



HAL
open science

Three essays on climate and air pollution mitigation policies

Marion Leroutier

► **To cite this version:**

Marion Leroutier. Three essays on climate and air pollution mitigation policies. Economics and Finance. Université Panthéon-Sorbonne - Paris I, 2021. English. NNT : 2021PA01E040 . tel-03547345

HAL Id: tel-03547345

<https://theses.hal.science/tel-03547345>

Submitted on 28 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

UNIVERSITÉ PARIS 1
PANTHÉON-SORBONNE

ÉCOLE D'ÉCONOMIE DE PARIS



ÉCOLE DOCTORALE: ED 465 – Économie Panthéon Sorbonne

THÈSE

Pour l'obtention du grade de docteur en Sciences Économiques
de l'Université Paris 1 Panthéon-Sorbonne

Présentée et soutenue publiquement le 5 juillet 2021 par

Marion LEROUTIER

TROIS ESSAIS SUR LES POLITIQUES D'ATTÉNUATION DU CHANGEMENT CLIMATIQUE ET DE LA POLLUTION DE L'AIR

Sous la direction de :

Philippe QUIRION

Katheline SCHUBERT

Composition du jury :

Président du jury

Mouez FODHA Professeur à l'Université Paris 1 Panthéon-Sorbonne, PSE

Rapporteurs

Olivier CHANEL Professeur à Aix-Marseille School of Economics (AMSE)

Ulrich WAGNER Professor à University of Mannheim

Examineurs

Laure DE PREUX Assistant Professor à Imperial College London

Mirabelle MUÛLS Assistant Professor à Imperial College London

Directeurs

Philippe QUIRION Directeur de recherche au CNRS, CIRED

Katheline SCHUBERT Professeur à l'Université Paris 1 Panthéon-Sorbonne, PSE

UNIVERSITÉ PARIS 1
PANTHÉON-SORBONNE

ÉCOLE D'ÉCONOMIE DE PARIS



DOCTORAL SCHOOL: ED 465 – Économie Panthéon Sorbonne

PHD THESIS

Submitted to Université Paris 1 Panthéon-Sorbonne
For the Degree of Doctor of Philosophy in Economics

Prepared and defended on Monday 5 July 2021 by

Marion LEROUTIER

THREE ESSAYS ON CLIMATE AND AIR POLLUTION MITIGATION POLICIES

Thesis Advisors:

Philippe QUIRION

Katheline SCHUBERT

Jury :

President

Mouez FODHA Professor, Université Paris 1 Panthéon-Sorbonne, PSE

Reviewers

Olivier CHANEL Professor at Aix-Marseille School of Economics

Ulrich WAGNER Professor at University of Mannheim

Examiners

Laure DE PREUX Assistant Professor at Imperial College London

Mirabelle MUÛLS Assistant Professor at Imperial College London

Advisors

Philippe QUIRION Senior researcher at CNRS, CIRED

Katheline SCHUBERT Professor, Université Paris 1 Panthéon-Sorbonne, PSE

Remerciements - Acknowledgements

Chaque fois qu'on m'a demandé qui étaient mes directeurs de thèse, la réaction a été remarquablement constante lorsque j'ai donné la réponse : "La chance !". La chance, en effet. Philippe, merci d'avoir été disponible chaque fois que j'avais besoin de conseils, et d'avoir su me redonner de l'énergie à chacun de nos points. Ta connaissance pointue du secteur de l'électricité m'a grandement aidée dans les tâtonnements de mon premier chapitre, et je me réjouis à l'idée de continuer à travailler avec toi sur notre article "IDF". Merci aussi pour ton humanité et ton engagement au quotidien, qui sont sources d'inspiration. Katheline, merci d'avoir accepté de m'encadrer alors que j'avais quitté le circuit académique, et pour tes précieux conseils, en particulier sur les aspects plus théoriques de mes travaux en cours. Merci pour ta bonne humeur communicative. Merci pour l'opportunité d'enseigner le TD de ton cours, expérience qui m'a beaucoup plu. A tous les deux, je vous suis très reconnaissante de m'avoir soutenue dans les différentes étapes de cette thèse.

I am grateful to Olivier Chanel and Ulrich Wagner for accepting to be the referees of this dissertation. I thank Laure de Preux, Mouez Fodha and Mirabelle Muûls for accepting to be the examiners. A special thank you to Ulrich Wagner and Mirabelle Muûls, for accepting these roles on a short notice. I thank all the jury members for thoroughly reading the different chapters of my dissertation, and for the precious feedback during my pre-defense.

Je remercie vivement François Libois, qui a fait partie de mon comité de thèse, pour tous ses conseils, aussi bien sur mes projets que sur l'après-thèse, sans oublier l'introduction passionnante aux données géo-codées. I deeply thank Nicolas Koch for welcoming me at MCC, for his feedback on my research, and for his support on the job market.

Merci beaucoup à mes co-auteurs, Léo, Marie-Abèle et Philippe, qui ont contribué à faire de cette thèse une expérience collective et m'ont beaucoup appris. Nos échanges ont été d'autant plus importants pour moi lors des confinements successifs, en maintenant un lien par-delà les murs de nos appartements. Merci Léo pour tes nombreuses explications dans mes débuts sur R, et pour m'avoir sensibilisée au souci de la reproductibilité dans la recherche. Merci Marie-Abèle pour tes conseils tout au long du projet, ta bonne humeur et tes encouragements constants.

Travailler sur la pollution de l'air implique de côtoyer les travaux d'autres disciplines comme les sciences atmosphériques, l'épidémiologie ou la santé publique, parfois complexes à appréhender depuis l'économie. Je remercie profondément les personnes qui ont pris le temps de m'expliquer certains aspects techniques : Isabelle Coll, Augustin Collette, Olivier Perrussel, Alexia Baudic, et Sabine Host. J'en profite pour saluer tout le travail d'Airparif et d'Atmosud en faveur d'une meilleure connaissance de la qualité de l'air. Merci également au SDES du Ministère de la Transition Ecologique et Solidaire de permettre l'exploitation des données à des fins académiques, et en particulier à Etienne Kouevi pour sa gestion attentive de l'accueil des chercheurs et des conventions.

I thank all the researchers whom I had the chance to interact with at conferences and seminars, and during my stay at MCC. In particular, I thank Michael Pahle for inviting me at PIK and to the workshop on EU climate policies, Linus Mattauch for the stimulating discussions where I always learn a lot, and the FAERE for its great work connecting the environmental economics community in France.

I thank those who helped me decide to pursue a PhD and find a topic I like: Prof. Ian Gough for organizing the LSE seminar on Climate Change, Inequality and Social Policy, and for our discussion. Merci à Eric Maurin, mon directeur de mémoire en 2015, de m'avoir le premier aiguillée vers la littérature sur la pollution. I thank Lutz Sager for taking the time to share his PhD experience, and my friend Sugandha for our numerous discussions on whether we should start a PhD. Merci à Gaëtan F, Quentin P et Clara M-T de m'avoir encouragée, sans doute sans le savoir, à entamer une thèse, en me renvoyant une image si positive de la vie et du travail de doctorant.

Avoir partagé mon temps entre deux institutions aussi complémentaires que PSE et le CIRED a été une grande chance, et promet une liste de remerciements d'autant plus fournie. Je remercie le groupe Régulation et Environnement de PSE et ses différents membres pour le soutien matériel, ainsi que pour les interactions pendant et après les séminaires. Merci en particulier à Hélène Ollivier, avec qui je me réjouis de collaborer, Katrin Millock, Mireille Chiroleu-Assouline, Fanny Henriet, Thomas Douenne, Stéphane Gauthier, Angelo Secchi et David Martimort. Je remercie également Laurent Gobillon et Miren Lafourcade pour leurs éclairages sur la méthode SCM et l'économie urbaine. Un grand merci à la directrice du programme doctoral de PSE, Sylvie Lambert, pour son attention au bien-être des doctorants, et à Véronique Guillotin pour son implication remarquable dans tous les aspects administratifs de la thèse. Merci aux délégués des doctorants, et en particulier à Sarah Schneider-Strawczynski pour ses conseils.

Merci aux R3-treiziens et R3-treiziennes actuels et passés : Fanny, Simon, Paul D, Paul B, Georgia, Youssef, Marion et Pauline. Paul B, je n'aurais peut-être pas commencé cette thèse si on ne s'était pas retrouvés de l'autre côté de la Manche en 2017, et, qui sait, je ne l'aurais peut-être pas finie si on n'avait pas été dans le même bureau de ce côté-ci ! Cette thèse te doit beaucoup. Youssef, j'espère que nos discussions sur la mixité à l'école et la recherche se prolongeront dans l'après-thèse. Marion, partager avec toi les dernières lignes droites de la thèse m'aura aidé à appréhender cette période plus sereinement. Un grand merci pour ta relecture de mon introduction, et d'avoir mis un peu de vert dans notre bureau. Merci aussi à toutes celles et ceux qui font de PSE un endroit où il fait bon déjeuner : les anciens, Clara, Clément, Mariona, Laura, Lisa, Rozenn, Alessandro, Pepe, Malka, et puis Irène, Justine, Mélanie, Claire, Sarah, Juliette, Lennart, Martin, Benjamin, Manon, Sofia, Paolo, Antton, Sara, Yaz, Caroline, Max. Une mention spéciale à Sophie et Paul de la CAAFEEP, pour leur contribution au bien-être de tous les amateurs de caféine. Je remercie tout particulièrement Mariona, pour nos journées de télétravail qui m'ont tellement redonné le moral cette dernière année.

Côté CIRED, je remercie toute l'équipe, qui contribue à faire du CIRED un endroit unique et si convivial. Comme trois ans de thèse passent vite quand on est à cheval sur deux institutions, et confiné six mois ! J'espère avoir l'occasion de revenir vous voir au jardin tropical. Merci à ceux qui m'ont aidée dans mes recherches, en particulier à Gaëtan

Giraudet, Vincent Viguié - qui m'a, le premier, initiée aux problématiques croisant ville et environnement - et Fabien Leurent. Merci à Aurélien S pour nos discussions à différentes étapes de la thèse. Merci aux organisateurs du séminaire CIRED, Aurélie et Antoine. Merci Arancha pour l'énergie que tu déploies autour de la communication du CIRED. Merci Estelle pour l'attention exceptionnelle que tu portes à l'accueil des nouveaux, pour ta gentillesse et tes conseils de lecture. Merci Carine pour la gestion de mon contrat avec le CNRS. Merci Meriem pour les airs fredonnés dans le couloir. Merci à Thierry, mon premier co-bureau et fournisseur de chocolat, et à Améline, Claire, Emilien et Nicolas pour leurs encouragements sur les dernières étapes de la thèse. Merci à tous les YOLO, Améline, Aurélie, Léa, Vivien, Behrang, Basile, Nicolas, Claire, Quentin, Adrien, Tamara, pour les fous rires. Merci aussi à mes co-bureaux Lauriane et Julien, et à Philippe, Laurent, Anne, Auriane, Samuel, Naceur, Antoine, Vincent et Florian pour les discussions du midi.

Je remercie Marie et Jacques pour tous les week-ends passés au Vaudoué ces dernières années, parenthèses bucoliques toujours rafraîchissantes. Un grand merci aussi de nous avoir prêté votre appartement parisien pendant le premier confinement. Ce luxe nous a permis de vivre cette période relativement sereinement, j'en mesure la chance.

Je remercie tous mes autres amis pour les moments d'évasion, et ceux passés à refaire le monde : Ségolène, pour le chemin parcouru ensemble depuis la prépa ; Perrine, pour ta folie et ta soif d'aventures contagieuses ; Arthur, pour nos discussions si enrichissantes et nos randonnées ; Marie C, pour nos interludes musicaux - un jour, les carottes râpées perceront ; Elsa, qui sais si bien répandre l'optimisme autour de toi ; Sylviane, pour nos promenades et baignades au canal ; Mathilde et Félicien, pour les dîners qui passent si vite en votre compagnie ; Marie S-M, pour les goguettes à Anères, Paris, et au Havre ; Marine, pour les soirées aux Cascades, au théâtre et au bois de Vincennes ; Sophie, pour les randonnées, pas encore à jeun ; Katia, pour notre aventure macédonienne ; Laure, pour les Skypes transatlantiques et les balades parisiennes ; Alma, pour les glaces entre Paris et Berlin ; Luke, pour les discussions sur l'écologie et l'économie ; Wen et Raphaëlle, pour votre joie de vivre ; Sugandha, for everything we have shared since our London flat ; Aymeric, pour les films qui font quand même un peu peur ; Juliette, pour ton écoute toujours attentive ; Thomas, pour les découvertes musicales et ton si beau cadeau d'anniversaire ; Eléonore, Marion, Sarah, pour les bons moments pendant et après Kimso ; les Jeudis, pour les galettes des rois parisiennes, les Nouvel An bruxellois et les pacs bretons ; Marie C et Pierre, pour m'avoir initiée aux trésors du bois de Vincennes, qui ne s'arrêtent pas au JATP ; Merci aux radis aussi, vous avez beaucoup compté.

Merci à celles et ceux avec qui j'ai chanté, des karaokés dans les resto chinois aux églises du 14ème, des salons d'appartement aux cafés du village, merci à Anaël, Julia, à toutes les Oréades, aux goguetters, chanter avec vous m'a apporté beaucoup de joie.

Je remercie du fond du cœur ma famille, qui m'a toujours encouragée dans mes études et mes choix. Merci à mon père d'avoir toujours été disponible pour m'aider quand j'en avais besoin, et pour son attention minutieuse à l'actualité sur la pollution de l'air. Je ne sais comment remercier ma mère pour son écoute bienveillante, son humour, et son soutien à toute épreuve. Merci de me rappeler régulièrement qu'il faut travailler pour vivre, et non pas vivre pour travailler. Merci à mon beau-frère Gaëtan d'être toujours serein et rassurant. Merci à ma sœur Fanny d'être si gentille et d'avoir su m'encourager

toutes les fois (nombreuses !) où j'étais indécise. Merci à mon grand-père d'être là, tout simplement.

Enfin, je remercie Paul, protagoniste essentiel de ces trois dernières années. Merci d'avoir été là dans les peines et les joies de la thèse. Merci pour les tiramisus, les trajets en vélo, les codes Python, les chants révolutionnaires chiliens, les relectures de papier tard le soir, les parades d'escalade, et tant d'autres choses.

Marion Leroutier, Paris, le 04 juin 2021

Résumé de la thèse

Cette thèse est composée de trois chapitres indépendants, visant d'une part à mieux comprendre l'efficacité des politiques d'atténuation de la pollution, et d'autre part à estimer quels sont les secteurs et les individus contribuant à la pollution.

Le premier chapitre, intitulé *Carbon Pricing and Power sector Decarbonisation : Evidence from the UK*, examine l'efficacité d'une taxe carbone introduite dans le secteur électrique britannique en 2013, le CPS, qui est passé de 5,9 euros par tonne de CO₂ en 2013 à 26 euros par tonne de CO₂ en 2017. Depuis son introduction, le secteur électrique britannique a connu une transformation rapide, avec une diminution de la part du charbon dans la production d'électricité de 40% à 7% entre 2012 et 2017. Compte tenu de la forte intensité d'émission du charbon déplacé, les émissions de CO₂ du secteur électrique ont diminué de 57% sur cette période. J'évalue l'impact causal du CPS sur les émissions du secteur de l'électricité au Royaume-Uni en utilisant la méthode du contrôle synthétique (*Synthetic Control Method* ou SCM en anglais): je compare la trajectoire des émissions du Royaume-Uni à celle d'un Royaume-Uni contrefactuel constitué d'une combinaison pondérée d'autres pays européens. Je trouve que le CPS a conduit à une baisse des émissions du secteur électrique estimée à entre 20,5 à 26% par an en moyenne entre 2013 et 2017, selon les hypothèses retenues pour les potentiels facteurs de confusion. Je mets en évidence trois mécanismes par lesquels la taxe a affecté les émissions : premièrement, la taxe carbone a réduit l'intensité d'émission des centrales actives sur le marché électrique pendant toute la période, vraisemblablement par le biais d'un report du charbon vers le gaz ; deuxièmement, la taxe carbone a induit la fermeture de certaines centrales très émettrices ; troisièmement, les centrales déjà menacées de fermeture en raison des réglementations européennes sur les émissions industrielles ont eu une probabilité plus élevée d'effectivement fermer que dans le reste de l'Union Européenne.

Le deuxième chapitre, intitulé *Estimating the Causal Effects of Cruise Traffic on Air Pollution using Randomization-Based Inference*, est un travail conjoint avec Léo Zabrocki et Marie Abèle Bind. Nous examinons la contribution des navires de croisière à la pollution de l'air à Marseille, l'une des plus grandes villes portuaires européennes, dans un contexte d'inquiétude croissante quant aux effets du trafic maritime sur la pollution et la santé des résidents. Nous combinons un nouvel algorithme de *matching* par paire avec des données de séries temporelles à haute fréquence pour créer des expériences hypothétiques aléatoires, dans lesquelles seul le trafic de navires de croisière varie. Nous estimons l'effet de cette variation de trafic sur la pollution ambiante à l'échelle de la ville. Nous quantifions l'incertitude avec une méthode d'inférence basée sur la randomisation et construisons des intervalles de Fisher à 95% (*Fisherian intervals* ou FI). L'arrivée de bateaux de croisière dans le port à une heure donnée augmente les concentrations horaires de dioxyde d'azote (NO₂) de 4,7 µg/m³ (FI 95% : [1,4; 8,0]), les concentrations de dioxyde de soufre (SO₂) de 1,2 µg/m³ (FI 95% : [-0,1; 2,5]) et celles de particules (PM₁₀) de 4,6 µg/m³ (FI 95% : [0,9; 8,3]). L'entrée d'un navire de croisière supplémentaire dans le port un jour donné augmente les concentrations moyennes journalières de SO₂ de 0,7 µg/m³ (FI 95% : [0,1; 1,4]), soit une augmentation de 30% par rapport à la moyenne observée pour Marseille.

Les concentrations de PM_{10} et $PM_{2,5}$ sont également plus élevées de respectivement $3,5 \mu\text{g}/\text{m}^3$ (FI 95% : [0,5; 6,5]) et $2,5 \mu\text{g}/\text{m}^3$ (FI 95% : [0,2; 4,9]) le jour suivant, soit une augmentation de 13 à 16%, qui peut en partie refléter une augmentation du trafic routier. Nos résultats suggèrent que des expériences randomisées hypothétiques bien conçues constituent une approche prometteuse pour mieux comprendre les externalités négatives du trafic maritime.

Le troisième chapitre, intitulé *Tackling Transport-Induced Pollution in Cities : A Case Study in Paris*, est un travail conjoint avec Philippe Quirion. Nous examinons dans quelle mesure les individus contribuent aux émissions de CO_2 et de polluants atmosphériques locaux dans leur mobilité quotidienne, en nous basant sur des données d'enquête de mobilité en Ile-de-France, une aire urbaine parmi les plus polluées en Europe. Nous étudions les causes des inégalités en matière d'émissions et la façon dont les émissions pourraient être réduites. Nous documentons de grandes inégalités, les 20% plus gros émetteurs contribuant à 75 à 85% des émissions un jour de semaine représentatif, en fonction du polluant. Nous étudions les facteurs associés aux émissions du top de deux manières : premièrement, dans une analyse de décomposition exacte, nous montrons que la distance, les choix modaux et l'intensité des émissions contribuent de manière égale à expliquer les émissions du top pour les polluants locaux, tandis que pour le CO_2 , ce sont surtout les distances élevées et les choix modaux qui contribuent aux émissions du top ; deuxièmement, dans une analyse de régression, nous mettons en évidence l'association entre certaines caractéristiques d'emploi et les distances totales parcourues, la part modale de la voiture et son intensité d'émission. Nous montrons également les différentes associations entre le revenu du ménage et l'intensité d'émission des véhicules utilisés pour se déplacer, selon l'intensité d'émission porte sur les polluants locaux ou sur le CO_2 . Enfin, nous formulons des scénarios de potentiel de report modal basés sur des temps de trajet contrefactuels. Dans notre scénario central, 53% des déplacements en voiture pourraient être transférés vers les transports publics ou - pour la plupart - vers le vélo, en particulier électrique. Cela permettrait d'économiser 214 millions d'euros par an en CO_2 et en pollution locale, ce qui représente 19 à 21% du coût total de la pollution induite par la mobilité quotidienne. Nous discutons de ce qui peut entraver ou encourager un tel report modal, ainsi que des options alternatives pour ceux qui n'ont pas d'option de report.

DISCIPLINE : Sciences Économiques

MOTS-CLEFS : taxe carbone ; politiques climatiques ; pollution de l'air ; électricité ; trafic maritime ; transport urbain ; inégalités environnementales

Thesis summary

The three chapters of this dissertation aim at better understanding the effectiveness of pollution mitigation policies on the one hand, and who contributes to pollution at the sectoral and individual level, on the other hand.

The first chapter, *Carbon Pricing and Power sector Decarbonisation: Evidence from the UK*, examines the effectiveness of a carbon tax introduced in the UK power sector in 2013, the CPS, which increased from €5.9 per ton of CO₂ in 2013 to €26 per ton of CO₂ in 2017. Since its introduction, the UK power sector has undergone a rapid transformation, with a decrease in the share of coal in electricity generation from 40% to 7% between 2012 and 2017. Given the high emission-intensity of the displaced coal, power sector CO₂ decreased by 57% over the same period. I evaluate the causal impact of the CPS on power sector emissions in the UK with the Synthetic control method: I compare the trajectory of UK emissions to that of a counterfactual UK made of a weighted combination of other European countries. I find the CPS led to a decrease in UK emissions by between 20.5 and 26 percent on an average year between 2013 and 2017, depending on the assumptions made for potential confounding factors. I highlight three mechanisms via which the tax affected emissions: first, the carbon tax decreased the emission-intensity of plants staying in the market over the entire period, allegedly via a fuel switch from coal to gas; second, the carbon tax induced the net closure of some high-emission plants; finally, plants already at risk of closure due to European regulations on industrial emissions had a higher probability to effectively close than in the rest of the EU.

The second chapter, *Estimating the Causal Effects of Cruise Traffic on Air Pollution using Randomization-Based Inference*, is a joint work with Léo Zabrocki and Marie Abèle Bind. We examine the contribution of cruise vessels to air pollution in Marseille, one of the largest European port cities, in a context of rising concerns over the effects of maritime traffic on pollution and residents' health. We combine a new pair-matching algorithm with high-frequency time series data to create hypothetical randomized experiments where only cruise vessel traffic varies across matched pairs. We estimate the effect of this variation in traffic on city-level pollutant concentrations. We quantify uncertainty with randomization-based inference and build 95% Fisherian intervals (FI). The arrival of cruise vessels in the port increases hourly concentrations of nitrogen dioxide (NO₂) by 4.7 µg/m³ (95% FI: [1.4, 8.0]), of sulphur dioxide (SO₂) by 1.2 µg/m³ (95% FI: [-0.1, 2.5]), and of particulate matter (PM₁₀) by 4.6 µg/m³ (95% FI: [0.9, 8.3]). Having one additional cruise vessel entering the port on a given day increases city-level daily SO₂ by 0.7 µg/m³ (95% FI: [0.1, 1.4]), a 30% increase compared to the daily average. City-level PM₁₀ and PM_{2.5} are also higher by respectively 3.5 µg/m³ (95% FI: [0.5, 6.5]) and 2.5 µg/m³ (95% FI: [0.2, 4.9]) on the following day, a 13-16% increase which may partly capture an increase in road traffic. Our results suggest that well-designed hypothetical randomized experiments provide a principled approach to better understand the negative externalities of maritime traffic.

The third chapter, *Tackling Transport-Induced Pollution in Cities: A Case Study in*

Paris, is a joint work with Philippe Quirion. We examine how much individuals contribute to emissions of CO₂ and local air pollutants in their daily mobility based on detailed mobility survey data from the Paris region, an urban area amongst the most polluted in Europe. We investigate what drives inequalities in emissions and how emissions could be reduced. We document large inequalities, with the top 20% of emitters contributing 75-85% of emissions on a representative weekday, depending on the pollutant. We investigate factors associated with high emissions in two ways: first, in an exact decomposition analysis, we show that distance, modal choices and emission intensity contribute equally to explaining top local pollutant emissions, while for CO₂ emissions, high distances and modal choice drive top emissions the most; second, in a regression analysis, we highlight the association between some employment characteristics and total distances travelled, the car modal share and its emission intensity. We also show the different associations between income and vehicles' local vs CO₂ emission intensity. Finally, we formulate scenarios of modal shift potential based on counterfactual travel times. In our central scenario, 53% of car trips could be shifted to public transport or - for the most part - cycling, in particular e-cycling. This would save an annual €214m of avoided CO₂ and local pollution, which represents 19-21% of the total cost of daily mobility-induced pollution. We discuss what may hinder or encourage such modal shift, and alternative options for those unable to shift.

DISCIPLINE: Economics

KEYWORDS: carbon tax; climate policy; air pollution; power sector; maritime traffic; urban road transport; environmental inequalities

Sommaire général

- Remerciements - Acknowledgements** **5**
- Résumé de la thèse** **9**
- Thesis summary** **11**
- Introduction générale** **19**
- General Introduction** **39**
- 1 Carbon Pricing and Power sector Decarbonisation: Evidence from the UK** **57**
 - 1.1 Introduction 58
 - 1.2 The UK carbon tax: context and expected effects 61
 - 1.2.1 The UK Carbon Price Support 61
 - 1.2.2 Descriptive evidence 63
 - 1.2.3 Potential confounders 66
 - 1.3 Empirical strategy 68
 - 1.3.1 The synthetic control method 68
 - 1.3.2 The Data 70
 - Power plant-level emission data: 71
 - Country-level power sector characteristics: 71
 - 1.3.3 Selecting the predictors 72
 - 1.3.4 Selecting countries entering the donor pool 74
 - 1.4 Results 75
 - 1.4.1 Upper Bound 75
 - 1.4.2 Lower bound 78
 - Potential confounders and emission decomposition: 78
 - Lower bound: counterfactual emissions of plants converted to biomass if they had not converted: 80
 - 1.4.3 Inference 82
 - In-Time placebo 82
 - Leave-one-out test 83
 - Permutation test 83
 - 1.4.4 Risk of spillovers 86
 - Risk of spillover via increased electricity imports 86
 - Risk of spillover via a waterbed effect 89
 - 1.5 Discussion 90

1.5.1	Channels contributing to emission reduction	90
1.5.2	Comparison of the results with existing estimates	92
1.6	Conclusion	93
2	Estimating the Causal Effects of Cruise Traffic on Air Pollution using Randomization-Based Inference	95
2.1	Introduction	96
2.2	Materials and Methods	98
2.2.1	Data	98
	Vessel Traffic Data	98
	Air Pollution and Weather Data	98
	Road traffic Data	100
2.2.2	Method	101
	Stage 1: Formulating Plausible Interventions on Vessel Traffic . . .	101
	Stage 2: Designing the Hypothetical Randomized Experiments . . .	103
	Stage 3: Analyzing the Experiments using Randomization-based Inference	104
	Point estimate.	104
	Randomization-based quantification of uncertainty.	105
	Stage 4: Interpreting the Results	106
2.3	Results	106
2.3.1	Matching Results	107
2.3.2	The Effects of Cruise Vessel Traffic on Air Pollutants	111
2.4	Discussion	114
2.4.1	Putting our Results into Perspective	114
2.4.2	Reflection on the Methods	116
2.4.3	Potential Paths for Future Research	117
2.4.4	Concluding Remarks	118
3	Tackling Transport-Induced Pollution in Cities: A Case Study in Paris	119
3.1	Introduction	120
3.2	Air pollution and transport emissions in Paris	123
3.3	Data and methodology	124
3.3.1	The Data	124
	Individual transport:	124
	Emission factors	126
	Counterfactual travel time data	126
	Charging stations for Electric Vehicles	128
3.3.2	Methodology	128
	Building individual measures of contribution to pollution. . .	128
	Exact factor decomposition analysis	130
	Individual characteristics associated with high emissions . . .	132
3.4	Results	133
3.4.1	How unequal are contributions to emissions?	133
3.4.2	Are high emissions mostly due to high distances, high-emission modal shares or highly polluting cars?	136

3.4.3	Who emits pollution?	138
3.4.4	What are the options to reduce emissions?	142
	Shift to low-emission modes:	145
	Avoid travelling by teleworking:	149
	Improve : shift to an Electric Vehicle:	150
3.5	Discussion	152
3.5.1	A 80-20 rule?	152
3.5.2	Traditional and less traditional factors associated with emissions . .	152
3.5.3	From modal shift potential to actual modal shift	153
3.5.4	Limits	156
3.6	Conclusion	157
	Bibliography	159
	Appendices	181
	A Appendix to Chapter 1: Carbon Pricing and Power sector Decarbonisation: Evidence from the UK	183
A.1	Appendix	183
A.1.1	Evolution of per capita emissions, demand, trade and emission intensity in the UK and other countries	183
	Per capita power sector emissions	183
	Demand, trade and intensity channels	184
A.1.2	Potential confounders	187
	European level: LCP and IED Directives	187
	UK level: support to biomass conversion	188
	UK level: support to renewable energy	188
	UK level: capacity market	190
A.1.3	Identification of power installations	191
A.1.4	Summary statistics for the country-level dataset	192
A.1.5	Data sources	194
A.1.6	Common Support for predictors	196
A.1.7	Counterfactual emissions in the absence of biomass conversion for Drax and Lynemouth plants	198
A.1.8	Lower bound removing emissions from plants converted to biomass	199
A.1.9	Sensitivity of the results to the choice of predictors	201
A.1.10	Sensitivity of the results to the choice of the donor pool	206
A.1.11	Sensitivity of the results to a longer pre-treatment period	207
A.1.12	Estimation of the spillovers removing all interconnected countries from the donor pool	209
	B Appendix to Chapter 2: Estimating the Causal Effects of Cruise Traffic on Air Pollution using Randomization-Based Inference	211
B.1	Appendix	211
B.1.1	Reproducibility	211
	Replication Materials	211
	Caveats	211

B.1.2	Data Sources	211
B.1.3	Data Wrangling	212
	Formatting Cruise Traffic Data	212
	Merging Cruise Traffic Data with Air Pollution and Weather Data	212
	Merging the Data with Road Traffic Data	213
	Adding Calendar Indicators and Imputing Missing Values	213
B.1.4	Exploratory Data Analysis	213
	Vessel Traffic	214
	Air Pollutants	214
	Weather Parameters	214
	Road Traffic	222
B.1.5	Matching Procedure Diagnostics	223
	Hourly Experiment on Cruise Vessels' Arrivals	223
	Hourly Experiment on Cruise Vessels' Departures	237
	Daily Experiment on Cruise Traffic	251
B.1.6	Results - Analysis of the Matched Data	262
	Toy Example for Understanding Randomization Inference	262
	Hourly Experiment on Cruise Vessels' Arrivals	269
	Fisherian Intervals for the subset of hourly pairs with wind blowing from the West	272
	Hourly Experiment on Cruise Vessels' Departures	273
	Fisherian Intervals for the subset of hourly pairs with wind blowing from the West	276
	Daily Experiment on Cruise Vessel Traffic	277
	Dose-response	279
	Fisherian Intervals for pairs on days with wind blowing from the West	280
	Effect on induced road traffic	281
B.1.7	Sensitivity Analyses	282
	Alternative Test Statistic: Wilcoxon Signed-Rank Statistic	282
	Fisherian Intervals for Non-Imputed Observations	284
	Neymanian Inference for the Average Treatment Effect	288
	Statistical Power, Type M and S errors	289
B.1.8	Regression Analysis of the Initial Data	291
	Hourly Experiments	291
	Daily Experiment	292
B.1.9	Road Traffic and NO ₂ Concentrations	296

C Appendix to Chapter 3: Tackling Transport-Induced Pollution in Cities: A Case Study in Paris **299**

C.1	Appendix	299
C.1.1	Assumptions on NO _x , PM _{2.5} and CO ₂ emissions by transport mode	299
	Buses	299
	Cars and two-wheelers owned by the household	300
	Taxis and cars and two-wheelers not owned by the household	301
C.1.2	Method to retrieve counterfactual transport time with Google API .	302

C.1.3	Estimation of journey stage-specific emissions accounting for cold starts	304
C.1.4	Additional tables and figures	306

Introduction générale

Cette introduction est structurée en quatre parties : je commence par présenter les deux problématiques environnementales abordées dans ce mémoire : le changement climatique et la pollution de l'air extérieur. J'expose ensuite les différentes manières dont ces deux externalités pourraient être régulées selon la théorie économique. Une évaluation des défis pratiques rencontrés dans le choix des instruments politiques et l'estimation des dommages suit. Enfin, je présente les objectifs, les méthodologies et les contributions de cette thèse.

Le changement climatique et la pollution atmosphérique locale, deux problèmes environnementaux majeurs

Dommages et tendances des émissions

Il est désormais bien établi que le changement climatique d'origine anthropique représente un risque majeur pour les humains et les écosystèmes (IPCC, 2018). Sans une diminution drastique des émissions de gaz à effet de serre (GES) au cours des prochaines décennies, l'augmentation de la température moyenne et du nombre d'événements météorologiques extrêmes menacera les moyens de subsistance, et augmentera probablement la pauvreté et les inégalités entre et au sein des pays (Hallegatte and Rozenberg, 2017; Diffenbaugh and Burke, 2019). Dans le même temps, les engagements climatiques déclarés des pays sont en deçà des ambitions de l'Accord de Paris (Rogelj et al., 2016). Après une réduction temporaire en 2020 en raison de la pandémie de Covid-19 et des mesures de distanciation sociale associées (Le Quéré et al., 2020), les émissions mondiales de dioxyde de carbone (CO₂, le principal gaz à effet de serre) ont rebondi au premier trimestre 2021 (IEA, 2021).

La pollution atmosphérique est un deuxième enjeu environnemental majeur, identifié par l'Organisation mondiale de la santé (OMS) comme le "plus grand risque sanitaire d'origine environnementale" au monde (WHO, 2014). Elle est causée par la concentration de plusieurs types de polluants, distincts des gaz à effet de serre, dans l'air ambiant : les polluants les plus importants - qui seront abordés dans cette thèse - sont les particules PM_{2.5} et PM₁₀¹, le dioxyde d'azote (NO₂), l'ozone (O₃) et le dioxyde de soufre (SO₂). La pollution de l'air extérieur est responsable d'environ 4,2 millions de décès par an dans le monde (WHO, 2014), tandis que la pollution par les PM_{2.5} est liée à elle seule à environ 379 000 décès dans l'Union européenne (UE) (European Environment Agency, 2020) et 40

¹Les particules sont classées en fonction de leur taille. Les PM_{2.5} ne comprennent que les particules inférieures à 2,5 µm, tandis que les PM₁₀ comprennent toutes les particules inférieures à 10 µm.

000 décès en France (Santé Publique France, 2021). Des réglementations sur la pollution atmosphérique ont été mises en œuvre dans la plupart des pays à revenu élevé depuis les années 1950, à la suite du Grand Smog de Londres de 1952 et du Smog de Donora de 1948 aux États-Unis. Ces réglementations sont devenues de plus en plus strictes au fil du temps. Depuis, les niveaux de pollution ont considérablement diminué (voir US EPA (2016) pour les États-Unis, et Sicard et al. (2021) pour l'Europe).

La pollution atmosphérique reste une préoccupation importante pour au moins quatre raisons : premièrement, malgré la diminution observée dans les pays à revenu élevé, les concentrations restent souvent supérieures aux valeurs recommandées par l'OMS pour les particules, l'ozone et le dioxyde d'azote. Deuxièmement, les niveaux de pollution augmentent dans de nombreuses villes des pays émergents et à bas revenu (WHO, 2016). Troisièmement, des travaux récents suggèrent qu'il n'existe pas de seuil en dessous duquel l'exposition à la pollution est sans danger pour la santé physique (Di et al., 2017). Quatrièmement, les conséquences de l'exposition à la pollution atmosphérique vont au-delà de son impact sur la santé physique, des travaux récents en économie mettant en évidence des dégradations de la santé mentale, du capital humain et de la productivité induites par la pollution (voir, par exemple, Roth (2017); Chang et al. (2019)).

Responsabilité sectorielle des dommages causés par la pollution

La plupart des polluants d'origine anthropique rejetés dans l'atmosphère proviennent de processus de combustion. Les secteurs économiques contribuant aux émissions de CO₂ émettent généralement aussi des polluants atmosphériques locaux, et inversement. Pour illustrer cela, la figure 0.0.2 montre la contribution des différents secteurs aux émissions de CO₂, NO_x² et PM_{2.5} dans l'Union européenne en 2018³. La production d'électricité et de chaleur contribue à 30% des émissions totales de gaz à effet de serre de l'UE, mais aussi à 9% des émissions de NO_x, 2% des émissions de PM_{2.5}, et 20% des émissions de SO_x⁴. De même, le transport contribue à 23% des émissions totales de gaz à effet de serre de l'UE, mais aussi à 66% des émissions de NO_x, 22% des PM_{2.5} et 19% des émissions de SO_x. Les deux secteurs économiques considérés dans cette thèse, le secteur de la production d'électricité et de chaleur (chapitre 1), et le secteur des transports (chapitres 2 et 3), contribuent ensemble à 53% des émissions de gaz à effet de serre de l'UE en 2018, 75% de ses émissions de NO_x, 24% de ses émissions de PM_{2.5} et 39% de ses émissions de SO_x.

Si les chiffres ci-dessus nous permettent de déduire la responsabilité sectorielle dans

²NO_x est la formule désignant les oxydes d'azote, une catégorie d'émissions produisant entre autres le NO₂.

³comprenant le Royaume-Uni, qui a quitté l'UE en janvier 2020.

⁴SO_x est la formule désignant les oxydes de soufre, une catégorie d'émissions produisant entre autres le SO₂.

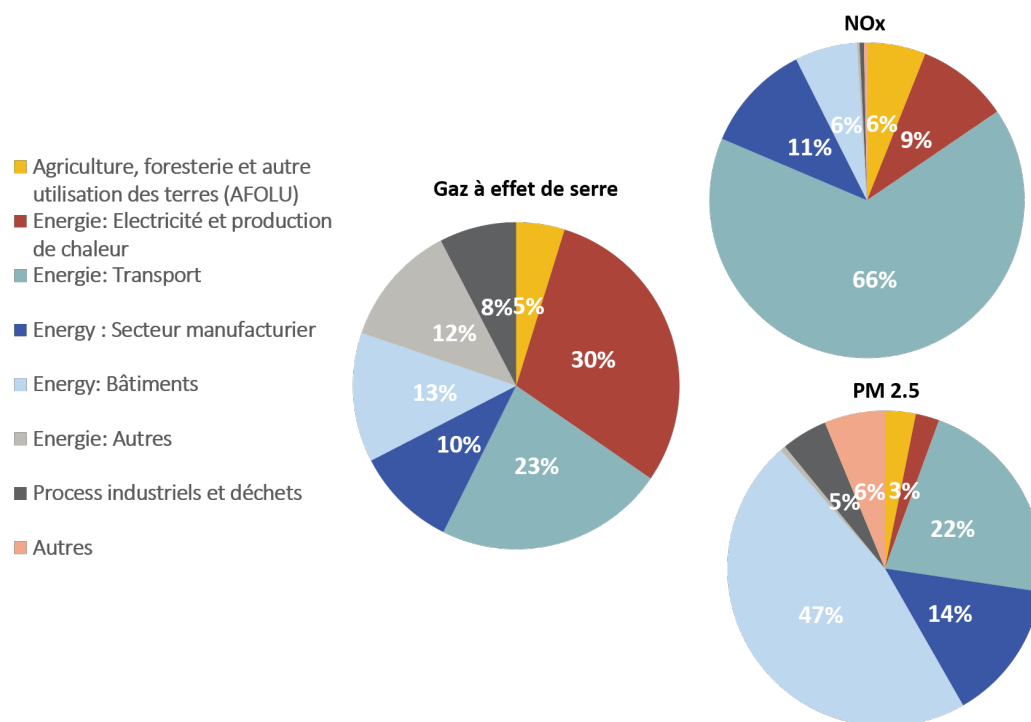


Figure 0.0.1 – Répartition des émissions de gaz à effet de serre, NOx et PM_{2.5} par secteur au sein de l’UE 28 en 2018

Sources : données ClimateWatch par secteur et pays pour les émissions GES, et données EMEP par secteur et pays pour les émissions de polluants de l’air locaux. La répartition sectorielle de ClimateWatch est basée sur la nomenclature définie par le GIEC. La répartition sectorielle de l’EMEP est basée sur la nomenclature GNFR-19, communément utilisée pour les polluants de l’air locaux. La plupart des secteurs GNFR peuvent facilement être rapprochés d’un secteur GIEC, à part les catégories “Autres” et “Naturelles”, qui ont été allouées au secteur GIEC “Autres”. Le secteur GNFR “Autres installations de combustion” inclut les émissions des installations de combustion dans les secteurs résidentiel, commercial, institutionnel et agricole. Il est considéré comme correspondant au secteur “Energie : bâtiments” de la nomenclature GIEC.

les dommages climatiques induits par le CO₂, ils ne permettent pas de conclure sur la responsabilité sectorielle dans les dommages liés à la pollution atmosphérique, mais donnent juste un indice. Cela s’explique par le fait que toutes les émissions de CO₂ contribuent de la même manière au changement climatique. En revanche, le lien entre les émissions de polluants atmosphériques et les dommages causés au bien-être humain s’opère en deux étapes : premièrement, les émissions provenant de la source de pollution se traduisent en concentrations dans l’atmosphère. Ce phénomène peut être influencé par une série de facteurs météorologiques et topologiques. Par exemple, une vitesse de vent plus élevée a un effet “nettoyant” et réduit la pollution, tandis que les épisodes d’inversion de température - pendant lesquels une masse d’air chaude empêche l’air froid près du sol de monter dans l’atmosphère, et les polluants de se disperser - sont associés à des concentrations plus élevées (Arceo et al., 2016). Deuxièmement, les concentrations de pollution se traduisent

en dommages en fonction du degré d'exposition de la population locale à cette pollution (des niveaux de pollution élevés dans une zone rurale à faible densité de population causeront moins de dommages aux humains que les mêmes niveaux de pollution dans une ville à forte densité de population), et de sa vulnérabilité. Le chapitre 2 de cette thèse s'intéresse à cette relation non-triviale entre les émissions et les dommages sanitaires.

L'approche économique de la pollution

La pollution en tant qu'externalité à réguler

Les économistes ont depuis longtemps identifié la pollution comme une défaillance du marché à laquelle il faut remédier. La pollution au sens large, y compris les contaminants de l'eau, du sol ou de l'air, est considérée comme une externalité négative (Pigou, 1920) : c'est un sous-produit de la production ou de la consommation d'une personne donnée, qui crée un dommage pour une tierce personne. Pourtant, la première personne ne supporte pas le coût du dommage causé, car il n'est pas reflété dans les prix du marché. En conséquence, la fourniture de biens polluants est trop importante par rapport à ce qui serait optimal. Dans l'économie néoclassique, les externalités sont l'une des trois défaillances du marché identifiées comme étant à l'origine d'inefficacité et justifiant l'intervention des pouvoirs publics. Savoir quels sont les meilleurs instruments de politique publique pour réduire l'externalité de pollution est une question clé en économie de l'environnement depuis plusieurs décennies (Bretschger and Pittel, 2019).

Juger les mérites des différents instruments de régulation

Il existe deux types de politiques publiques pour réglementer la pollution : les régulations de type "command-and-control", qui consistent à fixer des normes sur le type de technologie à utiliser (dans le cas de mandats technologiques) ou sur le niveau maximal de pollution (dans le cas de normes de performance), et à contrôler leur application ; et les solutions fondées sur le marché ou sur des incitations, qui consistent à fixer un prix à la pollution pour inciter les pollueurs à réduire leurs émissions. Les solutions fondées sur le marché peuvent encore être distinguées entre celles qui réglementent la quantité d'émissions (comme les marchés de plafonnement et d'échange fixant un plafond au nombre de permis d'émission en circulation) et celles qui réglementent leur prix (comme les taxes pigouviennes).

Les mérites de différentes réglementations peuvent être jugés sur la base de différents critères. Goulder and Parry (2008) citent comme critères d'évaluation potentiels

l'efficacité économique, le rapport coût-efficacité, la répartition des bénéfices et des coûts, la capacité à faire face aux incertitudes et la faisabilité politique.

En ce qui concerne l'efficacité, les preuves théoriques suggèrent que les solutions basées sur le marché sont plus rentables que les réglementations "command-and-control" dans le premier cas : cela est dû au fait que les pollueurs diffèrent généralement dans leur capacité à réduire la pollution - mesurée par leur coût marginal de réduction - (Newell and Stavins, 2003), et que seuls les instruments basés sur des incitations permettent d'égaliser les coûts marginaux et les bénéfices marginaux pour tous les émetteurs si les pollueurs connaissent leurs coûts de réduction mais pas le décideur politique. Un autre avantage théorique des instruments fondés sur le marché est qu'ils exploitent tous les canaux de réduction de la pollution, alors que les réglementations de type "command-and-control" en négligent certains, en particulier la réduction de la production (Spulber, 1985). Enfin, les instruments fondés sur des incitations génèrent généralement des recettes (sauf lorsque les permis d'émission sont distribués gratuitement) qui, si elles sont recyclées de manière appropriée, peuvent réduire d'autres taxes distorsives et générer des gains d'efficacité. Ce recyclage de la taxe pourrait ainsi créer un double dividende, en améliorant la qualité environnementale tout en réduisant la perte de bien-être nette associée à la politique environnementale (Baumol and Oates, 1988; Pearce, 1991; Chiroleu-Assouline and Fodha, 2014).

Toutefois, en présence de multiples défaillances du marché, la supériorité des instruments fondés sur le marché peut être remise en question, et une combinaison de différents instruments peut être appropriée, y compris des normes fondées sur le marché, mais aussi des subventions publiques. Par exemple, lorsque les coûts administratifs de surveillance des émissions sont élevés, les mandats peuvent être supérieurs (Goulder and Parry, 2008). Goulder et al. (2016) montrent également qu'en présence de distorsions préexistantes du marché des facteurs, les normes d'énergie propre sont plus rentables que les instruments fondés sur les prix, car elles représentent une taxe implicite plus faible sur les facteurs de production.

Le soutien à l'innovation dans les technologies à faible émission de carbone, outre la tarification carbone, se justifie également par la nature de bien public de l'innovation : dans un contexte où les gains de l'innovation ne peuvent pas être entièrement appropriés, l'investissement est trop faible sans soutien public (Fischer and Newell, 2008). Le secteur du transport routier illustre bien la multiplicité des défaillances du marché et la nécessité d'adopter des politiques de second choix, puisque le transport routier contribue à de multiples externalités : émissions de CO₂, pollution atmosphérique locale, mais aussi bruit, accidents et congestion (Parry et al., 2007).

Concernant les instruments de prix par rapport aux instruments de quantité, Weitz-

man (1974) suggère que leurs mérites relatifs dépendent de la pente de la fonction de dommage marginal lorsqu’il existe une incertitude sur les coûts globaux de la réduction de la pollution - ce qui est souvent le cas dans le monde réel. Les instruments de prix sont supérieurs lorsque la courbe de dommage marginal est plate et les instruments de quantité sont supérieurs lorsque la courbe de dommage marginal est raide. L’intuition qui sous-tend ce résultat est qu’il est d’autant plus important de choisir la bonne quantité lorsque les coûts des dommages augmentent considérablement pour une petite variation de la pollution. Comme la courbe des dommages marginaux du changement climatique est plutôt plate, les économistes ont préféré une taxe sur le carbone à un marché du carbone pour ces raisons théoriques. Au niveau international, Weitzman (2015) plaide également en faveur d’une taxe carbone uniforme plutôt que de permis négociables, en raison de l’incertitude concernant les trajectoires de coût d’abattement au niveau pays.

En ce qui concerne l’équité, des problèmes de répartition peuvent se poser parce que certains individus supportent un coût disproportionné de la réglementation. Par exemple, on a souvent constaté que les taxes sur le carbone étaient régressives dans les pays à haut revenu si les recettes de la taxe n’étaient pas redistribuées (Poterba, 1991). Cela s’explique par le fait que les ménages les plus pauvres consacrent une part plus importante de leurs dépenses de consommation aux biens à forte intensité de carbone. Des études ont montré que les taxes sur le carbone peuvent devenir progressives si les recettes fiscales sont redistribuées sous la forme de transferts forfaitaires à tous les ménages (Metcalf, 2009a; Cronin et al., 2018; Douenne, 2020; Berry, 2019). Cependant, il est peu probable que les transferts forfaitaires corrigent les problèmes d’équité horizontale, en raison de la forte hétérogénéité de l’incidence fiscale au sein d’un décile de revenu donné (Sallee, 2019; Douenne, 2020; Berry, 2019). Dans certains cas, il a été constaté que les normes sont plus régressives qu’une taxe carbone accompagnée de transferts forfaitaires (Davis and Knittel, 2018; Levinson, 2018) ; à l’inverse, Zhao and Mattauch (2020) montrent que les normes sont plus équitables lorsque les consommateurs affichent une préférence pour les attributs technologiques à forte teneur en carbone - et vérifient que c’est le cas aux Etats-Unis.

Le soutien politique et la faisabilité seront probablement affectés par les propriétés distributionnelles objectives des réglementations proposées. Au-delà de ces propriétés objectives, cependant, des recherches récentes soulignent l’importance de la perception des politiques (Douenne and Fabre, 2021; Maestre-Andres et al., 2019), et de facteurs contextuels tels que la confiance politique (Rafaty, 2018). Passant en revue les travaux sur les effets de différents instruments de régulation Goulder and Parry (2008) concluent qu’“aucun instrument n’est clairement supérieur sur toutes les dimensions pertinentes pour le choix de politiques publiques ; même le classement sur une seule dimension dépend

souvent de la nature de l'instrument et des circonstances" (Goulder and Parry (2008), p 153, ma traduction). Comme ces "circonstances" sont susceptibles de varier d'un contexte à l'autre, cette conclusion appelle à davantage de preuves empiriques sur les effets des réglementations réellement mises en œuvre à travers le monde. Le chapitre 1 de cette thèse prend en compte cet enjeu et fournit des éléments de preuves sur l'effet d'une taxe carbone introduite au Royaume-Uni dans le secteur de production d'électricité.

La régulation de la pollution atmosphérique et des émissions de CO₂ en pratique

Une variété d'instruments

Dans la pratique, les polluants locaux et les émissions de CO₂ ont été réglementés par une combinaison d'instruments. Les instruments de type "command-and-control" ont été historiquement plus courants et ont consisté en des normes avec des exigences spécifiques pour utiliser les meilleures technologies disponibles (*Best Available Technology* ou BAT en anglais) ou d'autres mandats de technologies spécifiques (Metcalf, 2009b). Aujourd'hui, différents types de réglementations coexistent en fonction des juridictions et du polluant : les polluants locaux des installations industrielles (y compris le secteur électrique) sont réglementés par des normes dans l'Union européenne (directives sur les émissions industrielles) et aux États-Unis (Clean Air Act) ; par des programmes de plafonnement et d'échange dans certains États américains (California Regional Clean Air Incentives Market (RECLAIM)) ; par des taxes dans certains pays européens (la TGAP en France ou les redevances NO_x en Suède ou en Norvège (Bonilla et al., 2018)). Les émissions de CO₂ sont également réglementées via une combinaison d'instruments de type "command-and-control", par exemple sous la forme de normes d'efficacité énergétique (comme le CAFE américain) ou de normes d'émission (dans l'Union européenne) dans le secteur des transports, et d'instruments fondés sur le marché, par exemple le marché européen du carbone (EU ETS) ou l'ETS chinois. Les instruments basés sur le marché ont augmenté au cours de la dernière décennie, avec 64 instruments de tarification carbone à travers le monde en 2021 contre seulement 21 dix ans auparavant en 2011 (World Bank, 2021).

Preuves de l'efficacité des politiques existantes

L'efficacité de plusieurs des réglementations susmentionnées a été estimée de façon empirique (par exemple, Fowlie et al. (2012) pour RECLAIM ; Colmer et al. (2020) pour le SCEQE ; Currie and Walker (2019) pour le Clean Air Act ; Andersson (2019) pour la

taxe carbone suédoise). Cependant, ces évaluations sont rares, en particulier concernant les instruments de tarification carbone. Une des raisons est simplement que ces instruments sont plutôt récents. Une autre raison est d'ordre méthodologique : au-delà de la quantité de données requises, l'estimation causale de l'impact d'une tarification carbone sur les émissions nécessite de surmonter le "problème fondamental de l'inférence causale" (Holland, 1986a), selon lequel, dans les pays où une politique est en place, nous ne pouvons observer l'évolution des émissions qu'en présence de la politique, mais pas dans la situation contrefactuelle où la politique est absente.

Ce problème est d'autant plus prononcé dans le cas des instruments de tarification carbone, qui ciblent généralement un secteur ou une région entière. Sur les 21 évaluations empiriques des prix du carbone existants répertoriées par Rafaty et al. (2020), quatre évaluations portent sur le secteur manufacturier, cinq sur le seul secteur des transports, six sur les secteurs de l'électricité et de l'industrie manufacturière⁵ et trois mettent en commun plusieurs instruments nationaux sur différents secteurs. Seuls trois articles (y compris le chapitre 1 de cette thèse) se concentrent sur le seul secteur de l'électricité et ils considèrent tous le même instrument. Il est donc nécessaire d'obtenir davantage de preuves de l'efficacité de la tarification carbone en général, mais aussi dans le secteur de l'électricité en particulier, étant donné sa forte contribution aux émissions mondiales. En outre, il existe une grande hétérogénéité dans les estimations de l'efficacité de la taxe carbone Rafaty et al. (2020); Green (2021). Il est donc nécessaire de disposer de données plus systématiques sur les facteurs expliquant ces différences, qu'elles soient dues à des facteurs contextuels ou à des différences méthodologiques dans la stratégie d'estimation et le champ d'application considérés.

Comment prendre en compte deux externalités : co-bénéfices et arbitrages

Le fait que certains secteurs émettent à la fois une pollution locale et mondiale peut impliquer que les réglementations s'attaquant à un type de polluant affecteront également l'autre. Le fait qu'une rigueur accrue des politiques d'atténuation entraîne une diminution de la pollution atmosphérique locale dépend de l'élasticité de substitution entre les deux (Ambec and Coria, 2013). Si les polluants locaux se substituent aux émissions de CO₂, la baisse des émissions de CO₂ augmentera les émissions de polluants locaux. S'ils sont complémentaires, la baisse de CO₂ réduira également les émissions de polluants locaux.

Dans de nombreux cas, il s'est avéré que les polluants locaux et mondiaux sont complémentaires : dans ce cas, la politique climatique réduisant les émissions de CO₂ améliorera

⁵articles évaluant l'effet du marché ETS

également la qualité de l'air et pourra générer des bénéfices pour la santé. Ces bénéfices indirects de l'atténuation climatique font partie des effets secondaires positifs plus larges de la politique climatique, regroupés sous le concept de co-bénéfices (Nemet et al., 2010; Intergovernmental Panel on Climate Change, 2015). Des études antérieures suggèrent que les co-bénéfices de la qualité de l'air découlant des politiques d'atténuation peuvent représenter une part substantielle des bénéfices monétisés et égaler, voire dépasser, les coûts d'atténuation (Karlsson et al., 2020; Rauner et al., 2020).

Des cas où les polluants locaux se substituent aux émissions de CO₂ ont également été identifiés : c'est notamment le cas dans le secteur du transport automobile, où l'incitation à réduire l'intensité des émissions de CO₂ des voitures risque de stimuler la part de marché des voitures diesel, qui émettent plus de polluants locaux (Durrmeyer, 2018; Linn, 2019). Cela implique qu'une politique conçue pour internaliser l'une des externalités peut involontairement en augmenter une autre, ce qui complique la régulation.

Seules quelques études ont examiné empiriquement les implications des co-bénéfices pour la conception des politiques. Wagner and De Preux (2016) (2016) étudient la perte potentielle d'efficacité et d'équité de l'échange de droits d'émission, lorsqu'il donne lieu à un échange implicite de pollution locale entre des zones dont les dommages marginaux liés à la pollution atmosphérique sont hétérogènes. Durrmeyer (2018) montre que les propriétés distributives d'un bonus-malus automobile portant sur les émissions de CO₂ pourraient être améliorées en tenant compte de son impact sur l'émission de polluants locaux au niveau individuel. Dans le secteur du transport routier, des préférences individuelles différentes pour les attributs technologiques à forte intensité de carbone *versus* forte intensité de pollution locale pourraient affecter ces propriétés redistributives, mais cette question reste, à ma connaissance, largement inexplorée.

Estimation des dommages

Les résultats théoriques sur la régulation optimale dépendent de paramètres tels que la forme de la courbe du coût marginal de réduction et de la courbe des dommages marginaux. Cependant, l'estimation des dommages induits par les externalités est particulièrement difficile. Les économistes s'appuient généralement sur les prix du marché pour effectuer une analyse du bien-être. Mais, par définition, les externalités ne sont pas tarifées. C'est pourquoi une partie importante de la littérature économique sur le changement climatique et la pollution atmosphérique est consacrée à l'estimation des dommages de ces externalités. L'estimation des dommages liés au changement climatique est plus difficile que celle des dommages liés à la pollution de l'air, car les dommages climatiques se produiront principalement dans le futur et ne sont pas observés.

Traditionnellement, les dommages liés au changement climatique ont été estimés à

l'aide de modèles d'évaluation intégrée (*Integrated Assessment Models* ou IAM en anglais) et présentés sous la forme du coût social du carbone (*Social Cost of Carbon* ou SCC en anglais), qui représente le dommage marginal lié à l'émission d'une tonne supplémentaire de CO₂ à un moment donné. Ce calcul implique un très grand nombre d'hypothèses de modélisation et a été fortement critiqué au motif que les dommages estimés qui en résultent sont trop faibles ou peu fiables et ne reflètent pas bien la dynamique du changement climatique (Pindyck, 2013). En effet, les estimations du coût social du carbone peuvent changer de manière significative lorsque cette dynamique est prise en compte (Taconet et al., 2021), ou lorsque les inégalités dans les impacts climatiques sont modélisées (Dennig et al., 2015). Récemment, une littérature florissante sur l'"économétrie climatique" (Hsiang, 2016) a tenté d'approcher des fonctions de dommages climatiques en estimant l'impact des changements météorologiques passés et en les projetant dans le futur. Cette littérature est également confrontée à un certain nombre de défis, tels que la prise en compte de l'adaptation ou l'incorporation des dommages liés aux événements catastrophiques (Auffhammer, 2018). Les estimations nationales de coût social du carbone, bien qu'imparfaites, peuvent néanmoins être utilisées dans les évaluations d'impact réglementaire ou pour des calculs de coin de table estimant l'impact sur le bien-être de mesures données, comme nous le faisons au chapitre 3 de cette thèse.

Le défi que représente l'estimation des dommages causés par la pollution atmosphérique est lié à la dimension spatiale de cette externalité, plutôt qu'à sa dimension temporelle. Comme expliqué ci-dessus, il ne suffit pas de savoir où la pollution est émise pour estimer la quantité de dommages. La chaîne causale entre émissions et dommages est faite d'une série complexe de relations, en quatre étapes : 1) comment les émissions se traduisent en concentrations, 2) qui est exposé aux concentrations et 3) quelle est la vulnérabilité des groupes exposés à une dose donnée 4) quelle valeur monétaire peut être attribuée à la mesure de la vulnérabilité. Les défis liés à l'estimation de l'impact causal de la pollution atmosphérique comprennent l'exposition non aléatoire à la pollution, les erreurs de mesure et le biais de variable omise (Graff Zivin and Neidell, 2013). Pour contourner ces difficultés, la riche littérature économique empirique sur la pollution de l'air s'est appuyée sur deux approches : premièrement, des expériences naturelles (par exemple, Chay and Greenstone (2003) ou Lavaine and Neidell (2017)) reliant des variations quasi-exogènes de l'activité économique à des variations de pollution de l'air et des résultats en santé ; deuxièmement, des phénomènes météorologiques générant une variation exogène dans les concentrations de pollution, dans une approche de variable instrumentale (par exemple, Deryugina et al. (2019) ou Arceo et al. (2016)). Ce faisant, cette littérature s'est attachée à estimer les dommages causés par une unité de concentration de pollution (étape 2), 3) et 4)), plutôt que les dommages causés par les émissions ou une activité économique

donnée. La littérature économique n'a pas souvent modélisé la façon dont les émissions se traduisent en concentrations (les exceptions incluent Muller and Mendelsohn (2007) ou Holland et al. (2016)), ou à quel point des secteurs donnés contribuent aux concentrations dans un scénario statu quo (plutôt que sous l'influence de chocs quasi-exogènes, comme dans le cas des expériences naturelles). Cette question a plutôt été laissée aux sciences de l'atmosphère et aux disciplines connexes, qui s'appuient sur des modèles de dispersion ou des méthodes de répartition des sources. Dans le chapitre 2 de cette thèse, nous proposons une approche reposant sur des données observationnelles pour estimer la contribution d'une activité économique donnée aux concentrations de pollution observées au niveau d'une ville.

Une fois la quantité de dommages estimée, un prix monétaire lui est attribué. Les méthodes de valorisation monétaire comprennent la mesure directe des dépenses médicales ou de médicaments induites par la pollution, ou l'application de valeurs monétaires aux impacts estimés sur la mortalité ou la morbidité, en utilisant des valeurs établies pour des indicateurs tels que la valeur d'une année de vie (VOLY) ou l'année de vie corrigée de la qualité (QALY). À partir de ces différentes étapes reliant l'émission d'un polluant et les coûts monétaires, des estimations des coûts d'émission au niveau national ont été proposées pour les études d'impact de la réglementation (voir par exemple CE Delft (2018)).

Le lien entre l'équité et les inégalités environnementales

Pour effectuer une analyse du bien-être des instruments de politiques publiques, il faut comprendre comment les coûts et les bénéfices de cette politique sont distribués. Comme le soulignent Hsiang et al. (2019), comprendre comment les dommages sont distribués dans la période précédant la mise en œuvre de la politique donne la distribution des bénéfices potentiels de la politique. De même, comprendre qui contribue aux émissions dans le contexte pré-politique donne la distribution des coûts potentiels pour les individus ciblés. Bien entendu, cela dépend des catégories d'émetteurs et - concernant la pollution de l'air - des catégories de victimes de la pollution visés par la politique. La littérature sur les inégalités environnementales peut apporter un éclairage sur ces trois dimensions.

Cette littérature distingue trois grands types d'inégalités environnementales. Un premier type d'inégalités environnementales concerne l'exposition aux dommages. Savoir qui est le plus exposé aux dommages environnementaux (en se concentrant ici sur les dommages liés au changement climatique et à la pollution atmosphérique) sert également deux objectifs : d'une part, cela permet de déduire qui est susceptible de bénéficier des politiques d'atténuation ; d'autre part, en relation avec la littérature sur la justice climatique ou la justice environnementale, cela permet de prendre une position normative sur

l'équité des questions environnementales, d'autant plus lorsque les inégalités d'exposition sont combinées à des inégalités de contribution. Une littérature croissante s'intéresse aux inégalités d'impact du changement climatique et aux inégalités d'impact de la pollution atmosphérique (Hsiang et al., 2019).

Deuxièmement, et plus proche du sujet de cette thèse, il existe des inégalités dans la contribution aux émissions. Cette question a été largement étudiée dans le cas des inégalités mondiales d'émissions de CO₂ (Chancel and Piketty, 2015) et des émissions de CO₂ au sein des pays (Sager, 2019; Ivanova and Wood, 2020). En ce qui concerne la pollution atmosphérique locale, la littérature est plus rare et s'est généralement concentrée sur les inégalités entre pays dans le contexte de l'hypothèse de la courbe de Kuznets environnementale, qui postule une relation en forme de U inversé entre le développement économique et les flux de pollution (Dinda, 2004). Les inégalités en matière d'émissions polluantes locales au niveau individuel ont été beaucoup moins étudiées (à l'exception de Levinson and O'Brien (2018)), peut-être en relation avec la limite soulignée ci-dessus : connaître la contribution individuelle aux émissions n'est pas suffisant pour déduire la responsabilité individuelle dans les dommages causés par la pollution atmosphérique, et la plupart des données sur la consommation individuelle manquent de détails spatiaux. Outre leur rôle dans la compréhension de qui peut supporter le coût des politiques de réduction, les inégalités de contribution aux émissions sont parfois abordées d'un point de vue normatif, afin d'évaluer qui devrait payer le coût des politiques d'atténuation. Cet aspect moral est particulièrement important dans le débat sur le climat et dans les négociations internationales, où le principe des "responsabilités communes mais différenciées", formalisé lors du Sommet de la Terre de Rio en 1992, identifie les pays occidentaux comme ayant une responsabilité historique dans les émissions et le devoir de payer une plus grande part des coûts d'atténuation et d'adaptation.

Un troisième type d'inégalités environnementales est directement lié aux effets des politiques d'atténuation. Ces types d'inégalités ont été mentionnés ci-dessus lors de l'examen de la régressivité ou de la progressivité des différents instruments. Le revenu a été considéré comme une dimension clé permettant d'évaluer l'incidence des taxes sur le carbone ou d'autres instruments politiques. Toutefois, pour un niveau de revenu donné, de nombreuses autres caractéristiques peuvent influencer le profil des personnes qui supporteront la charge des politiques d'atténuation. Une dimension clé est le potentiel de substituer un bien propre au bien polluant étant taxé ou interdit. Par exemple, dans le cas des politiques visant à lutter contre les voitures polluantes - qu'il s'agisse d'émissions de CO₂ ou d'émissions locale -, les caractéristiques géographiques sont un déterminant important des différences de potentiels de substitution : Gillingham and Munk-Nielsen (2019) montrent l'importance de la distance au lieu de travail et de la disponibilité des transports publics

pour expliquer l'élasticité du prix du carburant et la demande de kilomètres en voiture. Douenne (2020) souligne également que pour un décile de revenu donné, les ménages vivant dans les zones rurales ou périurbaines ont une incidence fiscale plus élevée que les ménages urbains. Cependant, son analyse suggère que l'hétérogénéité de l'incidence de la taxe carbone n'est que faiblement expliquée par les caractéristiques observables des ménages. Dans le cas des émissions de carbone provenant des transports, cela peut s'expliquer par le fait qu'au sein d'une zone donnée, il existe encore une hétérogénéité considérable dans les comportements de mobilité. Une analyse sur des données de mobilité individuelle, telle que celle conduite au chapitre 3 de cette thèse, pourrait être nécessaires pour estimer le potentiel de substitution de la voiture à un niveau plus granulaire et mieux comprendre cette hétérogénéité.

Cette thèse

Cette thèse rassemble trois travaux de recherche empiriques portant sur l'atténuation du changement climatique (chapitre 1), la pollution de l'air (chapitre 2), ou les deux (chapitre 3). Les méthodes utilisées mobilisent des données provenant de trois contextes différents, tous en Europe : le Royaume-Uni (chapitre 1), et deux villes/aires urbaines françaises, la ville de Marseille (chapitre 2), et la région parisienne (chapitre 3).

L'objectif de cette thèse est double : premièrement, estimer l'efficacité de différentes options politiques pour réduire les émissions, en examinant l'effet d'une taxe carbone existante sur les émissions de CO₂ (chapitre 1), et de différentes options pour réduire à la fois les émissions de CO₂ et de pollution locale dans le contexte du transport urbain (chapitre 3) ; deuxièmement, comprendre qui contribue à la pollution, soit au niveau sectoriel et en se concentrant sur la pollution atmosphérique locale, en examinant le cas spécifique du trafic des navires de croisière (chapitre 2), soit au niveau individuel, en étudiant les émissions induites par les déplacements à l'échelle de la ville (chapitre 3). Ci-dessous, je résume la contribution de chaque chapitre, je reviens sur les approches méthodologiques, et j'esquisse quelques recommandations de politique publiques et des pistes pour les recherches futures.

Contributions par chapitre

Le premier chapitre, intitulé *Carbon Pricing and Power sector Decarbonisation: Evidence from the UK*, examine l'efficacité d'une taxe carbone introduite dans le secteur de l'électricité au Royaume-Uni en 2013, le CPS, qui est passée de 5,9 €/tCO₂ en 2013 à 26 €/tCO₂ en 2017. J'évalue l'impact causal de le CPS sur les émissions du secteur de l'électricité au Royaume-Uni avec la méthode du contrôle synthétique. Je constate que les

émissions du secteur électrique britannique ont diminué de 20,5 à 26% par an en moyenne entre 2013 et 2017 par rapport à un Royaume-Uni synthétique composé d'une combinaison d'autres pays européens. La principale contribution de ce chapitre est d'ajouter à la rare littérature empirique évaluant l'impact des instruments régionaux et nationaux de tarification carbone. Une nouveauté de mon article est de se concentrer sur le secteur de l'électricité, très peu étudié jusqu'à présent mis à part des études sur l'effet du marché carbone européen, qui comprend à la fois des installations du secteur électrique et manufacturier. Contrairement à deux autres articles examinant également la CPS, par Abrell et al. (2019) et Gugler et al. (2020), l'utilisation de la méthode de contrôle synthétique me permet de capturer deux types d'effets non analysés auparavant : premièrement, l'effet de la politique sur les mécanismes de long terme tels que la fermeture d'usines ; deuxièmement, l'interaction de la taxe carbone avec les réglementations existant au niveau européen sur les émissions de polluants locaux d'origine industrielle. Une deuxième contribution est d'analyser la décarbonation rapide du secteur électrique britannique - très commentée dans les médias ⁶ - en estimant la contribution spécifique d'une taxe carbone à cette réduction. Enfin, ce chapitre apporte une contribution à la littérature économique environnementale en utilisant la méthode du contrôle synthétique. Il innove en combinant des données agrégées, utilisées dans la plupart des articles sur la SCM, avec des données au niveau des installations. Cela me permet d'isoler des chocs spécifiques et d'étudier les mécanismes qui conduisent à la réduction : la diminution de l'intensité d'émission des installations qui restent sur le marché, et la fermeture des installations, en distinguant les installations qui risquent déjà de fermer en raison d'autres réglementations, et les autres installations.

Le deuxième chapitre, intitulé *Estimating the Causal Effects of Cruise Traffic on Air Pollution using Randomization-Based Inference*, est un travail conjoint avec Léo Zabrocki et Marie Abèle Bind. Nous examinons la contribution du trafic maritime de bateaux de croisière aux concentrations de pollution à Marseille, l'une des plus grandes villes-port européennes. Ce chapitre a une portée interdisciplinaire et est lié à la littérature en statistique appliquée et à en science atmosphérique. Nous considérons que notre principale contribution est d'ordre méthodologique, car nous combinons une approche de *matching* par paire adaptée à l'analyse de données de séries temporelles à haute fréquence, avec une inférence basée sur la randomisation. Cette méthodologie pourrait être appliquée à une variété de contextes à la fois en économie, en science atmosphérique, en épidémiologie et dans d'autres disciplines cherchant à répondre à des questions causales. Une deuxième contribution consiste à mettre en lumière la contribution du trafic des bateaux de croisière - une activité qui devrait se développer dans les années à venir - aux concentrations de

⁶voir par exemple <https://www.ft.com/content/a05d1dd4-dddd-11e9-9743-db5a370481bc>

pollution à l'échelle d'une ville. Cette estimation est une première étape dans la quantification des dommages marginaux associés au trafic maritime, et peut aider à estimer les bénéfices des régulations mises en œuvre dans ce secteur.

Le troisième chapitre, intitulé *Tackling Transport-Induced Pollution in Cities : A Case Study in Paris*, est un travail conjoint avec Philippe Quirion. Nous examinons dans quelle mesure les individus contribuent aux émissions de CO₂ et aux polluants atmosphériques locaux dans leur mobilité quotidienne, en nous basant sur des données d'enquête détaillées de la région parisienne, l'une des villes européennes les plus polluées. Nous étudions les causes des inégalités de contribution aux émissions, et la manière dont celles-ci pourraient être réduites. Ce chapitre contribue à la littérature de plusieurs façons : tout d'abord, nous ajoutons à la vaste littérature sur les inégalités environnementales en matière de contribution aux émissions. La nouveauté de notre article est qu'il se concentre sur l'échelle urbaine - une échelle pertinente pour l'analyse des effets distributifs de politiques de transport locale - et ce avec un large échantillon représentatif de la population francilienne. Deuxièmement, nous contribuons à la littérature en transports portant sur le potentiel de réduction des émissions. Nous innovons en nous appuyant sur des données contre-factuelles de temps de trajet provenant d'une interface de programmation d'applications (API), et en examinant entre autres le potentiel du vélo électrique, un mode de transport peu étudié dans la littérature, mais dont les ventes sont en forte croissance sur la période récente. Une contribution moins centrale est de souligner que les politiques ciblant le CO₂ et les polluants locaux peuvent avoir des implications distributives différentes : les individus appartenant à des ménages à bas revenus ont tendance à utiliser des véhicules à forte intensité de pollution locale, tandis que individus appartenant à des ménages à haut revenu ont tendance à utiliser des véhicules à faible intensité de pollution, mais à forte intensité d'émission de CO₂.

Données et méthodes utilisées

Mener des recherches empiriques sur la pollution et les émissions de CO₂ soulève deux défis importants : un premier défi est spécifique aux questions étudiées et consiste à trouver des données mesurant les émissions à l'échelle appropriée. Un deuxième défi est spécifique à la recherche empirique en microéconomie appliquée et consiste à trouver un contrefactuel approprié pour répondre à des questions causales.

Données

Le résultat envisagé dans les trois chapitres consiste à imputer la pollution (soit les émissions de CO₂, soit la pollution locale) à différents secteurs (chapitres 1 et 2), à des in-

stallations industrielles (chapitre 1) ou à des individus (chapitre 3). Le principal défi de l'utilisation des données de pollution pour la recherche empirique est que la disponibilité des données dépend souvent de l'existence de réglementations sur ce polluant. En ce qui concerne les émissions de carbone, elles sont publiquement disponibles à l'échelle nationale et sectorielle dans la plupart des pays développés, suite aux initiatives de la Convention-cadre des Nations unies sur les changements climatiques et des Conférences des Parties pour harmoniser les inventaires de GES. Il est beaucoup plus difficile de trouver des données sur les émissions au niveau micro, que ce soit du côté de la production ou de la consommation. Pour le chapitre 1, je tire parti de la collecte de données mise en œuvre pour vérifier la conformité sur le marché européen du carbone (EU ETS). Cela me permet d'accéder aux données d'émissions annuelles au niveau des installations pour toutes les installations soumises à l'ETS, y compris les installations du secteur de l'électricité. Pour le chapitre 3, nous reconstruisons les émissions de gaz à effet de serre au niveau individuel, en combinant les informations individuelles sur les distances parcourues et le choix modal avec d'une part les informations sur les véhicules possédés par les ménages, et d'autre part les facteurs d'émission par type de véhicule produits par les institutions nationales ou européennes.

Pour les polluants locaux, comme nous l'avons déjà mentionné, il existe deux types de mesures : les mesures d'émission, qui ont une unité de masse (généralement des kilogrammes), reliant une source polluante à une quantité de polluant libérée dans l'atmosphère, et les concentrations, qui ont une unité de masse par volume (généralement des microgrammes par mètre cube - $\mu\text{g}/\text{m}^3$), mesurant la densité des polluants dans l'atmosphère. Comme pour le CO_2 , les données sur les émissions sont disponibles publiquement par pays et par secteur pour différents polluants dans la plupart des pays développés, alors qu'au niveau micro, elles ne sont disponibles que lorsque les réglementations existantes imposent une déclaration obligatoire (par exemple, les directives européennes sur les émissions industrielles). Dans le chapitre 3, nous estimons les émissions de polluants locaux au niveau individuel en utilisant une méthode similaire à celle utilisée pour les émissions de CO_2 , en combinant les données de l'enquête de mobilité et les facteurs d'émission.

En ce qui concerne les concentrations de pollution atmosphérique, elles ont historiquement été mesurées à l'aide de stations de mesure dispersées sur un territoire donné. Récemment, les chercheurs se sont également appuyés sur une source relativement nouvelles, les données de télédétection, qui sont notamment disponibles pour les régions et les périodes où les données des stations de surveillance sont rares ou peu fiables. Cependant, lorsque l'échelle considérée est la ville, comme dans le chapitre 2, les avantages des données d'imagerie satellite sont plus limités. Nous utilisons donc les données des stations de mesure. L'un des inconvénients de cette source de données, là encore, est qu'elle

ne surveille que les polluants qui sont réglementés au niveau national. Alors que des preuves épidémiologiques récentes suggèrent que les particules les plus nocives sont les plus petites⁷, en France, les premières réglementations de type “command-and-control” concernaient les concentrations de PM_{10} , de sorte que la plupart des stations de surveillance ne mesurent que les PM_{10} , et non les $PM_{2.5}$, et encore moins les PM_1 . La mesure des particules plus fines est très récente et concerne peu de stations. Par conséquent, les chercheurs empiriques ayant besoin d’une certaine profondeur historique ne peuvent que partiellement s’appuyer sur les stations du réseau de mesure pour évaluer l’impact des activités polluantes et des politiques d’atténuation sur la qualité de l’air.

Les autres données utilisées dans cette thèse comprennent des données météorologiques et de transport à l’échelle de la ville de Marseille, sous forme de séries temporelles à haute fréquence (chapitre 2), des données de panel sur différentes caractéristiques du secteur électrique au niveau pays pour les pays de l’UE (chapitre 1), des données en coupe sur la mobilité provenant d’une enquête (chapitre 1) et des données innovantes sur les temps de trajet contrefactuels provenant d’une API (chapitre 3).

