



Pollution and Environmental Policies

Martin Jégard

► To cite this version:

Martin Jégard. Pollution and Environmental Policies. Economics and Finance. Université Panthéon-Sorbonne - Paris I, 2022. English. NNT : 2022PA01E038 . tel-03977651

HAL Id: tel-03977651

<https://theses.hal.science/tel-03977651>

Submitted on 7 Feb 2023

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.

UFR02 Économie - Paris Jourdan Sciences Économiques
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 29 août 2022 par

Martin JÉGARD

Pollution et Politiques Environnementales

Sous la direction de :
Hélène OLLIVIER
Katheline SCHUBERT

Composition du jury :

Rapporteurs

Isabelle MÉJEAN
Francesco VONA

Professeur à Sciences Po
Professeur à l'Université de Milan

Examineurs

Ariell RESHEF

Robert ELLIOTT
Antoine DECHEZLEPRÊTRE

Directeur de Recherche au CNRS
Professeur à l'Université Paris 1 Panthéon-Sorbonne, PSE
Professor à l'Université de Birmingham
Associate Professorial Research Fellow
au Grantham Research Institute (LSE) - OCDE

Directrices de thèse

Hélène OLLIVIER
Katheline SCHUBERT

Professeur à PSE - Chargée de Recherche au CNRS
Professeur à l'Université Paris 1 Panthéon-Sorbonne, PSE

UFR02 Économie - Paris Jourdan Sciences Économiques
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 the 29th of August 2022 by

Martin JÉGARD

Pollution and Environmental Policies

Thesis Advisors :
Hélène OLLIVIER
Katheline SCHUBERT

Jury :

Reviewers

Isabelle MÉJEAN
Francesco VONA

Professor at Sciences Po
Full Professor at the University of Milan

Examiners

Ariell RESHEF

Senior researcher at CNRS
Professor at Paris 1 Pantheon-Sorbonne & PSE

Robert ELLIOTT

Professor at the University of Birmingham

Antoine DECHEZLEPRÊTRE

Associate Professorial Research Fellow
at the Grantham Research Institute (LSE) - OECD

Advisors

Hélène OLLIVIER

Research Fellow at CNRS - Professor at PSE

Katheline SCHUBERT

Professor at Université Paris 1 Panthéon-Sorbonne - PSE

Remerciements - Acknowledgements

Je souhaite tout d'abord remercier Hélène dont le soutien et la bienveillance ont toujours été sans faille. Je souhaite à chaque doctorant·e d'avoir la chance d'être encadré·e par un·e chercheur de cette qualité. Je souhaite également remercier Katheline et Ariell qui ont toujours répondu présents pour aiguiller mes travaux pendant ces années de thèse. Leurs conseils avisés m'ont aidé à y voir plus clair sur les objectifs donnés à mes différents projets de recherche. Un grand merci à Isabelle Méjean, Francesco Vona, Robert Elliott et Antoine Dechezleprêtre pour avoir accepté de faire partie de mon jury. En particulier, merci à Isabelle et Francesco pour vos retours lors de ma pré-soutenance. I also want to thank Geoffrey Barrows for his support. I learned a lot by working on our co-authored projects and I am indebted to him for teaching me how to think about an empirical research project.

Un grand merci à mes camarades doctorants, en particulier Manon et Benjamin. Nos conversations m'ont mis les baskets aux pieds et je vous dois d'avoir trouvé l'excuse pour raccourcir mes journées à PSE afin aller courir sur les quais ou au bois de Vincennes. Merci aussi à Max, Xavier, Irène, Léo, Marion, Thomas, Paul, Morten, Eva, Rind, Claire, Rosanne et tous les autres. La thèse prend du temps et ce temps semble parfois long. Heureusement, il peut aussi passer plus vite lorsque l'on est bien entouré. Merci à l'administration de PSE, ainsi que ses professeurs. Merci en particulier à Roxana pour son aide inestimable pendant le *job market*, à Véronique pour son soutien pendant les longs confinements successifs, à Nicolas pour nos discussions durant cette dernière année et à François pour son cours de *Geocoded Datasets* qui m'a permis de mener une bonne partie de mes projets empiriques.

Merci à Eloi, mon camarade ch'timi égaré lui aussi en région parisienne, pour ton exemple de chercheur investi et pour nos bières aux Fontaines ou devant Notre-Dame. Merci à Benjamin pour les coupures sauvages et salvatrices entre chacune de ces quatre années. J'espère que nous continuerons d'arpenter ensemble de nombreux autres chemins de traverse. Pour les (trop rares) moments passés dans le Nord, merci à mon ami de toujours, Marius. La vie parisienne serait aussi bien fade sans tous les membres de PàP et en particulier Marie, Cécile et le docteur Lesage. Merci à vous de me rappeler régulièrement qu'il est bon de sortir de chez soi pour profiter des terrasses parisiennes !

Merci à mes sœurs et mes frères, Marie, Laure, Théophile et Henri, mais aussi Charles et Clara dont l'arrivée dans la famille ces dernières années nous aura tous remplis de joie. Je n'oublie pas non plus la petite Hortense, son parrain attend avec impatience de pouvoir lui expliquer pourquoi préférer le vélo à la voiture. Je remercie aussi ma grand-mère Bernadette, dont la porte a toujours été ouverte au cours de ces longues années d'étude qui se terminent (enfin) et qui a toujours su prendre soin de ses petits-enfants. Merci à mère à qui je dois tout. Merci à mon père et à mes grands-parents qui, de là où ils sont, continuent de nous soutenir dans toutes nos ambitions.

Résumé de la thèse

Les trois chapitres de cette thèse sont indépendants et contribuent à l'étude de l'économie environnementale par différentes approches.

Le premier chapitre, *Inter-Firm and Intra-Firm Spillover Effects of Industrial Regulation*, co-écrit avec Geoffrey Barrows, Raphael Cael et Hélène Ollivier, présente une nouvelle méthode d'évaluation d'une réglementation en présence d'effets indirects inter- et intra-entreprises. Notamment lorsque la concurrence est imparfaite, une réglementation affectant directement certaines entreprises uniquement est aussi susceptible d'affecter indirectement les entreprises non réglementées. Ainsi, l'approche traditionnelle des "doubles différences" n'intègre pas l'intégralité des effets. Pour surmonter cette limitation, nous développons un modèle structurel de concurrence imparfaite et d'entreprises multi-établissements et mettons au point une approche non biaisée pour estimer les effets d'une réglementation affectant une sous-population d'établissements. Nous appliquons notre méthode au cas du Système européen d'échange de permis d'émission en France et montrons que cette réglementation a permis de réduire les émissions de CO₂ des installations réglementées, ainsi que celles du secteur manufacturier dans son ensemble, et d'augmenter les revenus des entreprises réglementées.

Si la qualité de l'air doit être améliorée dans les lieux les plus densément peuplés, des politiques trop strictes peuvent aussi nuire aux industries locales et aux revenus des travailleurs. Dans le deuxième chapitre, *An Optimal Distribution of Polluting Activities Across Space*, j'analyse ces forces contradictoires à l'aide d'un modèle d'économie spatiale. Lorsque les entreprises n'internalisent pas les effets de la pollution et que les travailleurs migrent hors des lieux pollués, les plus grandes villes peuvent alors être trop petites par rapport à l'optimum. Imposer un prix d'émission de pollution relativement plus élevé dans ces endroits conduirait à un bien-être plus important. L'application empirique aux villes françaises montre que les politiques actuelles imposent des prix d'émissions plus élevés dans les grandes villes, mais que des gains de bien-être sont encore possibles.

Les effets de la pollution de l'air sur la productivité des travailleurs sont plutôt bien documentés. Cependant, leurs conséquences sur les performances des entreprises sont encore peu étudiées. Dans le troisième chapitre, *The Effects of Air Pollution on Exports: Evidence from PM_{2.5} & French Firms*, co-écrit avec Geoffrey Barrows et Hélène Ollivier, nous examinons les effets d'augmentations de la concentration aérienne en particules fines sur les exportations des entreprises françaises. Notre approche repose sur des données très désagrégées, à la fois dans l'espace et dans le temps. Afin d'estimer l'impact causal de la pollution sur les exportations, nous instrumentons les concentrations de pollution en utilisant les variations exogènes de la direction du vent au cours du temps. Nous constatons que des niveaux de pollution plus élevés entraînent une baisse des quantités exportées et de leurs valeurs.

DISCIPLINE : Sciences Économiques

MOTS-CLEFS : Pollution, Politiques Environnementales, Industrie

Thesis summary

The three chapters of this thesis are independent and contribute to environmental economics through different approaches.

The first chapter, *Inter-Firm and Intra-Firm Spillover Effects of Industrial Regulation*, co-authored with Geoffrey Barrows, Raphael Caelel, and Hélène Ollivier, presents a new method for estimating the impact of a regulation in the presence of inter- and intra-firm spillover effects. In particular, when competition is imperfect, a regulation targeting some firms is also likely to affect unregulated firms. Thus, the traditional difference-in-differences approach will not evaluate the full impact of the regulation. To overcome this limitation, we develop a structural model and an empirical strategy to estimate the effects of a regulation affecting a sub-population of industrial facilities. We apply our method to the European Emission Trading System and show that it reduced carbon emissions from regulated facilities as well as from the French manufacturing sector as a whole, but increased revenues of regulated French firms.

While pollution should be reduced in more densely populated locations, stringent regulations may also harm local industries and reduce workers' incomes. In the second chapter, entitled *An Optimal Distribution of Polluting Activities Across Space*, I analyze these conflicting forces using a spatial economics model to identify which cities should be targeted by air quality policies. When the effects of pollution are not internalized by firms and workers respond to poor air quality by migrating away, then the largest cities may be too small compared to the optimum. Imposing a higher pollution price in these locations relative to the rest of the country would lead to higher welfare. The empirical application on French cities shows that current policies impose higher emission prices in large cities, but welfare gains are still possible.

There is a lot of evidence that air pollution has negative effects on the productivity of industrial workers. However, little is known about its consequences on the economic performance of firms that employ these workers. In the third chapter, *The Effects of Air Pollution on Exports: Evidence from PM_{2.5} & French Firms*, co-authored with Geoffrey Barrows and Hélène Ollivier, we examine the effects of temporary and local increases in PM_{2.5} air concentration on the exports of French firms. Our approach relies on highly disaggregated data, both in space and in time. To estimate the causal impact of pollution on exports, we instrument local pollution concentrations using exogenous variations in wind direction over time. We find that higher pollution levels lead to lower export values and quantities.

DISCIPLINE: Economics

KEYWORDS: Pollution, Environmental Policies, Industry

Sommaire général

Remerciements - Acknowledgements	5
Résumé de la thèse	7
Thesis summary	9
Résumé substantiel en français	15
General Introduction	35
1 Inter-Firm and Intra-Firm Spillover Effects of Industrial Regulation	45
1.1 Introduction	45
1.2 Data and Reduced-Form Evidence of Spillover Effects	50
1.2.1 The European Emissions Trading Scheme	50
1.2.2 Firm-Level Data	52
1.2.3 Plant-Level Data	53
1.2.4 Reduced-Form Evidence of Inter-Firm Spillovers	55
1.2.5 Reduced-Form Evidence of Intra-Firm Spillovers	56
1.3 Model	59
1.3.1 Demand	59
1.3.2 Production, Emissions and Regulation	60
1.3.3 Profit Maximization and Equilibrium	62
1.3.4 The Effect of Regulation on Revenues	63
1.3.5 The Effect of Regulation on Emissions	65
1.3.6 Extension to Multi-plant Firms	65
1.4 Estimation Strategy	67
1.4.1 Estimating Parameters	67
1.4.2 Treatment Effects	70
1.4.3 Monte Carlo Experiments	71
1.5 Moving from Model to Data	76
1.5.1 Export Markets	76
1.5.2 International Imports	76
1.5.3 Imputing Emissions	77
1.5.4 Matching	77
1.6 Results	78
1.6.1 Estimating Reduced-form Coefficients	79
1.6.2 Estimating Average and Aggregate Treatment Effects	81
1.7 Conclusion	84

2	An Optimal Distribution of Polluting Activities Across Space	87
2.1	Introduction	87
2.2	Optimal Environmental Policy in a Simple Spatial Model	93
2.3	General Spatial Model	97
2.3.1	Setup	98
2.3.2	Equilibrium and Optimal Set of Emission Taxes	102
2.4	Estimation of the Model	102
2.4.1	Data	103
2.4.2	Estimation of Congestion and Pollution Damage Elasticities	104
2.4.3	Estimation of Sector-Specific Parameters	107
2.4.4	Trade Costs Matrix	110
2.4.5	Recovering Local Characteristics	111
2.5	Spatially Heterogeneous Emission Regulations	112
2.5.1	Evidence of Spatially Heterogeneous Emission Regulations	112
2.5.2	Welfare Impacts of Spatially Heterogeneous Emission Regulations	114
2.5.3	Optimal Mean-Preserving Set of Heterogeneous Regulations	115
	Reallocating Population Across Space	118
	Reallocating Emissions Across Space	119
2.6	Conclusion	123
3	The Effects of Air Pollution on Exports: Evidence from PM2.5 & French Firms	125
3.1	Introduction	125
3.2	Data	129
3.2.1	Atmospheric Pollution in France	129
3.2.2	Wind Direction and Meteorological Controls	131
3.2.3	Trade Data	132
3.2.4	Plants Geo-localization	134
3.3	Empirical Strategy	139
3.4	Results	141
3.4.1	Baseline Results	141
3.4.2	Longer-term effects of PM2.5	142
3.4.3	Heterogeneity Analysis	144
3.4.4	Robustness Tests	145
	Restricting to Single-plant Firms:	145
	Controlling for the effects of other harmful atmospheric pollutants:	146
	Placebo test:	146
3.5	Conclusion	147
	Bibliography	148
A	Appendix to Chapter 1	165
A.1	Further Empirical Results	165
A.2	Data Appendix	170
A.2.1	FICUS-FARE	170
A.2.2	EACEI	170

A.2.3	EUTL	172
A.2.4	Imports and Exports	173
A.2.5	Variable Definitions	174
Revenue	174
Employment	174
Energy consumption	174
A.2.6	Imputing Emissions	174
A.3	Theoretical Appendix	176
A.3.1	Proof of the exact CES price index change	176
A.3.2	Deriving the Local Approximation Estimation Equation	178
A.4	Comparing to other models	178
A.4.1	Analytical Results on Bias in DiD and LA Models	178
A.4.2	Monte Carlo Experiments	182
A.5	Extension to Oligopoly	189
A.5.1	Model with Bertrand-Nash competition	189
A.5.2	Monte-Carlo experiments for Oligopoly	190
B	Appendix to Chapter 2	195
B.1	Descriptive Statistics on Pollution from the Industry	195
B.2	Air Quality Regulations in France	198
B.3	Illustrative Framework Appendix	201
B.4	Relaxing the Free Migration Assumption	205
B.5	Estimation Appendix	207
B.5.1	Building Plant-Level Pollutants Emissions from Energy Surveys . .	207
B.5.2	Estimation of the Sector Specific Emission Intensities	207
B.6	Application to France Appendix	210
B.6.1	Sources of Equilibrium Data	210
Empirical Distribution of Wages Across Commuting Zones:	210
Empirical Distribution of Populations Across Commuting		
Zones:	210
Empirical Distribution of Emissions Across Commuting Zones:	210	
B.6.2	Descriptive Statistics on Equilibrium Data	211
B.6.3	Descriptive Statistics on Computed Local Characteristics	211
C	Appendix to Chapter 3	219
C.0.1	Timing of Exports	219
C.0.2	Extended Tables - Baseline Results	221
C.0.3	Heterogeneity Across Manufacturing Activities	221

Résumé substantiel en français

Notre société est confrontée aujourd’hui à de nombreux dérèglements environnementaux causés par les activités humaines. Deux de ces enjeux majeurs sont abordés dans ma thèse : la diminution des émissions de gaz à effet de serre (GES), afin de limiter le réchauffement climatique, et les conséquences de la pollution de l’air sur la santé et les activités humaines. Le rôle des secteurs industriels est particulièrement important sur ces deux enjeux, justifiant la mise en place de nombreuses politiques publiques visant à réduire leurs émissions de GES et d’autres polluants locaux. Ces politiques doivent être élaborées de manière cohérente, efficace et socialement acceptable. La recherche en économie peut contribuer à cette démarche en étudiant rétrospectivement et prospectivement les politiques possibles. Dans un contexte de rationalisation des politiques publiques, leur évaluation est nécessaire afin de mesurer l’impact de l’action publique. En effet, une connaissance précise des conséquences de la mise en place de réglementations visant à améliorer la qualité de l’air et à réduire les émissions de GES est un prérequis à leur acceptabilité par la société ainsi qu’à leur amélioration future.

Dans le premier temps de cette introduction, je présente le contexte environnemental et politique dans lequel s’inscrit cette thèse. Dans un second temps, je propose une synthèse de chacun des trois chapitres.

Changement climatique, pollution de l’air et politiques environnementales

Deux enjeux environnementaux majeurs

Gaz à effet de serre et changement climatique : Ces dernières décennies, des preuves solides de l’impact des activités humaines sur le climat ont été apportées. La combustion d’énergies fossiles, comme le pétrole ou le gaz, la production de biens, les déplacements ou encore la production d’électricité émettent des gaz à effet de serre (GES) qui contribuent au réchauffement de l’atmosphère au fur et à mesure de leur accumulation dans l’atmosphère (IPCC, 2021a). En pratique, la température moyenne à la surface terrestre entre 2011 et 2020 a augmenté de plus de 1°C par rapport à la période 1850-1900 (IPCC, 2021a). Le dernier rapport du Groupe d’experts intergouvernemental sur l’évolution du climat (GIEC) montre qu’ “[...] à moins de réductions immédiates, rapides et massives des émissions de gaz à effet de serre, la limitation du réchauffement aux

alentours de 1,5°C, ou même à 2°C, sera hors de portée.” (IPCC, 2021b). Au-delà des températures, ce changement climatique impacte d’autres variables, comme l’humidité et la sécheresse, les vents, la neige et la glace, les zones côtières et les océans. En particulier, la fréquence des catastrophes climatiques (ouragans, inondations, canicules et sécheresses) pourrait ainsi augmenter si l’évolution des températures n’est pas maîtrisée.

