DIGITALES ARCHIV

ZBW – Leibniz-Informationszentrum Wirtschaft ZBW – Leibniz Information Centre for Economics

Boss, Konstantin; Sileci, Lorenzo; Testa, Alessandra et al. Other Persons: Dalmazzone, Silvana

Conference Paper Climate economics and (its) knowledge

Provided in Cooperation with:

Centro di Ricerca e Documentazione Luigi Einaudi

Reference: Boss, Konstantin/Sileci, Lorenzo et. al. (2024). Climate economics and (its) knowledge. Torino: Centro di Ricerca e Documentazione Luigi Einaudi. https://www.centroeinaudi.it/images/abook_file/Quaderno_Giorgio_Rota_12.pdf.

This Version is available at: http://hdl.handle.net/11159/701640

Kontakt/Contact

ZBW – Leibniz-Informationszentrum Wirtschaft/Leibniz Information Centre for Economics Düsternbrooker Weg 120 24105 Kiel (Germany) E-Mail: rights[at]zbw.eu https://www.zbw.eu/

Standard-Nutzungsbedingungen:

https://savearchive.zbw.eu/termsofuse

Dieses Dokument darf zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden. Sie dürfen dieses Dokument nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, aufführen, vertreiben oder anderweitig nutzen. Sofern für das Dokument eine Open-Content-Lizenz verwendet wurde, so gelten abweichend von diesen Nutzungsbedingungen die in der Lizenz gewährten Nutzungsrechte.

Terms of use:

This document may be saved and copied for your personal and scholarly purposes. You are not to copy it for public or commercial purposes, to exhibit the document in public, to perform, distribute or otherwise use the document in public. If the document is made available under a Creative Commons Licence you may exercise further usage rights as specified in the licence.







QUADERNI DEL PREMIO «GIORGIO ROTA»

N. 12, 2024

CLIMATE ECONOMICS AND (ITS) KNOWLEDGE

Iniziativa realizzata con il sostegno di



Gli autori di questo Quaderno:

Konstantin Boss

Universitat Autònoma de Barcelona, Barcelona School of Economics

Lorenzo Sileci

Grantham Research Institute on Climate Change and the Environment, London School of Economics and Political Science

Alessandra Testa

Università di Torino e Collegio Carlo Alberto, Università di Bologna e Universitat Autònoma de Barcelona

Costanza Tomaselli

Imperial College Business School Department of Economics

Silvana Dalmazzone

Dipartimento di Economia e Statistica "Cognetti de Martiis", Università di Torino

Centro di Ricerca e Documentazione Luigi Einaudi Corso Re Umberto 1 • 10121 Torino • segreteria@centroeinaudi.it www.centroeinaudi.it

Copyright © 2024 by Centro di Ricerca e Documentazione Luigi Einaudi, Torino. Tutti i diritti sono riservati. Nessuna parte di questa pubblicazione può essere fotocopiata, riprodotta, archiviata, memorizzata o trasmessa in qualsiasi forma o mezzo – elettronico, meccanico, reprografico, digitale – se non nei termini previsti dalla legge che tutela il diritto d'autore.

Copyright © 2024 by Centro di Ricerca e Documentazione Luigi Einaudi, Turin. All rights reserved. No part of this publication may be photocopied, reproduced, stored, memorized or transmitted in any form or by any means - electronic, mechanical, reprographic, digital – other than in accordance with the law protecting copyright.





TABLE OF CONTENTS

Il Premio «Giorgio Rota» / The «Giorgio Rota» Award	5
Chi era Giorgio Rota / Giorgio Rota's Profile	9
Silver a Dalmarana	11
Silvana Dalmazzone	1.1
Introduction	
Lorenzo Sileci	
Carbon Pricing with Regressive Co-Benefits: Evidence from	
British Columbia's Carbon Tax	23
1. Introduction	23
2. Policy context, data and descriptive statistics	27
3. Empirical strategy	32
4. Results	34
5. Robustness checks	37
6. Mechanisms	40
7. Environmental justice dynamics	44
8. Health gains	49
9. Conclusions	53
References	54
Appendix	59
Alessandra Testa e Konstantin Boss	
What Goes Around Comes Around: The US Climate-Economic Cycle	83
1. Introduction	83
2. Data	86
3. Econometric methodology	89
4. Results	95
5. Discussion	105
6. Conclusion	111
References	112
Appendix	117

Costanza Tomaselli

Green or Greed? Unveiling The Environmental Impact of Market	
Consolidation on Carbon Emissions	125
1. Introduction	125
2. Literature review	128
3. Theoretical framework	130
4. Data and descriptive statistics	135
5. Empirical methodology and results	137
6. Conclusions	144
References	146



IL PREMIO «GIORGIO ROTA»

L'intento del Premio «Giorgio Rota» è di riprendere l'attività di ricerca annualmente condotta dal Comitato / Fondazione Giorgio Rota prima della sua inclusione nel Centro Einaudi, sulla relazione tra il pensiero e l'agire economico e un aspetto (ogni anno diverso) del vivere in società, mantenendo vivo il ricordo e l'insegnamento dell'economista Giorgio Rota, uno dei primi animatori del Centro, prematuramente scomparso.

Dal 2012 il Centro Einaudi ha dunque raccolto questa eredità rinnovando la formula della ricerca: è stato perciò istituito questo premio annuale dedicato a giovani ricercatori, con una qualificazione accademica nei campi dell'economia, sociologia, geografia, scienza politica o altre scienze sociali. I paper possono essere presentati sia in italiano che in inglese, e non devono essere stati pubblicati prima della data della Conferenza Rota, l'evento pubblico nel quale i vincitori hanno modo di presentare il loro lavoro.

La prima edizione aveva per tema *Contemporary Economics and the Ethical Imperative* e la Conferenza Giorgio Rota si è tenuta presso il Centro Einaudi il 25 marzo 2013 con keynote speech di Alberto Petrucci, LUISS Guido Carli, Roma.

La seconda edizione è stata su *Creative Entrepreneurship and New Media con Conferenza* Giorgio Rota presso il Centro Einaudi, 14 aprile 2014 e keynote speech di Mario Deaglio, Università di Torino.

La terza edizione ha analizzato il tema *The Economics of Illegal Activities and Corruption*, con Conferenza Giorgio Rota presso il Centro Einaudi, 15 giugno 2015. Keynote speech di Friedrich Schneider, Johannes Kepler University (Linz, Austria).

La quarta edizione verteva su *The Economics of Migration*. Il 20 giugno 2016 si è tenuta la Conferenza Giorgio Rota presso il Campus Luigi Einaudi, in collaborazione con FIERI. Keynote speech di Alessandra Venturini, Università di Torino. Dal 2016 inoltre il Premio è sostenuto dalla Fondazione CRT.

La quinta edizione trattava di *Economic Consequences of Inequality*, e i saggi vincitori sono stati presentati alla Conferenza Giorgio Rota del 4 maggio 2017, tenutasi presso il Campus Einaudi in collaborazione con il Dipartimento di Economia e Statistica "Cognetti de Martiis". L'Introduzione è di Andrea Brandolini, Banca d'Italia.



La sesta edizione del Premio è incentrata sul tema *The Economics of Health and Medical Care*. I paper vincitori sono stati presentati alla Conferenza Giorgio Rota tenutasi il 1° giugno 2018 presso il Campus Einaudi, in collaborazione con il Dipartimento di Economia e Statistica "Cognetti de Martiis". L'Introduzione è di Fabio Pammolli, Politecnico di Milano.

La settima edizione del Premio è incentrata sul tema Rural Economies, Evolutionary Dynamics and New Paradigms. I paper vincitori, riportati qui, sono stati presentati alla Conferenza Giorgio Rota il 6 maggio 2019 presso il Campus Einaudi, in collaborazione con il Dipartimento di Economia e Statistica "Cognetti de Martiis". Gli autori sono introdotti da un intervento di Donatella Saccone, docente di Economia politica all'Università di Scienze gastronomiche di Bra.

Digital Transformation: Analysis of Economic Impact and Potential è il titolo dell'ottava edizione del Premio. I paper vincitori sono stati presentati alla Conferenza Giorgio Rota l'11 maggio 2020 che a causa della pandemia da Covid si è tenuta online, in collaborazione con il Dipartimento di Economia e Statistica "Cognetti de Martiis". Gli autori sono stati introdotti alla Conferenza e nel volume da un intervento di Pietro Terna, ex Professore ordinario di Economia dell'Università di Torino e consigliere Centro Einaudi.

La nona edizione del Premio è stata sul tema *Main Economic Tendencies in the Contemporary World Economy*. I paper sono stati presentati il 26 maggio 2021 alla Conferenza Giorgio Rota che si è ancora tenuta per via telematica. Gli autori sono introdotti nel volume da un con tributo di Jack Birner, Università di Trento e Comitato scientifico del Centro Einaudi.

La decima edizione del Premio aveva per titolo *Labor, value, robots.* I paper vincitori, durante la conferenza tenutasi il 18 maggio 2022 al Campus Luigi Einaudi, sono stati presentati da Elisabetta Ottoz – direttrice del Dipartimento di Economia e Statistica "Cognetti de Martiis" dell'Università di Torino – che introduce anche questo volume.

Urban Economies as Complex Systems è il titolo dell'undicesima edizione del Premio, i cui vincitori – Luca Favero, Ilaria Malisan, Giacomo Rosso e Léa Bou Sleiman – sono stati premiati in occasione della XI Giorgio Rota Conference il 30 maggio 2023 al Campus Luigi Einaudi. Il volume che raccoglie i saggi vincitori è introdotto da Francesca Silvia Rota dell'Università di Torino IRCrES CNR.

I paper vincitori della dodicesima edizione del Premio – presentati alla Conferenza Giorgio Rota tenutasi il 15 maggio 2024 al Campus Luigi Einaudi – sono raccolti in questo volume, dal titolo *Climate Economics and (its) Knowledge.* I saggi di Lorenzo Sileci, Alessandra Testa e Konstantin Boss, e Costanza Tomaselli sono introdotti da una presentazione di Silvana Dalmazzone del Dipartimento di Economia e Statistica "Cognetti de Martiis" dell'Università di Torino.

THE «GIORGIO ROTA» AWARD

The intent of the «Giorgio Rota» Best Paper Award is to resume the research activity annually conducted by the Giorgio Rota Committee/Foundation before its inclusion in the Centro Einaudi. The focus is on the relationship between economic thought and action and a different aspect of living in society, keeping alive the memory and teaching of economist Giorgio Rota, one of the early members of the Centro, who died prematurely.

Since 2012, the Centro Einaudi has therefore taken up this legacy by renewing the research formula: this annual prize dedicated to young researchers with an academic qualification in the fields of economics, sociology, geography, political science or other social sciences has therefore been established. Papers may be submitted either in Italian or English, and must not have been published before the date of the Rota Conference, the public event at which the winners have the opportunity to present their work.

The first edition's theme was Contemporary Economics and the Ethical Imperative and the Giorgio Rota Conference was held at the Centro Einaudi on 25 March 2013 with keynote speech by Alberto Petrucci, LUISS Guido Carli, Rome.

The second edition, was on Creative Entrepreneurship and New Media with Conference Giorgio Rota at Centro Einaudi, 14 April 2014 and keynote speech by Mario Deaglio, University of Turin.

The third edition analysed the topic The Economics of Illegal Activities and Corruption, with Giorgio Rota Conference at Centro Einaudi, 15 June 2015. Keynote speech by Friedrich Schneider, Johannes Kepler University (Linz, Austria).

The fourth edition focused on The Economics of Migration. The Giorgio Rota Conference was held on 20 June 2016 at the Einaudi Campus, in cooperation with FIERI. Keynote speech by Alessandra Venturini, University of Turin. Since 2016, the Prize has also been supported by the Fondazione CRT.

The fifth edition dealt with Economic Consequences of Inequality, and the winning essays were presented at the Giorgio Rota Conference on 4 May 2017, held at the Einaudi Campus in collaboration with the Department of Economics and Statistics 'Cognetti de Martiis'. Introduction by Andrea Brandolini, Bank of Italy.

The sixth edition of the Prize, held in 2018, focused on the theme: The Economics of Health and Medical Care. The winning papers were presented at the Giorgio Rota Conference held on 1 June 2018 at the Einaudi Campus, in collaboration with the 'Cognetti de Martiis' Department of Economics and Statistics. Introduction by Fabio Pammolli, Politecnico of Milan.



The seventh edition of the Prize focuses on the theme Rural Economies, Evolutionary Dynamics and New Paradigms. The winning papers were presented at the Giorgio Rota Conference on 6 May 2019 at the Einaudi Campus, in collaboration with the 'Cognetti de Martiis' Department of Economics and Statistics. Introductory talk by Donatella Saccone, Professor of Political Economy at the University of Gastronomic Sciences in Bra.

Digital Transformation: Analysis of Economic Impact and Potential is the title of the eighth edition of the Award. The winning papers were presented at the Giorgio Rota Conference on 11 May 2020, which was held online due to the Covid pandemic, in collaboration with the 'Cognetti de Martiis' Department of Economics and Statistics. The authors were introduced at the conference and in the volume by a speech by Pietro Terna, former Professor of Economics at the University of Turin and Centro Einaudi advisor.

The ninth edition of the Award was on the theme Main Economic Tendencies in the Contemporary World Economy. The papers were presented on 26 May 2021 at the Giorgio Rota Conference online. The authors are introduced in the volume by a contribution by Jack Birner, University of Trento and Centro Einaudi Scientific Committee.

The tenth edition of the Prize was entitled Labor, value, robots. The winning papers, during the conference held on 18 May 2022 at the Einaudi Campus, were presented by Elisabetta Ottoz – Director of the Department of Economics and Statistics 'Cognetti de Martiis' at the University of Turin – who also introduced this volume.

Urban Economies as Complex Systems is the title of the eleventh edition of the Giorgio Rota Prize, whose winners — Luca Favero, Ilaria Malisan, Giacomo Rosso and Léa Bou Sleiman — were awarded at the XI Giorgio Rota Conference on 30 May 2023 at the Luigi Einaudi Campus. The volume collecting the winning essays is introduced by Francesca Silvia Rota of the University of Turin and IRCrES CNR.

The winning papers of the twelfth edition of the Prize – presented at the Giorgio Rota Conference held on 15 May 2024 at the Luigi Einaudi Campus – are collected in this volume, entitled Climate Economics and (its) Knowledge. The essays by Lorenzo Sileci, Alessandra Testa and Konstantin Boss, and Costanza Tomaselli are introduced by Silvana Dalmazzone Department of Economics and Statistics 'Cognetti de Martiis' at the University of Turin.





Giorgio Rota (1943-1984) è stato professore di Economia politica presso l'Università di Torino e consulente economico. Per il Centro Einaudi, è stato coordinatore agli studi e membro del comitato di direzione di «Biblioteca della libertà».

Le sue pubblicazioni scientifiche abbracciano diversi temi: l'economia dei beni di consumo durevoli, l'economia del risparmio, il mercato monetario e finanziario, l'inflazione e la variazione dei prezzi relativi, il debito pubblico. Ricordiamo tra esse: Struttura ed evoluzione dei flussi finanziari in Italia: 1964-73 (Torino, Editoriale Valentino, 1975); L'inflazione in Italia 1952/1974 (Torino, Editoriale Valentino, 1975); nei «Quaderni di Biblioteca della libertà», Passato e futuro dell'inflazione in Italia

(1976) e Inflazione per chi? (1978); Che cosa si produce come e per chi. Manuale italiano di microeconomia, con Onorato Castellino, Elsa Fornero, Mario Monti, Sergio Ricossa (Torino, Giappichelli, 1978; seconda edizione 1983); Investimenti produttivi e risparmio delle famiglie (Milano, «Il Sole 24 Ore», 1983); Obiettivi keynesiani e spesa pubblica non keynesiana (Torino, 1983).

Tra le sue ricerche va particolarmente citato il primo Rapporto sul risparmio e sui risparmiatori in Italia (1982), risultato di un'indagine sul campo condotta da BNL-Doxa-Centro Einaudi, le cui conclusioni riscossero notevole attenzione da parte degli organi di stampa. Da allora il Rapporto sul risparmio, ora Indagine sul risparmio, continua a essere pubblicato ogni anno.

GIORGIO ROTA'S PROFILE

Giorgio Rota (1943-1984) was a professor of Political Economy at the University of Turin and an economic consultant. For the Centro Einaudi, he was coordinator of the Study Committee and member of the editorial board of «Biblioteca della libertà».

His scientific publications cover various topics: the economics of consumer durables, the economics of savings, the money market and the financial market, inflation and public debt. Among his publications: Struttura ed evoluzione dei flussi finanziari in Italia: 1964-73 (Turin, Editoriale Valentino, 1975); L'inflazione in Italia 1952/1974 (Turin, Editoriale Valentino, 1975); in «Quaderni di Biblioteca della libertà»: Passato e futuro dell'inflazione in Italia (1976) and Inflazione per chi? (1978); Che cosa si produce come e per chi. Italian Handbook of Microeconomics, with Onorato Castellino, Elsa Fornero, Mario Monti, Sergio Ricossa (Turin, Giappichelli, 1978; second edition 1983); Productive Investments and Household Savings (Milan, «Il Sole 24 Ore», 1983); Keynesian Objectives and Non-Keynesian Public Expenditure (Turin, 1983). Particular mention must be made of the first Report on Savings and Savers in Italy (1982), the result of a field survey conducted by BNL-Doxa-Centro Einaudi, whose conclusions received considerable attention from the press. Since then, the Savings Report, now Report on the Italians' Savings and Financial Choices, has continued to be published every year.







Giorgio Rota (1943-1984) è stato professore di Economia politica presso l'Università di Torino e consulente economico. Per il Centro Einaudi, è stato coordinatore agli studi e membro del comitato di direzione di «Biblioteca della libertà».

Le sue pubblicazioni scientifiche abbracciano diversi temi: l'economia dei beni di consumo durevoli, l'economia del risparmio, il mercato monetario e finanziario, l'inflazione e la variazione dei prezzi relativi, il debito pubblico. Ricordiamo tra esse: Struttura ed evoluzione dei flussi finanziari in Italia: 1964-73 (Torino, Editoriale Valentino, 1975); L'inflazione in Italia 1952/1974 (Torino, Editoriale Valentino, 1975); nei «Quaderni di Biblioteca della libertà», Passato e futuro dell'inflazione in Italia

(1976) e Inflazione per chi? (1978); Che cosa si produce come e per chi. Manuale italiano di microeconomia, con Onorato Castellino, Elsa Fornero, Mario Monti, Sergio Ricossa (Torino, Giappichelli, 1978; seconda edizione 1983); Investimenti produttivi e risparmio delle famiglie (Milano, «Il Sole 24 Ore», 1983); Obiettivi keynesiani e spesa pubblica non keynesiana (Torino, 1983).

Tra le sue ricerche va particolarmente citato il primo Rapporto sul risparmio e sui risparmiatori in Italia (1982), risultato di un'indagine sul campo condotta da BNL-Doxa-Centro Einaudi, le cui conclusioni riscossero notevole attenzione da parte degli organi di stampa. Da allora il Rapporto sul risparmio, ora Indagine sul risparmio, continua a essere pubblicato ogni anno.

GIORGIO ROTA'S PROFILE

Giorgio Rota (1943-1984) was a professor of Political Economy at the University of Turin and an economic consultant. For the Centro Einaudi, he was coordinator of the Study Committee and member of the editorial board of «Biblioteca della libertà».

His scientific publications cover various topics: the economics of consumer durables, the economics of savings, the money market and the financial market, inflation and public debt. Among his publications: Struttura ed evoluzione dei flussi finanziari in Italia: 1964-73 (Turin, Editoriale Valentino, 1975); L'inflazione in Italia 1952/1974 (Turin, Editoriale Valentino, 1975); in «Quaderni di Biblioteca della libertà»: Passato e futuro dell'inflazione in Italia (1976) and Inflazione per chi? (1978); Che cosa si produce come e per chi. Italian Handbook of Microeconomics, with Onorato Castellino, Elsa Fornero, Mario Monti, Sergio Ricossa (Turin, Giappichelli, 1978; second edition 1983); Productive Investments and Household Savings (Milan, «Il Sole 24 Ore», 1983); Keynesian Objectives and Non-Keynesian Public Expenditure (Turin, 1983). Particular mention must be made of the first Report on Savings and Savers in Italy (1982), the result of a field survey conducted by BNL-Doxa-Centro Einaudi, whose conclusions received considerable attention from the press. Since then, the Savings Report, now Report on the Italians' Savings and Financial Choices, has continued to be published every year.



SILVANA DALMAZZONE


INTRODUCTION

Thank you for inviting me to open this year's Giorgio Rota Conference dedicated to climate change economics and its knowledge.

What I'm going to do is to offer you a feeling of what's inside the larger, fast-evolving box of climate change economics. A very recent survey, conducted in 2023 by the Policy Institute of the King's College in London as part of the PERITIA Horizon project, investigating public perceptions about climate change in six countries in Europe plus the UK, confirms that there are very wide misperceptions on the scientific knowledge of climate changes.

Regarding the survey's question "to the best of your knowledge, what percentage of climate scientists have concluded that human-caused climate change is happening?", we see that according to a large sample of interviewed Europeans the average proportion of scientists convinced that climate change is taking place and is human-induced is around 68%: enormously lower than the reality which is today of 99.9%.

FIGURE 1 • REPLIES TO PERITIA SURVEY'S QUESTION "TO THE BEST OF YOUR KNOWLEDGE, WHAT PERCENTAGE OF CLIMATE SCIENTISTS HAVE CONCLUDED THAT HUMAN-CAUSED CLIMATE CHANGE IS HAPPENING?"



Source: https://peritia-trust.eu/.



From another question it emerges that, on average, three-quarters of people, about 74%, say that climate change is mainly caused by human activities, which means that 26% think that it is not. Things go a little bit better in Italy, where the percentage is about 82%. Surprisingly, they go worse in Northern European countries, with 61% of the Norwegian respondents being convinced climate change is human-caused.

Despite remaining misperception and lack of knowledge, the large majority of people are however worried about the impact of global warming on future generations: 81% of people on average, according to PERITIA's survey. The answers also show that 80% of interviewed individuals say they are worried about the impact of global warming on humanity in general. Most people also think climate change is harmful now or will be harmful within the next 10 years. And a remarkable 62% of people say so not just about the world, including developing countries, but also about their own country.

49% Total 18% 16% 59% 15% 2% 5% 2% 1% Italy It is causing harm right now Ireland 52% 18% 15% Within the next 10 years Within the next few decades Germany 52% 18% 15% 3% 6% Don't know Whitin the next century In a century or more Poland 48% 21% 16% 46% 19% 16% 36% 7% 13%

FIGURE 2 • REPLIES TO PERITIA'S SURVEY'S QUESTION "WHEN, IF EVER, DO YOU THINK CLIMATE CHANGE WILL START TO HARM HUMANITY?"

Source: https://peritia-trust.eu/.

So, the perception is that there is a serious problem. Yet, actions to mitigate climate change don't match the public perception of how serious the problem is under current policies. There is a remarkable emission gap between the current policies scenario and the emission pathways that would keep global warming under the stated objectives of the last ten years' international agreements (Figure 3).






FIGURE 3 • GLOBAL GHG EMISSIONS UNDER DIFFERENT SCENARIOS AND THE EMISSIONS GAP IN 2030 AND 2035 (MEDIAN ESTIMATE AND TENTH TO NINETIETH PERCENTILE RANGE)

Source: UNEP, 2023.

The light turquoise line shows the trajectory we need to remain within +2°C of warming (Paris Agreement). The orange trajectory is the one to stay within the +1.8°C warming. And the last one, the dark turquoise one, is the +1.5°C trajectory. Under current policies, we are heading in the graph far higher than all these three trajectories. This means that the emission gap is pretty large and remains larger even with the new pledges made by countries under the Glasgow Climate Pact, plus all the officially announced mitigation pledges for 2030, which were added later on.

Until last year, the Glasgow Climate Pact plus all the official 2030 mitigation pledges reduced the projected emissions to 2030 indicated in the previous unconditional Nationally Determined Contributions (NDCs) by only 7.5%, whereas staying within the Paris Agreement's objectives of +2°C and +1.8°C would require a reduction of 30% with respect to previous pledges. The Glasgow objective of +1.5°C of warming would require a 55% reduction. Indeed, we are moving very, very shyly.



So, why so little? Why so late? Economics may help. Climate change economics does make an effort to help us understand why.

The first reason we are doing so little and going so slowly is obviously the real cost of leaving fossil fuels in terms of jobs, purchasing power, and welfare effects.

A second important answer is that climate change is a global problem, but damage is disjoint from the location of emission sources. The incentive for individuals, and even for individual countries, to refrain from emitting greenhouse gases is very small, and the effort of one country in isolation can be nullified by unmitigated growth of emissions in another country. In addition, countries that unilaterally engage and adopt stricter regulations may face carbon leakage and lose competitiveness in international markets.

A third reason why it is so difficult to act on mitigating climate change has to do with the fact that this particular environmental issue has very intricate distributive implications, between the North and the South and between the West and the rest of the world. Much more intricate, for example, than the ozone layer problem, which we managed to tackle quite effectively and quickly through international cooperation. Industrialized Western countries are mostly responsible for the cumulative emissions that over the last 150 years caused today's greenhouse gas concentrations in the atmosphere. Furthermore, developing countries are going to suffer most and first from the impacts. So, the developing world has so far been very reluctant to accept restrictions. But many of today's richest countries resist the idea of bearing all the costs without the participation of at least those giant Asian countries that now challenge Western economies on world markets.

The last answer that climate change economics is putting forward to explain why it is so difficult to act, pertains to human psychology. Human psychology has an incurable bias towards the present. Decades of behavioral and experimental economics have consistently found evidence of time inconsistency and the use of irrationally high discount rates in people's welfare. In making decisions with impacts that are obviously adverse to people's welfare, we know that individuals are willing to accept smaller immediate rewards over larger delayed rewards, even when the delayed rewards are objectively more valuable. Irrationally high discount rates, or the dictatorship of the present, emerge in individual preferences, corporate decisions, and political processes. For example, we have to force compulsory saving for retirement. There are issues of competitiveness that induce very short-sighted corporate strategies. Even in political processes, which is where a collective and long-term perspective should prevail, there is a structure of rewards that has a short-term view and the pressure from financial markets is high.



Therefore, if we humans find it so difficult even to act rationally in front of our now-to-do questions, such as doing our homework, let's imagine how it works when facing a massive ethical transformation of our socioeconomic system.

Nevertheless, there is a lot that climate change economics can do for us in matter of designing and evaluating climate policies.

First, cost-benefit analyses: we need to know the cost of carrying on business-as-usual and bearing the impacts, and on the other hand the cost of doing what it takes to avoid (or limit) those impacts. Economic assessments of the forecasted socioeconomic impacts of climate change rely on several different approaches. Older studies tend to be enumerative: they consider the largest feasible number of expected impacts in terms of their physical units, multiply them by the evaluated unit cost, and add them up.

Econometric studies make a step forward, allowing us to take into account interactions (for example, due to price changes) and the dynamic aspect of socioeconomic impacts of climate change in the future. A limitation is that they assume that differences in climate existing now between different places, at different latitudes, are a good proxy for differences in climate that will emerge in the future. One advantage is that they do not have to assume anything about, for example, adaptation behaviors, because they observe what has actually taken place over the past in different parts of the world.

More and more popular in recent years are also the socioeconomic valuations of the impacts of climate change based on computable general equilibrium (CGA) models. These have the advantage of being able to look at the whole economic system and include dynamic changes and interactions between prices and sectors. One issue is that they tend to be based on national accounts, and hence they tend to omit impacts on human health and on ecosystems. Also, they tend to express the results in terms of percentage loss of consumption or GDP, leaving aside other kinds of welfare impacts.


Another important part of the literature relies on integrated assessment models (IAMs), simplified mathematical descriptions of reality that integrate knowledge from two or more disciplinary domains (e.g., climate sciences and economics). They constitute the base, for example, of the evaluations included in the IPCC assessment records.

Lastly, elicitation studies are simply evaluations built on large sets of interviews with experts, investigating and statistically describing their opinions on the dimensions of expected future impacts of climate change.



Recent meta-analyses, for example the one published by Tol (2024) on *Energy Policy*, produces estimates of average negative impacts from all these different approaches and estimates a GDP impact from climate change by 2050 in the range of -2% and -3% losses, if we stay within a $+2.5^{\circ}$ C increase of the temperature; up to -11% by 2100, assuming we are going to stay under a $+3.0^{\circ}$ C degrees threshold by that time.

FIGURE 4 • THE ECONOMIC IMPACT OF CLIMATE CHANGE FOR A 2.5 °C WARMING RELATIVE TO PRE-INDUSTRIAL TIMES FOR COUNTRIES AS A FUNCTION OF THEIR INCOME (TOP PANEL) AND TEMPERATURE (BOTTOM PANEL).



Source: Tol, 2024: 11.



These results are very sensitive to discount rates, and they generally do not include catastrophic, acute risks (such as those from extreme weather events), but only the chronic impact of raising temperatures on factors productivity. Estimates by insurance companies, which do take into account catastrophic risks, tend to be much higher.

A recent estimate by Swiss Re Institute¹ points a loss of 14% of global GDP by 2050 under a +2.6°C warming scenario and of –18% of global GDP by 2050 under a +3.2°C scenario, and a loss for the European Union between 8 and 10.5% of GDP by 2050.

Certainly, these are global average values. If we look at how these losses are going to be distributed across the world, we find very large disparities between high-income countries and developing countries. Figure 4 (upper diagram), which on the x axis measures GDP per capita and on the y axis the percentage losses of GDP, shows how the losses would be scattered. Countries of the Global South are going to suffer up to –20, –25, and –30% loss of GDP. Poorest countries are going to suffer impacts much higher than the world average. An analogous result emerges from the bottom diagram, showing the correlation between expected percentage losses of GDP and a country's average temperature. Again, cooler countries (the Global North, to simplify a bit) tend to have smaller losses or even moderate gains, whereas warmer and tropical countries suffer much more severe losses.

Mitigation can make a difference. For Europe, decarbonization policies limiting warming to +2°C degrees would reduce the welfare losses by 70% compared to a +3.0 °C degree scenarios, while limiting warming to +1.5 °C would lower welfare losses by 90%. We are late and slow, but still on time to avoid the worst.


Decarbonizing our economies involves mitigation direct costs as well as technological, sectoral, and macroeconomic welfare costs, whose assessment relies, again, on an array of methods and models. An important one is the marginal abatement cost of carbon (MACC) curves. MACC curves look like the diagram in Figure 5. The diagram estimates the marginal cost of CO_{2e} emission abatement by technology, which is the height of the column, associated with the relative reduction potential, which is measured by the width of each column. All the technologies that you see in the lower, left-hand part of the diagram are associated to negative costs. It means that on the lifetime scale of these technologies the mitigation cost is negative, i.e., the investment would imply a net saving. Almost a quarter of the total abatement potential required to limit global warming under +2°C could be gained through measures with a zero or negative net life cycle cost. These include, for instance, more efficient lighting systems, motor system efficiency, installation of retrofits, and so on. Conversely, the technologies in the upper,

 $^{^1\} https://www.swissre.com/media/press-release/nr-20210422-economics-of-climate-change-risks.html$



right-hand part of the diagram cost more and more. Technological change is moving fast, and today's most costly technologies that we need to adopt to reach the +2°C objective have a cost of around 40 euro per ton of avoided CO_{2e} emissions.


FIGURE 5 • GLOBAL GHG ABATEMENT COST CURVE BEYOND BUSINESS-AS-USUAL (BAU) - 2030/US ENERGY SYSTEM MARGINAL ABATEMENT CURVE



Source: McKinsey, 2009: 7.

The diagram in Figure 6 comes from the last the 6th assessment report of IPCC, which offers an overview of mitigation options and their estimated ranges of costs and potential in 2030.

FIGURE 6 • GLOBAL MITIGATION COSTS FOR 2015 TO 2100 IN NET PRESENT VALUE (NPV)
DISCOUNTED AT A 5% DISCOUNT RATE AND EXPRESSED AS A SHARE OF THE BASELINE ECONOMY



Source: Clarke et al., 2014: 450.



The results in general equilibrium models show, not surprisingly, that both consumption and GDP losses increase as our steady state greenhouse gases concentration goal gets more stringent. Costs would be high for a 430-480 ppm target by 2100, implying, by the end of the century, between 2.0 and 5.7 GDP percentage loss. Slightly more in terms of consumption (between –2.2% and –5.8%).

These are mitigation costs calculated as if we started to reduce global emissions right now. But anyhow, what emerges from a pretty large set of literature is that the global economic benefit of limiting warming to +2°C is going to exceed the cost of mitigation in most of the studies assessed by the IPCC report, unless climate damages are going to materialise at the lower end of the range of possibilities, and unless future damages are discounted at very high rates.

Of course, all these estimates are very sensitive to the discount rate used. Economic models also tell us that mitigation costs increase sharply with delays in mitigation: on average, net mitigation cost increases by approximately 40% for each decade of delay in the moment global emissions will peak and start decreasing. They also tell us how much delay we can afford before missing the current climate targets, and how much costlier would it be to start later and accelerate with more stringent policies in subsequent years.


Another part of the climate change literature looks at adaptation. If we also include adaptation, then we can further reduce the losses. Mitigation and adaptation must be coordinated as they compete in the use of scarce resources. Economic theory tells us that the optimal level of adaptation is the one that equalises the marginal adaptation cost and the marginal adaptation benefits, i.e., where the adaptation curve has a slope of 45 degrees (Figure 7). But again, quantifying adaptation raises a lot of conceptual issues. For example, how should we account for the costs of adaptation if it also has other benefits beyond mitigating climate change, such as health benefits or welfare in other forms? What if it would have taken place anyhow? So, a problem of additionality does emerge.

Nevertheless, certainly also adaptation makes a difference. In the EU, for example, mitigating climate change to +1.5°C would half the damages from coastal and river floods, just to take one of the many forms of impact. Adding adaptation would reduce residual damages from river floods by 40 times and by 20 times those from coastal flooding.



FIGURE 7 • GRAPHICAL REPRESENTATION OF LINK BETWEEN THE COST OF ADAPTATION (ON THE X-AXIS) AND THE RESIDUAL COST OF CLIMATE CHANGE (ON THE Y-AXIS).

The left panel represents a case where full adaptation is possible, while the right panel represents a case in which there are unavoidable residual costs



Source: Chambwera et al., 2014: 953.