Méthodes

Les méthodes utilisées dans les deux premiers chapitres s’inspirent de méthodes empiriques de plus en plus répandues dans la littérature économique au cours des vingt dernières années, depuis la “révolution de la crédibilité en économie empirique” (Angrist-Pischke, 2010). Nous sommes confrontés au “problème fondamental de l’inférence” (Holland, 1986a) et devons trouver un groupe de contrôle approprié aux unités traitées (installations électriques dans le chapitre 1 et heures ou journées dans le chapitre 2). Dans les deux chapitres, le cadre conceptuel (*research design*) s’inscrit dans le modèle causal de Neyman-Rubin (Rubin, 1974a), largement utilisé pour étudier les questions causales en économie depuis les années 1990 (Athey and Imbens, 2017). Dans ce modèle, la question causale est formulée en termes de résultats potentiels. Pour un traitement binaire W avec deux valeurs, 0 (pas de traitement) et 1 (traitement), chaque unité i a deux résultats : $Y_i(0)$, décrivant la valeur du résultat en l’absence de traitement, et $Y_i(1)$, décrivant la valeur du résultat en présence du traitement. Les chercheurs observent le statut de traitement de chaque unité et le résultat correspondant, mais l’un des deux résultats potentiels n’est pas observé. Dans les deux chapitres également, la solution proposée implique du *matching*, qui consiste à élaguer les données afin de ne conserver que les unités traitées et de contrôle qui semblent similaires sur la base de caractéristiques observables (voir Imbens and Wooldridge (2009) pour une revue de ces méthodes). Dans le chapitre 1, j’utilise la méthode du contrôle synthétique développée par Abadie and Gardeazabal (2003), qui est

⁷ $PM_{2.5}$ are more harmful than PM_{10} , and PM_1 are more harmful than PM_1

considérée par Athey and Imbens (2017) comme “sans doute l’innovation la plus importante dans la littérature sur l’évaluation des politiques au cours des 15 dernières années” (Athey and Imbens (2017), p9, ma traduction). La méthode SCM combine des éléments des techniques de *matching* traditionnelles, comme l’exigence d’un support commun, avec des éléments de la méthode des différences de différences, en comparant l’évolution du résultat dans le temps pour les unités traitées et les unités de contrôle. Dans le chapitre 2, nous utilisons une méthode de *matching* exact par paire, qui est proche de la nouvelle classe de méthodes de *matching* proposée par Iacus et al. (2011) et appelée méthodes de “Monotonic Imbalance Bounding” (MIB). Par rapport à la classe existante de méthodes de *matching* désignée par Iacus et al. (2011) sous le nom de méthodes “Equal Percent Bias Reducing” (EPBR) (qui comprend le *matching* par score de propension), le nombre d’unités appariées n’est pas défini ex-ante mais après le *matching*, et est fonction de la “rigueur” du *matching*.

Dans le chapitre 3, nous utilisons différentes méthodes descriptives pour répondre à nos questions de recherche : les courbes de Lorenz, qui sont l’un des moyens les plus utilisés pour représenter des inégalités de distribution (Cowell, 2011) ; une technique de décomposition qui a été largement appliquée aux données d’émissions sectorielles, l’indice Log-Mean-Divisia (LMDI) (Ang, 2004), que nous adaptons pour analyser les différences d’émissions au niveau individuel ; et l’analyse de régression.

Recommandations politiques publiques et pistes pour la recherche future

La plupart des scénarios de trajectoire d’émissions compatibles avec l’Accord de Paris imposent une baisse drastique de l’utilisation du charbon dans la production d’électricité d’ici 2030, et une élimination totale du charbon d’ici à 2050 (Jewell et al., 2019). Mon premier chapitre montre le potentiel de la taxe carbone pour aider à réaliser cette transition. Trois facteurs ont sans doute permis d’atteindre des réductions d’émissions élevées au Royaume-Uni avec des fuites limitées en dehors du Royaume-Uni - un potentiel élevé de report du charbon vers le gaz, une interconnexion limitée avec d’autres pays, et des réglementations strictes sur les polluants atmosphériques mettant en péril la rentabilité de l’électricité au charbon. Les initiatives de taxe carbone dans le secteur de l’électricité pourraient être d’autant plus efficaces si ces conditions sont réunies. Les recherches futures pourraient garder le Royaume-Uni comme étude de cas et examiner si cette élimination progressive du charbon entraîne également des bénéfices significatifs pour la santé et l’environnement, comme l’ont montré des études de modélisation ex-ante à l’échelle du monde entier (Rauner et al., 2020). Une autre piste de recherche consisterait à étudier

les facteurs d'économie politique qui ont facilité cette transition, qui semblent contraster avec d'autres contextes tels que les cas allemand, polonais ou australien.

Les conclusions de notre deuxième chapitre appellent à la prudence dans l'évaluation des secteurs prioritaires à traiter en termes de pollution locale. Bien que nous montrons que le trafic des navires de croisière a un impact significatif sur les concentrations de pollution de l'air ambiant au niveau des villes, il n'est pas clair que la lutte contre cette source de pollution doive être prioritaire par rapport à d'autres sources de pollution telles que le trafic routier. La réponse dépend également des coûts d'abattement dans ces deux secteurs, qui devraient être examinés. Les recherches futures pourraient également évaluer l'impact du trafic maritime sur la santé en utilisant des données de santé au niveau des quartiers. Une autre piste serait d'appliquer un pipeline d'inférence causale similaire à un contexte où des réglementations de type "command-and-control" imposant une diminution de la teneur en soufre du carburant des navires ont été mises en œuvre et d'estimer l'efficacité de telles politiques. Enfin, dans notre analyse, il y a eu un compromis entre la réduction du déséquilibre des données (*imbalance*) et le fait d'obtenir des coefficients estimés précis. À la lumière de ce compromis, il serait intéressant de caractériser les conditions dans lesquelles le *matching* exact des paires est plus efficace que les autres méthodes de *matching*, par exemple en utilisant des méthodes de simulation.

Dans le troisième chapitre, nous montrons que plusieurs leviers sont nécessaires pour réduire la pollution liée à la mobilité quotidienne dans les villes. Compte tenu de l'importance des déplacements de longue distance dans les émissions totales, le transfert modal ne pourrait réduire les émissions que d'un cinquième pour un réseau de transport public donné. Notre travail suggère cependant que le vélo électrique a un potentiel parmi les plus élevés comme alternative à la voiture, et devrait être particulièrement ciblé par les politiques publiques. Encourager le passage aux véhicules électriques en affectant prioritairement les subventions à l'échat aux personnes incapables de changer de mode de transport est également jugé nécessaire. Enfin, nous montrons que certains groupes professionnels semblent être plus dépendants de la voiture et pourraient être affectés de manière disproportionnée par les restrictions de circulation. Les décideurs politiques locaux doivent être conscients que si des alternatives à l'utilisation de la voiture ne sont pas proposées, ce groupe pourrait former une puissante coalition s'opposant à toute restriction de l'utilisation de la voiture. Une suite naturelle à ce travail serait d'examiner les effets distributifs des politiques de transport existant ciblant une diminution de l'utilisation des voitures polluantes, telles que les zones à faibles émissions actuellement mises en place dans plusieurs villes françaises. Les recherches futures pourraient également mener des analyses similaires sur d'autres villes, afin de déterminer les caractéristiques qui influencent les inégalités en matière d'émissions et de potentiel de report modal au niveau ville.

General Introduction

This introduction is structured in four parts: I start by presenting the two environmental issues discussed in this dissertation: climate change and outdoor air pollution. I then expose the different ways in which these two externalities could be regulated according to economic theory. An assessment of the real-world challenges in choosing policy instruments and estimating damages follows. Finally, I present the goals, methodologies and contributions of my dissertation.

Climate change and local air pollution, two major environmental issues

Damages and Trends in emissions

It is now well-established that anthropogenic climate change represents a major risk for humans and ecosystems (IPCC, 2018). Without a drastic decrease in greenhouse gas (GHG) emissions in the next decades, the increase in mean temperature and in the number of extreme weather events will threaten livelihoods, and likely increase poverty and inequalities between and within countries (Hallegatte and Rozenberg, 2017; Diffenbaugh and Burke, 2019). At the same time, countries' declared climate pledges fall short of the Paris Agreement ambitions (Rogelj et al., 2016). After a temporary reduction in 2020 because of the Covid-19 pandemic and associated social distancing measures (Le Quéré et al., 2020), global carbon dioxide emissions (CO_2 , the main greenhouse gas) have rebounded the first quarter of 2021 (IEA, 2021).

Air pollution is a second major environmental issue, identified by the World Health Organization (WHO) as the world's "largest single environmental health risk" (WHO, 2014). It is caused by the concentration of several types of pollutants, distinct from greenhouse gases, in the ambient air: the most important pollutants - which will be discussed in this dissertation - are particulate matter $\text{PM}_{2.5}$ and PM_{10} ⁸, nitrogen dioxide (NO_2), ozone (O_3) and sulphur dioxide (SO_2). Outdoor air pollution accounts for an estimated 4.2 million deaths per year worldwide (WHO, 2014), while $\text{PM}_{2.5}$ pollution alone is linked to an estimated 379,000 deaths in the European Union (EU) (European Environment Agency, 2020) and 40,000 deaths in France (Santé Publique France, 2021).

Air pollution regulations have been implemented in most high-income countries since

⁸Particles are classified by particle size. $\text{PM}_{2.5}$ only include particles smaller than $2.5\mu\text{m}$, while PM_{10} include all particles smaller than $10\mu\text{m}$.

the 1950s in the aftermath of the 1952 Greater Smog of London and the US 1948 Donora Smog, and became increasingly stringent. Since then, there have been significant decreases in pollution (see US EPA (2016) for the US, and Sicard et al. (2021) for Europe). Air pollution remains an important concern for at least four reasons: first, despite the decrease observed in high-income countries, concentrations often remain above the WHO recommended limit values for particulate matter, ozone and nitrogen dioxide. Second, pollution levels are rising in many cities from emerging and low-income countries (WHO, 2016). Third, recent evidence suggest no threshold below which exposure to pollution is safe in terms of physical health (Di et al., 2017). Fourth, the consequences of exposure to air pollution go beyond its impact on physical health, with recent evidence pointing to pollution-induced degradations of mental health, human capital and productivity (see, e.g, Roth (2017); Chang et al. (2019)).

Sectoral responsibility over pollution damages

Most of the anthropogenic pollutants released in the atmosphere come from combustion processes. The economic sectors contributing to CO₂ emissions generally also emit local air pollutants, and the other way around. To illustrate this, figure 0.0.2 shows the contribution of different sectors to CO₂, NO_x⁹ and PM_{2.5} emissions in 2018 in the European Union¹⁰. Electricity and heat production contributes 30% of total EU greenhouse gas emissions, but also 9% of the NO_x emissions, 2% of the PM_{2.5} emissions, and 20% of the SO_x¹¹. Similarly, transport contributes 23% of total EU greenhouse gas emissions, but also 66% of the NO_x emissions, 22% of the PM_{2.5} emissions, and 19% of the SO_x emissions. The two economic sectors considered in this dissertation, the electricity and heat generation sector (Chapter 1), and the transport sector (Chapters 2 and 3), together contribute 53% of the EU greenhouse gas emissions in 2018, 75% of its NO_x emissions, 39% of its SO_x emissions and 24% of its PM_{2.5} emissions.

While the above figures allow us to infer sectoral responsibility in CO₂-induced climate damages, they do not allow us to infer sectoral responsibility in air pollution-induced damages; rather, they only give a hint. This is because all CO₂ emissions contribute equally to climate change. In contrast, the link between air pollutant emissions and damages for human well-being is mediated by two steps: first, *emissions* coming from the pollution source translate into *concentrations* in the atmosphere. This process can be affected by a variety of meteorological and topological factors. For example, a higher wind speed has a cleaning effect, while thermal inversion episodes - where pollutants are

⁹NO_x stands for nitrogen oxide and is a generic category of emissions producing among others NO₂.

¹⁰including the United Kingdom, who left the EU in January 2020.

¹¹SO_x stands for sulphur oxide and is a generic category of emissions producing among others SO₂.




Figure 0.0.2 – Sectoral breakdown of greenhouse gases, NO_x and PM_{2.5} emissions in the EU 28 in 2018

Sources: ClimateWatch data by sector and country for GHG emissions, and EMEP data for local pollutant emissions by sector and country. ClimateWatch sectoral breakdown is based on the IPCC. EMEP’s sectoral breakdown is based on the GNFR 19 categorization, commonly used for air pollutants. Most GNFR sectors can easily be mapped to IPCC sectors, except for the “Other” and “Natural” categories, allocated to “Other”. The GNFR sector “Other Stationary Combustion” includes emissions from stationary plants in the residential, commercial, institutional and agricultural sectors, and is considered as the equivalent of the IPCC “Energy: Building” sector.

trapped below a warmer layer of the atmosphere - deteriorate air quality (Arceo et al., 2016). Second, pollution concentrations translate into damages depending on how much the local population is exposed to such pollution (high pollution levels in a rural area with a low population density will cause less harm to humans than the same pollution levels in a densely populated city), and how vulnerable to pollution she is. The second chapter of this thesis speaks to this non-trivial relationship between emissions and potential health damages.

The economic approach to pollution

Pollution as an externality to be regulated

Economists have long identified pollution as a market failure that needs to be fixed. Pollution in a broad sense, including water, soil or air contaminants, is considered as a negative externality (Pigou, 1920): it is a by-product of the production or consumption of a first party which makes another party worse off. Yet, the first party does not bear the cost of the harm done because it is not reflected in market prices. As a consequence, there is too much provision of polluting goods compared to what would be optimal. In neo-classical economics, externalities are one of three identified market failures causing inefficiency and justifying government intervention. What are the best policy instruments to reduce the pollution externality has been a key question in environmental economics for several decades (Bretschger and Pittel, 2019).

Judging the merits of different regulations

There are broadly two types of public policies to regulate pollution: command-and-control regulations, which consist in setting standards on the type of technology to be used (in the case of technology mandates) or on the maximum level of pollution (in the case of performance standards), and monitoring enforcement; and market-based or incentive-based solutions, which consist in putting a price on pollution to incentivize polluters to reduce emissions. Market-based solutions can be further distinguished between those regulating the quantity of emissions (such as cap-and-trade markets setting a cap on the number of emission permits in circulation) and those regulating their price (such as Pigouvian taxes).

The merits of different regulations can be judged based on different criteria. Goulder and Parry (2008) cite economic efficiency, cost-effectiveness, the distribution of benefits and costs, the ability to address uncertainties and political feasibility as potential evaluation criteria.

Regarding efficiency, theoretical evidence suggests that market-based solutions are more cost-effective than command-and-control regulations in the first best: this is because polluters generally differ in their ability to reduce pollution - measured by their marginal abatement cost - (Newell and Stavins, 2003), and only incentive-based instruments allow to equate marginal costs and marginal benefits across all emitters if polluters know their abatement costs but the policy-maker does not. Another theoretical advantage of market-based instruments is that they exploit all pollution reduction channels, while command-and-control regulations neglect some, in particular output reduction

(Spulber, 1985). Finally, incentive-based instruments generally generate revenues (except when emission permits are freely distributed), which, if appropriately recycled, can reduce other distortionary taxes and create efficiency gains. This recycling could generate a “double dividend”, both improving environmental quality and reducing the net welfare cost of environmental policy (Baumol and Oates, 1988; Pearce, 1991; Chiroleu-Assouline and Fodha, 2014).

However, in the presence of multiple market failures, the superiority of market-based instruments can be questioned, and a combination of different instruments may be appropriate, including market-based, standards but also public subsidies. For example when the administrative costs of monitoring emissions is high, mandates may be superior (Goulder and Parry, 2008). Goulder et al. (2016) also show that in the presence of pre-existing factor market distortions, clean energy standards are more cost-effective than price-based instruments because they represent a smaller implicit tax on factors of production. Public support to innovation in low-carbon technologies, beside carbon pricing, is also justified by the public good nature of innovation: in a setting where the gains from innovation cannot be fully appropriated, investment is too low without public support (Fischer and Newell, 2008). The road transport sector illustrates well the multiplicity of market failures and the need for second-best policies, since road transport contributes to multiple externalities: CO₂ emissions, local air pollution, but also noise, accidents, and congestion (Parry et al., 2007).

As for the choice between price and quantity instrument, Weitzman (1974) suggests that the relative merits of price vs. quantity instruments depend on how steep the marginal damage function is, when there is uncertainty on the aggregate costs of pollution reduction - which is often the case in the real world. Price instruments are superior when the marginal damage curve is flat and quantity instruments are superior when the marginal damage curve is steep. The intuition behind this result is that it is all the more important to get the quantity right when damage costs increase a lot for a small change in pollution. Because the marginal damage curve of climate change is rather flat, economists have favoured a carbon tax over a carbon market on these theoretical grounds. At the global level, Weitzman (2015) also argues in favor of a uniform carbon tax rather than internationally tradeable permits, because of uncertainty regarding country-specific abatement cost profiles.

Regarding equity, distributional concerns may arise because some individuals bear a disproportionate cost of the regulation. For example, carbon taxes have often been found to be regressive in high-income countries if the receipts of the tax are not redistributed (Poterba, 1991). This is because poorer households dedicate a higher share of their consumption expenditures to carbon-intensive goods. Studies have found that carbon taxes

can become progressive if tax receipts are redistributed in the form of lump-sum transfers to all households (Metcalf, 2009a; Cronin et al., 2018; Douenne, 2020; Berry, 2019). However, lump sum transfers are unlikely to correct horizontal equity issues, because of the high degree of heterogeneity in tax incidence within income deciles (Sallee, 2019; Douenne, 2020; Berry, 2019). Standards have been found to be more regressive than a carbon tax with lump-sum transfers in some cases (Davis and Knittel, 2018; Levinson, 2018), while Zhao and Mattauch (2020) show that standards are more equitable when consumers exhibit a preference for high-carbon technology attributes - and they verify that this is the case in the US.

Political support and feasibility will likely be affected by the objective distributional properties of the proposed regulations. Beyond these objective properties, however, recent research highlights the importance of the perception of policies (Douenne and Fabre, 2021; Maestre-Andres et al., 2019), and of contextual factors such as political trust (Rafaty, 2018).

Reviewing evidence on the effects of different policy instruments, Goulder and Parry (2008) conclude that *"No single instrument is clearly superior along all the dimensions relevant to policy choice; even the ranking along a single dimension often depends on the circumstances involved."* (Goulder and Parry (2008), p 153). Since these "circumstances" are likely to vary from one context to the other, this conclusion calls for more empirical evidence on the effects of actual regulations implemented across the world. The first chapter of this thesis takes this recommendation to heart and provides evidence on the effect of a carbon tax implemented in the UK power sector.

Regulating air pollution and CO₂ emissions in practice

A variety of instruments

In practice, local pollutants and CO₂ emissions have been regulated via a combination of instruments. Command-and-control instruments have historically been more common and have consisted in standards with specific requirements to use the best available technology (BAT) or other specific technology mandates (Metcalf, 2009b). Nowadays, different types of regulations co-exist depending on the jurisdictions and on the pollutant: local pollutants from industrial installations (including the power sector) are regulated by standards in the European Union (Directives on industrial emissions) and in the US (Clean Air Act); by cap-and-trade programmes in some US states (California Regional Clean Air Incentives Market (RECLAIM)); by taxes in some European countries (the TGAP in France or the NO_x fees in Sweden or Norway (Bonilla et al., 2018)). Carbon emissions are also

regulated via a combination of command-and-control regulations, for example in the form of fuel efficiency standards (such as the US CAFE) or emission standards (in the European Union) in the transport sector, and market-based instruments, for example the European carbon market (EU ETS) or the Chinese ETS. Market-based instruments have increased in the past decade, with 64 carbon pricing instruments across the world in 2021 against only 21 ten years before in 2011 (World Bank, 2021).

Evidence on the effectiveness of existing policies

The effectiveness of several of the above-mentioned regulations has been estimated empirically (e.g, Fowlie et al. (2012) for RECLAIM; Colmer et al. (2020) for the EU ETS; Currie and Walker (2019) for the Clean Air Act; Andersson (2019) for the Swedish carbon tax). However, such evaluations are relatively scarce, in particular concerning carbon pricing instruments and their effectiveness. One reason is simply that these instruments are rather recent. Another reason is methodological: beyond the amount of data required, causally estimating the impact of a carbon pricing on emissions requires overcoming the “fundamental problem of causal inference” (Holland, 1986a), whereby in countries with a policy in place, we only observe the evolution of pollution in the presence of the policy, but not in the counterfactual situation where the policy is absent. This problem is all the more pronounced in the case of carbon pricing instruments, typically targeting an entire sector or region. Of the 21 empirical evaluations of existing carbon prices listed in Rafaty et al. (2020), four evaluations are on the manufacturing sector, five on the transport sector alone, six on both the power and manufacturing sectors¹² and three pool several country-level instruments on different sectors. Only three papers (including the first chapter of this thesis) focus on the power sector alone and they all consider the same instrument. Thus, more evidence is needed on the effectiveness of carbon pricing in general, but also in the power sector in particular, given its high contribution to worldwide emissions. In addition, there is a wide heterogeneity in estimates of carbon price effectiveness (Rafaty et al., 2020; Green, 2021). This calls for more systematic evidence on the factors explaining such differences, whether these are due to contextual factors or methodological differences in the estimation strategy and scope considered.

Accounting for two externalities: co-benefits and trade-offs

The fact that some sectors emit both local and global pollution may imply that regulations tackling one type of pollutant will also affect the other. Whether an increased stringency of mitigation policies leads to a decrease in local air pollution depends on the elasticity of

¹²papers evaluating the EU ETS

substitution between the two (Ambec and Coria, 2013). If local pollutants are substitute to CO₂ emissions, abating CO₂ will increase the emission of local pollutants. If they are complements, abating CO₂ will also reduce local pollutant emissions.

In many cases, local and global pollutants have been found to be complements: in this case, climate policy reducing CO₂ emissions will also improve air quality and may generate health benefits. These indirect benefits are part of broader positive side-effects of climate policy, grouped under the concept of co-benefits (Nemet et al., 2010; Intergovernmental Panel on Climate Change, 2015). Prior studies suggest that air quality co-benefits from mitigation policies can represent a substantial share of the monetized benefits and equal or even exceed mitigation costs (Karlsson et al., 2020; Rauner et al., 2020).

Cases where local pollutants are substitute to CO₂ emissions have also been identified: this is in particular the case in the car transport sector, where incentivizing a reduction in the CO₂ emission intensity of cars risks boosting the market share of diesel cars, which emit more local pollutants (Durrmeyer, 2018; Linn, 2019). This implies that a policy designed to internalise one of the externalities may unintentionally increase another, which complicates regulation.

Only a few studies have empirically examined the implications of co-benefits for policy design. Wagner and De Preux (2016) investigate the potential efficiency and equity loss from emission trading, when it gives rise to implicit local pollution trading across areas with heterogeneous marginal damages from air pollution. Durrmeyer (2018) shows that the distributional properties of a feebate - a combination of fees and rebates schemes - tackling CO₂ emissions could be enhanced by taking into account its impact on the emission of local pollutants at the city-level. In the road transport sector, different individual preferences for technology attributes with a high carbon vs. high local pollution intensity may affect these distributional properties, but this question remains, to my knowledge, largely unexplored.

Estimating damages

Theoretical results on optimal regulation depend on parameters such as the shape of the marginal abatement cost curve and the marginal damage curve. However, estimating damages induced by externalities is particularly challenging. Economists would usually rely on market prices to conduct welfare analysis. But, by definition, externalities are not priced. This is why an important part of the economic literature on climate change and air pollution is devoted to estimating the damages of these externalities. Estimating the damages from climate change is more challenging than those from air pollution, because climate damages are mostly going to occur in the future and are not observed.

Traditionally, climate change damages have been estimated with Integrated-Assessment-

Models (IAM) and presented in the form of the Social Cost of Carbon (SCC), which represents the marginal damage of emitting an additional ton of CO₂ at a certain point in time. This calculation involves a very large number of modelling assumptions and has been heavily criticized on the grounds that the resulting estimated damages are too low or unreliable and do not reflect well the dynamics of climate change (Pindyck, 2013). Indeed, estimates of the SCC can change significantly when these dynamics are accounted for (Taconet et al., 2021), or when inequalities in climate impacts are modelled (Dennig et al., 2015). Recently, a burgeoning "climate econometrics" (Hsiang, 2016) literature has sought to recover climate damage functions by estimating the impact from past changes in weather and projecting them into the future. This literature also faces challenges, such as accounting for adaptation or incorporating damages from catastrophic events (Auffhammer, 2018). National estimates of SCCs, although imperfect, can still be used in regulatory impact assessments or for back-of-the-envelope computations of the welfare impacts of policies, as we do in the third chapter of this thesis.

The challenge of estimating air pollution damages has to do with the spatial dimension of this externality, rather than its time dimension. As explained above, knowing where pollution is emitted is not sufficient to estimate the quantity of damages. A complex series of relationships make up the causal chain between emissions and damages. Its building blocks are fourfold: 1) how emissions translate into concentrations, 2) who is exposed to concentrations and 3) how vulnerable exposed groups are to a given dose 4) what monetary value can be put on the vulnerability measure.

Challenges in estimating the causal impact of air pollution include non-random exposure to pollution, measurement errors and omitted variable bias (Graff Zivin and Neidell, 2013). To circumvent these challenges, the rich empirical economic literature on air pollution has relied on two approaches: first, natural experiments (e.g, Chay and Greenstone (2003) or Lavaine and Neidell (2017)) linking quasi-exogenous variations in economic activity to variations in air pollution and health outcomes; second, meteorological phenomena generating exogenous variation in pollution concentrations, in an instrumental variable approach (e.g, Deryugina et al. (2019), Arceo et al. (2016)). Doing so, this literature has focused on estimating the damages caused by one unit of pollution concentration (step 2, 3) and 4)), rather than the damages caused by emissions or one given economic activity. The economic literature has not often modelled how emissions translate into concentrations (exceptions include Muller and Mendelsohn (2007) or Holland et al. (2016)), or how given sectors contribute to concentrations in a business-as-usual scenario - rather than under quasi-exogenous shocks like in natural experiments. This question has rather been left to atmospheric science and related disciplines, which rely on dispersion models or source apportionment methods. In the second chapter of this thesis, we suggest an

approach relying on observational data to estimate the contribution of a given economic activity to city-level pollution concentrations.

Once estimated, the quantity of damages linked to emissions are given a monetary price. Valuation methods include directly measuring medical or drug expenditures induced by pollution, or applying monetary values to estimated mortality or morbidity impacts, using established values for indicators such as the Value of a Life Year (VOLY) or the Quality-Adjusted Life Year (QALY). Combining the different steps linking the emissions of a pollutant and monetary costs, country-level estimates of emission costs have been proposed for regulatory impact assessments (see e.g, CE Delft (2018)).

The link between equity and environmental inequalities

Conducting a welfare analysis of policy instruments requires to evaluate the distribution of costs and benefits. As underlined in Hsiang et al. (2019), understanding how damages are distributed in the pre policy period yields the distribution of potential benefits from the policy. Similarly, understanding who contributes to emissions in the pre-policy setting yields the distribution of potential costs borne by the individuals to whom the policy applies. Of course, this depends on which groups of emitters and - for local air pollution - which group of pollution victims are targeted. The literature on existing environmental inequalities speaks to these three dimensions.

This literature distinguishes between three main types of environmental inequalities. A first type of environmental inequalities is in the exposure to the damages. Knowing who is most exposed to the environmental damages (focusing here on climate change and air pollution damages) also serves two purposes: on the one hand, it allows to infer who is likely to benefit from mitigation policies; on the other hand, in relation to the climate justice or environmental justice literature, it allows to make a normative stance on the fairness of environmental issues, all the more when inequalities in exposure are combined with inequalities in contribution. A growing literature has been focusing on inequalities in climate change impacts and inequalities in air pollution impacts (Hsiang et al., 2019).

Second, and closer to the topics studied in this dissertation, there are inequalities in contributions to emissions. This question has been extensively studied in the case of global carbon inequalities (Chancel and Piketty, 2015) and within-country carbon emissions (Sager, 2019; Ivanova and Wood, 2020). Regarding local air pollution, the literature is scarcer and has typically focused on between-country inequalities in the context of the Environmental Kuznets Curve hypothesis, positing an inverted U-shaped relationship between economic development and pollution flows (Dinda, 2004). Inequalities in local pollution emissions at the individual level have been far less studied (with the exception of Levinson and O'Brien (2018)), perhaps in relation to the limitation pointed above: know-

ing individual contribution to emissions is not sufficient to infer individual responsibility in air pollution damages, and most data on individual consumption lack spatial details. Beside their role in helping understanding who may bear the cost of abatement policies, inequalities in contribution to emissions are sometimes framed in a normative framework of responsibility, to assess who *should pay* the burden of climate policies. This moral aspect is particularly important in the climate debate and in international negotiations, where the “common but differentiated responsibilities” principle, formalized in the 1992 Rio Earth Summit, identifies Western countries as having a historical responsibility in emissions and a duty to pay a larger share of mitigation and adaptation costs.

A third type of environmental inequalities relate directly to the effects of mitigation policies. These type of inequalities have been mentioned above when discussing the regressivity or progressivity of different instruments. Income has been examined as one key dimension against which to evaluate the incidence of carbon taxes or other policy instruments. For a given level of income, however, many other characteristics may influence who bears the burden of mitigation policies. A particularly key dimension in this respect is the potential to substitute away from the pollution-intensive good being taxed or banned. For example, in the case of policies tackling pollution-intensive cars - whether pollution means CO₂ emission or local pollutants emissions -, geographic characteristics are an important driver of differences in substitution potentials: Gillingham and Munk-Nielsen (2019) show the importance of distance to workplace and public transport availability in explaining fuel price elasticity and demand for driving. Douenne (2020) also underlines that for a given income decile, households living in rural or peri-urban areas have a higher tax burden than urban households. His analysis suggests that the heterogeneity of the carbon tax incidence is only poorly explained by observable household characteristics, though. For the case of carbon emissions from transport, this may be because within a given area, there is still considerable heterogeneity in mobility behaviours. An analysis using individual mobility data, such as what we do in the third chapter of this thesis, may be required to estimate the potential to substitute away from cars at a more granular level and better understand this heterogeneity.

This dissertation

This dissertation collects three empirical research papers focusing on climate change mitigation (Chapter 1), air pollution (Chapter 2), or both (Chapter 3). The methods used mobilize data from three different contexts, all in Europe: the UK (Chapter 1), and two French cities/urban areas, the city of Marseille (Chapter 2), and the Paris area (Chapter 3).

The goal of this dissertation is twofold: first, estimating the effectiveness of different policy options to abate emissions, by examining the effect of an existing carbon tax on CO₂ emissions (Chapter 1), and of different options to reduce both CO₂ and local pollution emissions in the context of urban transport (Chapter 3); second, understanding who contributes to pollution, either at the sectoral level and focusing on local air pollution, by examining the specific case of cruise vessel traffic (Chapter 2), or at the individual level, by investigating travel-induced emissions at the city scale (Chapter 3). Below, I summarize the contribution of each chapter, I come back on the methodological approaches, and I outline some policy recommendations and fruitful avenues for future research.

Contributions by chapter

The first chapter, *Carbon Pricing and Power sector Decarbonisation: Evidence from the UK* examines the effectiveness of a carbon tax introduced in the UK power sector in 2013, the CPS, which increased from €5.9/tCO₂ in 2013 to €26/tCO₂ in 2017. I evaluate the causal impact of the CPS on power sector emissions in the UK with the Synthetic control method. I find that emissions from the UK power sector declined by between 20.5 and 26 percent on an average year between 2013 and 2017 compared to a Synthetic UK made of a combination of other European countries. The main contribution of this chapter is to add to the scarce empirical literature evaluating the impact of existing regional and national carbon pricing instruments. A novelty of my paper is to focus on the power sector, little studied so far except for few studies examining the impact of the EU ETS, which includes plants from the power and manufacturing sectors. In contrast to two other papers also examining the CPS, by Abrell et al. (2019) and Gugler et al. (2020), the use of the Synthetic method allows me to capture two types of effects not analysed before: first, the effect of the policy on long-term mechanisms such as plant closure; second, the interaction of the carbon tax with existing air pollution regulations at the European level. A second contribution is to analyse the rapid decarbonisation of the UK power sector - much commented in the media¹³ - by estimating the specific contribution of a carbon tax to this decarbonisation. Finally, this chapter adds to the environmental economics literature using the synthetic control method. It innovates by combining aggregate data, used in most SCM papers, with installation-level data. This enables me to isolate specific shocks and investigate the mechanisms leading to abatement: the decrease in the emission-intensity of plants staying in the market, and plant closure, distinguishing between plants already at risk of closure due to other regulations, and other plants.

The second chapter, *Estimating the Causal Effects of Cruise Traffic on Air Pollution using Randomization-Based Inference*, is a joint work with Léo Zabrocki and Marie

¹³see for example <https://www.ft.com/content/a05d1dd4-dddd-11e9-9743-db5a370481bc>

Abèle Bind. We examine the contribution of cruise vessel traffic to ambient air pollution concentrations in Marseille, one of the largest European port cities. This chapter has an interdisciplinary focus and is related to the applied statistics literature and to the atmospheric science literature. We see our primary contribution as methodological, as we combine a pair-matching approach suited to the analysis of high-frequency time series data, with randomization-based inference. This methodology could be applied to a variety of contexts both in economics, atmospheric science, epidemiology, and other disciplines seeking to answer causal questions. A second contribution is to shed light on the contribution of cruise vessel traffic - an activity expected to grow in the coming years - to city-wide pollution concentrations. This estimation is a first step in quantifying the marginal damages associated with vessel traffic and can help estimate the benefits of regulations implemented in this sector.

The third chapter, *Tackling Transport-Induced Pollution in Cities: A Case Study in Paris*, is a joint work with Philippe Quirion. We examine how much individuals contribute to CO₂ and local air pollutants in their daily mobility based on detailed survey data from the Paris area, one of the most polluted European cities. We investigate what drives inequalities in emissions and how emissions could be reduced. This chapter contributes to the literature in several ways: first, we add to the large literature on environmental inequalities in contribution to emissions. A novelty of our paper is to focus on the urban scale - a valid scale to examine the distributional impacts of local environmental and transport policies - with a large representative sample of residents. Second, we contribute to the transport policy literature examining the potential for emission reductions. We innovate by relying on counterfactual travel time data from an Application Programming Interface (API), and by examining among others the potential of electric bike, an under-investigated mode experiencing an important growth in sales in the recent period. A less central contribution is to highlight that policies tackling CO₂ and local pollutants may have different distributional implications: individuals from low-income households tend to use vehicles with a high pollution intensity, while individuals from high-income households tend to use vehicles with a low pollution intensity but a high CO₂ emission intensity.

Data and Methods used

Conducting empirical research on pollution and CO₂ emissions raises two important challenges: a first challenge is specific to the questions studied and consists in finding data measuring emissions at the appropriate scale. A second challenge is specific to empirical research in applied microeconomics and consists in finding a suitable counterfactual to answer causal questions.

Data

The outcome considered in all three chapters involves imputing pollution (either CO₂ emissions or local pollution) to different sectors (Chapters 1 and 2), industrial installations (Chapter 1) or individuals (Chapter 3). The main challenge in using pollution data for empirical research is that data availability often hinges upon having regulations on this pollutant. Regarding carbon emissions, they are publicly available at the national and sectoral scale in most developed countries, following initiatives from the UNFCCC and the Conference of the Parties to harmonize GHG inventories. It is much more difficult to find emission data at the micro-level, whether on the production or on the consumption side. For Chapter 1, I take advantage of the data collection implemented to verify compliance on the European carbon market (EU ETS). This gives me access to annual installation-level emission data for all the installations subject to the ETS, including power sector installations. For Chapter 3, we reconstruct individual-level greenhouse gas emissions by combining individual information on distances travelled and modal choice with household information on the vehicles owned and emission factors by vehicle type produced by national or European institutions.

For local pollutants, as mentioned earlier there are two types of measures: emission measures, having a mass unit (typically kilograms), connecting a polluting source with a quantity of pollutant released in the atmosphere; and concentrations, having a mass-per-volume unit (typically micrograms per cubic meter - $\mu\text{g}/\text{m}^3$), measuring the density of pollutants in the atmosphere. Like for CO₂, emission data are publicly available by country and sector for different pollutants in most developed countries, while at the micro-level, they only exist when regulations impose a mandatory reporting (e.g, European directives on industrial emissions). In Chapter 3, we estimate local pollutant emissions at the individual level using a similar method to that used for CO₂ emissions, combining mobility survey data and emission factors.

Regarding air pollution concentrations, they have historically been measured with monitoring stations scattered across a given territory. Recently, scholars have also relied on a relatively new data source, remote sensing data, which are available for regions and time periods where monitoring station data are scarce or unreliable. However, when the scale considered is the city such as in chapter 2, the advantages of satellite data are more limited. We therefore use monitoring station data. One disadvantage of this data source, again, is that it only monitors pollutants which are regulated at the national level. While recent epidemiological evidence suggests that the most harmful particles are the smaller ones¹⁴, in France the first command-and-control regulations were on PM₁₀ concentrations, such that most monitoring stations only measure PM₁₀, and not PM_{2.5},

¹⁴PM_{2.5} are more harmful than PM₁₀, and PM₁ are more harmful than PM₁₀

let alone PM_1 . The measure of these finer particles is very recent in France and only few monitoring stations do it. Therefore, empirical researchers can only partially rely on monitoring stations to evaluate the impact of polluting activities and mitigation policies on air quality.

The other data used in this dissertation include high-frequency time series data of transport activity and weather in Marseille (Chapter 2), a panel of country-level power sector characteristics for EU countries (Chapter 1), cross-sectional survey data on mobility (Chapter 1), and innovative counterfactual travel times data from an API (Chapter 3).

Methods

The methods used in the first two chapters draw on empirical methods which have been increasingly used in economics in the past two decades, since the “Credibility revolution in empirical economics” (Angrist and Pischke, 2010). We are facing the “fundamental problem of inference” (Holland, 1986a) and need to find a suitable control group to the treated units (power installations in Chapter 1 and hours or days in Chapter 2). In both chapters, the research design is embedded in the Neyman-Rubin Causal Model (Rubin, 1974a), the framework widely used to study causal questions in economics since the 1990s (Athey and Imbens, 2017). In this framework, the causal question is framed in terms of potential outcomes. For a binary treatment W with two values, 0 (no treatment) and 1 (treatment), each unit i has two potential outcomes: $Y_i(0)$, describing the outcome value under no treatment, and $Y_i(1)$, describing the outcome value under treatment. Researchers observe the treatment status of each unit and the corresponding outcome, but one of the two potential outcomes is not observed. In both chapters as well, the proposed solution involves some matching, which consists in pruning the data in order to only keep the treated and control units that look similar based on observable characteristics (see Imbens and Wooldridge (2009) for a review of those methods). In Chapter 1, I use the synthetic control method developed by Abadie and Gardeazabal (2003), which is considered by Athey and Imbens (2017) as “*arguably the most important innovation in the policy evaluation literature in the last 15 years*” (Athey and Imbens (2017), p9). The SCM method combines elements of traditional matching techniques, such as a requirement for common support, with elements from the difference-in-difference method, by comparing the evolution of the outcome over time for treated and control units. In Chapter 2, we use an exact pair matching method, which is close to the new class of matching methods proposed in Iacus et al. (2011) and called “Monotonic Imbalance Bounding” (MIB) methods. Compared to the existing class of matching methods referred to by Iacus et al. (2011) as “Equal Percent Bias Reducing” (EPBR) methods (which includes propensity score matching), the number of matched units is not defined ex ante but after

the matching, and is a function of the “strictness” of the matching.

In both chapters 1 and 2, inference relies on permutation tests. This approach can be linked to design-based (also called randomization-based) inference, one of two types of inference outlined in (Abadie et al., 2020). In design-based inference, the source of uncertainty comes from the fact that we only observe one of the two potential outcomes. In the other type of inference, sampling-based inference, more traditionally used in applied economics, the source of uncertainty comes from the fact that only a sample of the population is observed. In both chapters, design-based inference is more relevant given the nature of the research design: in chapter 1, the unit of analysis is the entire country and it is hard to think of the data as a random sample from a broader population. In chapter 2, we construct the experiment in such a way that treatment allocation across pairs can be considered as random. Concretely, permutation tests consist in randomly re-allocating the binary treatment across the different units a large number of times. The observed treatment effect is then compared to the distribution of treatment effects obtained under the set of random permutations.

In Chapter 3, we use different descriptive methods to answer our research questions: Lorenz curves, which are one of the most famous way of representing inequalities in distribution (Cowell, 2011); a decomposition technique which has been widely applied to sectoral emission data, the Log-Mean-Divisia Index (Ang, 2004), which we adapt to analyse differences in emissions at the individual level; and regression analysis.

Policy recommendations and paths for future research

Most emission pathway scenarios consistent with the Paris Agreement impose drastic declines in coal power by 2030 and full phase-out by 2050 (Jewell et al., 2019). My first chapter shows the potential of carbon pricing in helping to achieve this transition. Three factors arguably enabled to achieve high emission reductions in the UK with limited leakage outside the UK - a high fuel switching potential, a limited interconnection with other countries, and stringent air pollution regulations making coal use uneconomical. Carbon pricing initiatives in the power sector could be all the more successful if these conditions are met. Future research could keep the UK as a case study and examine whether such coal phase-out also entails significant health and environmental benefits, as found in ex ante modelling studies on the entire world (Rauner et al., 2020). Another research avenue would consist in investigating the political economy factors that facilitated this transition, which seem to contrast with other contexts such as the German, Polish or Australian cases.

The conclusions from our second chapter call for caution in assessing the priority sectors to tackle in terms of local pollution. Although we show that cruise vessel traffic

has a significant impact on city-level ambient air pollution concentrations, it is unclear whether tackling this source of pollution should be a priority over other sources of pollution such as road traffic. The answer also depends on the abatement costs in these two sectors, which should be examined. Future research could also evaluate the impact of maritime traffic on health using neighbourhood-level health data. Another avenue would be to apply a similar causal inference pipeline to a context where command-and-control regulations mandating a decrease in the sulphur content of vessel fuel has been implemented and estimate the effectiveness of such policies. Finally, in our analysis there was a trade-off between reducing data imbalance and obtaining precise estimates. In light of this trade-off, it would be interesting to characterize the conditions under which exact pair matching fares better than other matching methods using simulation methods.

In the third chapter, we show that several levers are necessary to reduce pollution from daily mobility in cities. Given the importance of high-distance trips in total emissions, modal shift could only reduce emissions by a fifth for a given public transit network. Our work however suggests that electric cycling has among the highest potential as an alternative to car, and should be targeted by public policies. Encouraging a shift to electric vehicles (EVs) by allocating EV adoption subsidies in priority to those unable to shift modes is also deemed necessary. Finally, we show that some professional groups appear to be more reliant on cars and could be disproportionately affected by driving restrictions. Local policy-makers should be aware that if alternatives to car use are not offered, this group could form a powerful coalition opposing any restriction in car use. A natural follow-up of this work would be to examine the distributional impacts of existing transport policies tackling car use, such as the low-emission-zones being currently rolled out in several French cities. Future research could also run similar analyses on other cities to characterize the city-level factors influencing inequalities in emissions and modal shift potential.

Chapter 1

Carbon Pricing and Power sector Decarbonisation: Evidence from the UK

Abstract: Decreasing greenhouse gas emissions from electricity generation is crucial to tackle climate change. Empirically, however, little is known about the effectiveness of existing economic instruments in the power sector. This paper examines the impact of the UK Carbon Price Support (CPS), a carbon tax implemented in the UK power sector in 2013. Relative to a synthetic control unit built from other European countries, emissions from the UK power sector declined by 26 percent on an average year between 2013 and 2017. Bounds on the effects of potential confounding policies in the UK and several placebo tests suggest that the carbon tax caused at least 80% of this decrease. Three mechanisms are highlighted: a decrease in emissions at the intensive margin; the closure of some high-emission plants at the extensive margin; and a higher probability of closure than in the synthetic UK for plants at risk of closure due to European air quality regulations. This paper shows that a carbon tax on electricity generation can lead to successful decarbonisation.

Acknowledgements: I am grateful to Philippe Quirion, Katheline Schubert, Nicolas Koch, Ulrich Wagner, Mirjam Kosch, Jan Abrell, Francois Libois, and three anonymous referees for their comments and suggestions. I thank seminar participants at PSE, LSE GRI, MCC, PIK and SSE, participants to the 2018 OECD environmental micro-data workshop, conference participants at Mannheim Energy conference, EAERE, FAERE, and participants to the Marseille Green Econ Spring School and CIRED summer school for useful feedback. I thank Ember (formerly Sandbag) for sharing their ETS data with me, Lorenzo Montrone for giving me access to the Global Coal Plant Tracker database, and Jan Abrell and Mirjam Kosch for sharing data on monthly power production.

1.1 Introduction

Every country in the world must reduce their greenhouse gas emissions in order to mitigate climate change. In the past two decades, virtually all governments have implemented a variety of abatement policies to address this challenge, including economic instruments in the form of carbon taxes and carbon markets (World Bank and Ecofys, 2018). Although carbon pricing is widely regarded by economists as the most cost-effective way to reduce emissions, ex-post evaluations of carbon pricing policies implemented in different sectors are still scarce (Green, 2021). This general observation is particularly true in the case of the power sector (Martin et al., 2016), which represented 25% of worldwide emissions in 2010 (IPCC, 2014).

In this paper, I examine the impact of a carbon tax introduced in the UK power sector in 2013, the Carbon Price Support (CPS), on carbon emissions. At that time and during the period of analysis considered in the paper, the UK was part of the European carbon market (European Union Emission Trading System, EU ETS) implemented in 2005. The carbon tax was introduced in response to the low prices prevailing on the European carbon market, while the UK was facing binding emission reduction targets under the 2008 Climate Change Act. The tax rate increased from around £5 (€5.9) per ton of equivalent carbon dioxide (hereafter tCO₂e) in 2013 to £18 (€26) in 2017. During the same period, the UK power sector experienced a remarkable transition: between 2012 and 2017, the share of coal in electricity generation decreased from 40% to 7%, gross electricity consumption decreased by 6%, and power sector greenhouse gas emissions decreased by 57% (Source: Eurostat). The rapid transformation of the UK power sector received significant coverage in the media and in policy reports (Evans, 2019; Brown, 2017), but how much the UK carbon tax contributed to such transformation is to date unclear.

To estimate this contribution, I apply the synthetic control method (Abadie and Gardeazabal, 2003; Abadie et al., 2010, 2015) to build a counterfactual UK with a weighted combination of European countries having power sectors with characteristic similar to the UK. I use countries which, like the UK, were part of the European Union during the period considered (2005-2017) as potential candidates to enter the counterfactual UK, as all these countries were subject to the same European climate and energy policies as the UK had been before the introduction of the CPS, in particular the EU ETS and European air quality regulations.

I estimate that the introduction of the CPS is associated with emissions reductions - or abatement - of between 141 and 191 million tons of equivalent carbon dioxide (hereafter MtCO₂e) over the 2013-2017 period, implying emission reductions of between 20.5% and 26% on an average year. This range depends on the assumed effect for three UK-specific

policies implemented around the same period - a subsidy to encourage the biomass conversion of coal plants, a new scheme for renewable support, a capacity market - and on the magnitude of CPS-induced spillovers. The upper bound assumes that biomass conversion is a consequence of the CPS, that the other two policies have a negligible impact over the 2013-2017 period, and that emission leakage from the UK to other European countries was negligible. The lower bound is more conservative: it puts a bound on the effect of the three UK policies and calculates the level of emissions from the synthetic UK which may be due to CPS-induced spillovers. Back-of-the-envelope calculations suggest that roughly a third of the lower bound impact was driven by UK plants facing a high carbon price responding differently to European air quality regulation. Another third was caused by the closure of a several high-emitting plants and the last third by a decrease in emissions from plants that remained in the market (likely due to fuel switching from coal to gas). A set of placebo tests suggest that this impact is causal, and these results are robust to several sensitivity analyses.

This paper contributes to several strands of the literature: first, it contributes to the growing empirical literature evaluating the impact of regional and national carbon pricing instruments (Martin et al., 2014; Rivers and Schaufele, 2015; Andersson, 2019; Colmer et al., 2020; Kim and Kim, 2016). Cropper et al. (2018) emphasized certain challenges involved with finding a suitable control group for the retrospective analysis of environmental regulation, including carbon pricing. This is especially true for the power sector, as almost all power plants are subject to the policy examined.¹ That both UK and non-UK power plants were subject to European-level energy policies while only UK plants were subject to the CPS offers an opportunity to compare the evolution of UK power sector emissions with that of an appropriately weighted average of European countries.

To my knowledge, two recent papers examine the effectiveness of the CPS: Abrell et al. (2019) estimate counterfactual electricity generation for each power plant subject to the CPS in the absence of the CPS using machine learning. They find that the CPS induced a total abatement of 26 MtCO₂e over the 2013-2016 period due to the short-term fuel switch from coal-fired to gas-fired plants². Gugler et al. (2020) rely on a Regression-Discontinuity-in-Time (RDiT) approach and exploits the annual change in the CPS tax

¹In the case of the ETS, the only exempted installations are those with a rated capacity of less than 20 Megawatt thermal input (MWth). In the UK, these facilities represented 0.2% of the installed capacity in 2015 (Source: Digest of United Kingdom Energy Statistics)

²Fuel switching occurs when carbon pricing increases the relative marginal cost of coal-fired plants compared to gas-fired plants due to the higher carbon intensity of the former. This change in costs modifies the short-term electricity supply curve, defined by the ranking of power plants by ascending marginal cost (the so-called “merit order”). As a result, the hourly output from high-emitting coal-fired plants’ increases while the hourly output from lower-emitting gas-fired plants decreases.

rate between 2013 and 2015. They estimate a cumulative abatement of 38.6 MtCO₂ over the 2013-2015 period compared to a no-policy scenario.

In contrast to these two papers, my paper uses less granular data but considers a longer post-treatment period and adopts a method allowing to take into account more mechanisms: a carbon tax on high-emitting input fuels may induce a decrease in emissions via fuel switching, but also via longer-term mechanisms such as plant closure and changes in demand or imports - although I find that demand and trade play a limited role compared to changes in the emission intensity of domestic production. Using a group of countries included in the European Union (EU) at the time in question as a control also enables me to control for the effect of both the EU ETS and a particularly important environmental regulation affecting all EU power plants at the period considered, the Large Combustion Plant Directive (LCPD). I also highlight the interactions between the UK carbon tax and these regulations. Finally, my work relies on open and freely accessible data, which facilitates replication.

Second, this paper contributes to the scarce literature examining the rapid decarbonisation of the UK power sector. Staffell (2017) links this decrease in emissions to the evolution of electricity demand, capacity, prices, the fuel mix, imports and exports in a descriptive approach. Wilson and Staffell (2018) insist rather on the significance of coal-to-gas fuel switching and underline the likely role of the CPS but do not quantify it. In contrast, the present work builds a comparison group and carefully examines potential confounding factors to estimate a plausibly causal impact of the CPS policy intervention on emission reduction.

Third, this paper is linked to a recent strand in the literature applying the synthetic control method to estimate the impact of environmental policies. The approach adopted here resembles that taken by Andersson (2019) who examines the impact of the Swedish carbon tax on transport sector emissions, but examines power sector emissions, where the carbon tax is levied on producers. Kim and Kim (2016) similarly assess the impact of carbon pricing in the power sector using the SCM approach, but in the context of the US Regional Greenhouse Gas Initiative (RGGI), where they focus on fuel switching rather than emission levels. Other recent works include Lee and Melstrom (2018), who estimate the impact of RGGI on electricity imports, and Isaksen (2020), who evaluates the effectiveness of international pollution protocols. A distinctive feature of my paper is that I build my outcome variable at the country level starting from disaggregated plant-level emission data, to allow me to account for shocks experienced by individual plants, and to document the channels through which the UK carbon tax may operate, an under-investigated area of research in the SCM literature according to Abadie (2021).

Beyond its academic contribution, this paper is relevant from a policy perspective. To

be in line with the 2015 Paris Climate Agreement and achieve net-zero emission targets, OECD countries must be coal-power free by 2030 (Rocha et al., 2016). The means necessary to make this transition are still under discussion. Some European countries are considering adopting a carbon price floor to hedge against variations in the ETS price (Newbery et al., 2019). Lessons can potentially be drawn from the UK situation analysed here.

The paper is organised as follows: Section 1.2 presents the background, potential effects of a carbon tax in the power sector, and descriptive evidence; Section 1.3 describes the empirical strategy; Section 1.4 presents the main results; Section ?? discusses them; Section 1.6 concludes.

1.2 The UK carbon tax: context and expected effects

1.2.1 The UK Carbon Price Support

The Carbon Price Support (CPS) was first introduced in April 2013. This domestic carbon tax was proposed in a double context of low prices on the EU carbon market, and the obligation for the UK to meet national targets for greenhouse gas emissions as defined in the 2008 Climate Change Act. The Climate Change Act set an emission target for 2050 and implemented a system of 5-year carbon budgets. Under the second carbon budget, running from 2013 to 2017, the UK had to reduce its total emissions by 236 MtCO₂e compared to the first carbon budget (covering 2008 to 2012). Low prices on the EU carbon market were perceived as potentially too low to effectively decrease emissions in the sectors covered by the ETS. In this context, the UK Government announced in March 2011 that a Carbon Price Floor (CPF) would be implemented in the power sector for the 2013/2014 budget year³. Under this price floor, power installations located in Great Britain (GB)⁴ would have to pay a tax called the Carbon Price Support (CPS), for which annual rates would reflect the difference between the desired level of carbon price floor and the expected carbon price on the EU ETS. The announced goal of the CPF was to tackle price uncertainty on the EU ETS and encourage investment in low-carbon technologies in the generation sector; in official communication documents, the CPF was labelled “support and certainty for low-carbon investment” (Hirst, 2018). The price floor was expected to increase over time, with a total carbon price target of £30 (around €35) by 2020.

³The budget year over which the annual tax rate is set runs from 1 April to 31 March of the next calendar year

⁴power generators located in Northern Ireland are integrated in a separate wholesale electricity market with the Republic of Ireland and are not subject to the policy.

The CPF was introduced as planned on April 1st, 2013. It was part of a broader reform called the Electricity Market Reform, which includes three other components which I describe in greater detail in the next section: a capacity market aiming at securing production capacity to back up intermittent renewable capacity; support for investments in renewable power capacity in the form of Contracts for Difference (CfDs)⁵; and Emission Performance Standards banning new coal-fired plants not fitted with Carbon Capture and Storage (CCS). The first rate of the CPS was set at around £5/tCO₂e. However, in 2014 the Government decided to freeze the CPS rate to £18/tCO₂e (€22 in 2016) until 2019/2020, after business representatives expressed concerns over the competitiveness of energy-intensive industries because of generators passing on the tax costs (Ares and Delebarre, 2016). Furthermore, actual ETS carbon prices turned out to be much lower than expected over the observed period. Because of the freeze and the difference between expected and actual carbon prices, the nature of the Carbon Price Support changed compared to how it was initially envisioned. It would effectively become a carbon tax with rates set several years in advance. Tax revenue go to the general budget.

The CPS applies to almost all power generators located in GB.⁶ The only exemptions are for stand-by generators used to provide emergency electricity supplies if a building's usual power supply is cut, and generators with a rated thermal input smaller than 2 MWth.

Table 1.2.1 shows the level of the tax rate confirmed for each period in 2016. Figure 1.2.1 overlays annual CPS rates with annual ETS carbon prices converted to British pound since 2009. The sum of the two gives the total carbon price paid by GB generators, which departs significantly from the level of carbon price floor initially envisioned. The CPS component nevertheless implies that GB power generators pay a much higher carbon price than non-GB power generators (only subject to the ETS price). In 2016, the relative difference reached a peak at five-fold. The rate of the tax depends on the carbon content of the input fuel used for power generation. The CPS rate on coal is about 70% higher than the tax on natural gas, in line with the much higher emission factor of coal. The CPS thus substantially increased the relative cost of coal-fired generation compared to gas-fired generation.

⁵CfDs guarantee a flat payment to low-carbon electricity generators: auctions determine the strike price, which reflects the long-term cost of generating low-carbon electricity for the awarded generators; when the electricity market price falls below the pre-determined strike price, contracted generators are then paid the difference between the strike price and market price; similarly, if the market price surpasses the strike price, contracted generators must pay this difference - See <https://www.emrsettlement.co.uk/about-emr/contracts-for-difference/> for more details.

⁶This includes conventional power plants, Combined Heat and Power (CHP) plants producing both electricity and heat (who only pay the CPS on the amount of fuel used to produce electricity for the grid), and auto-generators producing electricity for their own use (HM Revenue & Customs, 2017). Both CHP plants and auto-generators represent a negligible share of power production and emissions.

Period	CPS rate in £/tCO ₂ e
April 2013/March 2014	4.96
April 2014/March 2015	9.55
April 2015/March 2016	18.08
April 2016/March 2017	18
April 2017/March 2018	18
April 2018/March 2019	18

Table 1.2.1 – Level of CPS rate for each period in pound per ton of CO₂e

Source:Ares and Delebarre (2016)





Figure 1.2.1 – The Carbon Price Support and EUA price on the EU ETS

Notes: EUA stands for European Union Allowance and EUAs are carbon allowances traded on the ETS. Source for EUA price data: Ember website. Source for CPS prices: Hirst (2018). CPS prices adjusted with appropriate weights to reflect the calendar year rather than the April to March period. EUA price data converted to £ using yearly averages of monthly market exchange rates.

1.2.2 Descriptive evidence

Power sector emissions are the primary focus of this analysis. I define a variable of country-level per capita power sector emissions as the main outcome variable to facilitate country comparison⁷, expressed in tons of carbon dioxide equivalent or tCO₂e. Figure 1.2.2a shows de-meaned per capita power sector emissions for each European country between 2005 and 2017, using emission data described in section 1.3.2. I demean emission values by taking the difference between annual per capita power sector emissions and per capita power sector emissions averaged over the 2005-2012 period, which is the pre-treatment period before the introduction of the CPS in 2013. Most countries have relatively stable per

⁷The advantage of calculating emissions per capita rather than per MWh of electricity output is twofold: first, population as a variable is generally more stable over time than gross electricity production, so that the time variation in the outcome is mostly due to variations in emissions; second, decomposing emission per capita allows to analyse what happens to electricity demand and trade, rather than simply focussing on the emission intensity of domestic production.



(a) De-meaned per capita power emissions

(b) De-meaned per capita non-power emissions

Figure 1.2.2 – Evolution of per capita power and non-power sector emissions in European countries

Notes: For figure a (resp. figure b), per capita emission values were obtained by aggregating plant-level emission data for ETS participants identified as power installations (resp. non-power) at the country level, and dividing by annual country population. De-meaned per capita emissions were obtained by taking the difference between the annual value and the 2005-2012 average. “Other countries” include twenty European countries: Austria, Belgium, the Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Latvia, Lithuania, the Netherlands, Poland, Portugal, Slovakia, Spain, Sweden.

capita emissions, except for a few outliers, whose emissions are shown in dashed or dotted lines⁸. After 2012, UK emissions showed more of a decline relative to most other countries. Three other countries also showed decreasing emissions: Finland, Denmark and Greece. Emissions from Finland and Denmark varied significantly throughout the entire period (see Figure A.1.1 in Appendix). The decrease in emissions in Greece cannot, however, be traced back to a specific policy, but may be due to the large reforms implemented in all economic sectors around that period following the Greek debt crisis, combined with the deployment of a large amount of solar PV in 2011, 2012 and 2013 under an appealing feed-in-tariff program that was subsequently retroactively cut in 2014⁹. In contrast to the strong decrease in power sector emissions, UK per capita emissions in other ETS sectors follow the same path as other European countries, as shown on Figure 1.2.2b.

A variety of channels can trigger a pronounced decrease in power sector emissions. The following decomposition helps to understand the channels - for ease of reading, there are no indices, but all the variables should be interpreted as values for a given country c in a given year t . Calling P the country population and Q_{CO_2e} the quantity of emissions

⁸Estonia’s emissions are both high on average and with a high variance; the Czech Republic has the highest average after Estonia; Greece shows decreasing emissions after 2012; Finland and Denmark’s emissions have a high variance, likely due to the inter-annual variation in available hydro resources in Finland, and hydro and wind resources in Denmark.

⁹https://www.pv-magazine.com/2014/03/11/greece-brings-new-retroactive-measures-cuts-fit-by-30_100014491/

from the domestic power production¹⁰, $\frac{Q_{CO_2e}}{P}$ are per capita power sector emissions, and are the product of per capita electricity generation $\frac{Q_{elec}}{P}$ and the emission intensity of domestic power generation, $\frac{Q_{CO_2e}}{Q_{elec}}$:

$$\frac{Q_{CO_2e}}{P} = \frac{Q_{elec}}{P} \frac{Q_{CO_2e}}{Q_{elec}} \quad (1.1)$$

Q_{elec} can be rewritten as the difference between domestic gross electricity consumption C_g and net electricity imports, defined as electricity imports M minus electricity exports X , $(M - X)$. Gross electricity consumption is itself the sum of net consumption C_n (equivalent to demand), the amount of network losses, and the amount of electricity used by power generators. Grouping these two last components in the variable L , this leads to the following equation:

$$\frac{Q_{CO_2e}}{POP} = \left(\frac{C_n}{POP} + \frac{L}{POP} - \frac{(M - X)}{POP} \right) \frac{Q_{CO_2e}}{Q_{elec}} \quad (1.2)$$

From the right-end side of the equation, four different channels may lead to a decrease in per capita emissions: a decrease in consumption per capita $\frac{C_n}{POP}$ (the *demand* channel), a decrease in the amount of network losses and self-consumption of electricity by power generators $\frac{L}{POP}$ (the *network efficiency* channel), an increase in net imports per capita $\frac{(M-X)}{POP}$ (the *trade* channel), and a decrease in the average emission intensity of the domestic power sector (the *emission intensity* channel).

In appendix A.1.1, I show the evolution of the demand, trade and emission intensity channels¹¹. I find that UK electricity demand has been declining steadily since 2005 (Figure A.1.2a), and while UK net electricity imports per capita tended to increase, they remained very low compared to other countries (Figure A.1.2b). In contrast, the UK emission intensity of domestic power production can be seen following a similar pattern to total per capita emissions with a strong decrease after 2013 not observed in other European countries. These graphs suggest that the decrease in power sector emissions in the UK after 2013 is mostly due to a marked decline in the emission intensity of domestic production.

¹⁰Only fossil fuels used for electricity generation generate emissions, so Q_{CO_2e} is the sum of emissions from coal-, gas- and oil- fired power plants.

¹¹I leave aside the *network efficiency* channel, which is a technical component stable over time and unlikely to be influenced by carbon pricing. L can be estimated with Eurostat data as gross production Q_{CO_2e} plus net imports $(I - X)$ minus net consumption Y_n . For all European countries, L is constant over time in proportion of total gross production, at about 18%.

1.2.3 Potential confounders

Isolating the contribution of the CPS to this change in emissions poses certain challenges given that other policies were implemented both in the UK and across Europe during the same period. The fact that the UK and other countries in the European Union were all subject to the same policies enacted at the EU level, in particular the EU carbon market, air quality regulations, and the 2020 strategy setting targets for emission reductions and the deployment of renewable energy, allows me to differentiate out the effect of these policies by using other European countries as a counterfactual. Analysing the effects of UK-specific policies enacted at the same time as the CPS requires a different approach, prompting me to make various assumptions on the effects of these policies to bound their effects during the period considered. Four policies stand out as particularly important and are described in detail in Appendix A.1.2. Below I offer a brief summary of how each may have impacted UK emissions.

At the European level, the LCP directive (LCPD) is an air quality directive enacted in 2001, and made operational in 2008. It imposes emission limit values for local air pollutants for all combustion plants with a rated capacity above 50MWth. This required regulated plants to take the necessary steps to meet emission standards by 2008, or they could choose to opt out from the directive. Opt-out plants were exempted from the emission standards, but could not operate for more than 20,000 hours between January 1st, 2008 and December 31st, 2015 (European Commission, 2001), and were required to shut down once they had run for 20,000 hours or in 2015 (whichever came first). Plants had to decide by 2004 whether they wanted to opt-out or not¹². The UK had the highest share of opt-out capacity per capita in 2004, followed closely by Slovakia and Finland¹³. LCPD-induced plant closures could in part explain the decrease in emission levels observed in the UK compared to the averages in other EU countries. To avoid confounding the impact of the CPS and that of these two air quality directives, I will control for emission levels from LCP opt-out plants in my estimation strategy. The LCPD was replaced by the IED directive in 2016. The IED Directive offered a similar opt-out option, which required that plants decide by 2013. Given that the CPS had already been announced at that time, I consider the decision to opt out from the IED directive endogenous to the CPS. The UK has two IED opt-out plants, which have limited operating hours between January 1st 2016 and December 31st 2023 and must shut down completely by 2023.

At the UK level, three specific policies that fall under the Electricity Market Reform

¹²The decision to opt-out is made for each generating units, and some combustion plants only opted out some but not all of their generating units

¹³Own calculation based on EEA website: <https://www.eea.europa.eu/data-and-maps/data/large-combustion-plants-lcp-optimized-out-under-article-4-4-of-directive-2001-80-ec-4>

may have contributed to the decrease in emissions after 2013. First, the UK government subsidised the conversion of coal-fired power plants to biomass starting from 2012, and two plants representing 15% of UK emissions in 2012 benefited from a Contract for Difference for the biomass conversion of part (for Drax plant) or all of their units (for Lynemouth plant). Whether the subsidy for biomass conversion was decided to facilitate the conversion of coal-fired plants facing the CPS, or whether it was independent from the CPS remains somewhat ambiguous. In section 1.4.2, I develop a strategy to exclude the emission reduction induced by biomass conversion from my estimation.

Second, the Contracts for Difference system introduced in 2014 and its 2012 predecessor, the Financial Investment Decision (FID) Enabling for Renewables, could have impacted the fuel mix by increasing the share of renewable energy in the UK electricity production sector over the 2013-2017 period (outside the specific case of biomass generated by former coal-fired plants). However, available data on the projects being awarded a CfD in 2014, 2015 or 2017 reveal that only few of them were operational over the 2013-2017 period. Given the date when the projects became operational, the amount of clean electricity generated between 2013 and 2017 that can be imputed to the CfD projects represents only 0.4% of electricity generated with renewable sources (including waste and biofuels) and 0.1% of total electricity generated in the UK over the 2013-2017 period. If this electricity had been produced by coal-fired plants, the associated CO₂e emissions would have been 1.4 MtCO₂e (see Appendix A.1.2 for more details on the calculations).

Third, the capacity market introduces payments for electricity generators being awarded a capacity contract, in exchange for providing generation capacity at a pre-determined period of time. Since most of the capacity secured is for after 2018, this policy could only reduce UK emissions over the 2013-2017 period if two conditions are met: first, if the prospective capacity payment incentivised new capacity to be rolled out ahead of the capacity delivery year; second, if this new capacity had a lower emission intensity than existing plants. Using available public data on plants being awarded capacity contracts and new-build conventional plants in the UK between 2014 and 2017, I estimate that at most 2,590 GWh of electricity was generated over the 2013-2017 period from plants meeting the following conditions: 1)being awarded a capacity contract between 2014 and 2017 2)starting generation after 2014 and 3)not having been planned before 2014 based on available evidence¹⁴. The plants meeting these conditions are all fired with municipal solid waste and their generation over the 2013-2017 period represented only 0.6% of electricity generated with renewable sources and 0.2% of total electricity generated in the UK over the 2013-2017 period. If this electricity had been produced by coal-fired plants, the

¹⁴The largest new gas-fired plant which opened in 2016 and won several capacity contracts, Carrington power station, started being constructed in 2009 https://en.wikipedia.org/wiki/Carrington_Power_Station

associated CO₂e emissions would have been 2.3 MtCO₂e (see Appendix A.1.2 for more details on the calculations). Overall, the Contract for Differences and the capacity market are unlikely to have triggered important changes in power sector emissions over the 2013-2017 period (except for the impact of CfDs on the biomass conversion of coal-fired plants, considered separately).

1.3 Empirical strategy

1.3.1 The synthetic control method

To estimate the impact of the Carbon Price Support from other factors, I use the synthetic control method (SCM) developed in Abadie and Gardeazabal (2003) and Abadie et al. (2010, 2015). This method consists in building a counterfactual UK power sector by applying appropriate weights to the set of other European countries' power sectors. Providing that the obtained "synthetic" UK accurately reflects what the UK power sector would have looked like without the CPS, this method allows to estimate the causal impact of the CPS on per capita power sector emissions, and more generally on absolute abatement. The SCM method is particularly appropriate in the context of the CPS since the "treatment" applies to one country only, and within the country it affects almost all power installations, without time variation in treatment. Within the UK, there is then no obvious group of installations that could serve as counterfactual for how treated power plants would have evolved absent the policy.

Using the notation from the Neyman-Rubin Causal Model (Rubin, 1974b), the challenge is to estimate β_{UKt} when $t \geq 2013$, defined as:

$$\beta_{UKt} = Y_{UKt}^1 - Y_{UKt}^0 = Y_{UKt} - Y_{UKt}^0 \quad (1.3)$$

Y_{UKt}^1 designates, at each period, UK per capita power sector emissions in the presence of the CPS policy. Y_{UKt}^0 designates, at each period, UK per capita power sector emissions in the absence of the policy. β_{UKt} designates the difference between the two. Y_{UKt} designates the observed outcome. The challenge to estimate β_{UKt} , or "fundamental problem of causal inference" (Rubin, 1974b), comes from the fact that Y_{UKt}^1 is observed when $t \geq 2013$ but Y_{UKt}^0 is not.

Let us assume, following Abadie et al. (2010), that the outcome in the absence of intervention Y_{ct}^0 can be modelled as the following linear factor model, for each country c and period t :

$$Y_{ct}^0 = \delta_t + Z_c \alpha_t + f_t' \lambda_c + \epsilon_{ct} \quad (1.4)$$

δ_t is a time fixed effect, Z_c is a vector of observed exogenous country characteristics, α_t is a vector of unknown parameters, f_t is a vector of unobserved common factors (and f_t' its transpose), λ_c is a vector of unobserved country-specific effects or factor loadings, and ϵ_{ct} is an error term with mean 0 (typically capturing transitory shocks at the country level).

Such a model is more flexible than the typical difference-in-difference (DID) model because time effects and individual (country) time-invariant effects are allowed to interact. It is assumed that there is no permanent additive difference between the treated and control units (Doudchenko and Imbens, 2016). Abadie et al showed that under this specification, it is possible to use a function of outcomes observed post-treatment in other countries as an estimator of β_{UKt} :

$$\hat{\beta}_{UKt} = Y_{UKt} - \sum_{j=1}^J w_j^* Y_{jt} \quad (1.5)$$

Where $\sum_{j=1}^J w_j^* Y_{jt}$ is a weighted combination of the outcome for J countries having not implemented the policy, and the vector $W^* = (w_1^* \dots w_J^*)'$ should satisfy the following conditions:

$$\left\{ \begin{array}{l} w_j^* \geq 0 \quad \forall j = 1..J \\ \sum_{j=1}^J w_j^* = 1 \\ \bar{Y}_{UK}^K = \sum_{j=1}^J w_j^* \bar{Y}_j^K \\ Z_{UK} = \sum_{j=1}^J w_j^* Z_j \end{array} \right.$$

With \bar{Y}_{UK}^K a linear combination of pre-intervention outcomes in the UK and \bar{Y}_j^K a linear combination of pre-intervention outcomes for country j (The linear combination is defined by the vector $K = (k_1, \dots, k_{T_0})'$. For example, it can be the simple mean of pre-intervention outcomes $\bar{Y}_j^K = 1/T_0 \sum_{t=1}^{T_0} Y_{jt}$). Abadie et al also show that the estimator gets closer to the true parameter β_{UKt} when the number of pre-treatment periods is high compared to the scale of transitory shocks affecting countries.

In practice, to find the appropriate W vector I rely on an algorithm created by Abadie et al. The algorithm minimizes the distance between a vector of pre-intervention characteristics (also called predictors) in the treated country, X_{UK} (with dimensions $K \times 1$) and a weighted matrix of pre-intervention characteristics in the non-treated countries, $X_0 W$ (with dimensions $K \times K$). Pre-intervention characteristics are of two types: 1) the linear combinations of pre-intervention outcomes \bar{Y}_j^K , and 2) the country characteristics Z_j not affected by the intervention. To obtain the W vector, the programme starts with a positive and semi-definite matrix V that defines a dot product. The distance between X_{UK} and $X_0 W$ can then be written as

$$X_{UK} - X_0W = \sqrt{(X_{UK} - X_0W)'V(X_{UK} - X_0W)} \quad (1.6)$$

The goal is to find the vector $W^*(V)$ that minimizes this distance. Such minimization comes down to finding the right V matrix, which can be shown to be equivalent to a diagonal matrix assigning weights to linear combination of characteristics in X_{UK} and X_0W . Like Abadie and Gardeazabal (2003), I choose the V minimizing the mean squared prediction error (MSPE)¹⁵ of the outcome variable in the pre-treatment periods. Formally, let Y_{UK} be the (8×1) vector of pre-2013 power sector emissions from 2005 to 2012 for the UK and Y_j be the $(8 \times J)$ matrix of pre-2013 power sector emissions for the J other European countries. Then V^* is chosen among the set V of all non-negative diagonal $(K \times K)$ matrices, such that:

$$V^* = \operatorname{argmin}(Y_{UK} - Y_jW^*(V))'(Y_{UK} - Y_jW^*(V)) \quad (1.7)$$

The ability to build a good synthetic control can be assessed with at least two criteria: first, pre-intervention characteristics of the treated unit should be close to those of the synthetic unit. This depends on how well these characteristics predict the outcome and can be assessed by comparing pre-intervention characteristics for the treated and synthetic country. Second, the pre-intervention outcomes of the synthetic unit should be close to the pre-intervention outcomes of the treated unit. This can be checked graphically or by computing the MSPE. Compared to the difference-in-difference method, the number of pre-treatment periods should be large to limit the size of the bias, and relatively larger than transitory shocks affecting the countries (Abadie et al., 2010). As explained below, my main outcome variable is only available from 2005, which implies that my pre-treatment period has only eight years for the main specification. This is rather low compared to other published papers using the synthetic control method. I apply the same method on less precise aggregate data available since 1990 in appendix A.1.11 to assess whether the results change. The countries entering the synthetic UK are not the same, but the estimate of the impact is very close to the original one.

1.3.2 The Data

The empirical strategy relies on a comparison between the UK and other European countries and requires assembling a dataset at the country level. I first aggregate plant-level data on carbon emissions at the country-level. I then add different country-level power sector characteristics obtained from different sources.

¹⁵The MSPE gives the average of the squared difference between the treated unit's and the synthetic control's pre-intervention outcomes.

Power plant-level emission data: this data comes from the European Union Transaction Log (hereafter EUTL), the official register of the EU ETS managed by the EU Commission. The EUTL checks, records and authorises all transactions taking place between participants in the EU ETS. Every year since 2005, the start date of the EU ETS, participants have had to report their CO₂e emissions and surrender enough emission allowances to cover their emissions. Reported emissions are verified by an accredited verifier. Given that not all ETS installations are power installations¹⁶, the EUTL raw data contains both power plants and other types of plants.

One crucial step, thus, is to identify power installations. The main activity of each installation is publicly available, but there is no specific activity category for power installations. I rely on data provided by the UK-based think-tank Ember (formerly Sandbag) and a one-off file with more precise activity codes circulated by the EU Commission to identify power installations. Appendix A.1.3 describes the specific steps followed. I identify a total of 4,938 power plants, including 302 in total for the UK, with an average of 190 active power plants per year in the UK over the 2005-2012 period and 189 active power plants per year in the other twenty EU countries¹⁷, over the same period (See Appendix A.1.4 for summary statistics).

Almost all the UK power plants subject to the CPS are included in the data, except those with a rated thermal input between 2 and 20 MWth, not covered by the EU ETS. These small plants logically represent a very small share of total emissions. Two categories of UK plants present in the data are not subject to the CPS: power installations located in Northern Ireland, which represent a small share of UK power sector emissions (2.4% in 2012); and standby generators, also representing a small share of emissions¹⁸. I aggregate plant-level emissions at the country level, separately for power and non-power plants¹⁹, and obtain emission data for a panel of 21 European countries for the 2005-2017 period.

Country-level power sector characteristics: I add to this panel a set of annual country-level variables which I refer to in the descriptive analysis (see section A.1.1) and which I use in the empirical strategy: country population, installed capacity and power

¹⁶The ETS covers combustion installations with a rated capacity above 20 MWth, including power installations, and energy-intensive industries

¹⁷I exclude the countries which joined the ETS after 2005 (Romania, Bulgaria, Croatia), those which are not part of the European Union for the entire period considered (Slovenia, Norway, Liechtenstein and Iceland), and the three countries having less than ten power plants subject to the EU ETS: Luxembourg (only nine power plants), Cyprus (only three) and Malta (only two)

¹⁸Such generators are likely to be found in hospitals. In 2012, the six ETS power installations from the UK belonging to hospitals represent only 0.05% of UK power sector emissions

¹⁹Non-power plants are only used in figure 1.2.2b, to verify that the UK decrease in emissions only occurs in the power sector. Based on my categorization of power and non-power installations, there is a total of 9,127 non-power plants covered by the ETS, with an average of 618 active plants per year in the UK over the 2005-2012 period and 259 active plants per year in the other twenty EU countries.

generation by source, electricity consumption, electricity imports and exports, coal and gas prices, availability of lignite resources (a particularly polluting type of coal only used domestically), and average age of the coal-fired plants. Most of these data come from Eurostat. See appendix A.1.5 for details on each variable’s source. Table A.1.1 shows summary statistics by country for the main variables considered.