Afin de limiter cette augmentation des températures, le consensus est aujourd’hui qu’il faut réduire les émissions totales de GES. En 1997, 35 pays ont pris des engagements à l’occasion du sommet de Kyoto afin de réduire de 5 % leurs émissions de GES. En 2015, l’Accord de Paris sur le Climat est signé par plus de 50 pays (représentant plus de la moitié des émissions mondiales de GES) s’engageant à limiter l’augmentation de la température globale à 2°C.¹ En 2020, dans le contexte de la pandémie de Covid, les émissions mondiales ont connu un recul de 5,4 %. Cependant, il semble qu’elles soient aussitôt revenues aux valeurs “pré-Covid” et que la concentration de GES dans l’atmosphère continue d’augmenter à ce jour (UNEP, 2021). En particulier, les émissions de GES au cours de la dernière décennie (de 2010 à 2019) ont été plus importantes que jamais (IPCC, 2022). Cette quantité émise est du même ordre de grandeur que le “budget carbone”, c’est-à-dire la quantité maximum pouvant encore être émise tout en limitant l’augmentation de la température globale à 1,5°C. Pour le moment, la France s’est engagée à réduire de 40 % ses émissions de GES d’ici 2030 et à atteindre la neutralité carbone en 2050 (ce qui correspond à une division par 6 de ses émissions par rapport à 1990).² Les émissions françaises en GES ont d’ailleurs baissé de 20 % entre 1990 et 2019 (CITEPA, 2021). Ces engagements, ainsi que l’urgence climatique, posent donc aujourd’hui un défi et nécessitent de mettre en place une transition suffisante pour atteindre la neutralité carbone (Ademe, 2021).

Polluants atmosphériques et effets sur la santé : L’émission de particules polluantes, nocives pour la santé, est une autre conséquence négative des activités économiques. Elle est d’ailleurs reconnue comme étant “la plus importante menace environnementale pour la santé humaine” par l’Organisation mondiale pour la santé (OMS) dans la publication de ses dernières recommandations relatives à la qualité de l’air (WHO, 2021). Cette menace est due aux effets négatifs sur la santé d’une série de particules polluantes émises par diverses activités. Les principales sont les particules fines (PM10 et PM2.5), le dioxyde d’azote (NO2), l’ozone (O3) et le dioxyde soufre (SO2).³ Les effets de ces pol-

¹Les variations de températures sont mesurées par rapport au niveau de température préindustriel.

²La neutralité carbone consiste, d’une part, à diminuer ses émissions de GES au maximum, et d’autre part, à compenser les émissions restantes.

³Les particules fines se classent généralement en deux familles correspondant à la dimension principale de la particule : inférieure à 10µm pour les PM10 et à 2,5µm pour les PM2.5.

luants sont multiples et touchent différents paramètres physiologiques, comme le système respiratoire (Viehmann et al., 2015), le système cardiovasculaire (Brook et al., 2010) ou le cerveau (Costa et al., 2019). L'exposition précoce à un air pollué peut également avoir des effets latents, c'est-à-dire des effets qui ne se manifestent que plus tard dans la vie des enfants affectés (Bale et al., 2010; Almond & Currie, 2011; Colmer & Voorheis, 2020). Il est estimé que cette pollution de l'air est la cause de 4,4 millions de décès prématurés dans le monde tout polluants confondus (WHO, 2014), et que la pollution liée aux PM2.5 cause environ 379 000 décès prématurés dans l'Union européenne (UE) (European Environment Agency, 2021) et 40 000 en France (Santé Publique France, 2021). Une différence primordiale entre GES et polluants atmosphériques nocifs est que, pour ces derniers, le lieu d'émission est important pour en mesurer l'impact. En effet, si le réchauffement climatique est fonction de la quantité totale d'émissions de GES accumulées dans l'atmosphère - d'où la désignation de polluants "globaux" - les effets nocifs des particules polluantes se produisent lorsqu'ils sont respirés par les individus - d'où la désignation de polluants "locaux". Ainsi c'est leur concentration atmosphérique locale qui doit être contrôlée.

Depuis les premières catastrophes (le *Donora Smog* en 1948 aux Etats-Unis, le *Great Smog* de Londres en 1952 ou encore l'accident de Bhopal en 1984 en Inde) et les premières mises en évidence des effets de ce type de pollution, brève ou chronique, sur la santé humaine de nombreuses réglementations ont été mises en place afin de tenter de garantir une meilleure qualité de l'air. Au niveau international et depuis 1979, la Convention sur la pollution atmosphérique transfrontière longue distance, sous l'égide de l'Organisation des Nations unies (ONU), vise à réduire la pollution liée au SO₂, au NO₂ et aux particules fines, ainsi que certain métaux lourds et les polluants organiques persistants. Pour l'Union européenne, la stratégie pour une meilleure qualité de l'air se fonde sur trois piliers principaux : (1) la Directive sur la Qualité de l'Air Ambiant (2008/50/CE) qui met en place des standards de qualité de l'air, en pratique des niveaux de concentration à ne pas dépasser, ainsi que des obligations de mesure et de mise en place de plans dans les États membres (directive-cadre de 1996 concernant l'évaluation et la gestion de la qualité de l'air ambiant (96/62/CE)) ; (2) la Directive sur les engagements nationaux de réduction des émissions (directive 2001/81/CE puis 2016/2284) ; et (3) la mise en place de standards et limites spécifiques par source d'émission, ciblant notamment les secteurs industriels (avec par exemple la directive 2001/80/CE sur les émissions en provenance des grandes installations de combustion ou la directive 2010/75/UE sur les émissions industrielles).

Au cours des deux dernières décennies, la qualité de l'air s'est grandement améliorée et les concentrations d'un certain nombre de polluants ont fortement diminué dans les pays développés (aux Etats-Unis (US EPA, 2022) et dans l'Union européenne (AEE, 2013, 2020) par exemple). Cependant, l'OMS estime que 92 % de la population mondiale con-

tinue d'être exposée à des concentrations atmosphériques trop élevées (WHO, 2016). Bien que les concentrations les plus élevées soient observées dans des pays en développement, tels que l'Inde et la Chine, les valeurs sanitaires fixées par l'OMS ne sont toujours pas atteintes dans certains pays développés. Dans le cas de l'UE, malgré la réduction de la concentration d'un certain nombre de polluants ces dernières années, les concentrations en particules fines (PM2.5) et en ozone (O3) restent très élevées (Sicard et al., 2021). De plus, les PM2.5 ont aussi des effets sur la santé en dessous des seuils sanitaires fixés (Di et al., 2017). La France a été l'un des premiers pays à mettre en place une réglementation visant à limiter ces émissions de polluants (Massard-Guilbaud, 1999). Ainsi, la qualité de l'air extérieur s'est améliorée ces dernières décennies et les émissions de tous les polluants sont en baisse depuis 1990 (CITEPA, 2021). Pourtant, dans de nombreuses agglomérations, les standards européens ne sont pas encore atteints pour les particules fines (PM10) et le dioxyde d'azote (NO2). La France a d'ailleurs été condamnée par la Cour de justice de l'UE pour avoir "dépassé de manière systématique" les concentrations limites pour ces polluants (CJUE, 2022).

Le rôle de l'industrie dans l'émissions de ces pollutions : Les GES et polluants atmosphériques sont notamment émis par les activités humaines et principalement par la combustion d'énergie fossiles.⁴ En particulier, les secteurs des transports, de l'industrie, de l'agriculture et des bâtiments résidentiels et tertiaires jouent un rôle majeur dans ces émissions. Pour les GES, en 2019, ces secteurs représentaient respectivement 31 % des émissions totales pour les transports, 19 % pour l'industrie manufacturière et la construction, 19 % pour l'agriculture, 17 % pour le résidentiel et tertiaire, 10 % pour l'énergie, et 4 % pour le traitement de déchets (CITEPA, 2021). L'industrie de l'énergie a vu ses émissions se réduire de 46 % et l'industrie manufacturière de 41 %, contribuant ainsi fortement à la réduction des émissions nationales entre 1990 et 2019 (CITEPA, 2021). Selon le Réseau Action Climat, les émissions en provenance de l'industrie française ont diminué entre 1990 et 2015 du fait de l'amélioration de l'efficacité énergétique, mais aussi à travers une réallocation des émissions vers l'étranger, avec des fermetures d'usines, des délocalisations et une augmentation des importations. Ainsi, si l'empreinte carbone nationale s'est maintenue à un niveau stable depuis une vingtaine d'années, c'est que les "émissions importées" augmentent de façon continue et seraient même désormais plus importantes que les émissions sur le territoire français (HCC, 2020).⁵

Concernant la pollution de l'air, l'industrie manufacturière fait partie des contributeurs

⁴Certains polluants locaux sont aussi émis lors de processus spécifiques, comme la friction des pneumatiques sur la route dans les transports ou les émissions dues à des réactions chimiques dans l'industrie.

⁵Ces émissions correspondent aux émissions se produisant à l'étranger lors de la production des biens importés et consommés en France.

majeurs avec le secteur des transports. En 2019, l'industrie manufacturière française contribuait à 55 % des émissions de SO₂, 27 % des émissions de PM₁₀, 18 % de PM_{2.5} et 13 % des émissions de NO_x. Depuis 1990, les émissions de SO₂ ont diminué de 86 %, celles de PM₁₀ de 51 %, celles de PM_{2.5} de 58 %, et celles de NO_x de 55 % (CITEPA, 2021). Pour les Etats-Unis, les émissions de ces polluants ont également diminué considérablement depuis 1990. Shapiro & Walker (2018) a montré que cette réduction était principalement due à l'intensification des politiques environnementales (plus qu'à des effets de productivité ou de réallocation des flux commerciaux).

Ainsi, tant la réduction des émissions de GES que l'amélioration de la qualité de l'air continuent d'être, en France comme dans le reste du monde, un enjeu central pour les années à venir. Les politiques publiques mises en place par les gouvernements sont des instruments puissants pouvant résoudre ces problématiques (Shapiro, 2022). Cependant, de nombreuses questions se posent encore sur les avantages respectifs de ces instruments ainsi que sur leur efficacité en termes environnementaux mais aussi économiques.

Les politiques environnementales

L'émission de polluants dans l'atmosphère (GES, SO₂, NO₂, etc.) est une externalité négative : l'action d'un agent économique en affecte un autre sans qu'une compensation ait lieu. Cette défaillance du marché aboutit à un bien-être moindre des agents économiques car le coût des dommages causés par l'action en question n'est pas pris en compte par le pollueur dans ses décisions (quantités produites, investissements, etc.). La quantité de pollution émise est donc trop importante par rapport à l'optimum social. Dans cette situation, l'intervention publique peut permettre d'internaliser les dommages causés par la pollution et ainsi augmenter le bien-être social. Traditionnellement, ces interventions peuvent être séparées en deux familles distinctes. La première consiste à fixer des standards de technologie, ou des standards sanitaires, et à les faire respecter par les entreprises. Il s'agit de l'approche dite de *command-and-control*, utilisant des plafonds d'émissions par entreprise ou des standards technologiques par exemple. Cependant, les économistes considèrent que, dans certains cas, il est possible de corriger de manière plus efficace une externalité négative en fixant un prix adapté. Il s'agit de la seconde famille de politiques, dites *market-based*. Concernant l'émission de pollution, les deux instruments les plus répandus sont les taxes (ou subventions) et les marchés de permis d'émissions. Une taxe fixée au coût marginal des dommages liés à la pollution (changement climatique, impacts sanitaires des particules, etc.) permet aux entreprises d'internaliser l'impact de leur production, c'est à dire de tenir compte des effets environnementaux de leurs décisions (Pigou, 1920). Les marchés de permis d'émissions, ou instruments dits *cap-and-trade*,

peuvent également permettre cette internalisation. L'idée est de fixer la quantité totale d'émissions, par exemple pour respecter le budget carbone ou atteindre un certain niveau de réduction des émissions, puis de créer pour chaque unité de pollution pouvant être émise un droit d'émission que les entreprises peuvent s'échanger sur un marché créé à cet effet.

En pratique, le choix entre les différents instruments pour réglementer les externalités environnementales est aujourd'hui encore largement débattu (Bretschger & Pittel, 2020). Les approches en *command-and-control* ont souvent été privilégiées car perçues comme plus claires. Les technologies les plus polluantes sont bannies au profit de technologies plus propres. Par exemple, aux Etats-Unis c'est sur ce modèle que le *Clean Air Act* et le *Water Air Act* ont été pensés. C'est également le cas d'une partie de la réglementation européenne qui contraint les installations industrielles à adopter une série de standards technologiques appelés les "meilleures techniques disponibles" (MTD). De nombreux exemples de taxes sur les émissions de pollution existent également, comme la taxe sur les activités polluantes (TGAP) en France ou les redevances NOx en Suède et en Norvège (Bonilla et al., 2017). Cependant, lorsque les acteurs sont hétérogènes et qu'il existe de l'incertitude sur leurs coûts respectifs de diminution de leurs émissions, les instruments en *cap-and-trade* sont jugés plus efficaces (Newell & Stavins, 2003). Ce type d'instrument a été mis en place notamment aux Etats-Unis (avec par exemple l'*Acid Rain program* visant les émissions de SO₂ et de NOx, le programme *RECLAIM*⁶ pour les émissions de NOx à Los Angeles, ou encore le marché carbone californien), au Royaume-Uni et en Chine.

Le Système d'échange de quotas d'émission de l'Union européenne (SEQUE-UE), objet du premier chapitre de cette thèse, est aujourd'hui l'un des principaux exemples de ce type d'instrument. Ce système a été mis en œuvre par la Commission européenne et régit environ 11 000 installations dans 31 pays participants, dont la France. Le SEQUE-UE a été annoncé en 2000 dans un document de travail de la Commission européenne, puis officialisé par la directive établissant un système d'échange de quotas d'émissions de gaz à effet de serre en 2003 (2003/87/CE) et lancé le 1er janvier 2005. Le système réglemente toutes les installations de combustion dont la puissance thermique nominale est supérieure à 20 MW, ainsi que les autres processus de production dont la capacité ou la production est supérieure à des seuils prédéterminés propres à chaque secteur. Les entreprises réglementées doivent restituer des permis - appelés unités de quotas de l'UE - pour chaque tonne de GES émise au cours de l'année. La non-conformité entraîne une amende de 100 euros par tonne de GES dépassant le nombre de permis restitués. Au départ, la Commission européenne a planifié le SEQUE en trois phases : la première phase, conçue comme une période d'essai, s'est déroulée de 2005 à 2007 ; la deuxième phase s'est

⁶ *California Regional Clean Air Incentives Market*

déroulée de 2008 à 2012, correspondant à la première période d’engagement du protocole de Kyoto ; et la troisième phase s’est déroulée de 2012 à 2020. Les révisions du système en vue de la quatrième phase (2021-2030) sont aujourd’hui en cours. La mise en œuvre de la politique était relativement décentralisée au cours des deux premières phases : les pays participants soumettant à la Commission européenne des plans nationaux d’allocation de quotas (PNAQ) dans lesquels ils proposaient une liste d’installations à réglementer, des plafonds d’émission et les règles de distribution des quotas (Ellerman et al., 2015). Pendant ces phases, les permis ont été alloués gratuitement aux installations réglementées sur la base de leurs émissions historiques (selon la méthode du *grandfathering*).⁷ Pour la troisième phase, le système d’allocation a amorcé une transition vers un système d’enchères ouvertes, d’abord pour les installations de production d’énergie (à l’exception de quelques centrales en Europe de l’Est), puis progressivement pour les autres industries réglementées.

De nombreuses évaluations de l’impact de cette politique existent aujourd’hui (Petrick & Wagner, 2014; Löschel et al., 2018; Colmer et al., 2020b; Dechezleprêtre et al., 2018; Jaraite & Di Maria, 2018; Calligaris et al., 2018). Cependant, les révisions du système en vue des phases futures font l’objet d’intenses oppositions généralement motivées par l’idée qu’une réglementation trop ambitieuse pourrait faire peser un risque sur la compétitivité de l’industrie européenne (Demailly & Quirion, 2006; Branger et al., 2017; Joltreau & Sommerfeld, 2018). Les lobbies industriels, en particulier ceux représentant les intérêts des secteurs de l’acier, du ciment et des engrais, jouent un rôle important dans les discussions liées aux révisions du SEQUE et ces oppositions menacent aujourd’hui la pérennité de l’action environnementale (Thomas, 2021). Dans ce contexte, une évaluation du SEQUE-UE rendant compte non seulement de son impact sur les émissions de GES mais aussi de son impact sur les performances économiques des entreprises réglementées est nécessaire. Il s’agit de l’un des objectifs du premier chapitre de cette thèse. En effet, nous y calculons à la fois la réduction de GES permise par le SEQUE-UE en France mais nous montrons aussi que cette réglementation a eu un effet positif sur les revenus des entreprises réglementées.

De nombreuses questions se posent également au sujet des réglementations visant à améliorer la qualité de l’air en réduisant les émissions de polluants atmosphériques. Dans le cas de la France, cette réglementation est basée sur les normes européennes. L’UE fixe des plafonds de concentrations maximales dans l’air de différents polluants ainsi que des objectifs de réduction (directives 2004/107 et 2008/50/CE) qui doivent ensuite être appliqués dans chaque État membre. Ceux-ci doivent donc mettre en œuvre des plans d’action spécifiques pour rester au-dessous de ces plafonds et atteindre les objectifs de qualité de l’air. En France, ces actions sont menées à la fois par les autorités centrales et locales. Le gouvernement français réglemente les émissions polluantes par le biais

⁷En France, la Caisse des Dépôts était chargée de la mise en œuvre de la réglementation pour le pays.

d'instruments fiscaux et de normes réglementaires (Hauvuy & Riedinger, 2005; Bougon & Lavergne, 2019). Le principal instrument fiscal est la taxe générale sur les activités polluantes (*TGAP*) qui est une taxe nationale sur la quantité de polluants émise par an. En 2016, plus d'un millier d'installations étaient soumises à cette taxe (IGF, 2018). Cependant, la composante pollution atmosphérique de la *TGAP* est limitée et ne constitue pas l'instrument le plus efficace mis en œuvre pour limiter la pollution atmosphérique (Millock & Nauges, 2006; IGF, 2018).⁸ Si les grandes lignes de la réglementation des émissions industrielles sont établies au niveau national, leur mise en œuvre est généralement assurée par les autorités locales. Par exemple, le principal cadre réglementaire pour les activités polluantes est la réglementation *ICPE*, qui désigne les installations classées pour la protection de l'environnement. Le cadre *ICPE* est un ensemble de normes régissant les activités des usines polluantes en fonction de leur impact sur l'environnement. Elle impose notamment qu'une autorisation d'ouverture soit délivrée par l'autorité des préfets de département pour toute installation polluante sous condition de mise en œuvre de normes technologiques spécifiques et après réalisation d'une enquête locale auprès de la population. En outre, certaines de ces installations classées peuvent aussi relever de la directive européenne sur les émissions industrielles (IED).⁹ La principale obligation imposée par cette réglementation est d'adopter les MTD européennes : les autorisations d'installations industrielles sont conditionnées à l'utilisation des techniques les moins polluantes. À cet égard, la réglementation nationale actuelle se concentre principalement sur les plus grandes usines, mais toute usine polluante plus petite non couverte par la réglementation *ICPE* est réglementée au niveau de la municipalité. En outre, depuis la loi LAURE, adoptée en 1996, d'autres autorités locales peuvent prendre des mesures spécifiques pour améliorer la qualité de l'air locale et atteindre les objectifs nationaux en matière de qualité de l'air. Par exemple, depuis le début des années 2000, plusieurs "plans de protection de l'atmosphère" (PPA) ont été mis en œuvre dans différentes régions. Dans leur zone d'application respective, nombre de ces plans ont adopté des mesures différenciées pour les zones plus ou moins densément peuplées. Certains de ces plans ont introduit des normes environnementales plus strictes pour les industries manufacturières, notamment dans les agglomérations dépassant un certain seuil d'habitants. Par exemple, les trois *PPA* consécutifs de la région parisienne ont imposé des plafonds d'émissions industrielles de NO_x et de PM plus bas que les plafonds nationaux. Ces plans prévoient également des mesures d'urgence lorsque la concentration atmosphérique de certains pol-

⁸En 2016, 50 millions d'euros ont été collectés par rapport aux 3,8 milliards d'euros collectés pour la composante carbone de la taxe sur la consommation d'énergie en 2016 (DGEC, 2016)

⁹Depuis 2010, la directive sur les émissions industrielles a remplacé la directive sur la prévention et la réduction intégrées de la pollution. Ces directives sont similaires et visent à prévenir la pollution de l'air, de l'eau et du sol.

luants dépasse les normes nationales et européennes.