Setting objectives and setting targets is another crucial question. Where do we want to go? How ambitious need we to be in trying to mitigate climate change? Standard economics tells us that we should choose decarbonization targets so as to equalise marginal benefit and marginal damage. This would require, in practise, knowledge of marginal mitigation costs up to the different moments in time, of marginal economic value of impacts, and a collective choice of the discount rate to be used. All this kind of knowledge, which is, as we have seen, very difficult to quantify, makes it difficult to identify analytically at which concentration level, and therefore, given the decay rate, at which emission level, the discounted damage of an additional unit of GHG gases equals the cost of containment. This has been the object of debate in all international climate change conferences.

Eventually, what has been done in the real world is mostly to revert and define the objectives in terms of warming. We first need to decide whether, given the forecasted impact, we aim to stay within +2°C or +1.5°C warming. Then we convert those objectives into corresponding concentrations and thus of remaining carbon budget: to have a 50% chance of staying below 1.5°C, we can only emit 250 billion more tonnes of carbon. That's just six years of our current



emissions. For a 50% chance of staying below 2°C, the world could emit 1150 billion tonnes, around 28 years of current emissions.

Uncertainty plays a fundamental role in this, yet we often disregard it. The acceptable probability of meeting (or missing) the desired outcome makes a huge a difference, too. When we talk about staying within a +2.0°C scenario, we always assume to stay within that objective with a given probability. Limiting global warming to +2°C with a 67% probability, or with a 80% probability or with 50% probability, implies a dramatic change in the remaining carbon budget.

Finally – but I will leave this to the presentation of this year's papers winning the 2024 Giorgio Rota Award – climate change economics offers a lot also in terms of designing policy instruments, such as carbon pricing.

In conclusion, a final key point is the question of tackling distributive issues. The previous large scale societal transformations, led by the 19th century industrial revolution and the digital revolution of the last few decades, have been originated by technological change and have been driven by market forces. The ecological transition represents the first time we are facing the task of designing, implementing, and governing a deep, deliberate structural change that is not driven by market forces but will have to be driven by collective action. Also, this structural change is going to take place in a short period: the next 9-10 years, up to 2030, are going to be decisive. These necessary transformations will have no chance of success if they fail to build around them a sufficient strong legitimacy. And this is not going to happen if we are not going to deepen our knowledge of the distributional consequences of climate change and climate change policies, in order to ascertain the political feasibility and foresee the necessary compensatory actions.

REFERENCES

Clarke, L., Jiang, K., Akimoto, K., Babiker, M., Blanford, G., Fisher-Vanden, K., Hourcade, J.-C., Krey, V., Kriegler, E., Löschel, A., McCollum, D., Paltsev, S., Rose, S., Shukla, P.R., Tavoni, M., van der Zwaan, B.C.C., and van Vuuren, D.P., *Assessing Transformation Pathways*, in O. Edenhofer, R. Pichs-Madruga, Y. Sokona, E. Farahani, S. Kadner, K. Seyboth, A. Adler, I. Baum, S. Brunner, P. Eickemeier, B. Kriemann, J. Savolainen, S. Schlömer, C. von Stechow, T. Zwickel and J.C. Minx (eds), *Climate Change 2014: Mitigation of Climate Change*, Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, Cambridge (UK) – New York, Cambridge University Press, http://www.ipcc.ch/pdf/assessment-report/ar5/wg3/ipcc_wg3_ar5_chapter6.pdf



- Chambwera, M., Heal, G., Dubeux, C., Hallegatte, S., Leclerc, L., Markandya, A., McCarl, B.A., Mechler, R., and Neumann, J.E. (2014), *Economics of adaptation*, in C.B. Field, V.R. Barros, D.J. Dokken, K.J. Mach, M.D. Mastrandrea, T.E. Bilir, M. Chatterjee, K.L. Ebi, Y.O. Estrada, R.C. Genova, B. Girma, E.S. Kissel, A.N. Levy, S. MacCracken, P.R. Mastrandrea, and L.L. White (eds), *Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part A: Global and Sectoral Aspects*, Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, Cambridge (UK) New York, Cambridge University Press, pp. 945-977
- McKinsey (2009), Impact of the financial crisis on carbon economics: Version 2.1 of the global greenhouse gas abatement cost curve, https://www.cbd.int/financial/doc/Pathwaystoalowcarboneconomy.pdf Tol, R.S.J. (2024), A meta-analysis of the total economic impact of climate change, «Energy Policy», 185, 113922, https://doi.org/10.1016/j.enpol.2023.113922
- UNEP (2023), Emissions Gap Report 2023, https://www.unep.org/interactives/emissions-gap-report/2023/ section-5

LORENZO SILECI

CARBON PRICING WITH REGRESSIVE CO-BENEFITS: EVIDENCE FROM BRITISH COLUMBIA'S CARBON TAX

Abstract. This paper examines the impacts of carbon taxation on air quality co-benefits and environmental justice. Using high-resolution data and a synthetic difference-in-differences strategy, I find that the 2008 carbon tax in British Columbia has reduced PM_{2.5} emissions by 5.2-10.9%. The flow of monetised co-benefits from climate policy is large, corresponding to 40-81% of annual carbon tax revenues. While pollution reductions arise for all citizens, the tax widens pre-existing disparities in pollution exposure across income and racial diversity categories. The distribution of co-benefits from market-based climate mitigation instruments may be regressive, requiring additional policies targeting environmental inequalities.

Keywords. Carbon tax, air pollution, co-benefits, environmental justice

1. Introduction

The major sources of CO₂ emissions are the fossil fuel combustion processes which also release toxic air pollutants, making climate change and air pollution complementary externalities. Policy efforts to control CO₂ emissions by internalising the social cost of carbon are thus bound to give rise to significant health "co-benefits" associated with air quality improvements, with climate mitigation hailed as "the greatest global health opportunity of the 21st century" (Watts et al., 2015). Moreover, a substantial body of research has documented severe historical inequities in air pollution exposure across income and racial groups (Colmer et al., 2020; Jbaily et al., 2022), and recent work has reported the ambiguous impacts of market-based climate policy in closing these "environmental justice gaps" (Cain et al., 2024). In this paper, I jointly assess the air quality co-benefits and environmental justice implications of carbon taxation, leveraging as a case study the experience of the 2008 carbon tax in British Columbia, Canada.

Given the relative scarcity of long-tenured carbon pricing schemes, it is unsurprising that empirical evidence of their causal impact on local air pollution co-benefits is sporadic, and mostly limited to cap-and-trade schemes (Deschenes *et al.*, 2017; Hernandez-Cortes and Meng, 2023) with fewer studies focusing on fuel taxes (Basaglia *et al.*, 2023). On the contrary, there is a large and growing literature which, using theoretical insights (Parry *et al.*, 2015) and simulation models (Knittel and Sandler, 2011; Zhang *et al.*, 2021), has attempted to calculate the monetary



value of air pollution improvements due to carbon taxation and compare them with the cost of mitigation policies. In particular, net health co-benefits arising from carbon taxation are theorised to reach a high enough magnitude to partially or fully offset the mitigation costs for households at a national (Li *et al.*, 2018; Shindell *et al.*, 2016) and global (West *et al.*, 2013; Vandyck *et al.*, 2018) level, and may provide strong additional incentives for a swift transition to a low-carbon economy¹.

In light of the considerable size of projected air pollution co-benefits, it is fundamental to examine how carbon pricing policies may impact the spatial distribution of pollutants over affected populations, a theme also referred to as the "environmental justice question" (Banzhaf et al., 2019; Currie et al., 2023). While carbon taxation is expected to produce higher pollution reductions in areas with lower marginal abatement costs, this efficiency criterion is blind to equity considerations, and CO₂ abatement is not necessarily perfectly correlated with the dispersion of air pollutants (Hernandez-Cortes and Meng, 2023; Cain et al., 2024). It is thus paramount to inspect whether carbon taxation presents efficiency-equity trade-offs in the distribution of realised co-benefits, evidence of which is not unidirectional in the environmental economics literature (Fowlie et al., 2012; Boyce and Pastor, 2013; Grainger and Ruangmas, 2018; Shapiro and Walker, 2021; Currie et al., 2023; Sheriff, 2024).

The 2008 British Columbian carbon tax, covering approximately 75% of the Canadian province's CO₂ emissions, was initially introduced at a rate of \$10/tCO₂, and sequentially ramped up by \$5 per year until 2012, when it was frozen at \$30/tCO₂ until 2018. Importantly, no other Canadian Province introduced carbon pricing schemes between 2008 and 2018, when the tax was rolled out on a federal basis. I acquire high-resolution data on PM_{2.5}, based on a combination of satellite observations, geo-chemical models and ground-based monitoring stations, from Meng *et al.* (2019) and van Donkelaar *et al.* (2019), and combine them with granular socio-economic data at the Dissemination Area level², retrieved from the Canadian Census at 5-year intervals between 2001 and 2016. I exploit this highly disaggregated dataset to assess the effect of the carbon tax on air pollution co-benefits and the dynamics of the environmental justice gap.





¹ Reductions in morbidity and mortality due to improvements in air quality are likely not to capture the full extent of the local pollution externality: a large body of research has linked air pollution to non-health outcomes (see Aguilar-Gomez *et al.*, 2022, for a review). Studies have linked air pollution to negative educational outcomes (Ebenstein *et al.*, 2016; Wen and Burke, 2022), increase in crime rates (Bondy *et al.*, 2020) and suicides (Persico and Marcotte, 2022), reductions in labour productivity (Graff Zivin and Neidell, 2012), and in housing prices (Sager and Singer, 2024; Freeman *et al.*, 2019), suggesting that any attempt at quantifying the monetary impact of co-benefits based on health outcomes alone would, at best, provide a lower bound of the beneficial consequences of air quality improvements.

² Corresponding roughly to US Census tracts.

The central result of the paper is that the carbon tax has resulted in statistically significant reductions in PM_{2.5} concentrations, with a lower bound average estimate of -0.36 μ g/m³ and an upper bound average estimate of -0.89 μ g/m³, corresponding to a 5.2-10.9% reduction in particulate matter concentrations with respect to pre-treatment average levels. Importantly, this result is obtained by moving away from traditional difference-in-differences estimation, in light of a violation of the foundational parallel trends assumption: particulate matter trends between British Columbian and control Dissemination Areas diverge prior to the implementation of the carbon tax, thereby potentially biasing DID estimates. I rely on a family of estimators related to the synthetic control method (SCM) for comparative case studies (Abadie and Gardeazabal, 2003; Abadie, 2021), employing in particular the synthetic difference-in-differences (SDID) estimator by Arkhangelsky *et al.* (2021) as my preferred methodology.

I subsequently inspect the efficiency-equity trade off, examining whether air pollution reductions arise heterogeneously within British Columbian metropolitan areas. I split the pool of treated units in quintiles of pre-existing pollution, population density, median income levels and racial diversity, and estimate the impact of the tax on PM_{2.5} reductions for each quintile of these characteristics. While Pareto-optimal in the welfare dimension, with reductions in pollution across the board, the carbon tax is regressive in the environmental justice dimension: reductions are 1.6-2.2 times higher in the bottom quintile of pre-treatment air pollution, population density and racial diversity compared to the top quintile, and 1.7 times higher in the top median income quintile compared to the bottom quintile.

Finally, I convert my estimates of particle pollution reductions into mortality reductions³ and associated monetary gains, relying on the concept of the Value of a Statistical Life⁴. The median monetary health gains appear to be large, in the order of \$88-402/year per capita. The central estimate of \$198 is almost double the \$115.50 per capita Low-Income Climate Action Tax Credit, the carbon tax governmental rebate accruing to low-income individuals to mitigate the cost of carbon pricing. The total annual health gains are comparable to annual carbon tax revenues at its inception (Ministry of Finance, 2009) and amount to 40-81% of annual tax revenues at maturity (Ministry of Finance, 2013). Health gains exhibit a positive spatial correlation with income, corroborating the evidence on the increase in the environmental justice gap.

This paper contributes to the literature on three main fronts. First, I extend the recent evidence on the impact of carbon pricing on air pollution co-benefits, by providing the first study with

⁴ Following Fowlie et al. (2019) and Carozzi and Roth (2023).



³ Exploiting hazard rates adapted from the environmental health and epidemiology literature (Lepeule *et al.*, 2012; Krewski *et al.*, 2009).

an explicit empirical focus on carbon taxation instead of the frequently examined cap-and-trade schemes and fuel tax increases (Hernandez-Cortes and Meng, 2023; Basaglia *et al.*, 2023). I overcome known spatial and temporal selection problems connected with the use of sparse air quality monitors (Grainger and Schreiber, 2019) by relying on two sets of remotely sensed PM_{2.5} data (Meng *et al.*, 2019; van Donkelaar *et al.*, 2019) which provide full coverage of the spatial and temporal extent of my dataset. Further, I dispel the notion that the carbon tax has resulted in gasoline to diesel fuel substitution (Saberian, 2017), instead highlighting expected reductions in both fuels' total demand after the tax (Rivers and Schaufele, 2015; Bernard and Kichian, 2019). Moreover, by exploiting highly disaggregated census information on commute mode, I provide evidence on additional mechanisms underlying the air quality improvements: BC residents substitute high emissions trips with public transport and active commute modes following the implementation of the tax. My results are thus also consistent with the findings of Pretis (2022), who found that the 2008 carbon tax reduced CO₂ emissions in the transportation sector alone.

The second contribution regards the growing literature on the relationship between environmental policies and equity. I present the first ex post analysis of the effects of a carbon tax on the environmental justice (EJ) gap. I find that pricing carbon, while giving rise to widespread air quality co-benefits, may do so disproportionately with respect to pre-existing levels of air pollution, income, population density and racial diversity. My estimates thus add a data point to the nascent literature on ex post empirical evaluation of EJ effects from climate policy, which has so far reported mixed evidence (Cain et al., 2024). This result counterbalances some of the recent evidence of EJ implications of market-based instruments (Hernandez-Cortes and Meng, 2023), and highlights the potential for coupling climate mitigation policy with instruments targeting air pollution specifically, which have been shown to be effective in closing EJ gaps (Currie et al., 2023; Sager and Singer, 2024). While it is noteworthy that climate policy can give rise to significant air pollution and associated health co-benefits due to complementarities alone, improvements along the equity axis are not a necessary implication of efficiency-focussed instruments. In order to obtain the greatest gains across multiple independent policy targets, multiple policy instruments may be needed, a notion that economists have considered since the 1950s (Tinbergen, 1952).

Lastly, I contribute to the environmental policy evaluation literature by showing how the traditional DID estimator is susceptible of producing biased estimates, due to substantially diverging pre-treatment trends across treatment and control units. I solve this concern by exploiting SCM and the newly introduced SDID estimator (Arkhangelsky *et al.*, 2021) and exploiting, unlike recent studies in environmental policy evaluation (e.g. Andersson, 2019; Leroutier, 2022; Basaglia *et al.*, 2023) a subnational level treatment and a highly granular framework. In my setting, with multiple treated units and a large number of control units to



draw synthetic counterfactuals from, both the SCM and SDID perform well in addressing concerns about diverging pre-treatment trends and identify robust estimates of the impact of the carbon tax on PM_{2.5} levels, improving substantially upon traditional estimators and aggregate policy settings.

The remainder of the paper begins with a description of the carbon tax and the data sources in Section 2. In Section 3, I present the identification strategy, followed by the main results in Section 4. Section 5 shows the consistency of the main analyses to alternative specifications and mechanisms underlying the results are presented in Section 6. I examine environmental justice dynamics in Section 7, and estimate mortality reductions and associated monetary health gains in Section 8. Section 9 concludes the paper.

2. POLICY CONTEXT, DATA AND DESCRIPTIVE STATISTICS

2.1. The 2008 British Columbian Carbon Tax

The introduction of the British Columbia (BC) carbon tax was formally announced in the provincial budget plan in February 2008, catching the public off guard due to the unexpected nature of this move by the Liberal government (Harrison, 2012; Ahmadi *et al.*, 2022). The policy aimed to reduce emissions by a minimum of 33% below 2007 levels by 2020 (Azevedo *et al.*, 2023). Implemented on July 1, 2008, the initial tax rate was set at \$10/tonne CO₂eq and increased by \$5/tonne CO₂eq annually until it reached \$30 in 2012, establishing one of the highest carbon prices globally at the time (Murray and Rivers, 2015; Azevedo *et al.*, 2023). The carbon tax rate remained at \$30 until 2018, when it increased to \$35, with a subsequent annual increment of \$5 anticipated until it reached \$50/tonne in 2022. The tax, applicable to all fossil fuel purchases in BC, accounts for approximately 77% of the province's total greenhouse gas (GHG) emissions, underscoring the comprehensive scope of the policy (Murray and Rivers, 2015; Rivers and Schaufele, 2015; Ahmadi *et al.*, 2022; Azevedo *et al.*, 2023). Notably, the most affected sector is transportation, which contributed to 43.9% of the province's total CO₂ levels in 2007; exemptions cover exported fuels, non-combustion GHGs (e.g. landfill methane), and emissions generated outside BC⁵.

A key aspect of implementing the BC carbon tax is its commitment to revenue neutrality, serving as a crucial mechanism to secure public support and mitigate resistance to additional taxation, a notable challenge in the execution of carbon pricing schemes (Carattini *et al.*,

⁵ This excludes a significant portion of air transportation and non-metallic mineral manufacturing emissions. Additionally, non-fossil fuel sources like fugitive emissions and chemical processes are exempted, broadening the range of exclusions (Azevedo *et al.*, 2023).



2017; Carattini *et al.*, 2019)⁶. The revenue-neutral design of the tax involved returning funds to consumers and businesses through various means, including direct transfers to low-income individuals, income tax reductions, and corporate tax cuts (Murray and Rivers, 2015; Ahmadi *et al.*, 2022). In particular, the achievement of revenue neutrality in BC involves two primary mechanisms. Firstly, by initiating a 5% reduction in the bottom two income tax brackets, BC secured the lowest income tax rate in Canada for individuals earning up to \$122,000. This reduction was complemented by additional measures such as the "low income climate action" tax credit and the Northern and Rural Homeowner benefit (Azevedo *et al.*, 2023)⁷. Secondly, a series of reductions were applied to the general corporate tax rate, starting at 12% in 2008 and gradually decreasing to 11%, 10.5%, and 10% in 2010 and 2011, before returning to 11% in 2014. Simultaneously, the small business corporate income tax rate decreased from 4.5% to 2.5% in 2008 (Azevedo *et al.*, 2023)⁸. According to the Budget and Fiscal Plan, the carbon tax generated approximately \$1.2 billion in annual revenue since 2012 when the rate stabilized at \$30/tonne *CO*₂eq, with around \$1.4 billion returned to consumers (Ahmadi *et al.*, 2022; Azevedo *et al.*, 2023).

Given the popularity of the carbon tax, it is unsurprising that economists have conducted several analyses of its effectiveness across a range of measures. Focussing on the transport fuel market, Rivers and Schaufele (2015) and Lawley and Thivierge (2018) find 5-8% reductions in gasoline demand due to the tax implementation. Azevedo *et al.* (2023) investigate the employment response to the tax: the absence of aggregate effects masks heterogeneous impacts, with large emission-intensive firms negatively affected and small businesses benefitting from the policy. In terms of global pollutants, Ahmadi *et al.* (2022) detect emissions reductions in the manufacturing sector, while the multisectoral analysis of Pretis (2022) identifies significant reductions in transportation emissions with negligible effects on the remaining sectors of the economy.

⁸ Since 2008, various tax credits, ranging from the BC Seniors Home Renovation Tax Credit to the Film Incentive BC tax credit, have been implemented, contributing to the revenue redistribution.



⁶ Subsequent to the initial "Axe the tax" campaigns leading up to the 2009 provincial elections, polling data indicated a sustained increase in public approval of the tax until 2015 (Murray and Rivers, 2015). However, after 2012, there was a shift towards earmarking some revenues for specific sectors, creating a mixed system of redistribution (Murray and Rivers, 2015). Public opinion on the carbon tax was initially volatile, with campaigns against it leading up to the 2009 provincial elections, but sustained approval was observed until 2015 (Murray and Rivers, 2015). Recent studies, though, suggest that attitudes towards carbon pricing may be more influenced by partisan identities than updated information about potential rebates (Mildenberger *et al.*, 2022).

⁷ The low income climate action tax credit was initially set as \$100 per adult plus \$30 per child, and subsequently raised to \$115.50 per adult and \$34.50 per child (Ministry of Finance, 2009; Ministry of Finance, 2013). The Northern and Rural Homeowner Benefit amounts to \$200 but only applies to howeowners in areas outside the Capital (Victoria CMA), Greater Vancouver (Vancouver CMA) and Fraser Valley (Abbotsford CMA) regional districts. The appropriate rebate to compare to health gains is thus the low income climate action tax credit.

2.2. Data and descriptive statistics

In order to analyse the effect of British Columbia's 2008 carbon tax on air quality, I assemble and process information on local pollutants' concentrations, geographic characteristics, and socio-economic dynamics from multiple sources. The observational units used in the analysis are Dissemination Areas (DAs), the smallest standard geographic areas for which Canadian census data are disseminated. Since the paper is concerned with analysing the effect of carbon pricing on air quality in cities, I restrict the geographic scope of the dataset to 26 Canadian Census Metropolitan Areas (CMAs), thereby excluding rural areas and smaller towns⁹. Canadian census data is obtained from von Bergmann et al. (2022), while DA census boundaries are converted to common geographies based on von Bergmann (2021), and using DA administrative boundaries from the 2016 Canadian census as the target geography. My final dataset is thus comprised of 25,479 DAs observed over 19 years, from 2000 to 2018, across 26 CMAs. The main outcome variable employed in the paper is yearly average PM_{2.5} concentration from Meng et al. (2019), which combine information from satellite-retrieved Aerosol Optical Depth with simulations and groundbased observations obtained from monitoring stations readings. I extract the mean value of yearly PM_{2.5}, weighted by grid-cell level population counts obtained from Rose et al. (2020), onto the 25,479 DAs which constitute my dataset for every year between 2000 and 201810.

The main advantage of this source compared to data obtained from monitoring stations is its much wider spatial and temporal coverage, which also allows me to overcome the selection problem mentioned in Grainger and Schreiber (2019) relative to the location of monitoring stations within urban areas¹¹. The entity of data loss when using ground-based data is considerable: PM_{2.5} data from the National Atmospheric Surveillance Program (NAPS) is only available for 61 DAs in 2000, growing to 230 in 2018 as new monitoring stations get added every year (see Figure A.1). Nonetheless, the satellite-retrieved measurements from Meng *et al.* (2019), when restricted to the DAs with at least one PM_{2.5}

¹¹ Monitoring stations are likely to be located where air pollution is lower due to strategic behaviour and discrimination by local regulators, thereby introducing measurement error in an eventual empirical analysis.



⁹ The CMAs in the dataset are: St. John's, Halifax, Saint John, Quebec, Trois Rivieres, Sherbrooke, Montreal, Ottawa, Saguenay, Kingston, Toronto, Hamilton, St. Catharine's, Kitchener, London, Windsor, Sudbury, Thunder Bay, Winnipeg, Regina, Saskatoon, Calgary, Edmonton, Abbotsford, Vancouver, and Victoria. While the number of Canadian CMAs is 35 in the latest available census wave (2016), I only keep in the dataset those CMAs which were designated as such in the 2001 Census, in order to ensure compatibility across all waves.

 $^{^{10}}$ The resolution of the PM_{2.5} raster data is 0.01° x 0.01° , while population data is available for grid cells of dimension 0.0083° x 0.0083° , implying that the population raster had to be resampled at the resolution of the PM_{2.5} raster in order to be viable for use in the weighted mean calculation.

ground monitoring station, correlate well with the NAPS readings, as shown in Figure A.2 and Figure A.3.

I rely on the Meng et al. (2019) PM_{2.5} estimates in order to produce my main results. However, I also run the main analysis using PM2.5 concentration data from van Donkelaar et al. (2019). While the two estimates are highly related, with a Pearson correlation coefficient of 0.795 (see Figure A.4 and Figure A.5), the concentrations from Meng et al. (2019) are generally lower throughout the sample¹². In terms of relevant covariates and environmental justice dimensions, I first obtain population counts at the DA level from Rose et al. (2020), which are available for all years between 2000-2018¹³. Further, I employ four waves of the Canadian census (2001-2016 at 5-year intervals) to retrieve information on median income at the DA level, and on the racial composition of the DA population, calculating the share of population belonging to a visible minority and the Theil's Entropy Index (Iceland, 2004) for racial diversity. I also extract the 2006 Material Deprivation Index from Pampalon et al. (2012) for all DAs in my sample. If the carbon tax was successful in producing a behavioural adjustment in BC residents, an expected result would be higher take up of alternative means of transport within metropolitan areas. Therefore, I leverage the detailed information contained in the four waves of Canadian census data to retrieve DA-level data on commute mode shares. I divide commute modes in two different categories: high emissions (cars, taxis, and motorcycles), and low emissions (public transport, bicycles, and walking)14.

Figure 1 plots the baseline spatial distribution of the dependent variable and the main covariates over the Vancouver CMA, the most populated metropolitan area in the treated province of British Columbia. Time-varying variables are averaged over 2005-2007, the three years preceding the implementation of the carbon tax, while all variables retrieved from the Canadian Census are taken at their 2006 values, the last observation before the tax was instituted. The distibution of PM_{2.5} concentrations is highly spatially correlated with population density, as found e.g. in the US by Carozzi and Roth (2023) or Germany by Borck and Schrauth (2021). Confirming the insights of the environmental justice literature (Cain *et al.*, 2024) racial diversity and the inverse of median income are also highly spatially correlated with air pollution at the baseline. Baseline commute mode seems to be inversely related with the spatial distribution of PM_{2.5}: areas whose inhabitants are less reliant on cars,

¹⁴ I further decompose the low emissions category into public transport only and zero emissions commutes (cycling and walking).



_

¹² The choice of employing data from Meng *et al.* (2019) is conservative, as results using the van Donkelaar *et al.* (2019) dataset are generally higher in magnitude.

¹³ The dataset also contains population counts for all DAs extrapolated from Canadian censuses; however, this data is only available in 5-years intervals between 2001 and 2016.

taxis and motorbikes seem to be more polluted on average, a result probably due to their centrality with respect to the road networks and urban form¹⁵.




FIGURE 1 • DESCRIPTIVE STATISTICS AT THE BASELINE

Notes: Spatial distribution of PM_{2.5} and relevant covariates within the Vancouver CMA. Top row: PM_{2.5} and population density; Middle row: median income and visible minority population share; Bottom row: high emission and public transport commute mode shares.

Lastly, I obtain monthly information on the BC gasoline and diesel fuel markets, at the province level, for January 1991-December 2016. In particular, I extract the annual sales of transportation fuels (motor gasoline and diesel), from Statistics Canada (2021b), gasoline and diesel price data from Kalibrate (formerly Kent Group Ltd.) at the monthly level for the city of Vancouver, which I consider representative of the entire province, monthly after tax income and unemployment rate data from Statistics Canada (2021c), and the CAD-USD monthly exchange rate, retrieved from the Pacific Exchange Rate Service at University of British Columbia's Sauder School of Business.

¹⁵ Summary statistics for the whole sample, split across treatment and control CMAs, are presented in Table A.1 and Table A.2 for the pre-treatment and post-treatment periods, respectively.



3. EMPIRICAL STRATEGY

3.1. Simple and matched difference-in-differences

The core aim of my empirical strategy is to estimate the treatment effect of the 2008 British Columbian carbon tax on local air pollution, measured in terms of PM_{2.5} concentrations at the DA level. A traditional methodology for this estimation is a two-way fixed effects difference-in-differences (TWFE-DID) regression. The estimating equation takes the form:

$$PM_{2.5it} = \tau^{did} D_{it} + \theta_t + \eta_i + \varepsilon_{it}$$
 (1)

Where D_{it} is the DID binary indicator, taking value 1 for all treated units after the implementation of the carbon tax in 2008, and 0 for all other observations; θ_t and η_i are respectively time and unit specific fixed effects, ε_{it} is a time-varying idiosyncratic error term, and τ^{did} is the coefficient of interest, capturing the average effect of being exposed to the carbon tax.

In order for τ^{did} to be equal to the average treatment effect on the treated cohort (ATT), the identifying assumption is that parallel outcome trends between the treated and the control units hold, i.e. if the 2008 carbon tax had not been implemented in British Columbia, PM_{2.5} levels in British Columbian DAs would have followed the same trajectory as PM_{2.5} levels in DAs located in other Canadian provinces. Figure A.6 and Figure A.7 report the average PM_{2.5} trends for 2000-2016 and 2000-2018, respectively, for British Columbian and control DAs, together with the universe of PM_{2.5} observations. The parallel trends assumption is untestable by definition, but it is essential to inspect the pre-treatment outcome paths and the distribution of treatment and control observations around their mean pre-intervention trends. In both cases, there is reason to suspect that the DID estimator would fail to identify the correct ATT. A visual inspection pre-treatment trends suggests a violation of the parallel trends condition (more evidently in the case of Figure A.7), while a more formal placebo DID regression of PM_{2.5} on treatment status with data limited to 2000-2007 and treatment assigned in 2004 identifies a significant placebo divergence in trends in both cases (Table A.3). Moreover, the dispersion of control observations around their mean trends is much higher than for treatment units, revealing substantial heterogeneity: by giving equal weight to all control observations, DID will include units whose pre and post-treatment outcome paths fundamentally differ from those of DAs in British Columbia, likely introducing an upward bias in the coefficient.

A potential solution to the pre-treatment heterogeneity in levels and trends is matching treatment and control groups on the basis of baseline pollution levels and on covariates



which influence air quality. Restricting the analysis to DAs which experience similar outcomes and are exposed to similar pollution stressors can attenuate the pre-treatment dispersion in PM_{2.5} levels and divergence in trends, and ensure the sample is more balanced before performing the DID regression (Imbens, 2015). I use one-to-one matching, moving away from the traditionally employed propensity score algorithm and instead preferring Coarsened Exact Matching (CEM) (Iacus *et al.*, 2012). I perform two versions of this procedure: in the first one (MDID1), I match treatment and control units on the baseline (2005-2007) average level of PM_{2.5}. In the second one (MDID2), I add baseline averages of population density, median income, high emission commute mode share, and road density. I exploit CEM to pre-process and trim the sample before running a weighted TWFE-DID regression using the CEM matching weights $\widehat{\omega^{cem}}$ in the following form:

$$PM_{2.5it}\sqrt{\omega_i} = \tau^{did} D_{it}\sqrt{\omega_i} + \theta_t\sqrt{\omega_i} + \eta_i\sqrt{\omega_i} + \varepsilon_{it}\sqrt{\omega_i}$$
 (2)

3.2. Synthetic control method and synthetic difference-in-differences

The problem of diverging pre-treatment trends in empirical applications is often addressed through the SCM (Abadie and Gardeazabal, 2003; Abadie, 2021)¹⁶. In the BC carbon tax case, the SCM constructs a set of synthetic DAs as a weighted combination of control DAs by finding, for each treated unit i, a non-negative vector of weights ω^{sc} summing to one, which ensures that each convex combination of the i outcome variable for control units matches each outcome variable for the treated units for all periods up to the intervention date.

In order to combine the attractive features of both TWFE-DID (the inclusion of additive unit-specific and time-specific fixed effects), and SCM (reducing the reliance on the parallel trends assumption by weighting observations in order to ensure closely matched pre-intervention trends), Arkhangelsky *et al.* (2021) have introduced a new method, synthetic difference-in-differences (SDID), which employs time and unit (two-way) fixed effects in the regression function (as in TWFE-DID), together with unit-specific weights (as in SCM) and time-specific weights which lessen the role of time periods that are largely divergent from post-treatment time periods. In a nutshell, for each treated unit SDID estimates: (1) unit weights ω_i^{sdid} which underpin a synthetic control whose outcome is approximately parallel to the outcome for the treated unit; (2) time weights λ_t^{sdid} which ensure that the average post-treatment outcome for control units only differs by a constant from the weighted average of pre-treatment outcome for each of the control units – a synthetic pre-

¹⁶ Usually with a unique treated unit, but extensible to the case of multiple treated units.



treatment period using controls. Once unit and time weights are calculated, SDID estimates a TWFE regression on the resulting panel, identifying the SDID ATT τ^{sdid} by solving the minimisation problem¹⁷:

$$\left(\hat{\tau}^{sdid}, \hat{\mu}, \hat{\eta}, \hat{\theta}\right) = \underset{\tau, \mu, \eta, \theta}{\operatorname{argmin}} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} (PM_{2.5it} - \mu - \eta_i - \theta_t - \tau D_{it})^2 \widehat{\omega_i}^{sdid} \, \hat{\lambda}_t^{sdid} \right\}$$
(3)

In the remainder of the paper, I regard SDID as my preferred method in order to estimate the effect of the 2008 BC carbon tax on air pollution co-benefits, as the methodology allows me to overcome the apparent violation of the parallel trends assumption and pre-treatment outcome heterogeneity problems in conventional DID; nonetheless, I estimate my main regression and robustness checks using DID, MDID, SCM and SDID, in order to assess the direction of the potential bias. I calculate standard errors for SCM and SDID using the bootstrap variance estimation algorithm described in Arkhangelsky *et al.* (2021, p. 4109), with 200 replications. The procedure constructs a bootstrap dataset by sampling a portion of the original dataset with replacement, and computes the estimator $\tau^{(b)}$ on this subset for each iteration *b*. The variance is then defined as:

$$\hat{V}_{\tau}^{b} = \frac{1}{b} \sum_{b=1}^{B} (\hat{\tau}^{(b)} - \frac{1}{B} \sum_{b=1}^{B} \hat{\tau}^{(b)})^{2}$$
(4)

4. **RESULTS**

In Table 1, I report the results of the DID, MDID, SCM and SDID regressions, using the Meng *et al.* (2019) PM_{2.5} dataset. The simple DID regression is also reported graphically, alongside the outcome path plots for SCM and SDID, in Figure 2. It is immediate to note how the violation of the parallel trends assumption examined in Table A.3 and Figure A.6 results in a likely case of upward bias for the simple DID, confirmed by the positive and significant coefficient, $\hat{\tau}^{did} = 0.39 \mu g/m^3$, reported in column (1) of Table 1 and in the leftmost panel of Figure 2. This result would indicate that the carbon tax has resulted in an increase of PM_{2.5} emissions, which would contradict the findings by Rivers and Schaufele (2015) and Pretis (2022) on fuel consumption and CO₂ emissions. MDID1 and MDID2 results are instead obtained by pre-processing the sample using the CEM procedure

¹⁷ Section A2 presents a detailed formal comparison between TWFE-DID, SCM, and SDID, drawing on the seminal work of Arkhangelsky *et al.* (2021).



described in Section III, matching on baseline PM_{2.5} levels in column (2) and on baseline PM_{2.5}, population density, median income, high emission commute mode share and road density in column (3). The matching procedure produces estimation samples which are much more closely aligned (Figure A.8 and A.9), and reverses the sign on the simple DID estimates, with negative and significant results contained in the $\hat{\tau}^{mdid}$ = [-0.27, -0.35] $\mu g/m^3$ range that are a first indication of the incidence of carbon pricing on air quality cobenefits. Moving away from simple DID estimation seems to be an effective strategy in minimising the impact of diverging trends and unbalanced pre-treatment characteristics.