1.3.3 Selecting the predictors

Keeping the notation used in section 1.3.1, the set of predictors X_0 used to build the synthetic UK should be variables predicting country-level per capita power sector emissions, and which values are not affected by the CPS. Choosing characteristics’ values for the pre-treatment period ensures that the values are not affected by the CPS²⁰. The pre-intervention predictors chosen here are common drivers of emissions identified in the literature (Ellerman and McGuinness, 2008; Van den Bergh and Delarue, 2015; Lee and Melstrom, 2018). Appendix A.1.5 gives details on how each predictor variable is constructed.

In countries that, like the UK, rely both on coal- and gas-fired power plants for electricity generation, fuel switching has been identified as an important determinant of emissions variation. Fuel switching is influenced by the coal-to-gas price ratio (Ellerman and McGuinness, 2008), which is directly impacted by the CPS since the tax rate for coal is higher than for gas. I use country-level data on coal and gas prices to build a country-specific time-varying variable of coal-to-gas price ratio.

The coal price data are derived from trade statistics and do not take into account domestic coal resources. In particular, it does not take into account the availability of lignite, a low-quality type of coal with a very high emission intensity, used almost exclusively for power generation and mostly consumed domestically (Berghmans and Alberola, 2013). To account for the large differences in lignite resources across European countries - and its use for electricity generation -, I add a time-invariant predictor defined as a binary variable identifying the countries with large lignite resources: Germany, Poland, Hungary, Greece, and the Czech Republic²¹. Since the UK value is 0, the lignite variable constraints the programme to find a synthetic UK with as few countries with lignite reserves as possible.

Power sector emissions also depend on how much electricity demand needs to be covered by CO₂-emitting power plants. Residual load measures this amount of electricity

²⁰In theory, post-treatment values can also be included if the predictors are not affected by the treatment (Abadie et al., 2010).

²¹The lack of data on lignite reserves covering all Europe prompted me to build a binary rather than a continuous variable (such as the amount of proven reserves by country)

demand that requires using fossil fuels and biomass once generation from so-called “must-run” power generators (nuclear power plants) and those that generate with almost no marginal cost (solar, wind and hydro) are removed. I build a country-level time-varying variable of per capita residual load by taking the difference between electricity consumption and the generation from renewables and nuclear power plants, and dividing it by total population.

To account for the impact of the European air quality directives mentioned in section 1.2.3 and in Appendix A.1.2, I add one predictor measuring for each country the amount of emissions coming from installations that opted out of the Large Combustion Plant (LCP) Directive in 2004 and are expected to shut down in 2015. To build this variable, I first identify the name and location of those plants that opted-out of the LCP Directive based on the LCP data available on the European Environmental Agency’s website. I then manually identify these plants in the EUTL installation-level emission data²² to determine how much CO₂ these plants emit each year. For each country and each year, I calculate the sum of power sector emissions coming from LCP opt-out plants, and divide each sum by the country population to obtain a variable of per capita LCP opt-out emissions. The opt-out decision had to be made before the CPS was introduced. Thus the share of emissions coming from opt-out plants *before* the announcement of the CPS could not have been affected by the CPS. Using this predictor ensures that the synthetic UK will have about the same quantity of emissions from plants “at risk of closure” by 2015 as the UK. As a predictor, I take the value of per capita LCP opt-out emissions in 2009, shortly before the announcement of the introduction of the CPS in 2011.

The two last predictors I use are two lagged outcomes, which is standard in the SCM literature: these are per capita power sector emissions in 2005 and 2012, the first and last year of the pre-treatment period.

For the optimization, the residual load predictor is averaged for the period 2005-2012 and the coal-to-gas price ratio is averaged for the period 2007-2012 to ensure data consistency over time (see Appendix A.1.5 for more details). The remaining predictors are taken for one period only (lagged outcome, per capita opt-out emissions) or are time-invariant (lignite dummy) so they do not need to be averaged.

In a sensitivity analysis presented in Appendix A.1.9, I run the SCM with alternative sets of predictors. The magnitude of the results is unchanged when the installed capacity from plants using combustible fuels, the pre-treatment trend in renewables’ installed capacity, the number of heating degree days or the average age of coal-fired plants are included in the set of predictors (although some alternative sets of predictors satisfy less well the requirements of the synthetic control method).

²²The LCP data use a different installation identifier from the EUTL identifier

1.3.4 Selecting countries entering the donor pool

The “donor pool” designates the set of countries not affected by the CPS that will potentially enter the composition of the synthetic UK. The starting pool of countries consists of the twenty European countries included in the European Union (EU), other than the UK, described in the data section. Restricting the donor pool to EU countries rather than including other OECD countries has several advantages and one drawback. The main advantage is that over the period considered, the UK and other EU countries are subject to the same EU-level policies (in particular the EU ETS and the LCP directive, but also other energy policies). European countries would also have been more likely to be affected in a similar way by global shocks on the energy market, such as the 2011 US shale gas revolution. One drawback is that such geographic proximity and sectoral integration makes spillovers between treated and synthetic unit more likely.

Starting with this initial pool of twenty countries, it is important to discard the countries that are likely to be poor counterfactuals (Abadie et al., 2010). This essentially describes three country types: first, countries that suffered idiosyncratic shocks to the outcome of interest, either by directly introducing a policy targeting the power sector or via a more generic exogenous shock likely to affect the electricity sector; second, countries more likely to have been directly affected by the CPS; and third, countries with very different characteristics compared to the UK, which may cause severe interpolation biases.

By 2017, no other European country had adopted a carbon tax or a carbon price floor that would interact with ETS pricing in the power sector (Metcalf and Stock, 2020).²³ The most radical change in other European countries’ power sectors is the case of Germany, which unexpectedly decided to phase out nuclear energy following the 2011 Fukushima nuclear accident. I therefore exclude Germany from the donor pool. Since the European debt crisis significantly affected the Greek economic environment over the period, I also exclude Greece. However, including them in the donor pool does not change the results, as shown in appendix A.1.10.

Regarding the second country type, tension can occur between excluding countries from the donor pool countries whose outcomes are affected by the treated unit and identifying those sufficiently comparable to the treated unit (Abadie, 2021). While I do not exclude any country based on the risk of spillover, I do discuss this risk and offer an estimation the amount of potential spillovers in section 1.4.4.

Finally, to avoid including countries that differ too greatly from the UK, I eliminate Estonia, a country where high emissions per capita are due to the unusual use of oil

²³France and the Netherlands discussed introducing a carbon price floor as well (?), only the Netherlands have passed a concrete law in August 2018, and the Dutch CPF was scheduled to start in 2020.

shale for power generation, a high-emitting input fuel. I also exclude the two other Baltic countries, Latvia and Lithuania, which unlike the UK do not use coal for power generation (see Figure A.1.3). Since coal-to-gas fuel switching is expected to be an important driver of decarbonisation, it is relevant to restrict the analysis to countries with the capacity to do so.

In the end, the donor pool includes 15 EU countries. Appendix A.1.10 shows that changing the composition of the donor pool does change the composition of the synthetic UK and the estimates, but not their order of magnitude. To ensure that building a convex combination of countries (having positive weights) that closely reproduce the UK's values for predictors and emissions is possible, there needs to be common support between the distribution of the predictors in the donor pool and in the UK. I check that this is the case for all variables (See the histograms in appendix A.1.6).

1.4 Results

1.4.1 Upper Bound

I start by applying the SCM method using the emission outcome variable, donor pool and predictors exposed in the previous section. Figure 1.4.1a shows that the obtained synthetic UK (dashed line) reproduces well the trajectory of UK per capita power sector emissions (continuous line) before 2013, with a Mean Squared Prediction Error (MSPE) of 0.01. Compared to the average per capita power sector emissions for the donor pool (dotted line), the synthetic UK has a relatively close trajectory but higher per capita emissions. Table 1.4.1 shows the weights received by each country in the synthetic UK, which comprises five countries: Ireland (49.2%), Slovakia (25.6%), the Netherlands (13.7%), Finland (5.8%), and the Czech Republic (5.7%). The remaining potential control countries receive a weight of 0. The large weight observed for Ireland is not surprising: the two countries have close institutions and energy markets, and like the UK, Ireland has a substantial portfolio of coal- and gas-fired power plants. The Netherlands and Slovakia also showed a potential for coal-to-gas fuel switching (see figure A.1.3a). The Netherlands showed a residual load per capita close to the UK, and Slovakia and Finland showed, like the UK, a substantial amount of LCP opt-out emissions.

The good pre-treatment fit between the UK and synthetic UK suggests that after 2013, the synthetic UK accurately replicates the evolution of per capita emissions in the UK power sector absent the CPS (assuming no other UK-specific confounder). The fit is less good in 2012, where UK emissions peak. This is also the year where the share of coal in the UK fuel input mix is the highest, which can be partly explained by the low

coal-to-gas price ratio that year (lowest point since 2007) (BEIS, 2018). If power could easily be stored, the 2012 peak could also be interpreted as an anticipation effect of the CPS, which was announced in 2011. Coal-fired plants would then have an incentive to use their coal before being taxed in 2012, store the electricity, and sell it over subsequent years. But electricity cannot be stored, and production has to match demand at every point in time. The generation mix at each point in time depends on the merit order, that is, the ranking of plants' marginal costs. Anticipation can only materialize if some coal-fired plants alter the merit order by accepting to sell at a price lower than their marginal cost in order to get rid of their coal reserves. Power plants scheduled to close because of the LCPD may have had an interest in adopting such measures, especially if they had excess coal stocks that they wanted to get rid of before being taxed²⁴.

Table 1.4.2 shows the average value of each predictor for the UK, synthetic UK, and the average among the donor pool. The values of the predictors for the synthetic UK are close to the values for the actual UK - indicating that the synthetic UK is a relatively good counterfactual to the UK. The balance in predictors' values between the UK and synthetic UK is better than between the UK and the average taken from the donor pool for all predictors, further justifying the use of the SCM method.

Figure 1.4.1b shows the emission gap between the UK and synthetic UK for each period. The gap between the UK and synthetic UK widens significantly between 2014 and 2016, while UK emissions were slightly higher than synthetic UK emissions in 2013. This evolution is consistent with the timing of the introduction of the CPS (April rather than January 2013), with the strong increase in the CPS rate between 2013 and 2015, and the CPS freeze in 2015/2016. The corresponding annual abatement for each year $t \in [2013, 2017]$ can be calculated by multiplying the annual gap in per capita emissions by the UK annual total population. On an average year, emissions decrease by 26 percent, with an associated semi-elasticity of -1.65% of emissions per Euro of the tax on average. Adding up all annual abatements gives a total cumulative abatement of 191 million tCO₂e (MtCO₂e) over the 2013-2017 period. Emissions abatement was the most pronounced in 2017, when UK emissions became 50% lower than synthetic UK emissions.

²⁴Anecdotally, official data on annual coal consumption and stocks by electricity generators indicate that coal stocks as a share of stocks and consumption are lower in 2012 compared to previous periods (20% vs 27% on average over 2005-2012), although the difference is not large (BEIS (Department for Business, Energy & Industrial Strategy), 2019).

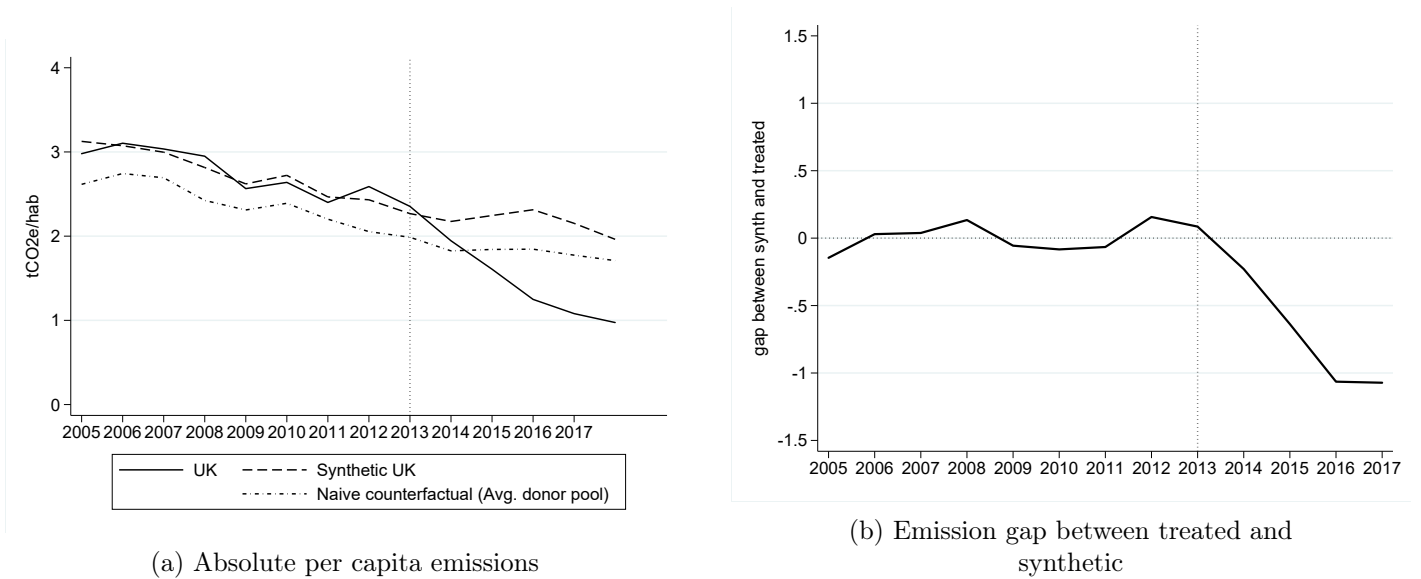


Figure 1.4.1 – UK and synthetic UK per capita emissions

Notes: For each period, the variable of per capita emissions corresponds to the sum of CO₂e verified emissions from power installations subject to the EU ETS, divided by the average country population that year. The vertical line is set in 2013, date where the CPS is introduced. The synthetic UK comprises five countries: Ireland (49.2%), Slovakia (25.6%), the Netherlands (13.7%), Finland(5.8%), and the Czech Republic (5.7%).

Country	Weight	Country	Weight
Austria	0	Ireland	0.492
Belgium	0	Italy	0
Czech Republic	0.057	Netherlands	0.137
Denmark	0	Poland	0
Spain	0	Portugal	0
Finland	0.058	Sweden	0
France	0	Slovakia	0.256
Hungary	0		

Table 1.4.1 – Country weights in Synthetic UK

Note: all weights are between 0 and 1 because the Synthetic control method imposes positive weights summing to 1.

Variable	UK	Synth. UK	Avg. Donor pool
Per capita residual load	4.29	4.30	3.37
Coal-gas price ratio	0.52	0.51	0.71
Per capita LCP opt-out emissions	0.29	0.24	0.22
Lignite dummy	0.00	0.06	0.20
Per capita emissions 2005	2.98	3.13	2.62
Per capita emissions 2012	2.59	2.43	2.05

Table 1.4.2 – Predictors’ values for the UK, synthetic UK and average of the donor pool

Notes: the per capita residual load is averaged for the period 2005-12, and the coal-to-gas price ratio for the period 2007-12. LCP opt-out emissions are taken in 2009, the lignite dummy is time-invariant. Outcome lags are taken in 2005 and 2012.

1.4.2 Lower bound

Potential confounders and emission decomposition: In the result presented above, the assumption is that the difference in emissions between the UK and synthetic UK after 2013 resulted solely from the Carbon Price Support. As mentioned in section 1.2.2, UK-based policies and European policies affecting the UK differently from other countries may have further contributed to the observed decrease in emissions in the UK.

Regarding European policies, the predictor of LCP opt-out emissions should guarantee that the actual and synthetic UK have approximately the same amount of emissions coming from plants facing a high risk of closure. Given the close values of the LCP opt-out predictor, any difference in the evolution of emissions from opt-out plants between the actual and synthetic UK is assumed to be caused by the CPS. For example, the CPS may have affected the way in which the remaining operating hours of each opted-out plant were distributed over the 2005-2015 period; it may also have motivated certain opt-out plants to lobby for government loopholes to allow them to remain operational in spite of set limits on operating hours. The decision to opt out from the IED directive occurred after the announcement of the CPS, so any difference observed in opt-out behaviour between UK and non-UK plants could be a consequence of the CPS.

In contrast, UK-specific policies implemented at the same time as the CPS cannot be controlled for in the SCM framework. In Appendix A.1.2, I estimate that both the Contracts for Differences and the capacity market likely had a limited impact on the fuel mix over the period considered, with estimated emissions reductions at a maximum of 1.4 MtCO_e for the CfD, and 2.3 MtCO_e at most for the capacity market. The situation is different for the biomass conversion policy: the largest UK coal-fired plant, Drax, converted half of its production units from coal to biomass between 2013 and 2016, and the smaller station Lynemouth stopped using coal in December 2015 to prepare for biomass

conversion.

To assess the role of the air quality directives and of the biomass conversion, I rely heavily on plant-level emission data. I decompose emissions into four categories for the UK and synthetic UK: emissions coming from LCP opt-out plants (light grey); emissions from IED opt-out plants (dark grey); emissions from UK plants having benefited from subsidies to convert to biomass (medium grey); and remaining emissions from other plants (black). Figure 1.4.2 shows the emission decomposition results in the UK and the synthetic UK.




Figure 1.4.2 – Per capita CO₂e emissions by source, UK and synthetic UK

Notes: The synthetic UK comprises five countries: Ireland (49.2%), Slovakia (25.6%), the Netherlands (13.7%), Finland(5.8%), and the Czech Republic (5.7%).

LCP opt-out emissions are higher in the UK than in the synthetic UK before 2008, but move closer to each other between 2009 and 2012, just before the CPS is introduced. After 2012, LCP opt-out emissions decrease sharply in the UK while they remain relatively constant in the synthetic UK. The evolution of the synthetic UK opt-out emissions may seem surprising: opt-out plants are expected to shut down by 2015 at the latest and we should have zero emissions from opt-out plants in 2016 and 2017, both in the UK and synthetic UK. Singhal (2019) confirmed that as many as 60% of opt-out plants actually continued to operate after 2015. The difference in the trajectory of UK and non-UK LCP opt-out plants suggests that the CPS intensified UK plants' response to the LCP opt-out option and accelerated their closure²⁵. The figure also confirms that the UK emission peak in 2012 mainly comes from LCP opt-out plants, which had an interest in using their polluting inputs before the introduction of the tax. IED opt-out emissions are relatively low in both the UK and synthetic UK just before the opt-out decision.

²⁵Such an interpretation would also confirm a Guardian journalist's statement that "[UK coal-fired] Plants have closed in recent years as EU pollution standards started to bite, but it was increases in the UK carbon tax that sealed their fate" (Vaughan, 2018).

Finally, the carbon emissions from UK plants that converted to biomass burning represent a substantial share of UK emissions, which decreased after 2013, in particular between 2015 and 2016. Drax power plant - responsible for more than 90% of the emissions from biomass conversion - had only half of its six units converted to biomass, suggesting that the decrease in emissions after 2013 must be partly explained by the impact of the CPS on the non-converted units. Furthermore, the introduction of the CPS may have influenced the willingness of UK plants to convert to biomass processing. The estimate above implicitly assumes that this was the case.

To avoid biomass conversion confounding the impact of the CPS on emissions, below I estimate counterfactual CO₂e emissions for the biomass converted plants if they had not converted to biomass. In Appendix A.1.8, I run a second test where I remove from the UK emissions variable all the emissions coming from biomass converted plants and generate a new synthetic UK based on this modified emission variable.


Lower bound: counterfactual emissions of plants converted to biomass if they had not converted: Appendix A.1.7 provides details on the imputation of emissions in the absence of biomass conversion. I summarize below the method used for each plant. Estimating what emissions from Drax plant would have looked like in the absence of biomass conversion is relatively straightforward because I observe the evolution of emissions for its coal units which did not convert. I first combine data on the monthly generation of the three Drax coal units over the 2009-2016 period, combined with their average emission intensity (kindly provided by Mirjam Kosch and Jan Abrell) to estimate the amount of CO₂ emissions coming from Drax coal units. I then subtract the emissions coming from the coal units from the total emissions reported for Drax in the EUTL data to estimate the emissions coming from the three units converted to biomass between 2013 and 2016. In 2016, the estimated CO₂ emissions for these units were close to zero, which makes sense given that the three units run entirely on biomass in 2016. I can then assume that their emissions were also zero in 2017, which means that all the emissions reported for Drax in the EUTL in 2017 came from the three coal units. Third, I assume that absent the biomass conversion, biomass converted units would have had a similar emission trend to that observed for the three coal units. Concretely, I start with their estimated CO₂ emissions value for 2012, and I apply the same annual percent change as the annual percent change for the three coal units, which provide me with "counterfactual" emissions for the three units converted to biomass. Finally, I add these "counterfactual" emissions to the actual emissions of the three coal units and obtain counterfactual emissions for Drax in the absence of the biomass conversion policy.

For Lynemouth plant, using the same method is not possible given that the plant as


a whole began its conversion in December 2015²⁶. I thus formulate a cruder hypothesis that absent the conversion to biomass, emissions in 2016 and 2017 would have been the same as in 2015. For the two plants, the estimated counterfactual emissions absent the biomass conversion policy only differ from the actual emissions for the 2013-2017 period. For the 2005-2012 period, counterfactual emissions equal actual emissions.

Finally, I generate a modified outcome variable for the UK, which include both Drax and Lynemouth's counterfactual emissions instead of their actual emissions. UK pre-treatment emissions remain the same as before with this modified variable, such that the synthetic UK obtained in the previous section still represent an appropriate comparison unit for the UK. Figure 1.4.3a shows the UK emission trajectory with this modified outcome variable (dark grey), overlaid to the actual UK and synthetic UK emission trajectories (in black). After "removing" the effect of biomass conversion, emissions are logically higher after 2013 for the modified UK emission variable. The gap between UK and synthetic UK is then reduced (Figure 1.4.3b). On an average year, emissions decrease by 22.5 percent, with an associated semi-elasticity of -1.41% per Euro of the tax. The total cumulative abatement amounts to 164 million of tCO₂e. Withdrawing the upper bound estimate for the effect of the capacity market and the CfD (≈ 4.7 MtCO₂e in total), I obtain a cumulative lower bound abatement of around 159 MtCO₂e. The estimated abatement is lower in Appendix A.1.8, where the emissions from the biomass converted plants are removed from UK emissions; the difference between the two abatement results corresponds roughly to the decrease in Drax and Lynemouth counterfactual emissions (in the absence of biomass conversion) between the pre- and post-treatment period.

²⁶<https://www.power-technology.com/projects/lynemouth-biomass-power-station-northumberland/>



(a) Absolute per capita emissions



(b) Gap between treated and synthetic

Figure 1.4.3 – SCM with counterfactual UK emissions w/o biomass conversion

Notes: The synthetic UK comprises five countries: Ireland (49.2%), Slovakia (25.6%), the Netherlands (13.7%), Finland(5.8%), and the Czech Republic (5.7%).

1.4.3 Inference

With the SCM method, inference can be derived from a set of placebo tests, which consists in applying the SCM method to untreated units or fake treatment dates (Abadie et al., 2010).²⁷ I am able to show that the results are likely driven by the causal impact of the Carbon Price Support by measuring (1) the likelihood of finding an effect of the same magnitude as what I find when I apply the method before 2013 (the in-time placebo test); (2) the likelihood that the result is driven by particular behaviour specific to one country in the donor pool (the leave-one-out test); (3) the likelihood of finding an effect of the same magnitude when I apply the method to other countries (the in-space placebo test or permutation test). The in-time and leave-one-out tests are run for the actual UK emissions used for the upper bound estimation from section 1.4.1. Results would be the same for the lower bound, given that the composition of the synthetic UK and the trajectory of emissions in the 2005-2012 period are the same. For the permutation test, I show results based both on the upper bound and on the lower bound estimation of abatement.

In-Time placebo One way to check that the results observed were indeed caused by the CPS policy is to assume that a similar policy was implemented at another date prior to 2013, apply the same method to generate a synthetic UK, and check that the UK and synthetic UK have similar per capita emissions before and after this artificial intervention

²⁷Having only one treated unit is insufficient for building confidence intervals as found in previous works by Gobillon and Magnac (2016) and Isaksen (2020).

date. Figure 1.4.4 shows the UK and synthetic UK outcome obtained when treatment is assumed to occur in 2010 rather than 2013. The synthetic UK closely resembles the UK emission trajectory before 2010, and there is no significant gap between treated and synthetic UK in 2011 and 2012.




Figure 1.4.4 – Gap between treated and synthetic UK, CPS assumed to start in 2010

Notes: Predictors are averaged for the period 2005-09, except for the coal-gas price ratio averaged for the period 2007-10. The lagged outcome is taken in 2009 instead of 2012. The synthetic UK comprises seven countries with a weight above 1%: Ireland (45.8%), Slovakia (23.1%), Finland (15.3%), the Czech Republic (3.3%), the Netherlands (3.1%), Sweden (2.3%) and Denmark (1.1%).

Leave-one-out test Another common test recommended by Abadie et al. (2010) is the leave-one-out test, which consists in running the synthetic control method again after iteratively removing each country that receive a positive weight in the synthetic UK baseline estimates. If the results change significantly when a country is removed, it means that the estimated effect may have been caused by the evolution of emissions for that country, rather than by the change in UK emissions. Figure 1.4.5 shows that the results change very little across the alternative donor pools. This test suggests that my estimate of abatement is not driven by the presence of a specific country in the donor pool.

Permutation test The permutation test consists in building a synthetic counterfactual for each country of the donor pool. Then, the gap between each country and its synthetic counterpart is compared with the gap obtained for the UK in the main results section. If for many countries the gap is as large as for the UK, it means that the gap obtained for the UK could have happened just by chance, rather than as a result of the introduction of the CPS. Figure 1.4.6a shows the gap between the treated and synthetic country for the UK and all the other countries in the donor pool. For the Czech Republic and France, having




Figure 1.4.5 – Leave-one-out test

Notes: Synthetic UK for the main specification: five countries: Ireland (49.2%), Slovakia (25.6%), the Netherlands (13.7%), Finland(5.8%), Czech Republic (5.7%). Specification without Ireland: four countries: the Netherlands (45.9%), Spain (38%), Denmark (8.6%), and the Czech Republic (7.5%). Specification without Slovakia: five countries: Ireland (52.5%), France (18.4%), the Netherlands (17.5%), the Czech Republic (6.2%) and Finland (5.4%). Specification without the Netherlands: four countries: Ireland (58.0%), Slovakia (26.2%), Finland (10.2%), and Poland (5.7%). Specification without Finland: thirteen countries: Ireland (47.9%), Slovakia (20.5%), the Netherlands (18.8%), Poland (11.1%), all other countries have a weight below 1%. Specification without the Czech Republic: five countries: Ireland (50.9%), Slovakia (22.3%), the Netherlands (14.2%), Poland (6.6%), and Finland (5.9%).

respectively the highest and lowest per capita emissions, and for Italy, it is impossible to find a convex combination of countries that will replicate the pre-2013 emissions. These countries are therefore not included. For Denmark and Finland, the pre-2013 fit is poor, with a pre-treatment MSPE more than 10 times greater than the UK.²⁸ Comparing the UK emission gap with these countries is less meaningful since the conditions for a good synthetic control are not met. Hence Figure 1.4.6b drops these two countries, as advised by Abadie et al. (2010). The UK stands out as having the largest decrease in per capita emissions after 2013.

To illustrate the difference in the magnitude of pre- and post-2013 emission gap between the UK and the other permutations, one can also compute the ratio of post to pre-MSPE for all countries (Abadie et al., 2010). We should expect to observe an unusually high ratio for the UK. Figure ?? shows that the UK ratio is indeed the largest, both with the upper bound and lower bound estimates of abatement. We can calculate the estimated probability to observe an effect as large as the one observed for the UK under a random permutation of the intervention on the data, by dividing the number of countries

²⁸As mentioned in appendix A.1.1, Denmark and Finland have a high variability in emissions, likely explained by the large inter-annual variations in renewable sources available for electricity generation.

having a higher ratio than the UK by the total number of countries (Abadie et al., 2010). Here the UK has the highest ratio amongst the 13 countries, so the associated probability is $1/13 = 7.7\%$, the lowest possible probability with this sample size.




Figure 1.4.6 – Permutation test

Notes: In both figures, the Czech Republic, France and Italy are not included: for these countries it is impossible to find a convex combination of countries replicating pre-2013 emissions. On figure b, the two countries with an MSPE 10 times higher than the UK, Denmark and Finland, are not included.




Figure 1.4.7 – Ratio of post to pre-MSPE

Notes: the Czech Republic, France and Italy are not included: for these countries it is impossible to find a convex combination of countries replicating pre-2013 emissions.

1.4.4 Risk of spillovers

For the synthetic control method to identify the causal impact of the intervention, candidate units for the synthetic control group should not be affected by the intervention. As an overlapping policy to an existing carbon market, The CPS could spill over to other European countries' power sectors via two channels highlighted by Perino et al. (2019): internal leakage, that is, an increase in UK net electricity imports from other European countries; or a waterbed effect, that is, an increase in emissions from European power plants not subject to the CPS, due to the negative effect of the CPS on ETS permit prices under a fixed emission cap. Quantifying the magnitude of these two effects for the EU carbon market as a whole goes beyond the scope of this paper, which focuses on the impact of the CPS on UK emissions. What I endeavor to assess is the risk of spillovers to countries entering the synthetic UK, given that they serve as a counterfactual for the evolution of UK emissions in the absence of a CPS.

I first estimate the amount of emissions from countries in the synthetic UK potentially caused by import spillovers. This amount is naturally bounded by the limited interconnection capacity of the UK with the rest of Europe. I then estimate the amount of emissions in the synthetic UK potentially caused by a waterbed effect. The two effects combined represent 11% of the estimated abatement of the lower bound.

Risk of spillover via increased electricity imports : UK net electricity imports per capita are generally low compared to other European countries (see Figure A.1.2b), representing 2% of gross electricity consumption in the 2005-2012 period. However, net imports increased to 5% of gross electricity consumption in the 2013-2017 period. If this


increase was caused by the CPS, it could threaten the identification strategy because two of the UK trading partners, Ireland and the Netherlands, have a combined weight of 63% in the synthetic UK. The increase in UK net imports would then increase the synthetic UK's emissions *as a result of the CPS* and contaminate the counterfactual. The question is how large in magnitude this contamination is, relative to the estimated abatement. I calculate the maximum amount of synthetic UK emissions that may have been directly caused by the CPS, considering that the increase in UK electricity imports from France, the Netherlands and Ireland after 2012 was entirely caused by the CPS²⁹. I estimate the emissions associated with these exports for Ireland and the Netherlands (those countries entering the synthetic UK). In Appendix A.1.12, I run another test where I exclude all interconnected countries from the donor pool to assess whether the presence of Ireland and the Netherlands in the synthetic UK would drive up the results. The estimated abatement is 14% lower without interconnected countries, but the balance in predictors' characteristics is also less good.

First, I calculate the excess electricity generation in the Netherlands and in Ireland which can be imputed to CPS-induced exports to the UK: to do so, I simply calculate, for every post-treatment year, the difference between electricity exports to the UK that year and average electricity exports to the UK in the pre-treatment period. I use electricity trade statistics from Ofgem to determine quarterly trade flow for each interconnector with the UK.³⁰ I estimate that on an average year between 2013 and 2017, the Netherlands produced an excess of 2,965 GWh, and Ireland produced an excess of 382 GWh, compared to the pre-treatment period.


Second, I calculate the emissions associated with this electricity generation. The emission intensity of this displaced generation depends on which technology is used for marginal generation. According to Guo and Newbery (2020), gas is the marginal fuel most of the time in the Netherlands. Furthermore, the emission intensity of fuel displaced by renewable energy in Ireland in 2012 was estimated at 0.43 tCO₂e/MWh (Sustainable

²⁹Guo and Newbery (2020) estimate that 0.9% of the CO₂ emission reduction taking place in the UK between 2015 and 2018 was undone by the increase in electricity imports from France and The Netherlands. I cannot use this estimate because the time period is different, the estimated UK emission reduction is based on a different method - a dispatch model of the 2015 GB power system, see (Kong Chyong et al., 2020) -, and the paper only considers France and the Netherlands (two markets coupled with the UK in 2014, which enabled cross-border electricity trading to take place as soon as market prices were different across the two sides of the interconnection) but not Ireland, which represents half of the synthetic UK.

³⁰Since the Netherlands-UK interconnector only became fully operational in April 2011, I average trade flows between the second quarter of 2011 and the fourth quarter of 2012 to get average trade flow pre-treatment (I include the first quarter of 2013 in the post-treatment period to be consistent with the rest of the analysis). For Ireland-UK trade, I consider separately the interconnectors between Ireland and Northern Ireland, open for the entire pre-treatment period and for which I average trade flows for the 2005-2012 period, and the interconnector between Ireland and Wales (East-West interconnector), which only opened in September 2012 and for which I only consider the trade flow of the last quarter of 2012.



(a) UK and synthetic UK w/o imports spillovers



(b) UK and synthetic UK w/o waterbed effect

Figure 1.4.8 – Spillover risk

Notes: The synthetic UK comprises five countries: Ireland (49.2%), Slovakia (25.6%), the Netherlands (13.7%), Finland(5.8%), and the Czech Republic (5.7%). UK emission values include estimated counterfactual emissions in the absence of biomass conversion for Lynemouth and Drax plants.

Energy Authority of Ireland, 2014), which is close to the emission intensity of gas in the UK. Assuming a marginal intensity of 0.43 tCO₂e/MWh both in the Netherlands and in Ireland, the excess emissions caused by the CPS are 6.4 MtCO₂e over the 2013-2017 period in the Netherlands, and 0.8 MtCO₂e in Ireland (exporting less to the UK than the Netherlands).³¹ Third, I remove these excess emissions from Dutch and Irish emission data over the 2013-2017 period. I then assess how the emission trajectory of the synthetic UK changes when these excess emissions are removed.³²

Figure 1.4.8a shows how the trajectory of the synthetic UK emission changes after removing these “leaked” emissions from Ireland and the Netherlands. Net imports from both the Netherlands and Ireland are higher than the pre-treatment average in 2015, 2016 and 2017, such that removing the estimated “leaked” emissions reduces emissions from the synthetic UK.³³ Overall, the gap between the UK and synthetic UK is smaller than when these spillovers are not accounted for, which is expected. The resulting cumulative abatement is smaller by 5% compared to that estimated in section 1.4.2.

³¹If I instead calculate emissions assuming that gas is the marginal fuel, with an emission intensity of 0.4 tCO₂e/MWh (which is the average for the UK, see (Abrell et al., 2019)), these excess emissions are 5.9 MtCO₂e in the Netherlands and 0.8 MtCO₂e in Ireland . If I assume that coal is the marginal fuel, with an emission intensity of 0.89 tCO₂e/MWh (which is the average for the UK, see (Abrell et al., 2019)), the excess emissions are 13.2 MtCO₂e for the Netherlands and 1.7 MtCO₂e for Ireland.

³²I do not impute these excess emissions back to the UK because the goal is not to estimate the impact of the CPS net of leakage, but rather to accurately estimate the impact of the CPS on UK emissions by making sure that the counterfactual does not include spillover effects.

³³Emissions slightly increased in the modified synthetic UK in 2013 and 2014, because net imports from Ireland decreased during this period and were not compensated by the increase in net imports from the Netherlands given the much higher weight of Ireland than the Netherlands in the synthetic UK.

Risk of spillover via a waterbed effect Theoretically, the waterbed effect designates the mechanism via which, under a common emission cap (in this case the cap set by the ETS carbon market), any emissions reduction in a given country only leads to an increase in emissions elsewhere (Böhringer et al., 2008; Goulder and Stavins, 2011; IPCC et al., 2014; Perino, 2018). The waterbed effect would arise because the CPS decreases demand for emission permits from UK installations subject to a higher carbon price. On the EU ETS market as a whole, the shift to the left of the demand curve can only be compensated by a price decrease, since the supply is fixed and perfectly inelastic because of the emission cap. With cheaper permits, individual installations subject to the ETS but not to the CPS can buy more allowances and emit more. Aggregate emissions remain unchanged.

The concern for the empirical strategy is a waterbed effect affecting the power installations included in the synthetic UK. This risk exists but the magnitude of the effect is likely to be small for two reasons. First, UK power installations represent only a small share of the total ETS market (in 2012 they represented 8.8% of total emissions covered by the EU ETS), so the demand-side shock from UK power installations is likely to be small. To illustrate this, Figure 1.4.8b simulates a 100% waterbed effect scenario. Under this scenario, the decrease in UK power sector emissions after 2013³⁴ is assumed to be caused by the CPS and be compensated by an equivalent increase in emissions coming from other ETS installations. Observed power sector emissions outside the UK therefore include a waterbed effect component, compared to what non-UK emissions would have looked like in the absence of the CPS. Assuming that the waterbed effect was spread across the different sectors and countries based on their share in ETS emissions in 2012, I estimate the waterbed effect component for each country's power sector. I then estimate each country's adjusted, *lower* emission value excluding the waterbed effect component³⁵. Figure 1.4.8b shows the emission trajectory for the modified UK emission variable (corresponding to the lower bound estimate), and for the synthetic UK after removing the hypothesized waterbed component. Given the low weight of UK power installations in the market's total emissions, the waterbed effect component is small once spread over all ETS countries, and the adjusted synthetic UK emissions are only slightly higher than in the main specification. The magnitude of this waterbed component is similar to that of the import spillover component estimated above: the cumulative abatement based on Figure 1.4.8b is also smaller by 5% compared to that estimated in section 1.4.2.

The second reason why a strong waterbed effect is unlikely linked to the specific

³⁴using the modified UK emission variable, the one including emissions from biomass converted plants if they had not converted to biomass.

³⁵To take a concrete example: in 2014, UK emissions were lower by 25 MtCO₂e than in 2013; Ireland represents 1.2% of ETS power sector emissions in 2012 (excluding the UK); power installations represent 66% of emissions in the whole ETS; the 2014 waterbed component for Ireland is estimated to be 1.2% × 66% × 25 = 0.2MtCO₂e, which represents 1.8% of Ireland's observed power sector emissions in 2014.

context of the EU ETS in the 2013-2017 period. At the time, there was a structural oversupply of allowances on the ETS, leading market participants to *bank* more allowances (Ellerman et al., 2016). In this context, if cheaper ETS permits were purchased by non-UK power installations, these permits are likely to have been banked for future use rather than transformed in contemporaneous emissions, leaving synthetic UK emissions uncontaminated.

Overall, the combined effect of internal leakage and the waterbed effect likely contributed to increasing emissions from other European countries, but this increase is deemed to be small. Emissions in the synthetic UK are thus slightly overestimated compared to what they would be in the absence of the CPS. Taking out the simulated import spillovers and waterbed components from the synthetic UK yields a gap between the UK and synthetic UK that is about 18 MtCO₂e lower than if we assume spillovers to be insignificant. Accounting for these spillovers and all potential confounding factors would decrease abatement from 159 MtCO₂ (lower bound estimation including the potential effect of the capacity and CfD policies) to 159-18=141 MtCO₂. The average annual abatement would be -20.5%.

1.5 Discussion

1.5.1 Channels contributing to emission reduction

My results suggest that in the absence of the CPS, UK power sector emissions would have been higher by between 141 and 191 MtCO₂e. The upper bound is the SCM estimate assuming that biomass conversion was caused by the CPS, that other UK-specific policies had a negligible impact, and that spillovers were negligible. The lower bound estimate extracts from the upper bound 1) the estimated effect of biomass conversion 2) an upper bound for the CfD and capacity market, and 3) an upper bound for import spillovers and the waterbed effect.

Using the plant-level emission data, I estimate the relative contribution of three mechanisms contributing to emission reductions, based on the lower bound estimation from section 1.4.2: 1) the decrease in the emission intensity of existing plants; 2) the closure of plants having not opted out from the LCP directive; and 3) the differentiated behaviour of UK LCP opt-out plants induced by the CPS. Figure 1.5.1 shows the results from this decomposition for the UK and synthetic UK: emissions from installations present in the EUTL data every year over the 2012-2017 period are in black; emissions from installations which appear in the EUTL data or disappear from it between 2012 and 2017 (which I interpret as a plant entry in the first case and a plant exit in the second case) are in

medium grey; and emissions from LCP opt-out installations are in light grey. The difference in black areas reflects the impact of the CPS at the intensive margin (excluding LCP opt-out plants), that is to say how much more or less existing plants emit as a result of the policy. The difference in medium grey areas reflects the impact of the CPS at the extensive margin (excluding LCP opt-out plants), or how much more or less power plants enter and exit the market as a result of the policy; finally, the difference in black areas between the UK and synthetic UK before and after 2013 captures the impact of the CPS on the emission trajectory of LCP opt-out plants.

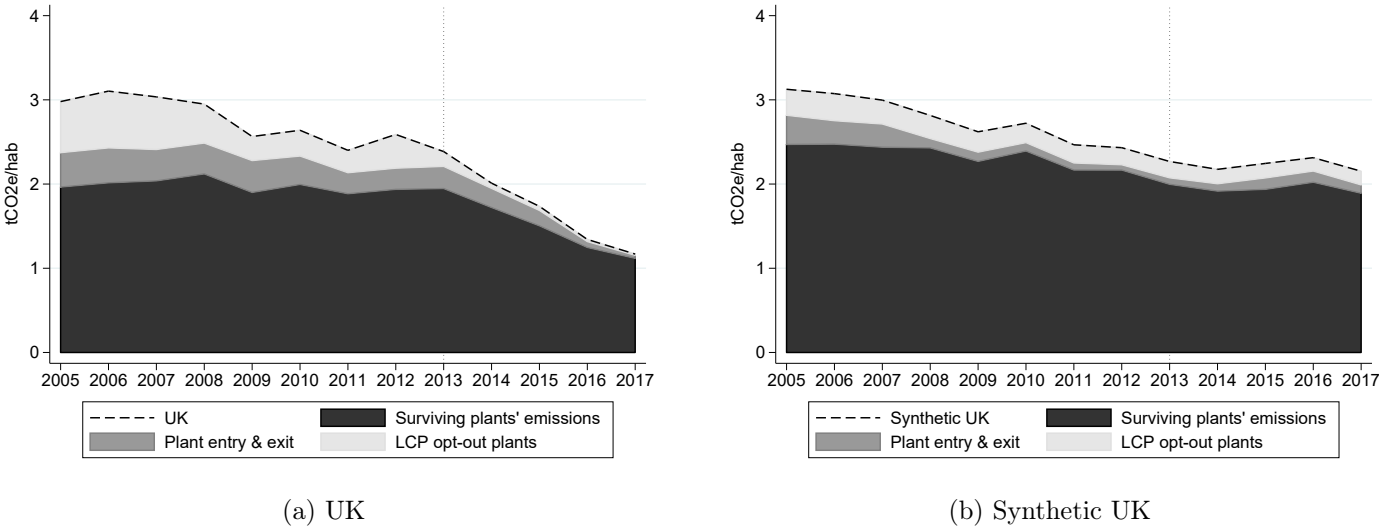


Figure 1.5.1 – per capita CO₂e emissions by source, UK and synthetic UK, lower bound estimation

Notes: For each period, the variable of per capita emissions corresponds to the sum of CO₂e verified emissions from power installations subject to the EU ETS except those in the UK converted to biomass, divided by the average country population that year. The synthetic UK comprises five countries: Ireland (49.2%), Slovakia (25.6%), the Netherlands (13.7%), Finland (5.8%), and the Czech Republic (5.7%).

Drawing on the difference-in-difference methodology, I calculate for each component the double difference between the average pre-treatment and post-treatment emissions, between the UK and synthetic UK. For the pre-treatment period, I take the average over the 2009-2012 period rather than the 2005-2012 period, given that UK and synthetic UK opt-out emissions were different before 2009. These back-of-the-envelope calculations suggest that the different behaviour of LCP opt-out plants contributed roughly 52 MtCO₂e over the 2013-2017 period (double difference in the light grey areas). The intensive margin excluding opt-out plants contributed roughly 45 MtCO₂e over the 2013-17 period (double difference in the black areas). The extensive margin excluding opt-out plants contributed roughly 50 MtCO₂e over the 2013-17 period (double difference in the medium grey areas). This extensive margin effect is the combined effect of having more emission reductions

from coal-fired plants shutting down in the UK than in the synthetic UK³⁶; and also fewer high-emission plants entering the market in the UK compared to the synthetic UK³⁷.

1.5.2 Comparison of the results with existing estimates

Two other working papers, by Abrell et al. (2019) and Gugler et al. (2020), assess the impact of the CPS on emissions reductions, using a different data coverage and different methods. In terms of data coverage, Gugler et al. (2020) consider the 2013-2015 period and Abrell et al. (2019) focus on the 2013-2016 period. Both papers also estimate emissions at the monthly or hourly level by combining hourly generation data with estimated plant-specific emission factors. Doing so, they are able to allocate emissions from January to March 2013 to the pre-treatment period, while I use annual emission data and have to allocate all the 2013 emissions to the pre-treatment period.

The annual results reported by Gugler et al. (2020) using a Regression Discontinuity in Time (RDiT) imply a cumulative abatement of 39 MtCO₂e over the 2013-2015 period. If I restrict my estimates to the 2013-2015 period, I find a cumulative abatement of between 36 MtCO₂e (lower bound from section 1.4.2) and 51 MtCO₂e (upper bound from section 1.4.1). Their results are thus included in my estimate range and are consistent with mine.

Abrell et al. (2019) estimate counterfactual generation for treated plants using a machine learning algorithm based on a short-run equilibrium model of the electricity wholesale market, and find a cumulative abatement of 26 MtCO₂e over the 2013-2016 period. If I restrict my own estimates to the same period, the abatement lies, in contrast, between 100 MtCO₂e (lower bound from section 1.4.2) and 120 MtCO₂e (upper bound from section 1.4.1). Abrell et al mention that their estimate "should best be viewed as providing a lower-bound empirical estimate of the environmental effectiveness of the UK carbon tax" (Abrell et al. (2019), p41), and the difference between our results is likely explained by the difference in data coverage and methodology. The main difference is that Abrell et al. only consider large gas- and coal- fired plants running for the whole period of analysis (from 2009 to 2016), while I include all types of power plants, including those closing or opening during the period in question. Doing so allows me to measure the impact of the

³⁶Two large coal-fired UK plants shut down in the period considered while they had neither opted-out from the LCP or IED directives: Rugeley power station closed in March 2012 and Longannet power station closed in March 2016. For Rugeley, the official reason was "a "continued fall in market prices" and increases in carbon costs" (Source: <https://www.bbc.com/news/uk-england-stoke-staffordshire-35526894>). For Longannet, the official reason was that it was "uneconomic to continue" [..]because of the high transmission charges and carbon taxes."(Source: <https://www.theguardian.com/environment/2016/mar/24/longannet-power-station-closes-coal-power-scotland>).

³⁷three new Dutch coal-fired plants entering the market in 2015 explain the increase in synthetic UK emissions at the extensive margin after 2014

CPS at the extensive margin and gauge how the CPS interacted with the LCP directive to accelerate closure, while the scope of Abrell et al is more specific to the intensive margin channel identified in section 1.5.1.

Based on the decomposition analysis from the previous section, I find an abatement of 14.5 MtCO₂e at the intensive margin over the 2013-2016 period. This estimate is closer to the order of magnitude estimated by Abrell et al. It is lower, which may be due to three factors: first, the difference in the method used; second, the difference in the emission scope: the emissions included in my "intensive margin" channels do not overlap perfectly with the emissions included in Abrell et al³⁸; third, the difference in the time frame, as my estimate for 2013 includes the Jan-March 2013 period, three winter months with a presumably high electricity consumption where the carbon tax was not yet in force.

1.6 Conclusion

My findings suggest that the carbon tax implemented in the UK power sector in 2013 resulted in a large decrease in carbon emissions due to several concurrent mechanisms: a decrease in carbon emissions of plants remaining in the market, a stronger response of opt-out plants to the LCP Directive than in countries with no carbon tax, and the closure of several high-emitting plants. While an advantage of the SCM method applied to aggregated plant-level data is that it takes into account different channels via which carbon pricing impacts emissions, it also comes with certain limitations: it does not allow for a precise estimation of the relative contributions of each of these different channels, nor for the estimation of heterogeneous treatment effects across the different treated plants.

From the point of view of its effectiveness, the CPS policy can be considered successful: the estimated abatement represents between 60% (for the lower bound estimate of 141 MtCO₂e taking into account biomass conversion, the Contract for Differences, the capacity market and spillovers) and 81% (for the upper bound abatement of 191 MtCO₂e) of the abatement necessary to achieve the targets set for the second carbon budget. While this is not the focus of this paper, other work suggests that this abatement has been achieved at a relatively low cost (Gugler et al., 2021).

Regarding the external validity of the results, it is important to keep in mind three factors that arguably enabled the tax to have such a high impact on abatement with limited carbon leakage: the relatively high potential for fuel switching from coal to gas, the relative isolation of the UK from other electricity markets that limited the risk of carbon

³⁸Some of the plants they consider are not included in my "intensive margin" emissions, but are instead in the "extensive margin" (for plants that closed in 2016, such as Rugeley) or in the "lcp opt-out" channels; essentially the other way around, some plants included in my "intensive margin" are not included in Abrell et al's analysis, such as the plants from Northern Ireland, not subject to the CPS.

leakage, and the context playing against new investments in high-emission generation. Several countries meet these criteria and could be good candidates to replicate the UK experience. Wilson and Staffell (2018) estimate that many European countries have sufficient idle gas capacity to completely eliminate coal via fuel switching, while Russia and the United States could switch 40-50% of their coal generation and China and India only 6-12%. Avoiding carbon leakage may be more difficult for countries with strong interconnections, but a solution could be to implement a carbon price floor at the regional level, as suggested by Newbery et al. (2019) for North-Western Europe. Finally, my results suggest that the interaction between increasingly stringent regulations of industrial emissions and a carbon price accelerated several plant closures. This context may also have encouraged companies operating multiple power plants to view the CPS favourably, as an instrument providing a clear price signal and making the case against coal more so than ever before. An evidence of this is that UK power companies supported the Carbon Price Floor (Hirst, 2018). Again, such a context is not likely to be found in other countries, given the global trend towards more stringent air quality regulations.

Chapter 2

Estimating the Causal Effects of Cruise Traffic on Air Pollution using Randomization-Based Inference

*co-written with Léo Zabrocki (PSE, EHES) and Marie-Abèle Bind
(Biostatistics Centre, Massachusetts General Hospital)*

Abstract: Local environmental organizations and media have recently expressed concerns over air pollution induced by cruise vessel traffic and its potential adverse health effects on the population of Mediterranean port cities. We explore this issue with unique high-frequency data spanning eleven years from Marseille, one of the largest European port cities. Using a new pair-matching algorithm designed for time series data, we create hypothetical randomized experiments and estimate the variation in air pollutant concentrations caused by a short-term increase in cruise vessel traffic. We carry out a randomization-based approach to compute 95% Fisherian intervals (FI) for constant treatment effects consistent with the matched data and the hypothetical intervention. At the hourly level, cruise vessels' arrivals increase concentrations of nitrogen dioxide (NO_2) by $4.7 \mu\text{g}/\text{m}^3$ (95% FI: [1.4, 8.0]), of sulfur dioxide (SO_2) by $1.2 \mu\text{g}/\text{m}^3$ (95% FI: [-0.1, 2.5]), and of particulate matter (PM_{10}) by $4.6 \mu\text{g}/\text{m}^3$ (95% FI: [0.9, 8.3]). At the daily level, having one additional cruise vessel entering the port increases city-level SO_2 by $0.7 \mu\text{g}/\text{m}^3$ (95% FI: [0.1, 1.4]), and PM_{10} and $\text{PM}_{2.5}$ by respectively $3.5 \mu\text{g}/\text{m}^3$ at Longchamp (95% FI: [0.5, 6.5]) and $2.5 \mu\text{g}/\text{m}^3$ (95% FI: [0.2, 4.9]) on the following day, which may partly capture an increase in road traffic. Our results suggest that well-designed hypothetical randomized experiments provide a promising approach to better understand the negative externalities of maritime traffic.

Acknowledgements: We thank Météo-France and Atmosud for their open data, the Marseille Port officials for sharing vessel traffic data, and Milena Suarez-Castillo for sharing road traffic data from the DIR Méditerranée. We are grateful to Tirthankar Dasgupta for his guidance on the computation of Fisherian interval and to Stéphane Shao for developing the matching algorithm. We thank Hélène Ollivier, Geoffrey Barrows, Philippe Quirion, Katheline Schubert, Francois Libois, Augustin Colette, Quentin Lippmann, Adam Rosenberg, Thiago Scarelli, and Georgia Thebault for their great feedback. Participants to the Paris School of Economics Applied Economics, PSIPSE and REM seminars as well as conference participants from the EAERE and FAERE provided useful comments and suggestions. Leo Zabrocki and Marion Leroutier acknowledge the support of the EUR grant ANR-17-EURE-0001.

2.1 Introduction

Particulate matter pollution induced by maritime traffic was estimated to cause 60,000 premature deaths worldwide in 2007, with the highest burden in the Mediterranean area (Corbett et al., 2007). In the past few years, local environmental organizations and media have increasingly raised concerns over air pollution induced by cruise ships' traffic (Friedrich, 2017; Chrisafis, 2018). Due to historical urban planning, many Mediterranean cities have their port in the city center and a large fraction of their population is exposed to vessels' emissions (AirParca, 2015). Besides, the Mediterranean region is not yet part of an Emission Control Area (ECA), unlike US coasts, where stringent regulations on fuel sulfur content have been implemented. In this context, understanding the influence of vessel traffic on air pollutant concentrations is a major issue of public health.

Our study focuses on the Mediterranean city of Marseille. In France, Marseille is the second largest city, with 870,000 inhabitants, and the second largest port, with 3 million passengers in 2019 (INSEE, 2020; GPMM - Port de Marseille Fos, 2020). Although pollution has decreased in Marseille in the past ten years, the concentration of several important pollutants remains high relative to the World Health Organization's recommendations and European legal standards. In 2019, Marseille was flagged by the European Commission as one of 11 French cities where daily limits of nitrogen dioxide (NO_2) have regularly been above the official threshold of $40 \mu\text{g}/\text{m}^3$, opening the way for possible sanctions. Over the 2008-2018 period, concentrations of $\text{PM}_{2.5}$ have also exceeded recommended thresholds 11% of the days, and 3.1% of the days for PM_{10} . Khomenko et al. (2021) estimate that 1.7% of total annual mortality in Marseille could be avoided if annual $\text{PM}_{2.5}$ levels decreased to the WHO recommended thresholds.

Among the causes of air pollution, car traffic has been identified for a long time. Pollution from maritime traffic is a more recent concern. A 2019 report suggested that maritime traffic had become a larger contributor to NO_2 emissions than road traffic (France Inter, 2019). This is due to the combination of two factors: on the one hand, emissions from road traffic have decreased due to the improvement in vehicles' pollution intensity - coming from a combination of stricter pollution standards at the European level and technological progress -, and some modifications in the local road network (Atmosud, 2020). On the other hand, maritime traffic intensity has increased by 50% between 2008 and 2018, as measured by the average tonnage flow into the port. This increase is driven by a three-fold increase in cruise vessel traffic, while the traffic of other types of vessels remained stable over the period.

According to emission inventories, maritime traffic contributes to 39% of the city's

nitrogen oxide (NO_x) emissions¹, 7.5% of local sulfur dioxide (SO₂) emissions, 13% of local particulate matter emissions with a diameter below 10 micrometers (PM₁₀), and 18% of local particulate matter with a diameter below 2.5 micrometers (PM_{2.5}) (AtmoSud, 2020). If maritime traffic contributes to the pollution exposure of residents in similar proportions to its emissions, it could be a key sector to target in order to improve ambient air quality. Furthermore, the neighbourhoods located close to the port are relatively deprived (Padilla et al., 2014). With recent research highlighting the heterogeneity in the causal impact of air pollution on health - poorer people being more affected than richer ones by the same pollution dose (Hsiang et al., 2019) -, the public health and welfare consequences of pollution from maritime activities could be particularly severe in Marseille and raise questions of environmental justice. Yet, isolating the contribution of vessels to observed air pollutant concentrations is methodologically challenging. Complex meteorological patterns can prevail along coastal sites, and ports are often located near major roads and industrial complexes (Sorte et al., 2020).

To identify the short-term causal effects of vessel traffic on ambient air pollution concentration in Marseille, we relate the variation in vessel traffic to the change in air pollutant concentration (Contini et al., 2011; Moretti and Neidell, 2011). We build a unique dataset combining high-frequency data on vessel traffic, weather patterns, and air pollutant concentrations over the 2008-2018 period. Our study is framed within the Neyman-Rubin Causal Model, which enables us to separate the design phase of the observational study from its statistical analysis (Rubin, 1974a; Holland, 1986b; Rubin, 2005). Using variation in vessel traffic, we emulate hypothetical randomized experiments designed to estimate the impact of an increase in vessel traffic on air pollutants.

We focus on cruise vessels, which have been particularly targeted by NGOs and the public debate around maritime traffic pollution (Chrisafis, 2018; Transport & Environment, 2019). A 2013 report on emissions by vessel type, combined with the emission inventory from that year, suggests that cruise vessels contribute around 4.5% of total NO_x emissions, less than 2.8% of SO₂ emissions and 2 to 5% of PM₁₀ and PM₂₅ emissions. Their contribution is likely higher in the recent period though, as cruise traffic has doubled while the traffic of other vessels has remained stable, and emissions from road traffic have decreased. We carry out two types of analysis: one at the hourly level and one at the daily level. We construct pairs of comparable periods using a new pair-matching algorithm designed for time series data, which allows us to adjust for observed covariates in a nonparametric manner (Sommer et al., 2018; Ho et al., 2007; Sommer et al., 2021). Randomization-based inference allows us to estimate the set of constant treatments effects that are consistent with our data and to avoid relying on asymptotic approximating

¹a generic category of pollutants including NO₂

distributions (Fisher et al., 1937; Rosenbaum et al., 2010). The causal inference pipeline we rely on can be considered as an alternative strategy to source apportionment and dispersion modelling methods (Sorte et al., 2020; Viana et al., 2014; Mueller et al., 2011; Piga et al., 2013; AtmoSud, 2018).

Section 2.2 presents the data and method used; Section 2.3 presents the results and section 2.4 discusses them.

2.2 Materials and Methods

2.2.1 Data

We built two datasets for the 2008-2018 period, one at the hourly level with 96,432 observations, and one at the daily level with 4,018 observations. Below we detail the data sources and variables used. See Appendices B.1.2, B.1.3 and B.1.4 for additional information and figures.

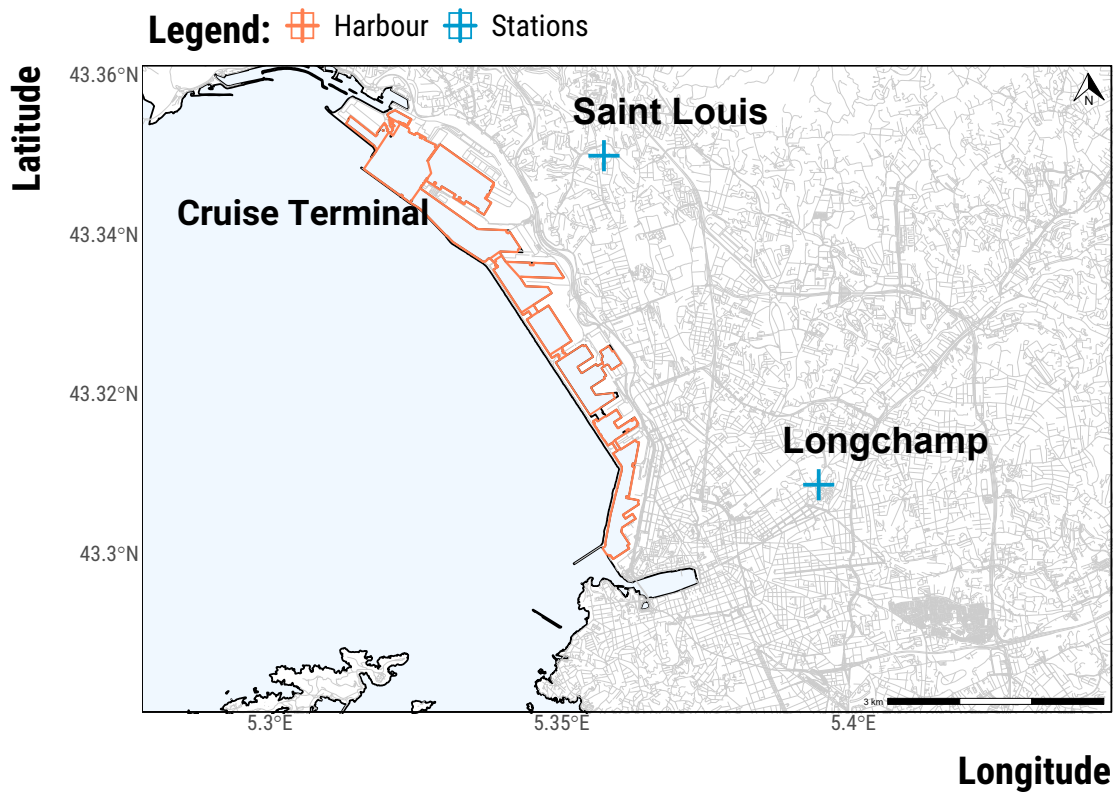
Vessel Traffic Data

We obtained data on 41,015 port calls from the Marseille Port authority. These represent the universe of all port calls between 2008 and 2018. For each vessel docking at the port, we know the exact date and hour of arrival and departure, as well as its name, its type, and its gross tonnage, which is a nonlinear and unitless measure of a vessel’s overall internal volume. This measure of a vessel’s volume can be related to its emissions of air pollutants and has been used in other studies as a proxy for the intensity of vessel traffic Contini et al. (2011); Moretti and Neidell (2011). Using information on vessel characteristics, we defined three broad categories: cruise, ferry, and other types of ships. We then calculated, for each vessel type, the total sum of gross tonnage at the hourly and daily levels. As shown in the Panel B of Figure 2.2.1, vessel traffic is regular: most vessels dock in the port in the morning and leave in the evening.

Air Pollution and Weather Data

We retrieved air pollution data from the two background monitoring stations managed by the local air quality agency AtmoSud and active for a sufficient number of periods. The first station, Saint-Louis, is the closest to the cruise terminal. It is located two kilometers away from the cruise terminal (North-Western extremity of the port) and six kilometers away from the ferry terminal (South-Eastern extremity of the port) (See Panel A of Figure 2.2.1). It only monitors NO_2 and PM_{10} . The second station, Longchamp, is located six kilometers away from the cruise terminal and three kilometers away from

(a)



(b)




Figure 2.2.1 – Localization of Marseille Port and Air Quality Monitoring Stations and Hourly Vessel Traffic Variation.

Notes: Panel A displays a map of Marseille city with its port and the two air quality monitoring stations located in Lonchamp and Saint-Louis neighborhoods. Grey lines represent the road network of the city. Panel B shows the average hourly variation in the gross tonnage of vessel arriving and departing the port. Gross tonnage is a unitless measure of the volume of a ship.

the ferry terminal (See Panel A of Figure 2.2.1). The Longchamp station monitors NO_2 , SO_2 , ozone (O_3), particulate matter with a diameter below $2.5 \mu\text{m}$ ($\text{PM}_{2.5}$) and PM_{10} ². Sulphur oxides (SOx), nitrogen oxides (NOx), and fine particulate matter are emitted to the atmosphere as a direct result of the combustion of maritime fuel (Sorte et al., 2020). SOx and NOx emissions directly produce NO_2 and SO_2 , and contribute to the formation of secondary pollutants particularly damaging for health, such as particulate matter of a larger size (i.e., $\text{PM}_{2.5}$ and PM_{10}), and O_3 (Viana et al., 2014).

Weather data come from Météo-France, the French national meteorological service. We downloaded data from the closest weather station, located 25 kilometers away from the city center, at Marseille airport. We calculated hourly and daily values for weather variables: rainfall height (mm), average temperature ($^{\circ}\text{C}$), humidity (%), wind speed (m/s), and wind direction measured on a 360 degrees compass rose where 0° is North.

Missing values were imputed using a chained random forest algorithm provided by the R package `missRanger` (Mayer, 2019).

Road traffic Data

We obtained hourly road traffic data over the 2011-2016 period from the regional *Direction Interdépartementale des Routes*, a decentralized state administration in charge of managing, maintaining, and operating roads. We selected hourly data for the six traffic monitoring stations with the best available recordings, two located North and four located East of the city. For each station, we measure three variables: the hourly flow of vehicles passing a meter; the hourly flow of vehicles longer than 9 meter passing a meter, which reflects road traffic for long vehicles such as trucks, city buses or tourist buses; and the average hourly occupancy rate, that is, the percentage of the time where vehicles are present on a given segment of the road. There is a quadratic relationship between road traffic flow and occupancy rate, where a low road traffic flow can correspond to both a low occupancy rate (few vehicles on the road) and a high occupancy rate (many vehicles on the road, congestion implies that the hourly flow is low). We averaged the three variables across the six stations to obtain average traffic data representative of road traffic in Marseille. Table 2.2.1 shows summary statistics of all the variables used in the analysis.

²Since the end of December 2018, another background monitoring station has been installed closer to the port, Verneuil station. We cannot use data from this station because our traffic data end in December 2018. In 2019, the average difference in NO_2 concentrations between Verneuil and Saint Louis stations is zero on average. In contrast, NO_2 concentrations are on average $7 \mu\text{g}/\text{m}^3$ lower in Longchamp than in Verneuil, and SO_2 concentrations are $1 \mu\text{g}/\text{m}^3$ higher on average.

Variable	Mean	SD	Minimum	Maximum
Vessel traffic				
Total gross tonnage - all vessels (unitless)	581,256	260,727	34,325	1,691,428
Total gross tonnage - cruise vessels (unitless)	189,417	231,782	0	1,245,396
Pollution				
NO ₂ (µg/m ³) at Longchamp	30.0	13.3	4	93
NO ₂ (µg/m ³) at Saint-Louis	36.5	16.3	3	91
O ₃ (µg/m ³) at Longchamp	53.8	23.0	2	121
PM ₁₀ (µg/m ³) at Longchamp	26.9	10.7	5	114
PM ₁₀ (µg/m ³) at Saint-Louis	30.4	15.2	3	176
PM _{2.5} (µg/m ³) at Longchamp	15.2	8.4	0	77
SO ₂ (µg/m ³) at Longchamp	2.3	2.5	0	32
Weather				
Average temperature (°C)	15.9	7.0	-3.5	31.3
Rainfall Indicator (% of observations)	25			
Average humidity (%)	65.2	12.4	30	96
Wind speed (m/s)	4.7	2.8	0.4	17.8
North-East Wind Direction (% of observations)	8			
South-East Wind Direction (% of observations)	24			
South-West Wind Direction (% of observations)	25			
North-West Wind Direction (% of observations)	43			
Road Traffic				
Hourly Road Traffic flow, all	2,163.0	381.8	4.7	4,392.0
Hourly Road Traffic flow, vehicles \geq 9 meters	30.3	15.0	0.0	60.5
Road occupancy rate	10.8	2.7	0.0	21.7

Table 2.2.1 – Descriptive Statistics at the Daily Level (2008-2018) (N=4,018).

2.2.2 Method

We conceptualize plausible but hypothetical randomized experiments to estimate the effects of an increase in vessel traffic on air pollutant concentrations in Marseille. We follow a “causal inference pipeline” (Sommer et al., 2018; Rosenbaum et al., 2010; Bind and Rubin, 2019; Sommer et al., 2021) conceived to analyze observational data in a rigorous and transparent manner.

Stage 1: Formulating Plausible Interventions on Vessel Traffic

We are interested in the following causal question: *Does cruise vessel traffic contribute to background air pollutant concentrations in Marseille?* The “ideal” experiment would randomly allocate hours or days to high versus low cruise vessel traffic. We could then confidently attribute the resulting differences in pollutant concentrations to vessel emissions.

In the absence of such randomized experiment in Marseille, we try to approximate an experimental setting by comparing pairs of short time series that are as similar as possible on a set of *observed* covariates but differ in their level of vessel traffic. We define below our hypothetical randomized experiments using the framework of the Neyman-Rubin Causal Model (Rubin, 1974a; Holland, 1986b; Rubin, 2005).

The units, which we index by t ($t = 1, \dots, T$), are either hours or days spanning over the 2008-2018 period, depending on the time scale of the experiment considered. At the hourly level, $V_{t,f}$ is the sum of the gross tonnage of cruise vessels docking or leaving Marseille port during hour t for flow direction $f \in \{\text{arrival, departure}\}$. We consider separately the total gross tonnage for cruise vessels' arrivals and departures as they may contribute to air pollution in the city with a different time frame. For example, cruise vessels leaving the port may start running their engines a few hours before effectively leaving, and therefore generate pollution over a long period of time. Cruise vessels entering the port may also take time to finish manoeuvring and generate emissions while they are docked. Our treatment indicator is $W_{t,f}$ and takes two values:

$$W_{t,f} = \begin{cases} 1 & \text{if } V_{t,f} > 0 \\ 0 & \text{if } V_{t,f} = 0 \end{cases} \quad (2.1)$$

Hourly units with $W_{t,f}$ equal to one are considered as "treated" while units with $W_{t,f}$ equal to zero belong to the control group. A treated hour is an hour with cruise vessels docking or leaving the port (depending on the value of f). A control hour is an hour for which there was no cruise vessel traffic for a specific flow.

At the daily level, we create an hypothetical randomized experiment which results are meaningful for policy-makers. We estimate the impact of having one additional cruise vessel entering the port, irrespective of its gross tonnage. We focus on vessels' entrances because vessels enter the port in the morning and tend to leave in the evening. By focusing on vessels' entrances, we are able to capture the impact of these vessels during the entire day where they remain docked, and not only their impact during the manoeuvring period. We define $N_{t, arrivals}$ as the number of cruise vessels entering Marseille port on day t . Our treatment indicator is W_t and takes two values:

$$W_t = \begin{cases} 1 & \text{if } N_{t, arrivals} - N_{t-1, arrivals} = 1 \\ 0 & \text{if } N_{t, arrivals} - N_{t-1, arrivals} = 0 \end{cases} \quad (2.2)$$

Daily units with W_t equal to one are considered as "treated" while units with W_t equal to zero belong to the control group. A treated day is a day with one additional cruise vessel entering the port compared to the day before. A control day is a day with a stable

cruise vessel traffic. The days with a decrease in the number of cruise vessels entering the port are discarded. In addition, we will make sure that treated and control units have the same number of cruise vessels entering the port at $t-1$. Thus, compared to control units, treated units will have the same number of cruise vessels entering the port on day $t-1$, but one additional cruise vessel entering the port on day t .

Finally, in our setting, each hourly and daily unit has two continuous potential outcomes - $Y(W_{t,f}=0)$ and $Y(W_{t,f}=1)$ for hourly units; $Y(W_t=0)$ and $Y(W_t=1)$ for daily units - whose values range in the set of plausible pollutant concentrations in $\mu\text{g}/\text{m}^3$.

Stage 2: Designing the Hypothetical Randomized Experiments

At the design stage, our goal is to obtain a sample of similar units for which the assignment to the treatment and control groups can be assumed to be unconfounded (Rubin, 1991). Formally, this unconfoundedness assumption states that the assignment to treatment is independent from the potential outcomes given a set of *observed* confounders. We use a pair-matching algorithm to obtain treated and control units with similar values for observed covariates. Matching is a nonparametric method which prunes the observations to limit the imbalance between treated and control units (Ho et al., 2007; Rubin, 2006; Stuart, 2010; Imbens, 2015). Concretely, let \mathbf{X}_t be the vector of observed covariates for each unit, with t the time indicator and $X_t^{(k)}$ the k^{th} covariate. Our algorithm matches a treated unit to a control unit only if the component-wise distances between their covariate vectors $(X_t^{(1)}, X_t^{(2)}, \dots, X_t^{(K)})$ are lower than pre-defined thresholds $(\delta_1, \delta_2, \dots, \delta_K)$. For a pair of covariate vectors \mathbf{X}_t and $\mathbf{X}_{t'}$, we use the following distance:

$$\Delta_{\mathbf{X}_t, \mathbf{X}_{t'}} = \begin{cases} 0 & \text{if } |X_t^{(k)} - X_{t'}^{(k)}| < \delta_k \text{ for all } k \\ +\infty & \text{otherwise} \end{cases} \quad (2.3)$$

To limit confounding, we select two sets of covariates. First, calendar variables (i.e., hour of the day, day of the week, bank day, holidays, month, and year) are related to both vessel traffic and air pollution. Second, weather covariates (i.e., average temperature, rainfall indicator, average humidity, wind direction blowing either from the East or West, and wind speed) could also influence both vessel traffic and air pollution. We use lags of these variables to ensure that treated and control units are as similar as possible before the treatment occurs. We define matching thresholds noting that they should be strict enough to make treated and control units comparable with each other, but not too strict to avoid reducing the sample size too much. Given this trade-off, the thresholds are stricter for the hourly experiment for which the sample size is 24 times larger.

At the hourly level, we match exactly on calendar variables over the current and two

previous hours before the treatment occurred (i.e., 0, 1, 2 lags) and allow a maximum distance of 30 days between treated and control units. For weather parameters, we carried out an iterative process, for which we tried different discrepancy values and kept the ones that led to balanced treated and control groups while resulting in enough matched pairs. We found that a maximum discrepancy of around half a standard deviation often yields a good balance. We match exactly for the East and West wind directions because they play an important role in the dispersion of air pollutants and in determining whether the monitoring station is downwind or upwind of the port terminal.

At the daily level, we create similar pairs of treated and control units over the current and previous day before the treatment occurred (0 and 1 lags). We relax some of the constraints from the hourly level to have enough matched pairs. We strictly match on the day of the week, bank days, and holidays over the two days of the series. We allow treated and control units to have up to two years of difference, but with a maximum of three months difference (for example, a Monday in December 2008 can be matched to a Monday in March 2010). For weather parameters, we match exactly on the wind direction on day t , and we allow a small discrepancy threshold for temperature and wind speed on t and $t-1$. We add as a constraint that the difference in the number of cruise vessels entering the port at $t-1$ should be zero. This way, within a matched pair, the treated unit has a similar cruise vessel traffic as the control unit at $t-1$, and exactly one more cruise vessel entering the port at t . We add matching thresholds for the pollutants for which we expect to detect an impact of maritime traffic, in order to reduce pairwise variation and improve the precision of our estimates. Finally, we add matching thresholds for road traffic flow and road occupancy rate at $t-1$ ³.

Based on these thresholds, each treated unit is matched to its closest control unit using a maximum bipartite matching algorithm (Micali and Vazirani, 1980). If no control unit is available to match a treated unit, it is discarded. Table 2.2.2 displays all thresholds values used in our matching procedure.

Stage 3: Analyzing the Experiments using Randomization-based Inference

Point estimate. As a point estimate of a Fisherian interval, we take the observed value of the average of pair differences in a pollutant concentration, which is also our test statistic. As argued by Keele et al. (Keele et al., 2012), if the assumption of a constant additive treatment effect is true, this difference in means is an unbiased estimator for the individual-level treatment effect.

³We only have road traffic data for the 2011-2016 period. We allow the algorithm to match on road traffic variables for every time period by imputing the average of road traffic variables for the years without data. This way, the road traffic matching constraint is only binding at the 2011-2016 period

	Hourly	Daily
Calendar Indicators		
Distance in days	30	720
Hour of the day in t	0	
Weekday, Bank Days and Holidays in t	0	0
Weekday, Bank Days and Holidays in $t-1$	0	0
Weekday, Bank Days and Holidays in $t-2$	0	
Month in t		3
Weather Parameters		
Average Temperature ($^{\circ}\text{C}$) in t	4	4
Average Temperature ($^{\circ}\text{C}$) in $t-1$	4	4
Average Temperature ($^{\circ}\text{C}$) in $t-2$	4	
Rainfall Dummy in t	0	
Rainfall Dummy in $t-1$	0	
Rainfall Dummy in $t-2$	0	
Average Humidity (%) in t	9	
Average Humidity (%) in $t-1$	9	
Average Humidity (%) in $t-2$	9	
Wind direction in two categories (East/West) t	0	0
Wind direction in two categories (East/West) $t-1$	0	
Wind direction in two categories (East/West) $t-2$	0	
Wind speed (m/s) in t	SD/2	2
Wind speed (m/s) in $t-1$	SD/2	3
Wind speed (m/s) in $t-2$	SD/2	
Number of cruise vessels arrivals in $t-1$		0
Average hourly road traffic flow in $t-1$		500
Average hourly road occupancy rate in $t-1$		3
NO ₂ at Saint-Louis in $t-1$		12
NO ₂ at Lonchamp in $t-1$		12
PM ₁₀ at Saint-Louis in $t-1$		12
SO ₂ at Longchamp in $t-1$		2

Notes: This table displays the maximum distance allowed for each covariate in the pair matching algorithm. Distances between treated and control units are presented for hourly and daily experiments. For example, it means that, for each matched pair, treated and control units must have the same values for weekday, bank days and holidays indicators in t . SD stands for standard deviation. If a discrepancy value is missing, it means that the associated covariate was not used in the matching procedure.

Table 2.2.2 – Maximum Discrepancies allowed for each Covariate between Treated and Control Units.