Cette différenciation des réglementations environnementales dans l'espace existe également dans d'autres régions. Aux Etats-Unis par exemple, le *Clean Air Act* prévoit des mesures spécifiques appliquées aux installations industrielles dans les comtés où les concentrations sont trop élevées. En Chine, dans le cadre de son plan quinquennal, le gouvernement fixe un objectif national de réduction des émissions industrielles qui est décomposé en objectifs infranationaux que les provinces et les villes se doivent d'atteindre. En particulier, les niveaux d'ambition de ces objectifs diffèrent beaucoup suivant les régions (Yamada, 2020). Dans ce contexte, le second chapitre de ma thèse s'intéresse aux conséquences de l'existence de réglementations différenciées régionalement. En effet, si des réglementations plus strictes peuvent contribuer à améliorer la qualité de l'air dans certaines villes, elles peuvent également réduire la productivité des industries locales et affecter négativement les revenus des travailleurs. J'aborde donc dans ce chapitre le compromis entre économie et environnement dans un contexte spatial afin d'identifier la meilleure façon de répartir les émissions de pollution entre des villes différentes.

Enfin, l'approche traditionnelle des politiques environnementales, dite "pigouvienne", consiste à estimer le coût que le dérèglement environnemental en question fait porter par la société. Dans le cas du changement climatique ou de la pollution de l'air, il s'agit du prix que les individus seraient prêts à payer afin d'éviter l'émission de GES ou de polluants atmosphériques nocifs. Ce prix correspond généralement à l'équivalent monétaire de la réduction d'utilité due à ces émissions. Dans le cas du changement climatique, des modèles alliant sciences climatiques et interactions économiques entre agents sur plusieurs générations ont notamment été élaborés afin de calculer un "coût social du carbone" correspondant au dommage marginal causé par l'émission d'une tonne de CO₂ supplémentaire.¹⁰ Cette approche doit alors permettre d'identifier le niveau adéquat d'une taxe carbone ou du prix des permis d'émission sur les marchés. Dans le cas de la pollution de l'air, l'approche est sensiblement différente. Les dommages sont liés aux concentrations de pollution, et non aux niveaux d'émissions. Actuellement, peu d'analyses font le lien entre le lieu d'émission, un modèle de circulation atmosphérique et la qualité de l'air locale. Il n'est donc généralement pas possible d'estimer un coût marginal d'émission de particules polluantes. L'approche traditionnelle est donc d'estimer le coût total qu'impliquent les concentrations observées, en calculant son impact monétaire sur la santé, les activités économiques et la biodiversité.¹¹ Cette valeur permet alors de fixer un niveau d'ambition des politiques mises en place. À l'heure actuelle de nombreuses analyses existent sur les

¹⁰ *Integrated Assessment Models* (IAM)

¹¹ Généralement, l'impact monétaire des effets sur la santé prend en compte à la fois les dépenses médicales impliquées et une forme de valorisation des pertes en espérance et en qualité de vie.

effets sanitaires de la pollution de l'air. Ces études montrent que la mauvaise qualité de l'air implique un coût important en termes de dépenses médicales mais aussi de qualité de vie dégradée et d'une réduction de l'espérance de vie (Waidyatillake et al., 2021). Cependant, une littérature récente s'attache à montrer que la pollution de l'air a également des effets directs sur les activités économiques. En particulier, une mauvaise qualité de l'air diminue la productivité des travailleurs industriels (Graff Zivin & Neidell, 2012; Chang et al., 2016; Adhvaryu et al., 2022; He et al., 2019). C'est dans ce contexte que s'inscrit le troisième chapitre de cette thèse qui estime les dommages causés par la pollution des particules fines (PM2.5) aux exportations françaises. L'estimation des conséquences de la pollution de l'air sur les entreprises pourrait avoir des implications en termes de politiques publiques. En effet, si les réglementations visant à améliorer la qualité de l'air sont généralement perçues comme un coût supplémentaire pour les activités économiques, la mise en lumière de bénéfices économiques directs pourrait amener à revoir à la hausse l'ambition de ces politiques.

Synthèse des chapitres

Chapitre 1 : Une évaluation du SEQE-UE

De nombreuses réglementations, telles que le SEQE-UE, ne s'appliquent qu'à certains sous-ensembles spécifiques d'entreprises au sein de secteurs industriels. Par exemple, certains gouvernements n'appliquent des normes environnementales qu'aux entreprises situées dans les régions les plus polluées (Greenstone, 2002; Fowlie et al., 2016), n'offrent des incitations à l'embauche qu'aux entreprises situées dans des zones en difficulté (Neumark & Simpson, 2015), et ciblent les subventions à l'activité vers les petites et moyennes entreprises (Martin et al., 2017; Rotemberg, 2019; Smagghue, 2019). La méthode la plus courante pour évaluer ces politiques consiste à comparer l'évolution des résultats des entreprises réglementées à ceux des entreprises non réglementées au moment du changement de politique - il s'agit de l'estimateur dit des "doubles différences" (Card & Krueger, 2000). Dans le cas où les résultats des entreprises réglementées suivent une tendance parallèle à celle des entreprises non réglementées avant la mise en place de la réglementation, on peut supposer que, en l'absence de cette réglementation, entreprises réglementées et non réglementées auraient suivi la même évolution (c'est l'hypothèse dite des "tendances parallèles"). Cependant, même si cette hypothèse est vérifiée, il est probable que, dans le cadre d'une concurrence imparfaite, le statut réglementaire d'une entreprise influence les résultats des autres entreprises de l'échantillon. Il s'agit d'une violation de l'une des hypothèses nécessaire à la validité de la méthode des doubles différences : la variable de

traitement d’une unité ne doit pas affecter le résultat d’une autre unité. Ce problème de contamination est largement reconnu dans la littérature, mais rarement traité empiriquement. Dans le cas du SEQE-UE, l’unité réglementée est l’installation, c’est à dire l’usine. En pratique, seules les installations dépassant certains seuils d’émissions de GES sont incluses et doivent justifier de leurs émissions annuelles en acquérant des permis d’émission. Ainsi, certaines entreprises peuvent posséder à la fois des installations réglementées et non réglementées. Cette caractéristique du SEQE-UE permet potentiellement la réallocation d’une partie de l’activité d’une usine à l’autre au sein d’une même entreprise. Il s’agit d’un second effet de contamination qui doit être pris en compte afin d’estimer les effets de cette politique.

Dans le premier chapitre de cette thèse, intitulé *Inter-Firm and Intra-Firm Spillover Effects of Industrial Regulation* et écrit conjointement avec Geoffrey Barrows, Raphael Calem et Hélène Ollivier, nous proposons une nouvelle méthode permettant d’estimer les effets d’une réglementation incomplète sur les ventes et les émissions de GES au niveau de l’entreprise, des usines et de l’ensemble de l’industrie, tout en tenant compte des effets de concurrence et de réallocation. Pour cela, nous spécifions un modèle d’offre et de demande de produits différenciés en adoptant des hypothèses standards : fonction d’utilité CES (*Constant Elasticity of Substitution*), fonction de production Cobb-Douglas, hétérogénéité de la productivité des entreprises et concurrence monopolistique.¹² Pour établir un lien entre le statut réglementaire à l’échelle de l’installation et les revenus à l’échelle de l’entreprise, nous adoptons également un modèle simple d’entreprises possédant plusieurs usines (chaque usine participe à la production d’un même bien à l’échelle de l’entreprise).¹³

Nous considérons l’introduction d’une politique ciblant le prix de l’un des intrants de production pour un certain sous-ensemble d’installations dans l’économie. Nous montrons ainsi que la variation de revenus d’une entreprise avant et après l’introduction de la politique dépend du statut réglementaire de cette entreprise (effet direct à travers les coûts de production de cette entreprise) mais aussi du statut réglementaire de toutes les autres entreprises du secteur (effet indirect se “propageant” à travers les indices de prix du secteur et des industries qui le composent). Dans un second temps, nous tirons parti de la struc-

¹²Ces hypothèses sont standards pour les modèles incluant l’ensemble des industries manufacturières (voir par exemple Atkeson & Burstein (2008)). Dans certains cas, lorsque l’étude porte sur une seule industrie, des modèles flexibles d’élasticités-prix inter-entreprises peuvent être estimés de manière structurelle en utilisant des informations détaillées sur les prix à la production et les caractéristiques des produits (Berry et al., 1995). Ces données ne sont généralement disponibles que pour des industries particulières. Ainsi, cette méthode n’est pas adaptée à l’estimation des effets agrégés de réglementations affectant des industries différentes (De Loecker & Syverson, 2021).

¹³Ce modèle implique que la concentration des ressources productives entre les usines d’une même entreprise peut soit augmenter soit diminuer les coûts à l’échelle de l’entreprise, selon que les forces de dispersion dominent les économies d’échelle.

ture du modèle pour identifier les coefficients qui déterminent les effets directs et indirects de changements du vecteur de réglementation sur les revenus et les émissions de GES de l'ensemble des entreprises. Cette étape consiste à estimer conjointement les effets direct et indirect de la réglementation à l'aide de la méthode des moments généralisée. L'effet direct est obtenu en régressant la variation de revenus à l'échelle de l'entreprise sur le changement de statut réglementaire de cette entreprise tout en incluant des effets fixes au niveau de l'industrie. Le coefficient estimé identifie alors l'effet direct de la réglementation sur les coûts de production (multiplié par une constante correspondant à une combinaison de paramètres de demande). En l'absence d'effets indirects, ce coefficient serait aussi l'effet moyen de la réglementation sur les entreprises réglementées (*Average Treatment on the Treated* ou ATT). Cependant, la compétition étant imparfaite (les variations de prix de vente d'une entreprise affectent les revenus de ses concurrents) cette régression ne permet pas d'identifier l'ATT. Les effets fixes inclus dans cette équation capturent les variations d'indice de prix à l'échelle de l'industrie. Il est alors possible d'estimer l'impact indirect de la réglementation en régressant ces effets fixes par industrie sur une moyenne pondérée des statuts réglementaires des entreprises de cette industrie (tout en contrôlant pour les caractéristiques du secteur).

Intuitivement, si la réglementation augmente les coûts, alors les industries ayant une densité de réglementation plus élevée devraient voir une croissance plus forte de leurs indices de prix. Des travaux antérieurs construisent des mesures similaires en utilisant la densité réglementaire pondérée (Cai & Szeidl, 2022; Muehlegger & Sweeney, 2021; Rotemberg, 2019). Toutefois, une caractéristique clé de notre procédure est l'utilisation de la théorie des indices des prix afin d'identifier une pondération correcte.¹⁴ Enfin, montrons comment résoudre l'équilibre dit "contrefactuel" correspondant à l'absence de réglementation à l'aide de ces seuls coefficients ainsi que de l'équilibre observé. La comparaison de ces deux équilibres, observé et contrefactuel, permet d'identifier l'impact sur chaque entreprise et installation industrielle de l'introduction de la politique. Il est alors possible de calculer des valeurs d'ATT et d'ATC (*Average Treatment on the Controls*, c'est-à-dire l'effet moyen de la réglementation sur les entreprises et installations non réglementées).

Après avoir montré à l'aide de simulations numériques que cette méthode produit des estimations non biaisées, nous l'appliquons à l'évaluation de l'impact du SEQE-EU sur l'industrie française entre 2005 et 2016. Pour cela, nous faisons appel à plusieurs sources de données. Tout d'abord, nous utilisons les déclarations fiscales (confidentielles) de revenus et de ventes à l'export de l'ensemble des entreprises manufacturières françaises (base de données FARE/FICUS). Chaque entreprise est identifiée par un code admin-

¹⁴En particulier, l'hypothèse de demande CES implique des poids spécifiques par entreprises reflétant à la fois la part de marché avant et après l'introduction de la réglementation (Sato, 1976; Vartia, 1976).

istratif unique, le SIREN, qui nous permet de la suivre dans le temps et de combiner des observations provenant de bases de données différentes. Ensuite, nous mesurons la consommation d'énergie à l'aide des enquêtes EACEI (*Enquête Annuelle sur la Consommation d'Energie dans l'Industrie*). Ces enquêtes comprennent toutes les dépenses et consommations liées à l'énergie, avec des détails sur plusieurs types d'énergie et de combustibles (quantités consommées et dépenses), au niveau de l'usine. Nous utilisons ces données, combinées aux facteurs d'émission par type d'énergie de l'Agence française de l'environnement et de l'énergie (ADEME), pour calculer les émissions de GES (en équivalent CO₂) des installations industrielles. Enfin, la Commission européenne publie la liste des installations incluses dans le SEQUE-UE et les informations leur correspondant sur une plate-forme publique dédiée (*l'European Union Transaction Log*). À partir de cette liste, nous obtenons le SIRET de chaque installation réglementée en France ainsi que la date initiale de l'application de la réglementation à cette installation.

L'utilisation de ces données et de la méthodologie décrite ci-dessus nous permet de montrer que le SEQUE-UE a augmenté les revenus des entreprises réglementées (c'est-à-dire des entreprises ayant au moins une usine réglementée) et a réduit les émissions des usines réglementées. Le fait que la réglementation ait augmenté les ventes peut être surprenant, car elle consiste, en pratique, en une augmentation du coût de l'énergie pour les entreprises réglementées. Toutefois, les réglementations environnementales peuvent aussi déclencher des investissements permettant de réduire les coûts et d'augmenter les ventes (Porter & van der Linde, 1995; Ambec et al., 2013).¹⁵ En particulier, les résultats de Cael & Dechezlepretre (2016) et Cael (2020) indiquent que le SEQUE-UE a permis l'augmentation de la R&D et de "l'innovation verte", ce qui pourrait expliquer ces effets positifs sur les revenus. Nous constatons également qu'au cours de ses deux premières phases, le SEQUE-UE a permis de réduire les émissions totales de CO₂ des installations françaises d'environ 5 à 12 % par an selon l'année.

Chapitre 2 : Comment réglementer la pollution industrielle dans l'espace?

Il est désormais largement reconnu que la pollution atmosphérique cause des dommages substantiels à la santé humaine.¹⁶ Pourtant, la plupart des gens vivent dans de grandes villes où se concentrent également certaines industries polluantes. Afin d'améliorer la qualité de l'air dans ces zones, de nombreux gouvernements ont mis en place des régle-

¹⁵Cette hypothèse est d'ailleurs étayée par d'autres analyses du SEQUE-UE (Löschel et al., 2018; Dechezlepretre et al., 2018; Calligaris et al., 2018).

¹⁶Les dommages à la santé ont été mesurés même lorsque les niveaux de pollution sont inférieurs aux normes réglementaires (Graff Zivin & Neidell, 2012, 2013; Deryugina et al., 2019).

mentations particulières dans les agglomérations les plus peuplées et les plus polluées. Cependant, des réglementations plus strictes peuvent nuire aux entreprises locales, augmenter les coûts de production et affecter le salaire réel local.¹⁷ Cet arbitrage local entre revenus et pollution soulève la question de la différenciation des réglementations environnementales dans des lieux hétérogènes?

Dans le deuxième chapitre de cette thèse, intitulé *An Optimal Distribution of Polluting Activities Across Space*, je modélise les choix de localisation d'entreprises polluantes et des travailleurs affectés par la pollution afin d'identifier l'ensemble des politiques qui mettent en œuvre une répartition optimale des travailleurs, des entreprises et de la pollution dans les villes. Dans ce modèle, les choix de localisation effectués par les entreprises et les travailleurs individuels dépendent (1) du niveau "d'aménités" offert par les villes,¹⁸ (2) des externalités d'agglomération qui procurent un avantage de productivité aux entreprises situées dans les zones plus peuplées et (3) des forces de congestion. En particulier, la pollution atmosphérique agit comme une force de congestion : l'augmentation des émissions de particules polluantes liée à l'activité économique locale diminue la qualité de l'air, et les travailleurs réagissent en s'éloignant des zones les plus polluées. Les entreprises, en revanche, n'internalisent cet effet négatif de leurs émissions sur la population locale que dans la mesure où elles doivent payer une taxe sur ces émissions. Elles émettent alors trop de pollution dans les plus grandes villes et cela conduit à une diminution de la main-d'œuvre disponible, car les travailleurs quittent ces villes trop polluées. Ainsi, les niveaux de production et la population sont inférieurs à leurs niveaux optimaux dans les villes les plus grandes et les plus productives. Du fait de ce mécanisme, lorsque les politiques environnementales ne ciblent pas les grandes villes en particulier, celles-ci sont trop petites par rapport à "l'optimum social". Je montre ensuite qu'un ensemble de politiques aboutissant à un coût d'émission de pollution plus élevé dans les grandes villes permettrait de rendre les activités locales plus "propres", la qualité de l'air meilleure et ces villes plus attrayantes. Enfin, j'applique une version généralisée de ce modèle au cas de la répartition des activités polluantes à travers les zones d'emploi françaises. Je constate que les réglementations actuelles sont plus strictes dans les grandes villes, conformément à ce que le modèle préconise. Néanmoins, le bien-être pourrait augmenter en ciblant davantage certaines grandes villes.

Dans ce modèle, je considère un pays unique avec un nombre discret de villes qui diffèrent par leurs niveaux de productivité et d'aménités. L'utilité des travailleurs habitant dans une ville est une fonction des aménités locales ainsi que de leur niveau de consommation

¹⁷C'est l'argument des "*jobs versus the environment*" (Morgenstern et al., 2002).

¹⁸Ce terme d'aménités désigne "un ensemble de valeurs matérielles et immatérielles attachées à nos territoires et qui « marquent » leur attractivité" (Rivière, 2010).

d'un bien industriel composite. Ce niveau de consommation dépend du salaire local ainsi que de l'indice de prix local. Suivant l'approche développée par Rosen (1979) et Roback (1982), je considère une population fixée de travailleurs identiques qui choisissent à l'équilibre dans quelle ville habiter afin de maximiser leur utilité. Dans chaque ville, ces travailleurs sont employés par des entreprises appartenant à des secteurs distincts et produisant des variétés différenciées. Ces entreprises utilisent la main-d'œuvre et les émissions polluantes comme intrants substituables pour la production.¹⁹ Les coûts de production dépendent donc des salaires locaux, de la productivité locale du travail ainsi que des taxes locales sur les émissions.

