental Economics and Management*, 56(1):69–82.
- Rosendahl, K. E. and Storrøsten, H. B. (2015). Allocation of emission allowances: Impacts on technology investments. *Climate Change Economics*, 6(03):1550010.
- Rozenberg, J., Vogt-Schilb, A., and Hallegatte, S. (2018). Instrument choice and stranded assets in the transition to clean capital. *Journal of Environmental Economics and Management*.
- Rudebusch (2021). Climate change is a source of financial risk. *FRBSF Economic Letter*.

- Runst, P. and Höhle, D. (2021). Quasi-carbon taxation-the german eco tax and its impact on co2 emissions. Technical report, ifh Working Paper.
- Sager, L. (2019). Estimating the effect of air pollution on road safety using atmospheric temperature inversions. *Journal of Environmental Economics and Management*, 98:102250.
- Samoli, E., Stergiopoulou, A., Santana, P., Rodopoulou, S., Mitsakou, C., Dimitroulopoulou, C., Bauwelinck, M., de Hoogh, K., Costa, C., Marí-Dell’Olmo, M., et al. (2019). Spatial variability in air pollution exposure in relation to socioeconomic indicators in nine european metropolitan areas: A study on environmental inequality. *Environmental Pollution*, 249:345–353.
- Sanderfoot, O., Bassing, S., Brusa, J., Emmet, R., Gillman, S., Swift, K., and Gardner, B. (2022). A review of the effects of wildfire smoke on the health and behavior of wildlife. *Environmental Research Letters*, 16(12):123003.
- Santos, R., Gravelle, H., and Propper, C. (2017). Does quality affect patients’ choice of doctor? evidence from england. *The Economic Journal*, 127(600):445–494.
- Sanyal, P. and Ghosh, S. (2013). Product market competition and upstream innovation: Evidence from the us electricity market deregulation. *Review of Economics and Statistics*, 95(1):237–254.
- Sarmiento, L. (2022). Air pollution and the productivity of high-skill labor: evidence from court hearings. *The Scandinavian Journal of Economics*, 124(1):301–332.
- Sato, M., Rafaty, R., Calel, R., and Grubb, M. (2022). Allocation, allocation, allocation! the political economy of the development of the european union emissions trading system. *WIREs Climate Change*, n/a(n/a):e796.
- Schlenker, W. and Walker, W. R. (2016). Airports, air pollution, and contemporaneous health. *The Review of Economic Studies*, 83(2):768–809.

- Schmidt, A., Ivanova, A., and Schäfer, M. S. (2013). Media attention for climate change around the world: A comparative analysis of newspaper coverage in 27 countries. *Global Environmental Change*, 23(5):1233–1248.
- Sen, S. and von Schickfus, M.-T. (2020). Climate policy, stranded assets, and investors' expectations. *Journal of Environmental Economics and Management*, 100:102277.
- Shaw, C., Hales, S., Howden-Chapman, P., and Edwards, R. (2014). Health co-benefits of climate change mitigation policies in the transport sector. *Nature Climate Change*, 4(6):427–433.
- Sijm, J., Neuhoff, K., and Chen, Y. (2006). Co2 cost pass-through and windfall profits in the power sector. *Climate policy*, 6(1):49–72.
- Steiner, V. and Cludius, J. (2010). Ökosteuer hat zu geringerer umweltbelastung des verkehrs beigetragen. *DIW Wochenbericht*, 77(13/14):2–7.
- Stern, T. (2007). Fuel taxes: An important instrument for climate policy. *Energy policy*, 35(6):3194–3202.
- Stern, T. (2012a). Distributional effects of taxing transport fuel. *Energy Policy*, 41:75–83.
- Stern, T. (2012b). *Fuel taxes and the poor: The distributional effects of gasoline taxation and their implications for climate policy*. Routledge.
- Stern, T. and Muller, A. (2008). Output and abatement effects of allocation readjustment in permit trade. *Climatic Change*, 86(1-2):33–49.
- Stock, J. and Yogo, M. (2005). Asymptotic distributions of instrumental variables statistics with many instruments. *Identification and inference for econometric models: Essays in honor of Thomas Rothenberg*, 6:109–120.
- Sun, C., Kahn, M. E., and Zheng, S. (2017). Self-protection investment exacerbates air pollution exposure inequality in urban china. *Ecological Economics*, 131:468–474.

- The World Bank (2023). State and trends of carbon pricing 2023. World Bank Group Report.
- Tobin, P., Schmidt, N. M., Tosun, J., and Burns, C. (2018). Mapping states' Paris climate pledges: Analysing targets and groups at COP 21. *Global Environmental Change*, 48:11–21.
- Treasury, H. (2016). The green book: Appraisal and evaluation in central government.
- UBA (2019). Methodological convention 3.0 for the assessment of environmental costs. *Umweltbundesamt*.
- UK-AIR (2023). Analysis of abatement options to reduce pm2.5 concentrations: Support for national air pollution control strategies. Available for download at <https://uk-air.defra.gov.uk/library/>.
- UK BEIS (2018). Energy consumption in the uk, 2018 update. Dataset.
- UK BEIS (2019). Energy intensive industries (eiis): Compensation for the indirect costs of the eu emission trading system and the carbon price support mechanism. Report.
- Umweltbundesamt (2012). Best practice cost rates for air pollutants, traffic, power generation and heat production. appendix b of the “method convention 2.0 for the evaluation of environmental costs”.
- US EPA (2005). Average carbon dioxide emissions resulting from gasoline and diesel fuel.
- van Benthem, A. A., Crooks, E., Giglio, S., Schwob, E., and Stroebel, J. (2022). The effect of climate risks on the interactions between financial markets and energy companies. *Nature Energy*, 7(8):690–697.
- van der Ploeg, F. and Rezai, A. (2020a). The risk of policy tipping and stranded carbon assets. *Journal of Environmental Economics and Management*, 100:102258.