Randomization-based quantification of uncertainty. We carry out a test-inversion procedure to build 95% Fisherian (also called “Fiducial”) Intervals (FI) for the constant unit-level treatment effect. We closely follow the procedure detailed by T. Dasgupta and D.B. Rubin in their forthcoming book (Dasgupta and Rubin, 2021). Instead of gauging a null effect for all units, we test J sharp null hypotheses $H_0^j: Y_t(1) = Y_t(0) + \tau_j$ for $j = 1, \dots, J$, where τ_j represents a constant unit-level treatment effect size. We test 201 sharp null hypotheses of constant treatment effects ranging from $-10 \mu\text{g}/\text{m}^3$ to $+10 \mu\text{g}/\text{m}^3$

with an increment of $0.1 \mu\text{g}/\text{m}^3$. For each constant treatment effect j , we calculate the upper p -value associated with the hypothesis $H_0^j: Y_t(1) - Y_t(0) > \tau_j$ and the lower p -value for $H_0^j: Y_t(1) - Y_t(0) < \tau_j$. We run 100,000 permutations for each hypothesis to approximate the null distribution of the test statistic. For each experiment, running the exact number of possible allocations - equal to 2^N , with N the number of matched pairs - is computationally too intensive. The sequence of J hypotheses $H_0^j: Y_t(1) - Y_t(0) > \tau_j$ forms an upper p -value function of τ , $p^+(\tau)$, while the sequence of alternative hypotheses $H_0^j: Y_t(1) - Y_t(0) < \tau_j$ makes a lower p -value function of τ , $p^-(\tau)$. To calculate the bounds of the $100(1-\alpha)\%$ Fisherian interval, we solve $p^+(\tau) = \frac{\alpha}{2}$ for τ to get the lower limit and $p^-(\tau) = \frac{\alpha}{2}$ for the upper limit. We set our α significance level to 0.05, and thus calculate two-sided 95% Fisherian intervals. This procedure allows us to get the range of constant treatment effects consistent with our data (Rosenbaum et al., 2010; Dasgupta and Rubin, 2015), and the hypothetical assignment mechanism we posit.

In Appendix B.1.6, we provide a detailed toy example for understanding each step of our procedure.

Stage 4: Interpreting the Results

The causal interpretation of our results relies on the plausibility of the hypothetical experiment and of the unconfoundedness assumption (Rubin, 1991). This is a strong assumption as it states that the treatment assignment probability is not a function of potential outcomes given observed and unobserved confounding factors (Sekhon, 2009).

In the case of the effect of cruise vessel traffic on ambient pollution, we also note that our causal estimates could capture the road traffic induced by cruise vessel passengers. This is part of the causal effect that we want to capture. We however note that road traffic flow appear to be relatively balanced across treated and control units in the matched samples of the three experiments (See appendix B.1.5).

Finally, we insist that our results apply to our matched sample rather than to the full initial sample.

2.3 Results

We first present covariate balance diagnostics on how our matching performed. At the hourly level, we take the experiment for cruise vessel arrivals as a representative example. The matching performances of the other experiments at the hourly and daily levels are similar, as discussed in appendix B.1.5. We then display results for the effects of hourly and daily cruise vessel traffic on air pollutant concentrations.

2.3.1 Matching Results

	Hourly cruise arriving	Hourly cruise departing	Daily cruise
N_{Total}	96,432	96,432	4,018
N_{Treated}	4,034	4,037	750
N_{Control}	92,396	92,393	1,386
N_{Pairs}	131	113	84

Notes: This table displays the total number of observations, N_{Total} for each experiment, the number of potential treated and controls units before matching, N_{Treated} and N_{Control} , and the number of matched pairs, N_{Pairs} .

Table 2.3.1 – Number of Matched Pairs by Experiment.

As shown in Table 2.3.1, our matching procedure at the hourly level results in few matched treated units, with less than 4% of treated units matched to similar control units for the experiments at the hourly level. Two main reasons explain this result. First, cruise vessel traffic is regular over time (see figure B.1.2), so that it is hard to find similar hours with and without traffic. Second, even if we relax our matching constraints, it is difficult to find treated and control units with similar weather and calendar covariates. The average difference in gross tonnage between treated and control units is about 65,000 for the hourly cruise experiment, which is the average gross tonnage of one cruise ship. In Figure 2.3.1, Panel A displays the average increase in cruise vessel arrivals at hour 0 and Panel C shows that, on average, treated and control units have similar vessel traffic for other vessel types and flows. The matching improves the balance of calendar and weather covariates (see Panel B of Figure 2.3.1 for cruise vessels’ arrivals and appendix B.1.5 for the other experiments).

In Figure B.1.23 and Figure B.1.22 show the difference in weather and calendar characteristics between the initial dataset (the full hourly sample) and the matched dataset. The matched hours reflect the time of the year and hour of the day where cruise vessels arrive in the port: matched hours are mostly around 7 am, and in the spring and summer seasons. This seasonal pattern explains the higher average temperature, lower wind speed and a lower occurrence of rain compared to the initial dataset. Tuesdays and the year 2013 are also over-represented compared to the initial data (see appendix B.1.5 for similar graphs on the other experiments).

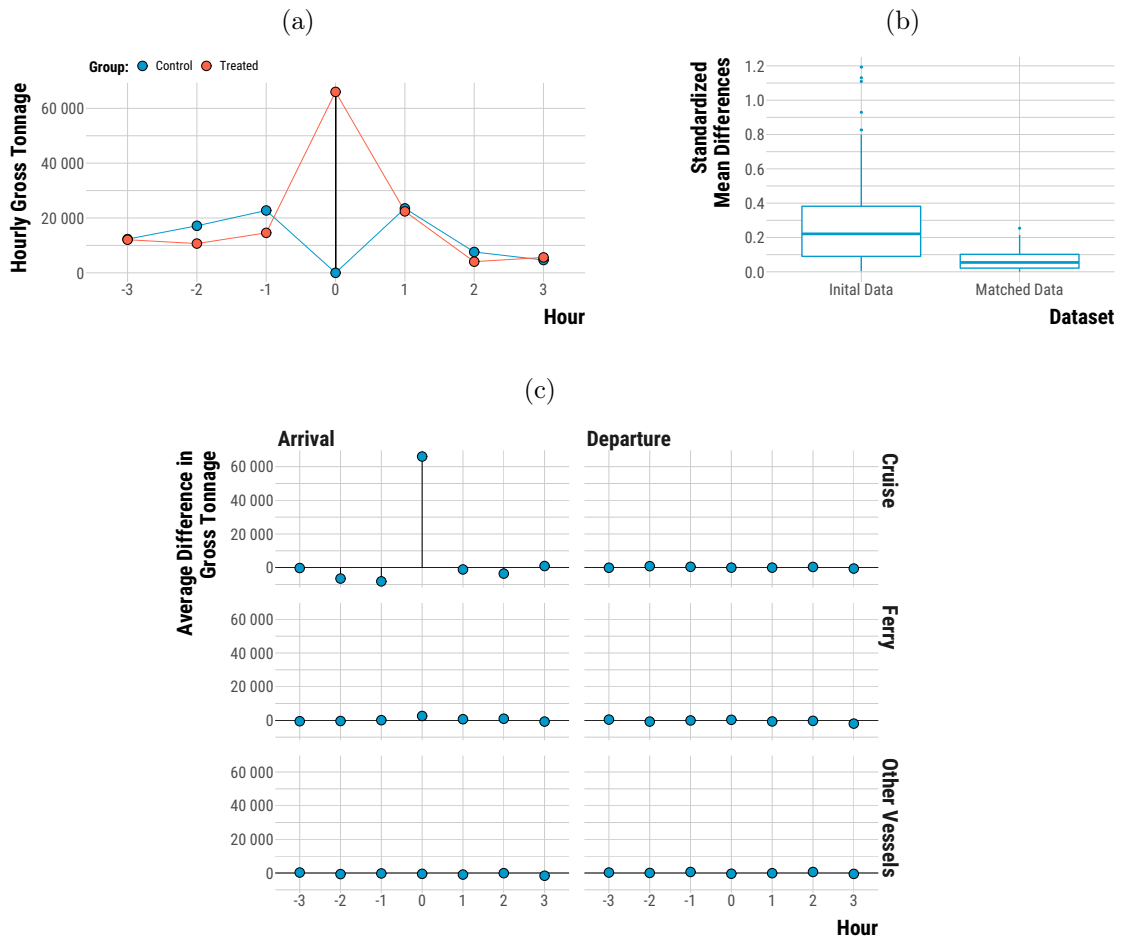


Figure 2.3.1 – Matching Diagnostics for the Cruise Arrivals Experiment.

Notes: Panel A shows the average hourly total gross tonnage for matched treated and control units in the arriving cruise experiment. Panel B displays the improvement in continuous covariates balance, measured as the standardized difference in means, for the arriving cruise experiment. Panel C plots the average difference in total gross tonnage between treated and control units by vessel type and flow for the arriving cruise experiment.

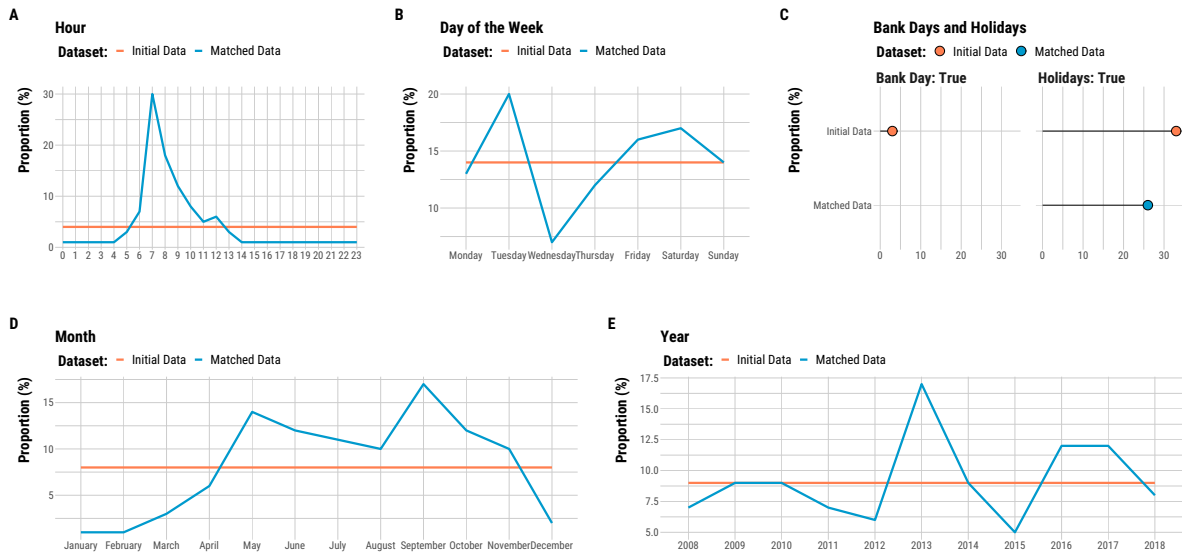


Figure 2.3.2 – Comparing Calendar Covariates Distribution for Initial and Matched Datasets.

Notes: Proportions of observations belonging to each hour of the day (Panel A), day of the week (Panel B), bank days and holidays (Panel C), month (Panel D) and year (Panel E) for the initial and matched datasets.

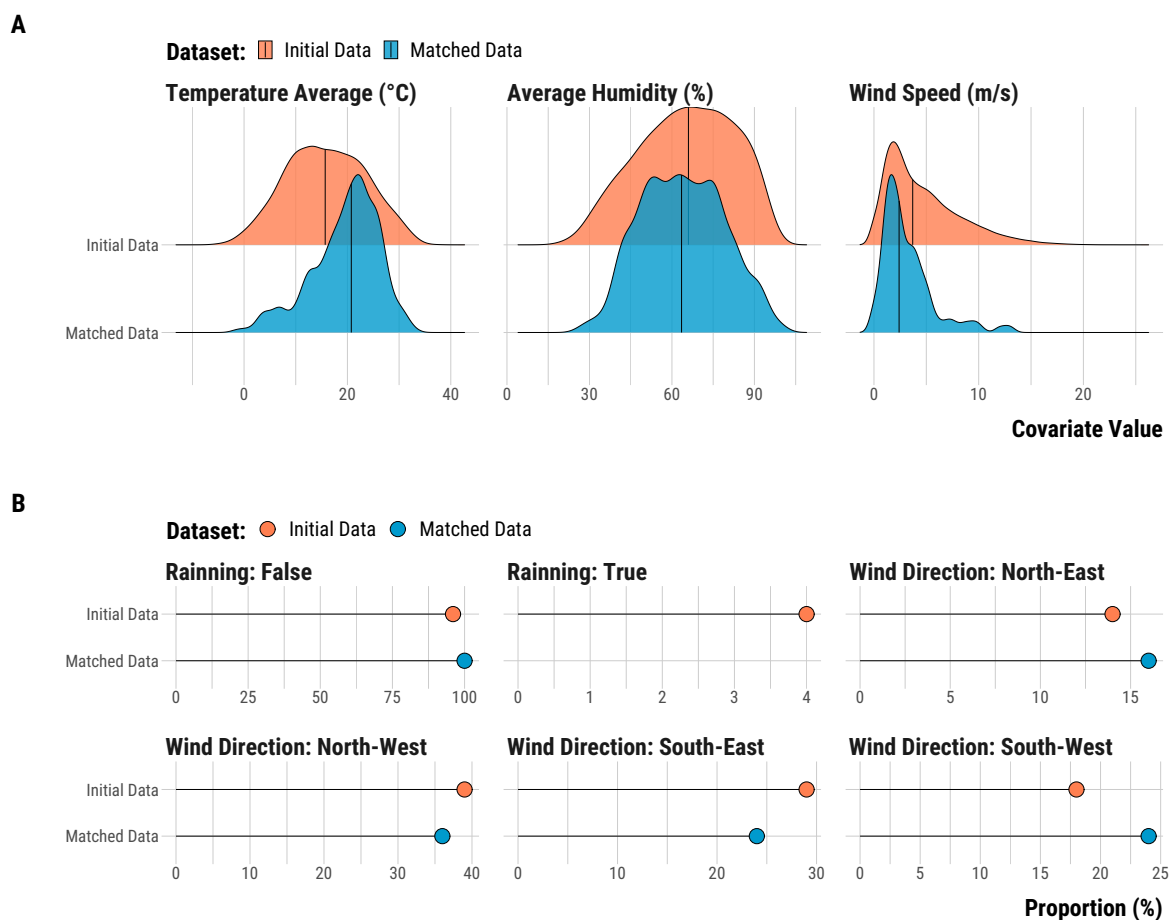


Figure 2.3.3 – Comparing Weather Covariates Distribution for Initial and Matched Datasets.

Notes: Panel A plots the density distributions of continuous weather covariates for the initial and matched datasets. Panel B displays the proportion of observations belonging to a particular category of discrete weather covariates.

At the daily level, we found 84 matched pairs, which means that 11% of the treated units were matched to similar control units. The average difference in gross tonnage between treated and control units is around 140,000 (Panel A, Figure 2.3.4). This amounts to a shock of two cruise vessels of an average size. This average intensity of treatment in terms of vessel tonnage flow likely reflects the fact that the additional cruise vessel entering the port in treated units also leaves the port in the evening of the same day. The variation in gross tonnage for other vessel types is similar across treated and control units (Panel C, Figure 2.3.4). Again, our matching procedure resulted in an improvement in covariate balance (Panel B, Figure 2.3.4). Figures B.1.50–B.1.51, we show the difference between the matched dataset and the initial dataset of potential treated and potential control units. The two datasets are similar in terms of weather covariates. Wednesdays and the years 2013-2016 and 2018 are over-represented, while weekends and the years

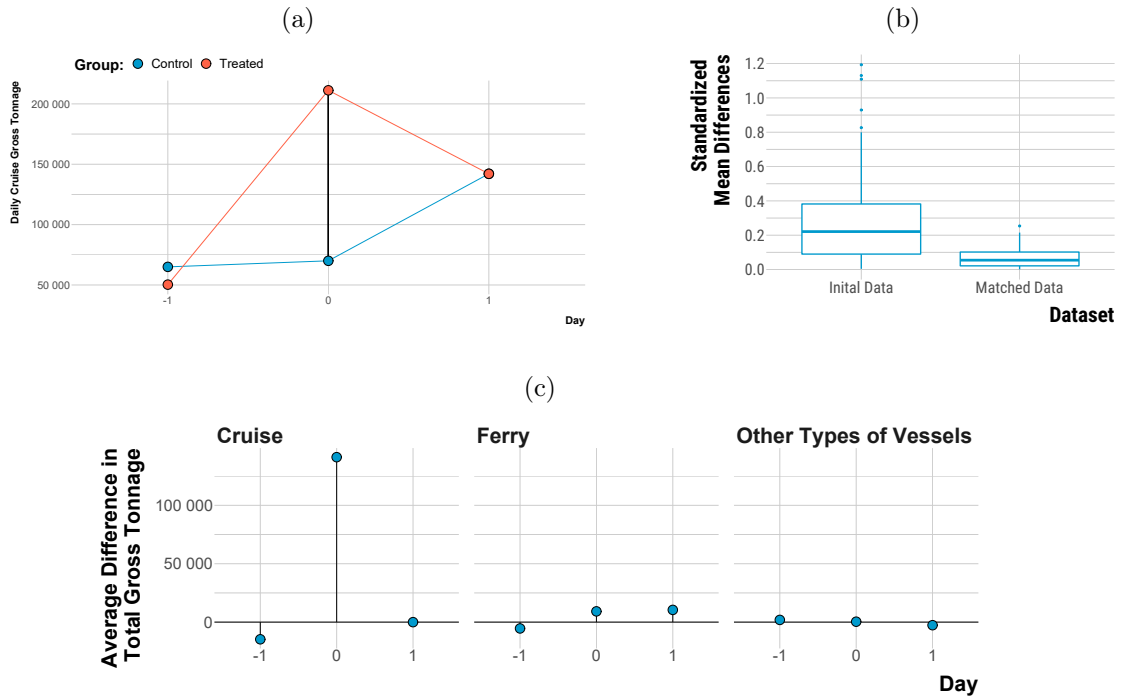


Figure 2.3.4 – Matching Diagnostics for the Daily Experiment.

Notes: Panel A shows the average daily total gross tonnage for matched treated and control units in the arriving cruise experiment. Panel B displays the improvement in continuous covariates balance, measured as the standardized difference in means, for the arriving cruise experiment. Panel C plots the average difference in total gross tonnage between treated and control units by vessel type.

2008-2010 are under-represented.

2.3.2 The Effects of Cruise Vessel Traffic on Air Pollutants

In Figure 2.3.5, we plot the point estimates and the 95% Fisherian intervals of the constant treatment effects on concentrations of NO_2 , SO_2 , and particulate matter that are consistent with our data for the two hourly experiments (see Figures B.1.60 and B.1.68 for results with more leads and lags and for ozone).

For NO_2 , results for experiments on cruise arrivals and departures are mixed. When ships dock, there is no clear pattern for NO_2 measured at Saint-Louis. We observe a clearer trend for Longchamp station, which is further away from the cruise terminal. At hour 0, the point estimate is $4.1 \mu\text{g}/\text{m}^3$ (95% FI: [0.9, 7.4]). There is, however, also a positive treatment effect before the treatment occurs: this could be consistent with the maneuvering phase of the ship as it enters the port or reveal the difficulty to obtain close pre-treatment outcomes. When vessels leave the port, We observe a decrease in NO_2 concentrations at Saint-Louis between $t-6$ and $t+1$, with a point estimate at $1.7\mu\text{g}/\text{m}^3$ at hour 0 (95% FI: [-2.2, 5.6]). The pattern is less clear for NO_2 concentrations at

Longchamp.

For SO_2 , we observe an increase in concentrations at hours 0, 1, and 2 following the arrival of cruise ships, with an increase in concentrations of $0.8 \mu\text{g}/\text{m}^3$ at hour 0 (95% FI: $[-0.3, 1.9]$). There is not a clear variation of SO_2 concentrations when cruise ships leave the port: data are consistent with small negative, null, and positive effects that are relatively high compared with the average hourly concentration of SO_2 .

For particulate matter, we see an increase in PM_{10} by $4.5 \mu\text{g}/\text{m}^3$ measured at Saint-Louis when cruise ships dock (95% FI: $[0.7, 8.3]$), and, by $2.2 \mu\text{g}/\text{m}^3$ when they leave the port (95% FI: $[-0.5, 4.9]$). We cannot distinguish a particular trend for PM_{10} and $\text{PM}_{2.5}$ measured at Longchamp station, although $\text{PM}_{2.5}$ concentrations seem to increase following the departure of cruise ships.




Figure 2.3.5 – Effects of Cruise Vessel Traffic at the Hourly Level.

Notes: Dots represent the point estimate of the unit-level treatment effect on a pollutant concentration. Lines are 95% Fisherian intervals of constant treatment effects consistent with the data. The effects are plotted from the third lag to the third lead. The treatment occurs at hour 0.

Figure 2.3.6 also shows mixed results for the daily experiment on cruise ships. At $t-1$, all the point estimates are close to zero and relatively precisely estimated because we matched on those lagged variables (except for PM_{10} and $\text{PM}_{2.5}$ measured at Longchamp). At t , we observe an increase in SO_2 concentrations: the point estimate of the 95% Fisherian interval for the effect of SO_2 at day 0 is equal to $0.7 \mu\text{g}/\text{m}^3$ (95% FI: $[0.1, 1.4]$). We also observe an increase in PM_{10} measured at Saint-Louis, but the Fisherian interval is wide: the point estimate is $2.3 \mu\text{g}/\text{m}^3$ (95% FI: $[-0.3, 4.8]$). Finally, NO_2 is increasing in both Longchamp and Saint Louis, but the Fisher exact p-value is relatively high and the intervals are wide. At $t+1$, all the particulate matter concentrations increase, which may reflect a delayed effect of cruise vessel on the following day, with the formation of

secondary particles. The point estimate for PM₁₀ is 4.2 µg/m³ at Saint Louis (95% FI: [0.7, 7.8]) and 3.5 µg/m³ at Longchamp (95% FI: [0.5, 6.5]). The point estimate for PM_{2.5} is 2.5 µg/m³ (95% FI: [0.2, 4.9]). However, for these increases at $t+1$, there could also be a confounding effect from the increase in road traffic flow observed at $t+1$ (see section B.1.6, figure B.1.79).

The two monitoring stations Saint Louis and Longchamp are located respectively to the East and South-East of the cruise terminal. In Figures B.1.61 ,B.1.69 and B.1.6, we test whether the effects at the hourly and daily levels are driven by the pairs for which wind blows from the North-West or South-West. We do so by estimating again the point estimate and Fisher intervals on the subset of these pairs. For the hourly experiment on cruise vessels' arrivals and for the daily experiment, we find that the point estimate is indeed slightly higher than on average for these pairs, but the precision does not improve.

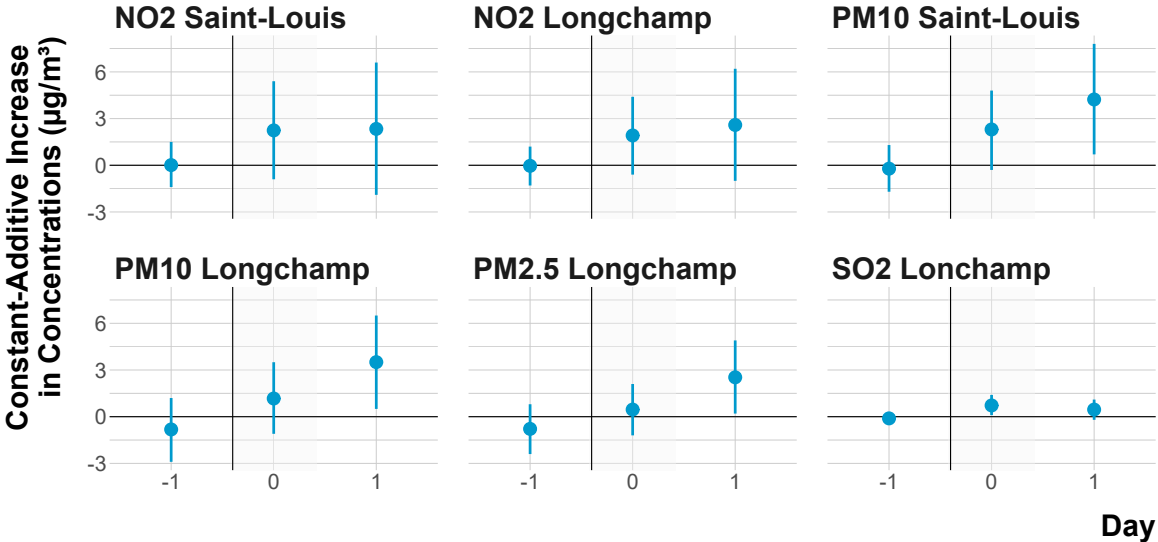


Figure 2.3.6 – Effects of Cruise Vessel Traffic on Pollutant Concentrations at the Daily Level.

Notes: Dots represent the point estimate of the unit-level treatment effect on a pollutant concentration. Lines are 95% Fisherian intervals of constant treatment effects consistent with the data. The effects are plotted from the first lag to the first lead. At Day 0, the treatment occurs.

For each experiment, we run two sensitivity analyses (see Appendix B.1.7). First, because the pair differences in pollutant concentrations were particularly disperse, we use the Wilcoxon signed-rank test statistic, known to be less sensitive to outliers. The 95% Fisherian intervals obtained with this test statistic are similar to those obtained with the average of pair differences. Second, we reproduce the analysis on non-missing concentrations because up to 20% of pollutant concentrations were imputed in our matched data.

We find similar results with slightly wider 95% Fisherian intervals.

2.4 Discussion

We start by discussing our results in view of the environmental science literature. We then reflect on the new statistical method used for our analyses. Finally, we suggest paths for future research assessing the causal impact of maritime traffic on pollution.

2.4.1 Putting our Results into Perspective

Our results point to a potential short-term effect of cruise traffic on the concentrations of NO₂, SO₂, and PM₁₀. For the hourly experiment on cruise ships' arrivals, our estimates for NO₂, SO₂, and PM₁₀ on the hour of treatment represent, respectively, 16% of the average NO₂ concentrations at Longchamp, 35% of the average SO₂ concentrations at Longchamp, and 15% of the average PM10 concentrations at Saint-Louis.

For the daily experiment, we find that having one more vessel cruise entering the port increases daily SO₂ concentrations measured at Longchamp by 30%. The point estimate for the increase in PM₁₀, for which we however fail to reject the sharp null hypothesis of no effect, represents 7.6% of average daily PM₁₀ concentrations measured at Saint-Louis. Our estimates for PM_{2.5} and PM₁₀ on the day following treatment represent around 13% of the average PM₁₀ and 16.5% of the average PM_{2.5} concentrations. However, our 95% Fisherian intervals are often wide, and the implied degree of randomization-based uncertainty can be quite large relative to the average concentration of these pollutants. On average, there is around one cruise vessel entering the port each day in the matched sample⁴. Therefore, the results from the daily experiment can be interpreted as reflecting the contribution of cruise vessel traffic on an average day of the matched sample.

Directly comparing our results to those found in the atmospheric science literature is difficult for several reasons. First, they are based on other methods - either source apportionment techniques or dispersion modelling - and usually only report average effects without comparable measures of uncertainty. Second, they often consider the entire traffic of vessels rather than isolating the impact of a pre-defined treatment focusing on one vessel type, as we do. Third, recent literature reviews have shown that contribution of vessel emissions to local air pollution depends highly on the port-city considered and the procedure carried out by researchers (Viana et al., 2014; Murena et al., 2018). We can nonetheless assess whether our causal estimates are of the same order of magnitude as estimates from the atmospheric sciences literature.

⁴with seasonal variation: the average is 0.4 in winter, 1.6 in spring, 1.4 in summer and 1.3 in autumn.

For gaseous pollutants such as NO_2 and SO_2 , the atmospheric science literature has mostly used approaches starting with emission inventories and inferring how emissions turn into concentrations using dispersion modeling Viana et al. (2014). The few studies on ports from the Mediterranean area find different contributions of maritime traffic to city-level concentrations depending on the size of the city, the location of the monitoring stations, the prevailing wind patterns, the type of boat considered and the assumptions used in the emissions inventory (Murena et al., 2018). These estimates typically take into account all the phases where a vessel may contribute to pollution, in particular the hotelling phase, while our experiment focuses on the navigation phase when vessels enter or leave the port. For NO_2 , estimates range from 1.2-3.5% for the contribution of cruise ships in summer in Naples, a city three times more populated than Marseille, to 32.5% for the contribution of all types of ships in the Italian city of Brindisi, much smaller than Marseille (Merico et al., 2016). In contrast, we failed to detect a significant increase in daily NO_2 concentrations following an increase by one cruise vessel entering the port. The estimates for SO_2 range from 1.5% for Naples in winter to 46% for Brindisi in summer. Our estimated contribution of a for SO_2 falls within this wide interval.

For particulate matter, atmospheric science studies often use source apportionment methods, consisting in monitoring particles near the port for a few months and tracing back the origin of different particles based on the statistical decomposition of chemical data acquired at monitoring sites (Sorte et al., 2020; Viana et al., 2008). Estimates for the contribution of vessels to PM_{10} concentrations in other Mediterranean ports range from 1.1% for Rijeka in Croatia up to 11% for Genoa in Italy (Merico et al., 2016; Bove et al., 2014). Our estimated contribution of cruise traffic to daily PM_{10} concentrations are broadly consistent with these studies.

The media and non-governmental organizations have insisted on the high contribution of vessel traffic, and in particular cruise vessel traffic, to city-level emissions as measured by emission inventories (France Inter, 2019; Transport & Environment, 2019). Our estimates, based on the specific impact of cruise vessels during the manoeuvring phase (and the hotelling phase on the day of arrival for the daily experiment), imply a relatively low contribution of cruise vessel traffic to concentrations. We contrast the relatively small contribution of cruise ships to NO_2 concentrations in our estimates with the contribution of road traffic, which can be inferred from a simple comparison between weekdays and weekends (see Appendix B.1.9). Because they are balanced in terms of weather covariates, the difference in observed concentrations between weekdays and weekends can be attributed to differences in economic activity only, and in particular to differences in road traffic. Road traffic decreases by 20% on average on Saturdays and Sundays. In parallel, NO_2 concentrations decrease by 20% of their average level at the Saint-Louis station.

Although other sources of pollution may be less intense on weekends, the road traffic and NO₂ time series follow an extremely similar pattern, suggesting a strong contribution of road traffic to ambient concentrations compared to maritime traffic. Beyond emission inventories informing on the relative contribution of different sectors to emissions, more systematic assessments are needed to understand the relative contribution of different sources to ambient concentrations in order to estimate the benefits of abatement in each sector and prioritize policies.

2.4.2 Reflection on the Methods

The causal inference pipeline we follow helps to clearly distinguish the design stage of our study — where we create hypothetical experiments — from its statistical analysis. Our pair-matching procedure has two notable advantages. First, it prunes treated units for which we cannot find a similar control unit, and thereby avoids extrapolating treatment effects for units without any empirical counterfactuals. In a way, a matching procedure reveals the common support available in the data from which we can draw our statistical inference. Second, our approach adjusts for covariates in a nonparametric way and achieves balance between treated and control units on observed covariates. This is another advantage, as it is often hard to guess what functional forms are needed to adjust for confounding factors (Cochran and Rubin, 1973; Ho et al., 2007; Imbens, 2015).

Yet, drawing randomization-based inferences from high-frequency observational data also poses inherent difficulties. Finding comparable treated and control units is challenging. At the hourly level, it is difficult to match a treated unit with a control unit because vessel traffic is very regular within different periods of the year. For instance, cruise vessels nearly always dock in the port at particular hours and days of the week—leaving few control hours without any cruise traffic. In addition, obtaining days with close weather patterns over several consecutive days is extremely difficult: at the hourly level, it was nearly impossible to find similar pairs over three lags of covariates.

Surprisingly, even if we strive to find similar pairs of treated and control units, we observe a wide heterogeneity in pair differences in pollutant concentrations, which makes it difficult to precisely estimate the potential contribution of vessel emissions. In our study, we are therefore confronted with a trade-off between the comparability of units within pairs and the sample size on which we base our statistical analysis. Matching on pollutant concentrations at $t-1$ as done for the daily experiment helps reducing this heterogeneity somewhat, but the variation in pair differences widens again at t , without being correlated with the intensity of the treatment (Figure).

Analyzing the full sample using a multivariate regression model delivers more precisely estimated effects and point estimates closer to 0 (see Appendix B.1.8). This regression-

based approach however does not provide an explicit imputation of the missing potential outcomes. It also relies on its ability to correctly model the functional form to adjust for confounders and to extrapolate treatment effects outside the support of the data. In the context of cruise vessel traffic, we believe that the assumptions underlying our pair matching algorithm are more plausible than those necessary for multivariate regression.

Regarding the statistical inference procedure, randomization-based inference allows us to avoid large-sample approximations and makes no assumption on the distribution of our test statistic under the sharp null hypothesis (Rosenbaum et al., 2010; Bind and Rubin, 2020). Given that we deal with small sample sizes and provided that our treatment effect assumptions are correct (e.g., constant and additive causal effect, unconfoundedness of the treatment), we believe that our procedure offers a more appropriate quantification of uncertainty in our estimates. However, randomization-based inference, as any inference mode, does not overcome issues implied by having a low statistical power to detect plausible effect sizes of cruise traffic on air pollution. In Appendix B.1.7, we carry out a *post hoc* design analysis for each experiment (Gelman and Carlin, 2014; Gelman et al., 2020). Given the size of our matched sample, we would have a low statistical power if the true effect of cruise vessel traffic on pollutant concentrations was low: estimated effects that are “statistically significant” would overestimate the true effect of vessel traffic on pollutants. Simulations should be implemented to better guide future research on this specific issue.

Last, our randomization-based inference procedure relies on the stringent assumption that the treatment is constant, while it might have been of interest to estimate the average treatment effect and quantify its uncertainty for each hypothetical experiment. We therefore provide results from a Neymanian inference perspective (Imbens and Rubin, 2015; Splawa-Neyman et al., 1990). In Appendix B.1.7, we calculate for each experiment the estimates of the average treatment effects and their associated 95% confidence intervals. Although based on a different interpretation of the data, results from Fisherian and Neymanian inference are substantively similar. We could also have implemented a Bayesian model-based approach, which explicitly imputes the missing potential outcomes given the observed data and can target a variety of estimands (Bind and Rubin, 2019; Imbens and Rubin, 2015; Rubin, 1978).

2.4.3 Potential Paths for Future Research

We see at least three main improvements for future research on the effects of maritime traffic on air pollution. First, it would be useful to exploit data on the duration vessels keep their engines running while docked at the port, as several studies indicate that a large share of air pollutant emissions occur during this phase (AirParca, 2015; Murena et al.,

2018). Second, monitoring stations in Marseille only measure some pollutants and are located relatively far away from the port. It would be useful to carry out similar analyses as ours in a port city where pollutants such as ultrafine particles are monitored and with receptors located in the port at different heights (Viana et al., 2014; Mocerino et al., 2020). Third, several areas have implemented regulations to decrease the sulfur content of vessel fuel. This type of policy is particularly well-suited for causal inference methods such as interrupted-time series, difference-in-differences, and synthetic control (Kotchenruther, 2017; Grange and Carslaw, 2019; Zhu and Wang, 2021). Researchers could estimate how pollutant concentrations evolved before and after the policy was implemented by comparing the treated area to control areas.

2.4.4 Concluding Remarks

Our study is a complementary approach to current source-apportionment and dispersion modeling methods. We provide very detailed replication materials in the hope that researchers could implement our method for other ports. We believe that well-designed observational studies relying on a causal inference pipeline could bring new insights on the environmental and health consequences of maritime activities.

Chapter 3

Tackling Transport-Induced Pollution in Cities: A Case Study in Paris

co-written with Philippe Quirion (CNRS, CIRED)

Abstract: Urban road transport is an important source of local pollution and CO₂ emissions. To tackle these externalities, it is crucial to understand who contributes to emissions today and what are the alternatives to high-emission trips. We estimate individual contributions to transport-induced emissions, by bringing together data from a travel demand survey in the Paris area and emission factor data for local pollutants and CO₂. We document high inequalities in emissions, with the top 20% of emitters contributing 75-85% of emissions on a representative weekday, depending on the pollutant. Top emissions result from a combination of high distances travelled, a high reliance on car and, mainly for local pollutants, a higher emission intensity of cars. We estimate with counterfactual travel times that 53% of current car drives could be shifted to electric bikes or public transport with a limited time increase. This would reduce the emissions from daily mobility by 19-21%, with corresponding annual health and climate benefits of around €214m.

Acknowledgements: We thank Nicolas Laruelle, Louis-Gaëtan Giraudet, Francois Bareille, Francois Li-bois, Laure de Preux, Olivier Chanel, Mouez Fodha, an anonymous referee from FAERE and participants from the CIRED and AgroParisTech seminars for helpful comments and discussion; from Airparif, Alexia Baudic and Cécile Honoré for giving us access to air pollution data, Geraldine Le Nir for her comments, and Olivier Perrussel for technical advice on the estimation of emissions; from Ile de France Mobilités, Olivier Mahieu for helpful comments and for giving us access to simulated distance data. We thank Paul Dutronc-Postel for helping us retrieve transport time with the Google Maps API. Marion Leroutier thanks ANR for the support of the EUR grant ANR-17-EURE-0001.

3.1 Introduction

Road transport is responsible for several well-documented environmental externalities (Parry et al., 2007). First, it contributes to outdoor air pollution, which has been identified by the WHO as the “world’s largest single environmental health risk” (WHO, 2014), accounting for an estimated 4.2 million deaths per year. Beside its impact on physical health, air pollution negatively impacts mental health (Bishop et al., 2018; Braithwaite et al., 2019), the formation of human capital (Currie et al., 2014) and productivity (Chang et al., 2019). Road transport also contributes to greenhouse gas emissions, mostly carbon dioxide (CO_2), with an increasing contribution relative to other economic sectors in most developed countries (IEA 2019). This trend needs to be reverted to achieve emission reductions consistent with the Paris agreement.

This paper focuses on local pollutant and CO_2 emissions from transport in urban areas, where emissions are both more detrimental to health and possibly easier to tackle than in rural areas. On the first point, many urban areas suffer from high levels of pollution, including in developed countries subject to relatively strict environmental regulations: in Europe, France, Germany and the UK were condemned in 2018 for failing to meet air quality standards in several cities (European Commission, 2018). On the second point, urban areas present more alternatives to cars: the higher density makes active modes more attractive, and public transport is more widespread (Creutzig et al., 2020). Yet, policy proposals aiming at restricting driving for polluting cars, whether motivated by air quality or climate mitigation concerns, are controversial (Viegas, 2001; Le Parisien, 2019; Delhaes and Kersting, 2019; Isaksen and Johansen, 2020). It is then crucial to understand who the high emitters are, and whether they could easily switch to a low-emission alternative.

In this paper, we estimate how much individuals contribute to transport-related pollution via their daily travels. To do so, we combine individual travel information from a large representative survey conducted in the Paris area with mode-specific and vehicle-specific emission factors. We focus on two local pollutants having detrimental effects on health, nitrogen oxide (NO_x) and fine particulate matter ($\text{PM}_{2.5}$), and the main greenhouse gas, carbon dioxide (CO_2). We find strong inequalities in emissions among individuals, with the top 20% of emitters contributing 75-85% of emissions on a representative weekday, depending on the pollutant.

We then investigate the characteristics associated with high emissions using two complementary methods: in a first step, we note that total emissions are the exact product of three channels: distance, modal choice, and emission intensity (per kilometer.passenger and within modes). We apply an exact factor decomposition analysis (LMDI) on emissions to understand how the respective contributions of these three channels to top emissions.

For local pollution, higher distances travelled, a higher reliance on car, and higher emission intensities within modes contribute about the same to top emissions. In contrast, for CO₂ top emissions are mostly explained by high distances and a high reliance of cars, and less by differences in emission intensities.

In a second step, we investigate the individual characteristics associated with each of the three channels, focusing on car - the most emitting mode - for modal choice and emission intensity. Beyond the characteristics already documented in the literature, distance to the centre or employment status, we highlight the important association of some employment characteristics with the reliance on car, such as being a manual or trades worker, a self-employed white-collar, working in a factory or with atypical working hours. We also show the ambivalent role of income, which is associated with higher distances, a higher probability to use a car and a higher CO₂ emission intensity of cars, but not with a higher NO_x and PM_{2.5} emission intensity.

Finally, we investigate the potential to reduce emissions. We use counterfactual transport time data from a transport Application Programming Interface (API) to estimate the modal shift potential. We find that 53% of current car drives could be shifted to regular, electric cycling or public transport against an increase in travel time of at most 10 min per trip and limited daily increase in travel time. Such modal shift would save 21% of the total NO_x emissions induced by passenger daily mobility, 19% of the total PM_{2.5} emissions and 19% of the total CO₂ emissions. We document, with less precision, the potential for teleworking (distance lever) and shifting to electric vehicles (emission intensity lever).

Our paper contributes to several strands of the literature: first, we contribute to the literature on environmental inequalities by investigating individual contribution to transport-related pollutants and CO₂ emissions. There is a vast literature examining cross-country inequalities in local pollution emissions - in relation to the Environmental Kuznets Curve hypothesis (Dinda, 2004) -, and a more limited literature examining inequalities at the individual or household level (Levinson and O'Brien, 2018). On CO₂ emissions, there is also flourishing literature looking at inequalities in individual carbon footprint at the country or regional scale (Sager, 2019; Ivanova and Wood, 2020). Most of the studies estimating individual emissions rely on input-output methodologies combined with micro-level consumer expenditure surveys, which provide very limited information on travel behavior (mostly the purchase of fuel and public transport tickets and subscription), and lack precise spatial information. As far as we know, the subset of papers specifically examining the incidence of carbon tax in relation to transport emissions also relies on consumer expenditure surveys as far as we know (Douenne, 2020; Cronin et al., 2018). Our paper is closer in spirit to studies from the transport and urban planning liter-

ature estimating (inequalities in) individual emissions from transport using detailed travel diaries from a sample of individuals (Brand and Preston, 2010; Yang et al., 2018; Brand et al., 2021). An important limitation of these studies, however, is to rely on low sample sizes, and, often, on non-representative surveys where highly educated individuals are over-represented. In contrast, we use a large representative survey ($N=23,690$). Finally, although we do not examine a policy in particular, our paper is connected to previous work having estimated the distributional impacts of different transport policies affecting car use in the Paris area, such as Bureau and Glachant (2008), Bureau and Glachant (2011) or Bou Sleiman (2021), analysing respectively the distributional impact of road pricing, of reducing the cost of public transit, and the displacement effect of closing urban expressways.

Second, we contribute to the literature examining the potential for emission reductions from transport, in particular the potential for modal shift (Javaid et al., 2020; Yang et al., 2018). By using data on trip duration by mode retrieved from a transport API, we are able to estimate the share of trips that can be done with another mode than car, based on observed individual travel data. Compared to previous work focusing on the potential for modal shift for short trips specifically (de Nazelle et al., 2010), or restricting the analysis to a modal shift to public transport or bike (Yang et al., 2018), we do not set a limit on trips' distances and we allow for substitution with an under-investigated transport mode, electric bike, which we show has a high potential.

Third, we contribute to the literature examining the trade-offs and complementarities in tackling both CO_2 and local pollution. Durrmeyer (2018) and Linn (2019) show that while effective in decreasing CO_2 emissions, CO_2 -based vehicle taxes are likely to increase the emission of damaging air pollutants (NO_x and $\text{PM}_{2.5}$), because they increase the share of diesel cars, less CO_2 -intensive but more intensive in NO_x and $\text{PM}_{2.5}$. The reverse trade-off may exist in the case of local transport policies driven by air pollution concerns, and low-emission zones indeed tend to be more restrictive for diesel cars than for gasoline cars. Our results suggest that a policy targeting cars' local pollutant emission intensity may also have different distributional impacts from a policy targeting the CO_2 emission intensity, since we find different associations between household income and the $\text{PM}_{2.5}$ vs. CO_2 emission intensity of car trips. At the same time, policies leading to a modal shift away from car would achieve both a reduction in local pollutant and CO_2 emissions.

The paper is organized as follows: Section 3.2 presents the local context of air pollution in the Paris area; Section 3.3 presents the data and methods used; Section 3.4 presents the results and section 3.5 discusses them.

3.2 Air pollution and transport emissions in Paris

We consider the Paris metropolitan area, which we define here as the administrative *region* of Ile de France (IdF), represented on Figure 3.2.1a - the *region* is the first level of administrative subdivision in France.¹ The IdF region has a population of 12.2 million inhabitants and is made of three layers: the city of Paris in the centre (red), a first layer around Paris called the “inner suburbs”, made of three small *départements* (blue) - the second level of administrative subdivision in France, and a second layer called the “outer suburbs”, made of four larger but less dense *départements* (yellow).

The Paris area is a typical monocentric city where most public transport lines converge to the centre. Air pollution levels regularly exceed recommended and legal thresholds. Figure 3.2.1b shows NO₂ concentrations in 2015 and shows that the legal threshold of 40µg/m³ is exceeded in Paris and the majority of the inner suburbs. While concentrations of the main regulated pollutants² have been decreasing over the past ten years, as shown on the maps in Figures C.1.4 and C.1.5, they remain high in the city centre for Nitrogen dioxide (NO₂) and PM_{2.5}³. Furthermore, despite the improvement in air quality, air pollution is the number one environmental concern in IdF according to a 2018 survey, and 61% of the respondents think that air pollution has increased in the past ten years.

In this paper, we focus on emissions of two local air pollutants: NO_x, a generic category of pollutants including NO₂, and PM_{2.5}. We choose these pollutants for two reasons: first, road traffic is a major contributor for these two pollutants: it is responsible for 56% of nitrogen oxides (NO_x) and 35% of the PM_{2.5} emissions of the region (Source: Airparif). Second, these two pollutants have detrimental effects on health: long-term exposure to NO₂ is associated with increases of bronchitis in asthmatic children and reduced lung function growth (World Health Organization, 2018). Exposure to PM_{2.5} has detrimental effects on health and increases mortality risk in the short- (Deryugina et al., 2019) and long-term (Lepeule et al., 2012), without evidence of a threshold below which exposure would be harmless (World Health Organization, 2018). We also study the emissions of CO₂ emissions, road traffic being responsible for 32% of the region’s total emissions (Source: Airparif).

To dampen local pollution from cars, several regional and local policies have been implemented. Short-term driving restrictions based on license plate numbers have been systematically imposed since 2014 during pollution peaks. Long-term measures advertised

¹The Paris metropolitan area as defined by the French statistical institute does not include all the IdF region; it excludes a small part of the outer suburbs. We consider the whole region because our transport data are representative of the population from the entire region

²nitrogen dioxide NO₂, ozone O₃, and particulate matter PM₁₀

³in contrast, ozone is higher in rural areas

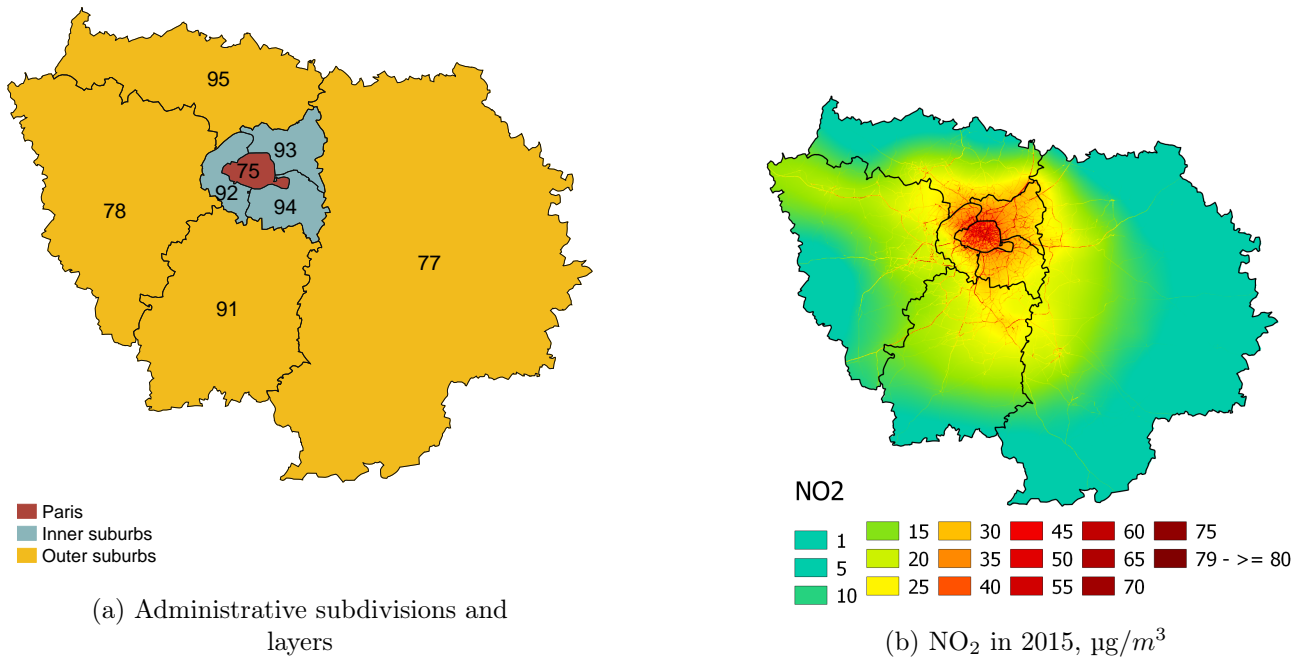


Figure 3.2.1 – The Paris area

Note: The black line shows the division of each area in *département*. The numbers are the administrative identifiers for each *département*: 75-Paris; 92-Hauts-de-Seine; 93-Seine-Saint-Denis; 94-Val-de-Marne; 77-Seine-et-Marne; 78-Yvelines; 91-Essonne; 95-Val d’Oise

by the regional authority include developing the public transport network, building more cycling lanes, reserving lanes for buses, clean vehicles and car-pooling, as well as speed reduction on the ring road (Région Ile de France, 2016). By far, the most ambitious policy specifically targeting air pollution is the Low Emission Zone (LEZ) projected to be rolled-out in Paris and the surrounding municipalities between 2017 and 2024, which should progressively ban all polluting vehicles - defined by their age and fuel type - from the city centre. Yet this policy has met political opposition from some municipal authorities (Le Parisien, 2019). To reduce both local air pollution and CO₂ emissions from cars, the Paris metropolitan area also announced the complete ban of diesel cars by 2024 and of gasoline cars by 2030 (Le Monde, 2018).

3.3 Data and methodology

3.3.1 The Data

We combine four types of data from different sources.

Individual transport: We use transport data from the 2010 wave of the EGT (*Enquête générale des transports* - EGT 2010-STIF-OMNIL-DRIEA), a survey conducted every 8

to 10 years in the IdF region. The 2010 wave was conducted between October 2009 and May 2010, and between October 2010 and May 2011. The survey contains detailed information on the transport choices of 35,175 individuals from 14,885 households⁴ on a given weekday⁵. The sample is representative of the IdF population (as characterized in the 2008 census) in terms of household size, type of housing and individual socio-economic and demographic profiles⁶. The EGT is also broadly representative of the 2011 IdF population (see Table C.1.3, comparing selected household characteristics for the entire EGT sample and from 2011 administrative data). The EGT also has detailed socio-demographic characteristics (see Table C.1.2 for descriptive statistics of the whole sample at the household level). For the present analysis, we use the subsample of adults having done at least one trip during the weekday (N=23,690). This represents 93.07% of the surveyed adults. Table 3.3.1 shows descriptive statistics for this subsample. In table C.1.4 we present a balancing test comparing mean observed characteristics for the full sample and the sample of adults with at least one trip recorded. The two samples are broadly balanced. There are small significant differences in the activity status, with a higher share of full-time employed individuals and lower shares of inactive individuals in the sample of mobile individuals. Mobile individuals are also more educated and have a higher income compared to the whole sample. We consider that this small selection bias is not an issue given the descriptive purpose of our analysis.

The survey records and geolocates all the places visited by each individual during the day with a grid size of 100 meters*100 meters. Within each trip defined by an origin and destination location, the data describes each journey stage, a journey stage being defined as a single travelling mode⁷. Only the trips starting or finishing within the IdF boundaries are recorded and geolocated. For all the trips starting (finishing) in the IdF region but finishing (starting) in another region, we do not know the departure (arrival) point's location, nor the trip distance. We use three variables in the analysis which are not readily available in the EGT data:

- **Actual distances travelled:** The EGT data only contains as-the-crow-flies distances for each trip and journey stage. We obtained data on actual distances from

⁴The sampling rate at the household level is 1/330. In 2010, the IdF region had a population of 11.79 millions inhabitants

⁵The respondents are asked about all their trips from the day preceding the interview, which can correspond to a day between Monday and Friday. We include survey day-of-week fixed effects in all our regression analyses.

⁶based on 30 categories combining gender, age, socio-professional category and main occupation

⁷For example, a work commuting trip by subway including one change will include four journey stages: the first stage is the journey by foot from home to the subway station; the second stage is the subway journey with the first metro line, finishing at the subway station where the commuter changes lines; the third stage is the subway journey with the second metro line, finishing at the subway station near the workplace; the fourth stage is the journey by foot from the subway station to the workplace.

the regional transport authority, estimated with a shortest-path algorithm.

- **Continuous income variable:** in the EGT data, household income is self-declared and interval-coded in nine income brackets, with a non-response rate of 6%. In order to estimate the relationship between income deciles and contribution to emissions, we estimate the full distribution of income using an interval regression imputation method. Since the method assumes an underlying normal model for the partially observed imputed variable (given other predictors) and the distribution of income is usually log-normal, we apply a log transformation to the income brackets declared in the EGT. We then estimate the continuous income variable by including several socio-economic factors known to be correlated with income in the interval-coded regression.⁸ For households with a missing income bracket, we use a predictive mean matching imputation method, using the same predictors and similarly predict their continuous income. Finally, we transform the obtained continuous variable of household monthly income into a variable of annual income per consumption unit (using the OECD equivalence scale). Table C.1.3 shows that the average income per consumption unit obtained with this imputation is close to the average income per consumption unit in IdF in 2011 obtained from administrative data.
- **Public transport stops within a one kilometre radius:** We create an indicator variable indicating whether a household lives less than one kilometer away from a public transport stop. To do so, we combine geocoded information on the location of each public transport stop in 2010 contained in a separate EGT file (including subway, regional train and streetcar), with information on households' place of residence.

Emission factors We use emission factor data by transport mode (and by type of vehicle for cars and two-wheelers) coming from a variety of sources, described in detail in the next section and in Appendix C.1.1.

Counterfactual travel time data To estimate modal shift options for car drivers, we estimated travel time for different transport modes for all the non-walking trips reported in the EGT data. This represents 68,110 trips made by 20,725 individuals, including 33,010 car drives made by 10,875 individuals. For each trip, we identified the departure and arrival points with the latitude and longitude of the centroid of the origin and destination

⁸List of predictors: age, age squared, gender, education level and socio-economic class of the household head; socio-economic category of her partner; number of household members working full-time and number working part-time; housing status of the household; dummy for whether the household is eligible to family allowances based on the number and age of children, to proxy for social transfers.

	Mean	Sd
Residence: Paris	0.21	0.41
inner suburbs	0.37	0.48
outer suburbs	0.42	0.49
Age	45.72	16.62
Education: Primary school	0.06	0.23
Secondary education	0.39	0.49
Higher education < 3 years	0.14	0.35
Higher education \geq 3 years	0.35	0.48
Still in education	0.07	0.25
SES: Farmers	0.00	0.03
Manual workers	0.11	0.31
Office workers	0.19	0.39
Intermediate professions	0.19	0.40
Traders and craftspeople	0.03	0.17
Managers and executives	0.20	0.40
Pensioner	0.20	0.40
Other	0.07	0.26
Estimated Net household income	40,786.73	25,901.23
Estimated Net household income per consumption unit	24,166.10	14,626.03
Distance to workplace (km)	10.57	10.69
Nb of trips prev. day	4.32	2.40
Modal share for trips: car	0.39	0.44
collective transportation	0.27	0.38
bicycle	0.02	0.11
two-wheeler	0.02	0.11
walking	0.31	0.37
other mode	0.00	0.05
Observations	23,690	

Table 3.3.1 – Summary statistics - Individuals ≥ 18 years old

Note: Source: EGT data. Data weighted with EGT individual-level sampling weights. SES stands for Socio-Economic Status. The eight categories follow the aggregate classification of the French Statistical Institute.

squares. We then used the Google Console Directions API to predict each trip’s duration for three different transport modes: driving, cycling and public transit. Our trip requests gave results for more than 99.9% of the cases for car and cycling trips and for 85% of the cases for public transit trips (34% of which suggested walking as the fastest way to arrive at destination). For the remaining 15% of trips, no public transit route exists between the departure and arrival point. Appendix C.1.2 provide more details on the exact data request and compares travel times declared in the survey to those given by the API. The predicted times of the API are 20-39% lower than those declared by individuals, depending on the mode. Since we base our calculations on the API times only, it should not bias our estimation of modal shift potential too much. We use the API times for cycling trips to estimate the duration of the same trips had they been done with electric bikes instead of regular bikes. We assume an average cycling speed of 15km per hour and an average electric cycling speed of 19km per hour, following the figures from a 2015 survey ⁹, and apply this constant speed factor to the estimated travel times of cycling trips.

Charging stations for Electric Vehicles We use a GIS software to identify all the households having at least one electric vehicle (EV) charging station less than 500 meters away from home. We did not find an exhaustive dataset of all charging stations located in the IdF region. We instead combine geocoded data from four different sources: OpenStreetmap¹⁰ (where many stations located in Paris centre are missing), the national open data service¹¹ (where many stations located in Paris centre are also missing), and subregional open data services providing data on two cities (Paris and Rueil-Malmaison).

3.3.2 Methodology

Building individual measures of contribution to pollution. We estimate individual- and trip-level contributions to local and global pollution based on the detailed information contained in the EGT. For local pollutants, we use NOx and PM_{2.5}. For global pollution, we use CO₂ emissions. The total emissions of pollutant P for individual i during the day are the sum of her emissions at the trip level, with T the total number of trips made during he day:

$$E_{P,i} = \sum_{t \in T} E_{P,i,t} \quad (3.1)$$

⁹The survey was conducted in four European countries including France <https://6-t.co/etudes/donnees-inedites-va-e-en-europe-panel/>

¹⁰https://geodatamine.fr/dump/charging_station_geojson.zip

¹¹<https://www.data.gouv.fr/fr/datasets/fichier-consolide-des-bornes-de-recharge-pour-vehicules-electriques/>

We calculate emissions at the trip level $E_{P,i,t}$ using information on each journey stage j that t is made of. For each journey stage j , we know the (estimated) journey distance in kilometers, the travel mode used, the emission factor associated with the mode, and the number of passengers if the mode used is a private vehicle (car or two-wheeler).

i 's emissions of Pollutant P on trip t made of J journey stages are defined as:

$$E_{P,i,t} = \sum_{j \in J} d_{j,i} e_{P,j,i} r_{j,i} \quad (3.2)$$

Where $d_{j,i}$ is the distance travelled by i on journey stage j , $e_{P,j,i}$ is the pollutant P 's emission factor associated with travel mode m used in journey stage j in grams per kilometre driven; $r_{j,i}$ is the inverse of the occupancy rate¹² of mode m for individual i for journey stage j . For all the journey stages done with a collective transport mode, the occupancy rate is set to one, as an average occupancy rate is included in the estimation of their emission factor.

The assumptions made to estimate emission factors for each mode are explained below and more extensively in Appendix C.1.1. Active modes (walking, cycling, skate-boarding, etc.) have a zero emission factor for all three pollutants. The train and subway have a zero emission factor for NOx and CO₂¹³, but not PM_{2.5}, due to the emissions from train brakes. For transportation modes with positive emission factors - buses, two-wheelers and cars for NOx and CO₂, plus electric public transport for PM_{2.5}, we use a combination of sources described in Appendix C.1.1.

Emission factors can exist in two versions for cars: the "true", on-road emission factor, which varies with the vehicle speed, quality of the road and driving conditions; and the type-approval values, given by car manufacturers and subject to emission standards under EU rules. We use on-road emission factors for NOx and PM_{2.5}, but type-approval values for CO₂, for several reasons: first, the discrepancy between type-approval and real-world emissions is much stronger for NOx than for CO₂ emissions¹⁴, so it matters more to correct the NOx emission factor than the CO₂ emission factor. Second, for cars' emission factors, there exists a rich vehicle-specific data source for type-approval CO₂ emission factors but not for NOx and PM_{2.5}. Using it allows us to estimate CO₂ emission factors based on all the car characteristics declared by the household, in particular horsepower,

¹²The occupancy rate is defined as the number of passengers in the vehicle.

¹³These modes embody some NOx and CO₂ emissions, but given our focus on air pollution mitigation in the Paris area, we think it is satisfying to focus on exhaust emissions only.

¹⁴Baldino et al. (2017) compare on-road and type-approval emission factors for recent diesel vehicles brought under the spotlight by the 2015 Volkswagen scandal. They find that average on-road CO₂ emissions are on average 30% higher than laboratory emission standards and type-approval values for a sample of Euro 5 and Euro 6 cars (registered after 2011), and report that the gap has been increasing over the 2001-2015 period. For NOx emissions, they find a much higher discrepancy with an average factor of 4 between the type-approval and real-world values.

a variable likely to be correlated with households' socioeconomic status. Since we seek to identify the socio-economic and spatial factors associated with emissions, this information is key. Third, for $\text{PM}_{2.5}$ specifically, using on-road emission factors allows us to take into account not only exhaust emissions, but also emissions from tyres and brakes, which represent a substantial share of emissions (OECD, 2020). The on-road emission factors for NO_x and $\text{PM}_{2.5}$ come from two sources, which both rely on the EU vehicle emissions calculator Copert (see EMEP/EEA (2018) for more details). The type-approval CO_2 emission factors come from the national environmental agency Ademe for cars, and from the French Ministry of the Ecology for other transport modes.

The NO_x , $\text{PM}_{2.5}$ and CO_2 emission factors by transport mode are summarized in table 3.3.2. The factors shown for car and taxi are those imputed when an individual travels with a car that she does not own or a taxi, for which we do not have vehicle characteristics. We then impute a constant emission value from a representative car (a 2008 diesel car of 7 hp). For taxis, we multiply the emission factor by two to reflect empty journeys, as suggested in (Ministère de la Transition écologique et solidaire, 2018). In reality, there is a large variation in the emission intensity of journey stages made with individual car in the survey (see Figures C.1.1, C.1.2 and C.1.3 showing the different values obtained for the emission intensities per passenger ($e_{P,j,i}r_{j,i}$), by transport mode and pollutant). The heterogeneity in NO_x emission factors for cars is the highest, with few extremely high values corresponding to old light-commercial vehicles. We use these emission factors to calculate $E_{P,i}$ for NO_x , $\text{PM}_{2.5}$ and CO_2 . Given the scope of the EGT survey, these estimated individual emissions only include emissions from trips made within the metropolitan area for a representative weekday.

The emission factors described here are per kilometre emission factors which do not vary with the distance or duration of the trip. In reality, for local pollutants, short car trips tend to have a higher emission intensity than long car trips. This is due to the fact that when the car starts and the engine is cold, cold starts contribute to additional exhaust emissions for a certain distance and duration, irrespective of the trip's total distance (Frank et al., 2000). Failing to take this into account may lead us to underestimate emissions from short trips. This matters when we estimate emission savings from modal shift, because modal shift is more feasible for short car trips. Thus, in section 3.4.4, we estimate journey stage-specific emission factors where we apply a higher emission factor to the first few minutes of the trip to reflect cold starts. Appendix C.1.3 details the methodology used.

Exact factor decomposition analysis Starting from equation (2), we re-write individual emissions in the form of an extended Kaya identity (see Wang et al. (2005);

Transport mode	Unit	NOx (mg/km)	PM _{2.5} (mg/km)	CO ₂ (g/km)
Walking and related modes	per passenger	0	0	0
Cycling	per passenger	0	0	0
Street-car	per passenger	0	7	0
Metro	per passenger	0	7	0
Train	per passenger	0	7	0
Bus	per passenger	242	5	117
Taxi	per passenger	1,178	127	332
Car not owned by the household	per vehicle	589	63	166
Two-wheeler not owned by the household	per vehicle	86	21	65

Table 3.3.2 – Assumed contribution to pollution emissions in different transport modes

Note: NOx and PM_{2.5} emission factors reflect on-road emissions and CO₂ emission factors reflect type-approval values. All the assumptions are explained in Appendix C.1.1

Mahony (2013); Bigo (2019) for other examples), as the product of distance, modal share and emission intensity by mode. Note D_i the total distance travelled by individual i , $S_{m,i}$ the modal share of mode m , and $I_{P,m,i}$ the average emission intensity of mode m used by individual i for pollutant P (using the notations from equation 3.2, $I_{P,m,i} = e_{P,m,i}r_{m,i}$). If we call $d_{m,i}$ the total distance travelled by individual i with mode m and $E_{P,m,i}$ the total emissions of pollutant P from using mode m , we have:

$$E_{P,i} = \sum_{m \in M} D_i \frac{d_{m,i}}{D_i} \frac{E_{P,m,i}}{d_{m,i}} = \sum_{m \in M} D_i S_{m,i} I_{P,m,i} \quad (3.3)$$

Given this multiplicative structure, we can use the Log Mean Divisia Index (LMDI) developed by Ang (2004, 2005) to decompose differences in individual-level emissions into differences in distance, modal choice, and the emission intensity by mode. We then calculate the contribution of each component in explaining the difference in emissions between an average individual from the middle quintile (the middle 20% of the distribution in emissions), and reference individuals from quintiles 1,2, 4 and 5 of emissions. The LMDI decomposition has been originally developed to explain changes in emissions over time and this is how it has been applied mostly in the literature. Ang et al. (2015) suggest that the LMDI is also an appropriate method to compare emissions between countries at a given point in time, since it combines ease of use with desirable properties of perfect decomposition and symmetry of decomposition. Some applications have used the LMDI for this purpose, using aggregate country-level data (Liu et al., 2017). Although the method has, to our knowledge, not been applied to individual-level data as we do here, our decomposition is mathematically equivalent to the cross-country case.

We proceed as follows: for each pollutant P , we define 5 quintiles of emissions, Q1 to Q5. We generate a reference individual for each quintile, that is, an individual having the average distance D_{Qk} , modal share $S_{m,Qk}$, and emission intensity $I_{m,Qk}$ of her quintile

Qk , $k = 1..5$ ¹⁵. For the reference individual of quintile Qk , the extended Kaya equation reads:

$$E_{P,Qk} = \sum_{m \in M} D_{Qk} S_{m,Qk} I_{P,m,Qk} \quad (3.4)$$

As recommended in Ang et al. (2015), we define a benchmark individual, here the reference individual from quintile 3, to which we compare the reference individuals from each quintile. We then apply the LMDI decomposition. The difference in emissions between Qk , $k = 1, 2, 4, 5$ and $Q3$ can be decomposed into the difference in the distance (D), modal share (S) and intensity (I) components:

$$E_{P,Qk} - E_{P,Q3} = \Delta E_{P,Qk-Q3,tot} = \Delta E_{P,Qk-Q3,D} + \Delta E_{P,Qk-Q3,S} + \Delta E_{P,Qk-Q3,I} \quad (3.5)$$

Following Ang (2005), this can be rewritten:

$$E_{P,Qk} - E_{P,Q3} = \sum_{m \in M} w_m \ln\left(\frac{D_{Qk}}{D_3}\right) + \sum_{m \in M} w_m \ln\left(\frac{S_{m,Qk}}{S_{m,3}}\right) + \sum_{m \in M} w_m \ln\left(\frac{I_{m,Qk}}{I_{m,3}}\right) \quad (3.6)$$

Where w_m is defined as:

$$w_m = \frac{E_{P,Qk} - E_{P,Q3}}{\ln(E_{P,Qk,m}) - \ln(E_{P,Q3,m})} \quad (3.7)$$

And $E_{P,Qk,m}$ are the emissions of pollutant P associated with mode m for quintile Qk .¹⁶

Individual characteristics associated with high emissions The LMDI decomposition is possible because individual emissions are defined as the exact product between total distance travelled, modal shares, and the emission intensity of different modes. These three components are not independent from each other and result from a complex chain of decisions taken at the individual or household level, including the choice of residence, workplace, vehicle bundle, and modal choice. Modelling all these decisions goes beyond the scope of this paper. We instead investigate in three separate regression analyses

¹⁵this average individual has emissions $E_{P,i}$ that differ from the average emissions of her quintile, given the multiplicative form of the decomposition formula: the product of averages is not the average of the product

¹⁶The modal share of bus, two-wheeler and car is 0 for the bottom quintile of NOx emissions. To be able to apply the log formula, we apply the "Small Value" strategy suggested in Ang and Liu (2007), that is, we replace the zero values by $\delta = 10^{-100}$

which individual characteristics are associated with distance, modal choice and emission intensity (focusing on car for modal choice and emission intensity).

To investigate the characteristics associated with distance travelled, we estimate a log-linear model. Defining $\ln(y)$ the natural logarithm of total distance travelled during the day, x the set of covariates, and ϵ an error term, we set:

$$\ln(y) = x\beta + \epsilon \quad (3.8)$$

We then examine the characteristics associated with using a car at least once during the day with a logit model. Defining S_{car} the modal share of car, the model writes:

$$Pr(S_{car} > 0|x) = \Lambda(x\beta_2) = \frac{\exp(x\beta_2)}{1 + \exp(x\beta_2)} \quad (3.9)$$

We finally examine the characteristics associated with the average emission intensity of car trips. We calculate the average emission intensity of car trips for each individual with a positive car modal share. We estimate a simple linear model, and our results should be interpreted conditionally on driving a car on that day. Defining $I_{P,car}$ the average emission intensity of the car trips for pollutant P , and μ an error term, we estimate the following model for the three pollutants NOx, PM_{2.5} and CO₂:

$$I_{P,car} = x\beta_3 + \mu \quad (3.10)$$

We run the models on two samples: the full sample of individuals, and the sample of individuals in employment, for whom we have rich information on employment characteristics. In all regressions, we control for survey day-specific effects: survey day-of-the-week (we do not have information on the exact survey date); whether the individual encountered a problem with taking transport that day (such as a car's breakdown, a public transport's strike, or bad weather conditions); whether the individual was on holidays or on sickness leave that day.

3.4 Results

3.4.1 How unequal are contributions to emissions?


Figure 3.4.1 illustrates the high inequalities in daily emissions at the individual level using Lorenz curves: on a representative weekday, the top 20% of NOx emitters contribute 85% of NOx emissions, the middle 48% contribute 15%, and the bottom 32% have a zero

contribution¹⁷ (figure 3.4.1a). The top 20% of PM_{2.5} emitters contribute 78% of PM_{2.5} emissions, the middle 62% contribute 22%, and the bottom 18% have a zero contribution (figure 3.4.1b). The top 20% of CO₂ emitters contribute 75% of emissions, the middle 48% contribute 25%, while 32% have a zero contribution (figure 3.4.1c).


Top emitters are not exactly the same across pollutants but the correlation is high¹⁸: the top 20% of NOx emitters contribute 70% of CO₂ emissions. Inequalities of contribution to emissions at the trip level (as defined by equation 3.2) are higher than at the individual level, reflecting the high dispersion of trip distances (see Figure C.1.6).

¹⁷Only individuals with at least one trip are in the sample, so those with zero emissions are the ones travelling only with active modes, electric collective transportation or electric car

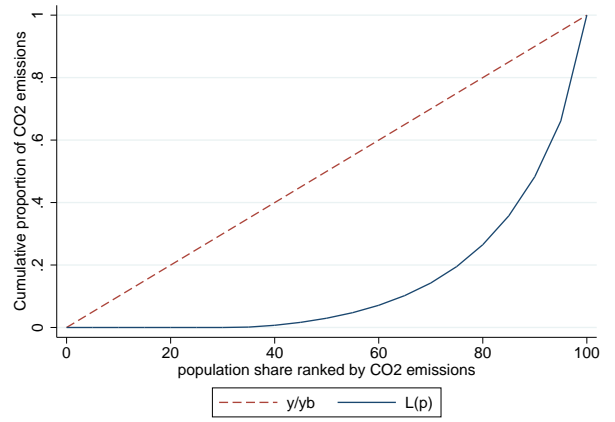
¹⁸the correlation coefficient between individual-level NOx and CO₂ emissions is 0.82



(a) NO_x emissions



(b) PM_{2.5} emissions



(c) CO₂ emissions

Figure 3.4.1 – Lorenz curves for contributions to emissions at the individual level

Note: the x-axis shows the percentiles of individual-level emissions and the y-axis shows the share of total emissions generated by all the individuals below that percentile. The red dotted curve shows how the distribution would look like if everyone contributed equally to emissions Source: EGT data.

Sample: all adults with at least one trip on the day


3.4.2 Are high emissions mostly due to high distances, high-emission modal shares or highly polluting cars?

Figure 3.4.2 show the results of the LMDI decomposition for NO_x, PM_{2.5} and CO₂ emissions. Tables C.1.5, C.1.6, C.1.7 and C.1.8 show the components' values for each quintile and the LMDI Deltas. For NO_x and CO₂ emissions, the lower emissions of the bottom two quintiles are mostly explained by a different modal share, which is expected given the zero emission factor of public transport and active modes, the only modes taken by 32% of the individuals. For PM_{2.5}, subway and train do not have a zero emission factor, such that distance plays a greater role in explaining the low emissions of the bottom two quintiles.


For NO_x and PM_{2.5}, emission intensity, distance and modal share contribute about the same way in explaining the difference between the Q5 and Q3. For example, for NO_x, differences in emission intensity contribute 36%, differences in distance 34%, and differences in modal share 30%. For this pollutant, the values for each component are about 2.5 times greater for Q5 than for Q3, with daily distances travelled of 62km, a car modal share of 92%, and an emission intensity of car trips of 794 mg/km (see Table C.1.5).

For CO₂ emissions, the role of emission intensity is less important than that for local pollutants. Distance and modal share are more important, especially for the top two quintiles. Differences in distances explain 58% of the difference in emissions between Q5 and Q3 for CO₂ (a contribution 24 percentage points higher than for NO_x). Differences in modal share explain 36% (6 percentage points higher than for NO_x). Differences in emission intensity explain only 6% (30 percentage points lower than for NO_x).

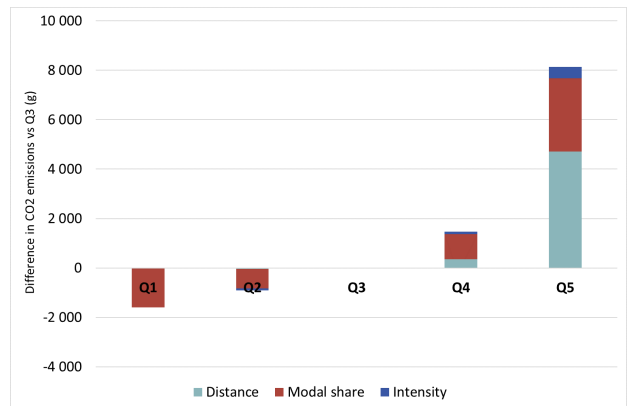
To summarize, the top 20% of NO_x and PM_{2.5} emitters are individuals combining high distances travelled, by car, and with high-emitting cars. In contrast, the top 20% of CO₂ emitters are individuals combining high distances travelled and car trips, with cars only slightly more emission intensive than the average car.



(a) NO_x



(b) PM_{2.5}



(c) CO₂

Figure 3.4.2 – Contribution of distance, modal choice and emission intensity to the differences in emissions, by pollutant

Note: These graphs show, for each pollutant, the difference in emissions between the average individuals from quintiles 1, 2, 4 and 5 and the benchmark individual from quintile 3 (total length of the bars), decomposed into differences in total distance travelled, modal shares, and the emission intensity of a given mode. The LMDI formula used is the additive decomposition (Ang, 2004).

3.4.3 Who emits pollution?

We now turn to the individual and household characteristics associated with higher emissions. Column 1 of table 3.4.1 shows the estimated coefficients for the distance regression (equation 3.8). Columns 2 and 3 show the average marginal effects from the logit estimation on the propensity to use a car (equation 3.9), before and after controlling for having a motorized vehicle available on that day¹⁹.

Spatial factors play an important role for distance and the propensity to use a car: living in central Paris is associated with distances shorter by 24%²⁰ and being 26 percentage point less likely to use a car compared to living in the inner suburbs (w/o controlling for car availability), while living in the outer suburbs is associated with distances longer by 48% and being 18 percentage point more likely to use a car. Living close to a transport stop is associated with distances shorter by 18%, probably partly capturing the fact that public transport stops are located in denser areas. Living close to a public transport stop is also associated with a decrease in the likelihood to use a car, an association that persists after controlling for the availability of a car²¹.

The employment status is a second important characteristic. Being unemployed or inactive is associated with shorter distances, with decreases ranging from -49% (for the unemployed) to -67% (for pensioners) compared to being employed, and a much lower propensity to use a car. Income is a third important characteristic, in part via the positive correlation between income and car ownership: being in the bottom decile is associated with a 22 percentage points lower probability to use a car compared to the six middle deciles when vehicle availability is not controlled for, but only a 6 percentage points lower probability after controlling for it. Symmetrically, being in the top two income deciles is not significantly associated with a higher probability to use a car once vehicle availability is accounted for. Even after including a rich set of socio-economic, spatial and demographic factors as well as controls relative to the survey day, the R-squared for the distance regression is quite low at 0.18, suggesting an important role for other, potentially unobserved factors driving mobility.

¹⁹The omitted categories for the categorical variables present in the model are: for the place of residence, we omit living in the inner suburbs; for gender, we omit male; for income deciles, we take as reference the middle 40% and report coefficients for the two bottom deciles D1 and D2 and the two top deciles D9 and D10. For the activity status, we omit employed individuals; for education, we omit the "primary or secondary education" category; for the type of car owned by the household, we omit the "No car owned" category.

²⁰for small values of estimated coefficients $\hat{\beta}$, a 1-unit change in X corresponds approximately to an expected increase in Y of $\hat{\beta}\%$, but for larger values, the exact interpretation is that a 1-unit change in X corresponds to an expected value of Y multiplied by $e^{\hat{\beta}}$. Most of the obtained coefficients are relatively high in magnitude, so we use the exponential formula to interpret the results.