Cette structure permet l'introduction des deux externalités spatiales qui permettent traditionnellement d'expliquer la répartition de l'activité économique dans l'espace. Premièrement, la productivité locale incorpore le phénomène d'économies d'agglomération : lorsque le nombre de travailleurs augmente, la productivité locale augmente également. Ces économies peuvent résulter d'effets de diffusion des connaissances, de la mise en commun de la main d'œuvre locale ou des économies d'échelle locales.²⁰ Deuxièmement, le bien-être des habitants d'une ville dépend d'effets de congestion : lorsque la population locale augmente *ceteris paribus*, le bien-être local diminue. Cette relation peut-être due, par exemple, à l'augmentation des prix de l'immobilier ou des temps de trajet sur le réseau routier. Potentiellement, cet effet incorpore également l'effet négatif de la pollution due aux secteurs non industriels, comme la consommation d'énergie résidentielle ou les transports, et qui ne sont pas explicitement modélisés. Enfin, je fais l'hypothèse que les travailleurs peuvent déménager d'une ville à l'autre en réaction à un changement d'équilibre, sans encourir aucun coût. La nouveauté de l'article repose sur l'hypothèse que la quantité totale d'émissions industrielles par ville affecte négativement le niveau local d'aménités.²¹

Comme je le montre dans ce deuxième chapitre, ce cadre théorique est particulièrement applicable au cas de la France. Premièrement, une grande partie des émissions de polluants nocifs résulte des activités industrielles en France. Deuxièmement, les concentrations atmosphériques locales de ces polluants dans plusieurs villes françaises sont souvent supérieures aux recommandations de l'Organisation mondiale de la santé (OMS). Troisièmement, les réglementations sur la qualité de l'air mises en œuvre au cours des dernières décennies en France ont tendance à être plus strictes dans les grandes villes.

¹⁹La pollution peut être modélisée de manière équivalente comme un intrant ou un sous-produit de la production industrielle dans une configuration similaire à Copeland & Taylor (2004).

²⁰Combes & Gobillon (2015) fournit un aperçu de la littérature empirique qui documente les économies d'agglomération.

²¹Il s'agit d'un effet qui a été identifié empiriquement (Chay & Greenstone, 2005; Davis, 2011; Khanna et al., 2021), mais pas encore inclus dans une configuration quantitative.

À l'aide d'un large jeu de données sur la répartition de l'activité à travers les zones d'emplois françaises et de données d'émission de pollution et de qualité de l'air, j'estime les élasticités clés qui déterminent l'intensité des externalités de pollution, les économies d'agglomération et les effets de congestion. L'utilisation de données à l'échelle des installations industrielles me permet également d'estimer, pour chaque secteur industriel, l'élasticité régissant la substitution entre émissions et main-d'œuvre dans la production industrielle. L'observation de la répartition actuelle de la main d'œuvre, des salaires ainsi que des émissions industrielles de PM2.5 me permet de calculer les vecteurs d'aménités, de productivité et de coût d'émission qui expliquent cette répartition. En particulier, je constate que le coût d'émettre des particules PM2.5 est plus élevé dans les villes les plus productives, mais pas nécessairement dans les villes dotées des aménités les plus importantes. Je caractérise ensuite les implications de cette distribution spatiale des coûts marginaux des émissions en termes de bien-être. Je montre que des réglementations plus strictes en matière d'émissions dans les grandes villes leur ont permis de réduire leurs émissions polluantes et de devenir plus grandes que ce qu'elles auraient été dans le cadre d'une réglementation uniforme. Pour finir, j'identifie enfin la distribution spatiale de coûts d'émissions qui maximiserait le bien-être des travailleurs : il serait optimal d'augmenter ces coûts dans les villes à haut niveau d'aménités.

Chapitre 3 : La pollution de l'air affecte-t-elle les exportations françaises?

La pollution de l'air n'affecte pas seulement la santé des individus (Aguilar-Gomez et al., 2022). En effet, une vaste littérature a montré les effets néfastes de la pollution sur la productivité des travailleurs (Graff Zivin & Neidell, 2012; Chang et al., 2016; He et al., 2019; Adhvaryu et al., 2022), sur leurs capacités cognitives (Allen et al., 2016; Heyes et al., 2016; Archsmith et al., 2018; Chang et al., 2019) et leur participation au marché du travail (Hanna & Oliva, 2015; Holub et al., 2021). Néanmoins, les conséquences de ces effets sur les performances économiques des entreprises dans lesquelles les individus affectés par la pollution travaillent sont encore peu connues. À notre connaissance, seules quelques études relient les chocs de productivité dus à la pollution aux résultats économiques agrégés (Dechezleprêtre et al., 2019; Fu et al., 2021; Khanna et al., 2021). Si la pollution affecte négativement la productivité des travailleurs et la quantité de main-d'œuvre disponible, la théorie prédit alors une augmentation des coûts de production des entreprises menant potentiellement à une réduction de leur production et de leurs revenus. La mise en lumière empirique de ce type de phénomènes pourrait donc révéler que l'amélioration de la qualité de l'air a un impact économique direct positif en plus des bénéfices sanitaires considérés

généralement.

Dans le troisième chapitre de cette thèse, intitulé *The Effects of Air Pollution on Exports: Evidence from PM2.5 & French Firms*, co-écrit avec Geoffrey Barrows et Hélène Ollivier, nous fournissons une estimation causale des effets de la pollution atmosphérique sur l'activité des entreprises françaises, et en particulier sur leurs exportations. Nos résultats indiquent que des concentrations plus élevées de PM2.5 ont un effet négatif à la fois sur la valeur totale et sur la quantité des exportations à l'échelle de l'entreprise. En particulier, nous nous intéressons à l'ensemble des entreprises françaises exportatrices, entre 2009 et 2015, et constatons qu'une augmentation de 10 % des concentrations de PM2.5 auxquelles sont confrontés les établissements détenus par une entreprise exportatrice entraîne une diminution d'environ 1,5 % de la valeur totale et de la quantité des exportations de cette entreprise.

Pour obtenir ces résultats notre analyse s'appuie sur des observations exhaustives et très désagrégées, du point de vue temporel et géographique, sur les données commerciales des entreprises et la qualité de l'air local. Nous observons les exportations mensuelles au niveau des entreprises françaises et les combinons avec des informations sur les coordonnées géographiques des établissements détenus par les entreprises exportatrices ainsi qu'avec des variables environnementales locales : les concentrations atmosphériques de PM2.5, la direction du vent, les précipitations et les températures.²² L'utilisation de données issues de la modélisation du transport et des processus chimiques relatifs aux particules polluantes présentes dans l'atmosphère permet une approche plus précise que celle développée dans de précédents travaux s'appuyant sur les mesures de pollution par des stations de mesure potentiellement localisées dans des lieux spécifiques et non représentatifs. De plus, notre panel d'entreprises exportatrices est issu des déclarations obligatoires aux Douanes françaises, ce qui signifie que nous observons l'ensemble des entreprises françaises exportatrices, et non une sous-population comme cela est fréquent dans la littérature. Un autre avantage de nos données commerciales est leur disponibilité à une fréquence mensuelle. En effet, la plupart des travaux sur les effets de la pollution à l'échelle individuelle montrent que ces effets se produisent sur plusieurs jours ou semaines. L'utilisation de données annuelles est susceptible de cacher des effets à court terme, en particulier si les entreprises compensent les effets négatifs à court terme de la pollution en augmentant les exportations au cours des périodes ultérieures.

Les deux obstacles principaux à l'estimation d'un effet causal de l'augmentation de la pollution sur l'activité économique sont l'*omitted variable bias*, ou "biais dû à une variable

²²La plupart des activités de production se déroulent en intérieur (Vette et al., 2001), nous nous concentrons donc sur la pollution due aux PM2.5 qui est composée de particules suffisamment fines pour pénétrer les bâtiments.

omise”, et la *reverse causality*, ou “relation de causalité inverse”. Dans notre contexte, l’*omitted variable bias* provient de facteurs locaux qui influencent à la fois les concentrations de pollution atmosphérique et l’activité économique. Une importante littérature a mis en lumière l’effet de la température sur la santé, la main-d’œuvre disponible et la productivité des travailleurs (Dell et al., 2014; Graff Zivin & Neidell, 2014; Addoum et al., 2020; Aguilar-Gomez et al., 2021). Cependant, la température influence également les concentrations de pollution. Cette double dépendance risque de biaiser les estimations si les variations de température ne sont pas incluses dans l’analyse. Par exemple, étant donné que des températures très élevées entraînent des concentrations plus importantes de polluants atmosphériques, l’omission de l’effet de la température entraînerait une surestimation de l’effet négatif de la pollution sur les exportations (une partie de l’effet négatif de la température sur les exportations serait attribuée à tort à la pollution plus élevée). Afin de contrôler l’influence des facteurs climatiques, nous combinons donc nos données de pollution avec des données sur les températures et les précipitations.

L’obstacle de la *reverse causality* est dû à l’endogénéité de la pollution de l’air : un choc positif de productivité dans un endroit donné provoque potentiellement des niveaux de production et de consommation plus élevés mais aussi des déplacements plus nombreux, émettant ainsi davantage de polluants et augmentant leur concentration dans l’atmosphère. Cela induit une corrélation positive entre productivité locale et pollution qui ne reflète pas l’effet d’un choc de pollution exogène sur l’activité économique. Afin de surmonter ce problème, nous mettons en place une stratégie de “variable instrumentale” où nous exploitons les variations exogènes de la direction du vent dans les zones où sont situées les établissements appartenant aux entreprises exportatrices. Nous montrons que la direction du vent influence fortement les concentrations locales de PM_{2.5} et, après avoir contrôlé les effets de saisonnalité, les tendances annuelles et les variables météorologiques, nous faisons l’hypothèse que les changements locaux et à court terme de la direction du vent n’affectent les activités économiques que par leurs effets sur les concentrations locales de pollution. En pratique, dans chaque lieu où se trouvent des établissements appartenant à des entreprises exportatrices, nous commençons par estimer des modèles dans lesquels les concentrations mensuelles de PM_{2.5} sont prédites linéairement par la direction du vent, la température, les précipitations, le mois civil (pour contrôler la saisonnalité) et l’année (pour inclure les tendances temporelles). Nous calculons ensuite les valeurs de pollution prédites par ces modèles en fonction de la direction du vent uniquement. Enfin, ces niveaux de pollution prédits sont utilisés comme instruments d’une stratégie *two-stages least squares* dans laquelle nous instrumentons les concentrations de PM_{2.5} auxquelles sont exposées les entreprises (calculées comme une moyenne pondérée des concentrations de PM_{2.5} observées dans les lieux où ces entreprises ont des établissements).

Sous réserve de contrôler les effets des autres variables météorologiques afin de résoudre l'*omitted variable bias* et satisfaire la condition d'*exclusion restriction* de notre instrument (l'instrument n'influence la variable expliquée que par le biais de la variable instrumentée), notre stratégie nous permet d'estimer un coefficient non biaisé mesurant l'effet causal de la concentration de PM2.5 sur les exportations.

General Introduction

Our society is confronted with numerous environmental unbalances caused by human activities. Two of these major issues are addressed in my thesis : the reduction of greenhouse gas (GHG) emissions, in order to limit global warming, and the consequences of air pollution on human health and activities. The role of industrial sectors is particularly important on these two issues, justifying the implementation of numerous public policies aimed at reducing their GHG emissions and other local pollutants. These policies must be developed in a coherent, efficient and socially acceptable way. Economics research can contribute to this process by studying policy options retrospectively and prospectively. In a context of rationalization of public policies, their evaluation is necessary in order to measure the impact of public action. Indeed, a precise knowledge of the consequences of the implementation of regulations aiming at improving air quality and reducing GHG emissions is a prerequisite for their acceptability by society, as well as for their future improvement.

Climate Change, Air Pollution and Environmental Policy

Two major environmental issues

Greenhouse Gases and Climate Change: There is now a strong body of evidence of the impact of human activities on the climate. The burning of fossil fuels, such as oil and gas, the production of goods, travel, or the generation of electricity emit GHG that contribute to the warming of the atmosphere as they accumulate in the atmosphere. In practice, the average temperature at the Earth's surface between 2011 and 2020 has increased by more than 1°C compared to the period 1850-1900 (IPCC, 2021a). The latest report from the Intergovernmental Panel on Climate Change (IPCC) shows that "...unless there are immediate, rapid, and massive reductions in greenhouse gas emissions, limiting warming to around 1.5°C, or even to 2°C, will be out of reach." (IPCC, 2021b). Beyond temperatures, this climate change impacts other variables, such as humidity and drought, winds, snow and ice, coastal areas and oceans. In particular, the frequency of climatic disasters (hurricanes, floods, heat waves, droughts, etc.) could increase if the temperature increase is not controlled.

In order to limit this temperature increase, the consensus today is that we must reduce total GHG emissions, or at least stop their increase. In 1997, 35 countries made commitments at the Kyoto Summit to reduce their GHG emissions by 5%. In 2015, the Paris Climate Agreement was signed by more than 50 countries (representing more than half

of the world’s GHG emissions) committing to limit the global temperature increase to 2°C.²³ In 2020, in the context of the Covid pandemic, global emissions have declined by 5.4%. However, they appear to have immediately returned to “pre-Covid” values and the concentration of GHGs in the atmosphere continues to increase to this day (UNEP, 2021). In particular, GHG emissions over the last decade (2010 to 2019) have been larger than ever before (IPCC, 2022). This amount emitted is of the same order of magnitude as the “carbon budget”, which is the maximum amount that can still be emitted while limiting the global temperature increase to 1.5°C. For the moment, France is committed to reducing its GHG emissions by 40% by 2030 and to achieving carbon neutrality by 2050 (which corresponds to a division by 6 of its emissions compared to 1990).²⁴ French GHG emissions (CO2 equivalent) have also decreased by 20% between 1990 and 2019 (CITEPA, 2021). These commitments, as well as the climatic emergency, pose a challenge today and require the implementation of a sufficient transition in order to reach carbon neutrality (Ademe, 2021).

Air Pollutants and Health Effects: Another negative consequence of economic activities is the emission in the atmosphere of pollutants that are harmful to health. This is recognized to be “the most important environmental threat to human health” by the World Health Organization (WHO) in its latest air quality guidelines. This threat is due to the negative health effects of a series of pollutants emitted by various activities. The main ones are fine particulate matters (PM10 and PM2.5), nitrogen dioxide (NO2), ozone (O3), and sulphur dioxide (SO2).²⁵ The effects of these pollutants are multiple and affect different physiological parameters, such as the respiratory system (Viehmann et al., 2015), the cardiovascular system (Brook et al., 2010) or the brain (Costa et al., 2019). Early exposure to polluted air may also have latent effects, that is to say effects that do not manifest themselves until later in the lives of affected children (Bale et al., 2010; Almond & Currie, 2011; Colmer & Voorheis, 2020). It is also estimated that this air pollution is the cause of 4.4 million premature deaths in the world, all pollutants combined, and that pollution linked to PM2.5 causes about 379,000 premature deaths in the European Union and 40,000 in France. A key difference between GHGs and harmful air pollutants is that, for the latter, the place of emission is important for measuring their impact on health. Indeed, while global warming is a function of the total amount of accumulated GHG emissions - hence the designation of “global” pollutants - the harmful effects of air pollutants occur when they are breathed in by individuals - hence the designation of “local”

²³Temperature changes are measured relative to pre-industrial temperature levels.

²⁴Carbon neutrality consists, on the one hand, in reducing GHG emissions as much as possible, and on the other hand, in offsetting the remaining emissions.

²⁵Fine particulate matters are generally classified into two families corresponding to the main dimension of the particle: less than 10 μ m for PM10 and 2.5 μ m for PM2.5.

pollutants. Thus it is their local atmospheric concentration that must be controlled.

Since the first disasters (*Donora Smog* in 1948 in the United States, the *Great Smog* of London in 1952 or the accident of Bhopal in 1984 in India) and the first evidence of the effects of this type of pollution, brief or chronic, many regulations were set up in order to try to guarantee a better air quality.²⁶ At the international level and since 1979, the Convention on Long-range Transboundary Air Pollution, opened under the aegis of the United Nations, aims to reduce pollution related to SO₂, NO₂ and PM, as well as certain heavy metals and persistent organic pollutants. For the European Union, the strategy for a better air quality is based on three main pillars: (1) the Ambient Air Quality Directive (2008/50/EC), which sets air quality standards, in practice concentration levels not to be exceeded, as well as measurement and planning obligations in the Member States (1996 Framework Directive on Ambient Air Quality Assessment and Management (96/62/EC)); (2) the Directive on National Emission Reduction Commitments (Directive 2001/81/EC and then 2016/2284); and (3) the establishment of specific standards and limits by emission source, targeting in particular the industrial sectors (with for example Directive 2001/80/EC on emissions from large combustion plants or Directive 2010/75/EU on industrial emissions).

Over the past two decades, air quality has greatly improved and concentrations of a number of pollutants have decreased significantly in developed countries (in the United States (US EPA, 2022) and the European Union (AEE, 2013, 2020), for example). However, WHO estimates that 92% of the world's population continues to be exposed to excessively high air concentrations of pollution (WHO, 2016). Although the highest concentrations are observed in developing countries, such as India and China, the health values set by the WHO are still not reached in some developed countries. In the case of the EU, despite the reduction in the concentration of a number of pollutants in recent years, concentrations of fine particles (PM_{2.5}) and ozone (O₃) remain very high (Sicard et al., 2021). Moreover, PM_{2.5} also have effects on health below the health thresholds fixed (Di et al., 2017). France was one of the first countries to set up regulations aiming at limiting this problem (Massard-Guilbaud, 1999). Thus, outdoor air quality has improved over the last decades and emissions of all pollutants have been decreasing since 1990 (CITEPA, 2021). However, in many cities, EU standards are not yet reached for fine particles (PM₁₀) and nitrogen dioxide (NO₂). France has been condemned by the Court of Justice of the EU for having systematically exceeded the concentration limits for these pollutants.