- van der Ploeg, F. and Rezai, A. (2020b). Stranded assets in the transition to a carbon-free economy. *Annual Review of Resource Economics*, 12(1).
- Vandyck, T., Keramidis, K., Tchung-Ming, S., Weitzel, M., and Van Dingenen, R. (2020). Quantifying air quality co-benefits of climate policy across sectors and regions. *Climatic Change*, 163(3):1501–1517.
- Venmans, F. M. J. (2016). The effect of allocation above emissions and price uncertainty on abatement investments under the EU ETS. *Journal of Cleaner Production*, 126(Supplement C):595–606.
- Verde, S. F. (2020). The impact of the eu emissions trading system on competitiveness and carbon leakage: the econometric evidence. *Journal of Economic Surveys*, 34(2):320–343.
- Verde, S. F., Graf, C., and Jong, T. (2019). Installation entries and exits in the EU ETS: patterns and the delay effect of closure provisions. *Energy Economics*, 78:508–524.
- Vermeulen, R., Schets, E., Lohuis, M., Kölbl, B., Jansen, D.-J., and Heeringa, W. (2018). An energy transition risk stress test for the financial system of the Netherland. Technical Report Volume 16 - 7, De Nederlandsche Bank.
- Wagner, U. J. and De Preux, L. (2016). The co-benefits of climate policy: Evidence from the eu emissions trading scheme.
- Wall Street Journal (2019). Economists statement on carbon dividends.
- Wang, Y., Xiong, L., and Tang, M. (2017). Toxicity of inhaled particulate matter on the central nervous system: neuroinflammation, neuropsychological effects and neurodegenerative disease. *Journal of Applied Toxicology*, 37(6):644–667.
- Ward, C. J. (2015). It’s an ill wind: The effect of fine particulate air pollution on respiratory hospitalizations. *Canadian Journal of Economics/Revue canadienne d’économique*, 48(5):1694–1732.

- Weiss, M. (2009). How do Germans react to the commuting allowance? Technical report, arqus Discussion Paper.
- Weitzman, M. (2015). Internationally-tradable permits can be riskier for a country than an internally-imposed carbon price. *Available at SSRN 2652796*.
- Weitzman, M. L. (1974). Prices vs. quantities. *The review of economic studies*, 41(4):477–491.
- Williams, A. M. and Phaneuf, D. J. (2019). The morbidity costs of air pollution: Evidence from spending on chronic respiratory conditions. *Environmental and Resource Economics*, 74(2):571–603.
- Wolff, H. and Perry, L. (2010). Trends in clean air legislation in Europe: Particulate matter and low emission zones. *Review of Environmental Economics and Policy*, 4(2):293–308.
- World Bank (2020). Carbon pricing dashboard.
- World Bank (2022). State and trends of carbon pricing 2022.
- World Health Organization (2016). Ambient air pollution: A global assessment of exposure and burden of disease.
- Wu, X., Braun, D., Schwartz, J., Kioumourtzoglou, M., and Dominici, F. (2020). Evaluating the impact of long-term exposure to fine particulate matter on mortality among the elderly. *Science Advances*, 6(29):eaba5692.
- Xu, Y. (2017). Generalized synthetic control method: Causal inference with interactive fixed effects models. *Political Analysis*, 25(1):57–76.
- Yamazaki, A. (2017). Jobs and climate policy: Evidence from British Columbia’s revenue-neutral carbon tax. *Journal of Environmental Economics and Management*, 83:197–216.

- Zetterberg, L. (2014). Benchmarking in the european union emissions trading system: Abatement incentives. *Energy Economics*, 43:218–224.
- Zhang, J. and Mu, Q. (2018). Air pollution and defensive expenditures: Evidence from particulate-filtering facemasks. *Journal of Environmental Economics and Management*, 92:517–536.
- Zhang, X., Chen, X., and Zhang, X. (2018). The impact of exposure to air pollution on cognitive performance. *Proceedings of the National Academy of Sciences*, 115(37):9193–9197.
- Zimmer, A. and Koch, N. (2017). Fuel consumption dynamics in europe: Tax reform implications for air pollution and carbon emissions. *Transportation Research Part A: Policy and Practice*, 106:22–50.
- Zivin, J. G., Neidell, M., and Schlenker, W. (2011). Water quality violations and avoidance behavior: Evidence from bottled water consumption. *American Economic Review*, 101(3):448–453.
- Zou, E. Y. (2021). Unwatched pollution: The effect of intermittent monitoring on air quality. *American Economic Review*, 111(7):2101–2126.
- Zvereva, E. L. and Kozlov, M. V. (2010). Responses of terrestrial arthropods to air pollution: a meta-analysis. *Environmental Science and Pollution Research*, 17:297–311.