²¹the vehicle availability variable is defined at the individual level and concerns the reference day, it is different from the variables of car ownership defined at the household level

	(1)	(2)	(3)
	ln dist	uses car	uses car
Inner Paris	-0.277*** (0.0273)	-0.262*** (0.0116)	-0.154*** (0.0116)
Outer suburbs	0.389*** (0.0209)	0.180*** (0.00764)	0.112*** (0.00682)
Public transport stop	-0.204*** (0.0218)	-0.166*** (0.00761)	-0.121*** (0.00670)
Motorized vehicle at hand			0.475*** (0.00736)
Female	-0.297*** (0.0159)	-0.0487*** (0.00531)	0.00571 (0.00479)
Household size	0.00405 (0.00680)	0.0167*** (0.00256)	0.00464* (0.00221)
D1	-0.222*** (0.0399)	-0.217*** (0.0146)	-0.0640*** (0.0145)
D2	-0.179*** (0.0379)	-0.124*** (0.0128)	-0.0275* (0.0111)
D9	0.181*** (0.0299)	0.0587*** (0.0102)	0.0132 (0.00952)
D10	0.196*** (0.0294)	0.0641*** (0.0103)	0.00180 (0.00950)
Pupil/Student	0.250*** (0.0329)	-0.192*** (0.0138)	-0.0445*** (0.0134)
Unemployed	-0.666*** (0.0478)	-0.0854*** (0.0141)	-0.0282* (0.0126)
Other inactive	-0.930*** (0.0260)	-0.0413*** (0.00806)	-0.0415*** (0.00733)
Pensioner	-1.092*** (0.109)	-0.182*** (0.0301)	-0.116*** (0.0316)
Higher education <3 years	0.271*** (0.0261)	0.0569*** (0.00884)	0.0121 (0.00802)
Higher education ≥3 years	0.217*** (0.0225)	0.0161* (0.00757)	-0.0216** (0.00685)
Constant	2.842*** (0.0372)		
N	23596	23600	23524
R-squared	0.1810		

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Standard errors clustered at the household level in parentheses. Columns (2) and (3) report the average marginal effects for each coefficient. All specifications also include survey-day fixed effects and indicator variables for problems with taking transport, being on leave or on sickness leave on the survey day. D1,...,D10: indicator for household income deciles

Table 3.4.1 – Estimated coefficients for distance and propensity to use a car - all individuals

Table 3.4.2 reports the estimated coefficients for the emission intensity regression (equation 3.10) for NO_x, PM_{2.5} and CO₂, before (columns 1, 3 and 5) and after (columns 2, 4 and 6) controlling for the type of vehicle owned by household. Some characteristics are associated with a higher emission intensity for all pollutants, such as living in Paris or owning a light-commercial vehicle. The other way around, being unemployed or inactive and, all else equal, having a higher education diploma, are associated with a lower emission intensity for all pollutants, all else equal.

Other characteristics have an ambiguous role, and are associated with an increase in the emission intensity for some pollutants and a decrease or no effect for others. In line with the well-documented differences in local pollution and CO₂ emission factors for diesel vs. gasoline cars, owning a diesel car is associated with a higher emission intensity for NO_x and PM_{2.5}, but a lower emission intensity for CO₂, compared to owning a gasoline car. Being in the top income decile is strongly associated with a higher CO₂ emission intensity, even after controlling for the type of vehicle owned by the household. This positive correlation between top income and CO₂ emission intensity can be explained by the fact that rich households generally own heavier, larger and more powerful cars, attributes that correlate positively with the CO₂ emission factor. On the other hand, being in the bottom two deciles is associated with a significantly higher PM_{2.5} and CO₂ intensity, and a higher NO_x intensity (but the coefficient is not significant). This may be due to the fact that the cars owned by poorer households are older and more often light-commercial vehicles, two attributes that correlate positively with the emission intensity across all pollutants, and more often powered with diesel, which is positively correlated with PM_{2.5} and NO_x intensity. Like for the distance regression, the explanatory power of the socio-economic, spatial and demographic factors included in the regression is low, at 0.15-0.16 when the type of car owned by the household is accounted for.

Table C.1.11 show the results of fitting similar models on the subsample of individuals in employment, after adding controls for the distance to work, type of commute, type of workplace and type of job²². As expected, the type of commute influences the distances travelled and propensity to use a car: an increase by 1% of the as-the-crow-flies commuting distance is associated with total distances travelled higher by 0.5%, controlling for the type of commute flow (defined by the combination of residence location (Paris/inner suburbs/outer suburbs) and workplace location (Paris/inner suburbs/outer suburbs)). Commuting *type* matters more than commuting *distance* for the propensity to use a car: having to commute from suburbs to suburbs (reference category) is associated with

²²The omitted reference categories are: for the place of residence combined with the place of work: individuals living in the suburbs and working in the suburbs (inner or outer); for the employment status: working full-time; for the workplace type: working in any other place than a factory; for socio-professional category: intermediate professions.

	(1)	(2)	(3)	(4)	(5)	(6)
	NOx/km, all	NOx/km, all	PM25/km, all	PM25 ₂ /km, all	CO ₂ /km, all	CO ₂ /km, all
Inner Paris	75.30*	68.62*	8.785***	8.096***	19.10***	19.18***
	(32.76)	(29.28)	(1.533)	(1.490)	(2.764)	(2.766)
Outer suburbs	-0.249	23.11	1.404*	2.876***	-3.564***	-3.365***
	(14.61)	(12.69)	(0.666)	(0.607)	(0.802)	(0.769)
Public transport stop	-6.142	-19.78	-0.720	-2.282***	1.516*	1.856**
	(13.66)	(12.64)	(0.632)	(0.590)	(0.732)	(0.706)
Female	-101.3***	-77.10***	-4.919***	-3.532***	-0.955	-0.724
	(11.35)	(10.64)	(0.510)	(0.480)	(0.718)	(0.708)
Household size	8.844	14.15*	0.699**	1.603***	-1.367***	-1.686***
	(6.351)	(5.760)	(0.263)	(0.242)	(0.296)	(0.287)
D1	69.77	12.12	5.815**	2.293	4.915*	4.247*
	(43.91)	(41.96)	(1.882)	(1.725)	(1.933)	(1.934)
D2	22.72	-12.96	7.292***	3.735**	3.073*	3.664*
	(26.15)	(26.35)	(1.542)	(1.441)	(1.541)	(1.529)
D9	-8.191	-7.537	-0.790	0.0830	3.058*	2.423*
	(21.14)	(19.83)	(0.894)	(0.841)	(1.241)	(1.198)
D10	-13.76	8.543	-1.475	0.990	7.681***	7.046***
	(22.07)	(20.22)	(0.871)	(0.822)	(1.281)	(1.254)
Pupil/Student	-113.7***	-79.07**	-5.967***	-2.465	2.654	1.976
	(22.55)	(24.96)	(1.434)	(1.401)	(1.884)	(1.915)
Unemployed	-80.51***	-63.81***	-4.410**	-3.763**	-6.428***	-6.039***
	(19.48)	(18.53)	(1.520)	(1.381)	(1.825)	(1.804)
Other inactive	-150.3***	-121.5***	-9.218***	-7.594***	-9.488***	-9.099***
	(13.82)	(12.66)	(0.726)	(0.671)	(0.985)	(0.965)
Pensioner	-55.32	-37.62	-0.838	-0.0630	-5.845	-5.437
	(35.30)	(43.42)	(3.978)	(3.921)	(5.265)	(5.094)
Higher education <3 years	-95.10***	-77.79***	-4.916***	-3.882***	-7.110***	-6.939***
	(18.03)	(16.23)	(0.801)	(0.744)	(0.999)	(0.965)
Higher education ≥ 3 years	-143.4***	-113.2***	-6.742***	-5.561***	-6.935***	-6.117***
	(17.09)	(15.02)	(0.708)	(0.653)	(0.896)	(0.863)
HH owns Diesel Car		129.0***		22.25***		-8.854***
		(7.261)		(0.623)		(0.857)
HH owns Gasoline LCV		1108.0***		29.02***		34.26***
		(129.6)		(2.867)		(3.134)
HH owns Diesel LCV		2171.3***		67.93***		68.58***
		(327.8)		(5.108)		(6.966)
Constant	696.4***	560.5***	58.96***	47.20***	160.5***	161.6***
	(29.49)	(25.73)	(1.284)	(1.196)	(1.504)	(1.498)
N	13097	13094	13097	13094	13097	13094
R-squared	0.0235	0.1514	0.0415	0.1642	0.0330	0.0803
Pseudo R-squared						

Standard errors clustered at the household level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: All specifications also include survey-day fixed effects and indicator variables for problems with taking transport, being on leave or on sickness leave on the survey day. D1,...,D10: indicators for household income deciles

Table 3.4.2 – Regression coefficients for the emission intensity of trips made by car - all individuals

an increase in the likelihood to use a car by 24 to 35 percentage points compared to commuting from Paris to the suburbs or Paris to Paris, probably reflecting the low density of the (radial) Parisian public transport network in the suburbs. The type of job does not affect distances travelled much once other spatial and socio-economic characteristics are taken into account. The R-squared of the distance regression is much higher than for the analysis of the whole sample, suggesting a high explanatory power of job location and employment characteristics.

While the type of occupation does not affect distances travelled much, it is strongly associated with the propensity to use a car: working in a factory is associated with an increase in the likelihood to use a car by 9.6 percentage points, as is having atypical working hours²³. Having a self-employed white-collar profession or being a trades worker are associated with an increase in the likelihood by 12-14 percentage points. Having a low-skilled profession such as personal domestic services, office clerk in the public sector or unqualified manual worker is associated with a lower propensity to use a car, an association seemingly mediated by the lack of car availability. Finally, being a qualified manual worker, craft worker or trades worker is associated with a higher emission intensity for all pollutants, which may be due to the more widespread use of light-commercial vehicles for these professions.

3.4.4 What are the options to reduce emissions?

We investigate options to reduce emissions from car trips specifically, which are responsible for more than 90% of travel-related emissions in our data²⁴. Options to reduce emissions may depend on the trip purpose. Figure 3.4.3 shows the distribution of trip purposes by number of car trips, distances travelled and emissions.²⁵ Work-related trips (com-

²³Atypical working hours are defined as going to work or coming back from work before 5am, or going to work after 4pm.

²⁴96% of the NOx emissions, 90% of the PM_{2.5} emissions, and 91% of the CO₂ emissions. In contrast, trips by metro or train are responsible for 0% of NOx, 7% of PM_{2.5} and 0% of CO₂ emissions, trips by bus are responsible for 4% of NOx, 1% of PM_{2.5} and 7% of CO₂ emissions, and trips by two-wheelers are responsible for 1% of NOx, 2% of PM_{2.5} and 2% of CO₂ emissions

²⁵We use information from the survey on the origin and destination motive (home/ workplace/study place/shopping...) to classify trips in 6 purposes: Commuting trips are those starting or finishing at the work or study place and finishing or starting at another place, except a work-related place. Other work trips are trips where the origin or destination motive is "Work other" (typically, this would be the location of a client meeting or a restaurant where the employee is having a lunch break), and the other motive is home, the workplace or the study place, as well as trips between a workplace and study place. Shopping trips are trips where the destination motive is shopping, or the origin motive is shopping and the destination is home or the work-related. Leisure trips are trips where the destination motive is leisure, or the origin motive is leisure and the destination is home or work-related. Escort trips are trips where the destination motive is escorting, or the origin motive is escorting and the destination is home or work-

muting or business trips) contribute to around 55-60% of emissions, and other purposes (shopping/leisure/escort) to 40-45%.

To assess the emission savings allowed by the different options, we shift to a measure of individual emissions taking into account the fact that short car trips tend to have a higher emission intensity than long car trips. This is due to the fact that when the car starts and the engine is cold, cold starts contribute to additional exhaust emissions for a certain distance and duration, irrespective of the trip's total distance. These cold exhaust emissions have a higher emission factor than the "hot" emissions emitted under normal driving conditions. Cold start emissions thus represent a higher share of the total emissions for short trips compared to long trips. To take this into account, we replace the per kilometre average emission factors from Airparif used in the previous sections with journey stage-specific emission factors for all car and LDV trips where the individual uses a vehicle owned by the household (this represents 90% of the car trips). We follow the methodology exposed in the EMEP/EEA guidance (Ntziachristos and Zissis, 2020), with some simplifications (Airparif used the same method to calculate the average emission factors that we use in the previous sections). The detailed calculations are explained in appendix C.1.3. In short, we impute a relatively higher emission factor or "cold" emission factor during the first 8 minutes of a trip, when the engine is assumed to be cold, and a relatively lower emission factor, or "hot" emission factor, to the remaining kilometres. For PM_{2.5} only, we also add a non-exhaust emission factor related to emissions from tyre and brake wear and a non-exhaust emission factor related to road surface wear, for the whole distance. In the main estimation of emissions, all these components were averaged into a per kilometre emission factor by Airparif. This method reflects more accurately that short trips contribute disproportionately to emissions. However, the results we obtain are very close to the ones we would have obtained with the average emission factors used for

related. We do not have information on the person being escorted, but typically this includes escorting children to school or after-school activities. A number of trips belong to chains: for example, the first trip starts at home and finishes at the children's school, and the second trip starts at the children's school and finishes at work. Given our classification, the first trip will be recorded as an escort trip and the second one as commuting. "Other trips" are all trips not covered by the previous purposes.

the rest of the analysis.²⁶




Figure 3.4.3 – Share of trip purposes in the number of trips, distances travelled and emissions

Note: the first bar chart shows the proportion of trip purpose in the number of trips, the second shows the proportion as a share of total distances driven, the third as a share of NOx emissions and the fourth as a share of CO₂ emissions. Source: EGT data. Sample: all trips made by car or taxi by individuals aged above 18

We consider different options to reduce emissions. According to the “Avoid-Shift-Improve” framework (Creutzig et al., 2018), policies to limit greenhouse gas emissions in the transport sector can be classified into measures aiming at 1) avoiding the need to travel, which in terms of the extended Kaya equation will tackle the distance component; 2) shifting travel to the lowest carbon mode, which will tackle the modal share component; and 3) improving vehicles to be more energy-efficient and fuels less carbon intensive, which will tackle the emission intensity component. The framework is also suited to examine options to abate emissions of local pollutants. We investigate in depth the second option of modal shift, and estimate the share of car trips that could be shifted to low-emission modes. Doing so, we abstract from general equilibrium effects such as the impact of modal shift on road congestion and the demand for driving, the impact of a reduction in commuting on housing prices, which could possibly generate a rebound effect. We also

²⁶The Lorenz curves obtained with these alternative emission factors are also extremely close to the ones obtained with per kilometre emission factors.


investigate the extent to which teleworking could reduce the need to travel (option 1), assuming that place of residence and workplace do not change. Finally, we examine the potential for a shift to less emission-intensive cars by estimating the share of residents who could access a public charging station for EVs or install one at home.

Shift to low-emission modes: We examine the share of car trips²⁷ that could easily be substituted with regular bicycle, electric bicycle, or public transit. Modal choice depends on several cost and preference parameters, and a model of modal choice goes beyond the scope of this paper. We focus here on two dimensions to examine feasibility of a modal shift: the travel time expressed in minutes, and the trip purpose. Based on these two dimensions, we formulate three scenarios of modal shift potential, with an increasing number of constraints. We compare the travel times with different modes for an existing trip using the counterfactual travel times from Google API. For each scenario, we calculate the proportion of possible modal shifts and the associated NO_x, PM_{2.5} and CO₂ emission savings.


Some constraints are common to the three scenarios. First, we impose that switching away from car is only possible if the travel time with the alternative mode is not longer by more than 10 min. Figure 3.4.4 shows the cumulative distribution function of the time difference between driving and cycling, driving and electric cycling, and driving and public transit for all the car trips in the sample. 62% of the trips would be at most 10 min longer by regular bike than by car (graph a, blue line), 73% by electric bike (graph a, red line), but only around 30% by public transit (graph b). Second, we impose an age constraint for cycling: we restrict modal shift to regular bikes to individuals below 60 and modal shift to electric bikes (requiring less effort) to individuals below 70²⁸. Third, we impose that modal shift is not possible when the purpose of the trip is likely to entail carrying heavy materials, which includes work-related driving round (for professions such

²⁷defined as trips using car as their main mode, although some of them may contain journey stages with other modes

²⁸These thresholds seem realistic given that 90% of the few cyclists observed in the EGT data are below 65 (97% of them use a non-electric bike).



(a) Cycling



(b) Public transit

Figure 3.4.4 – Cumulative Distribution Function of the difference in travel time between car and cycling

Note: Sample: all trips currently done by car. Source: Authors' calculations based on Google API outputs. For example, the intersection of the blue line and the left dashed blue line indicates that 62% of the trips currently made by car would last at most 10 minutes more if they were done with a regular bicycle (blue line)

as plumbers or electricians) or escorting someone to a transport stop.

Other constraints are specific to each scenario: In scenario 1, we only use the time, age and type of trip constraints. In scenario 2, which is our preferred scenario, we impose two additional constraints: first, shifting to cycling or e-cycling is only allowed if the resulting daily distance cycled is lower than 20km for regular bikes, and 40km for electric bikes; second, we impose that the additional time spent in transport during the day should not exceed 20 min. In scenario 3, we start with scenario 2 and impose an additional constraint for the type of trip: modal shift is not allowed for shopping trips to a large retail store or mall, which are likely to be associated with heavy loads to carry.

Table 3.4.3 reports, for each scenario, the share of trips that could be shifted to each mode, the share that could be shifted to at least one mode, and the associated NO_x, PM_{2.5} and CO₂ savings (expressed as a share of the total travel-related emissions in the sample). Adding the distance and time constraints between scenarios 1 and 2 decreases the share of car trips that could be cycled from 47% to 24%. The electric bike constraint is less binding, and the share of trips that could be e-cycled remains quite high in scenario

	Scenario 1	Scenario 2	Scenario 3
Switching to cycling possible	0.47	0.24	0.22
Switching to electric bike possible	0.66	0.47	0.43
Switching to public transport possible	0.21	0.21	0.19
At least one modal switch possible	0.68	0.53	0.49
NOx saved as a % of total	0.33	0.21	0.20
PM _{2.5} saved as a % of total	0.32	0.19	0.18
CO ₂ saved as a % of total	0.31	0.19	0.18
N	45,245	45,245	45,245

Table 3.4.3 – Possibility of Modal shift at the car trip level

Note: Source: EGT data. Observations without individual-level sampling weights.

2, at 47%. Finally, 21% of trips could be done with public transit. Overall, 53% of all car trips, representing 19% of NOx and CO₂ emissions and 18% of PM_{2.5} emissions, have at least one substitute under scenario 2. The share of emissions saved increases to 21% for NOx and CO₂ and 20% for PM_{2.5} emissions if we add the individual survey weights. This share is relatively low compared to the share of trips having a substitute because substitutable trips are shorter on average.²⁹

We estimate the monetary benefits associated with scenario 2, in terms of improved air quality and climate change abatement. For the unit cost of NOx and PM_{2.5} emissions, we use monetary values from the European Commission report on the external costs of transport EU Commission (2020)³⁰. We adjust the values for France given for the year

²⁹The results would be very close if we took the average emission factors rather than the journey stage-specific emission factors accounting for differences in the share of cold starts between short and long trips: running this analysis with the average emission factor, we would have obtained emission savings of 17% for NOx and PM_{2.5}, instead of 19% for NOx and 18% for PM_{2.5}

³⁰See Annex A and Annex C from EU Commission (2020), and pp59-67 of CE Delft (2018) for more details on the economic valuation of health and the assessment of air pollution costs. In short, the monetary values include the costs of air pollution in terms of individual health, crop losses, material and building damages, and biodiversity losses. The different cost factors are estimated in three steps, based on the methodology developed in the 2007 NEEDS project (NEEDS, 2007): first, emissions are translated into concentrations; second, concentrations are translated in health and environmental impacts using dose-response functions; third, health and environmental impacts are given a monetary value. Sources for the cost values include (NEEDS, 2007) and updates from more recent sources. For the health costs (which represent the largest share of costs), mortality and morbidity dose-response functions are based

2016 for inflation and we obtain a unit cost of €28.03 per kilogram of NO_x and €134.98 per kilogram of PM_{2.5} in 2020. The value for PM_{2.5} is conservative: we take the estimate for urban areas, which is three times lower than that for metropolitan areas, while part of the emission savings from Scenario 2 would occur in the Paris metropolitan area. For the unit cost of CO₂ emissions, we use the official value for the social cost of carbon in France in 2020 France Stratégie (2019), given in euro 2018, and adjust it for inflation to obtain a euro estimate in real terms. We obtain a unit cost of €84.5 per ton of CO₂.

We first calculate the external environmental costs of passenger transport in IdF, absent any modal shift. For this, we apply the individual EGT survey weights to estimate the total emissions generated by individuals given current modal choices. We then combine this total emission value with unit values of NO_x, PM_{2.5} and CO₂. We find that the daily mobility of residents generates an environmental cost of around €4.7m per day, of which €3.0m for local pollution and €1.7m for CO₂ emissions. Assuming that the survey is representative of the 220 annual working days³¹, the annual environmental cost of daily mobility in IdF is at least €1,000m.

We then estimate the monetary benefits that would be realised under scenario 2, after estimating the absolute quantity of emission savings with the survey weights. We obtain daily benefits of €1.1m, of which €0.61m from avoided local pollution and €0.36m from avoided CO₂ emissions. With 220 working days, the corresponding annual benefit is €214m. One caveat is that using emission factors from 2010 may overestimate the quantity of NO_x and PM_{2.5} emissions saved compared to what would be obtained in 2020: the vehicle fleet from 2010 was on average more polluting than the vehicle fleet from 2020 due to a rising stringency of European pollution standards. But having a conservative

on a WHO study (WHO, 2013). Mortality impacts are monetized using an estimate of VOLY (Value of a Life Year) of €70,000 per life year for the EU28, derived from a literature review. The EU-level VOLY value is translated into country-specific values using unit value transfers adjusting for income differences across countries. Morbidity impacts are estimated using a conversion table expressing illness and disability as partial mortality in a QALY (quality-adjusted life year) framework, assuming that 1 QALY=1/1.087 VOLY.

³¹The travel intensity reported in the survey is representative of an average weekday between October and May, where some individuals are on holiday but likely not a large share. There are probably fewer trips in IdF in July and August, two months where most people take several weeks of holidays in France

estimate for the unit cost of PM_{2.5} probably mitigates the risk of overestimation. We also note that by focusing on the benefits of modal shift in terms of air pollution reduction and CO₂ mitigation, we do not include the potential costs associated with modal shift (e.g, time lost, nor other types of benefits such as the health benefits from active mobility³².

At the individual level, under scenario 2 30% of the drivers could not shift any of their car trips, and 28% could shift modes for part of their car trips only. These drivers are more likely to live in the outer suburbs (72%, versus 63% in the entire sample of drivers), and drive longer trips on average (47km per day, vs. 32 km per day in the entire sample of drivers). For these individuals, other solutions are needed. Below, we investigate the potential for teleworking (the distance or "Avoid" component) and the potential for shifting to an electric vehicle (the emission intensity or "Improve" component).

Avoid travelling by teleworking: Teleworking could be all the more relevant since 42% of the employed individuals unable to shift modes use car for commuting. Work-related car trips (either commuting or business trips) also have a lower-than-average potential for modal shift: in scenario 2, only 50% of commuting trips and 38% of business trips have a modal shift option (vs. 53% on average).

The potential for teleworking has recently gained prominence in the public debate and in the literature in the context of the Covid-19 pandemic and associated social distancing measures (Dingel and Neiman, 2020; Alipour et al., 2020; Lennox, 2020). We combine information on the socio-professional category and the workplace to define a variable of potential to telework³³. We consider that teleworking is not possible for manual workers, farmers or traders, craftspeople, CEOs. For the other socio-professional categories, we consider that teleworking is possible for employees from the private and public sector as long as they work in an office³⁴. According to these criteria, 39% of all the car commuters

³²The health benefits of walking and cycling induced by the increase in physical activity have been shown to significantly outweigh the risks due to pollution inhalation and cyclists' accidents (Rojas-Rueda et al., 2011; Rabl and de Nazelle, 2012)

³³We cannot use the exact same definition of potential to telework as in the recent paper by Dingel and Neiman (2020) due to data limitations.

³⁴as opposed to working at a factory, at other people's homes, at a hospital or school, at a public

have a job type which could be done from home. If all the car commuters who cannot shift modes entirely worked from home, an additional 15% of NO_x, 14% of PM_{2.5} and 15% of CO₂ emissions would be saved from the avoided commuting and business trips. Assuming that these individuals could telework two days a week (two fifths of their time), this would save an additional 5% of PM_{2.5} and 6% of NO_x and CO₂ emissions compared to the emission savings achieved via modal shift only.

Improve: shift to an Electric Vehicle: Another alternative to modal shift is a shift to an electric vehicle (EV). There are well-documented monetary and non-monetary barriers to the uptake of EVs: cost of purchase, availability of a charging station, cultural habits (Oxford Institute for Energy Studies, 2019). To identify who may be likely to shift to an EV under some assumptions, we would need a model of car purchasing decisions which is beyond the scope of this paper. We simply note two points suggesting that the barriers associated with the purchasing cost and charging point availability may be overcome with adequate policies. First, among the 58% of current drivers unable to shift modes for all their car trips, only 9% are from the two bottom deciles of income, such that their budget constraint is less binding than for the whole population. Second, at least 17%³⁵ of them have a publicly available EV charging station at less than 500m from their place of residence in 2020, and 77% of them have a private parking space at their place of residence, where a charging station could be installed. Finally, less than 1% of them drive more than 200km per day (with the limitation that trips outside the IdF region are not recorded), such that the autonomy of the EV should not be an issue for this daily mobility. Alternative vehicles such as two-seat microcars or delivery tricycles (known in the broad category of "Little Vehicles" (Schneider, 2018)) may also provide travel services currently provided by traditional cars at a lower cost than electric vehicles.

Table 3.4.4 summarizes teleworking and EV shift options at the individual level, for the entire sample of drivers (first column) and for the 58% of drivers who cannot shift all
institution, or at a shop

³⁵This estimate is conservative because the data on EV charging stations appears not to be exhaustive.

their car trips to low-emission modes (second column). Available options are very close across the two samples.

	Share of drivers	Share of drivers who cannot fully shift to low-emission modes
Teleworking possible	0.28	0.29
Teleworking possible and current commute by car	0.13	0.13
Has a private parking spot	0.76	0.77
Has a public EV charging station within 500m	0.18	0.17
N	13,140	7,562

Table 3.4.4 – Teleworking and EV shift options for current drivers

3.5 Discussion

3.5.1 A 80-20 rule?

We find a strong concentration of local pollutant and CO₂ emissions as far as daily mobility is concerned. This result had, to our knowledge, not been reported before for a large city based on representative data of the residents. Brand and Preston (2010) report that the top 20% UK emitters contribute to 60% of CO₂ emissions from transport and mention a “60-20 rule”, but their analysis is based on a small sample of residents from one UK region and includes both daily mobility and long-distance trips. Based on a small non-representative sample from Beijing residents, Yang et al. (2018) report that 20% of the top emitters contribute 70% of emissions, on both weekdays and weekends. Our results for the Paris area, based only on weekday trips within the area, suggest a “80-20” rule on average across the pollutants considered. Only considering weekdays seems relevant to analyse the potential for air pollution mitigation in the Paris area, because ambient pollution tends to be higher on weekdays, where car traffic and economic activity are higher. For CO₂ emissions, examining long-distance trips and weekends seems necessary to get the full picture of carbon footprint inequalities: indeed, residents from the city centre (who in our analysis contribute substantially less to emissions than suburban residents) tend to take the plane more often and emit more during their long-distance trips (Pottier et al., 2020).

3.5.2 Traditional and less traditional factors associated with emissions

Four factors associated with daily distances travelled and modal choice have previously been highlighted in the literature: employment status, household income, household residence location vis a vis the city centre, and agglomeration size (Nicolas and David, 2009; Blaudin De Thé et al., 2020; Pottier et al., 2020). Our results are consistent with this

literature. We also highlight the role of the commute type in explaining the propensity to use a car, with suburbs-to-suburbs commutes being more reliant on cars.

A newer aspect of our work is to document the association between employment characteristics and the propensity to use a car, with individuals working at factories and manual, trades and crafts workers having a higher propensity to use a car. The high reliance on car of these professional categories could play a role in the political economy of opposition to policies regulating car use.

Another contribution is to highlight the different relationship between income and local pollutants' emission intensity on the one hand, and income and CO₂ emission intensity on the other hand. Two factors can drive up the emission intensity of individuals from the bottom decile: first, the NOx emission intensity of light-commercial-vehicles is much higher than the emission intensity of regular cars of the same age, and individuals with manual occupations from the middle and bottom deciles are more likely to have such cars; second, lower-income individuals have older cars on average, and the NOx and PM_{2.5} emission factors are determined by the age and fuel type of the vehicle. In contrast, in our data the CO₂ emission factor depends on the age, fiscal horsepower and energy of the car. Higher-income people tend to have a higher fiscal horsepower and newer cars, and CO₂ emissions increase with fiscal horsepower but they do not vary much with the age of the car. This distribution of vehicle characteristics across income groups suggest that policies based on the NOx emission intensity of vehicles, such as Low-emission zones, whose exclusion criteria depend on the age of vehicles, could be more regressive than policies regulating the CO₂ emission intensity of vehicles, such as feebates.

3.5.3 From modal shift potential to actual modal shift

The LMDI decomposition suggests that the emission intensity of vehicles is only one driver of emissions, and a minor one for CO₂ emissions. Policies tackling modal shift and demand for distance are also needed. Regarding modal shift, we document a relatively large potential based on travel time criteria. Adequate policies are required to fulfil this

potential: despite the potential of modal shift to cycling and electric cycling, its modal share is only 1.9% of total trips in 2018 in the Paris area (Omnil-Ile de France Mobilites, 2019).

Cost factors, in the form of monetary costs or travel times, have been extensively studied to explain modal choice and modal shift. Regarding modal choices, an important driver of travel times is the built environment, which also influences the quality of the travel via its impact on walkability or accessibility (Javaid et al., 2020). Regarding modal shift, it is influenced by the relative cost of different modes as captured by inter-modal cross elasticity measures, that is, how much demand for a given mode decreases when the monetary or time cost of alternative modes decreases. In a recent meta-analysis of the literature, Wardman et al. (2018) reports relatively low values of inter-modal cross elasticities for car demand: The highest elasticity is when the alternative mode is subway, and is 0.23. However, the meta-analysis does not include studies examining cycling as alternative mode. The other way around, there is a wide variation in estimates of the inter-modal cross elasticity for the demand of other modes when car use is made costlier, with estimates ranging from 0.01 to 0.86 depending on the alternative mode examined and the cost parameter (fuel cost, in-vehicle-time, etc).³⁶ To our knowledge, there is no estimate of inter-modal cross elasticities for the IdF region. Estimates for the simple price elasticity of car demand are around -0.25 (Direction générale du Trésor, 2011). Finally, cross-elasticities based on time attributes tend to be larger than those based on price attributes such as fuel cost. It means that enabling a 10% decrease in travel time may be more effective than a 10% decrease in subway fares to facilitate modal shift.

However, these individual cost factors are not the only ones influencing modal shift. Mattauch et al. (2016) highlight the role of cognitive factors such as statu-quo bias, overconfidence or framing effects to explain modal choices, implying that individuals do not necessarily analyse travel decisions on a trip-by-trip basis but rather rely on past decisions. The consequence is that modal choices are sticky.

³⁶This elasticity is higher when the initial market share of car is high and that of alternative modes is low.

The forced experiment of the Covid-19 crisis could be an opportunity for a permanent shift in habits against the statu-quo, as observed in the case of other disruptions in usual travel habits such as public transport strikes (Larcom et al., 2017). Given the behavioral factors influencing modal choice, rolling out cycling infrastructure in a disrupted time could also have a multiplier effect. Recent evidence suggests that pop-up bike lanes rolled out to facilitate social distancing during Covid-19 have increased cycling between 11 and 48% in the following months, depending on the city considered (Kraus and Koch, 2021). One key question for future research is whether these relatively large effects will persist over time.

Finally, for active modes, weather conditions may also play a role, with warm and dry weather conditions having a positive influence and rain, snow, wind, overly cold or hot weather having the opposite effect (Böcker et al., 2013). For electric bikes specifically, which we find enable a large part of the modal shifts, their relatively high cost and the risk of bike theft are other important factors hindering a wider adoption in the Paris area (Cazi, 2020), although sales have been increasing significantly since 2017 (Le Monde, 2021) (figures are at the national level). Electric bike-sharing options may be a good way to promote a higher take-up while addressing the monetary costs of electric bikes and the risk of theft.

For drivers without a modal shift option, reducing distance and emission intensity is needed. Only 13% of drivers combine commuting by car and being able to work from home. But the emission savings associated with them teleworking are relatively high given their high commuting distance. In the long-term, urban planning could play a role in reducing demand for car trips, for example by improving the diversity and design of the suburbs (see Blaudin De Thé et al. (2020) for a discussion of these dimensions) and making cities more polycentric. Regarding policies tackling the emission intensity of cars, such as subsidies to buy EVs or low-emission cars, they are all the more needed in the outer suburbs, where individuals are less likely to be able to shift modes. However, to date the means-tested subsidies for new car purchases introduced with the Parisian Low

Emission Zone are only available for household living within the planned LEZ boundaries (Paris and part of the inner suburbs), excluding households from the outer suburbs. One recommendation would be to open the subsidies to individuals living outside the Greater Paris area but working in the LEZ. Note that the per kilometer reduction in air pollution and CO₂ emissions allowed by electric vehicles is smaller than that allowed by shifts to active modes or electric public transport, due to higher lifecycle emissions of cars and the non-exhaust particulate emissions of electric cars (OECD, 2020), which are particularly damaging for health (Daellenbach et al., 2020).

3.5.4 Limits

The main limitation of our analysis is that we do not take into account the potential rebound effect of the different options to reduce emissions. In the case of modal shift, we imagine two possible types of rebound: first, rebound from individuals renouncing to have a car, who may spend the savings from not owning a car on carbon-intensive goods and services, as evidence in a study on Finland (Ottelin et al., 2017). A second type of rebound effect could occur via a reduction in congestion which would increase the marginal utility of driving. More research is needed to estimate the magnitude of such an effect, but it could be partially mitigated by a reduction of the space left to cars in the public space, proportional to the reduction of car use. In the case of teleworking, rebound may occur if people used the time freed up by the absence of commute for leisure travels. To our knowledge, the only empirical study estimating the impact of teleworking finds a net reduction in traffic and city-level pollution at the monthly level (Giovanis, 2018), but it does not measure the effects on long-distance trips.

Although we use data from 2010, we think that our results are still relevant to explain today's distribution of emissions in Paris. Preliminary results from the new wave of the EGT survey (planned to be carried out between 2018 and 2022, but currently stalled due to the Covid-19 crisis) suggest that the average number of trips, time and distances spent travelling have not changed since 2010 (Omnil-Ile de France Mobilites, 2019). The average

modal share changed only slightly, with a small decrease in car use (from 37.8% of the trips in 2010 to 34.4% in 2018), compensated by an increase in active transportation modes and collective transportation. Using data from 2010 may be more problematic to estimate the emission savings associated with our scenarios in absolute terms and the associated monetized benefits. But using conservative estimates for the unit cost of emissions likely counterbalances this risk.

3.6 Conclusion

We show that inequalities in contribution to transport-related emissions are large in the Paris area, with top emitters combining large distances travelled and a reliance on high-emitting cars. We document an important association between some employment characteristics and emissions. In a monocentric city like Paris, distance from the center and other spatial characteristics are also strongly associated with higher emissions. Although we report a high potential for modal shift in terms of the share of car trips where a low-emission substitute exists, the associated emission savings is only 20% because many long trips/trip with high-emission cars cannot be substituted. Policies encouraging a decrease in demand for travel and the adoption of low-emission cars are needed for the individuals unable to shift modes.

Regarding the external validity of our results, we expect that city size and density influence both the external cost of transport, as underlined by Carozzi and Roth (2019); Gaigné et al. (2012) and the potential for modal shift, as underlined by Nicolas and David (2009) and Brand et al. (2021). For the relationship between density and the environmental externalities from transport, the urban economics literature points to a potential trade-off between CO₂ emissions and local pollution, because one is a global externality and the other affects local residents only: compact (more dense) cities are associated with shorter distances and more public transport so they may reduce the quantity of polluting emissions (Gaigné et al., 2012). So compact cities may be good for

CO₂ emissions. At the same time, the cost of local air pollutants depends on how emissions translate into ambient concentrations and how many people are exposed to this pollution. Then, a higher density may lead to higher population-weighted pollution concentration, as evidenced by Carozzi and Roth (2019) in the US case, and also higher benefits from reducing local pollutant emissions. In contrast, the benefits from CO₂ emission reductions would be the same in all cities given the global nature of the climate change externality. Regarding the potential for modal shift, shifting to active modes may be easier in smaller cities with shorter distances travelled, but shifting to public transport may be harder, as the public transport network is usually less good in small cities.

We think that our results likely apply to other dense European cities with an important public transit network, such as London, Madrid or Rome, as well as other large French urban areas. In any case, it should be easy to replicate our analysis in other cities of the developed world, given the availability of transport survey data such as the one used in this paper in other cities (for example, the London Travel Demand Survey).

Bibliography

- Abadie, A. (2021). Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects. *Journal of Economic Literature*.
- Abadie, A., S. Athey, G. W. Imbens, and J. M. Wooldridge (2020). Sampling-Based versus Design-Based Uncertainty in Regression Analysis. *Econometrica* 88(1), 265–296. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA12675>.
- Abadie, A., A. Diamond, and J. Hainmueller (2010). Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California’s Tobacco Control Program. *Journal of the American Statistical Association* 105(490), 493–505.
- Abadie, A., A. Diamond, and J. Hainmueller (2015). Comparative Politics and the Synthetic Control Method. *American Journal of Political Science* 59(2), 495–510.
- Abadie, A. and J. Gardeazabal (2003). The Economic Costs of Conflict: A Case Study of the Basque Country. *American Economic Review* 93(1), 113–132.
- Abrell, J., M. Kosch, and S. Rausch (2019). How Effective was the UK Carbon Tax?—A Machine Learning Approach to Policy Evaluation. *Unpublished Working Paper*.
- AirParca (2015). Caimans report: Air quality impact and greenhouse gases assessment for cruise and passenger ships. Technical report.
- Alipour, J.-V., H. Fadinger, and J. Schymik (2020). My Home Is My Castle – The Benefits of Working from Home During a Pandemic Crisis: Evidence from Germany. Technical Report 14871, C.E.P.R. Discussion Papers.
- Ambec, S. and J. Coria (2013). Prices vs quantities with multiple pollutants. *Journal of Environmental Economics and Management* 66(1), 123–140.
- Andersson, J. J. (2019). Carbon Taxes and CO2 Emissions: Sweden as a Case Study. *American Economic Journal: Economic Policy*.
- Ang, B. W. (2004). Decomposition analysis for policymaking in energy:: which is the preferred method? *Energy Policy* 32(9), 1131–1139.
- Ang, B. W. (2005). The LMDI approach to decomposition analysis: a practical guide. *Energy Policy* 33(7), 867–871.
- Ang, B. W. and N. Liu (2007). Handling zero values in the logarithmic mean Divisia index decomposition approach. *Energy Policy* 35(1), 238–246.

- Ang, B. W., X. Y. Xu, and B. Su (2015). Multi-country comparisons of energy performance: The index decomposition analysis approach. *Energy Economics* 47, 68–76.
- Angrist, J. D. and J.-S. Pischke (2010). The Credibility Revolution in Empirical Economics: How Better Research Design Is Taking the Con out of Econometrics. *Journal of Economic Perspectives* 24(2), 3–30.
- Arceo, E., R. Hanna, and P. Oliva (2016). Does the Effect of Pollution on Infant Mortality Differ Between Developing and Developed Countries? Evidence from Mexico City. *The Economic Journal* 126(591), 257–280.
- Ares, E. and J. Delebarre (2016). The Carbon Price Floor. Technical report, House of Commons Library.
- Athey, S. and G. W. Imbens (2017). The State of Applied Econometrics: Causality and Policy Evaluation. *Journal of Economic Perspectives* 31(2), 3–32.
- AtmoSud (2018). Quelle qualite de l’air pour les riverains des ports de Nice et Marseille? Technical report, AtmoSud.
- AtmoSud (2020). Cigale: Consultation d’inventaires geolocalisés air climat energie. <https://cigale.atmosud.org/>.
- Atmosud (2020). Resume : Etat de la qualite de l’air à l’échelle de la ville de Marseille apres mise en service de la L2 (A507). Technical report.
- Auffhammer, M. (2018). Quantifying Economic Damages from Climate Change. *Journal of Economic Perspectives* 32(4), 33–52.
- Baldino, C., U. Titge, R. Muncrief, Y. Bernard, and P. Mock (2017). Road tested: Comparative overview of real-world versus type-approval NOX and CO2 emissions from diesel cars in Europe. *ICCT White paper*.
- Barbusse, S. (2005). Motocycles, cyclomoteurs - Energie et environnement. Technical report, Ademe.
- Baumol, W. J. and W. E. Oates (1988). *The Theory of Environmental Policy* (2 ed.). Cambridge: Cambridge University Press.
- Böcker, L., M. Dijst, and J. Prillwitz (2013). Impact of Everyday Weather on Individual Daily Travel Behaviours in Perspective: A Literature Review. *Transport Reviews* 33(1), 71–91. Publisher: Routledge _eprint: <https://doi.org/10.1080/01441647.2012.747114>.

- BEIS (Department for Business, Energy & Industrial Strategy) (2019). UK Coal statistics.
- Berghmans, N. and E. Alberola (2013). The Power Sector in Phase 2 of the EU ETS: Fewer CO₂ Emissions but just as much Coal. *Climate Report. paris CDC Climate research* (42).
- Berry, A. (2019). The distributional effects of a carbon tax and its impact on fuel poverty: A microsimulation study in the French context. *Energy Policy* 124(C), 81–94. Publisher: Elsevier.
- Böhringer, C., H. Koschel, and U. Moslener (2008). Efficiency losses from overlapping regulation of EU carbon emissions. *Journal of Regulatory Economics* 33(3), 299–317.
- Bigo, A. (2019). How to explain the past trends in transport CO₂ emissions in France?
- Bind, M.-A. C. and D. B. Rubin (2019). Bridging observational studies and randomized experiments by embedding the former in the latter. *Statistical methods in medical research* 28(7), 1958–1978.
- Bind, M.-A. C. and D. B. Rubin (2020). When possible, report a fisher-exact p value and display its underlying null randomization distribution. *Proceedings of the National Academy of Sciences* 117(32), 19151–19158.
- Bishop, K. C., J. D. Ketcham, and N. V. Kuminoff (2018). Hazed and Confused: The Effect of Air Pollution on Dementia. Technical Report 24970, National Bureau of Economic Research, Inc.
- Blaudin De Thé, C., B. Carantino, and M. Lafourcade (2020). The Carbon 'Carprint' of Suburbanization: New Evidence from French Cities.
- Bonilla, J., J. Coria, and T. Sterner (2018). Technical Synergies and Trade-Offs Between Abatement of Global and Local Air Pollution. *Environmental and Resource Economics* 70(1), 191–221.
- Bou Sleiman, L. (2021). Are car-free centers detrimental to the periphery? Evidence from the pedestrianization of the Parisian riverbank. Technical Report 2021-03, Center for Research in Economics and Statistics. Publication Title: Working Papers.
- Bove, M., P. Brotto, F. Cassola, E. Cuccia, D. Massabò, A. Mazzino, A. Piazzalunga, and P. Prati (2014). An integrated PM_{2.5} source apportionment study: positive matrix factorisation vs. the chemical transport model CAMx. *Atmospheric environment* 94, 274–286.

- Braithwaite, I., S. Zhang, J. B. Kirkbride, D. P. J. Osborn, and J. F. Hayes (2019). Air Pollution (Particulate Matter) Exposure and Associations with Depression, Anxiety, Bipolar, Psychosis and Suicide Risk: A Systematic Review and Meta-Analysis. *Environmental Health Perspectives* 127(12), 126002.
- Brand, C., T. Götschi, E. Dons, R. Gerike, E. Anaya-Boig, I. Avila-Palencia, A. de Nazelle, M. Gascon, M. Gaupp-Berghausen, F. Iacorossi, S. Kahlmeier, L. Int Panis, F. Racioppi, D. Rojas-Rueda, A. Standaert, E. Stigell, S. Sulikova, S. Wegener, and M. J. Nieuwenhuijsen (2021). The climate change mitigation impacts of active travel: Evidence from a longitudinal panel study in seven European cities. *Global Environmental Change* 67, 102224.
- Brand, C. and J. M. Preston (2010). ‘60-20 emission’—The unequal distribution of greenhouse gas emissions from personal, non-business travel in the UK. *Transport Policy* 17(1), 9–19.
- Bretschger, L. and K. Pittel (2019). Twenty Key Questions in Environmental and Resource Economics. Technical Report 19/328, CER-ETH - Center of Economic Research (CER-ETH) at ETH Zurich.
- Brown, G. (2017). British power generation achieves first ever coal-free day. *The Guardian*.
- Bureau, B. and M. Glachant (2008). Distributional effects of road pricing: Assessment of nine scenarios for Paris. *Transportation Research Part A: Policy and Practice* 42(7), 994–1007.
- Bureau, B. and M. Glachant (2011). Distributional effects of public transport policies in the Paris Region. *Transport Policy* 18(5), 745–754.
- California Air Resources Board (2021). Overview: Diesel Exhaust & Health.
- Carozzi, F. and S. Roth (2019). Dirty Density: Air Quality and the Density of American Cities. Technical Report dp1635, Centre for Economic Performance, LSE.
- Carslaw, D. C. and K. Ropkins (2012). Openair—an r package for air quality data analysis. *Environmental Modelling & Software* 27, 52–61.
- Cazi, E. (2020). Véligo profite de l’engouement pour le vélo en Ile-de-France.
- CE Delft (2018). Environmental Prices Handbook. Technical report.
- CGDD (2011). Les émissions de CO2 des véhicules utilitaires légers. Technical report, Ministère de l’Environnement, de l’Énergie et de la Mer.

- Chancel, L. and T. Piketty (2015). Carbon and inequality: From Kyoto to Paris. Technical report.
- Chang, T. Y., J. G. Zivin, T. Gross, and M. Neidell (2019). The Effect of Pollution on Worker Productivity: Evidence from Call Center Workers in China. *American Economic Journal: Applied Economics* 11(1), 151–172.
- Chay, K. Y. and M. Greenstone (2003). The Impact of Air Pollution on Infant Mortality: Evidence from Geographic Variation in Pollution Shocks Induced by a Recession*. *The Quarterly Journal of Economics* 118(3), 1121–1167.
- Chiroleu-Assouline, M. and M. Fodha (2014). From regressive pollution taxes to progressive environmental tax reforms. *European Economic Review* 69, 126–142.
- Chrisafis, A. (2018). 'I don't want ships to kill me': Marseille fights cruise liner pollution. *The Guardian*.
- Chyong, C.-K., B. Guo, and D. Newbery (2019). The impact of a Carbon Tax on the CO2 emissions reduction of wind. Technical Report 1904, Faculty of Economics, University of Cambridge.
- Cleveland, W. S. (1993). *Visualizing data*. Hobart Press.
- Cochran, W. G. and D. B. Rubin (1973). Controlling bias in observational studies: A review. *Sankhyā: The Indian Journal of Statistics, Series A*, 417–446.
- Colmer, J., R. Martin, M. Muûls, and U. J. Wagner (2020). Does Pricing Carbon Mitigate Climate Change? Firm-Level Evidence From the European Union Emissions Trading Scheme. Technical Report crctr224_2020_232, University of Bonn and University of Mannheim, Germany. Publication Title: CRC TR 224 Discussion Paper Series.
- Contini, D., A. Gambaro, F. Belosi, S. De Pieri, W. Cairns, A. Donateo, E. Zanotto, and M. Citron (2011). The direct influence of ship traffic on atmospheric PM2.5, PM10 and PAH in Venice. *Journal of Environmental Management* 92(9), 2119–2129.
- Corbett, J. J., J. J. Winebrake, E. H. Green, P. Kasibhatla, V. Eyring, and A. Lauer (2007). Mortality from ship emissions: a global assessment. *Environmental science & technology* 41(24), 8512–8518.
- Cowell, F. (2011). *Measuring Inequality* (Third Edition ed.). London School of Economics Perspectives in Economic Analysis. Oxford, New York: Oxford University Press.

- Creutzig, F., A. Javaid, N. Koch, B. Knopf, G. Mattioli, and O. Edenhofer (2020). Adjust urban and rural road pricing for fair mobility. *Nature Climate Change* 10(7), 591–594.
- Creutzig, F., J. Roy, W. F. Lamb, I. M. L. Azevedo, W. Bruine de Bruin, H. Dalkmann, O. Y. Edelenbosch, F. W. Geels, A. Grubler, C. Hepburn, E. G. Hertwich, R. Khosla, L. Mattauch, J. C. Minx, A. Ramakrishnan, N. D. Rao, J. K. Steinberger, M. Tavoni, D. Ürge Vorsatz, and E. U. Weber (2018). Towards demand-side solutions for mitigating climate change. *Nature Climate Change* 8(4), 260–263.
- Cronin, J. A., D. Fullerton, and S. Sexton (2018). Vertical and Horizontal Redistributions from a Carbon Tax and Rebate. *Journal of the Association of Environmental and Resource Economists* 6(S1), S169–S208.
- Cropper, M. L., R. D. Morgenstern, and N. Rivers (2018). Policy Brief—Facilitating Retrospective Analysis of Environmental Regulations. *Review of Environmental Economics and Policy* 12(2), 359–370. Publisher: Association of Environmental and Resource Economists.
- Currie, J. and R. Walker (2019). What Do Economists Have to Say about the Clean Air Act 50 Years after the Establishment of the Environmental Protection Agency? *Journal of Economic Perspectives* 33(4), 3–26.
- Currie, J., J. G. Zivin, J. Mullins, and M. Neidell (2014). What Do We Know About Short- and Long-Term Effects of Early-Life Exposure to Pollution? *Annual Review of Resource Economics* 6(1), 217–247.
- Daellenbach, K. R., G. Uzu, J. Jiang, L.-E. Cassagnes, Z. Leni, A. Vlachou, G. Stefenelli, F. Canonaco, S. Weber, A. Segers, J. J. P. Kuenen, M. Schaap, O. Favez, A. Albinet, S. Aksoyoglu, J. Dommen, U. Baltensperger, M. Geiser, I. E. Haddad, J.-L. Jaffrezou, and A. S. H. Prévôt (2020). Sources of particulate-matter air pollution and its oxidative potential in Europe. *Nature* 587(7834), 414–419.
- Dasgupta, T. and D. B. Rubin (2015). *STAT 240: Matched Sampling and Study Design*. Harvard university class.
- Dasgupta, T. and D. B. Rubin (2021). *Experimental Design: A Randomization-Based Perspective*. Unpublished Textbook.
- Davis, L. W. and C. R. Knittel (2018, November). Are Fuel Economy Standards Regressive? *Journal of the Association of Environmental and Resource Economists* 6(S1), S37–S63.

- de Nazelle, A., B. J. Morton, M. Jerrett, and D. Crawford-Brown (2010). Short trips: An opportunity for reducing mobile-source emissions? *Transportation Research Part D: Transport and Environment* 15(8), 451–457.
- Delhaes, D. and S. Kersting (2019). Yellow vests: Ban on diesel vehicles drives Germans to the street – to protest.
- Dennig, F., M. B. Budolfson, M. Fleurbaey, A. Siebert, and R. H. Socolow (2015). Inequality, climate impacts on the future poor, and carbon prices. *Proceedings of the National Academy of Sciences* 112(52), 15827–15832.
- Deryugina, T., G. Heutel, N. H. Miller, D. Molitor, and J. Reif (2019). The mortality and medical costs of air pollution: Evidence from changes in wind direction. *American Economic Review* 109(12), 4178–4219.
- Di, Q., Y. Wang, A. Zanobetti, Y. Wang, P. Koutrakis, C. Choirat, F. Dominici, and J. D. Schwartz (2017). Air Pollution and Mortality in the Medicare Population. *New England Journal of Medicine* 376(26), 2513–2522. Publisher: Massachusetts Medical Society _eprint: <https://doi.org/10.1056/NEJMoa1702747>.
- Diffenbaugh, N. S. and M. Burke (2019). Global warming has increased global economic inequality. *Proceedings of the National Academy of Sciences* 116(20), 9808–9813. Publisher: National Academy of Sciences Section: Social Sciences.
- Dinda, S. (2004). Environmental Kuznets Curve Hypothesis: A Survey. *Ecological Economics* 49(4), 431–455.
- Dingel, J. I. and B. Neiman (2020). How many jobs can be done at home? *Journal of Public Economics* 189(C).
- Direction générale du Trésor (2011). Document de Travail de la DG Trésor n 2011/02 - Opportunité socio-économique d’une hausse de prix des transports collectifs franciliens.
- Doudchenko, N. and G. W. Imbens (2016). Balancing, regression, difference-in-differences and synthetic control methods: A synthesis. Technical Report 22791, National Bureau of Economic Research, Inc.
- Douenne, T. (2020). The Vertical and Horizontal Distributive Effects of Energy Taxes: A Case Study of a French Policy. *The Energy Journal* 41(3).
- Douenne, T. and A. Fabre (2021). Yellow Vests, Carbon Tax Aversion, and Biased Beliefs. *American Economic Journal: Economic Policy*, forthcoming.

- Durrmeyer, I. (2018). Winners and Losers: The Distributional Effects of the French Feebate on the Automobile Market. Technical Report 18-950, Toulouse School of Economics (TSE).
- Ellerman, A. D., C. Marcantonini, and A. Zaklan (2016). The European Union Emissions Trading System: Ten Years and Counting. *Review of Environmental Economics and Policy* 10(1), 89–107.
- Ellerman, A. D. and M. McGuinness (2008). CO2 Abatement in the UK Power Sector: Evidence from the EU ETS Trial Period. Working Paper.
- EMEP/EEA (2018). EMEP/EEA air pollutant emission inventory guidebook.
- EU Commission (2020). Handbook on the external costs of transport : version 2019 – 1.1. Website. ISBN: 9789276181842 Publisher: Publications Office of the European Union.
- European Commission (2001). Large Combustion Plants Directive (2001/80/EC).
- European Commission (2018). Air quality: Commission takes action to protect citizens from air pollution.
- European Environment Agency (2020). Air pollution: how it affects our health.
- Evans, S. (2015). Old coal and gas plants won largest share of capacity market, final results confirm . *Carbon Brief*. [Online; accessed 12. Jan. 2021].
- Evans, S. (2019). Analysis: UK electricity generation in 2018 falls to lowest level since 1994. *Carbon Brief*.
- Fischer, C. and R. G. Newell (2008). Environmental and technology policies for climate mitigation. *Journal of Environmental Economics and Management* 55(2), 142–162.
- Fisher, R. A. et al. (1937). The design of experiments. *The design of experiments*. (2nd Ed).
- Fowlie, M., S. P. Holland, and E. T. Mansur (2012). What Do Emissions Markets Deliver and to Whom? Evidence from Southern California’s NOx Trading Program. *American Economic Review* 102(2), 965–993.
- France Inter (2019). A Marseille, les bateaux polluent désormais plus que les voitures.
- France Stratégie (2019). La valeur de l’action pour le climat. Technical report.

- Frank, L. D., B. Stone, and W. Bachman (2000). Linking land use with household vehicle emissions in the central puget sound: methodological framework and findings. *Transportation Research Part D: Transport and Environment* 5(3), 173–196.
- Friedrich, A. (2017). Heading to Venice? don't forget your pollution mask. *The Guardian*.
- Gaigné, C., S. Riou, and J.-F. Thisse (2012). Are compact cities environmentally friendly? *Journal of Urban Economics* 72(2), 123–136.
- Gelman, A. and J. Carlin (2014). Beyond power calculations: Assessing type s (sign) and type m (magnitude) errors. *Perspectives on Psychological Science* 9(6), 641–651.
- Gelman, A., J. Hill, and A. Vehtari (2020). *Regression and other stories*. Cambridge University Press.
- Gillingham, K. and A. Munk-Nielsen (2019). A tale of two tails: Commuting and the fuel price response in driving. *Journal of Urban Economics* 109, 27–40.
- Giovanis, E. (2018). The relationship between teleworking, traffic and air pollution. *Atmospheric Pollution Research* 9(1), 1–14.
- Gobillon, L. and T. Magnac (2016). Regional Policy Evaluation: Interactive Fixed Effects and Synthetic Controls. *The Review of Economics and Statistics* 98(3), 535–551.
- Goulder, L. H., M. A. C. Hafstead, and R. C. Williams III (2016). General Equilibrium Impacts of a Federal Clean Energy Standard. *American Economic Journal: Economic Policy* 8(2), 186–218.
- Goulder, L. H. and I. W. H. Parry (2008). Instrument Choice in Environmental Policy. *Review of Environmental Economics and Policy* 2(2), 152–174. Publisher: Association of Environmental and Resource Economists.
- Goulder, L. H. and R. N. Stavins (2011). Challenges from State-Federal Interactions in US Climate Change Policy. *American Economic Review* 101(3), 253–257.
- GPMM - Port de Marseille Fos (2020). Yearly figures. Technical report.
- Graff Zivin, J. and M. Neidell (2013). Environment, Health, and Human Capital. *Journal of Economic Literature* 51(3), 689–730.
- Grange, S. K. and D. C. Carslaw (2019). Using meteorological normalisation to detect interventions in air quality time series. *Science of The Total Environment* 653, 578–588.

- Green, J. F. (2021). Does carbon pricing reduce emissions? A review of ex-post analyses. *Environmental Research Letters*.
- Gugler, K., A. Haxhimusa, and M. Liebensteiner (2020). Carbon Pricing and Emissions: Causal Effects of Britain’s Carbon Tax. *Working Paper*.
- Gugler, K., A. Haxhimusa, and M. Liebensteiner (2021). Effectiveness of climate policies: Carbon pricing vs. subsidizing renewables. *Journal of Environmental Economics and Management* 106, 102405.
- Guo, B. and D. Newbery (2020). The Cost of Trade Distortion: Britain’s Carbon Price Support and Cross-border Electricity Trade. Technical Report 2014, Faculty of Economics, University of Cambridge. Publication Title: Cambridge Working Papers in Economics.
- Guo, B., D. M. Newbery, and G. Castagneto Gissey (2019). The Impact of Unilateral Carbon Taxes on Cross-Border Electricity Trading. Technical Report 1951, Faculty of Economics, University of Cambridge.
- Hallegatte, S. and J. Rozenberg (2017). Climate change through a poverty lens. *Nature Climate Change* 7(4), 250–256. Number: 4 Publisher: Nature Publishing Group.
- Healy, K. (2018). *Data visualization: a practical introduction*. Princeton University Press.
- Hirst, D. (2018). Carbon Price Floor (CPF) and the price support mechanism. Technical Report Number 05927, House of Commons Library.
- HM Revenue & Customs (2017). Excise Notice CCS 1/6: a guide to carbon price floor. Technical report.
- Ho, D. E., K. Imai, G. King, and E. A. Stuart (2007). Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Political analysis* 15(3), 199–236.
- Holland, P. W. (1986a). Statistics and Causal Inference. *Journal of the American Statistical Association* 81(396), 945–960. Publisher: [American Statistical Association, Taylor & Francis, Ltd.].
- Holland, P. W. (1986b). Statistics and causal inference. *Journal of the American statistical Association* 81(396), 945–960.
- Holland, S. P., E. T. Mansur, N. Z. Muller, and A. J. Yates (2016). Are There Environmental Benefits from Driving Electric Vehicles? The Importance of Local Factors. *American Economic Review* 106(12), 3700–3729.

- Hsiang, S., P. Oliva, and R. Walker (2019). The Distribution of Environmental Damages. *Review of Environmental Economics and Policy* 13(1), 83–103. Publisher: The University of Chicago Press.
- Hsiang, S. M. (2016). Climate Econometrics. Technical Report 22181, National Bureau of Economic Research, Inc.
- Iacus, S. M., G. King, and G. Porro (2011). Multivariate Matching Methods That Are Monotonic Imbalance Bounding. *Journal of the American Statistical Association* 106(493), 345–361. Publisher: Taylor & Francis _eprint: <https://doi.org/10.1198/jasa.2011.tm09599>.
- IEA (2021). Global carbon dioxide emissions are set for their second-biggest increase in history.
- Imbens, G. W. (2015). Matching methods in practice: Three examples. *Journal of Human Resources* 50(2), 373–419.
- Imbens, G. W. and D. B. Rubin (2015). *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press.
- Imbens, G. W. and J. M. Wooldridge (2009, March). Recent Developments in the Econometrics of Program Evaluation. *Journal of Economic Literature* 47(1), 5–86.
- INSEE (2020). Populations legales - 2017 - commune de Marseille (13055).
- Intergovernmental Panel on Climate Change (2015). *Climate Change 2014: Mitigation of Climate Change: Working Group III Contribution to the IPCC Fifth Assessment Report*. Cambridge: Cambridge University Press.
- IPCC (2014). Climate Change 2014: Impacts, Adaptation, and Vulnerability. Technical report.
- IPCC (2018). Summary for Policymakers — Global Warming of 1.5 °C. Technical report.
- IPCC, E. Somanathan, T. Sterner, T. Sugiyama, D. Chimanikire, N. K. Dubash, J. Essandoh-Yeddu, S. Fifita, L. Goulder, A. Jaffe, X. Labandeira, S. Managi, C. Mitchell, J. P. Montero, and F. Teng (2014). National and Sub-national Policies and Institutions. In *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 University Press, Cambridge, United Kingdom and New York, NY, USA.

- Isaksen, E. and B. Johansen (2020). Congestion pricing, air pollution, and individual-level behavioral responses.
- Isaksen, E. T. (2020). Have international pollution protocols made a difference? *Journal of Environmental Economics and Management* 103, 102358.
- Ivanova, D. and R. Wood (2020). The unequal distribution of household carbon footprints in Europe and its link to sustainability. *Global Sustainability* 3.
- Javaid, A., F. Creutzig, and S. Bamberg (2020). Determinants of low-carbon transport mode adoption: systematic review of reviews. *Environmental Research Letters* 15(10), 103002.
- Jewell, J., V. Vinichenko, L. Nacke, and A. Cherp (2019). Prospects for powering past coal. *Nature Climate Change* 9(8), 592–597. Number: 8 Publisher: Nature Publishing Group.
- Karjalainen, P., L. Pirjola, J. Heikkilä, T. Lähde, T. Tzamkiozis, L. Ntziachristos, J. Keskinen, and T. Rönkkö (2014). Exhaust particles of modern gasoline vehicles: A laboratory and an on-road study. *Atmospheric Environment* 97, 262–270.
- Karlsson, M., E. Alfredsson, and N. Westling (2020). Climate policy co-benefits: a review. *Climate Policy* 20(3), 292–316. Publisher: Taylor & Francis _eprint: <https://doi.org/10.1080/14693062.2020.1724070>.
- Keele, L., C. McConnaughey, and I. White (2012). Strengthening the experimenter’s toolbox: Statistical estimation of internal validity. *American Journal of Political Science* 56(2), 484–499.
- Khomenko, S., M. Cirach, E. Pereira-Barboza, N. Mueller, J. Barrera-Gómez, D. Rojas-Rueda, K. d. Hoogh, G. Hoek, and M. Nieuwenhuijsen (2021). Premature mortality due to air pollution in European cities: a health impact assessment. *The Lancet Planetary Health* 5(3), e121–e134. Publisher: Elsevier.
- Kim, M.-K. and T. Kim (2016). Estimating impact of regional greenhouse gas initiative on coal to gas switching using synthetic control methods. *Energy Economics* 59, 328 – 335.
- Kong Chyong, C., B. Guo, and D. Newbery (2020). The Impact of a Carbon Tax on the CO2 Emissions Reduction of Wind. *The Energy Journal* 41(1).

- Kotchenruther, R. A. (2017). The effects of marine vessel fuel sulfur regulations on ambient PM_{2.5} at coastal and near coastal monitoring sites in the US. *Atmospheric Environment* 151, 52–61.
- Kraus, S. and N. Koch (2021). Provisional COVID-19 infrastructure induces large, rapid increases in cycling. *Proceedings of the National Academy of Sciences* 118(15). ISBN: 9782024399117 Publisher: National Academy of Sciences Section: Social Sciences.
- Larcom, S., F. Rauch, and T. Willems (2017). The Benefits of Forced Experimentation: Striking Evidence from the London Underground Network. *The Quarterly Journal of Economics* 132(4), 2019–2055.
- Lavaine, E. and M. Neidell (2017). Energy Production and Health Externalities: Evidence from Oil Refinery Strikes in France. *Journal of the Association of Environmental and Resource Economists* 4(2), 447–477.
- Le Monde (2018). L’interdiction des véhicules les plus polluants étendue au Grand Paris à partir de juillet. *Le Monde.fr*.
- Le Monde (2021). Comment la révolution de la bicyclette bouleverse les villes françaises. *Le Monde.fr*.
- Le Parisien (2019). Grand Paris : moins d’une ville sur deux s’engage contre les véhicules polluants.
- Le Quéré, C., R. B. Jackson, M. W. Jones, A. J. P. Smith, S. Abernethy, R. M. Andrew, A. J. De-Gol, D. R. Willis, Y. Shan, J. G. Canadell, P. Friedlingstein, F. Creutzig, and G. P. Peters (2020). Temporary reduction in daily global CO₂ emissions during the COVID-19 forced confinement. *Nature Climate Change* 10(7), 647–653. Number: 7 Publisher: Nature Publishing Group.
- Lee, K. and R. T. Melstrom (2018). Evidence of increased electricity influx following the regional greenhouse gas initiative. *Energy Economics* 76, 127–135.
- Lennox, J. (2020). More working from home will change the shape and size of cities. Centre of Policy Studies/IMPACT Centre Working Paper, Victoria University, Centre of Policy Studies/IMPACT Centre.
- Lepeule, J., F. Laden, D. Dockery, and J. Schwartz (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), 965–970.

- Levinson, A. (2018). Energy Efficiency Standards Are More Regressive Than Energy Taxes: Theory and Evidence. *Journal of the Association of Environmental and Resource Economists* 6(S1), S7–S36.
- Levinson, A. and J. O’Brien (2018). Environmental Engel Curves: Indirect Emissions of Common Air Pollutants. *The Review of Economics and Statistics* 101(1), 121–133.
- Linn, J. (2019). Interactions between Climate and Local Air Pollution Policies: The Case of European Passenger Cars. *Journal of the Association of Environmental and Resource Economists* 6(4), 709–740.
- Liu, N., Z. Ma, and J. Kang (2017). A regional analysis of carbon intensities of electricity generation in China. *Energy Economics* 67, 268–277.
- Maestre-Andres, S., S. Drews, and J. v. d. Bergh (2019). Perceived fairness and public acceptability of carbon pricing: a review of the literature. *Climate Policy* 19(9), 1186–1204.
- Mahony, T. O. (2013). Decomposition of Ireland’s carbon emissions from 1990 to 2010: An extended Kaya identity. *Energy Policy* 59, 573–581.
- Martin, R., L. B. de Preux, and U. J. Wagner (2014). The impact of a carbon tax on manufacturing: Evidence from microdata. *Journal of Public Economics* 117, 1–14.
- Martin, R., M. Muûls, and U. J. Wagner (2016). The Impact of the European Union Emissions Trading Scheme on Regulated Firms: What Is the Evidence after Ten Years? *Review of Environmental Economics and Policy* 10(1), 129–148.
- Mattauch, L., M. Ridgway, and F. Creutzig (2016). Happy or liberal? Making sense of behavior in transport policy design. *Transportation Research Part D: Transport and Environment* 45, 64–83.
- Mayer, M. (2019). missranger: Fast imputation of missing values. r package version 2.1.0.
- Merico, E., A. Donato, A. Gambaro, D. Cesari, E. Gregoris, E. Barbaro, A. Dinoi, G. Giovanelli, S. Masieri, and D. Contini (2016). Influence of in-port ships emissions to gaseous atmospheric pollutants and to particulate matter of different sizes in a Mediterranean harbour in Italy. *Atmospheric environment* 139, 1–10.
- Metcalf, G. E. (2009a). Designing a Carbon Tax to Reduce U.S. Greenhouse Gas Emissions. *Review of Environmental Economics and Policy* 3(1), 63–83.

- Metcalf, G. E. (2009b). Market-Based Policy Options to Control U.S. Greenhouse Gas Emissions. *Journal of Economic Perspectives* 23(2), 5–27.
- Metcalf, G. E. and J. H. Stock (2020). The Macroeconomic Impact of Europe’s Carbon Taxes. Technical Report w27488, National Bureau of Economic Research.
- Micali, S. and V. V. Vazirani (1980). An $O(|V| |C| |E|)$ algorithm for finding maximum matching in general graphs. In *21st Annual Symposium on Foundations of Computer Science (sfcs 1980)*, pp. 17–27. IEEE.
- Ministère de la Transition écologique et solidaire (2018). Information GES des prestations de transport.
- Mocerino, L., F. Murena, F. Quaranta, and D. Toscano (2020). A methodology for the design of an effective air quality monitoring network in port areas. *Scientific reports* 10(1), 1–10.
- Moretti, E. and M. Neidell (2011). Pollution, health, and avoidance behavior evidence from the ports of los angeles. *Journal of human Resources* 46(1), 154–175.
- Mueller, D., S. Uibel, M. Takemura, D. Klingelhofer, and D. A. Groneberg (2011). Ships, ports and particulate air pollution-an analysis of recent studies. *Journal of Occupational Medicine and Toxicology* 6(1), 1–6.
- Muller, N. Z. and R. Mendelsohn (2007). Measuring the damages of air pollution in the United States. *Journal of Environmental Economics and Management* 54(1), 1–14.
- Murena, F., L. Mocerino, F. Quaranta, and D. Toscano (2018). Impact on air quality of cruise ship emissions in Naples, Italy. *Atmospheric Environment* 187, 70–83.
- NEEDS (2007). Final Report on the monetary valuation of mortality and morbidity risks from air pollution. Technical report.
- Nemet, G. F., T. Holloway, and P. Meier (2010). Implications of incorporating air-quality co-benefits into climate change policymaking. *Environmental Research Letters* 5(1), 014007.
- Newbery, D. M., D. M. Reiner, and R. A. Ritz (2019). The Political Economy of a Carbon Price Floor for Power Generation. *The Energy Journal* 40(1).
- Newell, R. and R. Stavins (2003). Cost Heterogeneity and the Potential Savings from Market-Based Policies. *Journal of Regulatory Economics* 23, 43–59.

- Nicolas, J.-P. and D. David (2009). Passenger transport and CO2 emissions: What does the French transport survey tell us? *Atmospheric Environment* 43(5), 1015–1020.
- Ntziachristos, L. and P. Boulter (2019). EMEP/EEA air pollutant emission inventory guidebook 2019 - 1.A.3.b.vi Road transport: Automobile tyre and brake wear. Technical report, European Environmental Agency.
- Ntziachristos, L. and S. Zissis (2020). EMEP/EEA air pollutant emission inventory guidebook 2019. Technical report, European Environmental Agency.
- OECD (2020). Non-exhaust Particulate Emissions from Road Transport: An Ignored Environmental Policy Challenge. Technical report.
- OFGEM (2013). Electricity interconnectors.
- Omnil-Ile de France Mobilites (2019). La nouvelle enquête globale transport Présentation des premiers resultats 2018. Technical report.
- Ottelin, J., J. Heinonen, and J. Seppo (2017). *Rebound effects for reduced car ownership and driving*. Routledge. Pages: 263-283 Publication Title: Nordic Experiences of Sustainable Planning.
- Oxford Institute for Energy Studies (2019). Electricity, Electric Vehicles, and Public Policy: Eight Key Takeaways. Technical report.
- Padilla, C. M., W. Kihal-Talantikite, V. M. Vieira, P. Rossello, G. Le Nir, D. Zmirou-Navier, and S. Deguen (2014). Air quality and social deprivation in four French metropolitan areas—a localized spatio-temporal environmental inequality analysis. *Environmental Research* 134, 315–324.
- Parry, I. W. H., M. Walls, and W. Harrington (2007). Automobile Externalities and Policies. *Journal of Economic Literature* 45(2), 373–399.
- Pearce, D. (1991). The Role of Carbon Taxes in Adjusting to Global Warming. *Economic Journal* 101(407), 938–48.
- Perino, G. (2018). New EU ETS Phase 4 rules temporarily puncture waterbed. *Nature Climate Change* 8, 262–264.
- Perino, G., R. A. Ritz, and A. v. Benthem (2019). Overlapping Climate Policies. Publication Title: NBER Working Papers.

- Piga, D., A. Armengaud, M. Deveze, M. Parra, N. Marchand, A. Detournay, and D. Salameh (2013). Synthèse du projet APICE - Marseille. Technical report, Technical Report.
- Pigou, A. (1920). *The Economics of Welfare* (MacMillan London ed.). Houndmills, Basingstoke, Hampshire ; New York: Palgrave Macmillan.
- Pindyck, R. S. (2013). Climate Change Policy: What Do the Models Tell Us? Technical Report w19244, National Bureau of Economic Research.
- Poterba, J. M. (1991). Is the Gasoline Tax Regressive? *Tax Policy and the Economy* 5, 145–164.
- Pottier, A., E. Combet, J.-M. Cayla, S. d. Lauretis, and F. Nadaud (2020). Qui émet du CO₂? Panorama critique des inégalités écologiques en France. Technical Report 2020.02, FAERE - French Association of Environmental and Resource Economists.
- Rabl, A. and A. de Nazelle (2012). Benefits of shift from car to active transport. *Transport Policy* 19(1), 121–131.
- Rafaty, R. (2018). Perceptions of Corruption, Political Distrust, and the Weakening of Climate Policy. *Global Environmental Politics* 18(3), 106–129. Publisher: MIT Press.
- Rafaty, R., G. Dolphin, and F. Pretis (2020). Carbon pricing and the elasticity of CO₂ emissions. Technical Report 20116, Faculty of Economics, University of Cambridge.
- Rauner, S., N. Bauer, A. Dirnaichner, R. V. Dingenen, C. Mutel, and G. Luderer (2020). Coal-exit health and environmental damage reductions outweigh economic impacts. *Nature Climate Change* 10(4), 308–312.
- Région Ile de France (2016). Changeons d’air en Ile de France: Plan régional pour la qualité de l’air (2016-2021). Technical report.
- Rivers, N. and B. Schaufele (2015). Salience of carbon taxes in the gasoline market. *Journal of Environmental Economics and Management* 74, 23–36.
- Rocha, M., B. Hare, P. Yanguas Parra, N. Roming, U. Ural, A. Ancygier, J. Cantzler, F. Sferra, H. Li, and M. Schaeffer (2016). Implications of the Paris Agreement for Coal Use in the Power Sector. Technical report, Climate Analytics.
- Rogelj, J., M. den Elzen, N. Höhne, T. Fransen, H. Fekete, H. Winkler, R. Schaeffer, F. Sha, K. Riahi, and M. Meinshausen (2016). Paris Agreement climate proposals need a boost to keep warming well below 2 degree celsius. *Nature* 534(7609), 631–639. Number: 7609 Publisher: Nature Publishing Group.

- Rojas-Rueda, D., A. d. Nazelle, M. Tainio, and M. J. Nieuwenhuijsen (2011). The health risks and benefits of cycling in urban environments compared with car use: health impact assessment study. *BMJ* 343, d4521. Publisher: British Medical Journal Publishing Group Section: Research.
- Rosenbaum, P. R. et al. (2010). *Design of observational studies*, Volume 10. Springer.
- Roth, S. (2017). Air pollution, educational achievements, and human capital formation. *IZA World of Labor*. Publisher: Institute of Labor Economics (IZA).
- Rubin, D. B. (1974a). Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies. *Journal of Educational Psychology*.
- Rubin, D. B. (1974b). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology* 66(5), 688.
- Rubin, D. B. (1978). Bayesian inference for causal effects: The role of randomization. *The Annals of statistics*, 34–58.
- Rubin, D. B. (1991). Practical implications of modes of statistical inference for causal effects and the critical role of the assignment mechanism. *Biometrics*, 1213–1234.
- Rubin, D. B. (2005). Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association* 100(469), 322–331.
- Rubin, D. B. (2006). *Matched sampling for causal effects*. Cambridge University Press.
- Sager, L. (2019). Income inequality and carbon consumption: Evidence from Environmental Engel curves. *Energy Economics* 84, 104507.
- Sallee, J. M. (2019). Pigou Creates Losers: On the Implausibility of Achieving Pareto Improvements from Efficiency-Enhancing Policies. Technical Report 25831, National Bureau of Economic Research, Inc.
- Santé Publique France (2021). Impact de pollution de l’air ambiant sur la mortalité en France métropolitaine. Réduction en lien avec le confinement du printemps 2020 et nouvelles données sur le poids total pour la période 2016-2019. Technical report.
- Schneider, B. (2018). Why Little Vehicles Will Conquer the City. *Bloomberg.com*.
- Sekhon, J. S. (2009). Opiates for the matches: Matching methods for causal inference. *Annual Review of Political Science* 12, 487–508.