²⁶Depending on the particles, an exposure to too high concentrations in a chronic or temporary way can have harmful effects on the respiratory system, the cardiovascular system, the skin, the fertility and the life expectancy

The role of the industry in these emissions: GHGs and air pollutants are emitted by human activities, mainly by the combustion of fossil fuels.²⁷ In particular, the transport sector, the industry, agricultural activities and residential and tertiary building sectors play a major role in these emissions. For GHGs, in 2019, these sectors represented respectively 31% of total emissions for transport, 19% for the manufacturing industry and construction, 19% for agriculture, 17% for residential and tertiary, 10% for energy, and 4% for waste treatment (CITEPA, 2021). The energy industry has seen its emissions reduced by 46% and the manufacturing industry by 41%, thus highly contributing to the reduction of national emissions between 1990 and 2019 (CITEPA, 2021). According to the *Réseau Action Climat*, emissions from the French industry have decreased between 1990 and 2015, due to improvements in energy efficiency but also through a reallocation of emissions to foreign countries, with plant closures, relocation and increased imports. Thus, if the national carbon footprint has remained at a stable level for the last twenty years, it is because “imported emissions” have been increasing continuously and are now even greater than emissions on French territory (HCC, 2020).²⁸

Concerning air pollution, the manufacturing industry is one of the major contributors, along with the transport sector. In 2019, the manufacturing industry contributed to 55% of national SO₂ emissions, 27% of PM₁₀ emissions, 18% of PM_{2.5} and 13% of NO_x emissions. Since 1990, SO₂ emissions have decreased by 86%, PM₁₀ by 51%, PM_{2.5} by 58%, and NO_x by 55% (CITEPA, 2021). For the US, emissions of these pollutants also decreased significantly since 1990. Shapiro & Walker (2018) showed that this reduction was mainly due to the intensification of environmental policies (more than to productivity or trade flow reallocation effects).

Thus, both the reduction of GHG emissions and the improvement of air quality continue to be, in France as in the rest of the world, a central issue for the years to come. Public policies implemented by governments are powerful instruments to solve these problems (Shapiro, 2022). However, there are still many questions about the respective advantages of these instruments and their effectiveness in environmental and social terms.

Environmental Policies

The emission of pollutants into the atmosphere (GHG, SO₂, NO_x, PM, etc.) is a negative externality: the action of one economic agent affects another without any compensation taking place. This market failure leads to a lower welfare of the economic agents because

²⁷Some local pollutants are also emitted during specific processes, such as the friction of tires on the road in transport or emissions due to chemical reactions in the industry

²⁸These emissions correspond to the emissions occurring abroad during the production of goods imported and consumed in France

the cost of the damages caused by the action in question is not taken into account by the polluter in her or his decisions (quantities produced, investments, etc.). The quantity of pollution emitted is therefore too high in relation to the social optimum. In this situation, public intervention can internalize the damage caused by pollution and thus increase social welfare. Traditionally, these interventions can be separated into two distinct families. The first consists of setting technology or health standards and making companies respect them. This is the so-called “command-and-control” approach (e.g., limits on emissions per company or technology standards). However, economists consider that, in certain cases, it is possible to correct a negative externality more effectively by setting an appropriate price. This is the second family of policies, known as “market-based” policies. With regard to pollution emissions, the two most widespread instruments are taxes (or subsidies) and emission permits markets. A tax set at the marginal cost of the damages linked to pollution (climate change, health impacts of air pollutants, etc.) allows firms to internalize the impact of their production, *i.e.* to take into account the environmental effects of their decisions (Pigou, 1920). Emission permit markets, called *cap-and-trade* instruments, are a second type of efficient instruments. The idea is to set the total quantity of emissions allowed, for example in order to respect the carbon budget or to reach a certain level of emission reduction, and then create for each unit of pollution that can be emitted an emission right that firms can trade on a market created for this purpose.

In practice, the choice between different instruments for regulating environmental externalities is still a key issue today (Bretschger & Pittel, 2020). The command-and-control approaches have often been favored because they are perceived as clearer. The most polluting technologies are banned in favor of cleaner technologies. For example, in the United States, this is the model on which the Clean Air Act and the Water Air Act were designed. This is also the case in the EU where industrial installations have to adopt specific technological standards called the “best available techniques” (BAT). Many examples of pollution taxes also exist, such as the tax on polluting activities (TGAP) in France or the tax on NO_x in Sweden or Norway (Bonilla et al., 2017). However, when the actors are heterogeneous and there is uncertainty about their respective costs of reducing their emissions, cap-and-trade instruments are considered more effective (Newell & Stavins, 2003). This type of instrument has notably been implemented in the United States (for instance with the Acid Rain program, the Los Angeles *RECLAIM* program²⁹, and the California carbon market), in the United Kingdom, and in China.

The European Union Emissions Trading Scheme (EU-ETS), the subject of the first chapter of this thesis, is also one of the main examples of such instruments to this day. It has been implemented by the European Commission, regulating approximately 11,000 power

²⁹ *California Regional Clean Air Incentives Market*

and manufacturing plants across 31 participating countries, including France. The EU-ETS was first announced in 2000 with a working paper from the European Commission, formalized with the Emissions Trading Directive in 2003, and initiated on January 1st, 2005. It regulates all combustion installations with a rated thermal input greater than 20MW, and other productive processes with capacity or output greater than predetermined industry-specific thresholds. Regulated firms must surrender permits – called EU allowance units – for each ton of CO₂ emitted during the year. Non-compliance triggers a fine of 100 euros per ton of CO₂ exceeding the number of permits surrendered. The European Commission initially designed three phases of the EU ETS. The first phase was designed as a trial period, running from 2005 to 2007; the second phase ran from 2008 to 2012, corresponding to the first commitment period under the Kyoto Protocol; and the third phase ran from 2012 until 2020.³⁰ During Phases I and II, the implementation of the policy was relatively decentralized, with participating countries submitting National Allocation Plans (NAPs) to the European Commission in which they proposed lists of installations to regulate, potential emission caps and allowance distribution rules (Ellerman et al., 2015). During these two phases, free permits were allocated to regulated installations on the basis of their historical emissions (according to the “grandfathering” method).³¹ In Phase III, allocation began a transition towards an open auction system, first for power installations (except for a few plants in Eastern Europe), and progressively for other regulated industries. Furthermore, a single emission cap (declining at a 1.74% rate per year) was adopted for the entire EU market and EU-wide sector-specific “benchmarks” were adopted for determining allowance distribution rules.

There are several impact evaluation studies of the EU-ETS (Petrick & Wagner, 2014; Löschel et al., 2018; Colmer et al., 2020b; Dechezleprêtre et al., 2018; Jaraite & Di Maria, 2018; Calligaris et al., 2018). However, revisions of the system for future phases are subject to intense oppositions, generally motivated by the idea that a stringent regulation could threaten the competitiveness of European industries (Demailly & Quirion, 2006), (Branger et al., 2017), (Joltreau & Sommerfeld, 2018). Industrial lobbies, in particular those representing the interests of the steel, cement and fertilizer sectors, play an important role in the discussions related to the revision of the system and these oppositions now threaten the sustainability of the environmental action (Thomas, 2021). In this context, an evaluation of the EU-ETS that takes into account not only its impact on GHG emissions but also its impact on the economic performance of regulated firms is needed. This is one of the objectives of the first chapter of this thesis. Indeed, we calculate both the GHG reduction implied by the EU-ETS in France but we also show that this regulation

³⁰Revisions of the regulation in the expectation of Phase IV (2021-2030) are already well under way.

³¹In France, the Caisse des Dépôts was in charge of the implementation of the regulation for the country.

has had a positive effect on the revenues of regulated firms.

Air quality regulations aiming at decreasing industrial pollutants emissions are also complex and there is a large debate on the consequences. In the case of France, air quality regulation is based on European standards. Limits for maximum air concentration of various pollutants and reduction objectives are set at the EU level (directives 2004/107 and 2008/50/CE) to be enforced in each Member States. In turn, they must implement specific action plans to remain within these limits and reach the air quality objectives. In France, these actions are taken both by central and local authorities. The French government regulates pollutant emissions through tax instruments and regulatory standards (Hauvuy & Riedinger, 2005; Bougon & Lavergne, 2019). The main tax instrument is the General Tax on Polluting Activities (*TGAP*), which is a national tax on the quantity of pollutant emitted per year. In 2016, more than a thousand installations were subjects to this tax (IGF, 2018). However, the air pollution component of the *TGAP* is limited,³² and does not constitute the most efficient instrument enforced to limit air pollution (Millock & Nauges, 2006; IGF, 2018). Frameworks regulating industrial emissions are mainly set at a national level. However, their implementation is usually enforced by local authorities. For instance, the main regulatory framework for polluting activities is the *ICPE* regulation, standing for Plants Classified for the Protection of the Environment. The *ICPE* is a set of norms governing polluting plants activities in relation with their impact on the environment. Notably, it requires that an opening permit is delivered by the authority of the departmental prefects for any polluting plant under the condition of implementing specific technology standards and after conducting a local survey of the local population.³³ Moreover, some of these classified plants may fall under the EU Directive on Industrial Emissions (IED).³⁴ This is the case if they are running polluting installations with capacities above thresholds set in the IED. The main obligation enforced under this regulation is the BAT: authorizations for industrial installations are conditional to the use of the least pollution intensive techniques. In that respect, the current national regulation focuses mainly on the largest plants, but any smaller polluting plant not covered by the *ICPE* regulation is regulated at the municipality level. Furthermore, since the LAURE law, adopted in 1996, other local authorities can use specific measures to improve local air quality and reach air quality national targets. For instance, starting in the early 2000's, several "Atmospheric Protection Plans" (*PPA*) have been implemented in different ar-

³²In 2016, €50 millions were collected compared to the €3.8 billions collected for the carbon component of the energy consumption tax in 2016 (DGEC, 2016)

³³These prefects are the regional authorities for the "*départements*" which are NUTS 3 geographic units (larger than commuting zones).

³⁴Since 2010, the Industrial Emissions Directive has replaced the Integrated Pollution Prevention and Control Directive. The guidelines are similar and aim at preventing air, water, and soil pollution.

eas. Within their respective application zone, many of these plans adopted differentiated measures for more versus less densely populated areas. Some of these plans introduced more stringent environmental standards for the manufacturing industries, especially in agglomerations above a certain threshold of inhabitants. For instance, the three consecutive *PPAs* for the Paris region mandated lower industrial NO_x and PM emission caps relative to the national caps. These plans also implement emergency responses when air concentration of certain pollutant exceed national and European standards.

This differentiation of environmental regulations across space also exists in other regions. In the US, for example, under the Clean Air Act industrial plants in counties where air pollution is too high have to implement specific actions. In China, as part of the five-year plan, the government sets a national industrial emissions reduction target that is broken down into sub-national targets that provinces and cities must meet. In particular, the ambition levels of these targets differ greatly across regions (Yamada, 2020). In this context, the second chapter of my thesis focuses on the consequences of such spatially differentiated regulations. Indeed, while stringent regulations may help improve air quality in some cities, they may also reduce the productivity of local industries and negatively affect workers' income. Therefore, in this chapter, I address the trade-off between the economy and the environment in a spatial context in order to identify the best way to allocate pollution emissions across different cities.

Finally, the traditional environmental economics Pigouvian approach, consists of estimating the cost to society of the environmental disruption in question. In the case of climate change or air pollution, this is the price that individuals would be willing to pay to avoid the emission of GHGs or harmful air pollutants. This price is generally the monetary equivalent of the utility reduction due to these emissions. In the case of climate change, models combining climate science and economic interactions between agents over several generations have been developed in order to calculate a "social cost of carbon" corresponding to the marginal damage caused by the emission of an additional ton of CO₂.³⁵ This approach should then make it possible to identify the appropriate level of a carbon tax or the price of emission permits on the markets. In the case of air pollution, the approach is significantly different. Damages are related to pollution concentrations, not to emission levels. Currently, few analyses link the location of emissions, air flow patterns and local air quality. It is therefore generally not possible to estimate a marginal cost of emitting air pollutants. The traditional approach is therefore to estimate the total cost of the observed concentrations, by calculating its monetary impact on health, economic activities and biodiversity.³⁶ This value can then be used to set a level of ambition for the policies

³⁵These models are called *Integrated Assessment Models* (IAM).

³⁶Generally, the monetary impact of health effects takes into account both the medical expenses involved

put in place. To this day, there are many analyses on the health effects of air pollution. These studies show that poor air quality implies an important cost in terms of medical expenses but also a degraded quality of life and a reduction in life expectancy (Waidyatillake et al., 2021). However, a recent literature has also focused on showing that air pollution has direct effects on economic activities. In particular, poor air quality reduces the productivity of industrial workers (Graff Zivin & Neidell, 2012; Chang et al., 2016; Adhvaryu et al., 2022; He et al., 2019). It is in this context that the third chapter of this thesis estimates the damage caused by particulate matter pollution (PM2.5) on French exports. Estimating the consequences of air pollution on the economic performances of firms could have implications in terms of public policy. Indeed, if regulations aimed at improving air quality are generally perceived as an additional cost for economic activities, the identification of direct economic benefits could lead to an increase in the ambition of these policies.

and a form of valuation of the losses in life expectancy and quality of life.

Chapter 1

Inter-Firm and Intra-Firm Spillover Effects of Industrial Regulation

1.1 Introduction

Governments often regulate specific subsets of firms differently from others in the same industry. For example, governments frequently impose environmental standards only for firms in the most polluted regions (Greenstone, 2002; Fowlie et al., 2016), offer employment incentives only to firms in distressed neighborhoods (Neumark & Simpson, 2015), and reserve business-promotion subsidies only for small or midsize firms (Martin et al., 2017; Rotemberg, 2019; Smagghue, 2019). The objective of these regulations is to incentivize particular outcomes at regulated firms; but if firms compete on the output market, then regulatory effects likely spillover to unregulated firms through competitive forces, potentially undercutting the objective of the regulation. The problem is compounded when firms own multiple plants. Profit maximizing firms may shift resources towards or away from regulated plants, in order to exploit or evade the regulation.

In this paper, we propose a method for estimating treatment effects of regulation on firm-level, plant-level, and aggregate outcomes in the presence of intra-firm and inter-firm spillover effects. The vast majority of empirical studies of regulation assumes no intra-firm or inter-firm spillover effects. Under this assumption, treatment effects can be estimated by standard difference-in-difference (DiD) techniques.¹ But if cost shocks spillover across or within firms – as they do in almost any model of imperfect competition – the stable unit treatment value assumption (SUTVA) fails, invalidating DiD estimators. A handful of recent papers explicitly address inter-firm spillover effects (Cai & Szeidl, 2022; Muehlegger & Sweeney, 2021; Rotemberg, 2019) and intra-firm spillover effects (Giroud & Mueller, 2019; Bartram et al., 2022; Gibson, 2019; Soliman, 2020) in isolation. This paper first fol-

¹The main empirical challenge confronted by DiD studies is the exogeneity of the regulation with respect to unobserved determinants of plant-level or firm-level outcomes. A recent group of papers also highlights the dependence of DiD estimation on the assumption of homogeneous treatment effects. If different plants become regulated in different years, and if the regulation induces heterogeneous treatment effects, then standard two-way fixed effect models can lead to biased estimates of average treatment effects (De Chaisemartin & d’Haultfoeuille, 2020).

lows their example in revealing reduced-form evidence of spillovers across firms within an industry and across plants within a regulated firm in the context of the European Union Emission Trading Scheme (EU ETS) regulating French manufacturers. To overcome the limitations associated with the fact that these exercises are performed separately, based on restrictive assumptions with respect to the evolution of market shares, productivity, and entry and exit, the paper then develops a structural model that encompasses both imperfect competition and intra-firm reallocation and proposes an estimator that separately identifies the effects of regulation on regulated units (plants or firms) and unregulated units.

Using confidential income statements for the universe of French manufacturing firms, as well as plant-level CO₂ emissions for a sample of establishments, we use our estimator to evaluate the effects of the EU ETS on revenues and emissions of French manufacturers. Launched in 2005, the EU ETS is the world’s largest cap-and-trade carbon emissions market and the EU’s flagship climate policy. An important feature of the EU ETS is that regulation applies at the sub-firm level - i.e. the plant. To minimize monitoring costs, the regulation targets plants with fossil fuel combustion capacity or output greater than industry-specific thresholds, thereby selecting larger plants owned by larger firms. We count only 465 firms and 699 plants ever regulated under the EU ETS in the French manufacturing sectors, among the 373,711 firms observed from 1994 to 2016. The regulation is thus incomplete on two dimensions: not all firms are regulated within each industry, and not all plants are regulated within a firm.

We specify a tractable model of demand and supply for differentiated products. We adopt standard assumptions for differentiated-product models: CES demand, Cobb-Douglas production, heterogeneous productivities, and monopolistic competition among single-plant firms.² We study the introduction of a regulation that targets the price of one of the inputs for some subset of firms in the economy. We show that in this framework, the evolution of firm-level revenues between a pre-regulation period and a post-regulation period depends on a firm’s own regulation status (through unit costs), as well as the regulation status of all firms in the sector (through the industry- and sector-wide price indices). We leverage the structure of the model to identify reduced-form coefficients using only information on the observed revenue changes and the observed vector of regulation. Armed with these reduced-form coefficients, we can compute average treatment effects on

²These assumptions are standard for models of the manufacturing sector as a whole. See for example Atkeson & Burstein (2008). In some cases, when researchers study a single industry, flexible patterns of across-firm price elasticities can be estimated structurally using detailed output price and product characteristic information (Berry et al., 1995). Data of this sort are only available for specific industries, hence this method is not well-suited to estimating aggregate effects of regulation (De Loecker & Syverson, 2021).

regulated firms and unregulated firms, as well as the aggregate effect on emissions.

We then extend our framework to multi-plant firms and consider a regulation that targets a subset of plants in the economy. To link plant-level regulation status with firm-level outcomes, we adopt a parsimonious model of across-plant aggregation. In the model, the concentration of resources across plants within the firm can either increase or decrease firm-level costs, depending on whether dispersion forces dominate economies of scope. We show how to identify the net effect of concentration from variation in labor concentration across plants within firms along with other reduced-form parameters to compute treatment effects.

For the empirical strategy, we first build a GMM procedure to estimate three reduced-form parameters that reflect (i) the direct effect of regulation on regulated firms' cost, scaled by demand-side parameters, (ii) the spillover effects of regulation on the industry-wide price index, which affects all firms within an industry, and (iii) the effect of the concentration of resources across plants on firms' cost. The direct effect of regulation on cost could also be obtained by regressing the log change in revenues on the change in firm-level regulation, absorbing all equilibrium effects with industry-by-year fixed effects. If costs do not influence the industry-wide price index, then this standard DiD estimation also identifies the average treatment effect on revenues for regulated firms (ATT). However, if costs pass through to the industry-wide price index, then the stable unit treatment value assumption fails, and this regressions does not identify the ATT. The model indicates that the industry-year fixed effects represent the log change in the (unobserved) industry-wide price indices. We thus exploit two moment conditions based on firm-level regulation status and concentration of resources across plants and two moment conditions based on industry-wide intensity of regulation and concentration, weighted by a function of pre-regulation and post-regulation firms' market shares.

Intuitively, if regulation increases costs, then industries with higher regulation density should see higher growth in their price indices. Previous work, which we designate as the local approximation (LA) method, uses base-year weighted regulation density (Cai & Szeidl, 2022; Muehlegger & Sweeney, 2021; Rotemberg, 2019). A key feature of our procedure is that we rely on price index theory to guide the construction of these weights. In particular, the CES assumption on demand indicates that firm-specific weights should reflect both pre-regulation market shares and post-regulation market shares (Sato, 1976; Vartia, 1976).³ Ignoring this pattern – using pre-regulation market shares only as weights

³This result is quite general. However, for the CES case, Sato (1976) and Vartia (1976) solved for the exact expression for the weights that makes the weighted average change in firm-level prices equal to the log change in the price index. For other functional form assumptions, we do not have analytical solutions for these weights.