Columns (4) and (5) confirm this insight by relying on the SCM and SDID methods, which also identify negative and statistically significant effects of the tax in reducing PM_{2.5} emissions. In the centre panel of Figure 2, I plot the average outcome path for the treated units and the traditional synthetic control. The improvement in pre-treatment fit is dramatic, with a minimal average deviation between British Columbian DAs and their controls, implying that the SCM performs well in giving positive weights to control units which best approximate treated DAs' outcome paths and zero weight to control units which exhibit different trends. Consistently with the hypothesised bias of the simple DID estimator, SCM indeed agrees with MDID in identifying an effect of opposite sign to DID, $\hat{\tau}^{sc} = -0.14 \,\mu\text{g/m}^3$. Results for the SDID estimator are graphically shown in the right-most panel of Figure 2. At the bottom of the panel, pre-treatment time-weights are represented in pink. Pre-treatment periods are weighted to match post-treatment levels (plus a constant) in the outcome variable for the control units. The SDID estimator does a particularly good job in imposing pre-treatment parallel trends in the years preceding the tax, even if weights λ_t are unevenly distributed over the pre-intervention period. However, negligible weights in 2007-2008 are reassuring, given that a standard caveat in event-study methodologies is the excessive reliance on the single period immediately preceding the intervention (Heckman and Smith, 1999). The SDID procedure is able to select control units which exhibit pre-treatment trends that are almost perfectly parallel to BC's outcome path, especially in the four-year window preceding the intervention. The estimated ATT is $\hat{\tau}^{sdid} = -0.36 \,\mu\text{g/m}^3$, corresponding to a 5.2% reduction with respect to pre-intervention mean pollution levels. I regard SDID as the preferred methodology due to its greater flexibility and to the selection of a sparser set of control DAs¹⁸. While SCM obtains a near-perfect fit pre-treatment, the outcome path of its synthetic unit

¹⁸ SDID selects indeed 6,258 control units among the untreated DAs and then performs DID on the matched sample with the inclusion of unit and time fixed effects to aid the estimation.



heavily depends on the particular set of units receiving positive weights, which in my highly disaggregated setting is not ideal¹⁹.

	(1)	(2)	(3)	(4)	(5)
	DID	MDID1	MDID2	SCM	SDID
$\hat{ au}$	0.3925	-0.2750	-0.3504	-0.1421	-0.3633
	(0.074)	(0.1495)	(0.1676)	(0.0809)	(0.0219)
Unit FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
ω_i		$\sqrt{\omega_i}$	$\sqrt{\omega_i}$	Yes	Yes
λ_t		•	•		Yes
N _{obs}	432939	305320	132430	432939	432939

TABLE 1 • THE 2008 CARBON TAX AND CHANGES IN PM_{2.5}

Notes: All point estimates represent the average impact of the 2008 carbon tax during the 2009-2016 post-treatment period. Standard errors in parentheses are calculated using the bootstrap variance estimation algorithm described in Arkhangelsky et al. (2021) with 200 replications for columns (1), (4) and (5), and are clustered at the CMA level for columns (2) and (3). In Column (2) the data is pre-processed by matching on coarsened bins of baseline PM_{2.5} levels. Column (3) additionally matches on population density, median income, high emissions commute mode share and road density at the DA level. All regressions use 2000-2016 data.

Notably, MDID, SCM and SDID all agree in identifying a negative and statistically significant effect of the 2008 carbon tax on PM_{2.5} emissions, contradicting "naive" DID estimates. The potential bias arising in the simple DID regression could be due to the diverging secular trends between treatment and control units, with treatment units on steeper declining trends prior to the implementation of the tax vis-à-vis control units. It is thus essential to address this concern in order to obtain a "clean identification" of the policy impact (Sager and Singer, 2024). Failing to do so would introduce a source of bias which could go as far as reversing the correct estimates. Finally, while I regard SDID as the preferred methodology over SCM due to its flexibility and its reliance on a larger portion of the control pool, it is crucial to note that the MDID in this instance obtains results which are similar in magnitude.

¹⁹ In Figure A.10, I aggregate all 6,258 DAs which receive positive weights to the CMA level, in order to obtain the composition of synthetic BC in terms of percentages of other Canadian CMAs, in a similar vein to the traditional SCM methodology of Abadie (2021).






FIGURE 2 • THE IMPACT OF THE 2008 CARBON TAX ON CHANGES IN PM_{2.5}

Notes: Graphical results from DID, SCM and SDID for PM_{2.5} concentrations, with Meng *et al.* (2019) data. The 2008 carbon tax is denoted by a black vertical line. Pre-treatment time weights λ_t are denoted in pink.

5. ROBUSTNESS CHECKS

5.1. Main results with van Donkelaar et al. (2019) PM_{2.5} data

I repeat the DID, SCM and SDID estimation using the van Donkelaar *et al.* (2019) PM_{2.5} dataset, which is available between 2000 and 2018. Notwithstanding the high correlation between the two outcome variables, as outlined in Figure A.4, both the treatment and control pre-intervention trends exhibit some differences with respect to the Meng *et al.* (2019) dataset²⁰. The violation of the parallel trends assumption is once again highlighted in a placebo DID regression (Table A.3), as well as in the graphical representation of the DID regression in Figure A.12 which, differently from the previous estimation, identifies a negative effect of the 2008 carbon tax on emissions of $\hat{\tau}^{did} = -0.5 \mu g/m^3$ (see Table A.4).

The SCM, represented graphically in the middle panel of Figure A.12, again obtains a good pre-treatment fit, signalling that each British Columbian DA's outcome path is best approximated by a convex combination of control DAs rather than equally weighted

 $^{^{20}}$ However, the temporal location of peaks and troughs is generally respected, as is the relationship between the BC and control units outcome path. Indeed, DAs located in British Columbia always exhibit lower average annual concentrations of particulate pollution, and their PM_{2.5} trend prior to 2008 appears to decline at an even faster pace than for control observations, barring some peaks in concentrations typical of the control provinces.



control units. Furthermore, as evidenced in Table A.4, the potential direction of the TWFE-DID bias is consistent with the main result: the SCM estimates a negative ATT of $\hat{\tau}^{sc}$ = $-0.71 \mu g/m^3$, therefore qualitatively reinforcing the SCM result of Table 1. A similar conclusion can be drawn from the results of the SDID estimation, presented in the rightmost panel of Figure A.12. The SDID procedure is able to select control units²¹ which exhibit pre-treatment trends that are almost perfectly parallel to BC's outcome path, with the exception of outlying time periods which receive zero-weights in the estimation. The estimated ATT of $\hat{\tau}^{sdid} = -0.89 \mu g/m^3$ is slightly lower, but qualitatively similar to the SCM ATT. In terms of magnitude, both the SCM and SDID regressions identify a substantial drop in PM_{2.5} concentrations with respect to 2000-2007 levels, corresponding to a reduction of 10.9% from the pre-intervention PM_{2.5} mean for British Columbia.

2.2. Accounting for measurement error in satellite-based estimates

The remotely sensed $PM_{2.5}$ datasets which I employ are gridded estimates of concentrations and may contain prediction error, which could substantially alter regression results (Fowlie et al., 2019). I assess the robustness of my estimates to this type of non-classical measurement error, by exploiting the geographic correspondence between gridded $PM_{2.5}$ data and DAs which contain at least one NAPS monitoring station, in order to construct a spatially matched dataset containing predicted and observed $PM_{2.5}$. I first calculate prediction error as the difference between satellite-derived $PM_{2.5}$ and monitor readings for the 1501 DA-year pairs which contain at least one NAPS monitoring station. As in Fowlie et al. (2019), I then regress prediction error $\Delta PM_{2.5}$ on a set of covariates at the DA level²² and I predict out of sample $\Delta PM_{2.5}$ for the entire dataset.

I adjust remotely sensed PM_{2.5} data by accounting for prediction error, and use the quantity $P\widehat{M}_{2.5it} = PM_{2.5it} + \Delta P\widehat{M}_{2.5it}$ to run SDID regressions using Meng *et al.* (2019) and van Donkelaar *et al.* (2019) data, respectively. The results, reported in Table A.5, are slightly lower though qualitatively similar to the main specifications, with $\hat{\tau}^{sdid} = -0.26\mu g/m^3$ using the Meng *et al.* (2019) dataset and $\hat{\tau}^{sdid} = -0.85\mu g/m^3$ using corrected van Donkelaar *et al.* (2019) data. The SDID estimator adequately identifies treatment effects by obtaining pre-treatment parallel trends in both instances (Figure A.13). The adherence between these results and the main specifications reinforces confidence about correctly measuring the policy effects. While the substantial difference between estimated treatment effects using the two gridded PM_{2.5} datasets remains, this is likely due to their calibration and prediction

 $^{^{22}}$ Namely, satellite-based PM_{2.5}, population density, nighttime lights, maximum and minimum temperature, and wind speed.



²¹ The composition of the donor pool, aggregated to the CMA level, is reported in Figure A.11.

procedures rather than prediction error. I conservatively adopt lower bound estimates using the Meng *et al.* (2019) as the main result, and regard all estimates using the van Donkelaar *et al.* (2019) product as the upper bound on my results.

5.3. Narrower temporal and spatial scope

Effect Dynamics: I restrict the estimation window to 2000-2013, in order to check whether the carbon tax ramp-up is the main mechanism behind the continuous reductions, and to what extent does the post-2013 tax rate freeze reverse the improvements²³. The results, presented in Figure A.14 and Table A.6 identify a higher ATT of $\hat{\tau}^{sdid} = -0.67$ µg/m³, which corroborates the hypothesis. The dynamics of the carbon tax phase-in are thus an important component of observed reductions: the effect is almost double in size in the first 5 years of the tax scheme, when tax rates increase step-wise every year. Air pollution improvements slightly reverse and stabilise at a lower level once the tax signal is kept constant.

Main CMA: I confine the treated pool to DAs within the Vancouver metropolitan area, excluding all DAs in the Abbotsford and Victoria CMAs. The resulting treatment cohort is comprised of 2874 DAs, vis-à-vis the 3490 DAs constituting the entire treatment unit pool; the control pool is kept the same, with 21989 control DAs. Perhaps unsurprisingly, given the relatively small number of DAs pertaining to the Abbotsford and Victoria CMAs, the results (reported in Figure A.15 and Table A.7) are qualitatively unchanged from the main regressions using the Meng *et al.* (2019) dataset.

NAPS Locations: I select DAs corresponding to the location of NAPS monitoring stations (see Figure A.1)²⁴. I thus consider just those locations in which pollution monitors have been established, thereby restricting the analysis to areas in which pollution is likely to be a greater concern. Here, the size of the dataset is considerably restricted: the cross-section of DAs kept in the treated pool counts just 25 observations, while 106 DAs are kept in the control pool. Once again, the results (presented in Figure A.16 and Table A.8), are qualitatively similar to the main specifications. Notably, the performance of the SDID estimator is not considerably worsened on this much smaller sample, achieving a reasonable pre-treatment fit, and therefore identifying a credible ATT.

²⁴ I match DAs with all monitoring stations in the dataset, regardless of the date of establishment of each monitoring station, in order to maximise observations.



²³ In 2012, the carbon tax was frozen at \$30/tCO₂ as reported in Section 2.

6. MECHANISMS

6.1. Reductions in transport fuel demand

The first candidate explanation for the observed reductions in particulate matter concentrations is a change in consumer behaviour regarding transportation fuel. Some evidence supporting this explanation is found in early analyses of the BC carbon tax (e.g. Rivers and Schaufele, 2015; Lawley and Thivierge, 2018), which use a limited post-intervention time period and only focus on gasoline consumption²⁵. On the contrary, fuel substitution away from gasoline and towards diesel is claimed to be a potential mechanism behind the PM_{2.5} increases found in Saberian (2017), notwithstanding the negative impacts found by the time series analysis of Bernard and Kichian (2019) and the strong prevalence of gasoline vehicles among BC car sales (see Figure A.19 and A.20).

I reconcile the evidence on the aggregate level effects of the carbon tax on transportation fuel demand by introducing a recently developed method for high-frequency time series analysis: the Causal-ARIMA (C-ARIMA) estimator of Menchetti *et al.* (2022). By exploiting features of ARIMA models, the method is especially appropriate to analyse complex seasonal, nonstationary processes such as gasoline and diesel sales observed monthly (see Figure A.17, panels A and B). C-ARIMA combines attractive features from the DID and SCM estimator for the case in which no suitable control unit is available²⁶ and when the number of pre-intervention time periods is large²⁷. Under standard assumptions²⁸, C-ARIMA is able to learn the treated unit's time series dynamics and forecast it after the shock takes place. By using the forecasted series as the treated unit's counterfactual outcome, the method identifies two main sets of causal effects: the temporal average causal effect and the cumulative treatment effect.

I run C-ARIMA separately for per capita monthly gasoline and diesel sales at the aggregate BC level between January 1991 and December 2016. The intervention date is July 2008, i.e. the specific month in which the BC carbon tax came into effect. In Table 2, I report the

²⁸ No temporal interference (i.e. absence of anticipation effects), covariates-treatment independence and conditional stationarity (Menchetti *et al.*, 2022).



_

²⁵ Which accounts for most of the residential vehicle fleet (see Figure A.19) but does not include heavy duty vehicles used in commercial and industrial operations (Bernard and Kichian, 2019).

²⁶ In my context, a pool of eligible control units is represented by other Canadian provinces. However, other provinces exhibit diverging pre-intervention trends in gasoline sales (see Figure A.18) when aggregating the TWFE-DID coefficients into an event study plot.

²⁷ As is the case in the monthly analysis of BC fuel consumption between January 1991 and December 2016, with 210 pre-intervention time periods.

results from estimations with and without a matrix of business cycle controls²⁹. Both the temporal average causal effect $\widehat{\tau}_t$ and the cumulative causal effect $\sum_{t=t_{int}}^T \widehat{\tau}_t$ are negative and statistically significant across all specifications, highlighting a successful impact of the BC carbon tax in decreasing fuel demand, consistently with Rivers and Schaufele (2015) and Bernard and Kichian (2019). In Figure A.17, the results from the estimation are reported graphically.

	Gasolin	e Sales	Diesel	Sales
$\hat{ au}_t$	(1) -3.883 (0.553)	(2) -4.675 (0.506)	(3) -1.756 (0.412)	(4) -0.912 (0.236)
$\sum\nolimits_{t=t_{int}}^{T}\widehat{\tau_{t}}$	-396.052 (56.453)	-818.405 (14.962)	-179.089 (42.066)	-92.983 (24.093)
Controls	-	Yes	-	Yes
N_{obs}	312	312	312	312

TABLE 2 • C-ARIMA: MONTHLY GASOLINE AND DIESEL DEMAND

Notes: The dependent variable is total monthly gasoline (diesel) sales per capita (in litres) recorded in British Columbia benween January 1991 and December 2016. Columns (2) and (4) include a matrix of monthly province-level covariates, namely consumer price index, gasoline (diesel) crude cost, population, unemployment rate, after tax income and the US-CAD exchange rate. Standard errors in parentheses are computed through 1000 bootstrap runs.

6.2. Commute mode switching

I analyse commute mode choices at the DA level as an additional mechanism driving the main results. While commute mode is an imperfect measure of the number and type of trips made by British Columbians, I can rely on the same administrative level to the one used in the main analysis by retrieving information from the 2001, 2006, 2011, and 2016 Canadian censuses, thereby preserving granularity. In Table 3 and Table 4, I report TWFE-

²⁹ Namely, provincial population, unemployment, after tax income, exchange rate and the cost of crude gasoline and diesel, respectively.



DID regression results³⁰ employing the share of commuters using high-emissions and public transport commute modes, respectively³¹.

In all tables, column (1) is the baseline specification, a simple DID regression with DA and year fixed effects and no controls, employing the full panel of DAs across census years. In column (4), I add weather controls for precipitation, maximum and minimum temperature, and wind speed, plus the natural logarithm of population and median income. When employing the full pool of control DAs, the first result of note is that British Columbian DAs experience an average 4.2% reduction in the use of cars, taxis, and motorcycles, which rises to 4.7% when adding controls. This reduction is almost specular to the increase in the share of commuters using public transport, biking and walking to reach their workplace (Table A.9). Moreover, as evidenced in Table 4 and Table A.10, most of this increase (3.5-3.9%) is due to a higher reliance on public transport, while a residual share of 0.5-0.7% is due to a switch to active commuting.

All results are confirmed and stronger in magnitude when considering more restrictive specifications: columns (2) and (5) restrict the specifications in (1) and (4) to the DAs which receive positive weights in the main SDID regressions, in order to establish whether the mechanisms are effectively retrieved when employing the same set of observations on which the main ATT is estimated. Results are higher in magnitude by about 1%, jumping to a 5.3% reduction in high-emission commute modes in the case without controls. Here, the inclusion of control variables slightly dampens the impact to 5.2%; nonetheless, the specularity with the increase in low-emission commute modes is preserved. Finally, in columns (3) and (6) I augment the DID regressions by retrieving an including the weights from the main SDID regressions. I weigh all treatment observations equally and all control observations according to the value of ω_i they receive after the data-driven SDID procedure. The magnitude of the increase in public transport commute share increases further, to 4.2% in the case without covariates and is again dampened to 4.1% in the case with covariates. The hypothesis of a behavioural adjustment by BC citizens in response to the carbon tax is thus confirmed; residents of BC's DAs switch away from high-emissions

³¹ As the low-emissions transport mode is the sum of public transport and zero-emissions modes, I only report the results for public transport in the main text and present the aggregate low emissions and the subsplit for zero-emissions in Table A.9, and Table A.10.



-

³⁰ Due to the structure of the data, collected at 5-year intervals, I am prevented from using the SCM and SDID methodology in this exercise; I thus resort to traditional TWFE-DID estimation of commute mode switching, analysing the data separately for each category of commute mode. Details on this estimation strategy are reported in Section B.

commute modes towards low-emissions ones, with public transport as the main container for these substitutions.

TABLE 3 • DID RESULTS FOR HIGH EMISSIONS COMMUTE MODE

		Hi	gh Emissions (Commute Mod	le	
	(1)	(2)	(3)	(4)	(5)	(6)
DID	-0.0417	-0.0527	-0.0549	-0.0466	-0.0519	-0.0516
	(0.0105)	(0.0095)	(0.0103)	(0.0102)	(0.0106)	(0.0109)
DA FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls				Yes	Yes	Yes
SDID control pool		Yes	Yes		Yes	Yes
SDID weights			Yes			Yes
R ²	0.87184	0.83989	0.84360	0.87595	0.84508	0.84847
Adj. R ²	0.82896	0.78629	0.79124	0.83400	0.79267	0.79721
N_{obs}	101358	38769	38769	100244	38348	38348

Notes: The dependent variable is the dissemination area level share of high emissions commutes. All regressions include dissemination area and year fixed effects. Columns (4)-(6) include controls for precipitation, maximum and minimum temperature, and wind speed, plus the natural logarithm of population and median income. Columns (2), (3), (5) and (6) restrict the control unit pool to DAs which receive positive weights in the main SDID regression. Columns (3) and (6) additionally include the estimated SDID unit weights ω as regression weights. Standard errors in parentheses are clustered at the CMA level.

TABLE 4 • DID RESULTS FOR PUBLIC TRANSPORT

		Public Transport Commute Mode				
	(1)	(2)	(3)	(4)	(5)	(6)
DID	0.0352 (0.0107)	0.0410 (0.0107)	0.0417 (0.0112)	0.0391 (0.0115)	0.0422 (0.0115)	0.0414 (0.0111)
DA FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls				Yes	Yes	Yes
SDID control pool		Yes	Yes		Yes	Yes
SDID weights			Yes			Yes
\mathbb{R}^2	0.83768	0.78668	0.78011	0.84196	0.79197	0.78571
Adj. R ²	0.78336	0.71526	0.70650	0.78851	0.72160	0.71322
N_{obs}	101358	38769	38769	100244	38348	38348

Notes: The dependent variable is the dissemination area level share of public transport commutes. All regressions include dissemination area and year fixed effects. Columns (4)-(6) include controls for precipitation, maximum and minimum temperature, and wind speed, plus the natural logarithm of population and median income. Columns (2), (3), (5) and (6) restrict the control unit pool to DAs which receive positive weights in the main SDID regression. Columns (3) and (6) additionally include the estimated SDID unit weights ω as regression weights. Standard errors in parentheses are clustered at the CMA level.



7. ENVIRONMENTAL JUSTICE DYNAMICS

In light of a growing literature in environmental justice (see Banzhaf *et al.*, 2019; Cain *et al.*, 2024), I examine efficiency-equity trade-offs in the realisation of co-benefits, inspecting whether the estimated air pollution reductions arise heterogeneously over metropolitan areas. In the main analysis, the parameter identifying the effect of the 2008 BC carbon tax on PM_{2.5} emissions has always been assumed as constant across treated units. Nonetheless, when dealing with disaggregated data within Census Metropolitan Areas, a homogeneously estimated ATT is likely to mask substantial heterogeneities across DAs which could be highly informative about the performance of different locations within metropolitan areas.

A first channel to explore is certainly that of pre-existing pollution levels: standard economic theory would in fact predict that emission abatement would happen first where the marginal cost of reducing emissions is lower, i.e. where pre-existing pollution is higher (that is, lower-hanging fruits would be picked earlier). This avenue is explored by Auffhammer *et al.* (2009) and Sager and Singer (2024), who find substantially higher reductions in PM_{2.5} and PM₁₀ due to the Clean Air Act in non-attainment US census tracts that are more polluted in the three years preceding the implementation of the policy. Nonetheless, the opposite result may also arise if the rate of vehicle replacement is higher in less polluted areas or if more polluted areas substitute more strongly towards less CO₂-intensive, but more PM_{2.5}-intensive vehicles such as diesel automobiles or, crucially, diesel-powered public transport³².

In light of the results of Borck and Schrauth (2021) and Carozzi and Roth (2023), it is also worth exploring whether heterogeneity in air pollution reductions arises at different levels of the population density distribution: indeed, while densely populated areas have been shown to experience higher concentrations of PM_{2.5} particulate, usually population density is higher in city centres, where greater opportunities for substitution away from cars may arise. Again, the nature of the eventual substitution plays a crucial role in determining the direction of the realised effect³³.

The most investigated avenue in studies focussing on air quality and environmental justice dynamics is certainly that of racial disparities (e.g. Colmer *et al.*, 2020; Jbaily *et al.*, 2022; Currie *et al.*, 2023; Sheriff, 2024; Sager and Singer, 2024). This is possibly due to a relatively higher weight of US-centric studies in the literature, with much fewer evidence on extra-

 $^{^{33}}$ To the point of determining detrimental impacts in case the substitute transport mode is more CO_2 efficient but emits greater $PM_{2.5}$ concentrations than the original one – such as the case of diesel-powered engines substituting for gasoline ones.



³² That is, if the implicit assumption of homogeneous abatement technology across locations is removed.

US contexts. Nonetheless, as shown in Figure 1, there is a high spatial correlation between the location of DAs with higher shares of visible minority population and the pre-intervention PM_{2.5} distribution. It is thus extremely important to (1) assess whether there is an ex ante racial disparity in PM_{2.5} exposure in British Columbian metropolitan areas, (2) examine whether ex post air quality racial EJ gaps fall or widen as a result of a market-based intervention such as a carbon tax.

Lastly, a large body of research has shown that carbon pricing is regressive along income and wealth dimensions (Poterba, 1991; Grainger and Kolstad, 2010; Sager, 2023), but its incidence depends on the destination of tax revenues (Metcalf, 2009; Davis and Knittel, 2019). While this insight is well understood, constituting for instance the basis for the BC low income climate action tax credit, there is much less certainty about the role of realised air quality co-benefits in mitigating or exacerbating income regressiveness. Indeed, if a disproportionately large share of co-benefits accrue to population in higher income brackets, there is an additional dimension of carbon tax inequality that is not factored in projected governmental revenue rebates.

In Figure 3, the long term changes in PM_{2.5} concentrations are mapped out for the three BC CMAs in my sample, to individuate areas which have withstood more pronounced improvements. Confirming the traditional insights of the EJ literature (Colmer *et al.*, 2020; Jbaily *et al.*, 2022) there appears to be a secular convergence in air pollution concentrations: dense, central areas with higher initial PM_{2.5} concentrations and higher levels of socioeconomic deprivation see greater absolute improvements in air quality.

These insights are only descriptive, as they do not take into account the potential for similar secular trends in control areas, nor the impact of the BC carbon tax on air pollution concentrations within a potential outcomes framework. Table A.11 and Table A.12 report ex ante and ex post gaps in observed PM_{2.5} in British Columbia and control CMAs along the top and bottom quintiles of three covariates describing environmental justice dimensions: (1) baseline population density, (2) baseline share of population belonging to a visible minority (as a measure of racial disparity), (3) baseline median income³⁴. It is straightforward to assess how, before the implementation of the carbon tax, EJ gaps manifest along every dimension in both subsamples. Denser, more racially diverse, and poorer areas of BC CMAs are comparatively more polluted at the start of the sample, by 1.9-2.2 3 µg/m³, or 27-32% of the pre-treatment PM_{2.5} average level in treated units. Similar EJ gaps are observed in control DAs, with generally higher pollution levels across quintiles.

³⁴ In Figure A.22, I also use Material Deprivation Index and the Theil Entropy Index for racial diversity.



For all DAs in treated and control CMAs, EJ gaps reduce substantially in the last three years of the sample, to 1-1.3 μ g/m³ in BC and 0.5-1.2 3 μ g/m³in control DAs.

Reading the statistics reported in Table A.11 and Table A.12 in conjunction, the pollution convergence hypothesis is not only confirmed for BC, but emerges as a common trend, with EJ gaps shrinking along every dimension across all Canadian DAs in the sample. It is thus important to examine whether the 2008 BC carbon tax has had heterogeneous contributions along these axes in order to determine winners and losers from climate mitigation policy. While the long term air quality improvements illustrated in Figure 3 accrue more strongly to relatively more deprived areas, this consideration is not causally determined in the absence of an appropriate counterfactual for each stratus of the EJ dimension under consideration. To estimate the impact of the carbon tax on PM_{2.5} emissions at different points of the distribution of EJ characteristics, I split the treatment sample into quintiles of baseline³⁵ PM_{2.5}, population density, median income and racial diversity. As the SDID methodology does not allow the inclusion of interaction terms in the estimation procedure, I then run SDID separately for each quintile, allowing the datadriven algorithm to select the combination of control DAs which best approximates the outcome path of each quintile split of the treated units (plus a constant). In Figure 4, I summarise the results graphically, reporting point estimates and 95% confidence intervals at each quintile of the baseline EJ characteristics' distributions.

Quintile-SDID results for baseline PM_{2.5} concentrations are presented in panel (A) of Figure 4. It is immediate to infer that greater reductions arise in DAs with lower pollution levels between 2005 and 2007. The bottom quintile of baseline pollution indeed experience 2.2 times larger reductions with respect to the top quintile. Panel (B), which shows the SDID effects for quintiles of baseline population density, is consistent with the results for baseline pollution levels.

³⁵ For time-varying covariates I use the average of the three years prior to treatment as the baseline value; for variables retrieved from the Canadian census, I use their 2006 values, i.e. the last observation prior to the implementation of the carbon tax.



 $\Delta PM_{2.5} (\mu g/m^3)$ $\Delta \text{ PM}_{2.5} (\mu g/\text{m}^3)$ $\Delta PM_{2.5} (\mu g/m^3)$

FIGURE 3 • SPATIAL DISTRIBUTION OF THE LONG-TERM CHANGE IN PM2.5 IN BC CMA

Notes: This figure plots the geographical distribution of changes in PM_{2.5} concentrations for the Vancouver (top panel), Victoria (middle panel) and Abbotsford (bottom panel) CMAs from 2000-2002 to 2014-2016 (three-year averages).

Denser locations within metropolitan areas see lower reductions of particulate matter with respect to less dense DAs, underpinning a worsening of the pollution-density gap. Taken in conjunction, these insights appear to confirm that the 2008 carbon tax was not effective in curtailing traffic in more central areas within British Columbian metropolitan areas, but rather had greater effect in peri-urban locations. More surprising are the results in panel (C), which highlight the fact that relatively better off DAs within metropolitan areas have experienced



greater reductions, possibly reflecting an inverse relationship between density and income, but more importantly signalling that the pollution-income gap has increased as a result of the carbon tax. This result is a clear confirmation of the hypothesis of co-benefits regressiveness, highlighting how there is an additional distributional dimension that needs to be considered when designing climate policy. Finally, panel (D) illustrates how metropolitan areas with a lower proportion of visible minorities, i.e. less racially diverse districts, experience greater gains in terms of realised air quality co-benefits, a result that is in agreement with the findings for baseline pollution and income, which exhibit high correlation with racial diversity. The positive findings of a burgeoning EJ literature focussed on the black-white exposure gap in the US and on command-and-control regulations (e.g. Currie *et al.*, 2023; Sager and Singer, 2024) are then not confirmed in the case of British Columbia and a market-based instrument.

It is important to notice how these results highlight pollution reductions that arise at every quintile of the baseline EJ characteristics, compared to a synthetic DID counterfactual of no carbon tax implementation. The carbon tax policy thus produces Pareto-optimal air quality co-benefits, since every group experiences improvements in pollution exposure after its implementation. Compared to a counterfactual scenario of no policy³⁶, the carbon tax is thus welfare-improving across the board, a result which could be considered beneficial also in terms of environmental justice, with least-advantaged groups observing important reductions in pollution exposure. However, the unequal distribution of realised PM_{2.5} reductions widens the EJ gap along every considered dimension. This confirms mixed evidence from the EJ literature, which claims that while market-based instruments for climate mitigation can give rise to inequality-improving air quality co-benefits (Hernandez-Cortes and Meng, 2023), they can also worsen pre-existing disparities (Grainger and Ruangmas, 2018; Shapiro and Walker, 2021; Cain et al., 2024) or result in no significant distributional changes (Fowlie et al., 2012). Economic instruments which specifically target air pollution, such as command-and-control regulation³⁷ have been shown to produce sustained EJ gains (Currie et al., 2023), and may thus be coupled with market-based climate mitigation policies³⁸ in order to reap the full set of benefits from regulation and reduce environmental inequality between groups, in addition to decreasing pollution exposure across groups.

³⁶ As opposed e.g. to the command-and-control counterfactual of Fowlie et al. (2012).

³⁸ Additional instruments can be aimed at internalising the congestion externality in urban centres and reducing local pollution (e.g. Gehrsitz, 2017; Pestel and Wozny, 2021; Sarmiento *et al.*, 2023), with a specific focus on policy impacts on disadvantaged communities. Further co-benefits can be generated by public transport electrification and incentives for alternative transport modes, as low-income and disadvantaged households are relatively more cash and credit constrained.



³⁷ A prime example of which are the National Ambient Air Quality Standards (NAAQS) enforced by the US Clean Air Act.




FIGURE 4 • QUINTILE-SDID RESULTS FOR ENVIRONMENTAL JUSTICE GAPS

Notes: Results of SDID regressions by quintile of baseline characteristics. Panel A) Quintiles of baseline PM_{2.5}; B) Quintiles of baseline population density; C) Quintiles of baseline median income; D) Quintiles of visible minority share. ATT point estimates reported in red, with 95% confidence intervals calculated with the Arkhangelsky *et al.* (2021) procedure with 200 bootstrap runs in grey shading.

8. HEALTH GAINS

In order to understand the magnitude of the economic co-benefits from air pollution reductions arising due to the 2008 carbon tax, I convert the quintile SDID PM_{2.5} estimates into a monetary quantification of the associated health gains. Notwithstanding the relatively low concentrations of particle pollution in the British Columbian context, where pretreatment air quality was of substantial better quality than in other North American locations (e.g. in the USA), it is important to note that the concept of "safe" thresholds for particle pollution concentrations is more normative than positive. Indeed, some studies (e.g. Krewski *et al.*, 2009) have highlighted that the marginal benefits from abatement may be nonlinear in baseline concentrations, with lower gains from abatement at higher levels of baseline air pollution. Hence, any improvement in air quality is likely to carry significant benefits in terms of reductions in mortality rates; moreover, the estimates reported in this section are a lower bound of the gains from local pollution reductions, as PM_{2.5} has been



shown to have a multidimensional impact, ranging from health to productivity, to cognition and the formation of human capital (Aguilar-Gomez et al., 2022).

Drawing from Fowlie *et al.* (2019) and Carozzi and Roth (2023), my approach consists of two steps. I first estimate the impact of a reduction in PM_{2.5} concentrations in terms of mortality reductions, using concentration-response ("hazard") functions derived from the environmental health literature. Second, I retrieve the central estimate of the willingness to pay (WTP) to avoid a premature death from Health Canada (2021) and Chestnut and De Civita (2009)³⁹, and multiply the mortality reductions estimated in the first step by the central estimate of the Value of a Statistical Life (VSL), equal to \$6.5 million in 2007 Canadian dollars, for each DA in the census metropolitan areas of Vancouver, Victoria, and Abbotsford.

The traditional form of the Cox proportional hazard model used in the environmental health literature is the log-linear regression reported in Fowlie *et al.* (2019):

$$\ln(\gamma) = \zeta + \alpha P M_{2.5} \tag{5}$$

Where $ln(\gamma)$ is the natural logarithm of mortality risk, $\zeta = ln(Z)$, and PM_{2.5} are the

local pollution concentrations. The term Z is a vector of covariates other than $PM_{2.5}$ which impact mortality, and can be rewritten as $Z = Z_0 + \exp(\beta_1 x_1 + \dots + \beta_n x_n)$, with Z_0 being the mortality risk when all covariates are zero. Indicating γ_{0i} as the baseline mortality risk, and rearranging terms⁴⁰, the change in mortality rate $\Delta \gamma_i$ can be related to the change in pollution levels $\Delta PM_{2.5i}$ with the following equation:

$$\Delta \gamma_i = \gamma_{0i} \left(1 - \frac{1}{e^{\alpha \Delta P M_{2.5i}}} \right) \tag{6}$$

In order to find the total number of deaths for each DA associated with the above change in mortality rate $\Delta \gamma_i$, this quantity needs to be multiplied by the population of each DA⁴¹:

$$\Delta Deaths_i = Population_i \left[\gamma_{0i} \left(1 - \frac{1}{e^{\alpha \Delta P M_{2.5i}}} \right) \right]$$
 (7)

⁴¹ I use the baseline population level, that is, the population of each DA in the year 2008.