- Shearer, C., N. Mathew-Shah, L. Myllyvirta, A. Yu, and T. Nace (2019). Boom and bust 2019. tracking the global coal plant pipeline. pp. 14.
- Sicard, P., E. Agathokleous, A. De Marco, E. Paoletti, and V. Calatayud (2021). Urban population exposure to air pollution in Europe over the last decades. *Environmental Sciences Europe* 33(1), 28.
- Singhal, P. (2019). Are Emission Performance Standards Effective in Pollution Control? Evidence from the EU’s Large Combustion Plant Directive. SSRN Scholarly Paper ID 3297528, Social Science Research Network, Rochester, NY.
- Sommer, A. J., M. Lee, and M.-A. C. Bind (2018). Comparing apples to apples: an environmental criminology analysis of the effects of heat and rain on violent crimes in boston. *Palgrave communications* 4(1), 1–10.
- Sommer, A. J., E. Leray, Y. Lee, and M.-A. C. Bind (2021). Assessing environmental epidemiology questions in practice with a causal inference pipeline: An investigation of the air pollution-multiple sclerosis relapses relationship. *Statistics in Medicine* 40(6), 1321–1335.
- Sommer, A. J., A. Peters, J. Cyrus, H. Grallert, D. Haller, C. L. Müller, and M.-A. C. Bind (2021). A randomization-based causal inference framework for uncovering environmental exposure effects on human gut microbiota. *bioRxiv*.
- Sorte, S., V. Rodrigues, C. Borrego, and A. Monteiro (2020). Impact of harbour activities on local air quality: A review. *Environmental Pollution* 257, 113542.
- Splawa-Neyman, J., D. M. Dabrowska, and T. Speed (1990). On the application of probability theory to agricultural experiments. essay on principles. section 9. *Statistical Science*, 465–472.
- Spulber, D. F. (1985). Effluent regulation and long-run optimality. *Journal of Environmental Economics and Management* 12(2), 103–116.
- Staffell, I. (2017). Measuring the progress and impacts of decarbonising British electricity. *Energy Policy* 102, 463–475.
- Stuart, E. A. (2010). Matching methods for causal inference: A review and a look forward. *Statistical science: a review journal of the Institute of Mathematical Statistics* 25(1), 1.
- Sustainable Energy Authority of Ireland (2014). Quantifying Ireland’s Fuel and CO2 Emissions Savings from Renewable Electricity in 2012.pdf. Technical report.

- Taconet, N., C. Guivarch, and A. Pottier (2021). Social Cost of Carbon Under Stochastic Tipping Points. *Environmental and Resource Economics* 78(4), 709–737.
- Timm, A. (2019). Retrodesign: Tools for Type s (Sign) and Type M (Magnitude) Errors. R package version 0.1.0.
- Transport & Environment (2019). One corporation to pollute them all. Technical report, Transport & Environment.
- Tukey, J. W. et al. (1977). *Exploratory data analysis*, Volume 2. Reading, Mass.
- US EPA (2016). Air Quality - National Summary.
- Van den Bergh, K. and E. Delarue (2015). Quantifying CO2 abatement costs in the power sector. *Energy Policy* 80, 88–97.
- Vaughan, A. (2018). UK government spells out plan to shut down coal plants. *The Guardian*.
- Viana, M., P. Hammingh, A. Colette, X. Querol, B. Degraeuwe, I. de Vlieger, and J. van Aardenne (2014). Impact of maritime transport emissions on coastal air quality in europe. *Atmospheric Environment* 90, 96–105.
- Viana, M., T. A. Kuhlbusch, X. Querol, A. Alastuey, R. M. Harrison, P. K. Hopke, W. Winiwarter, M. Vallius, S. Szidat, A. S. Prévôt, et al. (2008). Source apportionment of particulate matter in europe: a review of methods and results. *Journal of aerosol science* 39(10), 827–849.
- Viegas, J. M. (2001). Making urban road pricing acceptable and effective: searching for quality and equity in urban mobility. *Transport Policy* 8(4), 289–294.
- Wagner, U. and L. De Preux (2016). The Co-benefits of Climate Policy: Evidence from the eu Emissions Trading Scheme.
- Wang, C., J. Chen, and J. Zou (2005). Decomposition of energy-related CO2 emission in China: 1957–2000. *Energy* 30(1), 73–83.
- Wardman, M., J. Toner, N. Fearnley, S. Flügel, and M. Killi (2018). Review and meta-analysis of inter-modal cross-elasticity evidence. *Transportation Research Part A: Policy and Practice* 118, 662–681.
- Weitzman, M. (2015). Internationally-Tradable Permits Can Be Riskier for a Country than an Internally-Imposed Carbon Price. Technical report, Department of Economics, Harvard University.

- Weitzman, M. L. (1974). Prices vs. Quantities. *The Review of Economic Studies* October, 41(4), 477–491.
- WHO (2013). Health risks of air pollution in Europe – HRAPIE project. Recommendations for concentration–response functions for cost–benefit analysis of particulate matter, ozone and nitrogen dioxide.
- WHO (2014). 7 million premature deaths annually linked to air pollution.
- WHO (2016). WHO | Air pollution levels rising in many of the world’s poorest cities.
- Wilson, I. A. G. and I. Staffell (2018). Rapid fuel switching from coal to natural gas through effective carbon pricing. *Nature Energy* 3(5), 365–372. Number: 5 Publisher: Nature Publishing Group.
- World Bank (2021). State and Trends of Carbon Pricing 2021. Technical report, Washington, DC.
- World Bank and Ecofys (2018). State and Trends of Carbon Pricing 2018.
- World Health Organization (2018). Ambient (outdoor) air pollution.
- Wright, M. N. and A. Ziegler (2017). ranger: A Fast Implementation of Random Forests for High Dimensional Data in C++ and R. *Journal of Statistical Software* 77(1). arXiv: 1508.04409.
- Yang, Y., C. Wang, and W. Liu (2018). Urban daily travel carbon emissions accounting and mitigation potential analysis using surveyed individual data. *Journal of Cleaner Production* 192, 821–834.
- Zhao, J. and L. Mattauch (2020). When Standards Have Better Distributional Consequences Than Carbon Taxes. SSRN Scholarly Paper ID 3739546, Social Science Research Network, Rochester, NY.
- Zhu, J. and J. Wang (2021). The effects of fuel content regulation at ports on regional pollution and shipping industry. *Journal of Environmental Economics and Management* 106, 102424.

Appendices

Appendix A

Appendix to Chapter 1: Carbon Pricing and Power sector Decarbonisation: Evidence from the UK

A.1 Appendix

A.1.1 Evolution of per capita emissions, demand, trade and emission intensity in the UK and other countries

Per capita power sector emissions : Figure A.1.1 shows the evolution of per capita power sector emissions for the UK and twenty other European countries¹, using emission data from the European carbon market described in section 3.2. While the UK was among the top emitters before 2013, by 2017 it had joined the bulk of lower-emitting countries. The figure also shows that most countries tend to have stable emissions per capita, except for a few outliers, which emissions are shown in dashed or dotted lines².

The CPS may impact each of the three channels mentioned in section 2.2: demand, because the CPS increases the marginal cost of producing electricity and generators are likely to at least partially pass on this cost to consumers, as evidenced in Guo et al. (2019) and Ares and Delebarre (2016). Trade, because the CPS increases the relative cost of domestically produced electricity compared to imported electricity. Emission intensity, because the CPS can directly impact the fuel mix used for electricity generation, in the short-term and in the long-term. In the short-term, the higher tax on coal relative to gas increases the cost of coal-fired relative to gas-fired power generation. In the long-term, the CPS also makes it less profitable to run fossil fuel high-emitting plants, and might dampen investments in those plants to the benefit of low-carbon generation (Van den Bergh and Delarue, 2015). In the next paragraphs, I show how demand, trade and emission intensity

¹all 28 EU countries except Romania, Bulgaria, Slovenia, Croatia, Malta, Cyprus, and Luxembourg, which are not included in the empirical analysis (see section 3.2)

²Estonia's emissions are both high on average and with a high variance; Czech Republic has the highest average after Estonia; Greece has decreasing emissions after 2012; Finland and Denmark's emissions have a high variance, likely due to the inter-annual variation in available hydro resources in Finland, and hydro and wind resources in Denmark.




Figure A.1.1 – Per capita power sector emissions in European countries

Notes: emission values were obtained by aggregating plant-level emission data for ETS participants identified as power generators at the country level. Per capita emissions were obtained by dividing total emissions by the annual country population. “Other countries” include twenty European countries: Austria, Belgium, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Latvia, Lithuania, the Netherlands, Poland, Portugal, Slovakia, Spain, Sweden.

evolved over time in the UK compared to other European countries.

Demand, trade and intensity channels Figure A.1.2 shows the evolution of the three channels of demand, trade and emission intensity with the same method as Figure 1.2.2a. Demand is measured with power consumption per capita, trade is measured with net electricity imports per capita, and emission intensity is measured with emissions per gigawatt-hour of domestic power production. Demand decreases over the whole period in the UK compared to other European countries (figure A.1.2a). There is no obvious break in trend in 2013. This continuous decrease is consistent with the continuous improvement of energy efficiency in buildings and electric appliances in the UK since 2009 (Staffell, 2017). The lack of a visible link between the CPS and a change in demand can be linked to the finding by Chyong et al. (2019) that only about 60% of the CPS cost has been passed through to the GB day-ahead electricity market. It can also be explained by the financial compensation received by electro-intensive industries to cushion the price effect of the CPS and protect their competitiveness³(Hirst, 2018).

Regarding the trade channel, UK net imports per capita are low compared to other countries (figure A.1.2b). Being an island, the UK can only trade electricity via under-sea cables and has a limited trading capacity. At the period of interest, the UK can trade electricity with three countries only: via undersea interconnectors with France, the Netherlands and Ireland from Great Britain, and via ground connections to the Republic

³Electro-intensive industries have received a compensation of around £100 million for the period April 2013 to March 2015, and the support has been extended to 2019-2020. This compensation scheme is a specific component of a larger Energy Intensive Industries support introduced to compensate the cost increase induced by climate change policies (Hirst, 2018).

of Ireland from Northern Ireland. Between 2010 and 2012, the UK increased its trading capacity by 50%⁴, and UK net imports increased by from 2,661 GWh to 11,864 GWh. Net imports increased further between 2012 and 2015, but at a lower rate (from 11,864 GWh to 20,938 GWh), before decreasing again until 2017. Although imports started to increase before 2013, their increase after 2013 could be associated with the CPS. Using an econometric model of electricity trade, Guo and Newbery (2020) estimate that the CPS increased GB imports by 12,000 Gigawatt-hour per year between 2015 and 2018, after market coupling with France and the Netherlands in 2014. However, taken per capita and compared to other European countries, this increase in trade remains low: between the 2005-2012 and the 2013-2017 period, UK net imports per capita increase from representing 2% of gross electricity consumption to representing 5% of gross electricity consumption.

In contrast to demand and trade, the UK emission intensity of power generation decreased markedly after 2013 compared to most other European countries (Figure A.1.2c)⁵.

The emission intensity $Q_{CO_2e}Q_{elec}$ can be further decomposed as:


$$\frac{Q_{CO_2e}}{Q_{elec}} = \sum_i e_i q_i \quad (\text{A.1})$$

Where e_i is the average emission intensity of technology i used for electricity generation and q_i represents the share of gross electricity production covered by technology i . Power generation with renewable and nuclear energy sources is emission-free⁶. What matters in this equation, therefore, is the share of fossil fuel in total electricity generation on the one hand, and the emission intensity of fossil fuel generation on the other hand. Figure A.1.3 shows the technologies/fuels used for power generation for each European country in 2012 and in 2017. The countries are ranked by their 2012 coal share. The UK coal share fell by 30 percentage points between the two periods, while there was only little variation in most other European countries. The decrease in the coal share was compensated by an increase in the gas share (+ 14 percentage points (pp)), in the share of non-biomass renewables (+ 9 pp) in the biomass share (+ 5 pp) and in the nuclear share (+ 2 pp).


⁴GB became interconnected with the Netherlands in 2011, and in 2012 a new undersea interconnector with the Republic of Ireland was completed (OFGEM, 2013)

⁵The outlier with large variations in the emission intensity is Finland, again due to a large inter-annual variation in generation from renewables.


⁶these technologies embody some life-cycle emissions, but generation itself does not emit CO₂. An exception is for plants using biomass: they do release greenhouse gases, but are not bound to pay the ETS price nor the CPS because the carbon released when solid biomass is burned is expected to be re-absorbed during tree growth: see https://ec.europa.eu/clima/sites/clima/files/ets/docs/com_2018_842_final_en.pdf



(a) Demand per capita



(b) Net imports per capita



(c) Emissions per unit of electricity output

Figure A.1.2 – Channels: evolution of electricity demand, trade and emission intensity in the UK and other European countries

Notes: the variables appearing on these two graphs were obtained by taking the difference between the original variable and the 2005-2012 average. "Other countries" include twenty European countries: Austria, Belgium, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Latvia, Lithuania, the Netherlands, Poland, Portugal, Slovakia, Spain, Sweden.




Figure A.1.3 – Power sector’s input fuel mix in EU countries, 2012 and 2017

Notes: EU countries are ranked by ascending order of the share of coal in electricity generation in 2012, from left to right. LT: Lithuania; LV: Latvia; SE: Sweden; EE: Estonia; FR: France; BE: Belgium; AT: Austria; SK: Slovakia; FI: Finland; IT: Italy; HU: Hungary; ES: Spain; NL: The Netherlands; IE: Ireland; PT: Portugal; DK: Denmark; UK: United Kingdom; DE: Germany; CZ: Czech Republic; GR: Greece; PL: Poland. The legend from left to right corresponds to the histogram bar colors from bottom (coal and lignite) to top (biomass). Data come from Eurostat. Renewables include production from hydro, solar, wind, and tide, wave and ocean.

A.1.2 Potential confounders

Here I give more details on policies implemented at the EU or UK level around the same time period, which may have contributed to an emission decrease

European level: LCP and IED Directives Nine UK plants opted out from the LCP Directive and shut down between 2012 and 2015, and three other plants opted out partially (Source: EEA website). These 12 fully or partly opted out plants represented 11% of UK power sector emissions in 2011. The LCPD-induced plant closures could explain part of the pattern seen on figures A.1.2c and A.1.3 if the choice to opt-out and shut down occurred disproportionately more in the UK than in other European countries.

The LCPD was replaced by the IED directive in 2016. The IED Directive was enacted in 2010 and has a similar opt-out option as the LCPD. Plants had to decide by 2013 whether they wanted to opt-out or not under a limited lifetime derogation (LLD). Under the LLD opt-out option, plants are exempted from the emission standards but cannot operate for more than 17,500 hours between 1 January 2016 and 31 December 2023, and have to shut down once they have run for 17,500 hours or in 2023, whichever comes first. Two UK power plants opted out from the IED Directive. Given that plants had until 2013 to decide whether to opt-out or not and the CPS was announced in 2011, the IED opt-out decision was endogenous to the CPS.

UK level: support to biomass conversion The UK government supported the conversion of coal-fired power plants to biomass starting from 2012. This support first took the form of dedicated Renewables Obligation Certificates (ROCs). The ROCs were part of the broader Renewables Obligation scheme designed to support the deployment of large-scale renewable electricity generation; they created an obligation for electricity suppliers to source a proportion of their electricity from plants with ROCs. The ROCs were replaced by the FID Enabling for Renewables scheme in 2012 and then by the Contract for Difference scheme (CfD) introduced as part of the Electricity Market Reform in 2014. Two power plants received government support for conversion to biomass: Drax power station, representing 14% of UK power sector emissions in 2012, had already started to co-fire biomass in 2004. The company owning the station announced its intention to fully convert three of its six units to biomass in September 2012. The station benefited from Renewable Obligations Certificates for the conversion of its first unit, which was completed in 2013, and from Contracts for Differences for the conversion of its second and third units, which were completed respectively in 2014 and 2016. By the end of the period considered in this analysis, only three of the six units were converted to biomass, and the three remaining units continued to run with coal⁷. Lynemouth power station, a smaller plant, also received support under the FID Enabling for Renewables scheme. The station stopped burning coal and started the biomass conversion process in December 2015⁸.

The biomass conversion of these two plants, partial but early in the case of Drax, and full but later in the case of Lynemouth, led to a decline in their carbon emissions over the 2013-2017 period, since biomass is considered a zero-emission fuel. The dates of introduction of the CPS and of the support policies for biomass conversion are close. The UK government may have decided to subsidise conversion from coal to biomass partly to reduce the economic costs associated with the CPS for coal plant owners, and to facilitate the low-carbon transition. In this case, the biomass conversion could be viewed as a direct consequence of the CPS. But the two policies may also be independent, in which case the biomass conversion would have occurred even in the absence of the CPS.

UK level: support to renewable energy Second, The FID Enabling for Renewables and CfD programmes could have impacted the fuel mix more broadly than via its impact on the conversion of coal plants to biomass, by increasing the share of renewable energy in UK electricity production over the 2013-2017 period. I combine data on projects being awarded a Contract under the FID Enabling for Renewables or CfD programmes

⁷<https://www.drax.com/about-us/our-history/>

⁸https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/805441/LCP_Review_Lynemouth_DD-FP3137CG-V009-draftdecision.pdf

between 2013 and 2017⁹ and the renewable energy planning database put together by the UK Department for Business, Energy and Industrial Strategy (BEIS) monitoring all UK-based renewable projects¹⁰ to estimate the renewable capacity which was installed over the 2013-2017 period as a direct consequence of FID of the CfDs. Compared to the planned delivery year of projects mentioned in the CfD auction results, the actual date where a CfD project becomes operational is often delayed by several months. Given the actual date where each project becomes “operational” according to BEIS’s renewable energy planning database, only 1,035 MW of renewable capacity from CfD projects is operational by the end of 2017. This represents 2.4% of the total renewable capacity installed in the UK in 2017¹¹. I combine data on the date where each project becomes operational with data on running hours for each project (coming from the CfD auction results) to calculate the electricity generation associated with this installed capacity over the 2013-2017 period. I estimate that at most 1,588 GWh of electricity may have been generated by CfD projects - assuming that the theoretical number of hours given in the CfD auction results is the true number of hours. This represents just 0.4% of the total electricity generated with renewable energy sources¹² and 0.1% of total electricity generated in the UK over that period. (Abrell et al., 2019) report that the average capacity-weighted emission rate of UK coal-fired plants (taken before 2013) is 0.89 ton of CO₂ per Megawatt-hour of electricity produced. Such an emission rate implies that the 1,588 GWh of electricity produced by CfD projects would have caused emissions of $1,589.10^3 \times 0.89 = 1.4 \text{ MtCO}_2\text{e}$ over the 2013-2017 period if they had been generated by coal instead. This amount represents less than 1% of UK power sector emissions in 2013.

Other support policies to renewable energy exist in the UK, but the bundle of feed-in-tariffs, support to R&D for renewable energy, and regulatory instruments does not look fundamentally different from that implemented in other European countries, according to the IEA/IRENA Joint Policies and Measures Database listing the support policies implemented in each country since the 1970s¹³. All in all, the difference in emissions observed after 2013 in the UK does not seem driven by a renewable energy policy specific to the UK, apart from the support to biomass conversion described above.

⁹available at: <https://www.gov.uk/government/publications/contracts-for-difference/contract-for-difference>

¹⁰available at: <https://www.gov.uk/government/publications/renewable-energy-planning-data-base-monthly-extract>

¹¹The denominator includes waste and biofuels, since part of the CfD projects are for power generation from waste

¹²The denominator includes waste and biofuels, since part of the CfD projects are for power generation from waste

¹³The database can be accessed here: <https://www.iea.org/policiesandmeasures/renewableenergy/>. For the electricity sector only, there were more than 200 support measures in force in the 28 EU countries over the period considered. Examining each piece of legislation goes beyond the scope of this article.

UK level: capacity market The capacity market introduced in 2013 may impact emissions in two ways: first, before the capacity payment starts, securing a capacity contract can incentivize investing in new capacity, as auction payments can be seen as a subsidy for new power generation. Indeed, the capacity market was initially supposed to facilitate investments in new gas capacity (Evans, 2015). Second, once the capacity payment starts, the payment can keep a plant being economically profitable even with at low generation levels. The first auction took place in 2014 for capacity secured for 2018. Since my 2013-2017 period of analysis is before the auction payments start, the capacity market can only impact UK emissions via the first channel. I combine data on new-build plants being awarded a contract between 2014 and 2017¹⁴ and data listing all UK power plants with a capacity greater than 20MW with the year of commission or year generation began¹⁵, to gauge if the capacity market incentivized the construction of plants having a lower-than-average emission intensity over the 2013-2017 period. Only six plants were awarded a capacity contract in 2014, 2015, 2016 or 2017 and had a date of commissioning/where generation began between 2014 and 2017. One is a large gas-fired plant (CCGT), Carrington power station and the five other are smaller waste plants. The opening of Carrington power station cannot be imputed to the capacity market because the plant started being constructed in 2009¹⁶. On the other hand, the five waste plants may have opened as a direct consequence of being awarded a capacity contract, and have a lower-than-average emission intensity. These five plants represent an installed capacity of 1,141 MW. To estimate the associated power generation, I make several assumptions: first, I assume that the year of commissioning/where generation began indicated in the list of UK power plants is when generation began, and that generation began on January 1. Second, I take as a load factor the average load factor for conventional steam plants in the UK averaged over 2013-2017¹⁷, which is 35%. I obtain an upper bound of the low-carbon power generation imputable to the capacity market of 2,590 MWh over the 2013-2017 period. This represents 0.6% of electricity generated with renewable sources and 0.2% of total electricity generated in the UK over the 2013-2017 period. Such generation would have caused CO₂e emissions of $2,590 \cdot 10^3 \times 0.89 = 2,305$ tCO₂e over the 2013-2017 period if it had been generated with coal instead of renewable waste (similar calculations as the one used to estimate the emission reduction caused by the CfD). This amount represents just 1.5% of UK power sector emissions in 2013.

¹⁴available at: <https://www.emrdeliverybody.com/CM/Registers.aspx>

¹⁵available at: <https://www.gov.uk/government/statistics/electricity-chapter-5-digest-of-united-kingdom-energy-statistics-dukes>, DUKES 5.11 file

¹⁶https://en.wikipedia.org/wiki/Carrington_Power_Station

¹⁷data available at: <https://www.gov.uk/government/statistics/electricity-chapter-5-digest-of-united-kingdom-energy-statistics-dukes>, DUKES 5.10 file)

A.1.3 Identification of power installations

The plant-level emission data released by the EUTL does not provide information on which plant is a power generator. The UK-based think-tank Ember (formerly Sandbag) provided a database with total verified emissions data for 2008-2016 supplemented with a variable identifying all power plants. This identification has been performed internally by Ember in two steps: in the first step, Ember carried out an exact matching based on a file circulated by the European Commission in 2014 containing a list of individual participants with their sectoral classification. This classification is based on NACE rev2, the Statistical classification of economic activities in the European Community, which contains two-digit *divisions*, divided into three-digit *groups*, themselves divided into four-digit *classes*. Power installations are generally found in division 35 "Electricity, gas, steam and air conditioning supply", group 35.1 "Electric power generation, transmission and distribution", class 35.11 "Production of electricity". Ember classified as a power installation all the ETS participants with class 35.11. In the second step, Ember identified other power installations which either were not classified in class 35.11 (for example because they were part of an industrial site), or opened after the file was circulated by the European Commission, based on desk-based research and manual matching. For the verified emissions variable, the data provided by Ember are the same as the raw data retrieved from the EUTL.

To retrieve the power plants status of the few plants that shut down before 2008 (and are thus absent from the Ember dataset), I use the "Accounts to Firms Matching" dataset hosted by the Florence School of Regulation (FSR)¹⁸, listing participating installations until 2013 with their Nace rev 2 sectoral classification. I first match the FSR and Ember data to check the quality of Ember's power sector classification. Among the installations found in both the FSR and Ember data, 100% of the division 35 installations having sectoral class 35.11 (according to the FSR data) are classified as power installations by Ember. 96% of the division 35 installations with a sectoral class different from 35.11 are also classified as power installations by Ember. Installations from sectoral class 35.11 represent only 30% of the installations classified by Ember as power installations, but 80% of the carbon emissions. Installations from division 35 but with a different class from 35.11 represent 61% of Ember's power installations, but only 7% of emissions. Installations from division 35 but with a missing division or a missing class represent another 6% of Ember's power installations, and 11% of emissions. The remaining 3% of Ember's power installations, representing just 1% of emissions, have a division that is either different from 35, or missing.

¹⁸The dataset can be downloaded on this website: <http://fsr.eui.eu/climate/ownership-links-enhanced-eutl-dataset-project/>

To be consistent with Ember’s classification, I classify the few plants only present in the FSR but not in the Ember dataset as power installation when their group is 35. This way, I identify 314 additional power plants which shut down before 2008. Finally, I identify 4 additional installations having a missing sectoral division as power installation based on their name (containing “power station” or its equivalent in one of the European languages). There are no ETS installations opening in 2017, that is, present in the 2017 EUTL data but not in the 2008-2016 Ember data. After this additional matching, the power plant status is missing for only 3% of all EU ETS installations over the 2005-2017 period, with only a quarter of them having non-zero CO₂e emissions for at least one period.

A.1.4 Summary statistics for the country-level dataset

	(1)	(2)	(3)
	UK	Other countries	Donor pool
Nb. ETS power installations	190 (13)	190 (189)	193 (148)
Nb. ETS non-power installations	626 (82)	261 (269)	273 (229)
Population	62,047,417 (1,165,199)	20,142,519 (23,321,330)	20,229,390 (20,035,738)
CO ₂ e emissions from power installations (tCO ₂ e)	171,770,195 (13,562,672)	53,067,594 (78,052,009)	43,904,503 (47,414,350)
Per capita power sector emissions (tCO ₂ e per capita)	2.78 (0.26)	2.80 (2.10)	2.43 (1.48)
Final electricity consumption (GWh)	341,442 (13,197)	122,819 (147,040)	122,593 (119,366)
Gross Power production (GWh)	384,010 (13,570)	139,907 (173,760)	139,013 (143,684)
Proportion of renewables in production	0.05 (0.0183)	0.19 (0.184)	0.19 (0.175)
Proportion of nuclear	0.18 (0.0229)	0.21 (0.250)	0.24 (0.249)
Proportion of fossil fuel	0.776 (0.0316)	0.596 (0.271)	0.574 (0.271)
Proportion of coal	0.327 (0.0430)	0.226 (0.232)	0.235 (0.227)
Proportion of gas	0.398 (0.0613)	0.230 (0.185)	0.235 (0.191)
Coal price (€/kWh)	0.0101 (0.00209)	0.0152 (0.0166)	0.0170 (0.0188)
Gas price (€/kWh)	0.0201 (0.00420)	0.0254 (0.00572)	0.0258 (0.00527)
Installed capacity, fossil fuels (MW)	68,414 (2,869)	18,838 (22,481)	18,522 (18,353)
Installed capacity, wind and solar (MW)	4,732 (3,156)	4,480 (9,815)	3,363 (5,877)
Lignite resource dummy	0.000 (0.000)	0.250 (0.434)	0.200 (0.402)
Average age of coal-fired plants	36.04 (1.537)	28.57 (6.698)	28.86 (7.055)
Observations	8	160	120

Table A.1.1 – Summary statistics at the country level, average 2005-2012

Notes: standard deviations in parentheses; "Other countries" include twenty European countries: Austria, Belgium, Czech Republic, Denmark, Estonia, Finland, France, Germany, Greece, Hungary, Ireland, Italy, Latvia, Lithuania, the Netherlands, Poland, Portugal, Slovakia, Spain, Sweden; The "Donor pool" includes fifteen countries: Austria, Belgium, Czech Republic, Denmark, Finland, France, Hungary, Ireland, Italy, the Netherlands, Poland, Portugal, Slovakia, Spain, Sweden; for the UK, the values are averaged over the 2005-2012 period; for the "Other countries" and the "Donor pool", the values are averaged over the 2005-2012 period, then averaged across countries (without population weights).

A.1.5 Data sources

Coal-to-gas price ratio (predictor used in the main specification): To my knowledge, there is no harmonized series of country-level coal-to-gas price ratio for the period considered. I thus build a price ratio variable combining coal trade data from Eurostat and gas wholesale price data for large industrial consumers, also from Eurostat. For coal, I use annual trade data for imported coal from Comext, the official EU trade statistics. I aggregate the price and volume data for all the subcategories of coal that may be used for coal generation and obtain average nominal unit prices for imported coal. I fill the few data gaps by applying the growth rates from the closest non-missing data source, the IEA nominal coal price index for industry. The obtained coal price series compare well with the IEA price series in the electricity generation sector, for the few countries where both data are available.

For gas, I use Eurostat data on wholesale gas prices (excluding VAT and other recoverable taxes and levies) for the second largest consumption band of industrial consumers. This band corresponds to the average consumption of large gas-fired power plants as reported in the European Environmental Agency’s Large combustion plant database. I fill the few data gaps by imputing values from the third largest consumption band, or, if it is also missing, from the IEA gas price data. One drawback of this data source is that the consumption band categories and the methodology changed in 2007, which makes it difficult to build a consistent series of coal/gas price ratio before 2007. For this reason, I average the coal-to-gas price ratio predictor over the 2007-2012 period only (rather than the 2005-2012 period) when I apply the synthetic control method. I convert the obtained coal and gas price series to the same unit and combine them to build coal-to-gas price ratio for all European countries over the 2007-2016 period. The obtained price ratios compare well with the price ratios available from national statistical institutes for some countries¹⁹.

Lignite resources (predictor used in the main specification): Data on lignite resources in Europe come from the industry association Euracoal (European Association for Coal and Lignite; Source: <https://euracoal.eu/info/euracoal-eu-statistics/>) I create an indicator variable equal to 1 for countries with lignite resources greater than 0.5 Gt in 2012, and 0 otherwise. The variable is equal to one for Germany, Poland, Hungary, Greece, Czech Republic, and Bulgaria.

¹⁹For example, in the UK the Department for Business, Energy and Industrial Strategy publishes such data each quarter

Residual load per capita (predictor used in the main specification): Residual load is defined as the difference between electrical energy available for final consumption taken from Eurostat, and generation from renewables and nuclear power plants, also derived from Eurostat. Generation from renewables is the sum of total net electricity production from the five renewable sources hydro, tide wave and ocean, solar PV, solar thermal and wind²⁰. Generation from nuclear power plants is the sum of total net electricity production from nuclear power plants, including conventional plants, auto producers, and co-generation plants. This variable is then divided by the average population by country given by Eurostat.

Emissions from LCP opt-out plants in 2009 and IED opt-out plants in 2012 (predictor used in the main specification): The list of LCP and IED opt-out plants is available on the European Environmental Agency's website²¹. Since there is no common identifier between the EUTL and LCP and IED data, I manually matched the 172 LCP opt-out installations and the 140 opt-out installations under the LLD option located in the UK or in a country from the donor pool to the EUTL emission data (using information on the plant name and location). No match was found for two LCP opt-out plants, one from Finland and one from Poland, and for six IED opt-out plants, one from Poland, one from Slovakia, one from Czech Republic and one from the UK (but based on these installations' names, they are unlikely to be power installations except the installation from Poland). The LCP emission variable is obtained by aggregating CO₂e emissions from the LCP opt-out power plants at the country-level. The IED emission variable is obtained by aggregating CO₂e emissions from the IED opt-out power plants at the country-level.

Number of heating degree days (predictor used in the sensitivity analysis): Eurostat series "cooling and heating degree days by country - annual data".

Per capita capacity for combustible fuels, gas and coal (predictor used in the sensitivity analysis): The variable is derived from Eurostat data on electricity production capacities for combustible fuels by technology and operator. I aggregate all technologies and operators to get the total installed capacity.

Growth in per capita renewables capacity (predictor used in the sensitivity analysis): The variable is derived from Eurostat data on net electrical maximum ca-


²⁰geothermal, biomass and waste are not included since they are available on demand

²¹Source: <https://www.eea.europa.eu/data-and-maps/data/large-combustion-plants-lcp-opt-out-under-article-4-4-of-directive-2001-80-ec-4> for LCP opt-out and <https://www.eea.europa.eu/data-and-maps/data/lcp-9> for IED opt-out plants


capacity by renewable technology. I add up capacities for wind, solar, tide wave and ocean, hydro and geothermal and calculate the average annual growth rate between 2010 and 2012. 2010 is the year where the Europe 2020 strategy was adopted (including the target of increasing the share of renewable energy in final energy consumption to 20% by 2020), and 2012 is the last year before the introduction of the CPS.

Average age of operating coal-fired plants above 30 MWth (predictor used in the sensitivity analysis): The variable is derived from the Global Energy Monitor's "Global Coal Plant Tracker" (Shearer et al., 2019), a publicly available database categorizing every known coal-fired generating unit with a rated capacity above 30 MWth. I use information on the status of the unit (operating/retired/mothballed) and its commissioning date to build a country-level variable of the coal fleet's age, defined as the average capacity-weighted age of the coal-fired power plants operating every year.


A.1.6 Common Support for predictors




(a) Per capita CO₂e emissions (2005)




(b) Per capita CO₂e emissions (2012)




(c) Residual load per capita



(d) Lignite dummy



(e) LCP opt-out emissions per capita (2009)



(f) Coal-to-gas price ratio (2007-2012)

Figure A.1.4 – Histograms of main predictors. UK: black with transparent fill; donor pool: grey fill.

Notes: unless otherwise specified, all variables are averaged for the 2005-12 period

A.1.7 Counterfactual emissions in the absence of biomass conversion for Drax and Lynemouth plants

	(1)	(2)	(3)	(4)	(5)
	Total CO ₂ e	Generation coal units (MWh)	Estimated CO ₂ e coal units	Estimated CO ₂ e biomass units w/o conversion	Estimated CO ₂ e total w/o conversion
2005	20,771,624	-	-	-	20,771,624
2006	22,764,847	-	-	-	22,764,847
2007	22,160,413	-	-	-	22,160,413
2008	22,299,778	-	-	-	22,299,778
2009	19,851,702	11,584,366	10,425,929	9,425,773	19,851,702
2010	22,392,487	13,537,600	12,183,840	10,208,647	22,392,487
2011	21,465,607	15,093,899	13,584,509	7,881,098	21,465,607
2012	22,694,684	14,592,305	13,133,075	9 561 609	22,694,684
2013	20,319,513	14,398,937	12,959,044	9 434 905	22,393,948
2014	16,595,193	13,364,881	12,028,393	8 757 339	20,785,732
2015	13,192,780	13,808,137	12,427,324	9 047 783	21,475,107
2016	6,261,692	7,120,958	6,408,862	4,666,008	11,074,871
2017	6,215,220	-	6,215,220*	4,525,026	10,740,246

Table A.1.2 – Counterfactual CO₂ emissions in the absence of biomass conversion, Drax

Notes: emission data (column (1)) come from the EUTL. Generation data for the coal part over the 2009-2016 period (column (2)) come from Abrell et al. Emissions for the coal units over the 2009-2016 (column (3)) were estimated by applying an average emission factor of 0.9tCO₂/MWh, the average emission rate reported by Abrell et al for Drax plant. *Emissions for the coal units in 2017 are estimated to be the same as emissions for the entire plant as reported in column (1), since all biomass units are fully converted by then and emit zero CO₂. Emissions for the biomass units if they had not converted to biomass are estimated differently for the 2009-2012 and for the 2013-2017 period: for the 2009-2012 period, those units had not yet converted to biomass, so their emissions are simply the difference between EUTL emissions and emissions estimated for the coal units (column (1) - column (3)). For the 2013-2017 period, the units have started to convert to biomass, which is reflected in the total EUTL data. Emissions if those units had not converted to biomass are estimated assuming that they would have followed the same evolution as emissions from the coal units: they are obtained by multiplying the estimated emissions for the coal units each year (column (3)) by the ratio of emissions from the biomass units in 2012 (column (4)) over emissions from the coal units (column (3)) in 2012 (9,561,609/13,133,075=0.73)

	(1)	(2)
	Total CO ₂ e	Estimated CO ₂ e total w/o conversion
2005	2,685,512	2,685,512
2006	2,693,932	2,693,932
2007	2,695,748	2,695,748
2008	2,802,040	2,802,040
2009	2,543,017	2,543,017
2010	2,551,364	2,551,364
2011	2,612,450	2,612,450
2012	2,050,363	2,050,363
2013	2,284,177	2,284,177
2014	2,717,964	2,717,964
2015	1,287,305	1,287,305
2016	1,059	1,287,305 *
2017	2,421	1,287,305 *

Table A.1.3 – Counterfactual CO₂ emissions in the absence of biomass conversion, Lynemouth

Notes: emission data come from the EUTL. *It is assumed that in absence of biomass conversion, Lynemouth emissions would have been the same as in 2015 in 2016 and in 2017.

A.1.8 Lower bound removing emissions from plants converted to biomass


I estimate a second more conservative lower bound of the impact by applying the synthetic control method on a modified outcome variable, where emissions from UK plants having partly or fully converted to biomass are removed from the UK emissions for the entire period of analysis. This is conservative given that the largest plant converted to biomass, Drax, had half of its generating units still running with coal during the period of analysis, and we expect to see an emission decrease for these units caused by the CPS. In this second lower bound, this emission decrease will not be taken into account, nor the counterfactual emission decrease hypothesised in section 4.2 for the biomass converted units of Drax and Lynemouth. The new per capita emission variable is 15% lower for the UK after removing emissions from these plants while emissions from countries in the donor pool stay the same, so we expect the composition of the Synthetic UK to be different. With this modified outcome variable, it becomes harder to accurately build a synthetic UK using the initial set of predictors. I therefore add a third predictor of lagged outcome for the year 2010. Figure A.1.5 shows the original UK and synthetic UK emission trajectories (in black) and the UK and synthetic UK emission trajectories for the modified outcome variable (dark grey).

The new synthetic UK is made of seven countries: Italy (35.9%), Ireland (23.5%), Slovakia (21.3%), the Netherlands (15%), Finland (3%), Poland (0.4%), and Denmark

(0.9%). Country weights are shown in table A.1.4. Table A.1.5 shows that predictors' values are still closely aligned across the actual and synthetic UK, except for the residual load per capita - it makes sense that it is lower for the synthetic UK than for the UK, since UK emissions from fossil plants - which typically cover the residual load - have been made artificially lower than they truly are by removing the two plants converted to biomass.

The gap between the UK and synthetic UK per capita emissions is smaller than with the original outcome variable, which is expected. On an average year, emissions decrease by 18.1%. The total cumulative abatement is 132 million of tCO₂e. Withdrawing the upper bound estimate for the effect of the capacity market and the CfD ($\approx 4.7\text{MtCO}_2\text{e}$ in total), I obtain an abatement of around 127 MtCO₂e, which I consider the lower bound of the CPS impact. By 2017, the emission reduction caused by the CPS under this lower bound estimation represents 37% of the original synthetic UK level of emissions in 2017.

The difference between the total abatement from this lower bound and the one from section 4.2 is $164-132=32\text{ MtCO}_2\text{e}$, which is close to the hypothesised emission decrease for Drax and Lynemouth if all their units had continued to run with coal, estimated to 27MtCO₂e.



(a) Absolute per capita emissions

(b) Emission gap between treated and synthetic

Figure A.1.5 – Synthetic control method excluding emissions from plants having converted to biomass

Notes: the synthetic UK is made of the following countries for each specification: main specification: five countries: Ireland (49.2%), Slovakia (25.6%), the Netherlands (13.7%), Finland(5.8%), Czech Republic (5.7%). Lower Bound synthetic UK without emissions from plants converted to biomass: seven countries: Italy (35.9%), Ireland (23.5%), Slovakia (21.3%), the Netherlands (15%), Finland (3%), Denmark (0.9%) and Poland (0.4%).

Country	Weight	Country	Weight
Austria	0	Ireland	0.235
Belgium	0	Italy	0.359
Czech Republic	0	Netherlands	0.15
Denmark	0.009	Poland	0.004
Spain	0	Portugal	0
Finland	0.03	Sweden	0
France	0	Slovakia	0.213
Hungary	0		

Table A.1.4 – Country weights in Lower Bound Synthetic UK

Note: all weights are between 0 and 1 because the Synthetic control method imposes positive weights summing to 1.

Variable	UK	Synth. UK	Avg. Donor pool
Per capita residual load	4.29	4.27	3.37
Coal-gas price ratio	0.52	0.49	0.71
Per capita LCP opt-out emissions	0.29	0.29	0.22
Lignite dummy	0.00	0.004	0.20
Per capita emissions 2005	2.59	2.62	2.62
Per capita emissions 2010	2.24	2.26	2.39
Per capita emissions 2012	2.20	2.05	2.05

Table A.1.5 – Predictors' values for the UK, synthetic UK and average of the donor pool, lower bound

Notes: the per capita residual load is averaged for the period 2005-12, and the coal-to-gas price ratio for the period 2007-12. LCP opt-out emissions are taken in 2009, the lignite dummy is time-invariant. Outcome lags are taken in 2005, 2010 and 2012.

A.1.9 Sensitivity of the results to the choice of predictors

I test the sensitivity of the results from the upper bound to using four alternative sets of predictors to generate a synthetic UK. Note that the results would be the same if applied on the lower bound from part 4.2, since the composition of the synthetic UK is the same. In the first alternative set of predictors, I replace the per capita residual load with the annual number of heating degree days, a variable approximating the demand for energy needed for heating, likely to capture variations in peak power demand mostly covered by fossil fuels²². In the second alternative set of predictors, I add as predictor a measure of


²²In the EU, the number of heating degree days is measured as the number of days of the year where the average temperature is below a reference temperature of 15.5°C - under which energy for heating is needed - times the difference between this reference temperature and the temperature of the day. Compared with the average annual temperature used in other papers for predicting power demand, this variable better captures demand for power generation at low-temperature periods.

installed capacity for combustible fuels. This variable may influence the potential for fuel switching. Ideally, one would like to add specific variables for coal installed capacity and gas installed capacity, but these variables are not publicly available for all countries in the donor pool.


In the third alternative set of predictors, I add a predictor reflecting the growth in renewable capacity just before the introduction of the CPS. This variable can be considered as a proxy for the “business as usual” growth in renewables’ capacity, which would occur absent the CPS policy. I average this variable between 2010 and 2012. 2010 is the year of implementation of the Europe 2020 strategy, which sets a target for the share of renewables in final energy consumption to be reached by 2020 for EU countries. This announcement may be followed by a growth in renewable capacity in all European countries. 2012 is the last year of pre-treatment, which makes sure that the growth in renewable capacity is not affected by the CPS.

In the last alternative set of predictors, I add a predictor reflecting the age of the fleet of coal plants for each country. Newer coal-fired plants tend to be more efficient and produce less emissions per output of electricity, so we may expect the average age of coal power plants to influence a country’s emissions. Only plants with a capacity above 30 MWth are included in the calculation of the average age of coal plants at the country level (see appendix A.4) - but we expect these plants to be responsible for most of coal-based power generation.


Figure A.1.6 shows the common support for these four predictors. For each predictor, the UK value falls within the distribution of other countries’ values. Figure A.1.7 shows the emission paths for the different Synthetic UKs obtained with these alternative sets of predictors. The country weights are indicated below the figure for each Synthetic UK. The emission path of each new synthetic UK is relatively close to the original one. The fit is less good for the specification with the number of degree days and the one with the average age of coal-fired plants. Table A.1.6 shows the predictors’ values for the UK, for the original Synthetic UK, and for the four alternative Synthetic UKs. The alternative predictors’ values are close in the UK and in the different synthetic UKs. For the synthetic UK using the growth in renewable capacity as predictor (Alt. 3) and the synthetic UK using the average age of coal-fired plants (Alt. 4), having close values for the new predictor comes at the expense of a poorer predictor balance for some of the other predictors (lignite dummy and per capita opt-out emissions for Alt.3; lignite dummy and residual load per capita for Alt.4).




(a) Number of heating degree days



(b) Per capita capacity for combustible fuels




(c) Average annual growth in per capita renewables capacity (2010-2012)




(d) Average age of operating coal-fired power plants

Figure A.1.6 – Histograms of predictors used in the sensitivity analysis (UK: black with transparent fill; donor pool: grey fill)

Notes: unless otherwise specified, all variables are averaged for the 2005-12 period



(a) Absolute per capita emissions



(b) Emission gap between treated and synthetic

Figure A.1.7 – Sensitivity analysis: alternative set of predictors

Notes: the synthetic UK is composed in the following way for the different specifications: main specification: five countries: Ireland (49.2%), Slovakia (25.6%), the Netherlands (13.7%), Finland (5.8%), Czech Republic (5.7%). Specification with the number of heating degree days: five countries: Ireland (37.9%), Italy (31.6%), Finland (15.3%), Slovakia (12.9%) and Poland (2.4%). Specification with per capita combustion capacity: five countries: Ireland (40.8.7%), Slovakia (22.4%), the Netherlands (14.1%), Poland (6.7%) and Finland (6%). Specification with the growth in per capita renewables capacity: four countries: Italy (61.5%), Poland (20.7%), Ireland (16.1%), and Denmark (1.7%). Specification with the average age of the coal-fired power plants: six countries: Slovakia (35.2%), Czech Republic (31.2%), Sweden (12%), Hungary (11.2%), Spain (7.5%) and Finland (3%).

Variable	UK	Synth.UK				
		Original	Alt.1	Alt.2	Alt.3	Alt.4
Per cap. Residual load	4.29	4.30	X	4.44	4.27	2.01
Coal-gas price ratio	0.52	0.51	0.48	0.50	0.51	0.52
Per cap. LCP						
opt-out emissions 2009	0.29	0.24	0.29	0.29	0.36	0.31
Lignite dummy	0.00	0.06	0.024	0.07	0.21	0.42
Per cap. emissions 2005	2.98	3.13	2.98	3.12	3.06	2.94
Per cap. emissions 2012	2.59	2.43	2.32	2.45	2.49	2.45
Nb. of degree days	3020.30	X	3024.32	X	X	X
Per cap. combustible						
fuels capacity	1.10	X	X	1.15	X	X
Growth in per cap.						
renewable capacity	0.23	X	X	X	0.21	X
Avg. age of coal-fired						
power plants	36.04	X	X	X	X	36.07

Table A.1.6 – Predictors' values for the UK and each alternative synthetic UK

Notes: the per capita residual load, number of degree-days, per capita combustible fuels capacity and the average age of coal-fired power plants variables are averaged for the period 2005-2012. The coal-to-gas price ratio variable is averaged for the period 2007-2012. The growth in per capita renewable capacity is averaged for the period 2010-2012. LCP per capita opt-out emissions are taken in 2009, the lignite dummy is time-invariant. The outcome lags (per capita power sector emissions) are taken in 2005 and 2012.

A.1.10 Sensitivity of the results to the choice of the donor pool

Figure A.1.8 shows the sensitivity of the results from the upper bound to a different composition of the donor pool. Note that the results would be the same if applied on the lower bound from part 4.2, since the composition of the synthetic UK is the same. In a first test, I include Greece and Germany back in the donor pool, two countries that were previously excluded because they had experienced a shock in their power sector at the period of interest. In a second test, I include Germany, Greece, Latvia and Lithuania, all the countries that were previously excluded except Estonia (Adding Estonia makes it impossible to find a convex combination of countries replicating the trajectory of the UK, probably due to the too large discrepancy in emissions between Estonia and the other countries). In a third test, I exclude Denmark and Finland from the original donor pool. These two countries may be influencing the results substantially since they have large variations in per capita emissions and have a non-zero weight in the initial synthetic UK. The composition of the synthetic UK barely changes with different compositions of the donor pool. Figure A.1.8b shows that the emission reduction estimate obtained with the main specification is close to the estimates obtained with different compositions of the donor pool.

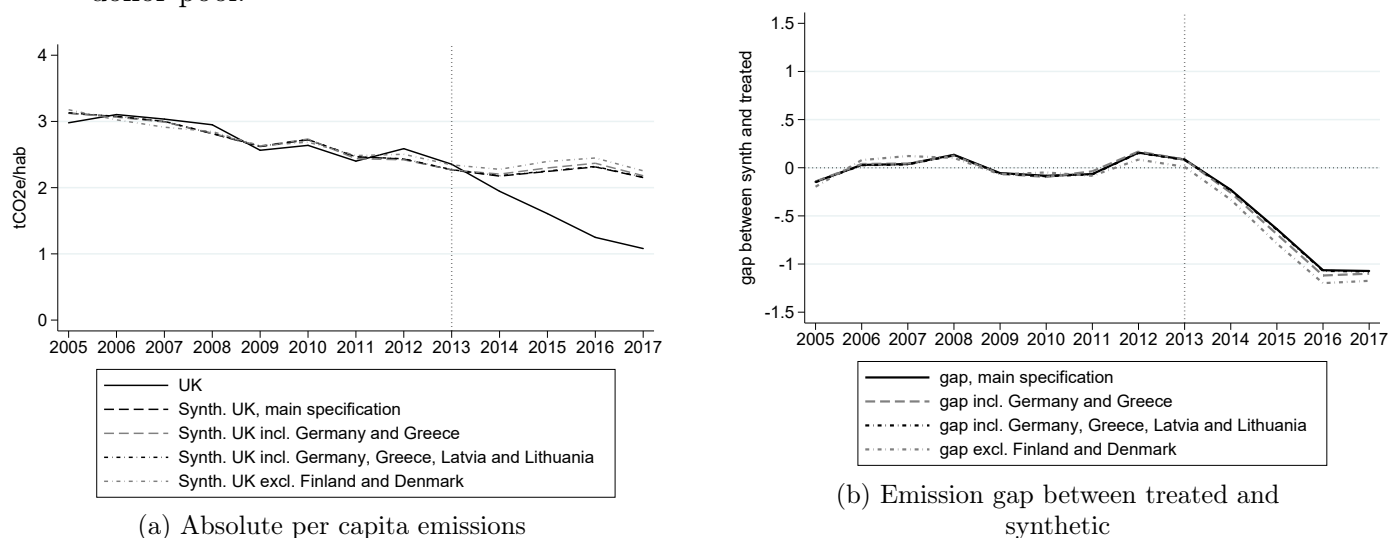


Figure A.1.8 – Sensitivity analysis: alternative donor pool

Notes: the synthetic UK is composed in the following way for the different specifications: main specification: five countries: Ireland (49.2%), Slovakia (25.6%), the Netherlands (13.7%), Finland (5.8%), Czech Republic (5.7%). Specification including Greece and Germany in the donor pool: four countries: Ireland (54.2%), the Netherlands (22.5%), Slovakia (16.6%), and Finland (6.6%). Specification including Greece, Germany, Latvia, Lithuania (entire donor pool except Estonia): five countries: Ireland (49.4%), Slovakia (25.2%), the Netherlands (14.2%), Finland (5.8%), Czech Republic (5.5%). Specification excluding Finland and Denmark: four countries: Ireland (51%), Slovakia (20.6%), the Netherlands (18.1%), Poland (10.3%).

A.1.11 Sensitivity of the results to a longer pre-treatment period

The emission data used in the main analysis are only available from 2005, which means that only eight years of pre-treatment data can be used to generate the synthetic UK and assess the validity of the method. Although there is no rule of the thumb for the minimum number of pre-treatment periods that can be deemed "safe" to apply the synthetic control method, Abadie et al mention that "the applicability of the method requires a sizeable number of pre-intervention periods" (Abadie et al., 2015). I test whether having a relatively short time period is likely to bias my estimate by applying the synthetic control method to a similar outcome variable available for twenty-three pre-treatment years (but which does not have information at the plant level). The new outcome variable is directly available at the country level on the European Environmental Agency (EEA)'s website, and includes all greenhouse gas emissions from the public electricity and heat production sector. I divide the variable by annual country population to obtain per capita emissions, and I use it over the 1990-2017 period. This aggregate variable does not allow to identify individual plants and isolate confounding factors like I do in the main analysis. But it would be reassuring to find an emission reduction close to the main result when I apply the synthetic control method to this longer dataset. I keep the same predictors as in the main specification.

Figure A.1.9 shows the results. The composition of the synthetic UK changes compared to the main result, with the new synthetic UK made of three countries: Poland (36.0%), Italy (44.2%), and Slovakia (19.7%). It means that the combination of countries best mimicking the evolution of UK emissions over the 1990-2012 period is not the same as the combination of countries best mimicking the evolution of UK emissions over the 2005-2012 period. The estimated cumulative abatement with this longer period is the same as the estimate for the upper bound using a shorter pre-treatment period: 191 MtCO_{2e} over the 2013-2017 period. While predictors' values were closely aligned between the UK and synthetic UK in the main result, this is no longer the case with the new synthetic UK: table A.1.7 shows that the predictors' values of the synthetic UK are further away from the UK than that of the average donor pool. This may indicate that averaging the predictors for the end of the pre-treatment period only is not appropriate to predict emission values from before 2005²³. I run a permutation test similar to the one performed in section 4.3. Figure A.1.10 shows that the decrease in emissions seen in the UK is not found with the

²³The other way around, keeping the original weighting and synthetic UK obtained in the main specification yields an emission trajectory which does not mimic well the emission trajectory for the UK before 2005. Applying the synthetic control method before 2005 could be inherently difficult in the case of power sector emissions compared to other sectors such as transport, due to the important change in European electricity markets over the 1990s and 2000s from heavily regulated industries to more liberalized and interconnected markets subject to a single carbon market.

same magnitude in other European countries. This test suggests that using a relatively short time period in my main results should not come at the cost of a large bias in the estimation

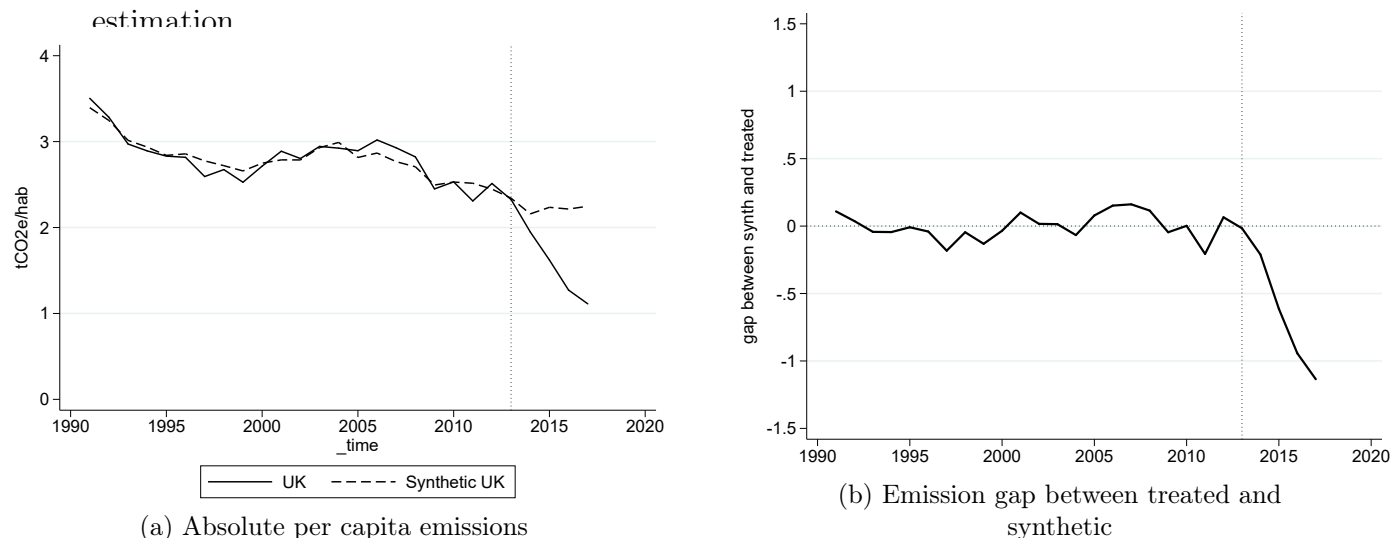



Figure A.1.9 – Sensitivity analysis: extended pre-treatment period with Eurostat greenhouse gas emissions by sector

Notes: the synthetic UK is made of three countries: Poland (36.1%), Italy (44.2%), Slovakia (19.7%).


Variable	UK	Synth. UK	Avg. Donor pool
Per capita residual load	4.29	3.34	3.37
Coal-gas price ratio	0.52	0.44	0.71
Per cap. LCP opt-out emissions 2009	0.29	0.64	0.22
Lignite dummy	0	0.36	0.20
Per cap. emissions 1990	3.57	3.57	2.70
Per cap. emissions 1998	2.67	2.72	2.77
Per cap. emissions 2012	2.51	2.45	2.14

Table A.1.7 – Predictors' values for the UK, synthetic UK and average of the donor pool, longer panel dataset

Notes: the per capita residual load is averaged for the period 2005-12, and the coal-to-gas price ratio for the period 2007-12. Per capita opt-out emissions are taken in 2009, the lignite dummy is time-invariant. Outcome lags are taken in 1990, 1998 and 2012.



(a) Permutation test with longer pre-treatment period




(b) Ratio of MSPEs with longer pre-treatment period

Figure A.1.10 – Sensitivity analysis: extended pre-treatment period with Eurostat greenhouse gas emissions by sector, permutation test


Notes: in both figures, France is not included: for this country it is impossible to find a stable diagonal V matrix

A.1.12 Estimation of the spillovers removing all interconnected countries from the donor pool

I run the synthetic control method again, after removing the countries interconnected to Ireland from the donor pool. When I do so, the pre-treatment fit becomes quite poor due to the high weight taken by Finland. I thus also run a specification where I also remove Finland from the donor pool. Figure A.1.11 shows the trajectory of these two alternative Synthetic UK compared to the trajectory of UK emissions. For the synthetic UK without the interconnected countries nor Finland, the gap in emissions is smaller and results in a cumulative abatement 22 MtCO_{2e} lower, of 142 MtCO_{2e}, 14% less than the lower bound estimate. Predictor balance is however less good, suggesting a trade-off, also pointed by Abadie (2021), between keeping in the donor pool countries sufficiently close to the treated unit, and having countries “too close” geographically and hence subject to spillovers from the treated unit. In particular, the amount of LCP opt-out emissions is not anymore aligned with the UK value and is greater in this new synthetic UK than in the original synthetic UK, which may partly explain why the emission gap is smaller than in the main specification.



(a) UK and synthetic UK



(b) gap between UK and synthetic UK

Figure A.1.11 – UK (incl. counterfactual emissions from biomass converted plants) and synthetic UK - no interconnected countries

Notes: the initial synthetic UK is made of five countries: Ireland (49.2%), Slovakia (25.6%), the Netherlands (13.7%), Finland(5.8%), Czech Republic (5.7%). The synthetic UK with Ireland, the Netherlands and France removed from the donor pool is made of five countries: Spain (58.9%), Finland (25.4%), Slovakia (7.6%), Czech Republic (6.5%) and Denmark (1.7%). The synthetic UK with Ireland, the Netherlands, France and Finland removed from the donor pool is made of four countries: Italy (72%), Poland (23.4%), Denmark (2.5%), and Czech republic (2.1%). UK emission values include estimated counterfactual emissions in the absence of biomass conversion for Lynemouth and Drax plants.

Appendix B

Appendix to Chapter 2: Estimating the Causal Effects of Cruise Traffic on Air Pollution using Randomization-Based Inference

B.1 Appendix

B.1.1 Reproducibility

Replication Materials

We strive to make our analysis reproducible. Our data are stored on the Open Science Framework at the following address [\[\]](#): A README file explains how to reproduce the statistical analysis. The folder is organized into two main sub-folders: one for the hourly analysis and one for the daily analysis. Each sub-folder contains the raw data and the **R** scripts that can be used to clean the data, run the matching procedure, and explore the results. The **R** Markdown scripts are compiled as HTML files, which makes it possible to see both the code and its outputs.

Caveats

The full replication of our study presents two important caveats. First, we were not allowed to share the weather data from Météo-France. We therefore added a small amount of noise to the original data. Researchers who reproduce our analysis will get different results. They can however easily check our coding procedure. Second, the matching procedure at the *hourly* level is computationally demanding. We had to rent an Amazon Web Services virtual computer (EC2 t3.2xlarge) to run the matching algorithm.

B.1.2 Data Sources

Our work combines data from several sources that can be accessed relatively easily:

- Marseille air pollution data can be downloaded from the website of the local air quality agency, AtmoSud. The address is <https://www.atmosud.org/donnees/telecharger>. The website is in French but a contact form is available to ask for help.
- Weather data can be downloaded from the website of the national meteorological agency Météo-France. To access the download service, a free research license must be created first, at <https://donneespubliques.meteofrance.fr/?fond=contact>.
- Vessel traffic data can be requested from the port authority of Marseille-Fos (GPMM). The contact details of the officers in charge of the data can be provided upon request.
- Road traffic were retrieved from the DIRMED/SPEP, which is the regional administration in charge of roads. The contact detail is spep.dirmed@developpement-durable.gouv.fr.

B.1.3 Data Wrangling

The hourly and daily scripts from the replication package entitled `script_data_wrangling.html` explain how to clean and merge the different datasets. We summarize the main steps below, which are similar for the hourly (N=96,432) and daily (N=4,018) analyses.

Formatting Cruise Traffic Data

The raw data contains information on each of the 41,015 port calls occurring between 2008 and 2018. The raw data contains information on each vessel's type, origin, destination, country flag, deadweight tonnage, gross tonnage and construction year. We only use two variables in the analysis, the gross tonnage and the vessel type. Out of the 41,015 port calls, 171 appear to have unrealistic gross tonnage values, which we set as missing. We use the chained random forest algorithm from the `missRanger` **R** package to impute missing values (Wright and Ziegler, 2017; Mayer, 2019). A simulation exercise available in the file `script_data_wrangling.html` of the replication folder shows that the imputation algorithm works well. For each hour or day, we then compute the total gross tonnage of cruise, ferry and other types of vessels coming to or leaving the port.

Merging Cruise Traffic Data with Air Pollution and Weather Data

We merge vessel traffic data with air pollution and weather data. All pollutants have missing readings but in different proportions. For instance, 3% of NO₂ concentrations and 10% of PM_{2.5} concentrations measured at Longchamp station are missing. Weather parameters have at most 1% of missing values.

Merging the Data with Road Traffic Data

The road traffic data are only available for the 2011-2016 period. The raw data contains information on the hourly traffic flow measured at 139 different monitoring stations scattered across Marseille main roads. For most stations, many readings are missing or have a zero value, which is likely to be wrong. After a visual inspection of the time series data from every traffic monitoring station, we discard observations from the monitoring stations with large breaks in the time series. We calculate the average traffic flow across the remaining six stations to obtain a city-level measure of the road traffic intensity. Four of them are located along the A50 highway, a West-East highway coming from the East of Marseille, with two measuring the traffic flow towards Marseille city centre and two the traffic flow in the opposite direction, towards the East of Marseille. The two other are located along the A7 highway, a North-South highway coming from the North of Marseille, with one measuring the traffic flow towards Marseille city center and one measuring the flow towards the North of Marseille.

Adding Calendar Indicators and Imputing Missing Values

We add calendar indicators for the day of the week, month, year, bank days, and holidays. Finally, we impute missing values for weather and air pollutant concentrations using the `missRanger` **R** package (Wright and Ziegler, 2017; Mayer, 2019). To gauge the performance of the imputation algorithm, we carry out a simulation exercise which reproduces the observed patterns of missing pollutant concentrations. Air pollution measuring stations often break down over several days. We retrieve the subset of data where all values are recorded and we erase values for observations that belong to four randomly sampled weeks. The performance of the algorithm was fair given the difficulty to guess missing values without observed lags of their values. For instance, at the daily level, the absolute average difference between real and imputed values is $5.1 \mu\text{g}/\text{m}^3$ for NO_2 concentrations measured at Longchamp station - the average daily NO_2 is $30 \mu\text{g}/\text{m}^3$. Additional simulation results are available in the hourly and daily `script_data_wrangling.html` files. Note that we make sure to check that incorrect imputations are not driving our results. In the section *Sensitivity Analyses* of this document, we remove units with missing values from the matched pairs and find similar results as with the imputation.

B.1.4 Exploratory Data Analysis

We carry out an extensive exploratory data analysis of our variables to better understand their patterns (Tukey et al., 1977; Cleveland, 1993; Healy, 2018). We identify the variables that are correlated with both cruise vessel traffic and air pollution, and that need to be

adjusted for in the matching procedure. This analysis also reveals how regular cruise traffic is, making it difficult to find similar hours or days with and without high traffic.

Vessel Traffic

Vessel traffic has long-term and seasonal patterns in Marseille, as shown in Panel A of Figure B.1.1. In the long term, traffic has increased over time. For a given year, traffic peaks in the summer, driven by an increase in cruise and ferry vessel traffic. At the hourly level, there are two peaks, one in the morning where vessels enter the port, and one in the evening when they leave (Panel B of Figure B.1.1). These patterns are important to take into account in the matching procedure as calendar variables are associated with changes in both cruise vessel traffic and air pollutant concentrations.

Air Pollutants

The exploratory data analysis of air pollutant concentrations reveals expected seasonal patterns as shown in Figure B.1.3. First, for most pollutants, concentrations have decreased over time between 2008 and 2018 (Panel A of Figure B.1.3). Second, air pollutant concentrations present a weekly pattern, with higher concentrations on weekdays compared to weekends (Panel B of Figure B.1.3). Third, air pollutant concentrations have an hourly pattern. They increase during the day and decrease in the evening (Panel C of Figure Figure B.1.3). Given these patterns, we need to include constraints on the year, season, day of the week and the hour of the day in the matching procedure.

Weather Parameters

Weather parameters also have monthly and hourly patterns, as displayed on Figure B.1.4 for hourly patterns. Wind patterns are particularly important because they can chase pollutants away from the city, or, on the contrary, bring them over the city. The polar plot of Figure B.1.5 shows that the wind mostly comes from the North-West and South-West in Marseille, and only rarely from the East.

In Figure B.1.6, we plot the daily predicted concentrations of each pollutant according to wind speed and wind direction using the **R** `openair` package (Carslaw and Ropkins, 2012). Concentrations are higher on days with a low wind speed. SO_2 is higher when the wind blows from the South-West, which could indicate a role for maritime traffic, since the port is located South-West from Longchamp station. It could also indicate pollution transport from the industrial zone of Fos-sur-Mer, located in the South-West of Marseille.

As weather parameters affect air pollutant concentrations, we take the average temperature, humidity, occurrence of rainfall, and wind speed and wind direction into account

in the matching procedure.

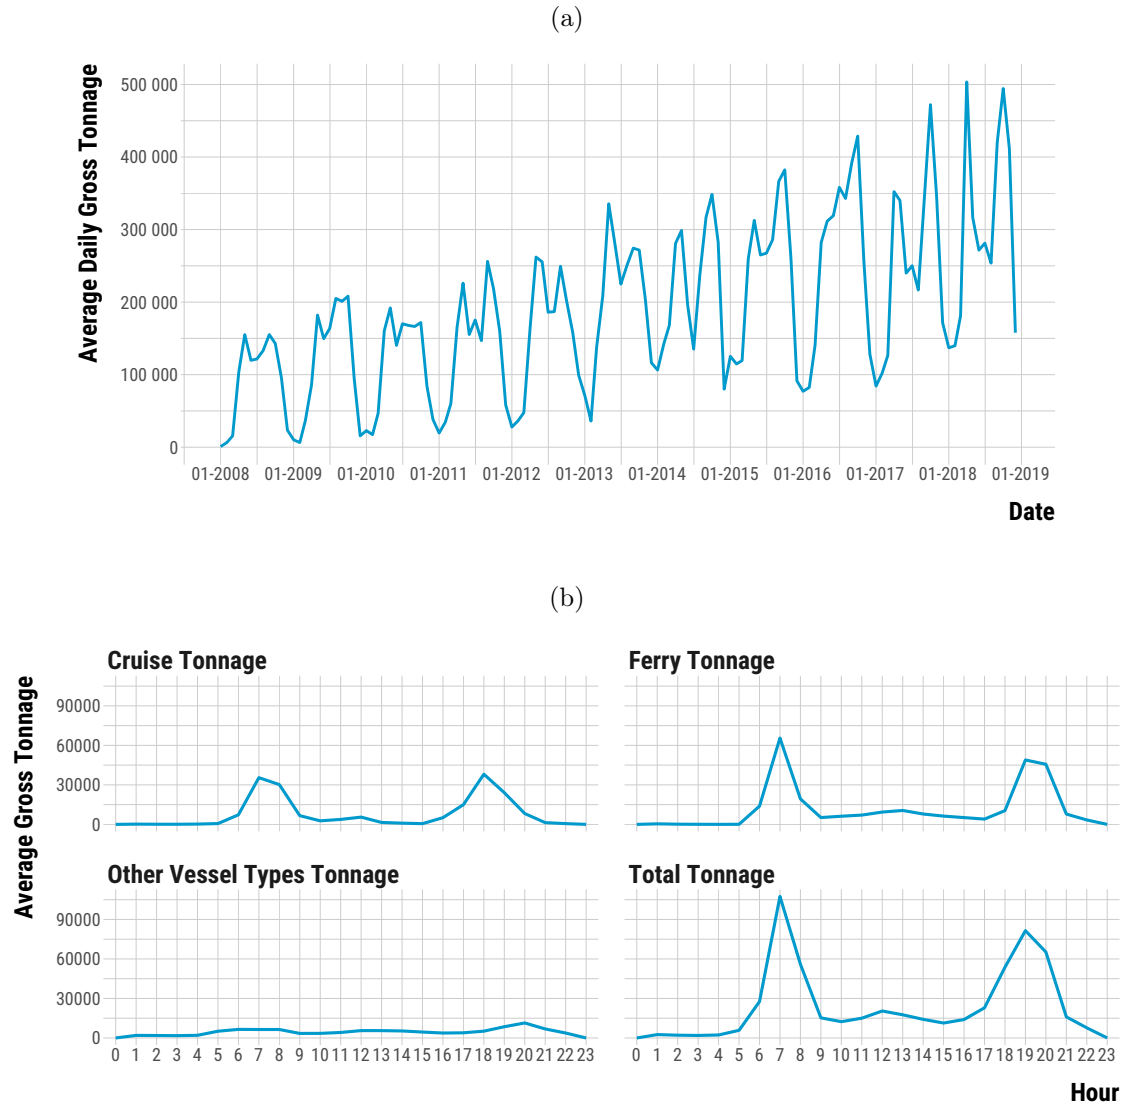


Figure B.1.1 – Seasonal and Hourly Variation in Vessel Traffic.

Notes: Panel A shows the average daily total gross tonnage of all vessel types by year-month over the 2008-2018 period. Panel B shows the average gross tonnage for each vessel type and the total across vessel types by hour of the day.

The exploratory analysis of cruise traffic data also shows how regular the traffic is. Figure B.1.2 displays the hourly gross tonnage of cruise ships entering the port each Monday of July and August in 2012. The variation in traffic is low, and hours with positive vessel traffic rarely have zero traffic on other Mondays. This regularity of the traffic partly explains why the matching procedure results in few matched pairs.


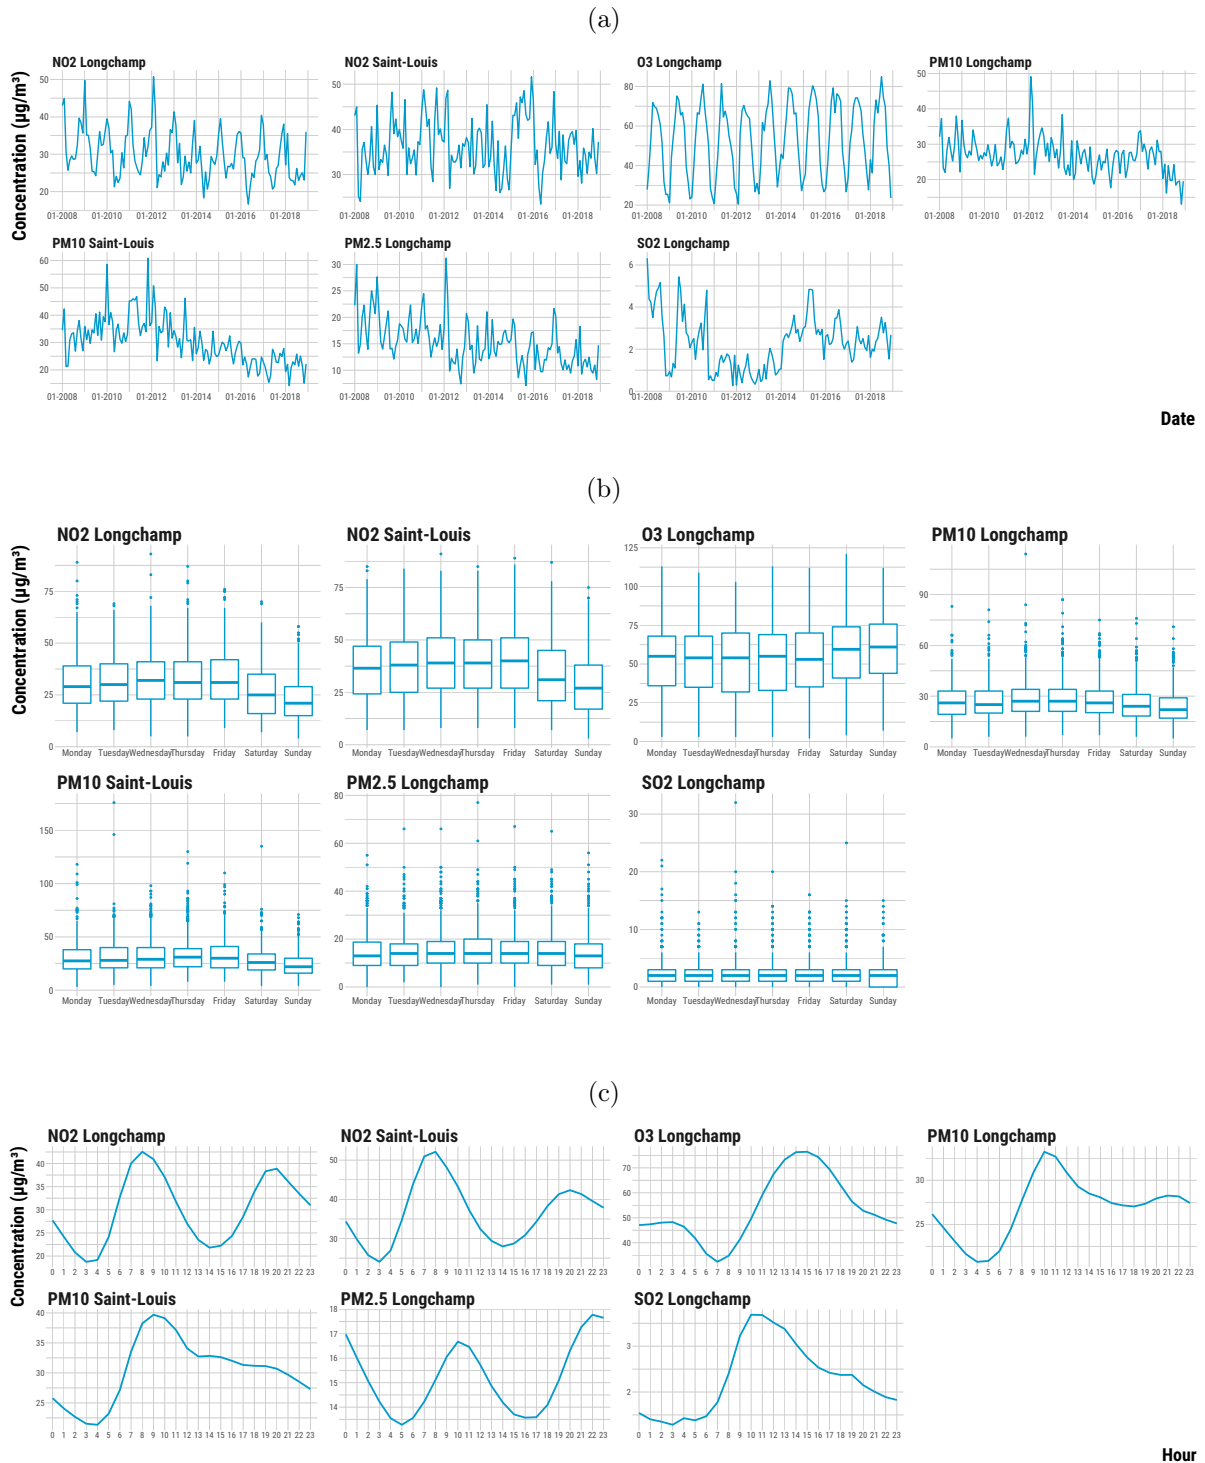


Figure B.1.2 – Hourly Traffic of Cruise Vessels Entering the Port on Mondays in July and August 2012.

Notes: Each panel represent a Monday of July or August 2012, indexed by its Julian date.



Notes: Panel A shows the evolution of monthly average concentrations over time. Panel B shows the boxplot of concentrations across days of the week, and Panel C shows the average concentrations by hour of the day.




Figure B.1.4 – Hourly Variations in Weather Parameters.

Notes: Each panel plots the hourly time series of the average value of a continuous weather parameter.




Figure B.1.5 – Polar Plot of Wind Directions.

Notes: Each blue line represents the average daily proportion of observations belonging to a particular wind direction. Wind direction are depicted on a 360° rose and proportions are inner circles whose scale is represented by the vertical numbers on the left of the main circle.

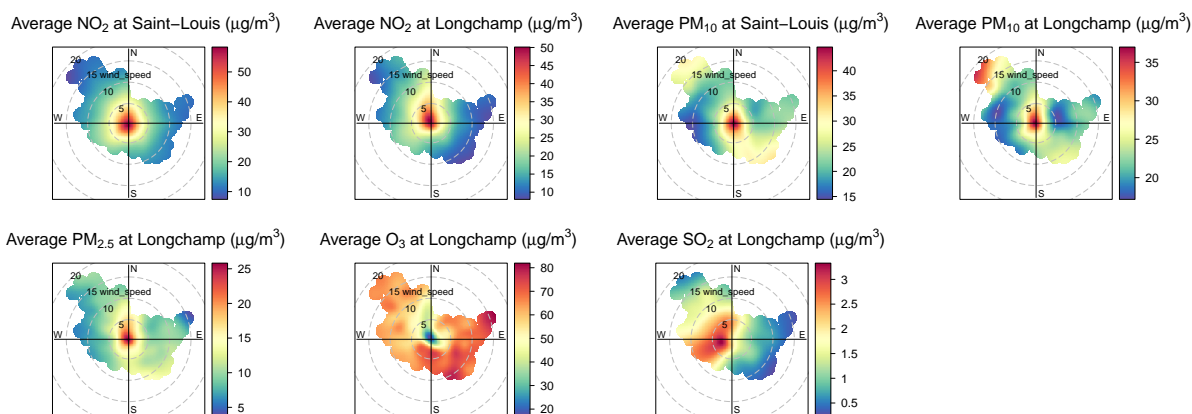


Figure B.1.6 – Polar Plot of Pollutant Concentrations according to Wind Speed and Direction.