– could induce an omitted variable bias.

Next, we estimate by OLS the input price effect of the regulation by regressing plant-level emissions on plant-specific treatment status and labor share, controlling for firm-year fixed effects. By exploiting within-firm variation, we capture the regulatory effect on regulated plants’ emissions, abstracting from changes in revenues at the firm level. Lastly, we compute treatment effects by comparing observed and counterfactual outcomes for each group of plants and firms. To do so, we solve for counterfactual revenues under the assumption that no firm was ever regulated.⁴ The counterfactual equilibrium is computed by solving a system of equations that relate counterfactual revenues to observed revenues, reduced-form parameters, observed regulation variables, and observed and counterfactual Sato-Vartia weights, which must update endogeneously. This system can be solved either with a nonlinear solver or a fixed point algorithm.

We run Monte Carlo simulations to evaluate the performance of our estimator in a finite-sample setting and compare it to the standard DiD and local approximation methods. Our estimator allows us to recover unbiased estimates of treatment effects in these experiments, while the DiD and LA methods can be severely biased and can even identify treatment effects of the wrong sign under some parameterizations.

We find that the EU ETS increased sales for regulated firms (firms with at least 1 regulated plant) and reduced emissions for regulated plants, relative to a completely unregulated counterfactual. The fact that regulation *increased* sales is perhaps surprising, as the regulation effectively increased the cost of energy for regulated firms. However, as Porter hypothesized, environmental regulation can trigger cost-saving investments that ultimately increase sales (Porter & van der Linde, 1995; Ambec et al., 2013).⁵ We also find that the policy reduced aggregate CO₂ emissions from French manufacturers generated in the production of goods for the French market by roughly 5-12% annually relative to an unregulated counterfactual in the first two phases of the policy (2005 - 2012), depending on the year. The DiD estimator overstates the effect on sales and reverses the sign of the effect on emissions.

⁴When output is denominated in value, it is not possible to separately identify the underlying structural parameters of the model; but we show that to compute counterfactual equilibria, it is sufficient to estimate just a few reduced-form parameters.

⁵There is in fact quite a lot of support for the Porter hypothesis, especially in the context of the EU ETS. See for example Löschel et al. (2018) Dechezleprêtre et al. (2018) and Calligaris et al. (2018). Evidence from Calem & Dechezlepretre (2016) and Calem (2020) indicate that the EU ETS increases R&D and green innovation, which could explain these positive effects on sales. All of these results are based on DiD estimators, so are potentially subject to the very critique we advance in this paper. Nevertheless, we document below that the sign of the effect on costs is the same as the sign of the DD estimate on revenues. So while the magnitude of the effects in these papers are likely biased, the sign is correct, according to our model.

This paper contributes to several strands of literature. First, we contribute to a burgeoning literature dedicated to estimating the effects of firm-level shocks on own-firm, competitor-firm, and aggregate outcomes. Recent examples include Cai & Szeidl (2022), Muehlegger & Sweeney (2021), and Rotemberg (2019) who study shocks to access to credit, shocks to input costs, and subsidies, respectively. All three papers derive linear estimation equations based on local approximations to the log change in industry-wide price indices.⁶ By contrast, we derive the exact change in revenues between a pre-regulation period and a post-regulation period using index price theory, and thus need not resort to approximations. We also allow for nonrandom assignment of regulation and non-random-walk productivity growth, which none of these papers do.⁷ In Monte Carlo experiments, we show that in some cases, the linear approximation method performs quite well, while in other cases, estimated effects can be biased.

The paper also contributes to a literature that structurally estimates the effects of environmental policies on emissions and competitiveness. Most of this literature studies homogeneous product markets (Fowlie et al., 2016; Hintermann, 2017; Ganapati et al., 2020; Fabra & Reguant, 2014), and often focuses on pass through from input prices to output prices (Muehlegger & Sweeney, 2021; Hintermann, 2017; Ganapati et al., 2020; Fabra & Reguant, 2014). A notable exception is Shapiro & Walker (2018), who study the effects of environmental regulations on emissions from US manufacturing. Our model shares many features with that of Shapiro & Walker (2018), but the context and empirical strategy is quite different. Shapiro & Walker (2018) back out the implied homogeneous environmental tax from aggregate data on sector-level emissions and revenues and study the effects of these implied taxes on total emissions. By contrast, we study the effect of a particular regulation on a discrete set of firms.

Third, our paper contributes to a handful of papers that studies how shocks propagate through intra-firm networks. A number of studies use DD estimators to estimate the effects of new regulations or other shocks that apply at the plant-level (Giroud & Mueller,

⁶Cai & Szeidl (2022) and Rotemberg (2019) consider differentiated product markets, while Muehlegger & Sweeney (2021) study a homogeneous product market (gasoline). Cai & Szeidl (2022) explicitly derives the estimation equation as a first order approximation (proposition 1). Muehlegger & Sweeney (2021) specify an estimation equation that is linear in own-firm cost and weighted average of competitor costs. The authors refer to their approach as a “reduced form...approximation” (footnote 14).

⁷Another group of papers studies equilibrium effects of demand and supply shocks in an economic geography framework (Adao et al., 2019; Franklin et al., 2021; Bergquist et al., 2019). These papers are similar in spirit to our work, though the emphasis in these papers is on labor supply as opposed to imperfectly competitive output markets. Also, in theoretical work, Sraer & Thesmar (2020) combine DiD estimation with a structural model to compute aggregate effects of firm-specific shocks, similarly to us. A key difference between Sraer & Thesmar (2020) and our work is that Sraer & Thesmar (2020) rely on specific assumptions on the distribution of productivities in the economy in order to aggregate over firms and build counterfactuals. By contrast, we solve a set of simultaneous equations for a discrete set of firms to retrieve incomes and market shares under different general equilibrium conditions.

2019; Bartram et al., 2022; Gibson, 2019; Soliman, 2020). We show how our method can incorporate a parsimonious model of multiplant firms, and we derive an estimation strategy that encompasses both inter-firm and intra-firm spillovers. As a result, our method can also be applied to study a broader set of regulations that target specific plants, by virtue of a plant’s size or location, for instance.

Finally, our paper contributes to the literature evaluating the effect of the EU ETS on firm-level emissions and competitiveness. Most of the studies so far have used DiD estimators at the firm-level, comparing firms with or without a regulated plant (Petrick & Wagner, 2014; Löschel et al., 2018; Colmer et al., 2020b; Dechezleprêtre et al., 2018; Jaraite & Di Maria, 2018; Calligaris et al., 2018)⁸. Identification in these papers is usually tied to variation in across-plant concentration of economic activity. Our model of multiplant firms provides rigorous micro-foundations for this identification strategy. Nevertheless, our Monte Carlo experiments and empirical application indicate that the DiD approach can lead to biased estimates as a result of failing to properly account for inter-firm and intra-firm spillovers.

1.2 Data and Reduced-Form Evidence of Spillover Effects

In this section, we discuss the empirical context related to the European Emissions Trading Scheme and the data for French manufacturing firms and plants. We provide reduced-form evidence of both inter-firm and intra-firm spillovers in our empirical context.

1.2.1 The European Emissions Trading Scheme

The EU ETS is a cap-and-trade carbon policy implemented by the European Commission, regulating approximately 11,000 power and manufacturing plants across 31 participating countries, including France. The EU ETS was first announced in 2000 with a working paper from the European Commission, formalized with the Emissions Trading Directive in 2003, and initiated on January 1st, 2005. The EU ETS regulates all combustion installations with a rated thermal input greater than 20MW, and other productive processes with capacity or output greater than predetermined industry-specific thresholds. Regulated firms must surrender permits – called EU allowance units – for each ton of CO₂ emitted during the year. Non-compliance triggers a fine of 100 euros per ton of CO₂ exceeding the number of permits surrendered.

⁸See Martin et al. (2015) and Joltreau & Sommerfeld (2018) for extensive reviews.

The European Commission initially designed three phases of the EU ETS. The first phase was designed as a trial period, running from 2005 to 2007; the second phase ran from 2008 to 2012, corresponding to the first commitment period under the Kyoto Protocol; and the third phase ran from 2012 until 2020.⁹ During Phases I and II, the implementation of the policy was relatively decentralized, with participating countries submitting National Allocation Plans (NAPs) to the European Commission in which they proposed lists of installations to regulate, potential emission caps and allowance distribution rules (Ellerman et al., 2015). During these two phases, free permits were allocated to regulated installations on the basis of their historical emissions (according to the “grandfathering” method).¹⁰ In Phase III, allocation began a transition towards an open auction system, first for power installations (except for a few plants in Eastern Europe), and progressively for other regulated industries. Furthermore, a single emission cap (declining at a 1.74% rate per year) was adopted for the entire EU market and EU-wide sector-specific “benchmarks” were adopted for determining allowance distribution rules.

The cap-and-trade mechanism effectively establishes a shadow price for carbon for regulated installations. Absent free permit allocations, firms must purchase permits on the open market for every ton of CO₂ produced at each regulated installation, or pay the fine. Even with free permit allocations, the opportunity cost of CO₂ produced at regulated installations increases. If the value of the marginal unit of CO₂ is less than the market price for allowances, firms have an incentive to abate the marginal unit and sell their (potentially freely acquired) permits on the open market. Hence, no matter the allocation of allowances, firms equate marginal benefit of CO₂ at regulated installations with the unique shadow price in equilibrium (Coase, 1969).

The European Commission publishes information for each regulated installation on the official ETS regulation platform – the European Union Transaction Log. From this website, we obtain the address of each regulated installation in France, the initial date of regulation, and the unique French tax identifier (SIREN) of the firm that owns the installation. We then match installations to plants by SIREN and street address (see Appendix A.2.3 for details). In total, we count 1,415 installations in France ever regulated under the EU ETS across 1,264 plants and 846 firms. We consider a plant regulated in a given year if it has at least 1 regulated installation in the year. Broken down by phase, we count 984 plants regulated under phase I, 953 plants regulated under phase II, and 982 plants regulated under phase III. The number of regulated plants and firms by sector are reported in Table 1.2.1.¹¹ We count 465 firms and 699 plants ever regulated in the manufacturing

⁹Revisions of the regulation in the expectation of Phase IV (2021-2030) are already well under way.

¹⁰In France, the Caisse des Dépôts was in charge of the implementation of the regulation for the country.

¹¹We aggregate 2-digit codes to standard sectoral classifications, as in Harrigan et al. (2018). For man-

	# Firms	# Plants
<i>Manufacturing</i>		
Motor vehicles & other transport equipment	11	13
Chemicals & Pharmaceuticals	85	126
Computer, Electronic & Optical products	4	4
Electrical equipment	2	2
Food, Beverages & Tobacco	104	148
Machinery & Equipment	7	9
Basic & Fabricated Metal products	40	57
Rubber, Plastic & Non-metallic mineral products	88	198
Textiles & Apparel	17	17
Wood & Paper products	99	116
Other manufacturing	8	9
Subtotal	465	699
<i>Non-Manufacturing/Unknown</i>	381	565
Total	846	1264

Notes:

Table 1.2.1: Treatment by Sector

sectors.

1.2.2 Firm-Level Data

We use confidential income statements and export sales for the universe of French manufacturing firms from the statistical agency INSEE and French Customs Bureau. All French firms are required to report total sales each year to the French tax authority, which INSEE then records in the FICUS database (for years 1994 - 2007) and FARE database (2008 - 2016).¹² Across all data sets, firms are identified with a unique administrative code called SIREN. In the surveys, firms declare a 4-digit activity code – the APE – from the national activity nomenclature (NAF).¹³ We take the 4-digit APE code as the “industry” of the firm since it identifies the main activity of the firm, and group aggregates of 2-digit APE code into “sectors”, as in Harrigan et al. (2018).

At the firm level, we observe that (i) the policy is incomplete: among the 373,711 French manufacturing firms, only 465 of them are regulated (i.e., firms with at least 1 regulated

ufacturing, we include all firms in all industries with 2-digit codes between 15 - 37 in the national activity nomenclature NAF Rev.1. We exclude extractive industries (codes 10 to 14) and energy production and distribution (40 and 41).

¹²This data is confidential and its access is subject to authorization by a Statistical Secret Committee.

¹³The APE (*Activité Principale de l'Etablissement*) corresponds to the main activity of the firm according to the national nomenclature denominated *Nomenclature des Activités Françaises* (NAF).

plant); (ii) the policy targeted large firms, as indicated by their higher sales, their average domestic market share of 9% (in contrast to 0.2% for the non-regulated firms), and their larger number of employees and plants. Descriptive statistics are reported in panel A of Table 1.2.2. Due to the size difference, it is likely that outcomes for regulated and unregulated firms would have followed different trends even without the policy.¹⁴ We thus match regulated to unregulated firms from the same industry on both levels and growth rates of sales in 2004, the year prior to the start of the EU ETS.¹⁵ Columns 3 and 4 show that regulated and unregulated firms in the matched sample, are closer in size, though regulated firms are still larger on average.

	Full Sample				Matched Sample			
	regulated		unregulated		regulated		unregulated	
	(1)		(2)		(3)		(4)	
<i>Panel A: Firm Level Dataset</i>								
Sales (Millions euros)								
Domestic	239.0	(1345)	3.585	(150.5)	198.8	(1343)	70.29	(386.4)
Export	110.1	(327.5)	0.910	(20.39)	92.62	(275.2)	21.21	(95.15)
Total	349.1	(1504)	4.495	(154.0)	291.4	(1489)	91.50	(438.2)
Market Share	0.093	(0.160)	0.002	(0.015)	0.061	(0.096)	0.016	(0.034)
# Workers	816.3	(1861)	22.23	(270.7)	613.2	(1060)	233.4	(517.7)
# Plants	3.719	(6.205)	1.167	(1.661)	2.905	(4.696)	2.191	(3.602)
# Regulated Plants	0.695	(1.310)			0.570	(0.997)		
# Firm-year Obs	7,652		3,126,194		6,128		23,252	
# Firms	465		373,711		314		1,314	
<i>Panel B: Plant Level Dataset</i>								
CO ₂ Emissions ('000 Kg)	91.92	(524.0)	3.047	(66.27)	75.35	(360.9)	11.84	(129.3)
# Workers	383.0	(504.6)	145.6	(282.9)	380.8	(446.7)	203.4	(376.5)
# Plant-year Obs	7,924		204,907		5,413		20,083	
# Firms	405		32,630		273		1,223	
# Plants	600		42,491		362		2,455	

Notes: Values indicate mean of annual observations, with standard deviations in parentheses. Data span 1994 - 2016 and includes all firms in all industries with 2-digit codes between 15 - 37 in the NAFRev.1.

Table 1.2.2: Descriptive Statistics

1.2.3 Plant-Level Data

We take plant-level information from two sources. First, we measure energy consumption using the EACEI (*Enquête Annuelle sur la Consommation d'Énergie dans l'Industrie*)

¹⁴Under a random walk assumption on productivity growth, then changes in revenues, for example, would be orthogonal to levels in 2004. However, random walk is a strong assumption that is unlikely to hold.

¹⁵We provide details on the exact matching algorithm in section 1.5.4.

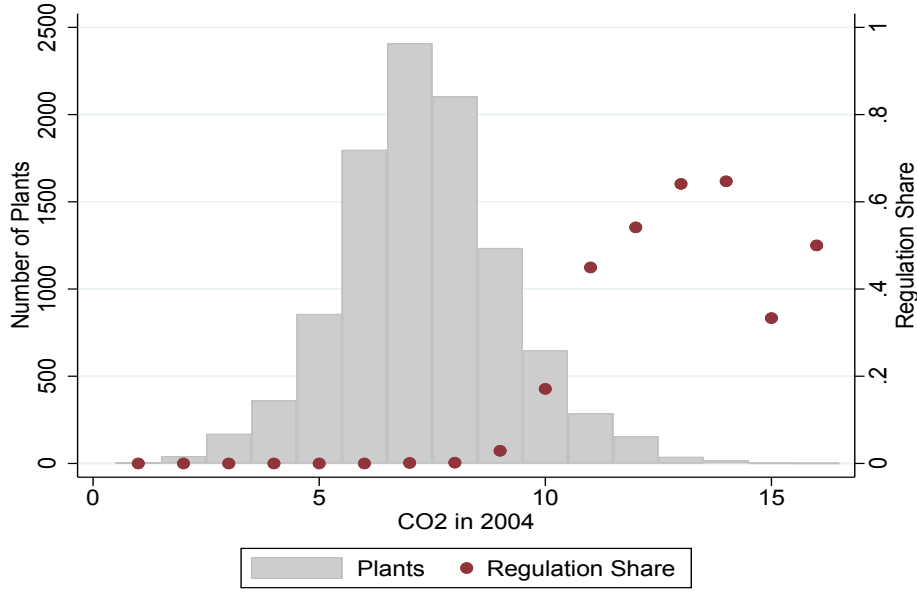


Figure 1.2.1: Regulation status by plant-level CO₂ emissions in 2004

surveys. These surveys include all energy-related expenditures, with details on several energy types and fuels (quantity consumed and expenditures), at the plant level. A plant is identified by a stable administrative code called SIRET in which the first 9 digits correspond to the SIREN of the firm the plant belongs to.¹⁶ All plants with more than 250 employees receive the EACEI survey each year from 1994 to 2016, with smaller plants sampled randomly. Firms report purchases of electricity (bought from the grid and self-generated), steam, natural gas and other types of gas, coal, lignite, coke, propane and butane, domestic and heavy fuels, oil and other types of petroleum products for surveyed plants. To convert energy consumption into CO₂ emissions, we use emission factors by energy type from the French Environment and Energy Agency (ADEME). Second, we take the number of employees for each plant in each year from the *Stock of Establishments*. Details about the data and the conversion factors can be found in the Data Appendix.

Descriptive statistics for both CO₂ and employment are reported in panel B of Table 1.2.2. We see that the EU ETS targeted large plants on these two dimensions: the number of workers and CO₂ emissions in an average regulated plant are higher by a factor 2.6 and a factor 30, respectively, than in an average non-regulated plant. Figure 1.2.1 shows that even though the regulation targeted the most polluting plants and did not regulate plants emitting less than the median CO₂ levels in 2004, the share of regulated plants within a bin of CO₂ emissions in 2004 never reaches 100%, even for the most polluting plants.

¹⁶The SIRET is a 14-digit code identifying a legal unit's plant. It is composed of the SIREN (9-digit identifier) of the corresponding legal unit and of 5 additional digits uniquely identifying plants within a legal unit.