_

³⁹ It must be noted that the reported estimate for the Value of a Statistical Life does not reflect directly the economic value of an individually identified person's life, but rather the aggregation of estimates of the WTP for a small reduction in mortality risk. Using the VSL central estimate of \$6,500,000, for example, the average Canadian would be willing to pay \$65 to reduce the risk of premature death by 1 out of 100,000.

⁴⁰ The derivation is as follows (Carozzi and Roth, 2023): $\Delta \gamma = Z \left(e^{\alpha P M_{2.5}^0} - e^{\alpha P M_{2.5}^1} \right) \rightarrow \Delta \gamma = Z e^{\alpha P M_{2.5}^0} [1 - e^{-\alpha (P M_{2.5}^0 - P M_{2.5}^1)}]$

And finally, the monetary health gains in terms of mortality reductions at the DA level, Δ Yi, are obtained by multiplying the above estimates by the VSL figure of \$6.5 million CAD obtained from Health Canada (2021):

$$\Delta Y_i = VSL * \Delta Deaths_i \tag{8}$$

Hence, in order to estimate the model outlined in Equation 6, and thus obtain mortality rate changes at the DA level, I first need to estimate the baseline mortality rate γ_{0i} . Consistently with the literature, I obtain data for deaths due to lung cancers, all circulatory diseases, and all respiratory diseases from the ICD.10 selected causes of death at the CMA level from Statistics Canada $(2021a)^{42}$. I divide total deaths due to the listed causes by total CMA population, and assign the resulting (baseline) mortality rates to all DAs in a given CMA. The parameter α is usually not directly indicated in epidemiology studies, which instead report the relative risk (RR) increase due to a given increase in PM_{2.5}. For instance, Lepeule *et al.* (2012) report an all-cause RR of 1.14 associated with a $\Delta PM_{2.5}$ of 10 µg/m³, while Krewski *et al.* (2009)'s estimate is 1.06. However, it is straightforward to retrieve α by exploiting the relationship between RR and $\Delta PM_{2.5}$: $\alpha = \ln(RR) / \Delta PM_{2.5}$.

I employ these two estimates, in combination with the estimated PM_{2.5} reductions for each quintile of the pre-intervention PM_{2.5} distribution, in order to calculate the gains from mortality reductions at the DA level for the three CMAs included in the treated sample: Vancouver, Victoria and Abbotsford. In Figure 5, I visually report the results of this exercise for each CMA, using RR = 1.14 as estimated by Lepeule *et al.* (2012)⁴³.

The left panel maps the estimated mortality reductions per 1000 people (estimated according to Equation 7), while the right panel shows the associated per capita health gains, estimated via Equation 8. The median per capita monetary gains due to the estimated reductions in PM_{2.5} are large: \$198 when using the Lepeule *et al.* (2012) RR and \$88 with the RR from Krewski *et al.* (2009)⁴⁴.

The monetary value of per capita air quality co-benefits from the BC carbon tax is 1.7 times of the per capita low income climate action tax credit, i.e. the carbon tax rebate for low-income families⁴⁵. Moreover, the total monetary value of co-benefits ranges between \$507.2 million and \$1.03 billion annually, or 40-81% of annual carbon tax revenues once

⁴⁵ For this comparison, I use the last revision of the low income climate action tax credit, amounting to \$115.50 per adult plus \$34.50 per child (Ministry of Finance, 2013).



⁴² This is the smallest geographic unit for which the data are available.

⁴³ Visual results using the RR estimate from Krewski et al. (2009) are reported in Figure A.23.

⁴⁴ The same gains are \$402 and \$178, respectively, if calculated using the ATT estimated with the van Donkelaar *et al.* (2019) PM_{2.5} dataset instead of Meng *et al.* (2019).

the tax reached its \$30/tCO₂ level in 2012 (Ministry of Finance, 2013). The spatial distribution of these gains shows substantial heterogeneity: in particular, it is once again striking how air pollution co-benefits seem to be concentrated in peri-urban areas and positively correlated with income (see also Figure A.24). The results confirms that co-benefits from carbon taxation appear to be regressively distributed in metropolitan areas, with greater air quality improvements arising in higher income, low pollution DAs, underpinning increasing environmental justice gaps, as also evidenced in Section VII.




FIGURE 5 • MORTALITY REDUCTION AND MONETARY HEALTH GAINS

Notes: Spatial distribution of mortality reductions per 1000 residents (left panel) and health gains per capita (right panel) using the RR estimates from Lepeule *et al.* (2012), for the Vancouver (top row), Victoria (middle row) and Abbotsford (bottom row) CMAs.



9. CONCLUSIONS

This paper connects two areas of extreme importance in environmental economics. Air pollution co-benefits from carbon taxation are likely to be as large in magnitude as to partially or fully offset the costs of climate mitigation. Incorporating the monetary value of air quality improvements in cost-benefit analyses of carbon taxes is essential in order to correctly calibrate them and enhance their attractiveness. Conversely, environmental justice implications of market-based instruments are an often overlooked dimension due to the focus on efficiency, rather than equity. Ignoring potentially regressive consequences in terms of the societal distribution of co-benefits could hinder public support towards climate policy.

I show that the introduction of carbon pricing can significantly improve local air quality. After the implementation of the 2008 carbon tax, PM_{2.5} levels dropped by 5.2-10.9% in British Columbian dissemination areas, compared to a no policy counterfactual obtained through the synthetic difference-in-differences estimator. The air quality improvement is driven by reductions in fuel demand and by transport mode switching, mostly in favour of public transport. In terms of environmental justice, alongside evidence of Pareto optimal improvements for all segments of the population, pollution reduction dynamics are significantly regressive, with greater effects found in less polluted, less dense, less racially diverse areas and in richer neighbourhoods. Finally, I convert the improvements in air quality into reductions in mortality rates and monetary health gains from co-benefits of carbon taxation. With a median estimate of \$198 per capita, the health gains are large and comparable to the rebates offered to low-income families in British Columbia to mitigate the impact of the tax on their disposable income, as well as to the total annual revenue from the carbon tax at maturity.

These results highlight an equity dimension of the regressive nature of carbon pricing, showing how environmental justice improvements are not a necessary consequence of market-based instruments. While addressing complementary global and local environmental externalities via a carbon tax can yield significant air pollution and health cobenefits in addition to climate mitigation, regressive outcomes ought to be considered. Instruments designed to attenuate inequitable effects may then be designed in advance of the deployment of carbon pricing in order to help closing environmental justice gaps and reap greater policy gains.



REFERENCES

- 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, pp. 493-505
- Abadie, A. and Gardeazabal, J. (2003), *The Economic Costs of Conflict: A Case Study of the Basque Country*, «American economic review», 93, 1, pp. 113-132
- Abatzoglou, J., Dobrowski, S., Parks, S., and Hegewisch, K. (2018), TerraClimate, a High-resolution Global Dataset of Monthly Climate and Climatic Water Balance from 1958-2015, «Nature Scientific Data», 5, p. 170191
- 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, pp. 403-425
- Ahmadi, Y, Yamazaki, A., and Kabore, P. (2022), How Do Carbon Taxes Affect Emissions? Plant-Level Evidence from Manufacturing, «Environmental and Resource Economics», 82, 2, pp. 285-325
- Andersson, J.J. (2019), Carbon Taxes and CO₂ Emissions: Sweden as a Case Study, «American Economic Journal: Economic Policy», 11, 4, pp. 1-30
- Arkhangelsky, D., Athey, S., Hirshberg, D.A., Imbens, G.W., and Wager, S. (2021), *Synthetic Difference-in-Differences*, «American Economic Review», 111, 12, pp. 4088-4118
- Auffhammer, M., Bento, A., and Lowe, S. (2009), Measuring The Effects Of The Clean Air Act Amendments on Ambient PM10 Concentrations: The Critical Importance of a Spatially Disaggregated Analysis, «Journal of Environmental Economics and Management», 58, pp. 15-26
- Azevedo, D., Wolff, H., and Yamazaki, A. (2023), *Do Carbon Taxes Kill Jobs? Firm-Level Evidence From British Columbia*, «Climate Change Economics», 14, 02, pp. 1-46
- Banzhaf, S., Ma, L., and Timmins, C. (2019), *Environmental Justice: The Economics of Race, Place, and Pollution*, «Journal of Economic Perspectives», 33, 1, pp. 185-208
- 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 10508
- Bernard, J.-T. and Kichian, M. (2019), The Long And Short Run Effects Of British Columbia's Carbon Tax On Diesel Demand, «Energy Policy», 131, pp. 380-389
- Bondy, M., Roth, S., and Sager, L. (2020), *Crime Is in the Air: The Contemporaneous Relationship between Air Pollution and Crime*, «Journal of the Association of Environmental and Resource Economists», 7, 3, pp. 555-585
- Borck, R. and Schrauth, P. (2021), *Population Density and Urban Air Quality*, «Regional Science and Urban Economics», 86, p. 103596
- Boyce, J. and Pastor, M. (2013), Clearing The Air: Incorporating Air Quality and Environmental Justice into Climate Policy, «Climatic Change», 120, 4, pp. 801-814
- Cain, L., Hernandez-Cortes, D., Timmins, C., and Weber, P. (2024), Recent Findings and Methodologies in Economics Research in Environmental Justice, «Review of Environmental Economics and Policy», 18, 1, pp. 116-142
- Carattini, S., Baranzini, A., Thalmann, P., Varone, F., and Vöhringer, F. (2017), *Green Taxes in a Post-Paris World: Are Millions of Nays Inevitable?*, «Environmental and Resource Economics», 68, 1, pp. 97-128



- Carattini, S., Kallbekken, S., and Orlov, A. (2019), How to Win Public Support for a Global Carbon Tax, «Nature», 565, 7739, pp. 289-291
- Carozzi, F. and Roth, S. (2023), *Dirty Density: Air Quality and The Density of American Cities*, «Journal of Environmental Economics and Management», 118, p. 102767.
- Chestnut, L.G. and De Civita, P. (2009), Economic Valuation of Mortality Risk Reduction: Review and Recommendations for Policy and Regulatory Analysis. Government of Canada, Policy Research Initiative
- Colmer, J., Hardman, I., Shimshack, J., and Voorheis, J. (2020), Disparities in PM 2.5 air pollution in the United States, «Science», 369, 6503, pp. 575-578
- Currie, J., Voorheis, J., and Walker, R. (2023), What Caused Racial Disparities in Particulate Exposure to Fall? New Evidence from the Clean Air Act and Satellite-Based Measures of Air Quality, «American Economic Review», 113, 1, pp. 71-97
- 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
- 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, pp. 2958-2989
- Doudchenko, N. and Imbens, G.W. (2016), Balancing, Regression, Difference-In-Differences and Synthetic Control Methods: A Synthesis, Working Paper 22791, National Bureau of Economic Research
- 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, pp. 36-65
- Fowlie, M., Holland, S.P., and Mansur, E.T. (2012), What Do Emissions Markets Deliver and to Whom? Evidence from Southern California's NOx Trading Program, «American Economic Review», 102, 2, pp. 965-993
- Fowlie, M., Rubin, E., and Walker, R. (2019), *Bringing Satellite-Based Air Quality Estimates Down to Earth*, «AEA Papers and Proceedings», 109, pp. 283-288
- Freeman, R., Liang, W., Song, R., and Timmins, C. (2019), Willingness uo Pay For Clean Air in China, «Journal of Environmental Economics and Management», 94, pp. 188-216
- Gehrsitz, M. (2017), The Effect Of Low Emission Zones On Air Pollution And Infant Health, «Journal of Environmental Economics and Management», 83, pp. 121-144
- Graff, Z.J. and Neidell, M. (2012), *The Impact of Pollution on Worker Productivity*, «American Economic Review», 102, 7, pp. 3652-3673
- Grainger, C. and Ruangmas, C. (2018), Who Wins from Emissions Trading? Evidence from California, «Environmental and Resource Economics», 71, pp. 703-727
- Grainger, C.A. and Kolstad, C.D. (2010), Who Pays A Price on Carbon?, «Environmental and Resource Economics», 46, pp. 359-376
- Grainger, C. and Schreiber, A. (2019), *Discrimination in Ambient Air Pollution Monitoring?*, «AEA Papers and Proceedings», 109, pp. 277-282
- Harrison, K. (2012), A Tale of Two Taxes: The Fate of Environmental Tax Reform in Canada, «Review of Policy Research», 29, 3, pp. 383-407
- Health Canada (2021), Health Impacts of Air Pollution in Canada: Estimates of morbidity and premature mortality outcomes 2021 Report



- Heckman, J.J. and Smith, J.A. (1999), The Pre-Programme Earnings Dip and the Determinants of Participation in a Social Programme. Implications for Simple Programme Evaluation Strategies, «The Economic Journal», 109, 457, pp. 313-348
- 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, p. 104786
- Iacus, S.M., King, G., and Porro, G. (2012), Causal Inference Without Balance Checking: Coarsened Exact Matching, «Political Analysis», 20, 1, pp. 1-24
- Iceland, J. (2004), The Multigroup Entropy Index (Also Known as Theil's H or The Information Theory Index), US Census Bureau
- Imbens, G.W. (2015), *Matching Methods in Practice: Three Examples*, «The Journal of Human Resources», 50, 2, pp. 373-419
- 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, pp. 228-233
- Knittel, C.R. and Sandler, R. (2011), Cleaning the Bathwater with the Baby: The Health Co-Benefits of Carbon Pricing in Transportation, Working Papers 1115, Massachusetts Institute of Technology, Center for Energy and Environmental Policy Research
- Krewski, D. et al. (2009), Extended Follow-Up and Spatial Analysis of The American Cancer Society Study Linking Particulate Air Pollution and Mortality, Research Reports, Health Effects Institute, Discussion 115-36.140, pp. 5-114
- Lawley, C. and Thivierge, V. (2018), Refining The Evidence: British Columbia's Carbon Tax and Household Gasoline Consumption, «The Energy Journal», 39, 2
- Lepeule, J., Laden, F., Dockery, D., and Schwartz, J. (2012), Chronic Exposure to Fine Particles and Mortality: an Extended Follow-Up of The Harvard Six Cities Study from 1974 to 2009, «Environmental Health Perspectives», 120, 7, pp. 965-970
- Leroutier, M. (2022), Carbon Pricing and Power Sector Decarbonization: Evidence from The UK, «Journal of Environmental Economics and Management», 111, p. 102580
- Li, M., Zhang, D., Li, C.-T., Mulvaney, K.M., Selin, N.E., and Karplus, V.J. (2018), *Air Quality Co-Benefits of Carbon Pricing in China*, «Nature Climate Change», 8, 5, pp. 398-403
- Menchetti, F., Cipollini, F., and Mealli, F. (2022), Combining Counterfactual Outcomes and ARIMA Models for Policy Evaluation, «The Econometrics Journal», 26, 1, pp. 1-24
- Meng, J., Li, C., Martin, R.V., Donkelaar, A. van, Hystad, P., and Brauer, M. (2019), Estimated Long-Term (1981–2016) Concentrations of Ambient Fine Particulate Matter across North America from Chemical Transport Modeling, Satellite Remote Sensing, and Ground-Based Measurements, «Environmental Science & Technology», 53, 9, pp. 5071-5079
- Metcalf, G.E. (2009), Designing a Carbon Tax to Reduce US Greenhouse Gas Emissions, «Review of Environmental Economics and Policy», 3, 1, pp. 63-83
- Mildenberger, M., Lachapelle, E., Harrison, K., and Stadelmann-Steffen, I. (2022), *Limited Impacts of Carbon Tax Rebate Programmes On Public Support For Carbon Pricing*, «Nature Climate Change», 12, 2, pp. 141-147
- Ministry of Finance (2009), Budget and Fiscal Plan 2009/10-2011/12, Tech. rep. British Columbia Ministry of Finance
- (2013), Budget and Fiscal Plan 2013/14-2015/16, Tech. rep. British Columbia Ministry of Finance



- Murray, B. and Rivers, N. (2015), British Columbia' S Revenue-Neutral Carbon Tax: A Review of The Latest Grand Experiment' in Environmental Policy, «Energy Policy», 86, pp. 674-683
- Pampalon, R., Hamel, D., Gamache, P., Philibert, M.D., Raymond, G., and Simpson, A. (2012), *An Areabased Material and Social Deprivation Index for Public Health in Qu'ebec and Canada*, «Canadian Journal of Public Health / Revue Canadienne de Sante'e Publique», 103, S17-S22
- 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, 4, p. 1550019
- Persico, C. and Marcotte, D.E. (2022), *Air Quality and Suicide*, Working Paper 30626, National Bureau of Economic Research
- Pestel, N. and Wozny, F. (2021), Health Effects of Low Emission Zones: Evidence from German Hospitals, «Journal of Environmental Economics and Management», 109, p. 102512
- Poterba, J. (1991), *Is the Gasoline Tax Regressive?*, «Tax Policy and the Economy», vol. 5, National Bureau of Economic Research, pp. 145-164
- Pretis, F. (2022), *Does a Carbon Tax Reduce CO₂ Emissions? Evidence from British Columbia*, «Environmental & Resource Economics», 83, 1, pp. 115-144
- Rivers, N. and Schaufele, B. (2015), Salience of Carbon Taxes in The Gasoline Market, «Journal of Environmental Economics and Management», 74, pp. 23-36
- Rose, A., McKee, J., Sims, K., Bright, E., Reith, A., and Urban, M. (2020), LandScan Global 2020, Oak Ridge National Laboratory, Oak Ridge.
- Saberian, S. (2017), The Negative Effect of the BC Carbon Tax on Vancouver Air Quality: A Good Climate for Bad Air?, Working Paper
- Sager, L. (2023), The Global Consumer Incidence of Carbon Pricing: Evidence from Trade, «Energy Economics», 127, p. 107101
- Sager, L. and Singer, G. (2024), Clean Identification? The Effects of the Clean Air Act on Air Pollution, Exposure Disparities and House Prices, «American Economic Journal: Economic Policy» [forthcoming].
- Sarmiento, L., Wagner, N., and Zaklan, A. (2023), *The Air Quality and Well-Being Effects of Low Emission Zones*, «Journal of Public Economics», 227, p. 105014
- Shapiro, J.S. and Walker, R. (2021), Where Is Pollution Moving? Environmental Markets and Environmental Justice, «AEA Papers and Proceedings», 111, pp. 410-414
- Sheriff, G. (2024), California's GHG Cap-and-Trade Program and the Equity of Air Toxic Releases, «Journal of the Association of Environmental and Resource Economists», 11, 1, pp. 137-170
- Shindell, D.T., Lee, Y., and Faluvegi, G. (2016), Climate and Health Impacts of US Emissions Reductions Consistent With 2°C, «Nature Climate Change», 6, 5, pp. 503-507
- Statistics Canada (2021a), Mortality and Potential Years of Life Lost, By Selected Causes of Death and Sex, Three-Year Average, Census Metropolitan Areas, Table 13-10-0741-01.
- Statistics Canada (2021b), Supply and Disposition of Refined Petroleum Products, Monthly, Table 25-10-0044-01
- Statistics Canada (2021c), Upper Income Limit, Income Share And Average Of Adjusted Market, Total And After-Tax Income By Income Decile, Table 11-10-0193-01
- Tinbergen, J. (1952), On the Theory of Economic Policy, Amsterdam, North-Holland Publishing Company.
 Vandyck, T., Keramidas, K., Kitous, A., Spadaro, J., Van Dingenen, R., Holland, M., and Saveyn, B.
 (2018), Air Quality Co-Benefits For Human Health And Agriculture Counterbalance Costs To Meet Paris
 Agreement Pledges, 9.4939




- Watts, N. et al. (2015), Health and Climate Change: Policy Responses to Protect Public Health, «The Lancet», 386, 10006, pp. 1861-1914
- Wen, J. and Burke, M. (2022), Wildfire Smoke Exposure Worsens Learning Outcomes, «Nature Sustainability», 5, pp. 920-921
- West, J., Smith, S., Silva, R., Naik, V., Zhang, Y., Adelman, Z., Fry, M., Anenberg, S., Horowitz, L., and Lamarque, J.-F. (2013), *Co-Benefits of Global Greenhouse Gas Mitigation for Future Air Quality and Human Health*, «Nature Climate Change», 3
- Zhang, W.-W., Zhao, B., Ding, D., Sharp, B., Gu, Y., Xu, S.-C., Xing, J., Wang, S.-X., Liou, K.-N., and Rao, L.-L. (2021), Co-Benefits of Subnationally Differentiated Carbon Pricing Policies in China: Alleviation of Heavy PM_{2.5} Pollution and Improvement in Environmental Equity, «Energy Policy», 149, p. 112060
- van Donkelaar, A., Martin, R.V., Li, C., and Burnett, R.T. (2019), Regional Estimates of Chemical Composition of Fine Particulate Matter Using a Combined Geoscience-Statistical Method with Information from Satellites, Models, and Monitors, «Environmental Science & Technology», 53, 5, pp. 2595-2611
- von Bergmann, J. (2021), tongfen: R Package to Make Data Based on Different Geographies Comparable, R package version 0.3.3
- von Bergmann, J., Shkolnik, D., and Jacobs, A. (2022), cancensus: R Package to Access, Retrieve, and Work with Canadian Census Data and Geography, R package version 0.5.3

APPENDIX


A1. DESCRIPTIVE STATISTICS AND GRAPHS

FIGURE A.1 • SPATIO-TEMPORAL PLACEMENT OF GROUND AIR POLLUTION MONITORS



Notes: Availability of PM_{2.5} readings in the National Atmospheric Surveillance Program database between 2000 and 2018. Lighter colours and larger dot sizes indicate higher availability of readings (monitoring stations which were added earlier).

FIGURE A.2 • SATELLITE AND GROUND PM_{2.5} READINGS



Notes: Scatterplot of satellite PM_{2.5} (Meng *et al.*, 2019) (y-axis) and PM_{2.5} from NAPS monitoring stations (x-axis). Both measures are in μ g/m3. The correlation coefficient is 0.597.






FIGURE A.3 • SATELLITE AND GROUND PM_{2.5} READINGS

Notes: Density plot of satellite $PM_{2.5}$ (Meng et al., 2019) (y-axis) and $PM_{2.5}$ from NAPS monitoring stations (x-axis). Both measures are in $\mu g/m3$.




FIGURE A.4 • TWO REMOTELY SENSED PM_{2.5} MEASURES

Notes: Scatterplot of satellite PM_{2.5} (Meng et al., 2019) (x-axis) vs Satellite PM_{2.5} (van Donkelaar et al., 2019) (y-axis). Both measures are in μ g/m3. The correlation coefficient is 0.729.






FIGURE A.5 • TWO REMOTELY SENSED PM_{2.5} MEASURES

Notes: Density plot of satellite PM_{2.5} (Meng et al., 2019) (x-axis) vs Satellite PM_{2.5} (van Donkelaar et al., 2019) (y-axis). Both measures are in μ g/m3.

TABLE A.1 • SUMMARY STATISTICS 2000-2007

	Co	ntrol Provin	ces	Br	tish Columb	ia
	N	Mean	SD	N	Mean	SD
PM _{2.5} (van Donkelaar et al. 2019)	175870	9.52	1.54	27920	8.06	1.19
PM _{2.5} (Meng et al. 2019)	175870	8.61	2.07	27920	6.95	1.39
Pop. Density (Rose et al. 2020)	175912	3358.26	3375.33	27920	3169.94	2136.98
Median Income	43978	26341.65	9088.59	6980	25055.65	8090.75
Minority Share	43709	18.07	20.7	6926	34.25	26.22
Theil's Diversity Entropy Index	43978	0.53	0.47	6980	0.79	0.44
Material Deprivation Index	20606	46.48	28.60	3313	43.57	28.12
High Emissions Commute %	43701	74.93	18.33	6926	77.64	16.37
Low Emissions Commute %	43701	24.52	18.26	6926	21.62	16.26
Public Transport Commute %	43701	17.02	14.48	6926	13.06	10.68
Zero Emissions Commute %	43701	7.50	9.43	6926	8.56	10.79
Precipitation (Abatzoglou et al. 2018)	175870	74.05	21.74	27920	131.20	37.06
Max Temp. (Abatzoglou et al. 2018)	175870	11.93	1.58	27920	14.55	0.66
Min. Temp (Abatzoglou et al. 2018)	175870	1.74	2.50	27920	6.46	0.62
Wind Speed (Abatzoglou et al. 2018)	175870	3.63	0.49	27920	2.98	0.16



Control Provinces British Columbia N Mean SD N Mean SD PM_{2.5} (van Donkelaar et al. 2019) 241865 8.15 1.51 38390 6.09 0.95 PM_{2.5} (Meng et al. 2019) 197888 7.35 1.73 31410 6.07 1.10 Pop. Density (Rose et al. 2020) 241879 3614.39 3365.36 38390 3478.58 2305.69 Median Income 43978 33324.06 11718.48 31772.63 9765.79 6980 Minority Share 43825 23.84 22.86 6958 40.2 27.15 Theil's Diversity Entropy Index 43978 0.65 0.51 6980 0.89 0.47 High Emissions Commute % 43806 74.39 20.20 6955 72.94 18.75 43806 25.06 20.12 6955 26.25 Low Emissions Commute % 18.61 Public Transport Commute % 18.85 15.89 6955 13.38 43806 18.38 Zero Emissions Commute % 43806 6.21 10.15 6955 7.87 11.32 Precipitation (Abatzoglou et al. 2018) 77.57 21.99 38390 134.58 37.33 241861

12.32

2.11

3.64

1.76

2.59

0.48

38390

38390

38390

14.58

6.56

3.00

0.86

0.80

0.19

TABLE A.2 • SUMMARY STATISTICS 2008-2018

A2. ADDITIONAL DETAILS ON THE EMPIRICAL STRATEGY

241861

241861

241861

A2.1. The DID parallel trends assumption

Max Temp. (Abatzoglou et al. 2018)

Min. Temp (Abatzoglou et al. 2018)

Wind Speed (Abatzoglou et al. 2018)

TABLE A.3 • PLACEBO DID REGRESSIONS ON 2000-2007 DATA

	(1)	(2)
	DID	DID
τ̂	-0.3754	-1.355
	(0.1362)	(0.2390)
Unit FE	Yes	Yes
Year FE	Yes	Yes
N_{obs}	203736	203736

Notes: All point estimates represent the placebo impact of a carbon tax assigned in 2004 during the 2005-2007 placebo post-treatment period. All regressions use 2000-2007 data. Column (1) uses the Meng et al. (2019) PM_{2.5} data, and column (2) uses the van Donkelaar et al. (2019) PM_{2.5} data. Standard errors in parentheses are clustered at the CMA level.






FIGURE A.6 • PM_{2.5} PRE-TRENDS DIFFERENCES (MENG ET AL., 2019)

Notes: Trends and observations of satellite $PM_{2.5}$ in British Columbia and an average of all control provinces, between 2000 and 2016. The implementation of the carbon tax in 2008 is highlighted by the dashed vertical line. Observed data is jittered around the observation year to enhance legibility. A placebo FE regression of $PM_{2.5}$ on the traditional DID indicator with data limited to 2000-2008 and treatment assigned in 2004 identifies a negative and significant effect, indicating a likely failure of the parallel trends hypothesis.




FIGURE A.7 • PM_{2.5} PRE-TRENDS DIFFERENCES (VAN DONKELAAR ET AL., 2019)

Notes: Trends and observations of satellite PM2.5 in British Columbia and an average of all control provinces, between 2000 and 2018. The implementation of the carbon tax in 2008 is highlighted by the dashed vertical line. Observed data is jittered around the observation year to enhance legibility.



A2.2 Comparison between TWFE-DID, SCM and SDID

In order to formally explain how SDID combines features from TWFE-DID and SCM, let me consider a balanced panel with N observations and T time periods. In the British Columbian case, the outcome variable is $PM_{2.5it}$, and the binary treatment is D_{it} . Let $i = 1, ..., N_{tr}$ be the treated DAs in BC, Let $i = N_{tr+1}, ..., N_{co}$ be the DAs in control provinces. The baseline TWFE-DID regression problem can be expressed as:

$$(\hat{\tau}^{did}, \hat{\mu}, \hat{\eta}, \hat{\theta}) = \underset{\tau, \mu, \eta, \theta}{\operatorname{argmin}} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} (PM_{2.5it} - \mu - \eta_i - \theta_t - \tau D_{it})^2 \right\}$$
(9)

Which is solved without the use of unit or time-specific weights, but with the inclusion of unit and time-specific fixed effects ηi and θt as also illustrated in Equation 2. The SCM estimator, instead, does not employ unit fixed effects, but includes time fixed effects and unit-specific weights ω_i^{sc} :

$$\left(\hat{\tau}^{sc}, \hat{\mu}, \hat{\eta}, \hat{\theta}\right) = \underset{\tau, \mu, \eta, \theta}{\operatorname{argmin}} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} (PM_{2.5it} - \mu - \eta_i - \theta_t - \tau D_{it})^2 \widehat{\omega_i^{sc}} \right\}$$
(10)

Finally, the SDID estimator combines features from Equation 9 and Equation 10. Unit weights ω_i^{sdid} are chosen such that the pre-treatment outcome path of control DAs are parallel to those of the treated units⁴⁶:

$$\omega_0 + \sum_{i=N_{tr}+1}^{N_{co}} \widehat{\omega}_i^{sdid} P M_{2.5it} \approx \frac{1}{N_{tr}} \sum_{i=1}^{N_{tr}} P M_{2.5it}$$
 (11)

Moreover, time weights λ 's did need to ensure that the pre-treatment levels for the control units differs from the post-treatment levels for the same units only by a constant. Letting t = 1, ..., T be the total length of the panel, T_{pre} be the number of pre-intervention periods, and T_{post} be the number of post-intervention periods, the condition can be expressed as:

$$\lambda_0 + \sum_{t=1}^{T_{pre}} \hat{\lambda}_t^{sdid} P M_{2.5it} \approx \frac{1}{T_{post}} \sum_{t=T_{pre}+1}^{T} P M_{2.5it}$$
 (12)

⁴⁶ When the intercept $ω_0$ and the regularisation parameter are set to 0, the unit weights $ω_i$ correspond to the SCM weights in Abadie *et al.* (2010). For further details on the procedure used to estimate ζ, please refer to Arkhangelsky *et al.* (2021).



Thus, the regression problem for the SDID estimator can be expressed as a weighted TWFE-DID problem which incorporates unit and time-specific fixed effects η_i and θ_t , plus unit and time-specific weights ω_i and λ_t , as illustrated in Equation 13:

$$\left(\hat{\tau}^{sdid}, \hat{\mu}, \hat{\eta}, \hat{\theta}\right) = \underset{\tau, \mu, \eta, \theta}{\operatorname{argmin}} \left\{ \sum_{i=1}^{N} \sum_{t=1}^{T} (PM_{2.5it} - \mu - \eta_i - \theta_t - \tau D_{it})^2 \widehat{\omega_i}^{sdid} \, \hat{\lambda}_t^{sdid} \right\}$$
(13)

A3. ADDITIONAL RESULTS AND ROBUSTNESS CHECKS

A3.1. Matched Difference-in-Differences Plots




FIGURE A.8 • PRE-TRENDS IN PM_{2.5} AFTER CEM (1)

Notes: Trends and observations of satellite PM_{2.5} (Meng *et al.*, 2019) in British Columbia and an average of all control provinces, between 2000 and 2016. The implementation of the carbon tax in 2008 is highlighted by the dashed vertical line. Observed data is jittered around the observation year to enhance legibility. The data results from a Coarsened Exact Matching procedure on baseline PM_{2.5} levels.






FIGURE A.9 • PRE-TRENDS IN PM_{2.5} AFTER CEM (2)

Notes: Trends and observations of satellite PM_{2.5} (Meng *et al.*, 2019) in British Columbia and an average of all control provinces, between 2000 and 2016. The implementation of the carbon tax in 2008 is highlighted by the dashed vertical line. Observed data is jittered around the observation year to enhance legibility. The data results from a Coarsened Exact Matching procedure on baseline PM_{2.5} levels and baseline levels of population density, median income, high emission commute mode share, and road density.

A3.2. Composition of SDID control pools




FIGURE A.10 • SYNTHETIC BC AT THE DA LEVEL (MENG ET AL., 2019)

Notes: Composition of the synthetic DID unit of Figure 2. Individual DA weights are aggregated up to the CMA level.








FIGURE A.11 • SYNTHETIC BC AT THE DA LEVEL (VAN DONKELAAR ET AL., 2019)

Notes: Composition of the synthetic DID unit of Figure A.12. Individual DA weights are aggregated up to the CMA level.

FIGURE A.12 • THE IMPACT OF THE 2008 CARBON TAX ON CHANGES IN PM2.5, VAN

A3.3. Main results with van Donkelaar et al. (2019) PM_{2.5} data





Notes: Graphical results from DID, SCM and SDID for PM2.5 concentrations, with van Donkelaar et al. (2019) data. Time weights λ_t are represented in light red at the bottom of the pre-intervention panel. The curved arrows graphically represent the ATT over the post-intervention period.




	(1) DID	(2) SCM	(3) SDID
î	-0.4954 (0.0085)	-0.7087 (0.1540)	-0.8896 (0.0300)
Unit FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
ω_i		Yes	Yes
λ_t			Yes
N_{obs}	483873	483873	483873

TABLE A.4 • THE 2008 CARBON TAX AND CHANGES IN PM_{2.5}

Notes: All point estimates represent the average impact of the 2008 carbon tax during the 2009-2018 post-treatment period. Standard errors in parentheses are calculated using the bootstrap variance estimation algorithm described in Arkhangelsky *et al.* (2021) with 200 replications. All regressions use 2000-2016 data.

A3.4. Accounting for measurement error à la Fowlie et al. (2019)

FIGURE A.13 • THE IMPACT OF 2008 CARBON TAX ON CHANGES IN PREDICTED $PM_{2.5}$ ACCOUNTING FOR MEASUREMENT ERROR



Notes: Graphical results from SDID for PM_{2.5} concentrations, with Meng *et al.* (2019) data (left panel) and van Donkelaar *et al.* (2019) data (right panel). Time weights λ_t are represented in light red at the bottom of the pre-intervention panel. The curved arrows graphically represent the ATT over the post-intervention period.