Notes: A higher distance from the center of a circle indicates a higher wind speed value. Wind directions indicate the directions from which the winds are blowing. The gradient of colors indicate the concentrations of a pollutant that were predicted using the wind components (Carslaw and Ropkins, 2012).

Given the specific wind patterns on the coastline, hourly wind direction and wind speed measured at the weather station near Marseille airport (off the coast) may not reflect wind direction and wind speed in Marseille city. In particular, Marseille is characterised by a breeze regime, where sea breeze blows from the sea to the city in the evening, and land breeze blows from the city to the sea in the morning. To account for this difference in wind patterns, the air quality agency rolled out weather sensors measuring wind speed and wind direction at Longchamp station in 2017. We obtained hourly data for the September 2017-December 2020 period. Below we compare wind direction and wind speed for the period common to the two datasets, September 2017-December 2018. Figure B.1.7 shows the difference between average daily wind speed measured at Marseille airport (main weather data source) compared to that measured at Longchamp. Wind speed is on average 3.6 m/s higher at the airport than that measured at Longchamp. Figure B.1.8 shows the polar plots of hourly wind direction at Longchamp and the Airport station. Compared to the airport data, wind comes more often from the North-East and from the South-West at Longchamp. This likely reflects land breeze patterns in the first case and sea breeze patterns in the second case. This difference in weather patterns means that we will have a measurement error on the hourly wind values. To understand how much it may impact the accuracy of our matching, in section B.1.5 we test whether units from 2017 or 2018 which were matched together based on the main wind variables indeed have the same wind patterns based on Longchamp wind data. This issue is only present at the hourly level, as breeze is a temporary phenomenon that tends not to affect daily dominant winds.





Figure B.1.7 – Difference in hourly wind speed between the wind measured at the airport and at Longchamp station.

(a) Longchamp station



(b) Airport weather station




Figure B.1.8 – Polar Plots of Wind Directions at Longchamp and at the airport weather stations.

Notes: Panel A shows the evolution of monthly average concentrations over time. Panel B shows the boxplot of concentrations across days of the week, and Panel C shows the average concentrations by hour of the day.

Road Traffic

Finally, our exploratory analysis of road traffic confirms that road traffic presents daily and hourly patterns, as shown on Figure B.1.9. In our matching procedure, we do not adjust for traffic flow for the hourly experiments, as forcing units to belong to the same hour and day of the week is enough. We show that road traffic is balanced after matching. We do adjust for the daily average of hourly flow in the daily experiment.

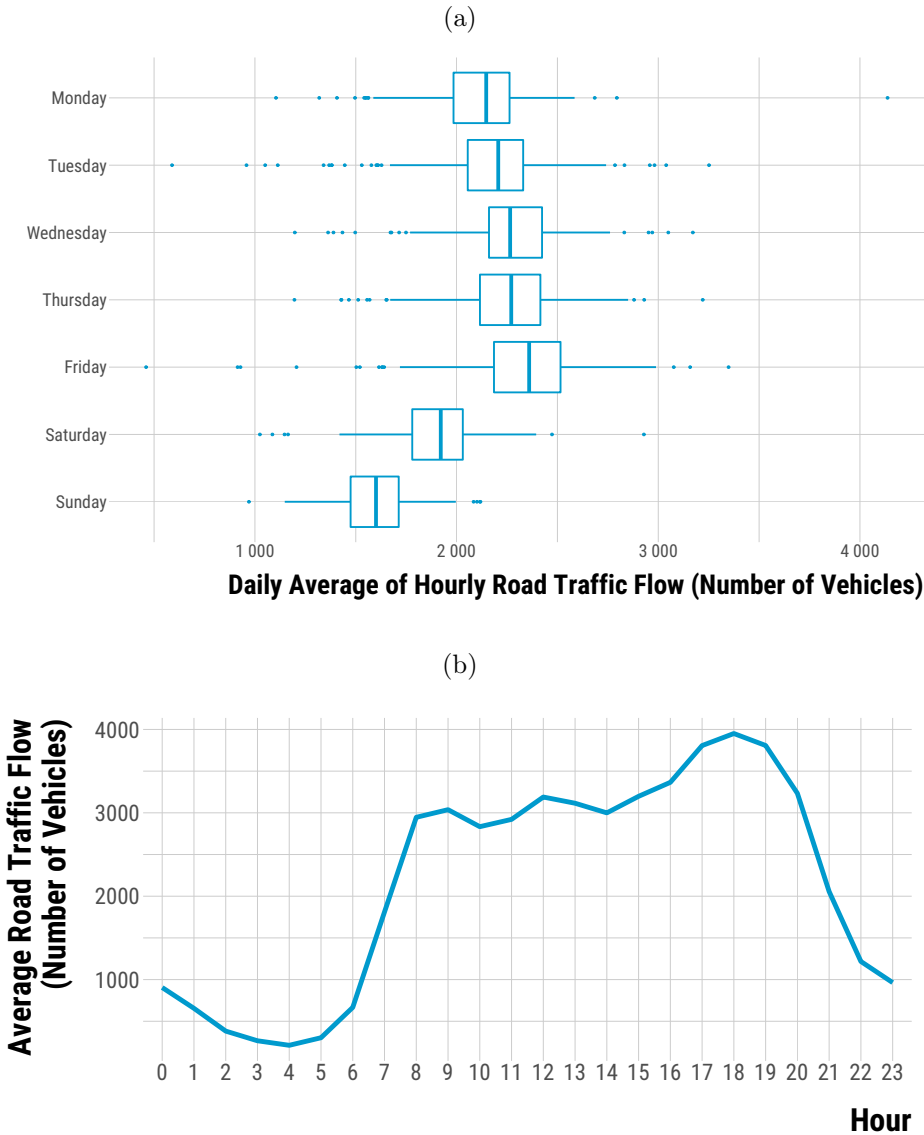


Figure B.1.9 – Daily and Hourly Variations in Hourly Road Traffic Flow.

Notes: Panel A shows the box plots of the daily average of hourly road traffic flow by day of the week. Panel B shows the average hourly traffic flow by hour of the day.

B.1.5 Matching Procedure Diagnostics

In this section, we assess the performance of the matching procedure for each of the three experiments (two hourly and one daily). First, we display the shock in gross tonnage implied by the treatment, that is, the average difference in gross tonnage between treated and control units for each vessel type. We make sure that if we define a treatment for cruise arrivals, the difference in gross tonnage is close to zero for other vessel types and flows. Second, we check whether covariates are balanced across treated and control units in the matched dataset. Third, we assess how balance improves in the matched dataset compared to the initial dataset made of potential treated and control units. Fourth, we compare the distributions of covariates in the matched data with those of the initial data.

Hourly Experiment on Cruise Vessels' Arrivals

Checking the Treatment Panel A of Figure B.1.10 shows the average hourly gross tonnage for treated and control units from 3 hours before the treatment occurs up to 3 hours after. On average, the traffic of cruise docking at the port is similar before and after the treatment takes place. Panel B plots the difference in average gross tonnage between treated and control units for each vessel type and flow. There is a strong increase in cruise vessel arrivals for treated units, but neither for departures nor for other vessel types.

Evidence of Covariates Balance for the Matched Pairs To assess covariates balance, we compare the density distributions and box plots for continuous covariates, and the proportions of each category for categorical variables, across treated and control units. In Figure B.1.11, we can see that continuous weather covariates are balanced across treated and control units from hours in t , $t-1$ and $t-2$. In Panel A of Figure B.1.12, we can see the distribution of wind direction is approximately similar for treated and control units. We match only on East and West direction, thus there remain small difference for the wind direction variable divided in four categories. For rainfall, the balance is perfect. In Figure B.1.13, we show that road traffic is balanced across treated and control units. In Figure B.1.14, we notice that pollutant distributions are very similar in hours $t-1$ and $t-2$. Finally, calendar variables such as the the hour of the day, bank days, holidays and year are balanced by design as we strictly match on these variables. Besides, the seasonality of variables is respected as a treated unit can be matched to a control with a maximum difference of 30 days.

Improvement of Covariates Balance after Matching We use love plots to show how covariate balance improves in the matched dataset compared to the initial dataset.

For continuous variables, we calculate the standardized mean difference between treated and control units. For categorical variable, we calculate the absolute difference in proportions. In Figure B.1.15, we can see that the balance of continuous weather covariates has improved overall after the matching. In Figures B.1.16–B.1.17, we note that the balance of wind direction categories and the rainfall dummy is relatively similar before and after matching. In Figure B.1.18, the balance has improved for most pollutants and lags. In Figure B.1.19, we see that the balance of gross tonnage for each vessel type and flow has generally improved over the 3 previous hours up to the 3 following hours. In Figure B.1.20, the balance of the road traffic variable have drastically improved after matching. Finally, we summarize in Figure B.1.21 the global improvement in covariates balance for continuous and categorical variables (for which we include calendar indicators). The matching procedure improves balance for continuous covariates but slightly increases the imbalance for categorical variables.

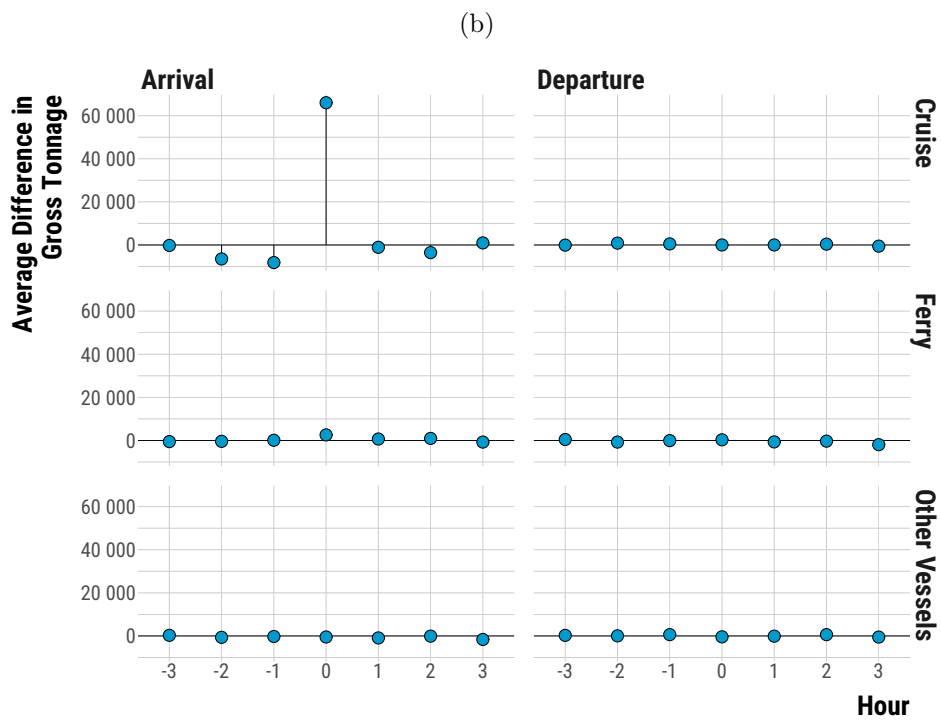



Figure B.1.10 – Visualizing the Occurrence of the Treatment.

Notes: Panel A shows the average gross tonnage for treated and control units from 3 hours before treatment up to 3 hours after. Panel B shows the average difference in gross tonnage between treated and control units by vessel type and flow.

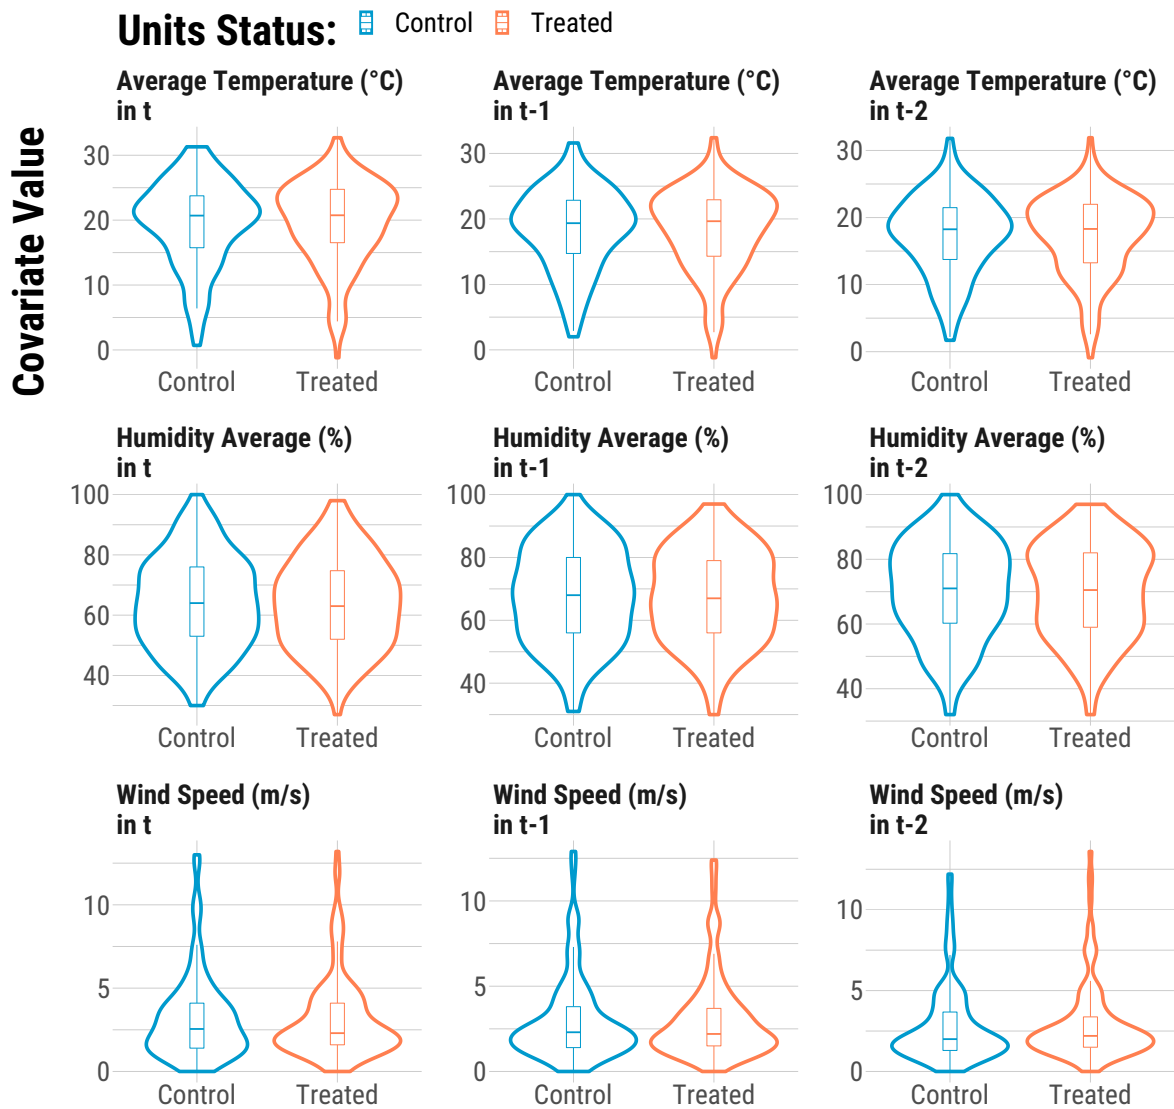


Figure B.1.11 – Continuous Weather Covariates Balance.

Notes: Each panel displays the density distribution and boxplot of weather variables by treatment status.




Figure B.1.12 – Categorical Weather Covariates Balance.

Notes: Panel A shows the proportion of distribution of wind directions by treatment status. Panel B displays the proportion of hours when rainfall occurs. In the hourly experiment of cruise arrivals, it only rained in the previous hour and for just for one percent of treated and control units.




Figure B.1.13 – Balance of Road traffic variables.

Notes: Each panel displays the density distribution and boxplot of road traffic variables by treatment status.

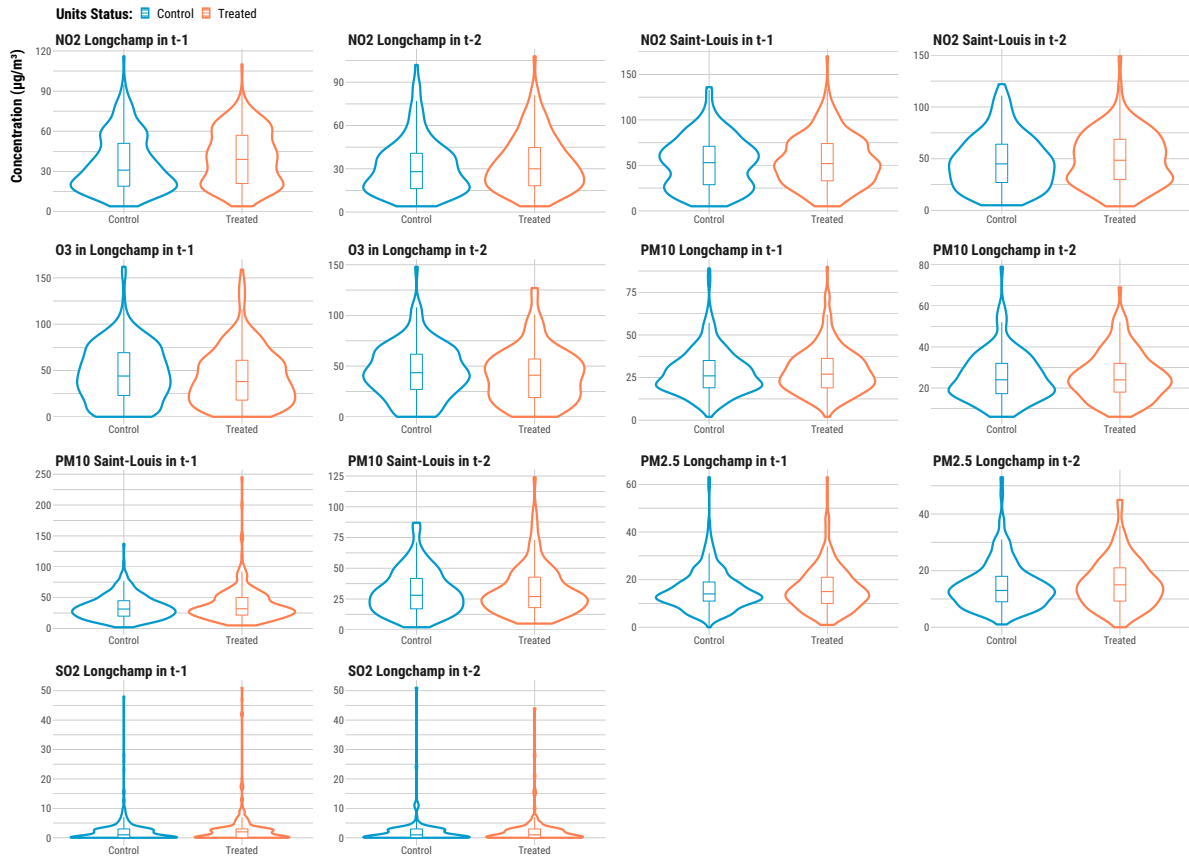


Figure B.1.14 – Balance of Air Pollutants.

Notes: Each panel displays the density distribution and boxplot of pollutants by treatment status.




Figure B.1.15 – Love Plots for Continuous Weather Covariates.

Notes: The standardized mean differences between treated and control units are plotted, before and after matching.




Figure B.1.16 – Love Plots for Wind Direction.

Notes: The absolute mean differences in proportions between treated and control units are plotted, before and after matching.




Figure B.1.17 – Love Plot for Rainfall.

Notes: The absolute mean differences in proportions between treated and control units are plotted, before and after matching. If a point is missing, it means that it was not raining for any units on that hour.




Figure B.1.18 – Love Plots for Air Pollutants.

Notes: The standardized mean differences between treated and control units are plotted, before and after matching.




Figure B.1.19 – Love Plots for Vessel Traffic.

Notes: The absolute standardized mean differences in gross tonnage between treated and control units are plotted, before and after matching.

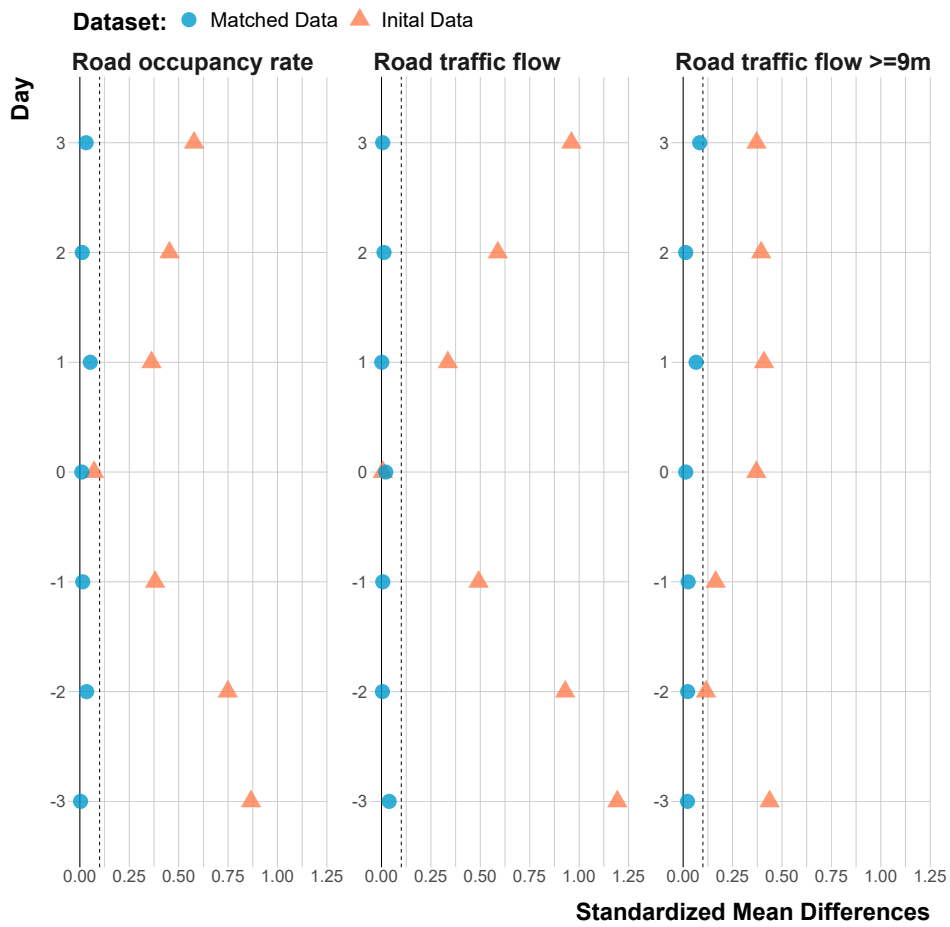


Figure B.1.20 – Love Plots for Road Traffic.

Notes: The absolute standardized mean differences in road traffic between treated and control units are plotted, before and after matching.


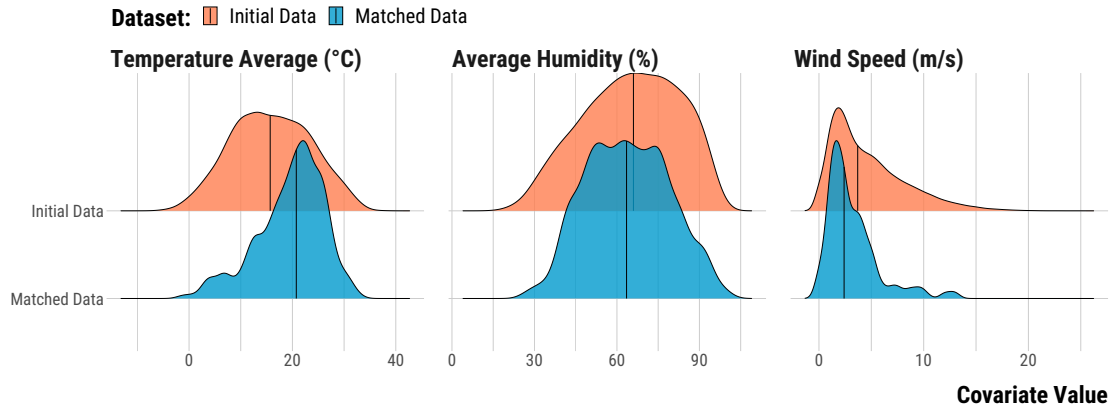


Figure B.1.21 – Overall Improvement in Covariate Balance.

Notes: Panel A shows the boxplot of standardized mean differences for continuous covariates before and after matching. Panel B shows the boxplot of absolute differences in percentage points for categorical weather covariates and calendar indicators before and after matching.

Comparing Initial Data to Matched Data In Figure B.1.22, we can see that the matched hours have a higher average temperature, a lower wind speed and a lower occurrence of rain. In Figure B.1.23, we notice that matched hours are mostly around 7 am and in the spring and summer seasons, with Tuesdays and the year 2010 over-represented compared to the initial data.

A



B




Figure B.1.22 – Comparing Weather Covariates Distribution for Initial and Matched Datasets.

Notes: Panel A plots the density distributions of continuous weather covariates for the initial and matched datasets. Panel B displays the proportion of observations belonging to a particular category of discrete weather covariates.

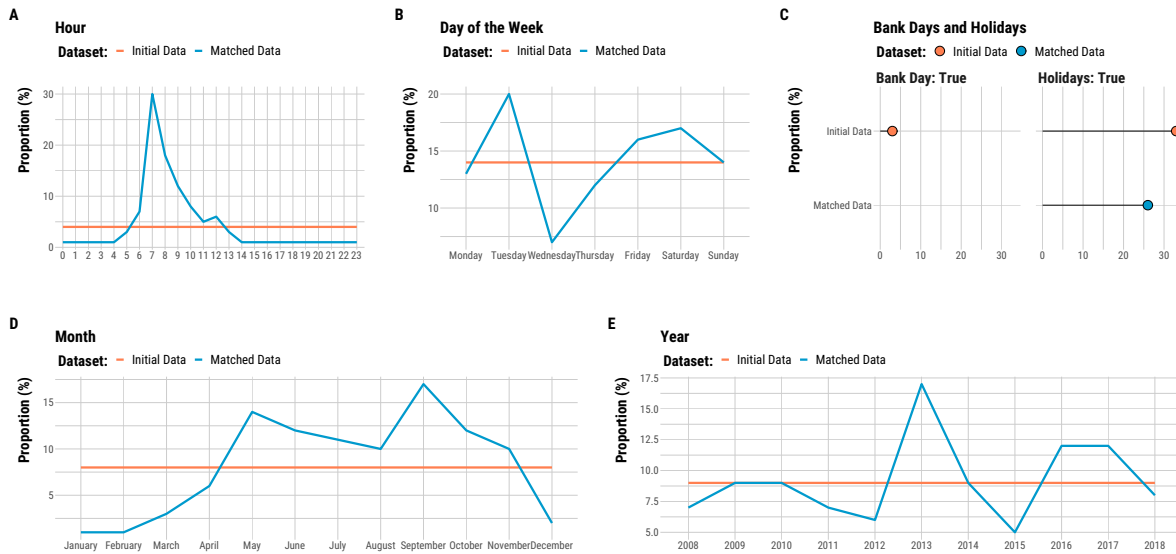


Figure B.1.23 – Comparing Calendar Covariates Distribution for Initial and Matched Datasets.

Notes: Proportions of observations belonging to each hour of the day (Panel A), day of the week (Panel B), bank days and holidays (Panel C), month (Panel D) and year (Panel E) for the initial and matched datasets.

Wind patterns across treated and control units based on Longchamp wind data

We restrict the matched sample to the 2017-2018 period to evaluate the balance in wind patterns based on Longchamp wind data (which are only available in 2017 and 2018). There are 18 pairs for which we have both wind data from the airport and at Longchamp station. All units have been matched on a binary wind direction indicator (East/West) based on wind direction measured at the airport. However, for 4 pairs, treated and control units do not have the same wind direction East/West based on the variable measured at Longchamp. If the measurement error is the same for the whole sample, $4/18=22\%$ of our matched pairs may actually have different wind direction in Marseille centre. For 3 other pairs, treated and control have the same wind direction based on Longchamp data, but the direction is opposite to what is measured at the airport. This is not an issue as long as we do not investigate treatment heterogeneity by wind direction. For wind speed, there is not much variation in wind speed as measured at Longchamp station, such that the pair difference in wind speed measured as Longchamp is always below the matching threshold.

Hourly Experiment on Cruise Vessels' Departures

See Section B.1.5 for more details on the graphs in each subsection.

Checking the Treatment On average, the traffic of cruise leaving the port is similar before and after the treatment takes place (Panel A of Figure B.1.24). There is a strong increase in cruise vessel departures for treated units, but neither for arrivals nor for other vessel types (Panel B).

Evidence of Covariates Balance for the Matched Pairs Similar to the experiment on vessel arrivals, all important covariates appear broadly balanced after matching.

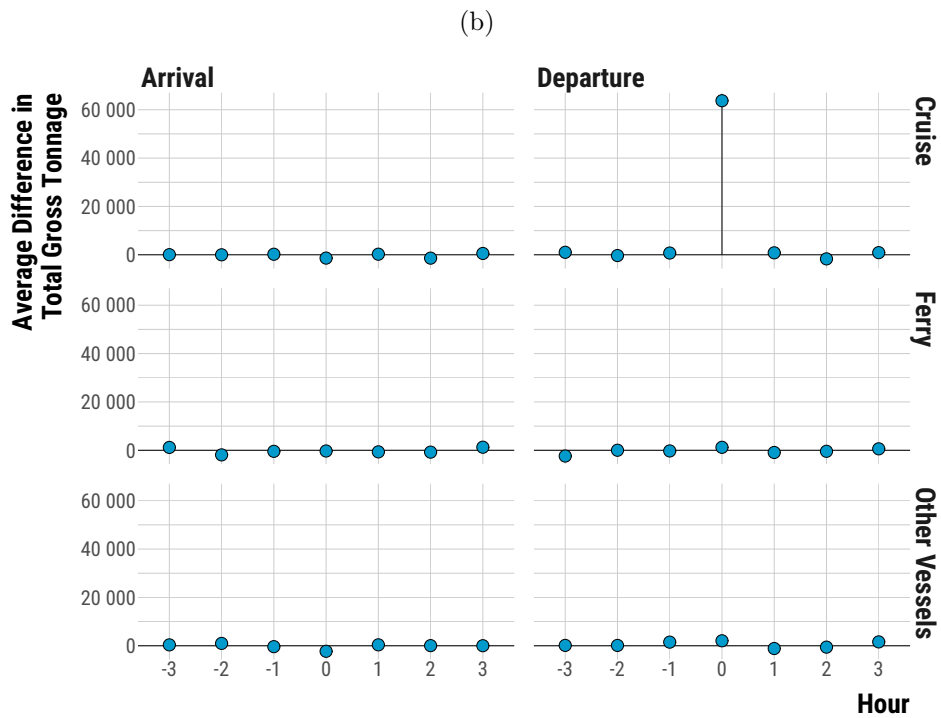



Figure B.1.24 – Visualizing the Occurrence of the Treatment.

Notes: Panel A shows the average gross tonnage for treated and control units from three hours before the occurrence of the treatment up to three hours after. Panel B shows the average difference in gross tonnage between treated and control units by vessel type and flow.

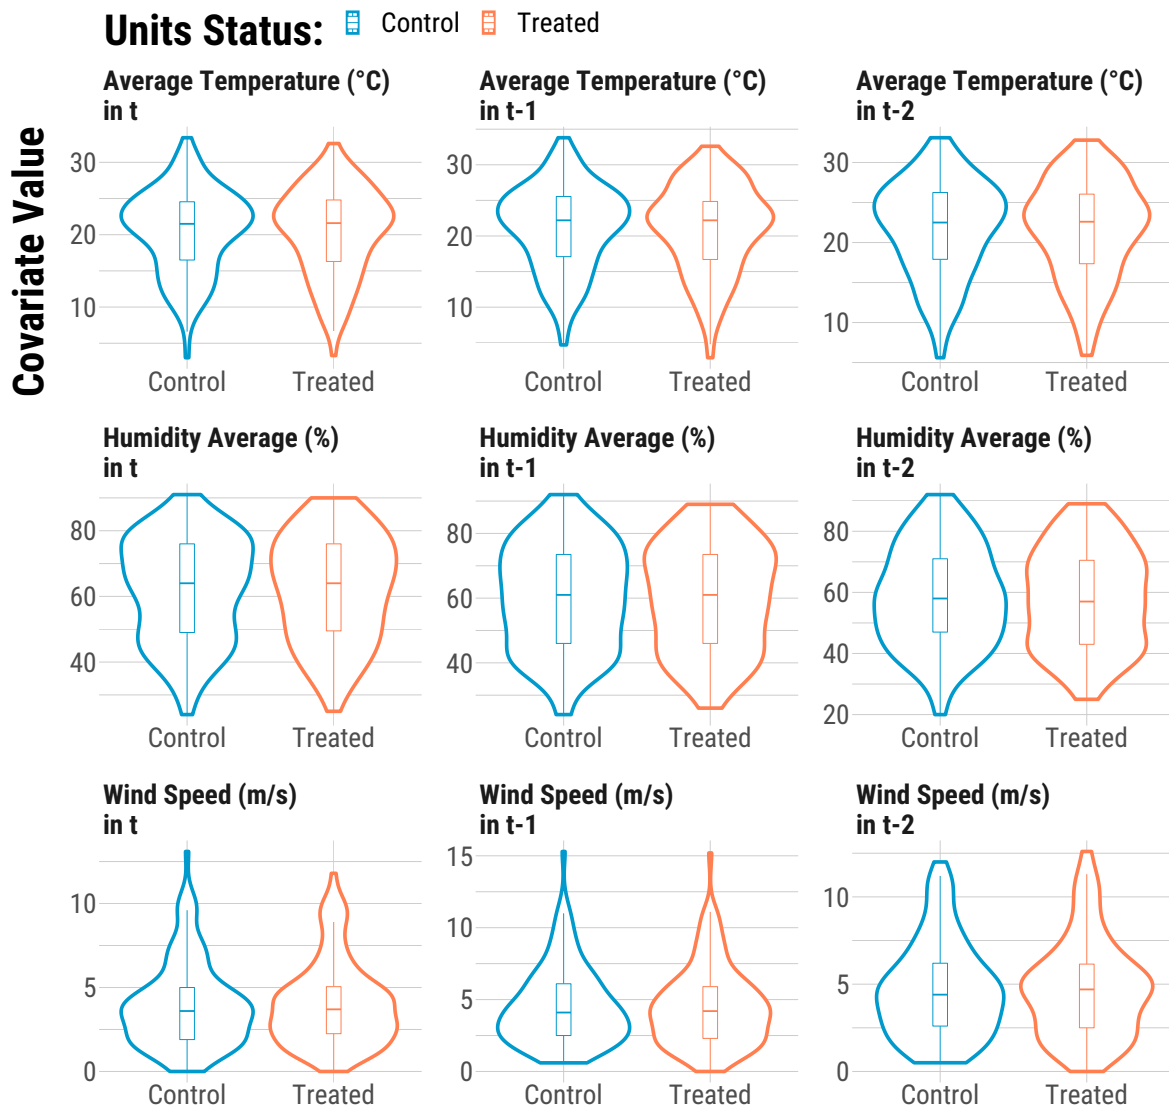


Figure B.1.25 – Continuous Weather Covariates Balance.

Notes: The density and boxplot distributions of continuous weather covariates are plotted for treated and control units.




Figure B.1.26 – Categorical Weather Covariates Balance.

Notes: The proportion of distribution of wind directions is displayed for matched treated and control units.

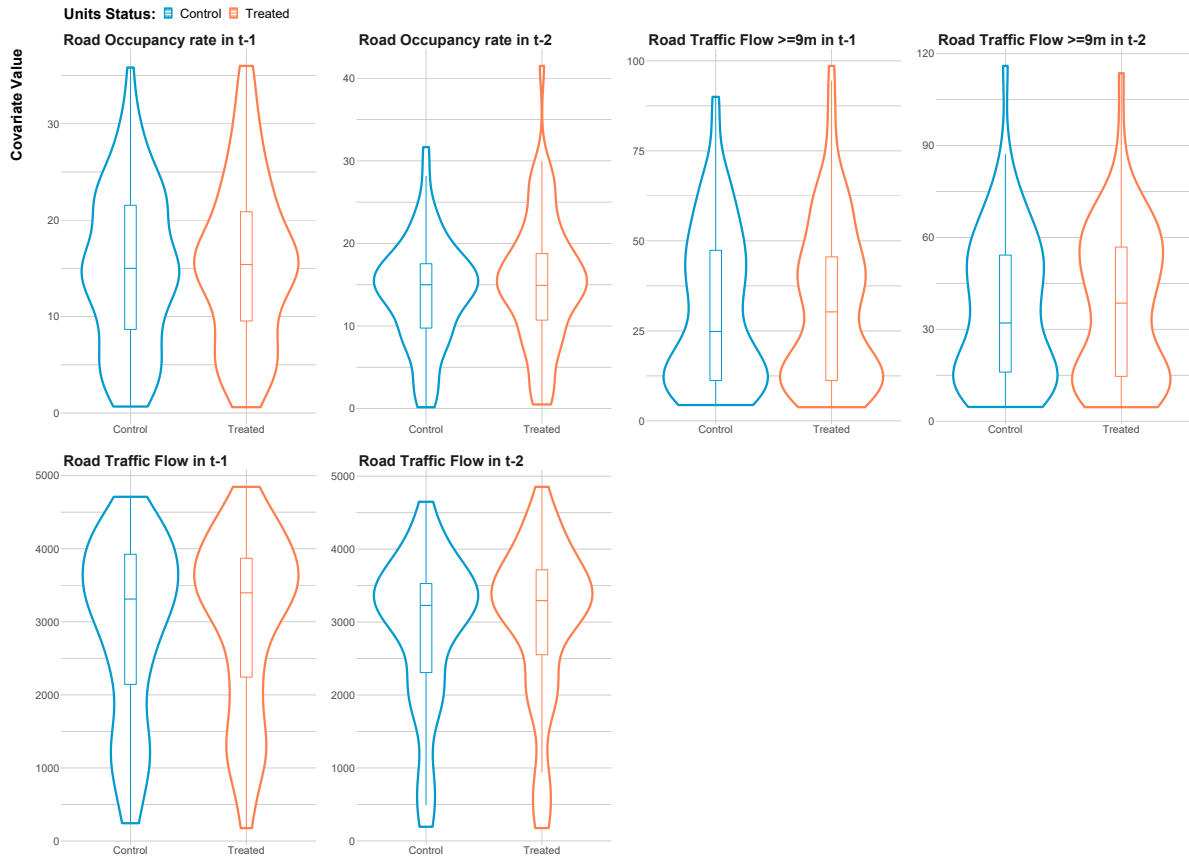


Figure B.1.27 – Balance of Road traffic variables.

Notes: Each panel displays the density distribution and boxplot of road traffic variables by treatment status.

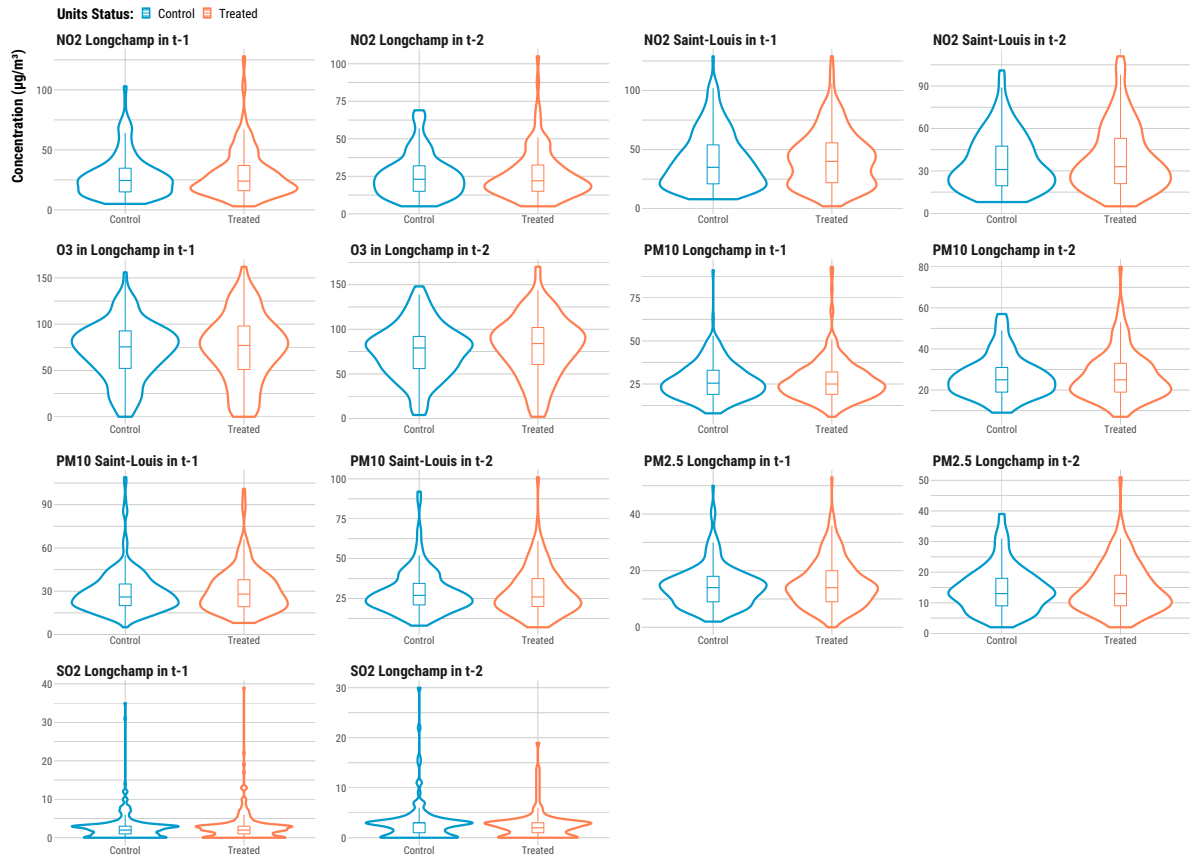


Figure B.1.28 – Balance of Air Pollutants.

Notes: Each panel displays the density distribution and boxplot of pollutants by treatment.

Improvement of Covariates Balance after Matching Covariate balanced improves after matching, as shown for the different variables (Figures B.1.29–B.1.34) and in the summary figure (Figure B.1.35).




Figure B.1.29 – Love Plots for Continuous Weather Covariates.

Notes: The standardized mean differences between treated and control units are plotted, before and after matching.




Figure B.1.30 – Love Plots for Wind Direction.

Notes: The absolute mean differences in proportions between treated and control units are plotted, before and after matching.




Figure B.1.31 – Love Plot for Rainfall.

Notes: The absolute mean differences in proportions between treated and control units are plotted, before and after matching. If a point is missing, it means that it was not raining for any units on that hour.

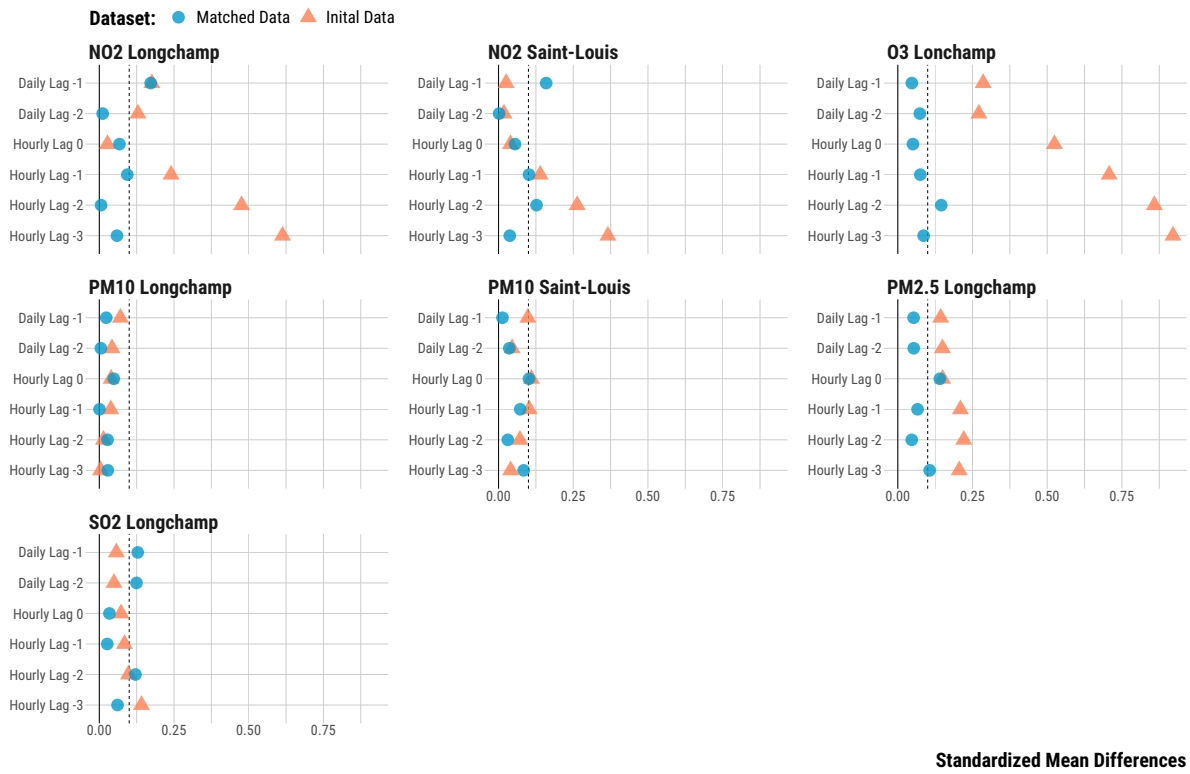


Figure B.1.32 – Love Plots for Air Pollutants.

Notes: The standardized mean differences between treated and control units are plotted, before and after matching.




Figure B.1.33 – Love Plots for Vessel Traffic.

Notes: The absolute standardized mean differences in gross tonnage between treated and control units are plotted, before and after matching.




Figure B.1.34 – Love Plots for Road Traffic.

Notes: The absolute standardized mean differences in road traffic between treated and control units are plotted, before and after matching.




Figure B.1.35 – Overall Improvement in Covariate Balance.

Notes: Panel A shows the boxplot of standardized mean differences for continuous covariates before and after matching. Panel B shows the boxplot of absolute differences in percentage points for categorical weather covariates and calendar indicators before and after matching.

Comparing Initial Data to Matched Data In Figure B.1.36, we can see that the matched hours have a higher average temperature, a lower humidity and a lower occurrence of rain. In Figure B.1.37, we notice that matched hours are mostly around 6 pm and in the spring and summer seasons, with Tuesdays and Wednesdays and the years 2010, 2017 and 2018 over-represented compared to the initial data.




Figure B.1.36 – Comparing Weather Covariates Distribution for Initial and Matched Datasets.

Notes: Panel A plots the density distributions of continuous weather covariates for the initial and matched datasets. Panel B displays the proportion of observations belonging to a particular category of discrete weather covariates.

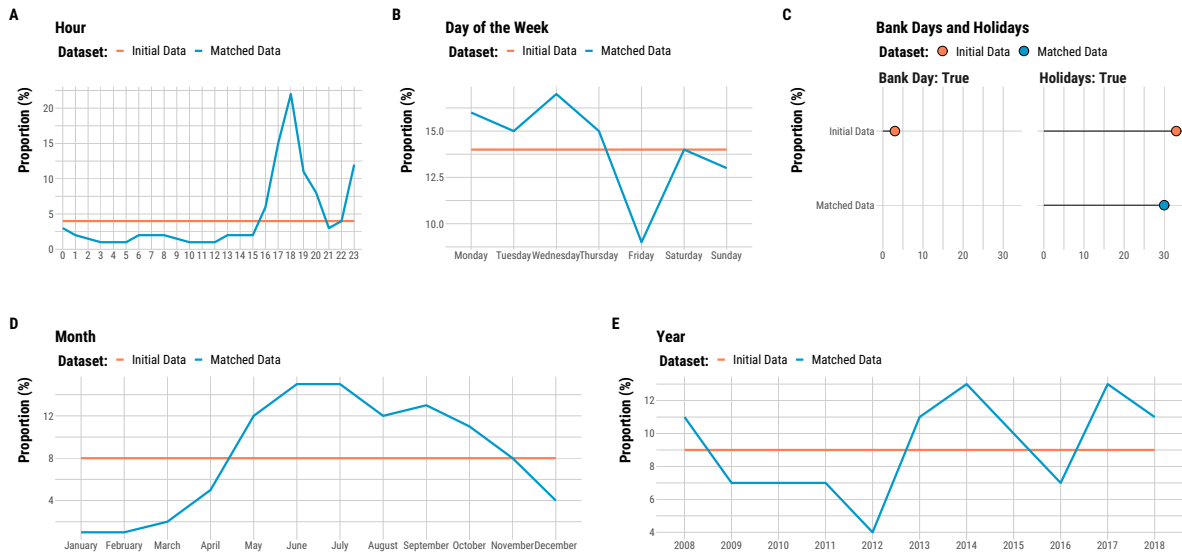


Figure B.1.37 – Comparing Calendar Covariates Distribution for Initial and Matched Datasets.

Notes: Proportions of observations belonging to each hour of the day (Panel A), day of the week (Panel B), bank days and holidays (Panel C), month (Panel D) and year (Panel E) for the initial and matched datasets.

Daily Experiment on Cruise Traffic

See Section B.1.5 for more details on the graphs in each subsection.

Checking the Treatment On average, the traffic of cruise is similar for treated and control units the day before the treatment occurs (Panel A of Figure B.1.38). There is a strong increase in cruise vessel gross tonnage in day 0 while the difference in tonnage for other vessel types is relatively small (Panel B).

Evidence of Covariates Balance for the Matched Pairs Similar to the experiment on vessel arrivals, all important covariates appear broadly balanced after matching.

Improvement of Covariates Balance after Matching In Figure B.1.43, we can see that the balance of continuous weather covariates has improved overall after the matching. In Figures B.1.44–B.1.45, we note that the balance of wind direction categories and the rainfall dummy is relatively similar before and after matching. In Figure B.1.46, the balance has improved for most pollutants. In Figure B.1.47, we see that the balance of gross tonnage for each vessel type has improved. In Figure B.1.48, the balance of the road traffic variable has improved after matching in t but not in $t-1$. Finally, we summarize in Figure B.1.49 the global improvement in covariates balance for continuous and categorical

variables (for which we include calendar indicators). The matching procedure improves balance for both types of covariates.




Figure B.1.38 – Visualizing the Occurrence of the Treatment.

Notes: Panel A shows the average gross tonnage for treated and control units from one day before the occurrence of the treatment up to one day after. Panel B shows the average difference in gross tonnage between treated and control units by vessel type.

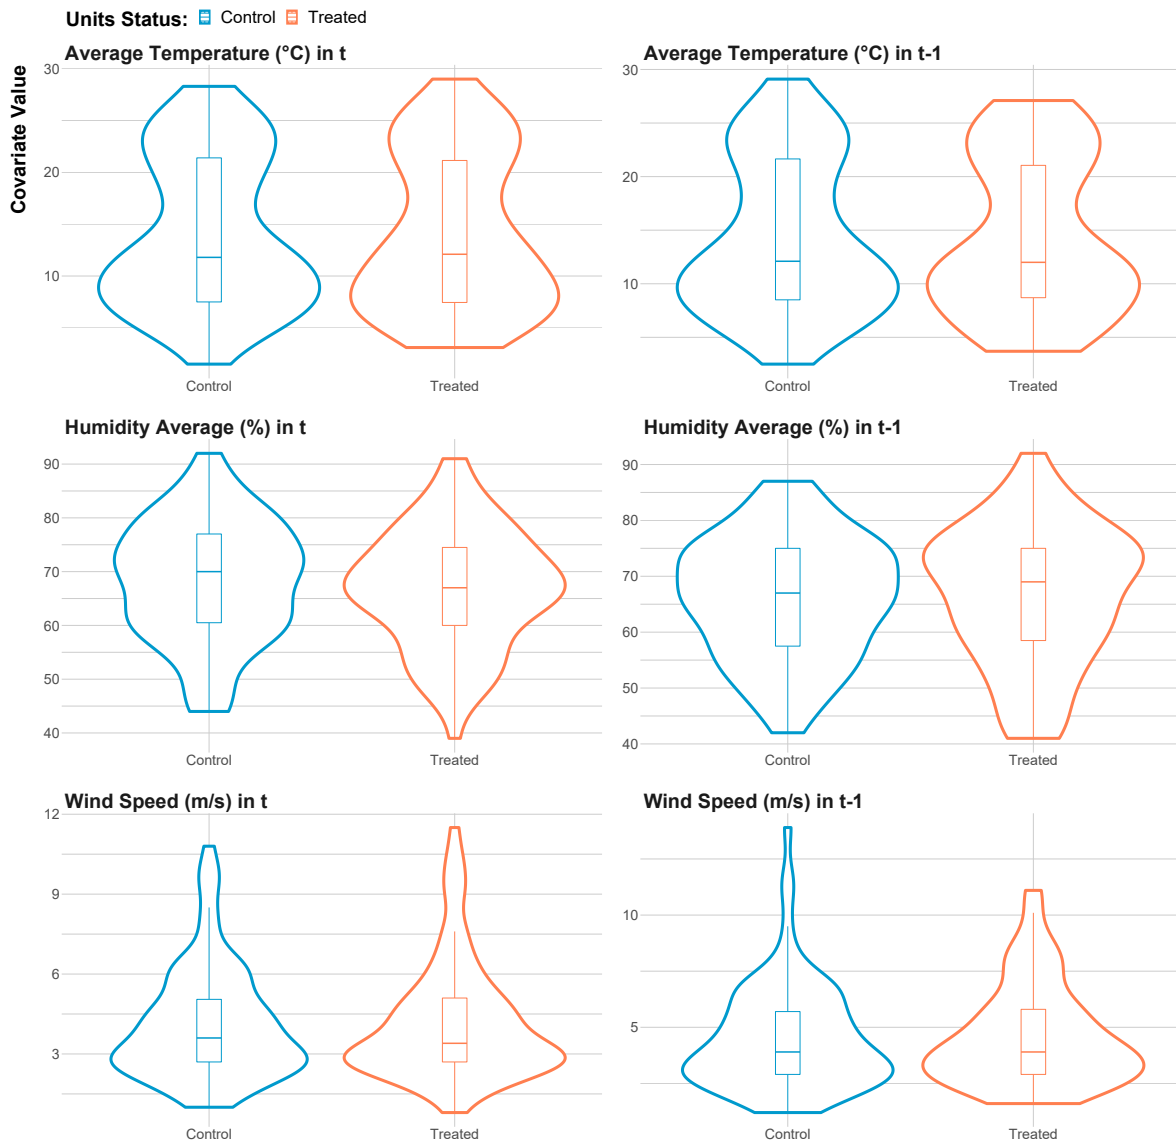


Figure B.1.39 – Balance of Continuous Weather Covariates.

Notes: The density and boxplot distributions of continuous weather covariates are plotted for treated and control units.




Figure B.1.40 – Balance of Categorical Weather Covariates.

Notes: Panel A shows the density distribution and the boxplot of continuous weather covariates for treated and control units. Panel B shows the average proportion of observations of categorical weather covariates for treated and control units.




Figure B.1.41 – Balance of Air Pollutants.

Notes: Each panel displays the density distribution and boxplot of each pollutant by treatment status on treatment day and for the two previous days.




Figure B.1.42 – Balance of Month and Year Indicators.

Notes: Panel A plots the distribution of the month variable for control and treated units. Panel B displays the distribution of the year variable for control and treated units.




Figure B.1.43 – Love Plot for Continuous Weather Covariates.

Notes: The standardized mean differences between treated and control units are plotted, before and after matching.




Figure B.1.44 – Love Plots for Wind Direction.

Notes: The absolute mean differences in proportions between treated and control units are plotted, before and after matching.




Figure B.1.45 – Love Plot for Rainfall.

Notes: The absolute mean differences in proportions between treated and control units are plotted, before and after matching.




Figure B.1.46 – Love Plots for Air Pollutants.

Notes: The standardized mean differences between treated and control units are plotted, before and after matching.




Figure B.1.47 – Love Plots for Vessel Traffic.

Notes: The absolute standardized mean differences in gross tonnage between treated and control units are plotted, before and after matching.




Figure B.1.48 – Love Plot for Road Traffic.

Notes: The absolute standardized mean differences in road traffic between treated and control units are plotted, before and after matching.




Figure B.1.49 – Overall Improvement in Covariates Balance.

Notes: Panel A plots the boxplots of the distribution of standardized mean differences for continuous covariates before and after matching. Panel B plots the boxplots of the distribution of absolute differences in percentage points for categorical covariates before and after matching.

Comparing Initial to Matched Datasets The initial dataset designates the dataset of the full time period ($N=4,018$). In Figure B.1.50, we can see that matched days have a lower average temperature, a higher average humidity, but are similar in terms of wind speed, wind direction and rainfall occurrence. In Figure B.1.51, we notice that matched days are mostly in the winter season and in August while spring is under-represented. Wednesdays are over-represented and weekends underrepresented.




Figure B.1.50 – Comparing Weather Covariates Distribution for Initial and Matched Datasets.

Notes: Panel A plots the density distributions of continuous weather covariates for the initial and matched datasets. Panel B displays the proportion of observations belonging to a particular category of discrete weather covariates.

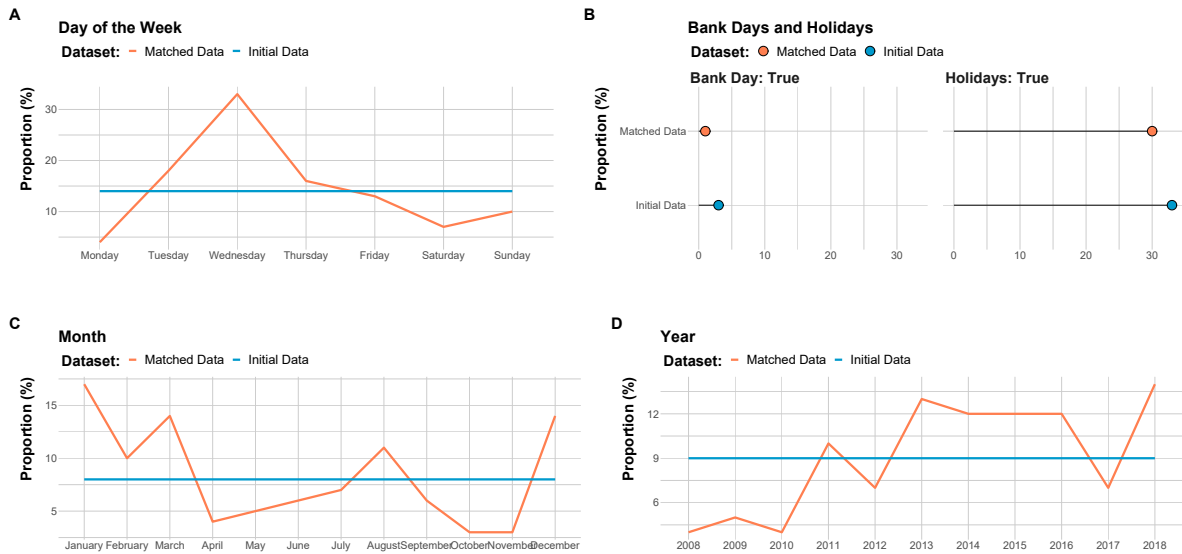


Figure B.1.51 – Comparing Calendar Covariates Distribution for Initial and Matched Datasets.

Notes: Proportions of observations belonging to each hour of the day (Panel A), day of the week (Panel B), bank days and holidays (Panel C), month (Panel D) and year (Panel E) for the initial and matched datasets.

B.1.6 Results - Analysis of the Matched Data

In this section, we first explain how to compute Fisherian intervals with a toy example. Then, for each experiment, we present the detailed results of our analysis: (i) the distribution of pair differences in concentration for each pollutant; (ii) the approximate p -values for the test of the sharp null hypothesis of no effect; and (iii) the 95% Fisherian intervals.

Toy Example for Understanding Randomization Inference

The Toy Example Considered In this toy example, we want to estimate the effect of cruise vessels docking at Marseille port on NO_2 concentrations. For simplicity, we imagine that our matching procedure resulted in 10 pairs of hours with similar weather and calendar characteristics. Treated hours are hours with cruise ships docking the port while control hours are hours without cruise vessels. The outcome of the experiment is the hourly NO_2 measured at a station in the city. The exposition of this toy example is inspired by those found in Rosenbaum’s textbook and Dasgupta and Rubin’s forthcoming textbook (Rosenbaum et al., 2010; Dasgupta and Rubin, 2021).

Science Table Table B.1.1 is the “Science Table” of our imaginary experiment, that is to say a table showing the treatment status and potential outcomes for each observation.

- The first column *Pair* is the indicator of the pair, taking values i from I to X (Roman numbers).
- The second column *Unit Index* is a within-pair index j for each unit.
- The third column W indicates the treatment allocation. $W = 1$ for treated units and $W = 0$ for controls.
- The fourth and fifth columns are the potential outcomes of each unit and represent the NO_2 concentrations measured in $\mu\text{g}/\text{m}^3$. $Y(W = 0)$ is the potential outcome when the unit does not receive the treatment and $Y(W = 1)$ is the potential outcome when the unit is treated. In this imaginary example, we know for each unit the values of both potential outcomes.
- The sixth column τ is the unit constant causal effect. Here, the causal effect is equal to $+3 \mu\text{g}/\text{m}^3$.
- The last column Y represent the observed outcome, which depends on treatment allocation: $Y_{i,j} = W_{i,j}Y_{i,j}(1) + (1 - W_{i,j})Y_{i,j}(0)$. Here, in each pair, the first unit does not receive the treatment and we observe $Y(0)$, while the second unit is treated and we observe $Y(1)$.

Pair	Unit Index	W	Y(0)	Y(1)	τ	Y
I	1	0	37	40	3	37
I	2	1	21	24	3	24
II	1	0	33	36	3	33
II	2	1	22	25	3	25
III	1	0	38	41	3	38
III	2	1	50	53	3	53
IV	1	0	41	44	3	41
IV	2	1	47	50	3	50
V	1	0	41	44	3	41
V	2	1	56	59	3	59
VI	1	0	33	36	3	33
VI	2	1	40	43	3	43
VII	1	0	23	26	3	23
VII	2	1	28	31	3	31
VIII	1	0	27	30	3	27
VIII	2	1	31	34	3	34
IX	1	0	27	30	3	27
IX	2	1	19	22	3	22
X	1	0	51	54	3	51
X	2	1	31	34	3	34

Table B.1.1 – Science Table.

Observed Data In reality, researchers do not have access to the Science Table but only to the Table B.1.2 below: they only have information on the pair indicator, the unit index, the treatment allocated and the observed NO_2 concentration. Our randomization inference procedure will be based only on this table.

Pair	Unit Index	W	Y
I	1	0	37
I	2	1	24
II	1	0	33
II	2	1	25
III	1	0	38
III	2	1	53
IV	1	0	41
IV	2	1	50
V	1	0	41
V	2	1	59
VI	1	0	33
VI	2	1	43
VII	1	0	23
VII	2	1	31
VIII	1	0	27
VIII	2	1	34
IX	1	0	27
IX	2	1	22
X	1	0	51
X	2	1	34

Table B.1.2 – Observed Table.

The intuition behind randomization-based inference is that the value for the estimated treatment effect (which is usually the average of pair differences) may have simply occurred by chance, rather than because of the treatment. If it happened only by chance, we would probably detect a similar effect size under a random permutation of the treatment allocation, where some of the “true” treated units are considered to be control units, and some of the “true” control units are considered to be treated units. We want to test how the treatment effect changes when treatment allocation is different from what it actually is, and we want to do this test for every possible permutation of the treatment allocation.

To conduct randomization-based inference, we need to:

- Know the number of unique treatment allocations. In a pair experiment, there are 2^I unique treatment allocations, with I the number of pairs. In this experiment, there are 1,024 unique treatment allocations.
- Define a test statistic. We will build its distribution under the sharp null hypothesis. Here, we use the average of pair differences as a test statistic.

Testing the Sharp Null of No Effect The sharp null hypothesis of no treatment states that $Y_{i,j}(0) = Y_{i,j}(1) \forall i, j$, that is to say the treatment has no effect for each observation. Under this hypothesis, we could impute the missing $Y(1)$ for control units and the missing $Y(0)$ for treated units. To create the distribution of the test statistic under this sharp null hypothesis, we could permute the treatment vector, express for each unit the outcome observed according to the permuted value of the treatment and then compute the average of pair differences. As this would be cumbersome in terms of programming, Paul Rosenbaum offers a more efficient procedure (chapter II of his textbook):

- For each unit i of each pair j , its observed outcome is equal to $Y_{i,j} = W_{i,j}Y_{i,j}(1) + (1 - W_{i,j})Y_{i,j}(0)$.
- The difference in outcomes for the pair i (i.e. the difference in outcomes between the treated and control units) is equal to $D_i = (W_{i,1} - W_{i,2})(Y_{i,1} - Y_{i,2})$
- Under the sharp null hypothesis of no effect, we have $Y_{i,j}(0) = Y_{i,j}(1)$ so that $D_i = (W_{i,1} - W_{i,2})(Y_{i,1}(0) - Y_{i,2}(0))$.
- If the treatment allocation within a pair is $(W_{i,1}, W_{i,2}) = (0,1)$, then $D_i = -(Y_{i,1}(0) - Y_{i,2}(0))$. If the treatment allocation is $(W_{i,1}, W_{i,2}) = (1,0)$, then $D_i = Y_{i,1}(0) - Y_{i,2}(0)$.
- Thus, under the sharp null hypothesis of no effect, the randomization of the treatment only changes the sign of the pair differences in outcomes.

In terms of programming, we proceed as follows:

1. We first compute the observed average of pair differences. We are now working with a table with 10 pair differences.
2. We then compute the permutations matrix of all possible treatment assignments. This is a matrix of 10 rows with 1024 columns (the number of possible permutations with two possible treatment values and 10 units).
3. For each vector of treatment assignment, we compute the average of pair differences.

Figure B.1.52 shows the distribution of the test statistic under the sharp null hypothesis. The vertical orange line is the value of the observed test statistic in Table B.1.2.




Figure B.1.52 – Distribution of the Test Statistic Under the Sharp Null Hypothesis.

Finally, we can compute a two-sided p -value. We again follow the explanations in Rosenbaum 2010 (chapter II) (Rosenbaum et al., 2010):

1. We calculate the proportion of permuted test statistics that are lower than the observed test statistic, and the proportion that are higher
2. We double the smallest proportion
3. We take the minimum between this value and one

In our example, the exact two-sided p -value is equal to 0.55.

Computing the Fisherian Interval To calculate a Fisherian interval, we follow the explanations provided in the second chapter of Dasgupta and Rubin’s forthcoming textbook (Dasgupta and Rubin, 2021):

- We test a set of K sharp null hypotheses $H_0^k: Y_{i,j}(1) = Y_{i,j}(0) + \tau_k$ for $k = 1, \dots, K$ and where τ_k is a constant unit-level treatment effect size. Here, we choose to test 81 sharp null hypotheses for 81 possible values for τ_k , from $-20 \mu\text{g}/\text{m}^3$ to $+20 \mu\text{g}/\text{m}^3$ with increments of $0.5 \mu\text{g}/\text{m}^3$.
- For each constant treatment effect τ_k , we calculate the upper p -value associated with the hypothesis $H_0^k: Y_{i,j}(1) - Y_{i,j}(0) > \tau_k$ and the lower p -value $H_0^k: Y_{i,j}(1) - Y_{i,j}(0) < \tau_k$.

- To test each hypothesis, we compute the distribution of the test statistic. The sequence of K hypotheses $H_0^k: Y_{i,j}(1) - Y_{i,j}(0) > \tau_k$ forms an upper p -value function of τ , $p^+(\tau)$, while the sequence of alternative hypotheses $H_0^k: Y_{i,j}(1) - Y_{i,j}(0) < \tau_k$ makes a lower p -value function, τ , $p^-(\tau)$. To calculate the bounds of the $100(1-\alpha)\%$ Fisherian interval, we solve $p^+(\tau) = \frac{\alpha}{2}$ for τ to get the lower limit and we solve $p^-(\tau) = \frac{\alpha}{2}$ for τ to get the upper limit. We set our α significance level to 0.05 and thus compute 95% Fisherian intervals. This procedure allows us to get the range of *constant* treatment effects consistent with our data.

As a point estimate of the Fisherian interval, we take the observed value of our test statistic, which is the average of pair differences in a pollutant's concentration. It is important to note that our test statistic is an estimate of the individual-level treatment effect of a hypothetical experiment rather than an estimate of an average treatment effect.

We could impute the missing potential outcomes for each hypothesis, randomly allocate the treatment, express the observed outcome, and compute the average of pair differences. We again employ a computational shortcut suggested by Rosenbaum (Rosenbaum et al., 2010).

- We start by making a sharp null hypothesis of a constant treatment effect τ such that $Y_{i,j}(1) = Y_{i,j}(0) + \tau$.
- For a pair i , recall that the observed pair difference in outcomes is $D_i = (W_{i,1} - W_{i,2})(Y_{i,1} - Y_{i,2})$.
- Under the sharp hypothesis, we have $D_i = (W_{i,1} - W_{i,2})((Y_{i,1} + \tau W_{i,1}) - (Y_{i,2} + \tau W_{i,2}))$.
- We rearrange the right-hand side expression and find that $D_i = \tau + (W_{i,1} - W_{i,2})(Y_{i,1}(0) - Y_{i,2}(0))$
- We have $D_i - \tau = (W_{i,1} - W_{i,2})(Y_{i,1}(0) - Y_{i,2}(0))$. This equation means that the observed pair difference in outcomes minus the hypothesized treatment effect is equal to $\pm(Y_{i,1}(0) - Y_{i,2}(0))$. We can therefore carry out the randomization inference procedure seen in the previous section from the vector of observed pair differences adjusted for the hypothesized treatment effect.

Figure B.1.53 shows the resulting p -value functions from our toy example. The orange line represents the alpha significance level, set at 5%, divided by two. The lower bound of the 95% interval is equal to $-7 \mu\text{g}/\text{m}^3$; its upper bound to $+11.5 \mu\text{g}/\text{m}^3$ and the point estimate is $+2.4 \mu\text{g}/\text{m}^3$. For this imaginary experiment, our point estimate is close to the true constant effect but the 95% Fisherian interval is wide: the data are consistent with both large negative and positive constant treatment effects.




Figure B.1.53 – Upper and Lower P -Value Functions.

Hourly Experiment on Cruise Vessels' Arrivals

Distribution of Pair Differences for each Pollutant over Time In Figures B.1.54–B.1.58, we display the distribution of pair differences in each pollutant's concentration from 3 hours before the treatment up to 3 hours after.




Figure B.1.54 – Distribution of Pair Differences in NO_2 Concentrations.




Figure B.1.55 – Distribution of Pair Differences in O₃ Concentrations.




Figure B.1.56 – Distribution of Pair Differences in SO₂ Concentrations.




Figure B.1.57 – Distribution of Pair Differences in PM₁₀ Concentrations.




Figure B.1.58 – Distribution of Pair Differences in PM_{2.5} Concentrations.

Sharp Null Hypotheses of No Effect Figure B.1.59 shows the two-sided p -values for the test of the sharp null hypothesis of no effect. The exact values of the two-sided p -values are available in the replication material of this experiment (script entitled `6_script_analyzing_results.html`).




Figure B.1.59 – Two-Sided P -Value for the Sharp Null Hypothesis of No Effect

Fisherian Intervals For each pollutant, we plot in Figure B.1.60 the 95% Fisherian Intervals for the constant treatment effects consistent with our data.




Figure B.1.60 – 95% Fisherian Intervals

Fisherian Intervals for the subset of hourly pairs with wind blowing from the West We estimate the point estimate and Fisherian Intervals again, keeping only the 82 matched pairs for which wind blows from West to East, such that the pollution monitoring stations are downwind from the cruise terminal. Figure B.1.61 the 95% Fisherian Intervals for the constant treatment effects consistent with our data. For most pollutants, the point estimates are higher than for the main estimates, suggesting a larger effect when wind blows from the cruise terminal to the monitoring stations.




Figure B.1.61 – 95% Fisherian Intervals

Hourly Experiment on Cruise Vessels' Departures

Distribution of Pair Differences for each Pollutant over Time In Figures B.1.62–B.1.66, we display the distribution of pair differences in each pollutant's concentration from 3 hours before the treatment up to 3 hours after.




Figure B.1.62 – Distribution of Pair Differences in NO₂ Concentrations.




Figure B.1.63 – Distribution of Pair Differences in O₃ Concentrations.




Figure B.1.64 – Distribution of Pair Differences in SO₂ Concentrations.




Figure B.1.65 – Distribution of Pair Differences in PM₁₀ Concentrations.




Figure B.1.66 – Distribution of Pair Differences in PM_{2.5} Concentrations.

Sharp Null Hypotheses of No Effect We plot in Figure B.1.67 the two-sided p -values for the test of the sharp null hypothesis of no effect. Interested readers can retrieve the exact value of the two-sided p -values in the replication material of this experiment (script entitled `6_script_analyzing_results.html`).




Figure B.1.67 – Two-Sided P -Value for the Sharp Null Hypothesis of No Effect

Fisherian Intervals For each pollutant, we plot in Figure B.1.68 the 95% Fisherian Intervals for the constant treatment effects consistent with our data.




Figure B.1.68 – 95% Fisherian Intervals

Fisherian Intervals for the subset of hourly pairs with wind blowing from the West We estimate the point estimate and Fisherian Intervals again, keeping only the 82 matched pairs for which wind blows from West to East, such that the pollution monitoring stations are downwind from the cruise terminal. Figure B.1.69 the 95% Fisherian Intervals for the constant treatment effects consistent with our data. For most pollutants, the point estimates do not change much, except for $\text{PM}_{2.5}$ and ozone, in contrast to the experiment on vessels' arrivals.




Figure B.1.69 – 95% Fisherian Intervals

Daily Experiment on Cruise Vessel Traffic

Distribution of Pair Differences for each Pollutant over Time In Figures B.1.70–B.1.58, we display the distribution of pair differences in each pollutant concentration from 1 day before the treatment up to 1 day after.




Figure B.1.70 – Distribution of Pair Differences in NO_2 Concentrations.




Figure B.1.71 – Distribution of Pair Differences in O₃ Concentrations.




Figure B.1.72 – Distribution of Pair Differences in SO₂ Concentrations.




Figure B.1.73 – Distribution of Pair Differences in PM₁₀ Concentrations.




Figure B.1.74 – Distribution of Pair Differences in $PM_{2.5}$ Concentrations.

Dose-response In Figure B.1.75, we show how the pair difference at t varies with the treatment intensity, that is, the difference in the total gross tonnage of cruise vessels entering the port between treated and control units. The heterogeneity in pair differences does not seem to reflect differences in treatment intensity.

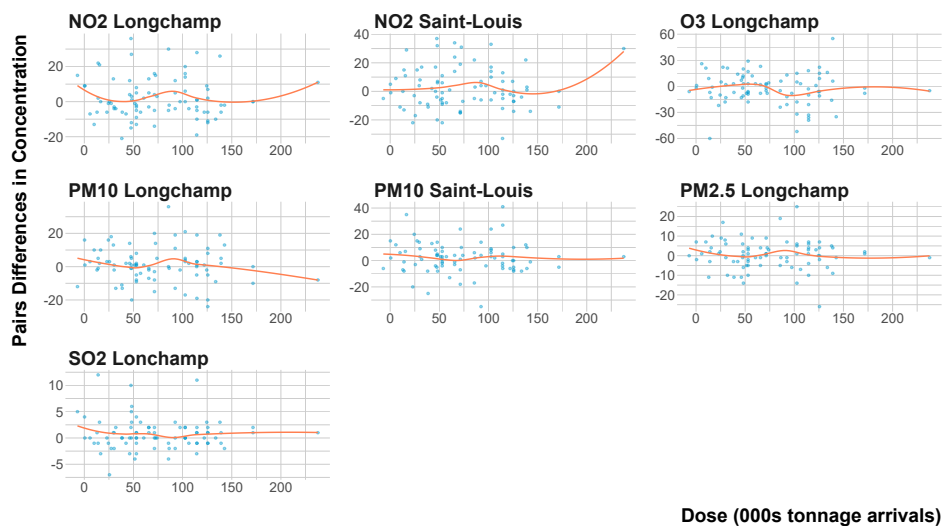


Figure B.1.75 – Pair Differences by difference in tonnage of entering cruise vessels.

Sharp Null Hypotheses of No Effect We plot in Figure B.1.76 the two-sided p -values for the test of the sharp null hypothesis of no effect. Interested readers can retrieve the exact value of the two-sided p -values in the replication material of this experiment (script entitled `6_script_analyzing_results.html`).