1.2.4 Reduced-Form Evidence of Inter-Firm Spillovers

A standard difference-in-differences (DiD) approach to evaluating treatment effects of an incomplete regulation such as the EU ETS abstracts from general equilibrium effects. For an outcome variable v , we estimate

$$\Delta v_{fist} = \beta_{v0}^{DD} R_{fist} + \delta_{ist} + \epsilon_{fist}, \quad (1.1)$$

where f is a firm that belongs to an industry i within a sector s , Δv_{fist} corresponds to the log difference between the outcome at time t relative to the pre-period $t_0 = 2004$, and variable v is defined as firm-level revenues. For years post-2004, we take the log change in total sales between a given post-period year and the year 2004 as the dependent variable. For years prior to 2004, we invert the ratio, so that the dependent variable represents the growth rate between a pre-period year and 2004. Treatment is identified through the indicator variable R_{fist} taking the value 1 if the firm owns any regulated plant at time t , and zero otherwise. The industry-year fixed effects δ_{ist} absorb all general equilibrium effects that are common within an industry. Thus, as long as we have $E[\epsilon_{fist}|R_{fist}, \delta_{ist}] = 0$, we should get an unbiased estimate of the direct treatment effect on firms' outcome v . To achieve this, we use a matched sample of regulated and non-regulated firms for which treatment assignment can be considered as good as random.

We would like to point out that interpreting β_{v0}^{DD} alone as an estimate of all treatment effects is misleading in this context. Indeed, some of the general equilibrium effects could be induced by the regulation. For instance, regulated firms may adjust their prices post regulation in such a way that price indices would be different in a counterfactual equilibrium without the regulation. In such context, the DiD framework correctly recovers the direct effect of regulation, β_{v0}^{DD} , but is insufficient to compute counterfactual changes in revenues since it misses the indirect effects.¹⁷

To investigate the indirect treatment effects that are channeled through competition forces and therefore result in changes in industry-wide price indices, we can follow the specifications from Cai & Szeidl (2022); Muehlegger & Sweeney (2021); Rotemberg (2019) by estimating

$$\Delta v_{fist} = \beta_{v0}^{LA} R_{fist} + \beta_{v1}^{LA} \sum_{k \in \Omega_{ist_0}} \theta_{kist_0} R_{kist} + \delta_{st} + \epsilon_{fist}, \quad (1.2)$$

¹⁷The same point is made by Baier & Bergstrand (2009) when observing that consistent estimation of trade elasticities requires controlling for multilateral resistance terms, similarly to how we include industry-by-year fixed effects to control for industry-wide price indices. They also argue that, in a counterfactual scenario in which trade costs change, the multilateral resistance terms would update as well, and hence need to be recomputed for the counterfactual scenario.

where the first term identifies the direct effect of being regulated and the second term the indirect effect induced by the intensity of regulation within an industry. This intensity of regulation is computed as the pre-period market share of all will-be-regulated firms within an industry, with θ_{kist_0} denoting the market share of firm k at period $t_0 = 2004$ and Ω_{ist_0} the set of firms observed at period t_0 . We call this equation the “local approximation” (LA) model and we show below that it follows from taking a local approximation to the price index changes around the pre-regulation equilibrium (see appendix A.3.2 for details). To estimate indirect treatment effects, equation (1.2) backs out from industry-time fixed effects and includes sector-time fixed effects, δ_{st} , instead. Thus, the term $\sum_{k \in \Omega_{ist_0}} \theta_{kist_0} R_{kist}$ somehow reflects how price indices vary within an industry due to the regulation and how their competitors’ regulatory statuses affect each firm.¹⁸

Figure 1.2.2 panel A shows that, for the DiD method to be valid, the population of firms must be restricted to the matched sample, otherwise the parallel trend assumption does not hold.¹⁹ It also reveals a positive direct effect on revenues of regulated firms. Yet, this estimate could suffer from a bias due to the potential violation of SUTVA if the outcomes of non-regulated firms are also influenced by the regulation. Panel B reveals that it is the case since the direct and indirect effects go in opposite directions and the indirect effect is statistically significant and of larger magnitude than the direct effect, thereby indicating that inter-firm spillovers exist. Yet, this approach abstracts from the multi-plant dimension of firms, which may induce biases in the estimated treatment effects on revenues and leave open the question of how to evaluate treatment effects on emissions at the plant level.

1.2.5 Reduced-Form Evidence of Intra-Firm Spillovers

By targeting large firms, which often own multiple plants, the policy may not only generate inter-firm spillovers, but also intra-firm spillovers. The EU ETS treatment is assigned at the plant level, and not at the firm level. As a result, firms may have a few plants regulated and other plants not regulated. To reveal the presence of across-plant reallocation forces within the firm, we estimate

$$\Delta Z_{jfist} = \beta_{Z0}^{DD} R_{jfist} + \beta_{Z1}^{DD} \sum_{\ell \in \Omega_{fist_0}} \alpha_{\ell fist_0} R_{\ell fist} + \delta_{ist} + \epsilon_{fist}, \quad (1.3)$$

¹⁸By assumption, competitors can only be French in this LA model since firm f ’s market share is defined as the ratio of firm f ’s total sales over the aggregate of all French firms’ revenues.

¹⁹We note that, even with the matched sample, we barely reject the existence of pre-trends in the post-announcement period. This is due to the fact that we use the same matched sample for all our empirical results, following the matching algorithm described in section 1.5.4. As our model will show, equation (1.1) is thus misspecified as it should define firm-level treatment based on the labor share of treated plants and include a control for the concentration of activities across plants.




Figure 1.2.2: The Regulation Effects on Firm-Level Revenues, with and without Inter-Firm Spillovers

Notes: Panel A presents estimates of β_{y0}^{DD} derived from OLS estimation of (1.1), with the full (matched) sample indicated with white (black) diamonds. Panel B presents estimates of β_{y0}^{LA} (white circles) and β_{y1}^{LA} (black circles) derived from OLS estimation of (1.2). Bars indicate 95% confidence intervals. For each period, we estimate a pooled regression, combining all years included in the period. The vertical line indicates the beginning of the regulation.

where the outcome variable Z_{jfst} corresponds to plant j 's CO₂ emissions, where plant j belongs to the set of plants Ω_{fst} owned by firm f at time t . Treatment is identified by both the plant-level regulation status R_{jfst} , which is a dummy variable that equals one if plant j is regulated and zero otherwise, and by the firm-level pre-period-weighted regulation intensity, where the weights are the plant-level labor shares.²⁰ This specification absorbs all general equilibrium effects that are common across firms within an industry in the industry-time fixed effects δ_{ist} and includes all plants. For comparison, we also estimate the standard DiD model abstracting from firm-level regulation intensity $\sum_{\ell \in \Omega_{fst0}} \alpha_{\ell fst0} R_{jfst}$ on a sample of plants that are either directly regulated or not regulated within a firm that owns no other regulated plant.

Figure 1.2.3 panel A shows the results from the standard DiD model using the full (matched) sample, which are indicated with black (white) diamonds. Here, we match plants based on their levels and growth rates of CO₂ emissions in 2004 (using propensity

²⁰In the French manufacturing data sets, we do not observe output at the plant level. Hence, we use labor shares instead of output shares to weight the regulation status of each plant.

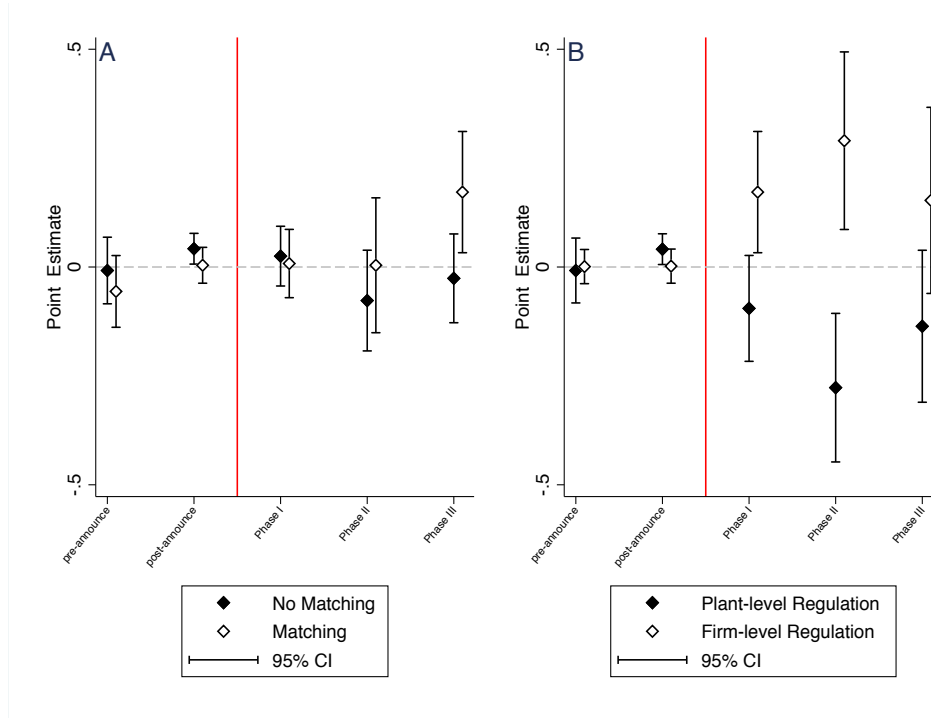


Figure 1.2.3: The Regulation Effects on Plant-Level Emissions, with and without Intra-Firm Spillovers

Notes: Panel A presents estimates of β_{Z0}^{DD} derived from the standard DiD equation (1.3) without firm-level regulation, using the full (matched) sample as indicated by black (white) diamonds. Panel B presents estimates of β_{Z0}^{DD} (black) and β_{Z1}^{DD} (white) from the OLS estimation of (1.3). Bars indicate 95% confidence intervals. For each period, we estimate a pooled regression, combining all years included in the period. The vertical line indicates the beginning of the regulation. Both panels include industry-by-year fixed effects.

score matching). Even though regulated and unregulated plants were trending in a parallel fashion prior to regulation, emissions from regulated plants grew slightly faster than unregulated plants after 2004. This effect is statistically insignificant for phases I and II, but in phase III, we can reject 0 at the 5% level in the matched sample. This model relies on across-firm variation rather than within-firm variation. Thus, these point estimates combine an increase in scale (that is, in revenues), which pushes emissions upward, and a potential change in emission intensity. Since we found positive effects of the regulation on firm-level sales in Figure 1.2.2, it is not surprising that regulated plants increased emissions faster than unregulated plants.

In panel B of Figure 1.2.3, we separately estimate the effect of firm-level regulation on plant-level emissions from the effect of plant-level regulation. Here, using the matched sample, we find that emissions from regulated plants grew slower post-regulation than unregulated plants, conditional on firm-level regulation. Moreover, we see that emissions from plants owned by firms with higher firm-level regulation values grew faster, post-

regulation. This provides evidence that firms reallocated their inputs, hence emissions, away from regulated plants and toward non-regulated plants, potentially undermining the overall effects of the policy.

Computing treatment effects on revenues and emissions that combine the inter-firm and intra-firm spillovers is particularly challenging absent a conceptual framework. This is why we now build a model that renders the mechanisms behind these spillovers explicit and then develop an empirical strategy that allows us to go beyond the limits of the reduced-form estimators.

1.3 Model

In the section, we present a tractable model of demand and supply for differentiated products. We start with a version in which all firms own and operate a single plant. We introduce a regulation that increases the price of one of the inputs of some subset of firms in the economy. We then extend the model to include multi-plant firms and a regulation that targets a subset of plants in the economy.

1.3.1 Demand

We consider an economy where consumers have access to differentiated varieties produced in several manufacturing sectors and an outside good (which corresponds to services). The representative consumer's preferences are described by the following three-tiered utility function:

$$U_t = (Q_{0t})^{a_0} \prod_s \left[\left(\sum_{i \in \Upsilon_s} (Q_{ist})^\nu \right)^{1/\nu} \right]^{a_s} \text{ where } Q_{ist} = \left(\sum_{f \in \Omega_{ist}} (Q_{fist})^\rho \right)^{1/\rho}. \quad (1.4)$$

The first tier is Cobb-Douglas across sectors, with $a_0 + \sum_s a_s = 1$, and the second and third tiers are Constant Elasticity of Substitution (CES) functions across the set of industries within a sector s , Υ_s , and across the set of varieties available in each industry i at time t , Ω_{ist} . The outside good consumption at time t is denoted by Q_{0t} while Q_{ist} represents the aggregate consumption in industry i and Q_{fist} the consumption level in variety f . Varieties f are imperfect substitutes within an industry, implying $0 < \rho < 1$. Similarly, we assume $0 < \nu < 1$. We also assume that two varieties within an industry are closer substitutes than goods from two different industries, $\rho > \nu$. Let I_t be the national income, which equals the expenditure level. Sector-specific expenditures Y_{st} (also for the outside good) are determined as follows: $Y_{st} = a_s I_t$.

Industry and sector price indices are defined in the usual way:

$$\Psi_{st} = \left(\sum_{i \in \Upsilon_s} (P_{ist})^{\frac{\nu}{\nu-1}} \right)^{\frac{\nu-1}{\nu}}, \quad P_{ist} = \left(\sum_{f \in \Omega_{ist}} (p_{fist})^{\frac{\rho}{\rho-1}} \right)^{\frac{\rho-1}{\rho}}. \quad (1.5)$$

Utility maximization then implies the following expression for expenses in variety f at time t :

$$y_{fist} = (p_{fist})^{\frac{\rho}{\rho-1}} (P_{ist})^{\frac{\nu-\rho}{(\nu-1)(1-\rho)}} (\Psi_{st})^{\frac{\nu}{1-\nu}} Y_{st}. \quad (1.6)$$

We can see that expenditures are decreasing in variety f 's price and in sector s 's price index, and increasing in industry i 's price index and in sector-wide expenditures.

1.3.2 Production, Emissions and Regulation

Each manufacturing firm produces a single differentiated variety, so we use f interchangeably to index varieties or firms. Production in manufacturing requires two factors of production: labor and energy, the latter being a source of CO₂ emissions when consumed (as it is the case for fossil fuels). The government sets an environmental regulation that has two implications: first, it raises the cost of energy for regulated firms to limit their CO₂ emissions; second, it may induce regulated firms to invest in energy-efficiency improvements, which would decrease the amount of CO₂ emissions per output. To save on implementation costs, the government only targets a few firms within each industry.²¹

We model production as a Cobb-Douglas combination of labor (L_{fit}) and energy (E_{fit}) with constant returns to scale:

$$Q_{fist} = (L_{fist})^{1-\gamma} (E_{fist})^{\gamma} \exp(\omega_{fist} + \gamma \mu_t^e R_{fist}), \quad (1.7)$$

where ω_{fist} denotes firm f 's Hicks-neutral productivity level at time t , $R_{fist} \in \{0, 1\}$ denotes the firm's regulatory status at time t , and μ_t^e captures the regulation's effect on energy efficiency at time t .²² To simplify notations, we are now suppressing the s subscript wherever possible. Productivity evolves according to a flexible Markov process:

$$\omega_{fit} = g(\omega_{fit-1}) + u_{fit}, \quad (1.8)$$

where $g(\cdot)$ is an arbitrary function of past firm-level productivity and u_{fit} is an i.i.d.

²¹In the theoretical model, we keep the assignment rule general. For instance, the government could assign treatment randomly or it could assign treatment with a probability increasing in initial size or emission levels. We will discuss specific assignment rules in the application section.

²²The energy efficiency effect of the policy varies with the stringency of the policy. Indeed, a more stringent policy will foster more technological investments and R&D, thereby lowering the emission intensity of production.

shock with mean zero. If $g(\omega_{fit-1}) = \omega_{fit-1}$, then productivity follows a random walk. Otherwise, productivity is path dependent. If $g(\cdot)$ is decreasing in ω_{fit-1} , more productive firms grow slower, whereas more productive firms grow faster if $g(\cdot)$ is increasing in ω_{fit-1} .

Production generates CO₂ emissions in proportion to the use of energy inputs. Emissions (Z_{fit}) emitted at time t by firm f are given by

$$Z_{fit} = \kappa_t E_{fit}, \quad (1.9)$$

where κ_t is the quantity of emissions per unit of energy. This conversion factor may vary over time due to changes in the national fuel mix and technological progress.

The outside good is produced under constant returns to scale using only labor. The term ω_{0t} denotes labor productivity in the outside sector at time t . The outside good is used as the numeraire, and it is produced under perfect competition. As a result, the wage rate w_{Lt} is pinned down by labor productivity in the outside good sector: $w_{Lt} = \exp(\omega_{0t})$. The manufacturing firms are therefore hiring labor at an exogenous wage rate w_{Lt} , which does not respond to the environmental regulation.

The manufacturing firms also buy energy at an exogenous price w_{Et} . Energy supply is assumed to be perfectly elastic at price w_{Et} , which is consistent with a situation in which a state-owned company supplies energy at a controlled price. If a firm is subject to the environmental regulation ($R_{fit} = 1$), then it also faces a price on its emissions, w_{Zt} , that is imposed by the government. For convenience, we express the price on emissions as a proportion of the exogenous energy price by imposing $w_{Zt} \equiv w_{Et}(e^{\mu_t^z R_{fit}} - 1)/\kappa_t$, so that the regulation equivalently raises the cost of energy for regulated firms or makes them pay for their emissions. The effect of the regulation on the energy price is summarized by μ_t^z , which can adjust over time to exogenous energy price changes and changes in κ_t .

Cost minimization yields the following unit cost function:

$$c_{fit} = (1 - \gamma)^{-(1-\gamma)} \gamma^{-\gamma} (w_{Lt})^{1-\gamma} (w_{Et})^\gamma \exp[-\omega_{fit} + \gamma(\mu_t^z - \mu_t^e) R_{fit}]. \quad (1.10)$$

The effect of the regulation on the unit cost of a regulated firm depends on the net regulation cost, $\tau_t \equiv \gamma(\mu_t^z - \mu_t^e)$. This tells us that the regulation will increase a firm's unit cost if the energy price effect dominates, $\mu_t^z > \mu_t^e$, but it will decrease its unit cost if the energy efficiency effect dominates, $\mu_t^z < \mu_t^e$.

1.3.3 Profit Maximization and Equilibrium

In each period t , the set of firms operating in a given industry (Ω_{it}) is exogenous and common knowledge. Thus, each industry is composed of an exogenous number of firms, N_{it} .²³ The motivation for this assumption is that, in our empirical application, we aim at evaluating the effect of a cost shock vector on active individual firms. We thus abstract from the entry/exit response to these cost shocks. However, exogenous industrial dynamics can lead some incumbent firms to exit and some new entrants to appear. As a result, the set of firms in each differentiated industry Ω_{it} varies exogenously over time.²⁴

Firms take their production decisions based on their regulatory status R_{fit} , their production function (1.7), and their productivity level (1.8). Given these characteristics, firms choose the price that maximizes their profits:

$$\max_{p_{fit}} \Pi_{fit} = p_{fit} Q_{fit} - c_{fit} Q_{fit} \quad (1.11)$$

where c_{fit} is described in (1.10). As is standard in models built for differentiated varieties, we assume that firms are monopolistic competitors, and therefore solve this maximization program without regard for how their decision affects the price index.²⁵ Due to constant returns to scale in production, pricing depends only on firm-level production costs and the standard Dixit-Stiglitz markup:

$$p_{fit} = \frac{c_{fit}}{\rho}. \quad (1.12)$$

In a competitive equilibrium, all consumers maximize their utility, all firms maximize their profits, and all markets clear. Hence, the representative consumer maximizes his utility by spending y_{fit} on each variety f as described by (1.6). In the outside sector, profit

²³Real-world data sets indeed contain finite numbers of firms. This fact has been explored in the model of Eaton et al. (2013), for instance. In contrast to this model, we keep the assumption that firms are small enough to not incorporate the effect of their decisions on price indices.