TABLE A.5 ● THE 2008 CARBON TAX AND CHANGES IN PM_{2.5}, ACCOUNTING FOR MEASUREMENT ERROR

	(1) SDID	(2) SDID
î	-0.260 (0.010)	-0.847 (0.027)
Unit FE	Yes	Yes
Year FE	Yes	Yes
ω_i	Yes	Yes
λ_t	Yes	Yes
N_{obs}	432939	432939

Notes: Point estimates represent the average impact of the 2008 carbon tax using data from Meng et al. (2019) (column 1) and van Donkelaar et al. (2019) (column 2) as the outcome, during the 2009-2016 post-treatment period. Standard errors are calculated using the bootstrap variance estimation algorithm described in Arkhangelsky et al. (2021) with 200 replications. All regressions use 2000-2016 data.

A3.5. Post-treatment period limited to 2013

FIGURE A.14 • THE IMPACT OF THE 2008 CARBON TAX ON CHANGES IN PM_{2.5}, 2000-2013 SAMPLE



Notes: Graphical results from DID, SCM and SDID for PM_{2.5} concentrations, with Meng *et al.* (2019) data, dataset restricted to 2013. Time weights λ_t are represented in light red at the bottom of the pre-intervention panel. The curved arrows graphically represent the ATT over the post-intervention period.




	(1)	(2)	(3)	
	DID	SCM	SDID	
î	0.0547 (0.0081)	-0.2723 (0.0803)	-0.6703 (0.0341)	
Unit FE	Yes	Yes	Yes	
Year FE	Yes	Yes	Yes	
ω_i		Yes	Yes	
λ_t			Yes	
N_{obs}	432939	432939	432939	

TABLE A.6 • THE 2008 CARBON TAX AND CHANGES IN PM_{2.5}, 2000-2013 SAMPLE

Notes: All point estimates represent the average impact of the 2008 carbon tax using data from Meng et al. (2019) as the outcome, during the 2009-2013 post-treatment period. Standard errors are calculated using the bootstrap variance estimation algorithm described in Arkhangelsky et al. (2021) with 200 replications. All regressions use 2000-2013 data.

A3.6. DAs in the Vancouver CMA

FIGURE A.15 • THE IMPACT OF THE 2008 CARBON TAX ON CHANGES IN PM_{2.5} FOR DAS IN THE VANCOUVER CMA



Notes: Graphical results from DID, SCM and SDID for PM_{2.5} concentrations, with Meng *et al.* (2019) data, dataset restricted to DAs in the Vancouver CMA. Time weights λ_t are represented in light red at the bottom of the pre-intervention panel. The curved arrows graphically represent the ATT over the post-intervention period.




TABLE A.7 • THE 2008 CARBON TAX AND CHANGES IN PM2.5, DAS IN THE VANCOUVER CMA

	(1)	(2)	(3)	
	DID	SCM	SDID	
î	0.4061 (0.0071)	-0.1062 (0.0702)	-0.3014 (0.0225)	
Unit FE	Yes	Yes	Yes	
Year FE	Yes	Yes	Yes	
ω_i		Yes	Yes	
λ_t			Yes	
N _{obs}	432939	432939	432939	

Notes: All point estimates represent the average impact of the 2008 carbon tax during the 2009-2016 post-treatment period, using the outcome variable from Meng et al. (2019) and restricting the sample to DAs in the Vancouver CMA. Standard errors are calculated using the bootstrap variance estimation algorithm described in Arkhangelsky et al. (2021) with 200 replications. All regressions use 2000-2016 data.

A3.7. DAs matching NAPS monitoring stations

FIGURE A.16 • THE IMPACT OF THE 2008 CARBON TAX ON CHANGES IN PM_{2.5} FOR DAS MATCHING NAPS LOCATIONS



Notes: Graphical results from DID, SCM and SDID for PM_{2.5} concentrations, with Meng *et al.* (2019) data, dataset restricted to DAs matching NAPS monitoring stations' locations. Time weights λ_t are represented in light red at the bottom of the pre-intervention panel. The curved arrows graphically represent the ATT over the post-intervention period.



	(1)	(2)	(3)
	DID	SCM	SDID
î	0.132	-0.865	-0.288
	(0.117)	(0.128)	(0.097)
Unit FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
ω_i		Yes	Yes
λ_t			Yes
N_{obs}	2227	2227	2227


TABLE A.8 • THE 2008 CARBON TAX AND CHANGES IN PM_{2.5}, DAS MATCHING NAPS LOCATIONS

Notes: All point estimates represent the average impact of the 2008 carbon tax during the 2009-2016 post-treatment period, using the outcome variable from Meng et al. (2019) and restricting the sample to DAs in the Vancouver CMA. Standard errors are calculated using the bootstrap variance estimation algorithm described in Arkhangelsky et al. (2021) with 200 replications. All regressions use 2000-2016 data.

A4. ADDITIONAL DETAILS ON MECHANISMS

A4.1. Fuel demand

FIGURE A.17 • C-ARIMA RESULTS FOR FUEL SALES



Notes: Graphical results from the C-ARIMA regressions on monthly gasoline and diesel sales. Panel (A) and (B) show the observed and forecasted gasoline and diesel sales time series for the full post-intervention horizon. Panels (C) and (D) represent the gap between observed and forecasted series for gasoline and diesel sales. 95% confidence intervals in grey shading using robust standard errors.






FIGURE A.18 • EVENT STUDY RESULT FOR FUEL SALES

Notes: Event study regression of monthly gasoline sales at the province level, 1990-2016 data. 95% confidence interval in grey shading using standard errors clustered at the province level.

A4.2. Commute mode switching empirical analysis

Ideally, when concerned with the estimation of PM_{2.5} reductions arising from the implementation of carbon pricing, I would look at DA-level reductions in motor fuel sales or in the quantity of vehicle kilometres travelled; however, these data are not available at the desired level of granularity for Canada between 2000 and 2018. Moreover, while Canadian province-level data on vehicle sales disaggregated by type of fuel is only available from 2011 onwards, the post-2011 trends in sales of diesel vehicles are relatively flat (See Figure A.20), and the landscape seems to be dominated by gasoline cars (See Figure A.19), suggesting that an eventual gas-to-diesel switch caused by the carbon tax incentive would have produced all of its results between July 2008 and January 2011 before bottoming out; the evidence for this conclusion is not very strong as a result. Another potential mechanism behind an increase in air pollution could derive from an exceptionally high rate of replacement in BC's car fleet with respect to other Canadian provinces, caused by the willingness of BC's residents to increase their cars' fuel efficiency and realise savings at the pump. If the savings per each tank refuel were sufficient to offset the increase in gasoline prices due to the carbon tax, British Columbian residents could have potentially travelled more kilometres than prior to the tax, thereby increasing road congestion and hence pollution due to a rebound effect. As shown in Figure A.21 there has indeed been a rapid increase in truck and SUV sales in British Columbia after 2008; however, this increase is



paralleled by similar jumps in truck sales in all large Canadian provinces⁴⁷, and it thus seems implausible to attribute it to the marginal effect of the carbon tax in raising fuel prices.

I instead exploit the information contained in the 2001, 2006, 2011 and 2016 waves of the Canadian census, which contains data on commute-to-work modes at the DA level for all Canadian CMAs. While the information on commute modes is not an exhaustive representation of all car trips made in each DA, the granularity of the data may shed light on whether residents of DAs located in British Columbia have adjusted their behaviour following the implementation of the carbon tax, substituting public transport or active commuting modes such as cycling and walking for car trips. In particular, I estimate the following equation:

$$Mode_{it} = \tau D_{it} + \theta_t + \eta_i + \varepsilon_{it} \tag{14}$$

Where $Mode_{it}$ is the share of each commute mode (high emission, low emission, public transport and zero emission), D_{it} is the carbon tax DID binary variable, θ_t and η_i are time and unit-specific fixed effects, and ε_{it} is an idiosyncratic error term. In additional specifications, I also add a vector of controls X_{it} which account for population density, median income, and weather covariates (precipitation, maximum and minimum temperature, and wind speed), hence the estimating equation becomes:


$$Mode_{it} = \tau D_{it} + \beta X_{it} + \theta_t + \eta_i + \varepsilon_{it}$$
 (15)

I initially run the TWFE-DID regressions for the whole sample, without trimming the control pool. In further specifications, I restrict the control sample to the units which receive positive ω_i weights in the SDID estimation of the main result, in order to ensure comparability across treatment and control cohorts and reduce the reliance on potentially violated parallel trends. Further, I retrieve the ω_i weights from the SDID estimation and weigh my restricted TWFE-DID regressions with the SDID weights, assigning equal weights $\frac{1}{N_{tr}}$ to the treatment cohort.

⁴⁷ Namely, Alberta, Ontario and Quebec.




FIGURE A.19 • NEW VEHICLE REGISTRATIONS IN BRITISH COLUMBIA BY MAJOR FUEL TYPES



Notes: The plot illustrates new vehicle registrations for gasoline and all other fuel types in British Columbia between 2011 and 2021.


FIGURE A.20 • NEW MINOR VEHICLE GROUPS REGISTRATIONS IN BRITISH COLUMBIA



Notes: The plot illustrates new vehicle registrations for low emissions (battery electric, plug-in hybrid, hybrid) and diesel fuel types in British Columbia between 2011 and 2021.







Notes: The plot illustrates 1990-2021 trends in passenger cars and truck and SUV sales, for the four largest Canadian Provinces. Top row: British Columbia and Alberta. Bottom row: Ontario and Quebéc.



TABLE A.9 • DID RESULTS FOR LOW EMISSIONS COMMUTE MODE

		Lo	w Emissions (Commute Mo	ode	
	(1)	(2)	(3)	(4)	(5)	(6)
DID	0.0408 (0.111)	0.0516 (0.0103)	0.0535 (0.0110)	0.0457 (0.0109)	0.0510 (0.0113)	0.0506 (0.0114)
DA FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls				Yes	Yes	Yes
SDID control pool		Yes	Yes		Yes	Yes
SDID weights			Yes			Yes
R ²	0.87321	0.84174	0.84532	0.87715	0.84674	0.84996
Adj. R ²	0.83078	0.78876	0.79354	0.83560	0.79490	0.79920
N_{obs}	101358	38769	38769	100244	38348	38348

Notes: The dependent variable is the dissemination area level share of low emissions commutes. All regressions include dissemination area and year fixed effects. Columns (4)-(6) include controls for precipitation, maximum and minimum temperature, and wind speed, plus the natural logarithm of population and median income. Columns (2), (3), (5) and (6) restrict the control unit pool to DAs which receive positive weights in the main SDID regression.

TABLE A.10 • DID RESULTS FOR ZERO EMISSIONS COMMUTE MODE

		Zero I	Emissions Con	nmute Mode		
	(1)	(2)	(3)	(4)	(5)	(6)
DID	0.0057 (0.0025)	0.0106 (0.0017)	0.0117 (0.0021)	0.0066 (0.0022)	0.0088 (0.0016)	0.0092 (0.0016)
DA FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls				Yes	Yes	Yes
SDID control pool		Yes	Yes		Yes	Yes
SDID weights			Yes			Yes
R ²	0.80811	0.80808	0.81877	0.81200	0.81355	0.82463
Adj. R ²	0.74390	0.74383	0.75810	0.74841	0.75047	0.76531
N_{obs}	101358	38769	38769	100244	38348	38348

Notes: The dependent variable is the dissemination area level share of zero emissions commutes. All regressions include dissemination area and year fixed effects. Columns (4)-(6) include controls for precipitation, maximum and minimum temperature, and wind speed, plus the natural logarithm of population and median income. Columns (2), (3), (5) and (6) restrict the control unit pool to DAs which receive positive weights in the main SDID regression.



A5. Environmental justice dynamics

A5.1. Descriptive statistics

TABLE A.11 • TRENDS IN PM_{2.5} BY QUINTILE OF EJ DIMENSIONS, BRITISH COLUMBIAN DAS

		EJ Dimension			
	Pop. Density	Diversity	Income		
		Baseline (2000-2002 Average)			
Top Quintile	8.66	8.83	6.68		
Bottom Quintile	6.48	6.83	8.53		
EJ Gap	2.18	2.01	-1.85		
		Post-treatment (2014-2016 Average)			
Top Quintile	6.63	6.67	5.47		
Bottom Quintile	5.32	5.46	6.49		
EJ Gap	1.31	1.21	-1.02		

Notes: All values are expressed in $\mu g/m^3$, for British Columbian DAs only. Baseline PM_{2.5} levels are calculated as 2000-2002 averages for all quintiles, post treatment PM_{2.5} levels are 2014-2016 averages. Quintiles are calculated on 2005-2007 levels for population density, and on 2006 Census levels for racial diversity and median income.

TABLE A.12 • TRENDS IN PM_{2.5} BY QUINTILE OF EJ DIMENSIONS, CONTROL DAS


		EJ Dimension	
	Pop. Density	Diversity	Income
		Baseline (2000-2002 Average)	
Top Quintile	9.88	9.67	8.69
Bottom Quintile	7.71	8.28	9.60
EJ Gap	2.17	1.39	-0.91
	Pe	ost-treatment (2014-2016 Avera	ige)
Top Quintile	7.05	7.10	6.41
Bottom Quintile	5.81	6.13	6.90
EJ Gap	1.24	0.97	-0.49

Notes: All values are expressed in $\mu g/m^3$, for British Columbian DAs only. Baseline PM_{2.5} levels are calculated as 2000-2002 averages for all quintiles, post treatment PM_{2.5} levels are 2014-2016 averages. Quintiles are calculated on 2005-2007 levels for population density, and on 2006 Census levels for racial diversity and median income.



A5.2. Additional results

FIGURE A.22 • ADDITIONAL QUINTILE-SDID RESULTS FOR ENVIRONMENTAL JUSTICE GAPS




Notes: Results of SDID regressions by quintile of baseline characteristics. Panel A) Quintiles of Material Deprivation Index; B) Quintiles of Theil's Diversity Entropy Index. ATT point estimates reported in red, with 95% confidence intervals calculated with the Arkhangelsky et al. (2021) procedure with 200 bootstrap runs in grey shading.



A6. POLLUTION, HEALTH AND DISTRIBUTIONAL IMPLICATIONS

A6.1. Estimates using RR from Krewski et al. (2009)

FIGURE A.23 • MORTALITY REDUCTIONS AND MONETARY HEALTH GAINS



Notes: Spatial distribution of mortality reductions per 1000 residents (left panel) and health gains per capita (right panel) using the RR estimates from Krewski et al. (2009), for the Vancouver (top row), Victoria (middle row) and Abbotsford (bottom row) CMAs.



A6.2. Health-income relationships

FIGURE A.23 • SPATIAL RELATIONSHIP BETWEEN HEALTH GAINS AND MEDIAN INCOME



Notes: Bivariate distribution of health gains using the RR from Lepeule et al. (2012) (left panel) and Krewski et al. (2009) (right panel) and median income for the Vancouver (top row), Victoria (middle row) and Abbotsford (bottom row) CMAs.



ALESSANDRA TESTA, KONSTANTIN BOSS¹

WHAT GOES AROUND COMES AROUND: THE US CLIMATE-ECONOMIC CYCLE

Abstract. We use a spatial data set of US temperatures in a factor-augmented VAR to quantify the contribution of the US economy to fluctuations in temperatures over the past 70 years. We show that there are at least five distinct sources of broad scale temperature fluctuations in the US and uncover a strong relationship of temperatures with aggregate productivity. Disentangling natural from anthropogenic effects, we find that economic expansions do not only lead to warming: technology improvements initially decrease temperatures, whereas investment and labor supply shocks increase them rapidly and persistently. This happens because the cooling effect of aerosol emissions initially outweighs the warming effect from greenhouse gases for technology shocks, but not for investment and labor supply shocks. Taken together, these economic shocks explain around 25% of long-term temperature variation in the US. In turn, temperature shocks induce small contractions in aggregate GDP, but can even be beneficial for the economy, when they predominantly hit the western states.

Keywords. Factor-augmented VAR, climate econometrics, temperature shocks, frequency domain identification

1. Introduction

The rise in global socio-economic activity and the accompanying increase in anthropogenic greenhouse gas (GHG) emissions that characterized the past century are known to be important causes of global warming. Worldwide average surface temperatures have already increased by 1.1°C since the industrial revolution and are projected to increase by between 1.4°C and 4.4°C until 2100 (IPCC, 2023). In turn, temperature increases can lead to lower agricultural yields (Deschênes and Greenstone, 2007), more premature deaths (Barreca et al.,

¹ Boss acknowledges funding under the FPI Grant PRE2020-093941 funded by MCIN/AEI/10.13039/501100011033 and by ESF Investing in your future. We would like to thank Selim Elbadri for a very helpful discussion. We also thank Luca Gambetti, Sarah Zoi, Marco Mazzali, Luca Sala, Tajana Dahlhaus and Christian Brownlees as well as seminar participants at the 64th Annual Conference of the Italian Economic Association (Società Italiana di Economia -SIE), the Banco de España, the Ruhr Graduate School Doctoral Conference, Dynamic Econometrics at Oxford, the JRC Ispra, the Centro Einaudi Torino as well as of the Bellaterra Macro Club for their very valuable comments. We are also grateful to the Centro Einaudi Torino for awarding one of the Giorgio Rota Best Paper Awards 2024 to this paper.



2015), and diminished productivity (Burke et al., 2015), resulting in potentially severe losses in welfare (Bilal and Känzig, 2024).

In this paper, we develop an empirical framework for the United States (US) to study how economic activity has affected temperatures and vice versa. We use a factor-augmented vector autoregression (FAVAR; Bernanke et al., 2005) to model the dynamics of US temperatures on a $0.5^{\circ} \times 0.5^{\circ}$ spatial grid together with key macroeconomic aggregates. To disentangle the effect of human activity on temperatures from the effect of temperatures on human activity, we rely on the notion of structural shocks that is common in causal macroeconomic inference (Ramey, 2016). We use partial identification techniques to pin down three well-established economic shocks in the frequency domain along the lines of Forni et al. (2023). First, a technology shock is identified as the main contributor to lowfrequency variation in utilization-adjusted TFP, similar to DiCecio and Owyang (2010) and Dieppe et al. (2021). Second, conditional on the technology shock, we identify an investment shock in the spirit of Justiniano et al. (2010, 2011) and Auclert et al. (2020) as the main driver of business-cycle fluctuations in investment. Third, similar to Shapiro and Watson (1988), we identify a labor supply shock as the main driver of the low-frequency component of hours worked, conditional on both the technology and investment shocks. On the other hand, we rely on statistical arguments to identify temperature shocks. As Angeletos et al. (2020) identify an economic "main business-cycle shock," we apply a similar reasoning to capture the main drivers of temperature fluctuations in specific geographic areas, such as the west coast, the east coast, the Gulf region, or the non-coastal states, as well as in specific frequency bands, for example, at the El Niño-La Niña periodicities. We then compute the impulse responses of US real GDP to these shocks.

Based on our analysis, we report the following qualitative results: First, it is insufficient to rely on a single measure of national temperatures such as (weighted) averages, as is frequently done in the literature (Dell et al., 2012; Burke et al., 2015; Acevedo et al., 2020). This is because there is a lower bound of five large shocks driving US temperatures. Average temperatures alone only reflect variation in the Midwest region and neglect temperature changes in the economically important coastal areas. This happens because the American Midwest is affected by strong cold air flows from the North and warm air flows from the South, leading to very high temperature variability (Kunkel et al., 2013). Geographic heterogeneity also matters for the effect of temperatures on aggregate GDP, a crucial relationship for environmental policy-making: If warming affects only the west of the country, this can be net positive for the economy, whereas temperature increases generally diminish output slightly. Second, we provide evidence for a relationship between temperatures and socio-economic activity mostly through changes in TFP. A loss in productivity is thought to be one of the main channels for the negative effects of temperature warming (Burke et al., 2015). We argue



along the lines of Pretis (2021) that it is important to properly distinguish if temperature fluctuations cause productivity changes or vice versa. In the case of the US, we find that the majority of the negative co-movement between temperatures and TFP is caused by economic shocks.

In addition, we contribute the following quantitative findings to the literature: First, on average, a quarter of the low-frequency component of US temperatures can be attributed to the three economic shocks, with technology shocks accounting for 10%, investment shocks for 11%, and labor supply shocks for 4%. In the east and south of the US, where manufacturing and natural resource processing are concentrated, the explained variation from technology shocks alone can be as high as 35%. High and medium cycle variations of temperatures, on the other hand, are not strongly explained by anthropogenic shocks. The economic shocks have small, yet persistent effects on temperatures. While technology shocks initially decrease temperatures in the industrial part of the country, this effect recedes in the long run despite the permanent effect on economic activity and emissions. Investment shocks and labor supply shocks lead to geographically homogeneous warming, in the area of 0.01°C, even though the economic expansion is mostly transitory. We argue that decreases in temperatures can be explained by a stronger effect of aerosol emissions than GHG emissions, whereas warming is observed when aerosols are removed and GHGs emitted. Second, central US and east coast-centered increases of 1°C lead to mild losses of aggregate GDP around 0.1%–0.13%. This is in line with the view that the US, for the most part, has been close to a bliss point where temperature warming has so far had essentially zero aggregate effects (see, e.g., Dell et al., 2012; Nath et al., 2023; Natoli, 2023). However, shocks that predominantly affect temperatures on the west coast can have expansionary effects. We find them to lead to up to 0.29% higher GDP after an initial decrease of around 0.32%. This is because when increases in temperatures occur in the west, they are accompanied by decreases in the east. The net effect of this is positive for aggregate real GDP. Temperature shocks are not persistent for temperatures anywhere in the US.

Comprehensive overviews of the climate-econometric literature are provided by Newell et al. (2021) and de Juan *et al.* (2022). The authors show that especially the estimates of economic damages from climate change vary substantially across methodologies. We relate to and expand the literature that quantifies the effect of temperatures on the US economy. Important contributions over the existing empirical literature are as follows: We identify the direct effect on temperatures of economic shocks that explain the bulk of macroeconomic fluctuations. This is necessary because policy-oriented models such as Cai and Lontzek (2019) focus on damages from temperature changes induced by such economic shocks on the economy, although usually relying on TFP shocks alone. In addition, we allow the data to determine the timing of the effects of emissions on temperatures rather than assuming that



economic activity translates into temperature changes with a delay of a year, as is customary in the literature (e.g., Donadelli *et al.*, 2017; Goulet Coulombe and Göbel, 2021), since this is not supported by climate research (e.g., Joos *et al.*, 2013; Forster *et al.*, 2020). Instead, we propose an identification based on statistical arguments with no implied timing restrictions.

Other studies in this area use mostly panel regressions without dynamic causal response estimates (e.g., Deryugina and Hsiang, 2014; Colacito et al., 2019; Gourio and Fries, 2020), which are less focused on the transmission mechanism of temperature fluctuations to the real economy. Kaufmann et al. (2013), Montamat and Stock (2020), and Stock (2020) discuss economic processes affecting climate forcing (and thus temperatures), but do not identify the stochastic processes explicitly. Empirical studies that compute the effects of economic shocks on US CO₂ emissions are Khan et al. (2019), Fosten (2019), and Bennedsen et al. (2021), however, no explicit connection to temperature changes is made. Since the effect of economic activity on temperatures is not exclusively driven by GHG emissions, but also other gases such as aerosols, Magnus et al. (2011), Storelymo et al. (2016), Phillips et al. (2020) provide a breakdown of the respective warming and cooling effects. We show that the aerosol cooling effect prevails for technology shocks, whereas other business cycle shocks lead to warming through a dominant impulse of GHGs. From a methodological view, our paper is closely related to Mumtaz and Marotta (2023), Berg et al. (2023), and Bastien-Olvera et al. (2022). The first two for the authors' use of a factor structure for temperature dynamics and the third one for the frequency domain decomposition of temperatures. While Mumtaz and Marotta (2023) use global data to characterize patterns of aggregate temperature movements, their study focuses on correlations with economic development indicators. We provide causal interpretations for the variations in temperature data and vice versa. Berg et al. (2023) consider only a single factor for their global data set, whereas we show that this captures a very localized temperature phenomenon. Bastien-Olvera et al. (2022) regress GDP growth onto the low-frequency component of average temperatures extracted using a low-pass filter. However, as we show, this component is substantially affected by economic shocks, for which the authors do not control.

The rest of the paper is organized as follows: Section 2 describes the temperature and economic data we use in the empirical model, Section 3 introduces the model and explains the identification methodology, Section 4 presents the findings, which are discussed in Section 5. Finally, Section 6 concludes.

2. DATA

Temperature data are obtained from the Terrestrial Air Temperature and Precipitation 1900–2017 Gridded Monthly dataset (Matsuura and Willmott, 2018), which provides monthly



mean temperatures over land at a 0.5×0.5 degree resolution for the entire globe. The authors compute the monthly average gridded data from daily weather station records, considering only stations for which no more than five daily data points in a given month are missing. The grid cell data are estimated from measurement station averages through spatial interpolation. Outliers and unrealistic values that might arise due to measurement error are removed by the authors.

3,325 of the grid points are located in the contiguous United States (i.e., excluding Alaska, Hawaii, and the US territories). We aggregate the monthly data to a quarterly frequency by taking the average over the three months in a quarter and seasonally adjust each time series using the deseason() function of the MATLAB Climate Data Toolbox (Greene *et al.*, 2019), which centers and linearly detrends each time series and then removes the climatology, i.e., the average of each given month in a year. In addition, we weight each grid point by the square root of the cosine of the latitude in the center of the cell. This is common practice in the literature that computes empirical orthogonal functions (EOFs) from climate data (Hannachi *et al.*, 2007) and serves as a means to account for the arc of the earth, which changes the size of degree-based grid cells further away from the equator relative to those that are closer to the equator. EOFs are, in essence, the loadings of the principal components computed for gridded climate data, which can be used to detect patterns such as the El Niño Southern Oscillation (ENSO) (Erichson *et al.*, 2020).

We use this method to summarize the information contained in the gridded land surface temperature dataset. To determine the number of principal components, we use the criterion of Alessi et al. (2010), which suggests using between 8 and 17 factors. For parsimony, we set the number of principal components to r=8r=8 and study the effect of choosing r=17r=17 in a robustness exercise. Figure 1 shows that the time series for average US temperature and the first principal component from our dataset are 96% correlated. In addition, Figure 2 shows that the first principal component – which carries the same signal as the average –explains temperature variation only in the Midwest of the US, while important economic centers such as the coastal areas are much less well explained. Expanding the information to r=8r=8 yields much higher explained variation, in the area of 80% almost everywhere in the US. Similar results appear in other large countries of the world, but are not reported here. Therefore, the information in average temperatures covered by a single principal component is clearly insufficient to capture the full temperature dynamics of the US. Any approach using only nationwide averages will likely miss important spatial temperature information.



FIGURE 1 • AVERAGE TEMPERATURES IN THE US AND FIRST PRINCIPAL COMPONENT.

Correlation is 96%





FIGURE 2 • R² FROM REGRESSION OF GRID CELL TEMPERATURES
ON PRINCIPAL COMPONENTS




The economic data we include are real GDP, real investment, nonfarm-business sector hours worked (obtained from FRED), and utilization-adjusted TFP (from Fernald, 2014). All economic variables enter the model in log-levels to account for the possibility of cointegration among economic and climate variables, as pointed out in Pretis (2020). We have checked the model in per-capita terms and found no major difference. A detailed account of all the economic data used in this paper and their construction is given in the Appendix. The sample we use for estimation of the baseline model runs at quarterly frequency from 1948 to 2017.



Figure 3 plots the economic data together with the trend in average US temperatures. Temperatures exhibit an initial decrease until around the 1970s, after which they trend upwards. The series appear to share a common trend as of the 1970s but diverge again after the Great Recession, where the growth rate in temperatures speeds up.

FIGURE 3 • HP-FILTERED TREND IN MEAN CONTIGUOUS US TEMPERATURES $(\lambda=160000) \text{ AND LOGARITHMIZED ECONOMIC TIME SERIES}$ Shaded areas are NBER recessions. All data are centered and scaled To have zero mean and unit variance



3. ECONOMETRIC METHODOLOGY

3.1. Reduced form data representation

Our estimation procedure is carried out in two steps, as in factor-augmented vector autoregressions (FAVAR) (e.g., Bernanke *et al.* (2005)) and dynamic factor models (DFM) (e.g., Forni *et al.* (2009)). These models have the advantage that they can accommodate datasets with many time series and allow for the straightforward identification of structural shocks and their propagation through the methods common in the literature on structural VARs (SVARs) (Ramey, 2016).

The model for the temperatures at grid cell i at time t is given by:

$$T_{it} = \lambda_i Y_t + \eta_{it} \tag{1}$$



where T_{it} are the raw temperatures and η_{it} is the idiosyncratic component. The vector of loadings λ_i captures the sensitivity of temperatures at grid cell i to the aggregate variables in the vector $Y_t = [f_t, y_t]'$. We combine the principal components f_t of the temperature data with the selected set of economic variables Y_t . This is a simple version of the model in Phillips et al. (2020), where we accommodate spatial dependence of temperatures on common factors. The reduced form model for Y_t is a VAR of lag order p:

$$A(L)Y_t = \mu + \epsilon_t, \quad \epsilon_t \sim WN(0, \Sigma)$$
 (2)

where μ is a constant term, A(L) is a matrix polynomial in the lag operator given by $A(L) = A_0 + A_1 L + A_2 L^2 + \cdots + A_p L^p$, and ϵ_t is a vector of reduced form white noise errors whose variance-covariance matrix is given by Σ . Treating the principal components f_t as observed, model (2) is efficiently estimated using OLS for each equation. The lag order is determined using the Akaike information criterion, which yields p = 2. Higher lag orders do not change our results substantially. The reduced form VAR in (2) is assumed to admit a moving average (MA) representation given by:

$$Y_t = C(L)\epsilon_t \tag{3}$$

where C(L) is obtained by inverting A(L), and we have dropped the constant as it is immaterial for our identification strategy and the model dynamics.

3.2. Identification

To identify economic and temperature shocks, we rely on techniques that have been proposed for the study of business cycle fluctuations. Most environmental models focus on aggregate productivity shocks as drivers of emissions (Annicchiarico and Di Dio, 2021). However, the recent contributions by Angeletos *et al.* (2020) and Forni *et al.* (2023) have shown that the economy, and by extension also emissions, fluctuates largely because of sources that are not purely related to movements in TFP. Therefore, our analysis is set up to provide evidence on alternative channels for the effect of socio-economic activity on temperatures, beyond RBC-style technology shocks alone. It is most common to distinguish fluctuations of high frequency, business cycle frequency, and low frequency. Table (1) shows the definitions of frequency bands we adopt for our purposes:

TABLE 1 • FREQUENCY BANDS ADOPTED FOR IDENTIFICATION

Frequency	Low	Business Cycle	High	Full Spectrum
Quarters	> 40	[6, 32]	(0, 6]	$(0,\infty)$

The business cycle frequency is between 6 (1.5 years) and 32 quarters (8 years), as is common in the economic literature (Angeletos *et al.*, 2020). This definition roughly coincides with medium cycles that are observable in climatic data as well. For example,



ENSO (El Niño-Southern Oscillation) influences global weather and occurs every 3-5 years and lasts for roughly a year (NOAA, 2023). The higher frequencies coincide with the strongest fluctuations in our temperature data. This component is most similar to the types of weather shocks usually identified in the literature. The low-frequency band is where we expect the strongest influence of socio-economic activity to show up, as it contains the slight upward trend in the data that is believed to be caused by human beings. Allowing the medium-cycle band to include a few more years (e.g., to include the 11-year solar cycles) does not affect our results.

The structural MA representation of (3) is given by

$$Y_t = C(L)SHu_t = D(L)Hu_t = K(L)u_t, \quad u_t \sim WN(0, I)$$
 (4)

where $SS' = \Sigma$, HH' = I, and $u_t = H'S^{-1}\epsilon_t$. Identification of the structural shocks boils down to pinning down columns of the orthonormal matrix H. The impulse responses of the economic variables (subindex E) and of temperatures (subindex T) are then given by

$$IRF_E = D_E(L)H \tag{5}$$

$$IRF_T = \Lambda D(L)H \tag{6}$$

The notation $D_E(L)$ is shorthand for selecting the rows from each of the matrices in D(L) which correspond to the entries of Y_t that belong to economic variables. Λ is the matrix containing the vectors of loadings λ_i for each grid cell.

3.2.1. Identification of economic shocks

We identify three economic shocks – a technology shock, an investment shock, and a labor supply shock. These are the three shocks proposed as the main business cycle drivers in Justiniano et al. (2010, 2011). To do this, we follow the procedure described in Forni *et al.* (2023), which identifies shocks according to their contribution to the cyclical variances of key variables.

Consider the structural representation of equation (4). The cyclical variance-covariance matrix of all variables in Y_t in the frequency band between $[\underline{\theta}, \overline{\theta}]'$ is given by:

$$V(\underline{\theta}, \overline{\theta}) = \int_{\underline{\theta}}^{\overline{\theta}} D(e^{-i\omega}) D(e^{i\omega})' d\omega$$
 (7)

where, for example, in the case of business cycle frequencies $\left[\underline{\theta}, \overline{\theta}\right]' = \left[2\pi/32, 2\pi/6\right]$ and i is the imaginary constant $i = \sqrt{-1}$. In practice, $V\left(\underline{\theta}, \overline{\theta}\right)$ can be obtained by computing the average over a grid of values between $\underline{\theta}$ and $\overline{\theta}$ and taking the real part of



this average (or computing the inverse Fourier transform of the right-hand side in the above equation). This returns the total variation of all variables in Y_t in the given frequency band as the diagonal elements of the matrix $V(\underline{\theta}, \overline{\theta})$.

To identify a particular shock instead, we use a single column h of the orthonormal matrix H to obtain:

$$\Psi(\underline{\theta}, \overline{\theta}) = \int_{\underline{\theta}}^{\overline{\theta}} D(e^{-i\omega}) h h' D(e^{i\omega})' d\omega \tag{8}$$

which is the variation of all variables in the given frequency band stemming from the shock associated with column h. For our identification strategy, we want to target only specific variables in a given band, so we select the rows of D that correspond to these variables. Suppose, for example, TFP is ordered second in Y_t then D_m for m=2 would select the corresponding row. As shown in Forni et al. (2023), this can easily be extended for multiple targets. This is discussed in more detail for the case of temperature shocks where we make use of this technique. We want to find the shock which contributed the majority of fluctuations in the given band to our target variable, so the column h is identified as:

$$h = \arg \max_{h} \left\{ \int_{\theta}^{\overline{\theta}} D_{m}(e^{-i\omega})' D_{m}(e^{i\omega}) d\omega \right\} \text{ s.t. } h'h = 1$$
 (9)

The h that solves this is the unit-length eigenvector corresponding to the largest eigenvalue of the matrix sandwiched in between h' and h in the above equation (as shown for the time domain in Uhlig, 2003).