Figure B.1.76 – Two-Sided P -Value for the Sharp Null Hypothesis of No Effect

Fisherian Intervals For each pollutant, we plot in Figure B.1.77 the 95% Fisherian Intervals for the constant treatment effects consistent with our data.




Figure B.1.77 – 95% Fisherian Intervals

Fisherian Intervals for pairs on days with wind blowing from the West We estimate the point estimate and Fisherian Intervals again, keeping only the 65 matched pairs for which wind blows from West to East, such that the pollution monitoring stations are downwind from the cruise terminal. Figure B.1.78 the 95% Fisherian Intervals for the constant treatment effects consistent with our data. The point estimates are generally slightly higher than for the main estimates, suggesting a larger effect when wind blows from the cruise terminal to the monitoring stations.




Figure B.1.78 – 95% Fisherian Intervals, only pairs where wind blows from the West

Effect on induced road traffic Passengers from cruise vessels are likely to spend the day visiting Marseille in a coach. We examine whether the effect we detect on pollutant concentrations are likely to be driven by the induced road traffic from cruise passengers. We expect that if this is the case, we will see an increase in the road traffic flow from vehicles longer than 9 meters (coaches are in this category) on treated days. Since we only have road traffic data for the 2011-2016 period, we only keep the pairs from this period. This reduces our matched sample to 38 pairs. Figure B.1.79 shows that none of the road traffic indicators increases at t . At $t+1$, there is an increase in road occupancy rate and the traffic flow of regular vehicle, but not of long vehicles. This increase in traffic flow is unlikely to be induced by the increase in vessel traffic at t , and may be due to the small sample size and caused by outliers. It may explain part of the increase in particulate matter concentrations at $t + 1$ on treated days.




Figure B.1.79 – 95% Fisherian Intervals for road traffic variables, 2011-2016 period

B.1.7 Sensitivity Analyses

In this section, we provide four types of sensitivity analyses for each experiment, as well as a retrospective design analysis of our experiment. First, we use the Wilcoxon signed-rank statistic as an alternative test statistic, because it is more robust to outliers and several matched pairs have a high difference in outcomes. Second, we calculate 95% Fisherian intervals for pairs after removing the missing values that were imputed. Third, we switch from Fisherian inference to Neymanian inference and estimate the average treatment effect and its associated 95% confidence intervals (Splawa-Neyman et al., 1990). Fourth, we compare our results from those found using a simple multivariate regression on the full dataset. Finally, we carry out a retrospective design analysis (Gelman and Carlin, 2014; Gelman et al., 2020) and calculate for each experiment the power of our analysis and the risk of type-S and type-M errors for different assumed true values of our estimand.

Alternative Test Statistic: Wilcoxon Signed-Rank Statistic

Let D_i be the observed difference in concentrations between the treated and control unit of pair i for a given pollution outcome. The Wilcoxon signed-rank statistic is defined as:

$$T = \sum_{i=1}^I \text{sgn}(D_i) \times q_i \tag{B.1}$$

where $\text{sgn}(D_i) = 1$ if $D_i > 0$ and $\text{sgn}(D_i) = -1$ if $D_i < 0$, and q_i is the rank of $|D_i|$ (Rosenbaum et al. (2010)).

In Figure B.1.80, Figure B.1.81 and Figure B.1.82, we plot for each experiment the 95% Fisherian intervals computed using the average pair differences and the Wilcoxon

signed-rank test statistics. Overall, the results are qualitatively similar, except for SO_2 for which the intervals are smaller when computed with the Wilcoxon signed-rank test statistic. This may be linked to the skewed distribution of SO_2 concentrations, with a high number of low values (between 0 and $2 \mu\text{g}/\text{m}^3$) and a few outliers (with values up to $30 \mu\text{g}/\text{m}^3$).




Figure B.1.80 – 95% Fisherian Intervals based Wilcoxon signed-rank Test Statistic.

Hourly Experiment - Cruise Vessels' Arrivals

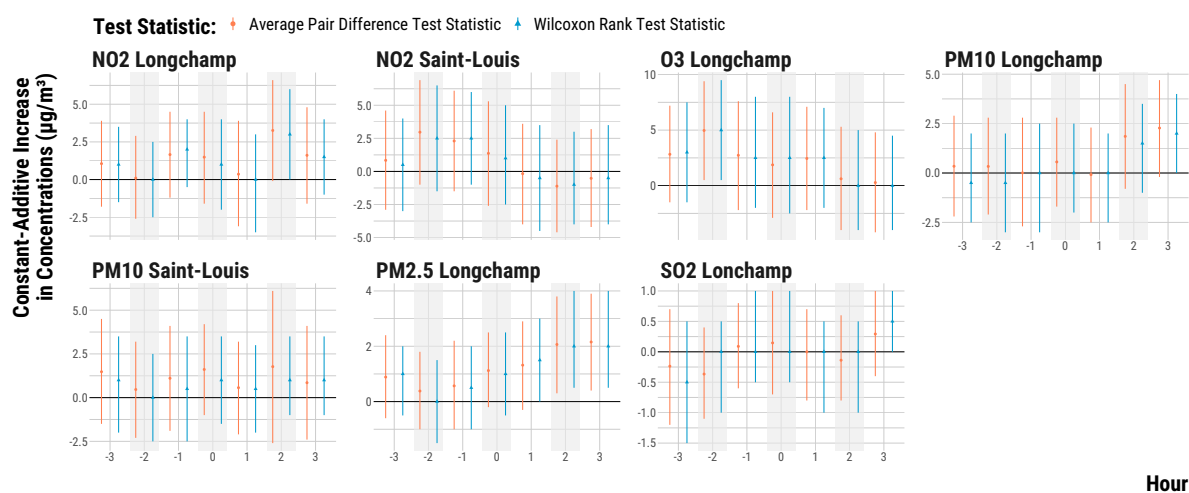


Figure B.1.81 – 95% Fisherian Intervals based Wilcoxon signed-rank Test Statistic.

Hourly Experiment - Vessel's Departures




Figure B.1.82 – 95% Fisherian Intervals based Wilcoxon signed-rank Test Statistic.

Daily Experiment

Fisherian Intervals for Non-Imputed Observations

Our simulation exercise (available in the replication package) shows that large imputation errors sometimes occur: for example, for a random sample of observed NO₂ values at Longchamp, the average of absolute differences between truly observed and imputed values is 7.4 µg/m³, which is 25% of the hourly average. We verify that our results are not driven by imputation errors by re-calculating 95% Fisherian intervals (using the Wilcoxon signed-ranked statistic) for each lead and lag after removing the pairs for which the value of the control or/and the treated unit was imputed on that lead/lag.

Hourly Experiment - Cruise Vessels' Arrivals Figure B.1.83 shows the number of pairs with missing pollutant concentrations for the hourly experiment on cruise vessels' arrivals. For PM₁₀ measured at Saint-Louis and PM_{2.5} measured at Longchamp, the number of pairs with missing values can represent up to 15% of the matched pairs.




Figure B.1.83 – Number of Pairs with Missing Values.

Figure B.1.84 shows the 95% Fisherian intervals calculated with the Wilcoxon signed-rank statistic, with and without missing values imputations. The intervals are slightly wider after removing pairs with missing values, but the results are similar to those including imputed concentrations.




Figure B.1.84 – 95% Fisherian Intervals based on Wilcoxon signed-rank Test Statistic for Non-Imputed Concentrations.

Hourly Experiment - Cruise Vessels' Departures Figure B.1.85 shows the number of pairs with missing pollutant concentrations for the hourly experiment on cruise vessels' departures. Like for the experiment on cruise vessels' arrivals, the number of pairs with missing values can represent up to 15% of the matched pairs.




Figure B.1.85 – Number of Pairs with Missing Values.

Figure B.1.86 shows the 95% Fisherian intervals calculated with the Wilcoxon signed-rank statistic, with and without missing values imputations. Like for the experiment on cruise vessels' arrivals, the intervals are slightly wider after removing pairs with missing values, but the results are similar to those including imputed concentrations.




Figure B.1.86 – 95% Fisherian Intervals based on Wilcoxon signed-rank Test Statistic for Non-Imputed Concentrations.

Daily Experiment Figure B.1.87 shows the number of pairs with missing pollutant concentrations for the daily experiment. For some pollutants, the number of pairs with missing values can represent up to 20% of the matched pairs.




Figure B.1.87 – Number of Pairs with Missing Values.

Figure B.1.88 shows the 95% Fisherian intervals calculated with the Wilcoxon signed-rank statistic, with and without missing values imputations. Like for the two other experiments, the intervals are slightly wider but the results are similar to those including imputed concentrations.




Figure B.1.88 – 95% Fisherian Intervals based on Wilcoxon signed-rank Test Statistic for Non-Imputed Concentrations.

Neymanian Inference for the Average Treatment Effect

To calculate the 95% confidence intervals with the Neymanian approach, we use the conservative sampling variance estimator of the pair-level average treatment effect found in (Imbens and Rubin, 2015), Chapter 10. If N denotes the number of matched pairs and $\hat{\tau}$ the estimate of the individual-level treatment effect, the estimator of the sampling variance writes:

$$\hat{V}^{pair}(\hat{\tau}) = \frac{1}{N(N-1)} \sum_{j=1}^N (\hat{\tau}^{pair}(j) - \hat{\tau}^{dif})^2 \quad (\text{B.2})$$

We then calculate 95% confidence intervals based on the assumption of an asymptotic normal distribution for τ :

$$CI^{0.95}(\hat{\tau}) = (\hat{\tau} - 1.96\sqrt{\hat{V}^{pair}(\hat{\tau})}, \hat{\tau} + 1.96\sqrt{\hat{V}^{pair}(\hat{\tau})}) \quad (\text{B.3})$$

In Figure B.1.89, Figure B.1.90 and Figure B.1.91, we show the 95% confidence intervals for the average treatment effect for each experiment. While they must be interpreted differently from Fisherian intervals, these confidence intervals provide results that are qualitatively similar to the 95% Fisherian intervals from our main results.




Figure B.1.89 – 95% Confidence Intervals for the Average Treatment Effect.

Hourly Experiment - Cruise Vessels' Arrivals




Figure B.1.90 – 95% Confidence Intervals for the Average Treatment Effect.

Hourly Experiment - Cruise Vessels' Departures




Figure B.1.91 – 95% Confidence Intervals for the Average Treatment Effect.

Daily Experiment

Statistical Power, Type M and S errors

In all experiments, our matching procedure resulted in few matched pairs. We may have a low statistical power. If our statistical power is low and we obtain a "statistically significant" effect, we have a higher chance that this estimate is of the wrong sign (Type S error) and/or overestimates the true effect of vessel traffic on pollution concentrations (Type M error). Here we carry out a retrospective design analysis (Gelman and Carlin,

2014; Gelman et al., 2020) to estimate the risk of making a type-S and type-M errors. While we do not know what the true effect of cruise vessels on air pollutants is, we can calculate the statistical power and probability to make type S and M errors under a set of plausible effect sizes. For each experiment, we calculate the standard error of the average treatment effect of cruise vessel traffic on NO₂ concentrations in *t* at Saint-Louis, and we compute the statistical power, type M and type S errors under plausible effect sizes using the `retrodesign` R package (Timm, 2019). Given that the matched sample sizes of the two hourly experiments are close, the results of the design analysis are also close across the two experiments and we only show the results for cruise vessels' arrivals.

Hourly Experiments In Figure B.1.92, we plot the statistical power, the exaggeration factor and the probability to make a type S error. First, we see that the average effect of cruise vessels on concentrations should be above 5 µg/m³ to reach a 80% statistical power. This would represent at least a 14% increase in the average hourly NO₂ concentration. For lower effect sizes, statistical power decreases and, conversely, we have a higher chance to overestimate the true contribution of vessels. For instance, for a true effect size of 2 µg/m³, we would overestimate the true effect by a factor just over 2. For effect sizes lower than 2 µg/m³, the chance of getting a "statistically significant" estimate of the wrong sign increases exponentially.




Figure B.1.92 – Statistical Power, Type M and S errors for the hourly experiment on cruise vessels' arrivals.

Daily Experiment In Figure B.1.93, we plot similar curves for the daily experiment. The effect size at which statistical power reaches 80% is around 4.5 µg/m³. This would represent at least a 13% increase in average daily NO₂ concentrations. For true effect sizes below 1 µg/m³, the risk to make a Type S sharply increases.




Figure B.1.93 – Statistical Power, Type M and S errors.

B.1.8 Regression Analysis of the Initial Data

Finally, we compare our results to estimates found using a simple multivariate regression on the initial dataset, made of the .

Hourly Experiments

For the two hourly experiments, we run the following model:

$$p_{t+j} = \alpha + \beta W_t + \mathbf{X}_t \gamma + \mathbf{C}_t \theta + \epsilon_t \quad (\text{B.4})$$

where:

- j is the index of the lag or lead
- t is either the hour (for the two hourly experiment) or the day index (for the daily experiment).
- p_t is the concentration of an air pollutant p .
- w_t is the treatment indicator.
- \mathbf{X}_t is the vector of weather covariates, which include the average temperature, the squared of the average temperature, an indicator for the occurrence of rainfall, the average humidity, the wind speed, the wind direction divided in the four principal directions (North-East, South-East, South-West, North-West).
- \mathbf{C}_t is the vector of calendar variables, which are indicators for the hour of the day (for the two hourly experiments), the day of the week, bank days, holidays, month, year and the interaction of these last two variables. ϵ_t is an error term.

We run this simple model from lag 3 to lead 3 of a pollutant. Figure B.1.94 and Figure B.1.95 show estimated coefficients for the treatment effects with their associated 95% confidence intervals for the two hourly experiments. Overall, the point estimates and confidence intervals obtained with the regression analysis are smaller than the point estimate and Fisherian intervals based on the parametric matching and a pruned dataset. Point estimates are also often of opposite sign.

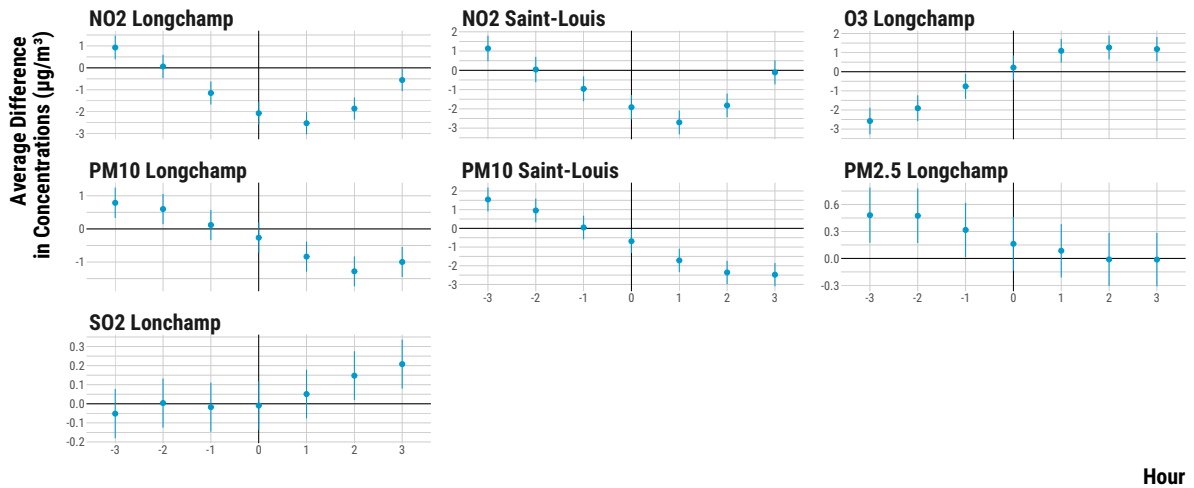


Figure B.1.94 – Regression Analysis of the Full Dataset.

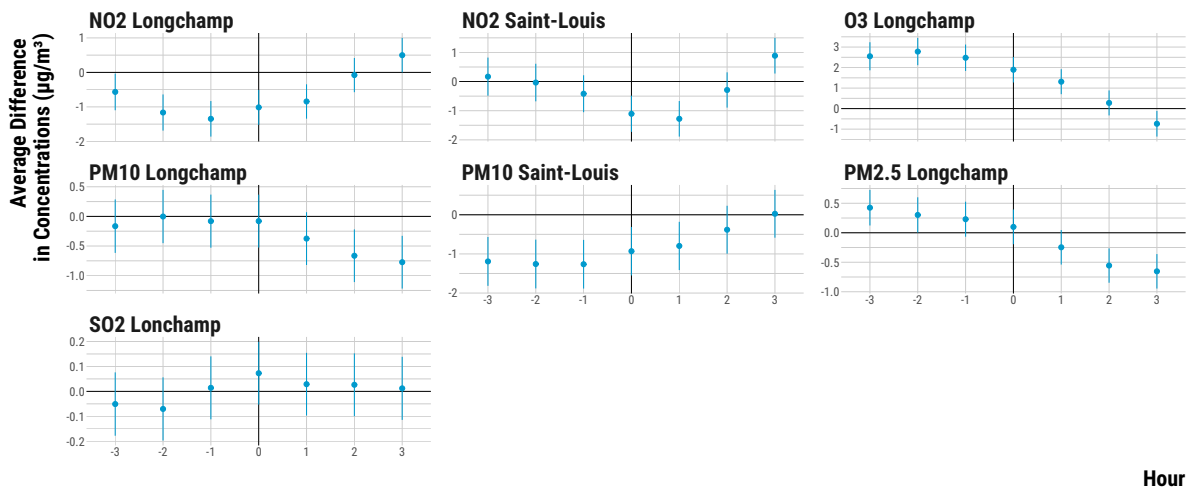


Figure B.1.95 – Regression Analysis of the Full Dataset.

Daily Experiment

For the daily experiment, we cannot run the regression model based on the binary treatment indicator, because the treatment is not meaningful without the exact pair matching:

without matching on the vessel traffic at $t-1$, the level of vessel traffic for treated units at t is not necessarily higher than that of control units. Therefore, we run the regression on the initial dataset of potential treated and potential control units ($N=2,136$) taking as a dependent variable the number of cruise vessels entering the port at t . The estimated coefficients of interest will measure the average impact on pollution of an increase by one in the number of cruise vessels entering the port. In that sense, the magnitude of the coefficient could be directly compared to the point estimate of the matched experiment, as they measure comparable effects.

We run the following model for $j = -1$, $j = 0$ and $j = 1$ (lag 1 to lead 1):

$$p_{t+j} = \alpha + \beta N_t + \mu V_t + \mathbf{X}_t \gamma + \mathbf{X}_{t-1} \delta + \mathbf{C}_t \theta + \epsilon_t \quad (\text{B.5})$$

where:

- j is the index of the lag or lead
- t is either the hour (for the two hourly experiment) or the day index (for the daily experiment).
- p_t is the daily concentration of an air pollutant p .
- N_t is the number of cruise vessel entering the port.
- β is the coefficient of interest
- V_t is the total tonnage of all types of vessel at $t - 1$.
- \mathbf{X}_t is the vector of weather covariates, which include the average temperature, the squared of the average temperature, an indicator for the occurrence of rainfall, the average humidity, the wind speed, the wind direction divided in the four principal directions (North-East, South-East, South-West, North-West).
- \mathbf{X}_{t-1} is the vector of the same weather covariates at $t - 1$
- \mathbf{C}_t is the vector of calendar variables, which are indicators for the day of the week, bank days at t and $t - 1$, holidays at t , month, year and the interaction of month and year. ϵ_t is an error term.

Figure B.1.96 plots the estimated $\hat{\beta}$, and Table B.1.3 shows the regression coefficients for all variables at t only. Generally speaking, the point estimates are lower in magnitude and more precisely estimated than in the matched pair experiment. Like in the daily experiment, we find than an additional cruise vessel entering the port is predicted to increase SO_2 concentrations at t , but the point estimate is much lower, at $0.17 \mu\text{g}/\text{m}^3$

(95% CI: [0.02,0.31]). We also see an increasing trend for PM_{10} and $PM_{2.5}$ measured at Longchamp, but the effect is not statistically different from 0, while there was a rather large impact in the paired experiment.

In contrast to the paired experiment, we detect an impact on NO_2 measured at both Saint Louis and Longchamp at t , with a point estimate 3 to 4 times lower than that of the paired experiment. The effect is short-lived as the point estimate is close to zero at $t + 1$. A similar pattern is visible for PM_{10} measured at Saint Louis, while there seemed to be an increase in NO_2 and PM_{10} at $t + 1$ in the paired experiment.

All in all, the directions of the effect are consistent between the regression analysis and the matched pair experiment, but the matched pair experiment reveals higher and less precisely estimated point estimates, as well as a lagged effect for particulate matter concentrations which is not found at Saint Louis station in the regression analysis.




Figure B.1.96 – Regression Analysis of the Full Dataset.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	NO ₂ SL	NO ₂ L	SO ₂ L	PM ₁₀ SL	PM ₁₀ L	PM _{2.5} SL	O ₃
Nb. cruise vessels arrivals t	0.881*** (0.305)	0.548** (0.257)	0.165** (0.073)	0.411 (0.401)	0.246 (0.304)	0.354 (0.220)	-0.625* (0.377)
Total tonnage t-1	0.011 (0.151)	-0.141 (0.127)	-0.034 (0.036)	-0.038 (0.198)	-0.061 (0.150)	-0.072 (0.109)	0.291 (0.187)
Temp t	-1.179*** (0.237)	-1.827*** (0.200)	-0.010 (0.057)	-1.004*** (0.312)	-2.061*** (0.236)	-2.016*** (0.171)	1.858*** (0.293)
Temp ² t	0.024*** (0.009)	0.041*** (0.007)	0.004* (0.002)	0.038*** (0.011)	0.068*** (0.009)	0.060*** (0.006)	-0.012 (0.011)
Temp t - 1	0.733*** (0.234)	0.406** (0.197)	0.057 (0.056)	-1.265*** (0.307)	-0.495** (0.233)	-0.472*** (0.169)	-0.098 (0.289)
Temp ² t-1	-0.025*** (0.009)	-0.015** (0.007)	-0.004** (0.002)	0.046*** (0.011)	0.023*** (0.009)	0.016*** (0.006)	0.0002 (0.011)
Rainfall Dummy t	-1.115** (0.494)	-1.449*** (0.416)	0.055 (0.118)	-3.595*** (0.649)	-3.340*** (0.491)	-2.290*** (0.356)	2.838*** (0.611)
Rainfall Dummy t-1	-1.416*** (0.500)	-0.744* (0.421)	-0.137 (0.120)	-5.431*** (0.657)	-4.341*** (0.497)	-2.891*** (0.361)	0.824 (0.618)
Humidity Average t	-0.163*** (0.031)	-0.103*** (0.026)	-0.036*** (0.007)	0.105*** (0.040)	0.146*** (0.031)	0.138*** (0.022)	-0.490*** (0.038)
Humidity Average t-1	-0.156*** (0.027)	-0.075*** (0.023)	-0.007 (0.007)	-0.025 (0.036)	-0.014 (0.027)	-0.002 (0.020)	-0.197*** (0.034)
S-E Wind t	1.553* (0.813)	-0.813 (0.685)	-0.212 (0.195)	-0.785 (1.069)	0.135 (0.810)	0.109 (0.587)	2.340** (1.006)
S-W Wind t	7.085*** (0.848)	3.352*** (0.714)	0.703*** (0.203)	-1.981* (1.115)	-0.047 (0.844)	-0.004 (0.612)	-0.280 (1.049)
N-W Wind t	1.540* (0.806)	5.928*** (0.679)	0.799*** (0.193)	-4.274*** (1.059)	0.623 (0.802)	0.755 (0.582)	-8.260*** (0.997)
S-E Wind t-1	-0.227 (0.815)	-1.598** (0.686)	-0.198 (0.195)	0.815 (1.071)	-0.215 (0.811)	-0.383 (0.588)	2.118** (1.008)
S-W Wind t-1	2.532*** (0.867)	-0.399 (0.730)	0.209 (0.207)	0.203 (1.139)	-1.448* (0.863)	-0.841 (0.625)	1.024 (1.072)
N-W Wind t-1	1.760** (0.818)	1.082 (0.689)	0.255 (0.196)	0.720 (1.075)	-0.552 (0.814)	-0.505 (0.590)	-2.651*** (1.012)
Wind speed t	-3.871*** (0.090)	-3.191*** (0.076)	-0.205*** (0.022)	-1.588*** (0.118)	-1.033*** (0.090)	-1.051*** (0.065)	2.003*** (0.111)
Wind speed t-1	-0.742*** (0.090)	-0.427*** (0.076)	-0.027 (0.022)	-0.798*** (0.119)	-0.547*** (0.090)	-0.410*** (0.065)	-0.412*** (0.112)
Holiday t	-1.503*** (0.504)	-0.529 (0.425)	0.047 (0.121)	-0.464 (0.663)	0.736 (0.502)	0.191 (0.364)	-0.371 (0.624)
Bank day t	-8.619*** (1.072)	-8.258*** (0.902)	-0.022 (0.256)	-5.352*** (1.409)	-2.584** (1.067)	-0.396 (0.773)	5.687*** (1.326)
Bank day t-1	-1.967* (1.057)	-2.304*** (0.889)	-0.225 (0.253)	-3.249** (1.389)	-2.332** (1.052)	-0.491 (0.762)	2.709** (1.307)
Constant	80.383*** (2.973)	74.872*** (2.503)	6.628*** (0.711)	52.909*** (3.908)	43.234*** (2.959)	36.314*** (2.145)	65.709*** (3.678)
Observations	2,135	2,135	2,135	2,135	2,135	2,135	2,135
R ²	0.778	0.775	0.399	0.588	0.528	0.609	0.831

Notes: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

All regressions include day of the week, month, and month \times year fixed effects

Table B.1.3 – Regression coefficients - Daily experiment

B.1.9 Road Traffic and NO₂ Concentrations

Figure B.1.97 shows that road traffic data have a strong weekly pattern, with stable and high levels of traffic on weekdays and lower levels of traffic on weekends. On average, there are 480 fewer vehicles per hour per road on Marseille main roads on weekdays compared to weekends. At the same time, NO₂ concentrations also decrease on weekdays (Figure B.1.98) while the total gross tonnage of vessel traffic (including all types of vessels) increases (Figure B.1.99). On average, NO₂ concentrations are lower by 7.5 µg/m³ at Longchamp and by 8.2 µg/m³ at Saint-Louis on weekdays compared to weekends. Since weekdays and weekends are very balanced in terms of weather covariates, as shown on Figure B.1.100, the decrease in NO₂ over the weekend cannot be due to different weather patterns. Although all economic activities are lower on weekends compared to weekdays, not just traffic, road traffic is known to be the main factor affecting city-level NO₂. We can thus consider that the decrease in NO₂ is the combined effect of changes in vessel traffic and changes in road traffic during that period. Finding that NO₂ concentrations decrease substantially on weekends may indicate that road traffic has a much higher responsibility in changes in ambient concentrations of NO₂ than maritime traffic.




Figure B.1.97 – Average Road Traffic Flow by Day of the Week.




Figure B.1.98 – Average NO₂ Concentrations over the Week.





Figure B.1.99 – Average Vessel Gross tonnage over the Week.

(a)



(b)

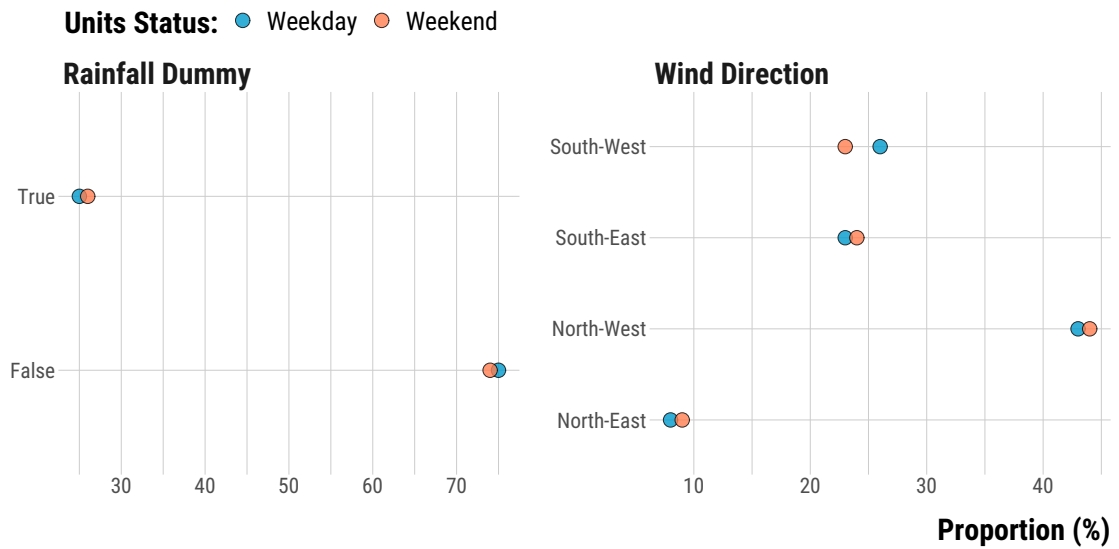


Figure B.1.100 – Weather Covariates Balance.

Notes: Panel A plots the density distribution and the boxplot of continuous weather covariates for weekends and weekdays. Panel B displays the average proportion of categorical weather covariates' observations for weekends and weekdays.

Appendix C

Appendix to Chapter 3: Tackling Transport-Induced Pollution in Cities: A Case Study in Paris

C.1 Appendix

C.1.1 Assumptions on NO_x, PM_{2.5} and CO₂ emissions by transport mode

For “polluting” modes (buses, cars, two-wheelers), the emission factor $e_{P,m}$ comes from different sources.

Buses For buses, the NO_x and PM_{2.5} emission factors per passenger are derived from the local air quality agency’s emission calculator¹. They give an emission factor of 180mg/km for an average bus in 2017. The average bus in France is 7.7 years old (Source: Observatoire de la mobilité), so the value for 2017 is for buses registered in 2009 on average. Assuming that the age of the fleet was the same in 2010, the average bus taken by the surveyed individuals in 2010 had been registered in 2002. We adjust for the difference in the years of the data by multiplying the Airparif bus emission factor for 2017 by the ratio of NO_x and PM_{2.5} emission factors for cars registered in 2002 compared to 2010, assuming that the improvement in emission factors was similar for buses and for cars over the period.

The CO₂ emission factor per passenger is derived from national values given in Ministère de la Transition écologique et solidaire (2018) and scaled down to adjust for the higher average number of passengers in IdF compared to other regions. The initial value

¹<http://www.airparif.fr/calculateur-emissions/>. Although the value given for particulate matter indicate a value in particulate matter of size below 10 microns (PM₁₀), most particles from engine combustion are actually smaller than 2.5µm: Karjalainen et al. (2014) mention that most exhaust particles from gasoline direct injection engines are around 0.1µm; California Air Resources Board (2021) mention that more than 90% of diesel particulate matter is less than 1µm in diameter. The EMEP/EEA Copert methodology from which Airparif emission factors are calculated also assumes that all PM from exhaust are PM_{2.5} (Ntziachristos and Zissis, 2020). A personal communication with the agency confirms that we can interpret the PM₁₀ emission factors as PM_{2.5}.

assumes 11 passengers by bus on average. Traffic data from the regional transport authority give an average of 14 passengers by bus in Ile de France, so we multiply the initial factor by 11/14.

Cars and two-wheelers owned by the household For two-wheelers and cars, the vehicle used is a vehicle owned by the household in 89% of the cases. We estimate the NO_x, PM_{2.5} and CO₂ emission factors of these vehicles based on their characteristics reported in the survey. For the NO_x and PM_{2.5} emission factors of cars, we use the information on the type of car (passenger car/LCV), the year of first registration and the fuel type. For the CO₂ emission factors of cars, we also use the information on the car's horsepower; For the NO_x and PM_{2.5} emission factors of two-wheelers, we use the year of first registration only, while for the CO₂ emission factor of two-wheelers, we also use the fuel type and type of two-wheeler (e.g, moped versus motorbike).

For cars, we use the NO_x and PM_{2.5} emission factors from the local air quality agency's emission calculator by type of fuel and date of registration of the car. The average speed, cold starts and horsepower of vehicles circulating in IdF are included as common parameters entering the calculation of emission factors for all fuel types and dates of registration. Regarding fuel type, the calculator distinguishes between diesel, gasoline, and electric cars. We assign LPG cars from the survey the same emission factor as a gasoline car from the same year. We assign hybrid cars from the survey the same emission factor as an electric car from the same year (this may underestimate emissions from hybrid cars, but they represent only 0.3% of the cars owned by households). The calculator does not have specific values for light-commercial vehicles. For these car types declared by the household, we proceed as follows: we take the emission factors for LCVs and cars from a different source, the Ominea database edited by an environmental agency called Citepa and giving reference values for emission factors for different economic sectors². We calculate the ratio of LCVs to car emission factors according to that source for each type of car and LCV defined by their fuel type and registration year (and taking the value for the "urban driving conditions" rather than "highway" or "rural"). We then multiply the NO_x and PM_{2.5} emission factors given for cars in the Airparif calculator by the OMINEA ratio, and obtain NO_x and PM_{2.5} emission factors respecting the relative difference of LCVs vs cars given in the OMINEA database. Particulate matter emissions from tyres and brakes are not taken into account in the OMINEA data, so we are assuming that the ratio of PM_{2.5} emission factors for LDVs over cars is the same for exhaust emissions and emissions from brakes and tyres.

For CO₂, we use data from the French Energy Agency (Ademe), which provides emis-

²<https://www.citepa.org/fr/omine/>

sion factors for all car models from 2001 to 2015. We build categories of car models defined by the same information as the one we have on the cars owned by households in the EGT data: year, fuel type (gasoline/petroleum/hybrid/electric/LPG), and administrative horsepower. Then, we calculate for each category the average CO₂ emission factor from the Ademe dataset, weighted by national-level market shares by brand³. We allocate to each car type from the EGT data the CO₂ emission factor from Ademe for the same car category. When the car owned by the household is older than 2001, we rely on data provided by Ademe⁴ giving average emission factors of cars sold in France by fuel type, for the years 1995-2018. We estimate emission factors for the period before 1995 by applying the same annual trend for emissions as for the 1995-2000 period. For electric cars, we assign a zero emission factor. The Ademe data reports emission factors for commercial vehicles only. For light-commercial vehicles owned by the household, we use the estimations given in CGDD (2011).

For two-wheelers, we use the NO_x and PM_{2.5} emission factors from the local air quality agency's emission calculator, scaled up to reflect 2010 values rather than 2019 ones. We apply the CO₂ emission factors from Barbusse (2005), which are differentiated by fuel type and by type of two-wheeler. The study dates back 2005 and the emissions are calculated for motorcycles first registered between 2003 and 2005. But this is a relatively good proxy for the median emission factor of the motorcycles owned by EGT households, which median first registration date is 2005. This single emission factor does not allow to reflect the heterogeneity in the registration year (from 1951 to 2011), but we do not think it is too much an issue given the low modal share of two-wheelers (< 1%).

Taxis and cars and two-wheelers not owned by the household When the vehicle used is a car not owned by the household or is a taxi, we impute the NO_x and PM_{2.5} emission factors of a 2008 diesel car (in 2010 most taxis were diesel vehicles⁵). We impute the CO₂ emission factor of a 2008 diesel car of 7 hp. We take values for recent vehicles because vehicles not owned by the household are likely to be company cars, which are often relatively new. For taxis, we multiply the emission factor by two to account for the fares driven without passengers, following the recommendations of Ministère de la Transition écologique et solidaire (2018). When the vehicle used is a two-wheeler not owned by the household, we impute the NO_x and PM_{2.5} emission factors of a Euro 3 two-wheeler from the Airparif calculator, and the CO₂ emission factor from a scooter. Table C.1.1 shows the unique emission factor obtained for buses, taxis, cars and two-wheelers not owned by

³we take the average of the registration market shares over the years 2000, 2005 and 2010 obtained from the French car manufacturer's association CFA.

⁴<http://carlabelling.ademe.fr/chiffrescles/r/evolutionTauxCo2>

⁵<https://www.auto-moto.com/actualite/environnement/faut-il-interdire-les-taxis-diesels-la-question-qui-fache-49587.html>

Transport mode	Unit	NOx (mg/km)	PM _{2.5} (mg/km)	CO ₂ (g/km)
Taxi	per passenger	1,178	127	332
Car not owned by the household	per vehicle	589	63	166
Two-wheeler not owned by the household	per vehicle	86	21	65

Table C.1.1 – Emission factors for private vehicles not owned by the household

Note: Authors’ calculations from Airparif, OMINEA, Ministère de la Transition écologique et solidaire (2018), Copert, Ademe

the household (here assuming one passenger per vehicle).




Figure C.1.1 – Distribution of NOx emissions per passenger, by transportation mode

Note: For each transportation mode, the box plot shows the distribution of NOx emissions across journey stages using this transportation mode. Call Q1 the 25th percentile, Q3 the 75th percentile, and IQR the interquartile range. The bar in each box shows the median value, the lower and upper hinges of the box respectively show Q1 and Q3, and the lower and upper lines show the lower and upper adjacent values defined at $Q1 - 1.5 \times IQR$ for the lower adjacent value, and $Q3 + 1.5 \times IQR$ for the upper adjacent value.

C.1.2 Method to retrieve counterfactual transport time with Google API

We pool together the trips likely to have exactly same duration based on Google’s prediction algorithm: for transit trips during the day (from 6 am to 9h59 pm) and cycling trips, changing the direction of the trip or its hour did not change the resulting duration based on a trial on a few trips. So we grouped together all the trips with the same or the




Figure C.1.2 – Distribution of $PM_{2.5}$ emissions per passenger, by transportation mode

Note: For each transportation mode, the box plot shows the distribution of $PM_{2.5}$ emissions across journey stages using this transportation mode. Call Q1 the 25th percentile, Q3 the 75th percentile, and IQR the interquartile range. The bar in each box shows the median value, the lower and upper hinges of the box respectively show Q1 and Q3, and the lower and upper lines show the lower and upper adjacent values defined at $Q1 - 1.5 \times IQR$ for the lower adjacent value, and $Q3 + 1.5 \times IQR$ for the upper adjacent value.

opposite point of departure and point of arrival, irrespective of the hour of departure. We are left with 49,242 trips with unique pair of {departure;origin}. We simulate day transit and all cycling trips so that they occur on a Tuesday morning. For driving trips, average traffic conditions are integrated in the algorithm, such that the hour of the trip and the direction of the flow can influence the trip duration. We group together trips with the same hour of departure, point of departure and point of arrival. We are left with 73,264 trips with unique point of departure X point of arrival X hour of departure. We simulate transit trips so that they occur on a Tuesday. Finally, we account for the fact that public transport is less frequent at night by estimating specific trip duration for public transport at night. For transit trips during the night (from 10 pm to 05h59 am), we group together trips the same way as for car trips. We are left with 2,844 trips. We simulate night transit trips so that they occur on a Monday evening.

A comparison of Google Maps' trip duration output and the trip durations self-reported by individuals in the EGT data reveal that Google Map's durations are lower for all the three modes: cycling trips are on average 39% shorter according to Google Maps (but this is based on a very small sample of cycling trips in the EGT), driving trips 32% shorter, and transit trips 20% shorter (the comparison is made for trips actu-




Figure C.1.3 – Distribution of CO₂ emissions per passenger, by transportation mode

Note: For each transportation mode, the box plot shows the distribution of CO₂ emissions across journeys using this transportation mode. Call Q1 the 25th percentile, Q3 the 75th percentile, and IQR the interquartile range. The bar in each box shows the median value, the lower and upper hinges of the box respectively show Q1 and Q3, and the lower and upper lines show the lower and upper adjacent values defined at $Q1 - 1.5 \times IQR$ for the lower adjacent value, and $Q3 + 1.5 \times IQR$ for the upper adjacent value.

ally using that mode in our data). Given the potential error in self-reported durations, the uncertainty margin of the API's estimations and the ten-year gap between the API request (2020) and the EGT data (2010), it is difficult to know which one is the true value, if any. What matters for us is that the relative time difference derived from the API's predictions for car, cycling and public transit trips reflects well the true relative difference in time. Given the higher discrepancy for cycling compared to driving and the lower one for transit compared to driving, we may underestimate the ability with which individuals switch from car to public transport and slightly overestimate the ability with which individuals switch from car to cycling.

C.1.3 Estimation of journey stage-specific emissions accounting for cold starts

Under the EMEP/EEA method explained in Ntziachristos and Zissis (2020), for each journey stage by car j made by individual i , NO_x emissions can be calculated as the sum of hot and cold exhaust emissions:

$$E_{NOx,i,j} = E_{NOx,i,j}^{hot} + E_{NOx,i,j}^{cold} \quad (C.1)$$

And PM_{2.5} emissions can be calculated as the sum of hot and cold exhaust emissions, plus emissions from tyre and brake wear, plus emissions from road surface wear.

$$E_{PM2.5,i,j} = E_{PM2.5,i,j}^{hot} + E_{PM2.5,i,j}^{cold} + E_{PM2.5,i,j}^{tyrebrake} + E_{PM2.5,i,j}^{roadsurf} \quad (C.2)$$

The amount of hot and cold emissions depends on the fraction of the distance driven with a cold engine. According to the EMEP/EEA guidance, this fraction is a function of the vehicle fuel, euro norm, trip distance and exterior temperature. To simplify, we instead set that the first 8 minutes of the trip are made with a cold engine, which is the assumption made by Airparif to calculate their average emission factors⁶.

Assuming that β is the share of the trip made with a cold engine, drawing on equation 3.2 we have:

$$E_{NOx,i,j} = d_{j,i}((1 - \beta)e_{NOx,j,i}^{hot} + \beta e_{NOx,j,i}^{cold})r_{j,i} \quad (C.3)$$

and

$$E_{PM2.5,i,j} = d_{j,i}((1 - \beta)e_{PM2.5,j,i}^{hot} + \beta e_{PM2.5,j,i}^{cold} + e_{PM2.5,j,i}^{tyrebrake} + e_{PM2.5,j,i}^{roadsurf})r_{j,i} \quad (C.4)$$

Noting $t_{j,i}$ the duration of the journey stage⁷, we set β to the maximum between 100% and $8/t_{j,i}$ to reflect that the first 8 minutes are made with a cold engine.

For $e_{NOx,j,i}^{hot}$ and $e_{PM2.5,j,i}^{hot}$, we use the EMEP-EEA values⁸ available by fuel, vehicle type and euro norm category. For the fuel type, vehicle type and euronorm, we take the characteristics of the vehicle taken for the trip (since we restrict the analysis to vehicles owned by the household, we have the vehicle characteristics). For each fuel \times vehicle type \times euro norm category, the emission factor depends on the average trip speed. We take a single value of 30km/h, corresponding to the average car trip speed in the EGT survey, derived from the travel time and distance results given by the Google console.

⁶This assumption is consistent with what is obtained in Ntziachristos and Zissis (2020) for a 10 kilometre trip with a diesel or old petrol car at an average speed of 25km/hour.

⁷The duration of each journey stage is not readily available in the EGT, which only gives the self-declared duration of the trip. More than 99% of the journey stages by car are included in 3-stage trips, with a first journey stage by foot, a second journey stage by car and a last journey stage by foot. We retrieve the duration of the journey stage by car assuming a walking speed of 4km/h and taking the difference between the trip's duration and the walking duration.

⁸accessible here: https://www.eea.europa.eu/publications/emep-eea-guidebook-2019/part-b-sectoral-guidance-chapters/1-energy/1-a-combustion/road-transport-appendix-4-emission/at_download/file

For $e_{PM_{2.5},j,i}^{tyrebrake}$ and $e_{PM_{2.5},j,i}^{roadsurf}$, we use the EMEP-EEA values from Ntziachristos and Boulter (2019), available by type of vehicle (passenger cars/light-duty trucks/heavy-duty trucks).

To obtain $e_{PM_{2.5},j,i}^{cold}$ and $e_{NOx,j,i}^{cold}$, we use the formula given in Ntziachristos and Zissis (2020) to calculate the ratios $e_{PM_{2.5},j,i}^{cold}/e_{PM_{2.5},j,i}^{hot}$ and $e_{NOx,j,i}^{cold}/e_{NOx,j,i}^{hot}$ (call the generic ratio e^{cold}/e^{hot} , and multiply this ratio by the hot emission factor. The ratio e^{cold}/e^{hot} depends on the pollutant, fuel type, vehicle type, euro norm, average outdoor temperature and average trip speed. For the fuel type, vehicle type and euronorm, we take the characteristics of the vehicle taken for the trip. We take 11.7°C as the average temperature, which is the average annual temperature for the IdF region.⁹ For the average speed, we do not estimate each journey stage's speed. Instead, we allocated to a journey stage the average speed of its origin-destination category, derived from the travel time and distance results given by the Google console. Trips starting and finishing in Paris have an average speed of 15km/h. Trips in Paris and the inner suburbs outside the Paris-Paris trips have an average speed of 22km/h. Trips starting and finishing in the outer suburbs have an average speed of 33km/h. Finally, trips starting or finishing in the outer suburbs and finishing or starting in Paris or the inner suburbs have an average speed of 40km/h. Note that $e_{PM_{2.5},j,i}^{cold}$ is available only for diesel cars and is assumed to be zero for petrol cars in Ntziachristos and Zissis (2020).

We re-estimate individual and trip-level NOx and PM_{2.5} emissions based on this alternative emission estimations for all the journey stages made with a vehicle owned by the household, and based on the method from section 3.3.2 for all the other modes.

C.1.4 Additional tables and figures

⁹Source: <https://fr.climate-data.org/europe/france/ile-de-france-301/>




Figure C.1.4 – Evolution of annual average NO₂ concentrations in IdF

Notes: Source: Airparif 2019 annual report. Concentrations are in $\mu\text{g}/\text{m}^3$. The legal threshold of $40 \mu\text{g}/\text{m}^3$ appears in red on the scale. It is also the threshold recommended by the WHO.




Figure C.1.5 – Evolution of annual average PM_{2.5} concentrations in IdF

Note: Source: Airparif 2019 annual report. Concentrations are in $\mu\text{g}/\text{m}^3$. The legal threshold is $25\mu\text{g}/\text{m}^3$, appearing in red on the scale. The threshold recommended by the WHO is $10\mu\text{g}/\text{m}^3$.

	Mean	Sd
Nb. household members	2.33	1.38
Residence: Paris	0.23	0.42
inner suburbs	0.37	0.48
outer suburbs	0.40	0.49
Housing: Social housing	0.23	0.42
Private tenants	0.23	0.42
Home-owners	0.51	0.50
Other housing status	0.03	0.17
Age, person of reference	49.58	15.98
Estimated Net income	37,571.06	24,535.46
Estimated Net income per consumption unit	24,655.83	14,640.12
Observations	14,882	

Table C.1.2 – EGT-Descriptive statistics at the household level

Note: Source: EGT data. Data weighted with EGT household-level sampling weights

	EGT	Administrative data
Nb. household members	2.33 (1.38)	2.48 (1.68)
Residence: Share living in Paris	0.23 (0.42)	0.22 (0.42)
Share living in the inner suburbs	0.37 (0.48)	0.37 (0.48)
Share living in the outer suburbs (%)	0.40 (0.49)	0.41 (0.49)
Share living in Social housing (%)	0.23 (0.42)	0.22 (0.41)
Housing: Share of private tenants	0.23 (0.42)	0.26 (0.44)
Share of home-owners	0.51 (0.50)	0.49 (0.50)
Share of other housing status	0.03 (0.17)	0.03 (0.18)
Age, person of reference	49.58 (15.98)	52.04 (17.10)
Net income per consumption unit	24,655.83* (14,640.12)	25,969.40** (85,486.92)
Observations	14,882	4,830,037

Table C.1.3 – Balance between EGT survey data and administrative data on selected household characteristics


Note: EGT observations weighted with household-level sampling weights. Source for the administrative data: Filocom data for 2011, an exhaustive census of housing units by January 1st 2011. *The income variable from EGT has been imputed using an interval regression imputation method. **The income variable from Filocom comes from fiscal sources and does not include non-taxable income sources such as housing or family benefits.

	1	2	3
	Individuals travelling	Full sample	(1)-(2)
Residence: Share living in Paris	0.143 (0.350)	0.140 (0.347)	0.00340
Share living in the inner suburbs	0.366 (0.482)	0.365 (0.482)	0.00107
Share living in the outer suburbs	0.490 (0.500)	0.495 (0.500)	-0.00447
Age, person of reference	45.20 (16.21)	45.69 (16.64)	-0.496***
Education: Primary school	0.0514 (0.221)	0.0588 (0.235)	-0.00735***
Secondary education	0.393 (0.488)	0.400 (0.490)	-0.00725
Higher education \leq 3 years	0.152 (0.359)	0.149 (0.356)	0.00248
Higher education $>$ 3 years	0.337 (0.473)	0.326 (0.469)	0.0104*
Still in education	0.0671 (0.250)	0.0654 (0.247)	0.00177
Socioprofessional category: Farmers	0.000756 (0.0275)	0.000711 (0.0267)	0.0000449
Manual workers	0.105 (0.307)	0.104 (0.306)	0.00102
Office workers	0.191 (0.393)	0.192 (0.394)	-0.000486
Intermediate professions	0.220 (0.414)	0.214 (0.410)	0.00628
Traders and craftspeople	0.0200 (0.140)	0.0198 (0.139)	0.000140
Managers and executives	0.197 (0.398)	0.190 (0.392)	0.00723*
Pensioner	0.198 (0.398)	0.213 (0.410)	-0.0155***
Other	0.0681 (0.252)	0.0669 (0.250)	0.00124
Activity status: Pupil/Student	0.0652 (0.247)	0.0633 (0.244)	0.00192
Part-time or full-time employed	0.648 (0.478)	0.624 (0.484)	0.0241***
Unemployed	0.0532 (0.224)	0.0578 (0.233)	-0.00454*
Other inactive	0.224 (0.417)	0.242 (0.429)	-0.0185***
Pensioner	0.00985 (0.0988)	0.0128 (0.113)	-0.00298**
Estimated Net income	40613.6 (25157.6)	40036.8 (24938.5)	576.8*
Estimated Net income per consumption unit	24051.2 (14327.1)	23725.4 (14262.1)	325.8*
Distance to workplace (km)	11.78 (11.99)	11.79 (12.02)	-0.0172
Observations	23690	25453	


mean coefficients; sd in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

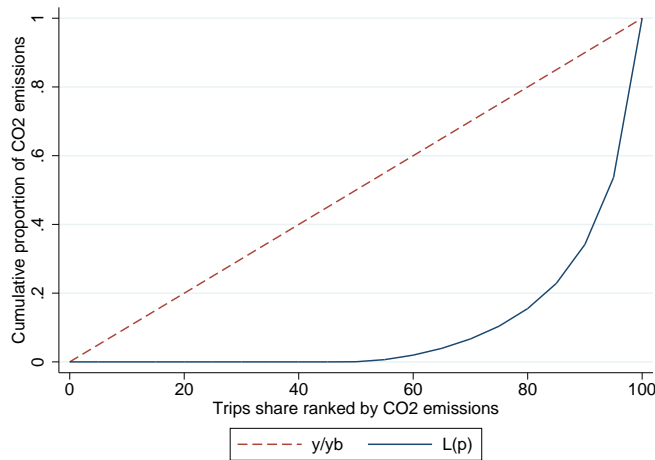
Table C.1.4 – Balancing test comparing the subsample of individuals with one trip recorder and the full sample



(a) NO_x emissions



(b) PM_{2.5} emissions



(c) CO₂ emissions

Figure C.1.6 – Lorenz curve, trip level

Note: the x-axis shows the percentiles of trip-level emissions and the y-axis shows the share of total emissions generated by all the trips below that percentile. The red dotted curve shows how the distribution would look like if everyone contributed equally to emissions Source: EGT data. Sample: all trips made by adults.

	NOx (mg)	Dist. (km)	Modal share (%)			Emission Intensity (mg/km)		
			<i>Bus</i>	Two- Wheeler	<i>Car</i>	<i>Bus</i>	Two- Wheeler	<i>Car</i>
	$E_{NOx,Qk}$	D_{Qk}	$S_{bus,Qk}$	$S_{tw,Qk}$	$S_{car,Qk}$	$I_{NOx,bus,Qk}$	$I_{NOx,tw,Qk}$	$I_{NOx,car,Qk}$
Q1	27	16.0	0.000	0.000	0.005	242	NA*	NA*
Q2	1,740	22.6	0.082	0.018	0.289	242	93.3	191.2
Q3	4,118	25.6	0.187	0.035	0.373	242	103.7	300.0
Q4	11,251	31.5	0.106	0.033	0.674	242	122.6	486.0
Q5	45,593	62.2	0.018	0.005	0.918	242	127.9	793.5

Table C.1.5 – Extended Kaya components by quintile of NOx emissions

*The modal share of these two modes is zero. In practice in the calculation of the LMDI, the same emission intensity of cars and two-wheelers as for Q3 has been imputed to Q1, such that these sub-components of modal share receive a 0 contribution to the difference compared to Q3

	NOx (mg)	Diff vs Q3 (mg)	Distance component (mg)	Modal share	Emission intensity
				component (mg)	component (mg)
	$E_{NOx,Qk}$	$\Delta E_{NOx,Q3,Qk,tot}$	$\Delta E_{NOx,Q3,Qk,D}$	$\Delta E_{NOx,Q3,Qk,S}$	$\Delta E_{NOx,Q3,Qk,I}$
Q1	27	-4,091	-303	-3,615	0.0
			(8%)	(92%)	(0%)
Q2	1,740	-2,377	-342	-1,150	-885
			(14%)	(48%)	(37%)
Q3	4,118	0	-	-	-
			-	-	-
Q4	11,251	7,134	1,422	2,886	2,826
			(20%)	(40%)	(40%)
Q5	45,593	41,475	14,215	12,295	14,965
			(34%)	(30%)	(36%)

Table C.1.6 – LMDI decomposition on NOx emissions at the individual level

	CO ₂ (g)	Dist. (km)	Modal share (%)			Emission Intensity (g/km)		
	$E_{CO_2, Qk}$	D_{Qk}	<i>Bus</i>	Two- Wheeler	<i>Car</i>	<i>Bus</i>	Two- Wheeler	<i>Car</i>
			$S_{bus, Qk}$	$S_{tw, Qk}$	$S_{car, Qk}$	$I_{CO_2, bus, Qk}$	$I_{CO_2, tw, Qk}$	$I_{CO_2, car, Qk}$
Q1	0	15.9	0.000	0.000	0.000	NA*	NA*	NA*
Q2	646.0	23.4	0.107	0.009	0.206	117	74.6	130.8
Q3	1,348	24.0	0.180	0.021	0.382	117	82.2	142.7
Q4	3,005	27.6	0.096	0.0033	0.708	117	94.0	149.1
Q5	9,810	67.2	0.023	0.019	0.908	117	104.9	158.6

Table C.1.7 – Extended Kaya components by quintile of CO₂ emissions

*The modal share of these two modes is zero. In practice in the calculation of the LMDI, the same emission intensity of cars and two-wheelers as for Q3 has been imputed to Q1, such that these sub-components of modal share receive a 0 contribution to the difference compared to Q3

	CO ₂ (g)	Diff vs Q3 (g)	Distance component(g)	Modal share component(g)	Emission intensity component(g)
	$E_{CO_2, Qk}$	$\Delta E_{CO_2, Qk-Q3, tot}$	$\Delta E_{CO_2, Qk-Q3, D}$	$\Delta E_{CO_2, Qk-Q3, S}$	$\Delta E_{CO_2, Qk-Q3, I}$
Q1	0	-1,596	-14 (1%)	-1,582 (99%)	- (0%)
Q2	646.0	-914	-33 (4%)	-798 (87%)	-84 (9%)
Q3	1,348	0	-	-	-
Q4	3,005	1,463	351 (24%)	1,016 (69%)	96 (6%)
Q5	9,810	8,134	4,711 (58%)	2,963 (36%)	460 (6%)

Table C.1.8 – LMDI decomposition CO₂ emissions at the individual level




Figure C.1.7 – Contribution of distance, modal choice and emission intensity to the differences in PM_{2.5} emissions

Note: This graph shows the difference between PM_{2.5} emissions from average individuals in quintiles 1, 2, 4 and 5 compared to the benchmark average individual in quintile 3, (total length of the bars), decomposed into differences in total distance travelled, modal shares, and the emission intensity of a given mode using the LMDI additive decomposition

	PM _{2.5} (mg)	Dist. (km)	Modal share (%)				Emission Intensity (mg/km)			
			<i>Metro</i>	<i>Bus</i>	Two- Wheeler	<i>Car</i>	<i>Metro</i>	<i>Bus</i>	Two- Wheeler	<i>Car</i>
	$E_{PM_{2.5},Qk}$	D_{Qk}	$S_{met,Qk}$	$S_{bus,Qk}$	$S_{tw,Qk}$	$S_{car,Qk}$	$I_{met,Qk}$	$I_{bus,Qk}$	$I_{tw,Qk}$	$I_{car,Qk}$
Q1	1.3	3.0	0.007	0.044	0.000	0.006	7.1	4.8	21.1	28.8
Q2	125	12.7	0.375	0.240	0.006	0.174	7.1	4.8	23.9	33.6
Q3	501	27.3	0.451	0.066	0.018	0.372	7.1	4.8	24.9	34.8
Q4	1,321	39.9	0.256	0.027	0.033	0.633	7.1	4.8	31.8	47.6
Q5	4,185	66.6	0.056	0.006	0.019	0.899	7.1	4.8	43.3	68.6

Table C.1.9 – Extended Kaya components by quintile of PM_{2.5} emissions

	PM_{2.5} (mg)	Diff vs Q3 (mg)	Distance component(mg)	Modal share component (mg)	Emission intensity component (mg)
	$E_{PM_{2.5},Qk}$	$\Delta E_{PM_{2.5},Q3,Qk,tot}$	$\Delta E_{PM_{2.5},Q3,Qk,D}$	$\Delta E_{PM_{2.5},Q3,Qk,S}$	$\Delta E_{PM_{2.5},Q3,Qk,I}$
Q1	1.3	-499	-169 (34%)	-312 (63%)	-18 (4%)
Q2	125	-376	-202 (54%)	-146 (39%)	-28 (7%)
Q3	501	0	-	-	-
Q4	1,321	820	316 (39%)	350 (43%)	154 (19%)
Q5	4,185	3,684	1,483 (40%)	1,284 (35%)	917 (25%)

Table C.1.10 – LMDI decomposition on PM_{2.5} emissions at the individual level

	(1)	(2)	(3)	(4)	(5)	(6)
	ln dist, all	uses car, all	uses car, all	NOx/km	PM25/km	CO ₂ /km
Outer suburbs	0.190*** (0.0201)	0.138*** (0.00972)	0.0773*** (0.00850)	-16.44 (19.71)	0.467 (0.839)	-4.792*** (1.018)
Public transport stop	-0.115*** (0.0198)	-0.156*** (0.00955)	-0.119*** (0.00834)	4.205 (16.77)	-0.261 (0.784)	0.707 (0.916)
Motorized vehicle at hand			0.476*** (0.0105)			
ln Commuting distance	0.528*** (0.0111)	0.0310*** (0.00352)	0.0228*** (0.00317)	11.71 (8.283)	2.303*** (0.347)	1.191** (0.414)
Res: Paris, Work: Paris	-0.159*** (0.0342)	-0.354*** (0.0172)	-0.239*** (0.0201)	169.0 (92.02)	14.46*** (3.141)	34.57*** (6.128)
Res: Paris, Work: Suburbs	-0.0493 (0.0331)	-0.256*** (0.0185)	-0.133*** (0.0192)	-52.27 (27.95)	0.679 (2.077)	4.046 (3.245)
Res: Suburbs, Work: Paris	-0.0393 (0.0216)	-0.236*** (0.0108)	-0.196*** (0.0105)	-31.58 (21.29)	-4.534*** (1.081)	0.905 (1.466)
D1	-0.0192 (0.0483)	-0.146*** (0.0228)	-0.0291 (0.0209)	74.24 (65.49)	3.444 (2.555)	1.639 (2.727)
D2	0.00162 (0.0363)	-0.106*** (0.0173)	-0.0264 (0.0149)	0.386 (29.94)	7.057*** (1.939)	2.759 (1.845)
D9	0.00679 (0.0259)	0.0352** (0.0121)	0.00665 (0.0117)	16.64 (29.42)	-0.797 (1.135)	0.959 (1.583)
D10	-0.0122 (0.0272)	0.0416** (0.0131)	-0.00242 (0.0123)	-10.55 (29.03)	-1.654 (1.144)	4.899** (1.710)
Work in Factory	-0.0173 (0.0358)	0.0957*** (0.0165)	0.0723*** (0.0135)	48.39 (40.11)	2.301 (1.587)	1.502 (1.754)
Work at individuals' home	0.169* (0.0797)	-0.0610 (0.0338)	-0.00591 (0.0302)	472.1* (196.8)	9.850* (4.477)	7.225 (5.058)
Work from home	1.208*** (0.138)	0.102** (0.0385)	0.0838* (0.0328)	-29.43 (69.03)	7.823 (4.469)	3.813 (5.306)
Work Other	0.0970** (0.0297)	0.0338** (0.0130)	0.0340** (0.0113)	22.14 (23.26)	3.554** (1.269)	2.398 (1.455)
Atypical working hours	0.0924* (0.0390)	0.0957*** (0.0235)	0.0982*** (0.0195)	-17.59 (42.69)	0.668 (2.141)	5.984** (2.138)
Works part time	0.0370 (0.0333)	-0.0225 (0.0138)	-0.00736 (0.0129)	-17.68 (24.30)	-0.0611 (1.349)	0.0324 (1.731)
Farmers	1.019 (0.530)	0.186 (0.150)	0.182 (0.199)	108.4 (133.4)	36.54 (19.32)	13.00 (13.08)
Qualified Manual workers	-0.0354 (0.0379)	-0.000418 (0.0186)	0.0317* (0.0153)	112.7** (37.12)	8.538*** (1.771)	6.921*** (2.029)
Unqualified Manual Workers	-0.105 (0.0557)	-0.133*** (0.0243)	-0.0145 (0.0211)	54.19 (52.71)	4.458 (2.451)	6.390* (2.651)
Office clerks public sector	-0.141*** (0.0320)	-0.0603*** (0.0143)	-0.00855 (0.0129)	50.27 (26.06)	2.605 (1.359)	0.0435 (1.579)
Office clerks private sector	-0.0566* (0.0284)	-0.0188 (0.0135)	0.0109 (0.0127)	14.93 (25.85)	0.186 (1.253)	1.946 (1.513)
Personal Domestic Services	-0.159** (0.0604)	-0.153*** (0.0271)	-0.0553* (0.0258)	-37.63 (53.25)	-0.166 (2.525)	3.822 (3.165)
Technicians	0.0385 (0.0336)	-0.0138 (0.0180)	-0.0122 (0.0156)	30.88 (28.84)	1.557 (1.523)	1.687 (1.856)
Craftworkers	0.0292 (0.185)	0.0509 (0.0517)	0.0189 (0.0461)	912.0* (445.1)	26.45** (9.279)	19.17* (9.117)
Trades workers	0.320* (0.125)	0.140** (0.0444)	0.108** (0.0341)	582.5* (243.7)	16.01*** (4.748)	15.11** (5.846)
CEOs	0.399** (0.140)	0.145* (0.0622)	0.0837 (0.0473)	79.62 (128.6)	-1.938 (4.226)	19.49** (6.828)
Self-employed white-collar	0.203* (0.0952)	0.126*** (0.0330)	0.0945** (0.0315)	-100.9* (42.72)	-5.962 (3.103)	2.624 (5.615)
Managers	0.00708 (0.0228)	-0.0195 (0.0114)	-0.0195 (0.0105)	-30.36 (16.70)	-0.810 (0.964)	0.600 (1.353)
Female	-0.133*** (0.0176)	-0.00118 (0.00815)	0.0301*** (0.00720)	-37.64** (14.42)	-2.789*** (0.762)	-1.153 (1.018)
Household size	-0.0139* (0.00692)	0.0206*** (0.00325)	0.00570* (0.00290)	12.92 (9.219)	0.578 (0.328)	-0.934* (0.382)
Constant	2.386*** (0.140)			526.5*** (111.3)	52.37*** (5.366)	156.7*** (6.318)
N	12793	12793	12753	7687	7687	7687
R-squared	0.4519			0.0401	0.0509	0.0571

Standard errors clustered at the household level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Columns (2) and (3) report the average marginal effects for each coefficient. All specifications also include survey-day fixed effects, variables for age and age squared, and indicator variables for problems with taking transport, being on leave or on sickness leave on the survey day. D1,...,D10: indicator for belonging to the first,...,tenth decile of household income.

Table C.1.11 – Regression coefficients for distance, propensity to use a car and emission intensity - workers

List of Figures

0.0.1	Répartition des émissions de gaz à effet de serre, NO _x et PM _{2.5} par secteur au sein de l'UE 28 en 2018	21
0.0.2	Sectoral breakdown of greenhouse gases, NO _x and PM _{2.5} emissions in the EU 28 in 2018	41
1.2.1	The Carbon Price Support and EUA price on the EU ETS	63
1.2.2	Evolution of per capita power and non-power sector emissions in European countries	64
1.4.1	UK and synthetic UK per capita emissions	77
1.4.2	Per capita CO ₂ e emissions by source, UK and synthetic UK	79
1.4.3	SCM with counterfactual UK emissions w/o biomass conversion	82
1.4.4	Gap between treated and synthetic UK, CPS assumed to start in 2010	83
1.4.5	Leave-one-out test	84
1.4.6	Permutation test	85
1.4.7	Ratio of post to pre-MSPE	86
1.4.8	Spillover risk	88
1.5.1	per capita CO ₂ e emissions by source, UK and synthetic UK, lower bound estimation	91
2.2.1	Localization of Marseille Port and Air Quality Monitoring Stations and Hourly Vessel Traffic Variation.	99
2.3.1	Matching Diagnostics for the Cruise Arrivals Experiment.	108
2.3.2	Comparing Calendar Covariates Distribution for Initial and Matched Datasets.	109
2.3.3	Comparing Weather Covariates Distribution for Initial and Matched Datasets.	110
2.3.4	Matching Diagnostics for the Daily Experiment.	111
2.3.5	Effects of Cruise Vessel Traffic at the Hourly Level.	112
2.3.6	Effects of Cruise Vessel Traffic on Pollutant Concentrations at the Daily Level.	113
3.2.1	The Paris area	124

3.4.1	Lorenz curves for contributions to emissions at the individual level . . .	135
3.4.2	Contribution of distance, modal choice and emission intensity to the differences in emissions, by pollutant	137
3.4.3	Share of trip purposes in the number of trips, distances travelled and emissions	144
3.4.4	Cumulative Distribution Function of the difference in travel time between car and cycling	146
A.1.1	Per capita power sector emissions in European countries	184
A.1.2	Channels: evolution of electricity demand, trade and emission intensity in the UK and other European countries	186
A.1.3	Power sector's input fuel mix in EU countries, 2012 and 2017	187
A.1.4	Histograms of main predictors. UK: black with transparent fill; donor pool: grey fill.	197
A.1.5	Synthetic control method excluding emissions from plants having converted to biomass	200
A.1.6	Histograms of predictors used in the sensitivity analysis (UK: black with transparent fill; donor pool: grey fill)	203
A.1.7	Sensitivity analysis: alternative set of predictors	204
A.1.8	Sensitivity analysis: alternative donor pool	206
A.1.9	Sensitivity analysis: extended pre-treatment period with Eurostat greenhouse gas emissions by sector	208
A.1.10	Sensitivity analysis: extended pre-treatment period with Eurostat greenhouse gas emissions by sector, permutation test	209
A.1.11	UK (incl. counterfactual emissions from biomass converted plants) and synthetic UK - no interconnected countries	210
B.1.1	Seasonal and Hourly Variation in Vessel Traffic.	215
B.1.2	Hourly Traffic of Cruise Vessels Entering the Port on Mondays in July and August 2012.	216
B.1.3	Long Term, Daily and Hourly Variations in Pollutant Concentrations.	217
B.1.4	Hourly Variations in Weather Parameters.	218
B.1.5	Polar Plot of Wind Directions.	218
B.1.6	Polar Plot of Pollutant Concentrations according to Wind Speed and Direction.	219
B.1.7	Difference in hourly wind speed between the wind measured at the airport and at Longchamp station.	220

B.1.8	Polar Plots of Wind Directions at Longchamp and at the airport weather stations.	221
B.1.9	Daily and Hourly Variations in Hourly Road Traffic Flow.	222
B.1.10	Visualizing the Occurrence of the Treatment.	225
B.1.11	Continuous Weather Covariates Balance.	226
B.1.12	Categorical Weather Covariates Balance.	227
B.1.13	Balance of Road traffic variables.	228
B.1.14	Balance of Air Pollutants.	229
B.1.15	Love Plots for Continuous Weather Covariates.	229
B.1.16	Love Plots for Wind Direction.	230
B.1.17	Love Plot for Rainfall.	231
B.1.18	Love Plots for Air Pollutants.	232
B.1.19	Love Plots for Vessel Traffic.	233
B.1.20	Love Plots for Road Traffic.	234
B.1.21	Overall Improvement in Covariate Balance.	235
B.1.22	Comparing Weather Covariates Distribution for Initial and Matched Datasets.	236
B.1.23	Comparing Calendar Covariates Distribution for Initial and Matched Datasets.	237
B.1.24	Visualizing the Occurrence of the Treatment.	239
B.1.25	Continuous Weather Covariates Balance.	240
B.1.26	Categorical Weather Covariates Balance.	241
B.1.27	Balance of Road traffic variables.	242
B.1.28	Balance of Air Pollutants.	243
B.1.29	Love Plots for Continuous Weather Covariates.	244
B.1.30	Love Plots for Wind Direction.	244
B.1.31	Love Plot for Rainfall.	245
B.1.32	Love Plots for Air Pollutants.	246
B.1.33	Love Plots for Vessel Traffic.	247
B.1.34	Love Plots for Road Traffic.	248
B.1.35	Overall Improvement in Covariate Balance.	249
B.1.36	Comparing Weather Covariates Distribution for Initial and Matched Datasets.	250
B.1.37	Comparing Calendar Covariates Distribution for Initial and Matched Datasets.	251
B.1.38	Visualizing the Occurrence of the Treatment.	252
B.1.39	Balance of Continuous Weather Covariates.	253

B.1.40	Balance of Categorical Weather Covariates.	254
B.1.41	Balance of Air Pollutants.	255
B.1.42	Balance of Month and Year Indicators.	256
B.1.43	Love Plot for Continuous Weather Covariates.	257
B.1.44	Love Plots for Wind Direction.	257
B.1.45	Love Plot for Rainfall.	258
B.1.46	Love Plots for Air Pollutants.	258
B.1.47	Love Plots for Vessel Traffic.	259
B.1.48	Love Plot for Road Traffic.	259
B.1.49	Overall Improvement in Covariates Balance.	260
B.1.50	Comparing Weather Covariates Distribution for Initial and Matched Datasets.	261
B.1.51	Comparing Calendar Covariates Distribution for Initial and Matched Datasets.	262
B.1.52	Distribution of the Test Statistic Under the Sharp Null Hypothesis. . .	267
B.1.53	Upper and Lower P -Value Functions.	269
B.1.54	Distribution of Pair Differences in NO ₂ Concentrations.	269
B.1.55	Distribution of Pair Differences in O ₃ Concentrations.	270
B.1.56	Distribution of Pair Differences in SO ₂ Concentrations.	270
B.1.57	Distribution of Pair Differences in PM ₁₀ Concentrations.	271
B.1.58	Distribution of Pair Differences in PM _{2.5} Concentrations.	271
B.1.59	Two-Sided P -Value for the Sharp Null Hypothesis of No Effect	272
B.1.60	95% Fisherian Intervals	272
B.1.61	95% Fisherian Intervals	273
B.1.62	Distribution of Pair Differences in NO ₂ Concentrations.	273
B.1.63	Distribution of Pair Differences in O ₃ Concentrations.	274
B.1.64	Distribution of Pair Differences in SO ₂ Concentrations.	274
B.1.65	Distribution of Pair Differences in PM ₁₀ Concentrations.	275
B.1.66	Distribution of Pair Differences in PM _{2.5} Concentrations.	275
B.1.67	Two-Sided P -Value for the Sharp Null Hypothesis of No Effect	276
B.1.68	95% Fisherian Intervals	276
B.1.69	95% Fisherian Intervals	277
B.1.70	Distribution of Pair Differences in NO ₂ Concentrations.	277
B.1.71	Distribution of Pair Differences in O ₃ Concentrations.	278
B.1.72	Distribution of Pair Differences in SO ₂ Concentrations.	278
B.1.73	Distribution of Pair Differences in PM ₁₀ Concentrations.	278
B.1.74	Distribution of Pair Differences in PM _{2.5} Concentrations.	279

B.1.75	Pair Differences by difference in tonnage of entering cruise vessels. . . .	279
B.1.76	Two-Sided <i>P</i> -Value for the Sharp Null Hypothesis of No Effect	280
B.1.77	95% Fisherian Intervals	280
B.1.78	95% Fisherian Intervals, only pairs where wind blows from the West . .	281
B.1.79	95% Fisherian Intervals for road traffic variables, 2011-2016 period . . .	282
B.1.80	95% Fisherian Intervals based Wilcoxon signed-rank Test Statistic. . .	283
B.1.81	95% Fisherian Intervals based Wilcoxon signed-rank Test Statistic. . .	283
B.1.82	95% Fisherian Intervals based Wilcoxon signed-rank Test Statistic. . .	284
B.1.83	Number of Pairs with Missing Values.	285
B.1.84	95% Fisherian Intervals based on Wilcoxon signed-rank Test Statistic for Non-Imputed Concentrations.	285
B.1.85	Number of Pairs with Missing Values.	286
B.1.86	95% Fisherian Intervals based on Wilcoxon signed-rank Test Statistic for Non-Imputed Concentrations.	286
B.1.87	Number of Pairs with Missing Values.	287
B.1.88	95% Fisherian Intervals based on Wilcoxon signed-rank Test Statistic for Non-Imputed Concentrations.	287
B.1.89	95% Confidence Intervals for the Average Treatment Effect.	288
B.1.90	95% Confidence Intervals for the Average Treatment Effect.	289
B.1.91	95% Confidence Intervals for the Average Treatment Effect.	289
B.1.92	Statistical Power, Type M and S errors for the hourly experiment on cruise vessels' arrivals.	290
B.1.93	Statistical Power, Type M and S errors.	291
B.1.94	Regression Analysis of the Full Dataset.	292
B.1.95	Regression Analysis of the Full Dataset.	292
B.1.96	Regression Analysis of the Full Dataset.	294
B.1.97	Average Road Traffic Flow by Day of the Week.	296
B.1.98	Average NO ₂ Concentrations over the Week.	297
B.1.99	Average Vessel Gross tonnage over the Week.	297
B.1.100	Weather Covariates Balance.	298
C.1.1	Distribution of NO _x emissions per passenger, by transportation mode .	302
C.1.2	Distribution of PM _{2.5} emissions per passenger, by transportation mode	303
C.1.3	Distribution of CO ₂ emissions per passenger, by transportation mode .	304
C.1.4	Evolution of annual average NO ₂ concentrations in IdF	307
C.1.5	Evolution of annual average PM _{2.5} concentrations in IdF	308
C.1.6	Lorenz curve, trip level	312

C.1.7	Contribution of distance, modal choice and emission intensity to the differences in PM _{2.5} emissions	315
-------	---	-----

List of Tables

1.2.1	Level of CPS rate for each period in pound per ton of CO ₂ e	63
1.4.1	Country weights in Synthetic UK	77
1.4.2	Predictors' values for the UK, synthetic UK and average of the donor pool	78
2.2.1	Descriptive Statistics at the Daily Level (2008-2018) (N=4,018).	101
2.2.2	Maximum Discrepancies allowed for each Covariate between Treated and Control Units.	105
2.3.1	Number of Matched Pairs by Experiment.	107
3.3.1	Summary statistics - Individuals ≥ 18 years old	127
3.3.2	Assumed contribution to pollution emissions in different transport modes .	131
3.4.1	Estimated coefficients for distance and propensity to use a car - all individuals	139
3.4.2	Regression coefficients for the emission intensity of trips made by car - all individuals	141
3.4.3	Possibility of Modal shift at the car trip level	147
3.4.4	Teleworking and EV shift options for current drivers	151
A.1.1	Summary statistics at the country level, average 2005-2012	193
A.1.2	Counterfactual CO ₂ emissions in the absence of biomass conversion, Drax .	198
A.1.3	Counterfactual CO ₂ emissions in the absence of biomass conversion, Lynemouth	199
A.1.4	Country weights in Lower Bound Synthetic UK	201
A.1.5	Predictors' values for the UK, synthetic UK and average of the donor pool, lower bound	201
A.1.6	Predictors' values for the UK and each alternative synthetic UK	205
A.1.7	Predictors' values for the UK, synthetic UK and average of the donor pool, longer panel dataset	208
B.1.1	Science Table.	264
B.1.2	Observed Table.	265
B.1.3	Regression coefficients - Daily experiment	295

C.1.1	Emission factors for private vehicles not owned by the household	302
C.1.2	EGT-Descriptive statistics at the household level	309
C.1.3	Balance between EGT survey data and administrative data on selected household characteristics	310
C.1.4	Balancing test comparing the subsample of individuals with one trip recorder and the full sample	311
C.1.5	Extended Kaya components by quintile of NO _x emissions	313
C.1.6	LMDI decomposition on NO _x emissions at the individual level	313
C.1.7	Extended Kaya components by quintile of CO ₂ emissions	314
C.1.8	LMDI decomposition CO ₂ emissions at the individual level	314
C.1.9	Extended Kaya components by quintile of PM _{2.5} emissions	315
C.1.10	LMDI decomposition on PM _{2.5} emissions at the individual level	316
C.1.11	Regression coefficients for distance, propensity to use a car and emission intensity - workers	317