We first identify the technology shock as the main driver of low-frequency variation in TFP as in Dieppe *et al.* (2021), which echoes the idea of Gali (1999) to identify technology shocks as the only long-run driver of labor productivity. Maximization does not imply that a single source is responsible for all long-run variation of TFP, but picks out the disturbance that contributes the most to its fluctuations. Dieppe *et al.* (2021) show this method to be more robust to interference from other shocks that typically occurs in variance maximization approaches such as Barsky and Sims (2011). Conditional on the identified technology shock, we then proceed to identifying the investment shock as the main driver of aggregate investment over the business cycle. Justiniano *et al.* (2010, 2011) show that such a shock can be interpreted as a shock to the marginal efficiency of capital, that is, how easily investment is converted to productive capital. The shock typically induces positive comovement between investment and consumption in both representative and heterogeneous agent models (Auclert *et al.*, 2020). The conditional shock is identified by finding another column of **H** call it **h**_j:



$$h_j = \arg \max_{h_j} \left\{ \int_{\theta}^{\overline{\theta}} D_m(e^{-i\omega})' D_m(e^{i\omega}) d\omega \right\} \text{ s.t. } h'_{tech} h_j = \mathbf{0} \text{ and } h'_j h_j = \mathbf{1} \quad (\mathbf{10})$$

Finally, the labor supply shock is identified similarly to the TFP shock as the main driver of hours worked in the low frequency, but conditional on both the technology shock and the investment shock. This identification is inspired by Shapiro and Watson (1988) with an analogy to the relationship between Dieppe *et al.* (2021) and Gali (1999). It is easy to extend the maximization constraints in the above equation to pin down this labor supply shock.

To check whether our approach delivers valid identification, we study it in a controlled experiment using the model of Justiniano *et al.* (2011). The approach correctly recovers the true IRFs to the economic shocks in the majority of cases as reported in the Appendix. Moreover, we check if the sequence of conditional identifications matters for the results in a robustness exercise.

3.2.2. Identification of temperature shocks

We use a similar method as for the economic shocks to identify temperature shocks. Conditional on the three economic drivers, we extract the maximizers of temperature fluctuations in our data set. Economic theory can inform the identification of economic shocks, whereas there is no clear guideline for the identifying traits of climate-related shocks. For example, zero restrictions using a recursive (Cholesky) or long-run neutrality (Blanchard-Quah) scheme seem appropriate, as these would have to hold at every temperature location in our data set, requiring an impossible number of zero responses to be enforced. Maximizing frequency variations of temperatures has the advantage of being statistically driven rather than theoretically and allows us to target many temperature series simultaneously rather than restricting individual variables.

To do this, we need to extend the above framework slightly. Call the IRFs of the temperature variables $\Omega(L) = \Lambda C(L)S$ and collect the columns of H which identify the economic shocks in $H_E = [h_{tech}, h_{inv}, h_{lab}]$. Then the maximization program is the following:

$$h_{Tj} = \arg\max_{h_{Tj}} \left\{ \int_{\underline{\theta}}^{\overline{\theta}} \Omega_m(e^{-i\omega})' W \Omega_m(e^{i\omega}) d\omega \right\} \text{ s.t. } h_{Tj}' H_E = [\mathbf{0}, \mathbf{0}, \mathbf{0}]' \text{ and } h_{Tj}' h_{Tj} = \mathbf{1}(\mathbf{11})$$

As before, h_{Tj} is a single column of H and can be found as the eigenvector of the matrix in the quadratic form in the above equation. W is a diagonal weighting matrix which contains the reciprocals of the square roots of the variances of the m targeted variables in the frequency band of interest. Given that all our data is measured in degrees Celsius, this is less of a concern, but is done for completeness.



We do not require the temperature shocks to be orthogonal to each other, only to the economic shocks, and inspect the resulting IRFs case by case. This is because the main identifying property these shocks have come from geography, which are hardly exclusive. Temperature fluctuations on the US west coast, for example, may be driven by additional impulses elsewhere in the country. Requiring these impulses to be orthogonal appears too restrictive. The targets and bands for identification are chosen as follows:

- I. Maximize the low frequency temperature variation everywhere
- II. Maximize the full spectrum temperature variation everywhere
- III. Maximize the full spectrum temperature variation for the West coast (states that border the Pacific Ocean)
- IV. Maximize the full spectrum temperature variation for the East Coast (states that border the Atlantic Ocean)
- V. Maximize the full spectrum temperature variation for the Gulf of Mexico states (Texas, Louisiana, Mississippi, Alabama, Florida)
- VI. Maximize the full spectrum temperature variation for non-coastal states
- VII. Maximize the business-cycle spectrum temperature variation everywhere to capture the ENSO pattern
- VIII. Maximize the high-frequency temperature variation everywhere to capture the weather shock predominantly used in the literature

The choice is motivated by the geographical patterns we observe in the data, which suggest important temperature commonalities in the Midwest, on the coastal regions, and the Gulf area. Moreover, the maximizer of low frequency temperature movements will likely pick up some non-US socio-economic shocks, and the full-spectrum maximizer is the closest to the temperature shock measured in an approach that uses average temperatures, only in this case, it is purged of US economic activity.

It is important to point out two properties of the shocks that are identified in our FAVAR framework. First, the shocks induce deviations of temperatures at many geographical locations in the US from their deterministic components. If the deterministic component of temperatures contains any trending behavior, a temperature shock constitutes a deviation from this trend. In that sense, explicitly computing the deviation of temperatures from some long-term trend and then using these deviations as a shock, as is done in Kahn et al. (2021), for example, is very similar, but skips the identification step that tries to pinpoint if the deviation comes from human sources or is of natural causes. Second, some climate econometric research stresses the importance of extreme weather events as more suitable measures of temperature shocks (Natoli, 2023). The shocks that we construct are precisely this: they are not predictable from past information about temperatures anywhere in the contiguous US and neither from information about GDP, TFP, investment, or hours worked. Whether this information set is sufficient is a difficult question to answer. Moreover, non-linearities or state-dependence may play an important role in the transmission of such shocks, all of which we consider to be important avenues for future research.



4. RESULTS

4.1. Descriptive results

We begin by summarizing the linkages between the US economy and temperatures through the lens of the model in equations (1) and (2). As a first exercise, we determine the number of shocks that drive US temperatures. In the macroeconometric literature, such shocks are sometimes referred to as deep shocks (Forni *et al.*, 2009). We do this by maximizing the full-spectrum fluctuations of all US temperature series without conditioning on other shocks. Notice that this is done on the spectral density matrix rather than the sample correlation matrix used for the computation of the principal components. We repeat the same exercise and target the full spectrum of variation in the four economic variables to see how these shocks affect temperatures. The outcomes of this are reported in Tables 2 and 3.

TABLE 2 • CUMULATIVE CYCLICAL VARIANCES EXPLAINED BY THE FIRST SIX SHOCKS
THAT MAXIMIZE THE FULL SPECTRUM VARIATION OF TEMPERATURES
AT GRID-CELL LEVEL IN THE US. ROUNDED TO TWO DECIMALS

			Low Fr	equenci	es			Business Cycles					High Frequencies					
	I	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6
Avg. Temp.	0.3	0.48	0.58	0.78	0.83	0.87	0.42	0.63	0.78	0.81	0.87	0.92	0.42	0.65	0.77	0.8	0.86	0.92
GDP	0	0.01	0.01	0.07	0.08	0.12	o	10.0	0.01	0.19	0.22	0.27	0.01	0.02	0.02	0.08	0.09	0.12
TFP	0	0.04	0.06	0.23	0.27	0.28	٥	0.03	0.07	0.52	0.59	0.6	0.01	0.02	0.04	0.29	0.3	0.31
Hours	0	0.01	0.01	0.08	1.0	0.14	0	10.0	0.02	0.25	0.31	0.36	0.02	0.03	0.05	0.18	0.2	0.22
Investment	0	0.02	0.02	0.04	0.06	0.08	0.01	10.0	0.02	0.13	0.16	0.2	0	0.01	0.03	0.06	0.07	0.09

TABLE 3 • CUMULATIVE CYCLICAL VARIANCES EXPLAINED BY THE FIRST SIX SHOCKS THAT MAXIMIZE THE FULL SPECTRUM VARIATION OF GDP, TFP, HOURS, AND INVESTMENT IN THE US. ROUNDED TO TWO DECIMALS

		L	ow Free	quencie	S		Business Cycles					High Frequencies						
	I	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6
Avg. Temp.	0.03	0.22	0.31	0.33	0.4	0.47	0.01	0.02	0.05	0.06	0.14	0.24	10.0	0.02	0.03	0.04	0.1	0.2
GDP	0.93	0.96	I	1	1	1	0.71	0.84	0.91	0.99	0.99	0.99	0.78	0.83	0.87	0.93	0.93	0.97
TFP	0.57	0.96	1	1	1	1	0.17	0.83	0.92	0.98	1	1	0.42	0.78	0.84	0.91	0.95	0.97
Hours	0.78	0.98	0.99	1	1	1	0.43	0.96	0.96	0.99	1	1	0.34	0.85	0.88	0.94	0.96	0.98
Investment	0.88	0.94	0.97	1	1	1	0.53	0.75	0.76	0.99	1	1	0.44	0.54	0.55	0.9	0.93	0.99

Two important new findings emerge from these tables. First, the common variation in US temperatures requires at least five shocks to reach more than 80% explained cyclical variance at all frequencies. After the fifth shock, the improvement in explained variance in any of the three bands of interest from adding another shock is below 5%. This number constitutes a lower bound for the actual number of exogenous temperature drivers, as the shocks here are not structurally identified, other than being mutually orthogonal variance



maximizers. Based on this result, reducing the effects of temperatures on economic aggregates to a single variable such as a (weighted) average, as is frequently done in the literature, is implausible.

Second, there is a connection between temperature and economic variation, mostly through TFP. The fourth temperature variance maximizer is responsible for a sizable share of TFP variation at all frequencies, particularly at the medium part of the spectrum. This seems intuitive: the low and medium frequencies are related to the trend in the temperature data and it is commonly believed that anthropological forces have contributed to this trend in the past half century. Since technology is an important ingredient for economic growth, we should expect it to correlate with the lower frequency components of temperatures. Moreover, we observe that, in line with the literature (e.g., Forni et al. (2023)), two shocks appear sufficient to capture a large share of the cyclical variation in key aggregate economic variables. In the low frequency and business cycle bands, hours, investment, and GDP are largely driven by the same shock, yet TFP is not. This echoes the findings of Angeletos et al. (2020) who also demonstrate a disconnection between TFP and business cycle fluctuations of GDP. Interestingly, investment fluctuations of high frequency appear to require more than three shocks to be accurately explained. Finally, we see that the second shock, which especially drives long-run TFP, is responsible for a large increase in the explained variance of average US temperatures.

The descriptive exercise does not allow us to tell apart the respective source of the fluctuation. Is the variation in temperatures due to climatic or economic shocks? What part of GDP variation is truly due to climatic shocks and which part just masquerades as interference from economic shocks? These questions go back to the cyclical nature of the climate-economic system, and we need the structural identification exercise explained in the preceding section for an answer.

4.2. Semi-structural results

4.2.1. Economic shocks

We begin by discussing the effects of the economic shocks on the economic variables. This is done to confirm that our identification procedure is indeed successful in selecting technology, labor supply, and investment related shocks as described in the macroeconomic literature. The impulse response functions for this are reported in Figure 4.






FIGURE 4 • IMPULSE RESPONSE FUNCTIONS FOR THE THREE STRUCTURAL ECONOMIC SHOCKS Shaded areas are bootstrapped 68% and 90% confidence bands

First, the technology shock leads to an immediate increase in TFP which is accompanied by an expansion of real GDP of around 0.4%. Hours initially decline (although this is statistically insignificant) and investment increases. These results are very similar to those found in Dieppe *et al.* (2021), who use labor productivity in a spectral identification exercise with a different VAR specification.

Second, the labor supply shock leads to a slowly-building increase in output of around 0.3%, a mildly hump-shaped response of hours after an initial increase, and an initial reduction in investment which is replaced by labor as an input to production. The TFP response is almost entirely insignificant, which is partially a result of conditioning on the technology shock. The slow-building GDP response is consistent with other studies that identify labor supply shocks such as Foroni *et al.* (2018) (for the US) and Peersman and Straub (2009) (for the euro area). The responses of hours and GDP are in line with the paper of Shapiro and Watson (1988), which we have used as motivation for the identification strategy.

Lastly, the investment shock creates hump-shaped expansions in investment, hours, and GDP and a hump-shaped decline in TFP. These responses are in line with the motivating paper of Justiniano *et al.* (2011). The decrease in TFP is also observed in Ben Zeev and Khan (2015) (although in their paper the response is insignificantly different from zero) for




investment-specific technology shocks. More inputs are used to produce only slightly more output, thus productivity must fall. We take these results as evidence that our proposed identification strategy can indeed correctly pick out empirically valid impulse responses in a joint identification framework, even though the identification approach is entirely built on spectral identification and does not exactly copy the approaches in the originally proposed papers.

Next, we describe the responses of US temperatures to the three expansionary economic shocks, a key result of this paper. It is important to note that the impact reactions (near impulse response horizon h = 0) of temperatures across the US to the shocks are difficult to measure accurately due to the high volatility of the temperature time series as opposed to the macroeconomic aggregates. We therefore prefer not to interpret temperature responses to economic shocks near the impact. The graphs in Figure 5 show the following picture: the technology shock has a cooling effect on temperatures in the east and the south of the US. Importantly, as the impulse horizon increases, the effect dissipates almost everywhere, which suggests that eventually, cooling and warming offset each other. The effect is persistently significant at the 68% confidence level even after 10 years. The investment shock leads to increases in temperatures almost throughout the US after 10 years, initially dominating in California, Arizona, near the Canadian border, and in the east. Finally, a similar pattern emerges for the labor supply shock, although the initial temperature responses are less pronounced compared to the investment and technology shocks. As far as the magnitudes of the responses are concerned, they range between -0.03and 0.01°C (technology shock), -0.01 and 0.02°C (labor supply shock), and -0.01 and 0.02°C (investment shock).²

² These values are computed across all horizons and grid cells as a single standard deviation around the mean response for each of the three shocks.



FIGURE 5 • GRID CELL TEMPERATURE IRFS AT GIVEN HORIZONS IN RESPONSE TO THE THREE ECONOMIC SHOCKS





Next, in Table 4 we report the relative importance of each of the three economic shocks in explaining average temperature movements, as well as the fluctuations of our economic variables at low, business cycle, and high frequencies.

TABLE 4 • INDIVIDUAL CYCLICAL VARIANCES EXPLAINED BY THE THREE IDENTIFIED ECONOMIC SHOCKS OVER THE THREE FREQUENCY BANDS

Numbers in parentheses are the 90% confidence bands associated with the percentage above. Rounded to two decimals

	I	.ow Frequenci	ies		Business Cycle	es	H	ligh Frequenci	es
	Tech.	Lab. Sup.	Invest.	Tech.	Lab. Sup.	Invest.	Tech.	Lab. Sup.	Invest.
Avg. Temp.	1.0	0.04	0.11	0.01	0.01	10.0	10.0	10.0	o
	(0.04,0.36)	(0.04,0.19)	(0.05, 0.22)	(0.01,0.06)	(0.01,0.06)	(0.02,0.07)	(0.01,0.05)	(0.01,0.05)	(0,0.03)
GDP	0.46	0.13	0.34	0.24	0.05	0.58	0.46	0.04	0.36
	(0.04, 0.93)	(0.01,0.47)	(0.02,0.58)	(0.04,0.75)	(0.01,0.21)	(0.13,0.73)	(0.11,0.79)	(0.01,0.26)	(0.06,0.54)
TFP	0.84	0.02	0.08	0.57	0.02	0.28	0.7	0.07	0.09
	(0.62,0.99)	(0,0.16)	(0,0.15)	(0.14, 0.83)	(0.01,0.27)	(0.04,0.46)	(0.19,0.72)	(0.01,0.29)	(0.02,0.21)
Hours	0.32	0.24	0.44	0.04	0.14	0.79	0.13	0.24	0.54
	(0.06, 0.83)	(0.07,0.5)	(0.04,0.65)	(0.02,0.53)	(0.03,0.21)	(0.26, 0.82)	(0.07, 0.59)	(0.04, 0.29)	(0.14,0.56)
Investment	0.52	0.08	0.39	1.0	0.03	0.85	0.21	0.11	0.56
	(0.11,0.89)	(0.01,0.33)	(0.06,0.64)	(0.02,0.51)	(0.01,0.07)	(0.42,0.92)	(0.06,0.42)	(0.01,0.16)	(0.26,0.66)


Taken together, the three economic shocks explain around 25% of the low-frequency movement of temperatures. Technology and investment shocks contribute the most (10% and 11%, respectively), while labor supply shocks contribute less (4%). We conclude from this that a non-negligible share of the trend- and long-cycle component of temperatures is caused by anthropological activity in the United States. The economic shocks are not important sources of average short-term temperature fluctuations, which we interpret as evidence that such fluctuations are mostly due to natural or non-US causes. The three shocks also appear to be reasonable choices to explain business cycle fluctuations in the economy. Together, they account for 87% of the business cycle (BC) variation in GDP, 87% of the variation in TFP, 97% of the variation in hours, and 89% of the variation in investment.

The spatial distribution of explained variances for the three shocks is presented in Figure 6. Given that there is hardly any variance arising at medium and short frequencies, we report this only for the low frequency. Patches of relevant fluctuations are observable in all three cases. For the technology shock, the variances explained are around 35% in the east and in the south, particularly in Texas. For the investment and labor supply shocks, the patterns emerge predominantly in the south and in the corridor across Colorado, Wyoming, and Idaho, where the labor supply shock was cooling. Explained variances for the investment shock are locally larger than 40% in some areas in the south, while they are lower in the case of the labor supply shock.




FIGURE 6 • GRID CELL LEVEL CYCLICAL VARIATION EXPLAINED AT LOW FREQUENCIES FROM THE THREE ECONOMIC SHOCKS


Technology Shock on LF Temperature Variation



Labor Supply Shock on LF Temperature Variation



Investment Shock on LF Temperature Variation






4.2.2. Temperature shocks

Next, we turn to the effects of the temperature shocks that are identified as described in section 3. For ease of interpretation we have normalized all shocks such that the impact response in average temperatures is scaled to 1 degree Celsius, as is customary. We are primarily concerned with the effect of temperature changes on GDP as all other economic variables were used for identification purposes. Figure 7 summarizes the resulting IRFs.

FIGURE 7 • IMPULSE RESPONSE FUNCTIONS FOR THE DIFFERENT TEMPERATURE SHOCKS.

Shaded areas are bootstrapped 68% and 90% confidence bands



All of the identified shocks lead to small and persistent GDP contractions between 0.1% and 0.2%, except for the shock that primarily affects the West Coast of the US. The confidence bands are consistently very close to the zero line. This result aligns with the majority of the literature, which finds substantial uncertainty in the estimates of temperature shocks in the US. For example, Newell *et al.* (2021) and Nath *et al.* (2023) find nearly zero effects for countries with an average temperature around 13 degrees Celsius, such as the



US. Negative effects of temperature shocks in the range of 0.1% are also reported by Natoli (2023) (although using an instrumental variable approach), and slightly more negative impacts are documented by Colacito et al. (2019). Dell et al. (2012) found insignificant effects of temperatures on output in rich countries. These results are consistent with the shock in our set that maximizes temperature variation across the entire US over all frequency bands. However, our analysis goes beyond this conclusion, revealing that more than one shock is needed to capture US temperature variation. In fact, without imposing orthogonality for this exercise, the West Coast shock is only 2% correlated with the low-frequency maximizer, 3% with the full spectrum maximizer, and a relatively low 33% with the East Coast shock. Interestingly, it produces a comparatively sizable expansion in aggregate GDP (although this is statistically insignificant). This effect would either be lost entirely or blended into average results obtained through conventional econometric techniques. As Table 4 suggests, the share of variation in the economic variables from temperature movements is very small, which is why we choose not to report them here.

For illustration of the spatial distribution of impulse responses, we focus on the full spectrum maximizer for temperatures everywhere and the West Coast shock. These two shocks are only 3% correlated, without the imposition of orthogonality. Figure 8 shows the signs of the responses. Clearly, the full spectrum maximizer without geographical constraints raises temperatures everywhere except for the West Coast. The shock that drives temperatures up on the West Coast simultaneously decreases them in the East. Due to the scaling of the average temperature to equal 1°C, the positive responses outweigh the negative ones. Both of these shocks are quantitatively important for temperature variations (38% and 16% on average, respectively, over all frequency bands). Importantly, we find no evidence of significant persistence in either of the temperature shocks considered here. After around three years, all effects on temperatures turn insignificant.



FIGURE 8 • GRID CELL TEMPERATURE IRFS AT GIVEN HORIZONS IN RESPONSE TO THE FULL SPECTRUM AND THE WEST COAST TEMPERATURE SHOCKS

Temperature IRFs to Full Spectrum Shock at Horizon: 0

Negative Insignificant.

Positive


Negative Insignificant Positive

Temperature IRFs to West Coast Shock at Horizon: 0


a) Full spectrum shock on impact

Temperature IRFs to Full Spectrum Shock at Horizon: 4

b) West coast shock on impact



c) Full spectrum shock after 1 year



d) West coast shock after 1 year



e) Full spectrum shock after 2 years



f) West coast shock after 2 years



g) Full spectrum shock after 3 years

h) West coast shock after 3 years



To summarize the semi-structural results, we observe that economic sources, particularly technology and investment shocks, are locally important drivers of temperature variations. These shocks lead to noticeable decreases (technology) and increases (investment, labor supply) in temperatures that persist for many years and are noticeable even relatively shortly after the initial shock. Treating temperatures as unaffected by anthropogenic forces even in the short run can thus lead to confounding causal effects, especially when annual data is used, as is customary in the literature. Moreover, it is crucial to distinguish the effects of temperature shocks on aggregate GDP based on the geographical location of the shock's epicenter. If the West Coast is predominantly affected, GDP may remain unaffected or even increase, while shocks in other parts of the country can lead to small contractions. This distinction is important for assessing the damages of temperature warming, which are incorporated into models used for policy decisions.

5. DISCUSSION

5.1. The effects of economic shocks on temperatures

The documented effects of the three economic shocks on temperatures across the US warrant closer inspection. The connection between economic activity and temperatures operates through the emission and storage of climate-active gases. Magnus *et al.* (2011) decompose the temperature effect of anthropogenic gas emissions into warming – through the emission of greenhouse gases (GHGs), most prominently CO₂ – and cooling – through aerosol emissions, most prominently SO₂. CO₂ is a long-lived, well-mixing gas that spreads through the Earth's atmosphere over time, while SO₂ produces quick but more short-lived localized cooling by reflecting incoming solar radiation.

There is increasing evidence from the natural sciences literature suggesting that emission impulses can lead to temperature effects within a short time span. Notably, Ricke and Caldeira (2014) and Zickfeld and Herrington (2015) suggest that CO₂ emission impulses can lead to significant warming relatively quickly – 93% of the peak warming effects materialize after 10-15 years following an emission impulse in their experiments, even considering potential non-linearities. Such horizons are well within the customary projection range for FAVAR models. Complementary to this, Joos *et al.* (2013) calculate average surface-temperature responses to CO2 emission impulses and find positive reactions contemporary with the initial impulse. Methane is another powerful GHG that develops much of its effects over a short horizon (Mar *et al.*, 2022). Therefore, our finding of quick temperature changes in the US after economic shocks aligns with results found in climatology research. Nevertheless, we want to emphasize that the very long run, where GHG effects are still active, may be less precisely estimated in our model.

Technology shocks induce cooling in parts of the US east and south. This suggests that the solar radiation effect from aerosol emissions outweighs the heating effect from GHG emissions at these locations, especially in the short run. We investigate this hypothesis further by running the following analysis: to the VAR consisting of GDP, TFP, investment,



and hours worked, we add time series for GHGs and SO₂ emissions in the US for the same sample we used in our previous analysis. The emissions data are available at a yearly frequency. The data for GHGs are retrieved from https://ourworldindata.org/greenhouse-gas-emissions and are based on Jones *et al.* (2023); the data for SO₂ are from Smith et al. (2011) until 1990 and from then on from the EPA (https://www.epa.gov/air-emissions-inventories/air-pollutant-emissions-trends-data). We estimate the VAR with a single lag and identify a technology shock and an investment shock in exactly the same fashion as before, using frequency domain techniques.

FIGURE 9 • IMPULSE RESPONSE FUNCTIONS OF LOG EMISSIONS TO TECHNOLOGY AND INVESTMENT SHOCKS IN THE YEARLY VAR (1) FOR ONLY ECONOMIC VARIABLES.

Identification in the frequency domain adapted to yearly measurements





Figure 9 shows the responses of SO₂ and GHG emissions to the two main expansionary shocks (technology and investment). The impulse response functions (IRFs) for the other economic variables are consistent with the quarterly exercise and are therefore not reported again. The permanent shock to TFP also leads to permanent increases in both SO₂ and GHG emissions. However, the increase in SO₂ emissions is about 2% initially and up to 6% after 15 years, while GHG emissions increase only between 0.5% on impact and slightly below 1.5% in the long run. We interpret this as evidence that what we observe in the quarterly FAVAR is cooling from increased aerosol emissions. This observation is consistent with the localized effects in the south-east of the country, which we discuss




further below. Importantly, as noted in Magnus *et al.* (2011), SO₂ is itself short-lived. Despite the sustained increase in SO₂ emissions, the cumulative warming effect from GHGs eventually neutralizes the cooling from aerosols in our quarterly FAVAR, which is why, as the IRF horizon increases, the cooling effects disappear or even turn into warming. For the investment shock, on the other hand, we observe impulses in both SO₂ and GHGs of equal magnitude. However, the SO₂ impulse is only mildly significant for about one year before emissions (insignificantly) decrease. GHG emissions increase strongly and persist for a longer period, leading to the rapid dissipation of the cooling effect and dominance of the warming effect from GHGs throughout the horizon in the quarterly FAVAR. This explains why temperature changes after the investment shock are observed across almost the entire country and remain significant even after 15 years – there is no sustained counteracting cooling effect.

FIGURE 10 • SO₂ AND CO₂ EMISSIONS ARE COMPUTED FROM EPA'S NEI 2020 DATA SET FOR SITE-SPECIFIC EMISSIONS (https://www.epa.gov/air-emissions-inventories/2020-national-emissions-inventory-nei-data). These include emissions from fossil fuel combustion, industrial processes and biomass (e.g. Wildfires), but exclude *onroad* emissions



a) Sulfur dioxide emissions 2020



b) Carbon dioxide emissions 2020



Curiously, the geographical pattern of temperature changes following a technology shock, as shown in Figure 6, roughly coincides with the locations of important parts of the American energy-producing, manufacturing, and natural resource processing industries. Figure 10 demonstrates that these areas are also centers of CO₂ and SO₂ emissions. Conley et al. (2018) study the responses of temperatures to the hypothetical removal of all US-based SO₂ emissions and document a very similar geographical pattern (with inverted signs, as they consider SO₂ removal rather than emission). Based on this observation, we are confident that our economic shocks lead to temperature-altering emissions in the expected parts of the country. Furthermore, given the localized nature of aerosol-related cooling, we interpret this spatial pattern as evidence that the channel we identify for our technology shock is indeed dominated by SO₂ emissions.

5.2. The effects of temperature shocks on GDP

Next, we turn to the discussion of the different effects of west coast-centered temperature shocks and the other temperature shocks we have identified. We focus on the full spectrum maximizer as a representative of the other shocks and recall that both shocks lead to a one centigrade increase in average US temperatures, but the GDP responses present opposite signs. Our reasoning for this finding is based on previous results in the literature.

First, consider sector-level responses. Increases in temperatures have been shown to reduce output in almost every industry, especially in agriculture and construction (Colacito *et al.*, 2019). The temperature increase that follows the full spectrum shock affects almost the entire US and thus essentially all industries (a notable exception being California), thus depressing aggregate GDP. Conversely, the west coast shock leads to increased temperatures on the west coast but is accompanied in the data by lower temperatures in the east. In our linear model, decreasing temperatures should be beneficial for output in those states. The heating in the west does not appear to offset this positive effect.

Second, we turn to geographical differences. Hsiang *et al.* (2017) provide estimates of the projected spatial distribution of climate effects for the US. They calculate a gain in agriculture from increased temperatures in the north-west of the country and project overall total damages to concentrate in the south-east of the country, whereas the north-western states experience positive effects from warming. The largest damages from temperature increases go through excess mortality in the densely populated east and the already warmer south of the US in their study, also reported by Carleton *et al.* (2022). Therefore, the warming in the west and cooling in the east we document after the west coast shock should benefit the western industries and lead to fewer deaths in the east, which sums to a net positive effect for aggregate GDP. The full spectrum shock, on the other hand, does not produce the warming gains in the north-west but leads to warming in the areas where excess mortality has been shown to be of high importance in the transmission of temperatures to GDP.



In light of these arguments, we carry out the following exercise to better understand how the shocks impact state-level income. We expect the full spectrum shock to be damaging almost everywhere and the west coast shock to be expansionary, at least in the eastern states, but potentially also in the west. To do this, we run the following local projections (Jordà, 2005) for each state in the contiguous US individually:

$$y_{t+h} = \mu_h + \beta_h \hat{s_t} + \gamma_h(L) y_{t-1} + \epsilon_{t+h}, \quad \text{for h} = 1, 2, ..., 40$$
 (12)

where y_{t+h} is the log of quarterly real personal income,³ μ_h is a constant, $\widehat{s_t}$ is alternatively the unit variance full spectrum or west coast shock estimated in the FAVAR, $\gamma_h(L)$ is a lag-polynomial of order two as in the FAVAR, and ϵ_{t+h} is a forecast error. The coefficient β_h measures the response to the shock of interest at each horizon h.


Figure 11 shows that the full spectrum temperature shock indeed decreases income in nearly all states, except for New York, which nonetheless experiences reductions in income for most of the horizon. The west coast shock, on the other hand, produces mixed impulse response functions (IRFs). The majority of economically large states (by share of national GDP) experience income increases, except for Colorado, Florida, and Texas, where the losses are relatively small. Big west coast economies such as California and Washington see long-run benefits from the shock, although these are small in magnitude. We take the evidence from this auxiliary model as supportive of the idea that temperature increases, in general, are detrimental for output, possibly by increasing mortality or lowering productivity. However, we caution that a measured increase of average US temperatures of one degree Celsius can come in different shapes, which produce different dynamics at the state level and then translate into different aggregate responses. We believe that our two example shocks are good representations of actual co-movement in temperatures experienced in the US. Any exercise focusing on the simple average temperature, which is similar to the full spectrum maximizer, will likely miss the effects induced by the west coast shock and may lead to incomplete conclusions for damage functions and policy implications.


³ Personal income data at the state level at quarterly frequency is collected from BEA table SQINC4 and deflated using the GDP deflator and alternatively the CPI. The sample spans Q1:1948 - Q4:2017.



FIGURE 11 • IMPULSE RESPONSES TO THE FULL SPECTRUM AND THE WEST COAST TEMPERATURE SHOCKS IDENTIFIED IN THE FAVAR

IRFs are obtained by means of a local projection of real personal income at the state level onto its own lags and the identified unit variance shock. The states with name tags are the largest 15 states by GDP. Blue lines indicate negative responses after 40 quarters. Red lines indicate positive responses after 40 quarters







6. CONCLUSION

We model an empirical joint climate-economic system to investigate the effect of economic shocks on temperatures in the US and vice versa. Using the principal components of a large, gridded dataset of US temperatures, we show that at least five shocks are necessary to accurately reflect temperature variations of different frequencies everywhere in the contiguous US, calling into question papers that use a single "climate shock" or focus on cross-sectional averages to reflect temperature warming. We show that a clear connection between the economy and temperatures exists, which is mostly driven by changes in Total Factor Productivity (TFP). We identify three economic shocks, arguably responsible for the bulk of business-cycle and long-term variation in the US economy and thus emissions of climate-active gases – a technology shock, a labor supply shock, and an investment shock. Identification in the frequency domain allows us to mix medium-term and long-term identification assumptions. There is clear evidence that economic activity has affected US temperatures. Together, the three shocks account for around 25% of the low-frequency component of US temperatures. Investment shocks increase temperatures on average, technology shocks decrease them, and we explore the reasons for this by showing a significant role for aerosol emissions that induce local, short-lived cooling and greenhouse gas (GHG) emissions that lead to slow-paced, encompassing warming.

On the other hand, the economic damages from changing temperatures are small and come with substantial uncertainty. We show that temperature changes that affect primarily the US west coast lead to small economic expansions, as they are accompanied by decreasing temperatures in the east and south. Shocks raising temperatures elsewhere are mildly recessionary, suggesting that the US has been well-adapted to temperature change in the past.



REFERENCES

- Acevedo, S., Mrkaic, M., Novta, N., Pugacheva, E., and Topalova, P. (2020), The Effects of Weather Shocks on Economic Activity: What Are the Channels of Impact?, «Journal of Macroeconomics», 65, 103207
- Alessi, L., Barigozzi, M., and Capasso, M. (2010), *Improved Penalization for Determining the Number of Factors in Approximate Factor Models*, «Statistics and Probability Letters», 80, 23-24, pp. 1806-1813
- Angeletos, G.-M., Collard, F., and Dellas, H. (2020), *Business-Cycle Anatomy*, «American Economic Review», 110, 10, pp. 3030-3070
- Annicchiarico, B., Carattini, S., Fischer, C., and Heutel, G. (2021), Business Cycles and Environmental Policy: Literature Review and Policy Implications, Cambridge (MA), National Bureau of Economic Research
- Auclert, A., Rognlie, M., and Straub, L. (2020), *Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated Bank Model*, National Bureau of Economic Research.
- Barreca, A., Clay, K., Deschênes, O., Greenstone, M., and Shapiro, J.S. (2015), Convergence in Adaptation to Climate Change: Evidence from High Temperatures and Mortality, 1900-2004, «American Economic Review», 105, 5, pp. 247-251
- Barsky, R. B., and Sims, E. R. (2011). *News Shocks and Business Cycles*, «Journal of Monetary Economics», 58, 3, pp. 273-289
- Bastien-Olvera, B. A., Granella, F., and Moore, F.C. (2022), Persistent Effect of Temperature on GDP Identified from Lower Frequency Temperature Variability, «Environmental Research Letters», 17, 8
- Bennedsen, M., Hillebrand, E., and Koopman, S.J. (2021), Modeling, Forecasting, and Nowcasting US CO₂ Emissions Using Many Macroeconomic Predictors, «Energy Economics», 96, 105118
- Ben Zeev, N., and Khan, H. (2015), *Investment-Specific News Shocks and US Business Cycles*, «Journal of Money, Credit and Banking», 47, 7, pp. 1443-1464
- Berg, K.A., Curtis, C.C., and Mark, N. (2023), GDP And Temperature: Evidence On Cross-Country Response Heterogeneity, National Bureau of Economic Research
- Bernanke, B.S., Boivin, J., and Eliasz, P. (2005), Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach, «The Quarterly Journal of Economics», 120, 1, pp. 387-422
- Bilal, A., and Känzig, D.R. (2024), *The Macroeconomic Impact of Climate Change: Global vs. Local Temperature*, National Bureau of Economic Research Working Paper Series
- Burke, M., Hsiang, S.M., and Miguel, E. (2015, *Global Non-Linear Effect of Temperature on Economic Production*, «Nature», 527, 7577, pp. 235-239
- Cai, Y., and Lontzek, T.S. (2019), *The Social Cost of Carbon With Economic and Climate Risks*, «Journal of Political Economy», 127, 6, pp. 2684-2734
- Carleton, T., Jina, A., Delgado, M., Greenstone, M., Houser, T., Hsiang, S., et al. (2022), Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits, «The Quarterly Journal of Economics», 137, 4, pp. 2037-2105
- Colacito, R., Hoffmann, B., and Phan, T. (2019), Temperature and Growth: A Panel Analysis of The United States, «Journal of Money, Credit and Banking», 51, 2-3, pp. 313-368



- Conley, A., Westervelt, D., Lamarque, J.-F., Fiore, A., Shindell, D., Correa, G., and Horowitz, L. (2018), *Multimodel Surface Temperature Responses to Removal of US Sulfur Dioxide Emissions*, «Journal of Geophysical Research: Atmospheres», 123, 5, pp. 2773-2796
- de Juan, A., Poncela, P., Rodríguez-Caballero, V., and Ruiz, E. (2022), *Economic Activity and Climate Change*, arXiv preprint arXiv, 2206.03187, https://arxiv.org/abs/2206.03187
- Dell, M., Jones, B.F., and Olken, B.A. (2012), Temperature Shocks And Economic Growth: Evidence From The Last Half Century, «American Economic Journal: Macroeconomics», 4, 3, pp. 66-95
- Deryugina, T., and Hsiang, S.M. (2014), Does the environment still matter? Daily temperature and income in the United States, National Bureau of Economic Research
- 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, pp. 354-385
- DiCecio, R., and Owyang, M. (2010), *Identifying Technology Shocks in the Frequency Domain*, Federal Reserve Bank of St. Louis Working Paper No. 2010-060A
- Dieppe, A., Francis, N., and Kindberg-Hanlon, G. (2021), The Identification of Dominant Macroeconomic Drivers: Coping With Confounding Shocks, ECB Working Paper No. 2021/2534
- Donadelli, M., Jüppner, M., Riedel, M., and Schlag, C. (2017), *Temperature Shocks and Welfare Costs*, «Journal of Economic Dynamics and Control», 82, pp. 331-355
- Erichson, N.B., Zheng, P., Manohar, K., Brunton, S.L., Kutz, J.N., and Aravkin, A.Y. (2020), *Sparse Principal Component Analysis Via Variable Projection*, «SIAM Journal on Applied Mathematics», 80, 2, pp. 977-1002
- Fernald, J. (2014), A Quarterly, Utilization-Adjusted Series on Total Factor Productivity. Federal Reserve Bank of San Francisco Working Paper No. 2012-19
- Forni, M., Gambetti, L., Antonio, G., Sala, L., and Soccorsi, S. (2023), An American macroeconomic picture, Mimeo, https://www.ssrn.com/abstract=3910770
- Forni, M., Gambetti, L., and Sala, L. (2014), No News in Business Cycles, «The Economic Journal», 124, 581, pp. 1168-1191
- Forni, M., Giannone, D., Lippi, M., and Reichlin, L. (2009), Opening the Black Box: Structural Factor Models With Large Cross Sections, «Econometric Theory», 25, 5, pp. 1319-1347
- Foroni, C., Furlanetto, F., and Lepetit, A. (2018). Labor supply factors and economic fluctuations, «International Economic Review», 59, 3, pp. 1491-1510
- Forster, P.M., Forster, H.I., Evans, M.J., Gidden, M.J., Jones, C.D., Keller, C.A., et al. (2020), *Current and Future Global Climate Impacts Resulting from COVID-19*, «Nature Climate Change», 10, 10, pp. 913-919
- Fosten, J. (2019), CO₂ Emissions and Economic Activity: A Short-To-Medium Run Perspective, «Energy Economics», 83, pp. 415-429
- Francis, N., and Ramey, V.A. (2009), Measures Of Per Capita Hours And Their Implications For The Technology-Hours Debate, «Journal of Money, Credit and Banking», 41, 6, pp. 1071-1097
- Gali, J. (1999), Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?, «American Economic Review», 89, 1, pp. 249-271
- Goulet Coulombe, P., and Göbel, M. (2021), On Spurious Causality, CO₂, And Global Temperature, «Econometrics», 9, 33



- Gourio, F., and Fries, C. (2020), Adaptation and the Cost of Rising Temperature for the US Economy, Mimeo, //www.ssrn.com/abstract=3696920
- Greene, C.A., Thirumalai, K., Kearney, K.A., Delgado, J.M., Schwanghart, W., Wolfenbarger, N.S., and Blankenship, D.D. (2019), *The Climate Data Toolbox for MATLAB*, Geochemistry, Geophysics, Geosystems
- Hannachi, A., Jolliffe, I.T., and Stephenson, D.B. (2007), Empirical Orthogonal Functions and Related Techniques in Atmospheric Science: A Review, «International Journal of Climatology», 27, 9, pp. 1119-1152
- Hsiang, S., Kopp, R., Jina, A., Rising, J., Delgado, M., Mohan, S., et al. (2017), Estimating Economic Damage from Climate Change in the United States, Science, 356, 6345, pp. 1362-1369
- IPCC (2023), AR6 Synthesis Report: Climate Change 2023. Intergovernmental Panel on Climate Change, https://www.ipcc.ch/report/sixth-assessment-report-cycle/
- Jones, M.W., Peters, G.P., Gasser, T., Andrew, R.M., Schwingshackl, C., Gütschow, J., and Le Quéré, C. (2023), National Contributions to Climate Change Due to Historical Emissions of Carbon Dioxide, Methane, and Nitrous Oxide Since 1850, «Scientific Data», 10, 1, p. 155
- Joos, F., Roth, R., Fuglestvedt, J.S., Peters, G.P., Enting, I.G., Von Bloh, W., et al. (2013), Carbon Dioxide and Climate Impulse Response Functions for the Computation of Greenhouse Gas Metrics: A Multi-Model Analysis, «Atmospheric Chemistry and Physics», 13, 5, pp. 2793-2825
- Jordà, O. (2005), Estimation and Inference of Impulse Responses by Local Projections, «American Economic Review», 95, 1, pp. 161-182
- Justiniano, A., Primiceri, G.E., and Tambalotti, A. (2010), *Investment Shocks and Business Cycles*, «Journal of Monetary Economics», 57, 2, pp. 132-145
- Justiniano, A., Primiceri, G.E., and Tambalotti, A. (2011), *Investment Shocks and the Relative Price of Investment*, «Review of Economic Dynamics», 14, 1, pp. 102-121
- Kahn, M.E., Mohaddes, K., Ng, R.N., Pesaran, M.H., Raissi, M., and Yang, J.-C. (2021), Long-Term Macroeconomic Effects of Climate Change: A Cross-Country Analysis, «Energy Economics», 104, 105624
- Kaufmann, R.K., Kauppi, H., Mann, M.L., and Stock, J.H. (2013), *Does Temperature Contain A Stochastic Trend: Linking Statistical Results to Physical Mechanisms*, «Climatic Change», 118, pp. 729-743
- Khan, H., Metaxoglou, K., Knittel, C.R., and Papineau, M. (2019), Carbon Emissions and Business Cycles, «Journal of Macroeconomics», 60, pp. 1-19
- Kilian, L. (1998), Small-Sample Confidence Intervals for Impulse Response Functions, «Review of Economics and Statistics», 80, 2, pp. 218-230
- Kunkel, K.E., Stevens, L.E., Stevens, S.E., Sun, L., Janssen, E., Wuebbles, D., et al. (2013), Regional Climate Trends and Scenarios for the US National Climate Assessment: Part 2. Climate Of The Southeast US, NOAA Technical Report, https://www.ncdc.noaa.gov/assessment-report
- Magnus, J.R., Melenberg, B., and Muris, C. (2011), Global Warming and Local Dimming: The Statistical Evidence, «Journal of the American Statistical Association», 106, 494, pp. 452-464



- Mar, K.A., Unger, C., Walderdorff, L., and Butler, T. (2022), Beyond CO₂ Equivalence: The Impacts of Methane on Climate, Ecosystems, and Health, «Environmental Science and Policy», 134, pp. 127-136
- Matsuura, K., and Willmott, C.J. (2018), Terrestrial Air Temperature: 1900-2017 Gridded Monthly Time Series, University of Delaware, Newark, http://climate.geog.udel.edu/~climate/html_pages/download.html#T2017
- Montamat, G., and Stock, J.H. (2020), *Quasi-Experimental Estimates of The Transient Climate Response Using Observational Data*, «Climatic Change», 160, 3, pp. 361-371
- Mumtaz, H., and Marotta, F. (2023), Vulnerability to Climate Change: Evidence from a Dynamic Factor Model, Mimeo, https://www.ssrn.com/abstract=3845743
- Nath, I.B., Ramey, V.A., and Klenow, P.J. (2023), How Much Will Global Warming Cool Global Growth?, Working Paper, https://doi.org/10.2139/ssrn.4358324
- Natoli, F. (2023), The Macroeconomic Effects of Temperature Surprise Shocks, SSRN 4160944. https://doi.org/10.2139/ssrn.4160944
- Newell, R.G., Prest, B.C., and Sexton, S.E. (2021), *The GDP-Temperature Relationship: Implications for Climate Change Damages*, «Journal of Environmental Economics and Management», 108, 102445
- NOAA (2023), What are La Niña and El Niño and why do they matter?, ENSO Fact Sheet, https://www.noaa.gov/enso
- Peersman, G., and Straub, R. (2009), Technology Shocks and Robust Sign Restrictions in a Euro Area SVAR, «International Economic Review», 50, 3, pp. 727-750
- Phillips, P.C., Leirvik, T., and Storelvmo, T. (2020), Econometric Estimates of Earth's Transient Climate Sensitivity, «Journal of Econometrics», 214, 1, pp. 6-32
- Pretis, F. (2020), Econometric Modelling of Climate Systems: The Equivalence of Energy Balance Models and Cointegrated Vector Autoregressions, «Journal of Econometrics», 214, 1, pp. 256-273
- Pretis, F. (2021), Exogeneity in Climate Econometrics, «Energy Economics», 96, 105122
- Ramey, V.A. (2016), Macroeconomic shocks and their propagation, in J.B. Taylor and H. Uhlig (eds), Handbook of Macroeconomics, Elsevier, vol. 2, pp. 71-162
- Ricke, K.L., and Caldeira, K. (2014), Maximum Warming Occurs About One Decade After a Carbon Dioxide Emission, «Environmental Research Letters», 9, 12, 124002
- Shapiro, M.D., and Watson, M.W. (1988), Sources of Business Cycle Fluctuations, in S. Fisher and C.S. Black (eds), NBER Macroeconomics Annual, Cambridge (MA), The Mit Press, vol. 3, pp. 111-148
- Smith, S.J., van Aardenne, J., Klimont, Z., Andres, R.J., Volke, A., and Delgado Arias, S. (2011), *Anthropogenic Sulfur Dioxide Emissions: 1850-2005*, «Atmospheric Chemistry and Physics», 11, 3, pp. 1101-1116
- Stock, J.H. (2020), Climate change, climate policy, and economic growth, in J.B. Taylor and H. Uhlig (eds), NBER Macroeconomics Annual, Cambridge (MA), The Mit Press, vol. 34, pp. 399-419). MIT Press.
- Storelvmo, T., Leirvik, T., Lohmann, U., Phillips, P.C., and Wild, M. (2016), *Disentangling Greenhouse Warming and Aerosol Cooling to Reveal Earth's Climate Sensitivity*, «Nature Geoscience», 9, 4, pp. 286-289



- Uhlig, H. (2003), *What Drives GNP?*, Unpublished manuscript, Euro Area Business Cycle Network, https://www.euro-area-bcn.org/papers/uhlig_gnp.pdf
- Wieland, V., Afanasyeva, E., Kuete, M., and Yoo, J. (2016), New Methods For Macro-Financial Model Comparison and Policy Analysis, in J.B. Taylor and H. Uhlig (eds), Handbook of Macroeconomics, Elsevier, vol. 2, pp. 1241-1319
- Wieland, V., Cwik, T., Müller, G.J., Schmidt, S., and Wolters, M. (2012), A New Comparative Approach to Macroeconomic Modeling and Policy Analysis, «Journal of Economic Behavior and Organization», 83, 3 pp. 523-541
- Zickfeld, K., and Herrington, T. (2015), The Time Lag Between a Carbon Dioxide Emission and Maximum Warming Increases With the Size of the Emission, «Environmental Research Letters», 10, 3, 031001

APPENDIX

A DATA CONSTRUCTION

We follow Angeletos et al. (2020) in constructing the economic variables.

TABLE 5 • ECONOMIC DATA SOURCES AND TRANSFORMATIONS

Data	FRED Mnemonic	Frequency/Transformation
Real gross domestic product per capita	A939RXoQ048SBEA	Q
Real Gross domestic product	GDPC1	Q
Share of GDP: gross private domestic investment	A006RE1Q156NBEA	Q
Share of GDP: personal consumption expenditures: durable goods	DDURRE1Q156NBEA	Q
Nonfarm business sector: average weekly hours	PRS85006023	Q
Employment Level	CE16OV	M2Q (EoP)
Total factor productivity (annualized Q-Q growth rate)	dTFPu (from Fernald)	Q

The variables enter the model as follows:

- 1. **Real GDP**: $\log(\text{GDPC1}) \times 100$
- 2. **Real investment**: log((DDURRE1Q156NBEA + A006RE1Q156NBEA) × GDPC1) × 100
- 3. **Hours**: $\log(PRS85006023 \times CE16OV) \times 100$
- 4. **TFP**: cumsum(dTFPU/400) \times 100
- 5. **Population**: GDPC1/A939RX0Q048SBEA

For checks, the variables real GDP, real investment, and hours can be transformed to per capita units by dividing by the population level as computed above before taking logs.

B BOOTSTRAP PROCEDURE

We compute confidence bands for the IRFs and the cyclical variances using the following bootstrap procedure:

- 1. 1.Use (2) to generate a new vector \mathbf{Y}_t by bootstrapping from the reduced form residuals.
- 2. Use the method of Kilian (1998) to correct the bias of the OLS estimates.
- 3. Use Λ to recompute the common component of temperatures, ΛY_t , and add the original idiosyncratic component, η_{it} , to get a new data set of US temperatures.
- 4. On this new data set, estimate r = 8 principal components, and re-estimate a bootstrap Λ_B .
- 5. Estimate the FAVAR in (2) again with p = 2.



- 6. Identify the shocks sequentially, compute IRFs and the cyclical variances.
- 7. Repeat this 1, 500 times to obtain bootstrap distributions of the IRFs and the cyclical variances.
- 8. Find the quantiles of the bootstrap distributions to get the 68% and 90% intervals.

C ROBUSTNESS CHECKS

To test the sensitivity of our results to the underlying assumptions, we conduct the following robustness checks:

1. Changing the number of temperature factors:

We originally used a statistical criterion to determine the number of factors to be extracted from the gridded temperature dataset, opting for r=8 for parsimony. The upper bound recommended by the criterion was r=17, which we also test. In this case, we set p=1 according to the Bayesian Information Criterion (BIC).

2. More lags:

Our results focus primarily on the low-frequency components of temperatures. To address potential inaccuracies due to a very short lag length, we test an increased lag length. In the baseline specification, p=2; here, we increase this to p=4. Given the frequentist approach to estimation, results become quite erroneous for even larger lag orders.

3. Sub-sample analysis (1970):

The dataset used spans from 1948 to 2017. The trend in temperatures attributed to human influence becomes very pronounced around 1970. Additionally, SO₂ emissions in the US start to decline from the 1970s. We repeat our analysis by excluding the first 22 years from the sample to observe any changes.

4. Potential interference from outside shocks:

Non-US shocks may drive business cycle (BC) and low-frequency (LF) variation in US aggregates, affecting temperatures. Although the US is typically considered to have frontier technology (Nath *et al.*, 2023), shocks from China might spill over and be misidentified as US shocks. Given the challenge of obtaining long quarterly time series for China, we use annual series for CO₂ emissions, which show a significant increase from 2000. We thus cut the sample at Q4:1999 to check for potential external influences from China.

5. Maximizing long-run IRFs instead of variances:

As an alternative to maximizing variances, we explore maximizing the long-run IRF of Total Factor Productivity (TFP) and hours, as suggested in Forni *et al.* (2014). This approach is crucial since the connection between the economy and temperatures largely runs through TFP, making accurate identification of the technology shock essential.



6. Variables in per capita terms:

Long-run economic dynamics may be affected by demographic changes (Francis and Ramey, 2009), which are not explicitly accounted for in our baseline specification. Population changes are a significant source of emission variations according to the Kaya identity. We check whether expressing economic variables (GDP, hours, and investment) in per capita terms alters our results.

Robustness Results:

The results are generally insensitive to changes in lag order, number of factors, or specification of variables in per capita terms. Minor changes are observed for sub-samples and when altering the long-run identification assumption, as detailed in robustness check 5. Figure 12 illustrates the IRFs for average US temperatures in response to economic shocks. The most notable differences occur when changing the sub-samples to post-1970 and pre-2000. In these cases, the technology shock leads to positive temperature responses due to the diminished role of SO₂ emissions and other aerosols in cooling temperatures after 1970. Similarly, excluding the more recent period attributes some cooling to the investment shock, as the reduction in SO₂ emissions has not yet fully materialized. These changes, while interesting, underscore the significance of this additional channel for the transmission of economic activity to temperature changes.

FIGURE 12 • IMPULSE RESPONSE FUNCTIONS OF US AVERAGE TEMPERATURES TO ECONOMIC SHOCKS FOR ROBUSTNESS CHECKS 1-6.







Figure 13, on the other hand, reports the IRFs of real GDP to the different temperature shocks for all robustness checks. We observe that changing the number of temperature principal components or the number of lags has negligible effects on the IRFs compared to our baseline specification. The same goes for taking the variables in per capita terms. Changes in the responses of GDP to the temperature shocks are slightly more pronounced if we use labor productivity instead of TFP or the maximal response identification strategy to obtain the technology shock and then condition the temperature shocks on it. All in all, the baseline specification lies roughly in the middle of the IRFs under the different robustness checks. We leave the robustness check IRFs of the economic variables to the economic shocks in the Appendix since the only minor difference arises when using the response maximization approach over the cyclical variance maximization approach.

FIGURE 13 • IMPULSE RESPONSE FUNCTIONS OF GDP TO TEMPERATURE SHOCKS FOR ROBUSTNESS CHECKS 1-6










FIGURE 14 • IMPULSE RESPONSE FUNCTIONS OF ECONOMIC VARIABLES TO ECONOMIC SHOCKS FOR ROBUSTNESS CHECKS 1-6.

Lastly, we check if the sequence of conditional identifications matters for our results. We therefore permute the identification order of the three economic shocks – technology (T), investment (I) and labor supply (H) – to allow for all possible orderings and report the economic and temperature IRFs.










FIGURE 16 • IMPULSE RESPONSE FUNCTIONS OF AVERAGE TEMPERATURES FOR DIFFERENT ORDERINGS OF IDENTIFICATION

Figures 15 and 16 show that while there are some differences in the responses of the economic variables if the investment shock is identified first, these do not translate to changes in the more important results for temperature changes following the economic expansions.

D SIMULATION EXERCISE

We simulate 1,000 instances of the model proposed by Justiniano *et al.* (2011) using the Macroeconomic Model Data Base in Dynare (Wieland *et al.*, 2012, 2016), adhering to the standard settings without modifications. Each simulation includes data for GDP, Total Factor Productivity (TFP), hours worked, and investment, along with additional series that are not considered for this exercise. For each of the 1,000 simulations, we extract the true Impulse Response Functions (IRFs) for neutral technology shocks, investment shocks, and wage markup shocks (which have a similar interpretation to our labor supply shocks). We then apply our sequential identification strategy to identify these three structural shocks in the frequency domain using a VAR(4) model with the four economic time series of interest.

In the model of Justiniano *et al.* (2011), the neutral technology shock is the sole driver of TFP growth, the wage markup shock is the primary factor influencing low-frequency changes in hours worked, and the investment shock predominantly affects investment variation in the business cycle band. Consequently, our identification approach is theoretically validated for this case.



FIGURE 17 • IMPULSE RESPONSE FUNCTIONS OF ECONOMIC VARIABLES TO ECONOMIC SHOCKS FROM SIMULATED DATA AS PER JUSTINIANO *ET Al.* (2011)




Figure 17 shows the bands resulting from the 1,000 identification exercises on simulated data as well as the theoretically true IRFs. Our VAR-based approach is very successful in capturing the correct dynamics in the vast majority of the simulation runs. This gives us confidence that it may also be useful in a purely applied setting.



COSTANZA TOMASELLI

GREEN OR GREED? UNVEILING THE ENVIRONMENTAL IMPACT OF MARKET CONSOLIDATION ON CARBON EMISSIONS

Abstract. This paper explores the relationship between market concentration and environmental performance, with a particular focus on the aftermath of mergers. Drawing from foundational economic principles, I hypothesize that increased market power, typically associated with reduced output relative to competitive market conditions, could similarly influence a firm's emissions profile, potentially lowering Greenhouse Gas (GHG) emissions. This hypothesis introduces a complex tension between two pivotal policy objectives: the reduction of emissions and the preservation of competitive market structures. Novel empirical findings suggest that mergers exhibit a comparable positive impact on environmental indicators. This insight paves the way for a broader discussion on the dual objectives of companies in merger scenarios—increasing their market power versus achieving environmental efficiency.

Keywords. Market concentration, climate risks, emissions, social welfare.

1. Introduction

Product market competition is of paramount importance for a well-functioning economy. It is a well-studied fact that competitors and new entrants push incumbent companies to set prices that reflect costs, which benefit customers. Firms with higher market power can set high prices, which has negative implications for society welfare, and resource allocation, can decrease the demand for labor and dampens investment in capital, it distorts the distribution of economic rents, and it discourages business dynamics and innovation (De Loecker, Jan Eeckhout, and Unger, 2020). For this reason, functioning of markets and the protection of consumer rights have been a priority for governments in the past decades. Specifically, countries have implemented competition policies aimed at regulating abuse of market power and protecting consumers.

On the other side, in recent years, another government priority has rapidly emerged. The increasing scientific evidence and the heightened frequency of extreme weather events underscore the urgent need for countries to rapidly decarbonize. Climate change poses a significant threat to economic stability, public health, and global ecosystems. The Intergovernmental Panel on Climate Change (IPCC) has highlighted the catastrophic consequences of failing to limit global warming to well below 2 degrees Celsius above pre-



industrial levels IPCC (2018). The economic impacts of climate change are profound, including reduced agricultural productivity, increased health care costs, and more frequent and severe natural disasters, which collectively threaten global economic growth Nordhaus (1991).

To mitigate these risks, governments have announced and implemented various environmental policies aimed at reducing carbon emissions. One of the most prominent initiatives is the European Union Emission Trading Scheme (EU ETS), which sets a cap on the total amount of greenhouse gases that can be emitted by covered entities and allows companies to buy and sell emission allowances Ellerman, Convery, and PERTHUIS (2010). This market-based approach incentivizes companies to innovate and reduce their emissions cost-effectively. Furthermore, studies have shown that carbon pricing mechanisms, such as the EU ETS, are essential tools in the transition to a low-carbon economy, as they internalize the external costs of carbon emissions and encourage investments in cleaner technologies Stern (2007).

This study explores the interplay between market concentration and environmental performance, with a particular emphasis on the aftermath of mergers. Given the scrutiny mergers attract from competition authorities due to potential market power implications and consumer harm, this research investigates a nuanced question: do mergers lead to a reduction in greenhouse gas (GHG) emissions for the consolidated entities compared to their premerger states? Drawing on fundamental economic theories, I hypothesized that increased market power, often resulting from mergers, may lead to reduced production levels. This reduction in output, when applied to the domain of environmental emissions, suggests that a more concentrated market could potentially lower GHG emissions. Alternatively, merged entities have access to better technology or better management, due to economies of scale and/or scope and improve their environmental footprints, without reducing their production levels. This hypothesis introduces a complex dynamic between the objectives of emissions reduction and the maintenance of competitive market structures.

Firstly, I set up a model which focuses on the nuanced dynamics between oligopolistic competition, environmental consciousness among consumers, and the impact of mergers on environmental emissions within a simplified economic model. The model, which encapsulates a scenario with two firms producing differentiated products amidst price competition, suggests a pivotal trade-off between prices and emissions in the presence of consumer environmental awareness. The analysis reveals that post-merger outcomes hinge on the magnitude of production efficiencies realized: significant efficiencies lead to increased output, lower prices, and higher emissions, whereas minimal efficiencies result in higher prices but lower emissions, showcasing the environmental benefits of reduced production. This trade-off underscores the complex relationship between competitive market behaviors and environmental impacts, highlighting how mergers can both exacerbate and mitigate



environmental damage depending on the resultant operational efficiencies. Furthermore, we introduce the potential for mergers to foster green innovation, proposing that beyond mere output adjustments, mergers may incentivize investments in environmentally friendly technologies, thus offering a pathway to reducing emissions. This dual-faceted view illuminates the intricate ways in which market consolidation can influence environmental outcomes, emphasizing the role of consumer preferences, technological innovation, and efficiency gains in shaping the ecological footprint of oligopolistic markets.

This paper employs two empirical strategies to investigate the impact of corporate mergers on Scope 1 absolute emissions. Firstly, a panel event study methodology is utilized to analyze the temporal effects of mergers on emissions, building on the works of Miller (2023) and Chaisemartin and D'Haultfoeuille (2020). This approach examines changes in emissions before and after the merger, controlling for unobserved heterogeneity across firms and over time. Secondly, a quasi-experimental design is adopted to address the endogeneity of merger selection. Inspired by Seru (2014), Bena and Li (2014), and Gugler *et al.* (2003), this approach compares firms that completed mergers (treatment group) with those that announced but subsequently cancelled their mergers (control group). By leveraging the difference-in-differences (DiD) framework, this strategy isolates the causal impact of mergers on emissions.

Overall, both empirical strategies provide evidence that mergers, particularly horizontal ones, in line with the hypothesis of a correlation between increased market power and reduced emissions, lead to a decrease in corporate emissions, highlighting the environmental benefits of market consolidation. This observation suggests that corporate consolidation might have implications for environmental performance, presenting a more complex picture than the traditional view that mergers primarily fulfil economic or financial goals. This analysis contributes to understanding how dynamics of market concentration, as a result of mergers and acquisitions, can impact a firm's environmental footprint. It highlights the importance of distinguishing between the sources of environmental benefits, advocating specifically for technological advancements as a key factor for improved environmental outcomes postmerger, rather than market concentration. Through this nuanced approach, the study adds to the ongoing conversation about the interplay between corporate strategy, market structure, and sustainability, emphasizing a balanced consideration of technological innovations alongside economic objectives.



2. LITERATURE REVIEW

Market concentration, which is also used as a substitute for competition intensity, can be defined by the extent to which market shares are concentrated between a small number of firms (OECD, 2018). Recent decades have seen a drastic change in market structure and concentration. The latest publications have noted a trend for increased industry concentration in the United States (Furman and Orszag, 2015; Autor *et al.*, 2020). On the contrary, the more current literature has not reached a consensus on the direction of concentration for Europe; Gutierrez and Philippon (2023) found that competition in Europe increased, due to independent regulators and appropriate competition policies, while Koltay, Lorincz, and Valletti (2023) observe a moderate increase in European industry concentration and a trend towards oligopolies.

On mergers, there is existing literature on the importance of mergers for market concentration and industry links (Ahern and Harford, 2014). Part of the literature on mergers focuses on their negative impact on consumer choice and whether merger threshold is appropriate, Nocke and Whinston (2022) find that current concentration levels are likely too permissive and could contribute to increase in prices which might harm consumers. If not screened properly mergers could also have other negative impact, i.e. on their own workforce, Berger *et al.* (2023) suggest that suggest that workers are harmed, on average, under the enforcement of the more lenient 2010 merger guidelines. Another side of the literature focuses on which companies merge, Crouzet and Eberly (2019) have found that companies with a high level of intangibles over their total assets, such as intellectual property and software, tend to have higher market power and increase market concentration over time, and whether mergers could be a positive incentive for innovation (Phillips and Zhdanov, 2013).

Mergers and acquisitions (M&A) are complex processes often fraught with various challenges that can lead to their failure even after being publicly announced. Several economic studies have explored the multifaceted reasons behind such outcomes. One significant factor is regulatory intervention. Regulatory bodies like antitrust authorities often scrutinize proposed mergers to ensure they do not create monopolistic entities that could harm consumers. For instance, Eckbo (1983) discuss how horizontal mergers are particularly prone to regulatory challenges due to potential anti-competitive effects. The study highlights that about 30% of proposed mergers fail due to regulatory rejections. Another critical reason is financing issues. Kaplan and Stromberg (2009) note that mergers often rely on significant amounts of debt financing. Adverse changes in credit markets or a re-evaluation of the target company's value can lead to financing shortfalls, causing the merger to collapse. The volatility of financial markets thus plays a crucial role in the



completion of M&A deals. Cultural clashes between merging entities also contribute to the failure of mergers. Weber, Shenkar, and Raveh (1996) emphasize that differences in corporate culture can lead to integration problems, resulting in operational inefficiencies and employee dissatisfaction. These cultural mismatches can become apparent during the due diligence process, leading to a reconsideration of the merger. Furthermore, changes in economic conditions can alter the strategic rationale for a merger. Shleifer and Vishny (2003) explain that stock market fluctuations can affect the perceived benefits of a merger. If the market conditions change significantly after the announcement, the acquiring company might find the merger less attractive, leading to its termination. In some cases, the due diligence process uncovers unforeseen liabilities or operational challenges. Krishnan, Hitt, and Park (2005) discuss how the discovery of such issues can cause acquiring firms to reassess the viability of the merger, often resulting in cancellation to avoid future financial burdens. Finally, shareholder opposition can also derail mergers. Shareholders of either the acquiring or target company may believe that the merger does not align with their financial interests. According to Mulherin and Boone (2000), active resistance from major shareholders can lead to the abandonment of the deal.

With respect to the literature on corporate emissions, several key drivers have been identified that influence firms' greenhouse gas outputs. One significant factor is the size and scale of the firm, as larger firms tend to have higher absolute emissions due to greater production volumes and energy consumption Cole and Elliott (2006). Additionally, industry-specific characteristics play a crucial role; sectors such as manufacturing and energy are typically more emission-intensive compared to service-oriented industries Duflo, Greenstone, and Hanna (2008). Regulatory environments and environmental policies are also critical drivers, as stricter regulations and effective enforcement can lead to significant reductions in emissions Kumar and Managi (2012). Firms' technological capabilities and innovation activities are another important determinant, with companies investing in green technologies often achieving lower emission levels Porter and Linde (1995). Furthermore, market pressures and consumer demand for sustainable practices can incentivize firms to adopt greener practices, thus reducing their carbon footprint Delmas and Montes-Sancho (2011). Finally, financial performance and access to capital markets can influence a firm's ability to invest in emission reduction technologies and practices, as better-performing firms are more likely to allocate resources towards sustainability initiatives Eccles, Ioannou, and Serafeim (2014).

This paper contributes to the more recent literature of the impact of non-market effects of market power and market concentration. With respect to the impact on policies Kang and Xiao (2023) find that a company's actions can significantly reduce government procompetitive policies, while Yue (2023) demonstrates how nascent industry can if organized



can nullify local regulations. Other articles have focussed on the impact of market concentration, specifically media, on elections and voters' availability of information (Martin and McCrain, 2023). Finally, on competition and the environment, Aghion *et al.* (2023) found that when consumers care about their environmental footprint, firms pursue greener products. This paper would extend the existing literature on market power and concentration to environmental considerations. My findings will shed light on how the two fields are linked, and whether policymakers need to be aware of such trade-offs when constructing policies in each field.

3. THEORETICAL FRAMEWORK

In this study, I propose a simplified model of oligopolistic competition where consumers are environmentally conscious. The model features two firms producing differentiated products, with production processes that result in emissions. A representative consumer purchases both goods, and emissions are considered harmful, leading to a scenario where, all else being equal, the consumer's demand for the two goods increases.

I explore the impact of a merger between these two firms on prices and emissions. It is posited that a merger could yield specific efficiencies from the combined production of the two goods. My analysis demonstrates that if these efficiencies are sufficiently large, the merger could lead to increased output, reduced prices, and heightened emissions. Conversely, in scenarios where the efficiencies are minimal, the merger leads to higher prices but benefits the environment through a reduction in emissions. Thus, our findings underscore a trade-off between prices and emissions in markets characterized by polluting production processes. My model is intentionally streamlined to underscore this trade-off and to articulate our underlying logic. I make certain assumptions regarding consumer preferences and the number of firms in the market. Nonetheless, these assumptions are not fundamental. The crucial assumptions are twofold: first, that demand decreases as prices increase, and second, that emissions escalate with increased output. Given these conditions, any model of competition would reveal a similar trade-off between pricing strategies and environmental preservation.

Toward the end of this section, I introduce the possibility of an alternative mechanism. Specifically, we argue that a merger could lead to a reduction in emissions not solely by diminishing output due to enhanced market power but also by fostering innovations in green technology.



Preferences and Technology There are two products $i \in \{1, 2\}$, and two firms. Each firm produces a different product. A representative consumer buys the two goods. The consumer has a Singh and Vives (1984) utility function:

$$u(q_1 + q_2) = q_1 + q_2 - \frac{1}{2}(q_1^2 + q_2^2) - \gamma q_1 q_2 - \emptyset z(q_1 + q_2)$$

where q_i is the quantity of product i, and the parameter $\gamma \in (0, 1)$ captures the degree of product differentiation. When $\gamma = 0$, products are completely unrelated, and firms act as local monopolists. When $\gamma = 1$, products are perfect substitutes, and Bertrand competition brings profits down to zero. We rule out both cases.

The function z ($q_1 + q_2$) describes the technology according to which total output ($q_1 + q_2$) generates emissions. We assume the following functional form:

$$z(q_1 + q_2) = (q_1 + q_2)^{\alpha}$$

When $\alpha > 1$ ($\alpha \le 1$), emissions are a convex (concave) function of output. For what follows, we assume a linear form: z ($q_1 + q_2$) = $q_1 + q_2$. The parameter $\varphi \ge 0$ captures the degree of environmental concern for the consumers. When $\varphi = 0$, the consumer does not care about emissions, for example, because the cost of pollution is sustained by people located in different locations or by future generations. Then, the utility function can be rewritten as:

$$u(q_1, q_2) = (1 - \emptyset)(q_1 + q_2) - \frac{1}{2}(q_1^2 + q_2^2) - \gamma q_1 q_2$$

The consumer's utility maximization problem results in the following demand functions:

$$q_i(p_i, p_j) = \frac{1 - \emptyset - p_i + \gamma(p_j + \emptyset - 1)}{1 - \gamma^2}$$

As expected, q_i (p_i , p_j) is increasing in p_i as goods are substitutes and decreasing in p_i as goods are normal. Interestingly, demand is also decreasing in φ . When the degree of environmental concern increases, the consumer reduces their consumption to reduce emissions. We assume that the two firms are equally efficient. Their marginal cost is $c \ge 0$. Profits can then be written as follows:

$$\pi_i(q_i, p_i, p_j) = (p_i - c)q_i = \frac{(p_i - c)(\gamma(p_j + \emptyset - 1) + 1 - \emptyset + p_j)}{1 - \gamma^2}$$

¹ Our results are qualitatively robust to changes in the parameter α . In particular, the quadratic case ($\alpha = 2$) is substantially equivalent to the linear case $\alpha = 1$. We stick to linearity for the sake of simplicity.



I now solve the game for two different states of the world $m \in \{0, 1\}$. If the state is m = 0, the two firms do not merge. If the state is m = 1, the two firms merge. Then, we will perform a welfare assessment of the merger.

Market Equilibrium Let us start from m = 0. Firms do not merge. Then, they set prices simultaneously and independently. The FOC for each firm implies:

$$p_i^*(p_j) = \frac{1}{2}(c + \gamma(\pi_i(p_j + \emptyset - 1) + 1 - \emptyset)$$

Intersecting the best responses, we obtain Nash Equilibrium (equilibrium henceforth) prices:

$$p_i^* = \frac{\gamma \emptyset + c - \gamma + 1 - \emptyset}{2 - \gamma}$$

Total emissions are:

$$z(q_1^* + q_2^*) = \frac{2(1 - c - \emptyset)}{(2 - \gamma)(\gamma + 1)}$$

Let us now turn to the case of m = 1. After a merger, firms set prices cooperatively. In particular, the merged entity chooses prices to maximize the joint sum of profits, that is:

$$\begin{split} \Pi \Big(q_i, q_j, p_i, p_j \Big) &= (p_i - \mu c) q_i + \Big(p_j - \mu c \Big) q_j \\ &= \sum_i \frac{(p_i - c) (\gamma \Big(p_j + \emptyset - 1 \Big) + 1 - \emptyset + p_j \Big)}{1 - \gamma^2} \end{split}$$

In this case, equilibrium prices are:

$$p_i^m = \frac{1}{2}(c\mu + 1 - \emptyset)$$

Total emissions are²:

$$z(q_1^m + q_2^m) = \frac{1 - c\mu - \emptyset}{\gamma + 1}$$

It is interesting to see that as ϕ increases, prices decrease for all m. As the degree of environmental concern increases, demand shrinks, and firms need to set lower prices.

² We assume that $\phi \leq 1-c$ so that output and prices are always positive for all m.



Merger, Prices and Emissions We are now ready to state our main prediction. The merger decreases prices if and only if

$$\mu < \frac{\gamma \emptyset + 2c - \gamma}{c(2 - \gamma)} \coloneqq \hat{\mu}$$

However, whenever $\mu < \hat{\mu}$, the merger increases emissions. The threshold $\hat{\mu}$ is increasing in φ and decreasing in γ . As in standard competition models, a merger presents a trade-off. On one hand, the merger increases market power, potentially leading to higher prices. On the other hand, the merger can generate efficiencies, allowing cost savings to be partially passed through to consumers. Thus, a merger results in higher prices if, and only if, the efficiencies are insufficiently large. Our model suggests a potential environmental "benefit" associated with price increases, as a reduction in output implies a reduction in emissions. Conversely, should the merger generate significant efficiencies, the merged entities may increase output (as production becomes more cost-effective), leading to higher emissions.

The threshold $\hat{\mu}$ increases with Φ . The more environmentally concerned the consumer, the less likely it is that the merger will decrease emissions. This counterintuitive outcome arises because an increase in Φ diminishes the consumer's willingness to pay, reducing firms' market power and making a pro-competitive outcome more probable. Conversely, the threshold $\hat{\mu}$ decreases with γ . A higher degree of product differentiation enhances the merger's ability to create market power, thereby reducing the likelihood of the merger being pro-competitive.

Green Innovation In this section, we explore how a merger can reduce emissions not only by decreasing output, which inevitably leads to higher prices, but also by encouraging investments in green innovations. We propose a modification to our model for this analysis. Suppose that before engaging in the Bertrand competition, each firm has the option to invest a cost of K > 0 in green technology. This technology, conceptualized as an emission abatement mechanism, enables firms to produce with minimal pollution. Given the consumer's environmental concerns, such innovation is likely to boost demand³. Firms will invest in innovation only if the anticipated increase in revenue outweighs the technology's cost, K. We examine how a merger influences firms' incentives to innovate.

 $^{^3}$ In scenarios where consumers are indifferent to environmental impact ($\phi = 0$), firms lack the incentive to invest in green technology. In the real world, the cost of emissions and the financial benefits derived from investments in abatement technologies often lead to cost reductions. The logic behind this alternative scenario parallels that of our initial model.



We specifically focus on equilibria where both firms choose to innovate⁴. Let $\Delta\pi$ (m) be the benefits for a single firm from the innovation as a function of market structure m (given that both firms innovate). To obtain these expressions, we compute firms' profits in the case of $\Phi = 0$, and we compare them with the profits that firms gain when $\Phi > 0$. Then,

$$\Delta\pi(0) = \frac{(\gamma - 1)\phi(\emptyset + 2c - 2)}{(\gamma - 2)^2(\gamma + 1)} > 0$$
$$\Delta\pi(1) = \frac{\phi(\emptyset + 2c\mu - 2)}{4(\gamma + 1)} > 0$$

For all m, both firms invest in the green technology if and only if the cost K is low enough.

$$\Delta \pi(0) \ge K \Rightarrow K \le \overline{K_0}$$

$$\Delta\pi(1)\geq K\Rightarrow K\leq \overline{K_1}$$

The merger increases the incentives to innovate as $\overline{K_1} > \overline{K_0}$.

If $K \in (\overline{K_1}, \overline{K_0}]$ both firms invest in the green technology if and only if the merger occurs (m = 1). The rationale behind this is straightforward. A merger enhances firms' incentives to innovate by increasing the returns on such investments. Innovation, particularly those that increase consumer demand through environmental benefits, becomes more financially appealing as it can elevate firms' profits. In the absence of a merger, however, competitive pressures may erode these additional profits. A merger mitigates this competition, enabling firms to allocate more resources towards innovation.

A merger can lead to a reduction in emissions through two distinct pathways. Firstly, by potentially reducing output, a merger might inadvertently raise prices, a scenario generally unfavourable to consumers. Secondly, and more constructively, it can encourage investments in green technologies. This dual-faceted outcome highlights the complex impact mergers can have on both market dynamics and environmental sustainability.

Due to data availability the section on green innovation is currently missing in the empirical results, future iterations of the paper might include it.

⁴ In the absence of a merger (m = 0), it is possible to find equilibria where only one firm innovates, leading to higher emissions compared to scenarios where both firms innovate (resulting in zero emissions). Our analysis concentrates on situations where both firms innovate, assessing whether a merger can amplify incentives for green innovation.



4. DATA AND DESCRIPTIVE STATISTICS

4.1. Data

Merger data is collated from S&P Capital IQ transactions on private and publicly listed firms globally from 2006 to 2022, I have to limit the sample to 2006 for transactions as early emissions data is only available from 2004. For each transaction I am provided with unique identifiers for the acquirer and target, their country of incorporation, and sector (SIC code). Fundamentals data is collected from Compustat and S&P, where available data on revenues, total assets and liabilities is matched to the merger database.

Firm-level carbon emission data is obtained from S&P Capital IQ. GHG scope 1 absolute emissions (emissions from directly emitting sources that are owned or controlled by a company) are used in this paper. Later iterations might include GHG scope 2 emissions (emissions from the consumption of purchased energy generated upstream from a company's direct operations) and GHG scope 1 intensity emissions (absolute emissions scaled by their sales or revenues). Transactions for which emission data is not available are excluded from the sample, so the merger figure might look smaller with respect to other papers that use the entirety universe of merger.

4.2. Descriptive statistics

Figure 1 shows that majority of the mergers in the sample are following 2020, this might not be aligned to usual samples in the merger literature however it is dictated by emissions data becoming more broadly available in recent years. Table 1 highlights how the majority of the mergers in the sample are in the manufacturing sector, the reasoning is bi-fold firstly the manufacturing sector is highly concentrated and historically had a significant merger activity, secondly, as the manufacturing sector is the most polluting, environmental regulation has usually applied mandatory disclosure and or targets for this sector before expanding it to the rest of the economy.






FIGURE 1 • MERGERS IN THE SAMPLE (MATCHED WITH EMISSIONS)

TABLE 1 • SECTOR DISTRIBUTION OF MERGERS (MATCHED WITH EMISSIONS)

	SIC sector	Mergers
1	Manufacturing	746
2	Financials	558
3	Information Technology	513
4	Consumer Discretionary	425
5	Health Care	325
6	Materials	310
7	Communication Services	284
8	Real Estate	228
9	Consumer Staples	218
10	Energy	175
11	Utilities	155
	Total	3,937

Table 2 provides detailed descriptive statistics for scope 1 absolute, revenues, assets and liabilities for the sample with all mergers. Overall, the descriptive statistics underscore the heterogeneity in financial metrics among mergers, with some companies exhibiting extreme values in emissions, revenue, assets, and liabilities. While Table 3 compares the mean values



of scope 1 absolute, revenues and assets5 between cancelled and successful mergers. As expected successful mergers have lower emissions, but they tend to have lower revenues and a smaller asset base.

TABLE 2 • DESCRIPTIVE STATISTICS: ALL MERGERS

Statistic	Total Scope 1 Absolute (CO2 Tonnes)	Total Revenues (USD Millions)	Total Assets (USD Millions)	Total Liabilities(USD Millions)
Mean	2,570,209.94	161.27	1,725.30	247.74
Standard Deviation	9,228,942.96	9,790.84	106,802.12	13,267.11
Min	792.57	-31.93	0.00	0.00
25th Percentile	6,327.38	8.05	27.12	4.97
Median	33,058.22	18.06	165.24	12.11
75th Percentile	238,669.33	49.19	365.38	55.63
Max	52,549,649.00	1,018,691.47	11,114,132.33	1,379,798.73

TABLE 3 • DESCRIPTIVE STATISTICS: COMPARISON BETWEEN SUCCESSFUL AND CANCELLED MERGERS

Mean	Cancelled Mergers	Successful Mergers	Difference
Total Scope 1 Absolute (CO2 tonnes)	2,570,209.94	3,380,567.49	-810,357.55
Total Revenues (USD Millions)	161.27	49.56	111.72
Total Assets (USD Millions)	1,725.30	270.30	1455.01

5. EMPIRICAL METHODOLOGY AND RESULTS

5.1. Panel event study

I firstly employ a panel event study methodology, in line with work of Miller (2023) and Chaisemartin and D'Haultfoeuille (2020), to investigate the impact of corporate mergers on total Scope 1 emissions. The panel event study framework allows for the analysis of temporal effects of merger events while controlling for unobserved heterogeneity across firms and over time. This approach builds upon the foundational work of seminal event studies by Fama *et al.* (1969) and MacKinlay (1997), which have been instrumental in

⁵ Unfortunately, the data for total liabilities is missing for several cancelled mergers, for this reason it has not been reported here.



examining the effects of corporate events on firm outcomes. The panel event study equation is specified as follows:

$$Y_{i,t} = \beta_0 + \sum_{\tau = -k}^{k} \beta_{\tau} D_{T,\tau} + X'_{i,t} \gamma + \alpha_i + \lambda_t + u_{i,t}$$

The dependent variable in the model is the log-transformed total scope one absolute emissions⁶, it is important to highlight that this is the sum of emissions of both the target company and the acquirer, cases where prior to the merger either company does not report their emission are excluded from the sample⁷. The coefficient of interest is $\beta \tau$, capturing the effect of the event at different time periods relative to the event. $D_{T,\tau}$ are indicator variables that take the value of 1 if time t is τ periods relative to the event (with $\tau = 0$ being the event period), and 0 otherwise.

 $X'_{i,\tau}$ is a vector of control variables for firm i at time, in this case revenues, assets and liabilities are used, these controls are used as studies (such as Hartzmark and Shue 2023) found that revenues, assets and liabilities might impact how much a company pollutes. To control for confounding factors, the model includes several fixed effects α_i is a fixed effect for the country of the acquiring firm, target firm (when they differ) and their sectors, emissions could be influenced by country specific policies or sector practices and/or specificity. Year fixed effects (λ_t) are included to control for macroeconomic trends and shocks that vary over time. The inclusion of fixed effects is crucial for controlling unobserved heterogeneity, thereby mitigating the risk of omitted variable bias.

5.2. Panel event study results

Table 4 summarizes the results of the event study of the impact of a mergers on Scope 1 absolute emissions. In column (1) no fixed effects or controls are included, this is also the same specification which is illustrated in Figure 2. In column 2-4 fixed effects are progressively added (sector, country and year fixed effects). Column (5) is the most comprehensive, including firm-level controls along with all fixed effects. Across all specification coefficients are negative (ranging from is -0.592 to -0.195) and statistically significant, indicating a robust negative effect of mergers on emissions after accounting for

⁷ I have spoken to the data provider on how emissions are categorised after the merger, usually they are reported only for the acquirer, if they are reported for both it means that the company acquired is still mandated to independently report their emissions (these are rare cases). Both instances are left in the sample.



⁶ The logarithmic form was adopted and the data were windsorized at the 2nd and 98th percentiles in line with other papers using emissions as their outcome variable Bolton and Kacperczyk (2021).


various factors. Figure 2 shows that the impact of a merger is a persistent decrease in the resulting company's emissions, which is still present three years following the merger.

TABLE 4 • RESULTS OF THE EVENT STUDY FOR SCOPE 1 ABSOLUTE EMISSIONS

	log(Scope 1 Absolute Emissions)				
	(1)	(2)	(3)	(4)	(5)
Post-Event	-0.589***	-0.592***	-0.587***	-0.195***	-0.200***
	(0.013)	(0.013)	(0.010)	(0.011)	(0.011)
Sector FE	N	Y	Y	Υ	Y
Country FE	N	N	Y	Y	Y
Year FE	N	N	N	Y	Y
Firm-level controls	N	N	N	N	Y
R^2	0.012	0.385	0.467	0.483	0.485
Adj. R^2	0.012	0.384	0.467	0.483	0.485
N	216,784	216,784	216,784	216,784	216,384

Note: The regression reports the combined companies' total emissions from the year of the merger to three years after. The controls are revenues, assets and, liabilities. The Fixed effects are SIC sector fixed effects, emission year, and companies' country. The decrease in the number of observations is because some companies are missing at least one control variable. The standard errors are clustered at firm level (regression without clustering leads to similar results).

FIGURE 2 • PERCENTAGE CHANGE IN SCOPE 1 ABSOLUTE EMISSIONS FOLLOWING A MERGER








FIGURE 3 • PERCENTAGE CHANGE IN SCOPE 1 ABSOLUTE EMISSIONS FOLLOWING
AN HORIZONTAL MERGER

TABLE 5 • RESULTS OF THE EVENT STUDY FOR SCOPE 1

ABSOLUTE EMISSIONS - HORIZONTAL MERGERS

log(Scope 1 Absolute Emissions)	
(4)	(5)
-0.195***	-0.0712***
(0.017)	(0.017)
Y	Y
Y	Y
Y	Y
N	Y
0.575	0.577
0.574	0.576
84,036	83,818
	Y N 0.575 0.574

Note: The regression reports the combined companies' total emissions from the year of the merger to three years after. The controls are revenues, assets and, liabilities. The fixed effects are SIC sector fixed effects, emission year, and companies' country. The decrease in the number of observations is because some companies are missing at least one control variable. The standard errors are clustered at firm level (regression without clustering leads to similar results).



In Table 5 I run the same specification as Table 4 but I limit my sample to horizontal mergers, by focussing on companies within the same SIC sector. Across all specification coefficients are negative (ranging from is -0.485 to -0.065) and statistically significant, indicating a robust negative effect of mergers on emissions after accounting for various factors. Similarly Figure 3 shows that the impact of a horizontal merger is a persistent decrease in the resulting company's emissions, which is still present three years following the merger.

5.3. Quasi-experiment

As selection into mergers is endogenous, the main complication is that the average treatment effect (ATE) where $ATE = E[y_i(C=1) - y_i(C=0)]$ are the emissions of firm i when it is (not) a part of merged company j=1 (j=0). This cannot be observed in the data, leading to a selection bias boxed below which creates issues with the estimates between merged and non-merged companies:

$$E[y_i(1)|C = 1] - E[y_i(0)|C = 0]$$

$$= E[y_i(1)|C = 1] - E[y_i(0)|C = 1]$$

$$+ E[y_i(0)|C = 1] - E[y_i(0)|C = 0]$$

In order to isolate the causal effect of merger on emissions I adopt a methodology similar to Seru (2014), Bena and Li (2014), and Gugler *et al.* (2003). In an ideal experimental setting, I could randomly assign firms with similar characteristics into merged and non-merged companies and remove this selection bias. To proxy for this ideal setting the empirical strategy in this section of the paper adopts a quasi-experiment involving cancelled mergers, i.e. mergers that were announced but failed to successfully complete, aiding to generate exogenous variation in acquisition outcomes of target firms. I hypothesize that the reasons for which the mergers failed to go through are unrelated to emissions of the target (control group).

Mergers could fail to complete after being announced due to a variety of reasons including regulatory hurdles (Eckbo, 1983), financing issues (Kaplan and Stromberg, 2009), cultural clashes (Weber, Shenkar, and Raveh, 1996), economic condition changes (Shleifer and Vishny, 2003), discoveries during due diligence (Krishnan, Hitt, and Park, 2005), and shareholder opposition (Mulherin and Boone, 2000). These factors should be unrelated to emissions of the target.

In my specification the treatment group is composed of firms in a completed merger and the control group is firms in a merger that was announced but subsequently cancelled. The two groups then form a sample in which the assignment of a firm to the acquirer role can



be considered random. This assumption allows me to eliminate any selection bias by comparing the emissions of firms in the treatment group before and after the merger with those in the control group (Seru, 2014).

The empirical strategy leverages the difference-in-differences (DD) framework to estimate the impact of mergers on corporate emissions. Specifically, we compare the logarithm of total Scope 1 absolute emissions (windsorized) between companies that completed their mergers (treatment group) and those that cancelled their mergers (control group). The specification is as follows:

$$Y_{it} = \alpha + \beta_1 \operatorname{After}_{it} + \beta_2 \left(\operatorname{After}_{it} \times T_i \right) + X'_{i,t} \gamma + \alpha_i + \lambda_t + u_{i,t}$$

where After is an indicator variable that takes a value of one for all the years after the event date and zero otherwise, and T is an indicator variable that takes a value of one for targets in the treatment group and zero for targets in the control group. Similar to the event study in this specification we have $X'_{i,\tau}$ a vector of control variables for firm i at time and several fixed effects, country of the acquiring firm and target firm (when they differ) and their sectors.

5.4. Quasi-experiment results

Table 6 shows the difference-in-differences results for the entire sample. The coefficient on the Post variable indicates the effect of the post-merger period on emissions. While it varies in significance across different specifications, it is consistently positive and significant in models (3) to (5), suggesting an increase in emissions post-merger when accounting for various fixed effects and controls. The interaction term Post*Treated is consistently negative and highly significant across all models, indicating that treated firms experienced a significant reduction in emissions compared to the control group after the merger. This finding is robust to the inclusion of sector, year, and country fixed effects, as well as firm-level controls, underscoring the validity of the observed effect.

Similarly Figure 4 demonstrates the presence of parallel trends, as evidenced by the overlapping confidence intervals for emissions of treated and control firms prior to the merger date. After the merger date, the treated companies show a decrease in emissions compared to the control group.

Table 7 focuses on horizontal mergers, examining the impact of mergers within the same industry. The coefficient on the Post-Event variable is negative and significant across all specifications, indicating a reduction in emissions for firms involved in horizontal mergers. The results are robust even after including various fixed effects and controls.



FIGURE 4 • DIFFERENCE IN PERCENTAGE CHANGE IN SCOPE 1 ABSOLUTE EMISSIONS

BETWEEN CANCELLED AND COMPLETED MERGERS




TABLE 6 • RESULTS OF THE DD SPECIFICATION FOR SCOPE 1 ABSOLUTE EMISSIONS

log(Scope 1 Absolute Emissions) (1) (2) (3)(4) (5) -0.485*** -0.481*** -0.065*** -0.195*** -0.0712*** Post-Event (0.022)(0.017)(0.015)(0.017)(0.017)Y Y Y Y Sector FE N Country FE Y Y N N Y Year FE N N N Y Y Firm-level controls N N N N Y R^2 0.008 0.433 0.577 0.554 0.575 Adj. R^2 0.008 0.433 0.554 0.574 0.576 N 84,036 84,036 84,036 84,036 83,818



TABLE 7 • RESULTS OF THE DD SPECIFICATION FOR SCOPE 1 ABSOLUTE EMISSIONS HORIZONTAL MERGERS

	log(Scope 1 Absolute Emissions)				
	(1)	(2)	(3)	(4)	(5)
Post	0.072	0.078	0.552***	0.452***	0.430***
	(0.077)	(0.058)	(0.056)	(0.051)	(0.051)
Post*Treated	-0.324***	-0.327***	-0.331***	-0.339***	-0.318***
	(0.080)	(0.060)	(0.058)	(0.052)	(0.053)
Sector FE	N	Y	Y	Υ	Y
Year FE	N	N	Y	Y	Y
Country FE	N	N	N	Y	Y
Firm-level controls	N	N	N	N	Y
R^2	0.002	0.433	0.483	0.572	0.575
Adj. R^2	0.002	0.433	0.483	0.571	0.574
N	89,710	89,710	89,710	89,710	89,544

6. CONCLUSION

This study delves into the intricate relationship between corporate mergers and environmental outcomes, specifically focusing on Scope 1 absolute emissions. The theoretical model explores the nuanced dynamics between oligopolistic competition, consumer environmental consciousness, and the impact of mergers on emissions. It highlights a trade-off where significant production efficiencies post-merger can lead to higher output, lower prices, and increased emissions, while minimal efficiencies result in higher prices but lower emissions, showcasing environmental benefits. The model also suggests that mergers may incentivize green innovation, offering a pathway to reducing emissions.

The findings from both empirical strategies highlight the reduction in emissions following a corporate merger. The first empirical strategy utilized a panel event study framework to assess the temporal effects of mergers on emissions. This approach, grounded in the work of Miller (2023) and Chaisemartin and D'Haultfoeuille (2020), allowed for the examination of emissions before and after the merger, while controlling for unobserved heterogeneity across firms and over time. The results consistently demonstrated a significant reduction in emissions following mergers, with the effect persisting up to three years post-merger. The



second empirical strategy adopted a quasi-experimental design to address the endogeneity of merger selection. Inspired by Seru (2014), Bena and Li (2014), and Gugler et al. (2003), this approach compared firms that completed mergers with those that announced but subsequently cancelled their mergers. Utilizing the difference-in-differences (DiD) framework, this methodology isolated the causal impact of mergers on emissions. The results indicated that firms in completed mergers experienced a substantial reduction in emissions compared to the control group. This effect was particularly pronounced in horizontal mergers, where firms within the same industry demonstrated significant emissions reductions. These results are aligned to the key hypothesis that an increase in market concentration is correlated with a decrease in emissions.

The findings highlight how corporate mergers improves the environmental performance of firms, particularly in terms of reduced emissions. This observation suggests a more nuanced picture of mergers, extending beyond their traditional economic or financial goals to include potential environmental advantages. The reduction in emissions post-merger could be attributed to several factors, including enhanced operational efficiencies, access to better technologies, and improved management practices. Moreover, the results emphasize the importance of distinguishing between the sources of environmental benefits, advocating for technological advancements as a key factor for improved environmental outcomes post-merger.

This study contributes to the broader literature on the non-market effects of market power and market concentration. By shedding light on the environmental implications of mergers, it provides valuable insights for policymakers and stakeholders. The results suggest that corporate consolidation, under certain conditions, can align with environmental sustainability goals. This has important implications for competition policy and environmental regulation, highlighting the need for a balanced approach that considers both economic and environmental objectives.

In conclusion, the study underscores the complex interplay between market consolidation and environmental performance. While mergers can lead to increased market power, they can also drive efficiencies that result in reduced emissions. The dual-faceted outcomes of mergers, in terms of both economic and environmental impacts, emphasize the need for integrated policy frameworks that promote sustainable business practices. Future research could further explore the mechanisms through which mergers influence environmental performance, as well as the long-term sustainability of these effects. This would provide a deeper understanding of how corporate strategies and market structures can be designed to support both economic growth and environmental protection.



REFERENCES

- Aghion, P., Bénabou, R., Martin, R., and Roulet, A. (2023), Environmental Preferences and Technological Choices: Is Market Competition Clean or Dirty?, «American Economic Review: Insights», 5, 1, pp. 1-19, DOI: 10.1257/aeri.20210014
- Ahern, K.R. and Harford J. (2014), The Importance of Industry Links in Merger Waves, «Journal of Finance», 69, 2, pp. 527-576, DOI: 10.1111/jofi.12122
- Autor, D. et al. (2020), The Fall of the Labor Share and the Rise of Superstar firms, «Quarterly Journal of Economics», 135, 2, pp. 645-709, DOI: 10.1093/qje/qjaa004
- Bena, J. and Kai, L. (2014), Corporate Innovations and Mergers and Acquisitions, «The Journal of Finance», 69, 5, pp. 1923-1960.
- Bolton, P. and Kacperczyk, M. (2021), *Do Investors Care About Carbon Risk?*, «Journal of Financial Economics», 142, 2, pp. 517-549, DOI: 10.1016/j.jfineco.2021.05.008
- Chaisemartin, C. de and D'Haultfoeuille, X. (2020), Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects, «American Economic Review», 110, 9, pp. 2964-2996, DOI: 10.1257/aer.20181169, https://www.aeaweb.org/articles?id=10.1257/aer.20181169
- Cole, M.A. and Elliott, R.J. (2006), Do Environmental Regulations Influence Trade Patterns? Testing Old and New Trade Theories, «The World Economy», 29, 8, pp. 1079-1097
- Crouzet, N. and Eberly, J.C. (2019), Understanding Weak Capital Investment: The Role of Market Concentration and Intangibles., «NBER working paper series»
- De Loecker, J., Eeckhout, J. and Unger, G. (2020), *The Rise of Market Power and the Macroeconomic Implications*, «The Quarterly Journal of Economics», 135, 2, pp. 561-644
- Delmas, M.A. and Montes-Sancho, M.J. (2011), *The Drivers of Greenhouse Gas Emissions Reductions*, «California Management Review», 54, 2, pp. 44-70
- Duflo, E., Greenstone M., and Hanna, R. (2008), *Indoor Air Pollution, Health and Economic Well-Being*, «S.A.P.I.E.N.S», 1, 1, pp. 7-12
- Eccles, R.G., Ioannou, I., and Serafeim, G. (2014), *The Impact of Corporate Sustainability on Organizational Processes and Performance*. «Management Science», 60, 11, pp. 2835-2857
- Eckbo, B.E. (1983), Horizontal Mergers, Collusion, and Stockholder Wealth, «Journal of Financial Economics», 11, 1-4, pp. 241-273
- Ellerman, A.D., Frank, J.C., and de Perthuis C. (2010), *Pricing Carbon: The European Union Emissions Trading Scheme*, Cambridge, Cambridge University Press
- Fama, E., Fischer, L., Jensen, M.C. and Roll, R. (1969), *The Adjustment of Stock Prices to New Information*, «International Economic Review», 10, 1, pp. 1-21
- Furman, J. and Orszag, P. (2015), A Firm-Level Perspective on the Role of Rents in the Rise in Inequality for us, https://gabriel-zucman.eu/files/teaching/FurmanOrszag15.pdf
- Gugler, K. et al. (2003), The Effects of Mergers: An International Comparison, «International Journal of Industrial Organization», 21, 5, pp. 625-653, https://www.sciencedirect.com/science/article/pii/S0167718702001078
- Gutierrez, G. and Philippon, T. (2023), How European Markets Became Free: A Study of Institutional Drift, «Journal of the European Economic Association», 21, 1, pp. 251-292. DOI: 10.1093/jeea/jvac071
- Hartzmark, S.M. and Shue, K. (2023), Counterproductive Sustainable Investing: The Impact Elasticity oaf Brown and Green Firms, https://ssrn.com/abstract=4359282
- IPCC (2018), Global Warming of 1.5°C. An IPCC Special Report on The Impacts of Global Warming of 1.5°C Above Preindustrial Levels and Related Global Greenhouse Gas Emission Pathways, Intergovernmental Panel on Climate Change
- Kang, K. and Xiao, M. (2023), Policy Deterrence: Strategic Investment in U.S. Broadband, Working Paper



- Kaplan, S.N. and Stromberg, P. (2009), The Effects of Leveraged Buyouts on Productivity and Related Aspects of Firm Behavior, «Journal of Financial Economics», 91, 1, pp. 1-23
- Koltay, G., Szabolcs L., and Valletti, T. (2023), Concentration and Competition: Evidence from Europe and Implications for Policy, «Journal of Competition Law & Economics», 19, 3, pp. 466-501
- Krishnan, H.K., Hitt, M.A., and Park D. (2005), *The Impact of Mergers on The Firms' Future Performance: Evidence from The Industry Consolidation*, «Strategic Management Journal», 26, 8, pp. 723-740
- Kumar, S. and Managi, S. (2012), *The Role of Environmental Regulation in Recessions*, «Environmental Economics and Policy Studies», 14, pp. 405-423
- MacKinlay, A.C. (1997), Event Studies in Economics and Finance, «Journal of Economic Literature», 35, 1, pp. 13-39
- Martin, G.J. and McCrain, J. (2023), *The Political Consequences of Media Consolidation*, Working Paper Miller, D.L. (2023), *An Introductory Guide to Event Study Models*, «Journal of Economic Perspectives», 37, 2, pp. 203-230, DOI: 10.1257/jep.37.2.203, https://www.aeaweb.org/articles?id=10.1257/jep.37.2.203.
- Mulherin, J.H. and Boone, A.L. (2000), *The Market for Corporate Control: The Empirical Evidence Since* 1980, «Journal of Economic Perspectives», 15, 1, pp. 67-90
- Nocke, V. and Whinston, M.D. (June 2022), Concentration Thresholds for Horizontal Mergers, «American Economic Review», 112, 6, pp. 1915-1948, DOI: 10.1257/aer.20201038
- Nordhaus, W.D (1991), To Slow or Not to Slow: The Economics of The Greenhouse Effect, «The Economic Journal», 101, 407, pp. 920-937
- Phillips, G.M. and Zhdanov, A. (2013), R&D and The Incentives from Merger and Acquisition Activity, «Review of Financial Studies», 26, 1, pp. 34-78, DOI: 10.1093/rfs/hhs109
- 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, pp. 97-118
- Seru, A. (2014), Firm Boundaries Matter: Evidence from Conglomerates and RD Activity, «Journal of Financial Economics», 111, 2, pp. 381-405, https://www.sciencedirect.com/science/article/pii/S0304405X1300278X
- Shleifer, A. and Vishny, R.W. (2003), *Stock Market Driven Acquisitions*, «Journal of Financial Economics», 70, 3, pp. 295-311
- Stern, N. (2007), The Economics of Climate Change: The Stern Review, Cambridge, Cambridge University Press
- Weber, Y., Shenkar, O., and Raveh, A. (1996), Corporate Culture Fit and Performance in Mergers and Acquisitions, «Human Relations», 49, 9, pp. 1181-1202

Climate Economics and (its) Knowledge è il titolo della dodicesima edizione del Premio Giorgio Rota, i cui vincitori – Konstantin Boss, Lorenzo Sileci, Alessandra Testa, Costanza Tomaselli – sono stati premiati in occasione della XII Giorgio Rota Conference il 15 maggio 2024 al Campus Luigi Einaudi. Il volume che raccoglie i saggi vincitori è introdotto da Silvana Dalmazzone del Dipartimento di Economia e Statistica "Cognetti de Martiis", Università di Torino.

Il Centro Einaudi nasce nel 1963 e si propone come punto di riferimento nel panorama culturale italiano: opera a due livelli, teoretico e di analisi empirica, tentando di cogliere e «anticipare» i tempi e i temi dell'analisi e del dibattito politico-economico nazionale e locale. Il Centro Einaudi conduce ricerca indipendente, organizza seminari, cura la formazione di giovani studiosi, pubblica libri e periodici su carta e online. L'approccio è orientato alle politiche, multidisciplinare, il riferimento è la tradizione liberale.

Climate Economics and (its) Knowledge is the title of the twelfth edition of the Giorgio Rota Prize, whose winners – Konstantin Boss, Lorenzo Sileci, Alessandra Testa, Costanza Tomaselli – were awarded at the XII Giorgio Rota Conference on 15 May 2024 at the Luigi Einaudi Campus. The volume collecting the winning essays is introduced by Silvana Dalmazzone of the Department of Economics and Statistics 'Cognetti de Martiis' of University of Turin.

Centro Einaudi was founded in 1963. It aims to be a reference in the Italian cultural panorama both on the theoretic and on the empirical level, by analysing and 'anticipating' the timing and subjects of political and economic public discussion, both nationwide and locally. The mission is to conduct independent research and to provide innovative recommendations for local and national policy makers. The approach is policy-oriented, multi-disciplinary, the reference is the liberal tradition.