²⁴Abstracting from the entry/exit response corresponds to a short-run equilibrium as defined in Melitz & Ottaviano (2008). Yet, we allow for exogenous variation in the set of active firms over time. This assumption entails that incumbent firms may exit with a constant exogenous probability ϱ from t to $t + 1$, and that, at the beginning of period t , an exogenous number M_{ist} of new entrants in industry i appear. M_{ist} can be interpreted as the number of entrepreneurs who can create a company in industry i at time t . Upon entry, each firm learns its productivity draw at no cost, as well as its regulation status. Given the draws and the absence of fixed cost to serve the economy, new entrants face the same constraints as the incumbent firms and can compute their own unit cost.

²⁵In Appendix A.5, we consider the alternative assumption of oligopolistic firms and derive prices and estimation equations in this setup. We also use the oligopoly equilibrium to perform an adversarial test of our estimation strategy—we generate simulated data as if firms are oligopolists and then apply our estimation strategy derived on the assumption of monopolistic competition. Even in that setting, our estimator reliably outperforms the alternatives.

maximization requires that $w_{Lt} = \exp(\omega_{0t})$, which pins down the wage. In the manufacturing sector, profit maximization requires that firms set their price under monopolistic competition as in (1.12). Finally, sector-level expenditures equates to the sum of sales of all firms producing in this sector at time t . For simplicity, we assume that national income I_t is exogenous, as in Helpman et al. (2008). In other words, we are considering a regulation which is not so large as to markedly affect national income. For each sector, we thus set

$$Y_{st} = \sum_{i \in \Upsilon_s} \sum_{f \in \Omega_{ist}} y_{fist}. \quad (1.13)$$

1.3.4 The Effect of Regulation on Revenues

If we observed technologies in such details that we could compute the cost function, it would be possible to identify the net regulatory cost, τ_t , directly from equation 1.10. Such details are unobserved in most applications, however, so it is useful to derive a method that allows us to estimate the net regulatory cost and treatment effects based on observable quantities.

From equation (1.6), we infer that firm f 's revenues depend on three components: (i) its own price, which is in turn determined by the unit cost function (1.10) given the pricing rule (1.12), (ii) its industry and sector price indices (1.5), and (iii) its sector-wide expenditures. The log difference in firm f 's revenues between time t and a base year t_0 prior to the regulation, denoted by Δy_{fit} , will thus depend on the change of these three components. We have $\Delta p_{fit} = \Delta c_{fit} = \tau_t R_{fit} + \Delta A_{fit} + (1 - \gamma)\Delta w_{Lt} + \gamma\Delta w_{Et}$, where $\Delta A_{fit} \equiv \omega_{fit_0} - \omega_{fit}$ represents the exogenous productivity growth. Sector-wide expenditures follow the exogenous variation in national income. The last component – the price indices – vary depending on the changes in individual firm's prices as follows:²⁶

$$\Delta P_{it} = \sum_{k \in \Omega_{it}^*} \phi_{kit} \Delta p_{kit} + \frac{1 - \rho}{\rho} \Delta \lambda_{it}, \quad \text{with } \phi_{kit} \equiv \frac{\frac{\vartheta_{kit} - \vartheta_{kit_0}}{\ln \vartheta_{kit} - \ln \vartheta_{kit_0}}}{\sum_{\ell \in \Omega_{it}^*} \frac{\vartheta_{\ell it} - \vartheta_{\ell it_0}}{\ln \vartheta_{\ell it} - \ln \vartheta_{\ell it_0}}}, \quad (1.14)$$

$$\Delta \Psi_{st} = \sum_{m \in \Upsilon_s} \Phi_{mt} \Delta P_{mt}, \quad \text{with } \Phi_{mt} \equiv \frac{\frac{\Theta_{mt} - \Theta_{mt_0}}{\ln \Theta_{mt} - \ln \Theta_{mt_0}}}{\sum_{h \in \Upsilon_s} \frac{\Theta_{ht} - \Theta_{ht_0}}{\ln \Theta_{ht} - \ln \Theta_{ht_0}}}, \quad (1.15)$$

where Ω_{it}^* denotes the set of varieties from industry i that are continuously sold between t and t_0 (also referred to as the continuing good set), and $\Delta \lambda_{it}$ indicates the log change in the market share of the common good set between t and t_0 , with $\lambda_{ist} \equiv \left(\sum_{\ell \in \Omega_{it}^*} y_{\ell it} \right) / \left(\sum_{\ell \in \Omega_{it}} y_{\ell it} \right)$. Terms ϕ_{kit} and Φ_{mt} indicate the weights applied to the price

²⁶See Appendix A.3.1 for the derivation of these expressions. We refer to Feenstra (1994) and Sato (1976); Vartia (1976) for this computation.

changes for firm k and industry m in sector s , which are computed following Sato (1976); Vartia (1976). The term $\vartheta_{kit} \equiv y_{kit} / \sum_{\ell \in \Omega_{it}^*} y_{\ell it}$ denotes the market share of firm k in the continuing good set from industry i at time t , and Θ_{mt} denotes the market share of industry m in sector s at time t (we assume that the set of industries does not change over time). These Sato-Vartia weights thus differ from simpler measures of market shares for firm f in industry i at time t , denoted by $\theta_{fit} \equiv y_{fit} / Y_{it}$, and for industry i in sector s at time t , denoted by $\Theta_{it} \equiv Y_{it} / Y_{st}$.

Combining the changes in the three components that determine firm f 's revenues (1.6) between post-regulation period t and pre-regulation base year t_0 yields

$$\begin{aligned} \Delta y_{fit} = & \frac{\rho \tau_t}{\rho - 1} R_{fit} + \frac{(\nu - \rho) \tau}{(\nu - 1)(1 - \rho)} \sum_{k \in \Omega_{it}^*} \phi_{kit} R_{kit} + \frac{(\nu - \rho)}{\rho(\nu - 1)} \Delta \lambda_{it} \\ & + \frac{\nu \tau}{1 - \nu} \sum_{m \in \Upsilon} \Phi_{mt} \sum_{\ell \in \Omega_{mt}^*} \phi_{\ell mt} R_{\ell mt} + \frac{\nu(1 - \rho)}{\rho(1 - \nu)} \sum_{m \in \Upsilon} \Phi_{mt} \Delta \lambda_{mt} + \Delta Y_t \\ & + \frac{\rho}{\rho - 1} \Delta A_{fit} + \frac{(\nu - \rho)}{(1 - \nu)(\rho - 1)} \sum_{k \in \Omega_{it}^*} \phi_{kit} \Delta A_{kit} + \frac{\nu}{1 - \nu} \sum_{m \in \Upsilon} \Phi_{mt} \sum_{\ell \in \Omega_{mt}^*} \phi_{\ell mt} \Delta A_{\ell mt}. \end{aligned} \quad (1.16)$$

The first term on the right hand side of equation (1.16) captures the direct effect of the regulation on revenues, holding all else constant. This term can be positive or negative, depending on the net regulation cost τ_t . The second and fourth terms capture the spillover effects of the regulation through its impact on price indices of the continuing good set within an industry and within a sector, respectively. Intuitively, these two terms depend on the combined market share of regulated firms within industry i and within sector s . The larger the market share of regulated firms, the larger the indirect effects. These spillover terms are common within an industry. The third term measures the change in the market share of the continuing good set in the industry, to account for entry and exit. The fifth term is just the sector-level analogue of the third term. The last line of equation (1.16) summarizes the revenue effects of the firm's own idiosyncratic productivity shock, as well as from the productivity shocks experienced by other firms in the same industry and sector. Input prices w_{Lt} and w_{Et} drop out of the equation, by virtue of being common to all firms.²⁷

²⁷Take any component of costs x , like input prices, that are the same for all firms. This component will cancel out of equation (1.16) since $\frac{\rho x}{1 - \rho} + \frac{(\nu - \rho) \sum_{k \in \Omega_{it}^*} \phi_{kit} x}{(\nu - 1)(\rho - 1)} + \frac{\nu \sum_{m \in \Upsilon} \Phi_{mt} \sum_{\ell \in \Omega_{mt}^*} \phi_{\ell mt} x}{\nu - 1} = x \left[\frac{\rho}{1 - \rho} + \frac{(\nu - \rho)}{(\nu - 1)(\rho - 1)} + \frac{\nu}{\nu - 1} \right] = 0$.

1.3.5 The Effect of Regulation on Emissions

The first order conditions for cost minimization yield the following expression for firm f 's emissions at time t :

$$Z_{fit} = \rho \gamma \kappa_t y_{fit} (w_{Et} e^{\mu^z R_{fit}})^{-1}. \quad (1.17)$$

The change in emissions from pre-period t_0 to period t can thus be written as:

$$\Delta Z_{fit} = -\mu_t^z R_{fit} + \Delta y_{fit} + \Delta \kappa_t - \Delta w_{Et}, \quad (1.18)$$

where Δy_{fit} corresponds to equation (1.16).

From this expression, we see that regulation affects the firm's emissions both by altering the scale of economic output, Δy_{fit} , and by encouraging substitution away from the polluting input whose price increases by a factor μ^z . We will refer to these as the “scale” and “technique” effects, respectively. The terms $\Delta \kappa_t$ and Δw_{Et} are not affected by the regulation, and thus move exogenously.

1.3.6 Extension to Multi-plant Firms

We now extend our theoretical framework to the case of multi-plant firms, where the regulation targets individual plants instead of firms. This will allow us to take account of intra-firm spillovers in addition to the inter-firm spillovers studied above.

Firm f still produces a unique variety, but it now owns J_{ft} plants at time t . These plants produce intermediate outputs using Cobb-Douglas combinations of labor and energy, and these intermediate outputs are then aggregated into a final output. This mapping of intermediate outputs to the final product is reminiscent of the approach of aggregating tasks to produce a final good, as in Acemoglu & Autor (2012). Presumably, firms own multiple plants because there exists complementarities across plants, which could stem from economies from vertical integration that improve productivity and lower overall costs of production (Alfaro et al., 2016) or from complementarities in intangible inputs, such as high-quality management, marketing know-how, or R&D capital (Atalay et al., 2014).

We assume that Hicks-neutral productivity, ω_{fit} , is common to all plants $j \in \{1, 2, \dots, J_{ft}\}$ within the firm. Regulation, however, can be plant-specific, with indicator R_{jfit} taking the value 1 if plant j is subject to regulation at time t and zero otherwise. Given these assumptions, plant-specific output and firm-level output are given by

$$Q_{jfit} = (L_{jfit})^{1-\gamma} (E_{jfit})^\gamma \exp(\omega_{fit} + \gamma \mu_t^e R_{jfit}) \quad (1.19)$$

and

$$Q_{fit} = \prod_{j=1}^{J_{ft}} \left(\alpha_{jfit}^{-\alpha_{jfit}} \right)^{1-\sigma} (Q_{jfit})^{\alpha_{jfit}}, \quad (1.20)$$

where α_{jfit} corresponds to the contribution of plant j to total output of firm f at time t , which is assumed to be exogenous. The term σ governs how the dispersion of economic activity across plants affects production. Its economic interpretation reflects a tension between the costs of dispersed activity, which dominate if $\sigma > 0$, and economies of scope (that reduce the unit cost), which dominate if $\sigma < 0$. Dispersion forces emerge whenever a firm adds new plants, since Cobb-Douglas aggregation function makes it more profitable to produce everything in a single plant. By contrast, economies of scope follow from the fact that the more similar-sized plants a firm own, the more complementarities it generate so that overall output increases. For a given amount of plants owned by a firm, the more concentrated production is, the smaller the gain from complementarities across plants.

Cost minimization yields the following firm-level unit cost function:

$$c_{fit} = (1 - \gamma)^{-(1-\gamma)} \gamma^{-\gamma} \left(\prod_{j=1}^{J_{ft}} \alpha_{jfit}^{-\alpha_{jfit}} \right)^{\sigma} (w_{Lt})^{1-\gamma} (w_{Et})^{\gamma} \exp(-\omega_{fit} + \tau_t R_{fit}). \quad (1.21)$$

This is almost identical to (1.10), except that we now define the firm's regulatory status as a function of plant-level regulation, $R_{fit} \equiv \sum_j \alpha_{jfit} R_{jfit}$, and include a term which reflects how the dispersion of economic activity across plants affects costs, $\prod_{j=1}^{J_{ft}} (\alpha_{jfit}^{-\alpha_{jfit}})^{\sigma}$. Since cost minimization implies $L_{jfit} = \alpha_{jfit} L_{fit}$, we can compute R_{fit} by using plant-specific labor shares combined with plant-level regulation.

Defining the firm-level dispersion term by $\alpha_{fit} \equiv \prod_{j=1}^{J_{ft}} \alpha_{jfit}^{-\alpha_{jfit}}$, and collecting structural parameters into constants, the changes in firm-level revenues between a pre-regulation period and a post-regulation period can be written as

$$\begin{aligned} \Delta y_{fit} = & \beta_0 R_{fit} + \beta_1 \Delta \alpha_{fit} + \beta_2 \sum_{k \in \Omega_{it}^*} \phi_{kit} R_{kit} + \beta_3 \sum_{k \in \Omega_{it}^*} \phi_{kit} \Delta \alpha_{kit} + \beta_4 \Delta \lambda_{it} \\ & + \beta_5 \sum_{m \in \Upsilon} \Phi_{mt} \sum_{\ell \in \Omega_{mt}^*} \phi_{\ell mt} R_{\ell mt} + \beta_6 \sum_{m \in \Upsilon} \Phi_{mt} \sum_{\ell \in \Omega_{mt}^*} \phi_{\ell mt} \Delta \alpha_{\ell mt} + \beta_7 \sum_{m \in \Upsilon} \Phi_{mt} \Delta \lambda_{mt} \\ & + \beta_8 \sum_{k \in \Omega_{it}^*} \phi_{kit} \Delta A_{kit} + \beta_9 \sum_{m \in \Upsilon} \Phi_{mt} \sum_{\ell \in \Omega_{mt}^*} \phi_{\ell mt} \Delta A_{\ell mt} + \Delta Y_t + \xi_{fit}, \end{aligned} \quad (1.22)$$

with $\xi_{fit} = \frac{\rho}{\rho-1} \Delta A_{fit}$.²⁸ Compared to equation (1.16), this multi-plant extension adds

²⁸Other parameters are defined as $\beta_0 \equiv \frac{\rho\tau}{\rho-1}$, $\beta_1 \equiv \frac{\rho\sigma}{\rho-1}$, $\beta_2 \equiv \frac{(\nu-\rho)\tau}{(\nu-1)(1-\rho)}$, $\beta_4 = \frac{(\nu-\rho)}{\rho(\nu-1)}$, $\beta_3 \equiv \frac{(\nu-\rho)\sigma}{(\nu-1)(1-\rho)}$

three terms: the exogenous change in firm f 's dispersion term, $\Delta\alpha_{fit}$, and the industry- and sector-level counterparts.

Cost minimization at the plant level implies

$$\Delta Z_{jfit} = -\mu^z R_{jfit} + \Delta\alpha_{jfit} + \Delta\kappa_t - \Delta w_{Et} + \Delta y_{fit}. \quad (1.23)$$

Plant-level emissions are thus affected by a firm-level scale effect and a plant-level technique effect $-\mu^z$, as before, but they also depend on the share of labor allocated to plant j .

1.4 Estimation Strategy

In this section, we develop a procedure to estimate the effects of regulation on revenues and emissions based on the model from section 1.3. We present our procedure for estimating reduced-form parameters and then show how to use these estimates to compute the treatment effects of the regulation. Finally, we compare our estimation procedure to existing methods in Monte Carlo experiments.

1.4.1 Estimating Parameters

The model indicates that the change in firm-level revenues between a pre-regulation period and a post-regulation period can be written as a function of regulation assignments, market shares, and unobserved productivity shocks (see eq. 1.22). Collecting all sector-wide variables into sector-time fixed effects, δ_{st} , yields

$$\begin{aligned} \Delta y_{fit} = & \beta_0 R_{fit} + \beta_1 \Delta\alpha_{fit} + \beta_2 \sum_{k \in \Omega_{it}^*} \phi_{kit} R_{kit} + \beta_3 \sum_{k \in \Omega_{it}^*} \phi_{kit} \Delta\alpha_{kit} + \beta_4 \Delta\lambda_{it} \\ & + \beta_8 \sum_{k \in \Omega_{it}^*} \phi_{kit} \Delta A_{kit} + \delta_{st} + \frac{\rho}{1-\rho} \Delta A_{fit}. \end{aligned} \quad (1.24)$$

All variables in the first line of equation (1.24) are observed, while variables in the second line are not. Under a random walk assumption for productivity growth, ΔA_{fit} s would be exogenous to the regulation, even if regulation depended on pre-period characteristics. In this case, it is tempting to group $\sum_{k \in \Omega_{it}^*} \phi_{kit} \Delta A_{kit}$ into the error term and estimate (1.24) by OLS. However, there are two identification problems. First, $\sum_{k \in \Omega_{it}^*} \phi_{kit} R_{kit}$ correlates with $\sum_{k \in \Omega_{it}^*} \phi_{kit} \Delta A_{kit}$. This is because the covariance of ΔA_{kit} and ϕ_{kit} depends on the density of regulation in the industry. For example, if the entire industry is unregulated,

$$\beta_5 \equiv \frac{\nu\tau}{1-\nu} \quad \beta_6 \equiv \frac{\nu\sigma}{1-\nu} \quad , \quad \beta_7 \equiv \frac{\nu(1-\rho)}{\rho(1-\nu)} \quad , \quad \beta_8 \equiv \frac{(\nu-\rho)}{(1-\nu)(\rho-1)} \quad , \quad \beta_9 \equiv \frac{\nu}{1-\nu}.$$

then ϕ_{kit} depends exclusively on A_s and α_s . By contrast, as regulation density in the industry increases, ϕ_{kit} depends less on A_s and α_s and more on R_{fits} . Hence, industries with high regulation density (high $\sum_{k \in \Omega_{it}^*} \phi_{kit} R_{kit}$) will have low correlations between ΔA_{kit} and ϕ_{kit} , and thus low values of $\sum_{k \in \Omega_{it}^*} \phi_{kit} \Delta A_{kit}$.²⁹ Secondly, industries with higher weighted-average productivity growth of continuing firms (high values of $\sum_{k \in \Omega_{it}^*} \phi_{kit} \Delta A_{kit}$) will mechanically have higher revenue share growth of continuing firms. Hence, we would expect $\Delta \lambda_{it}$ to correlate with $\sum_{k \in \Omega_{it}^*} \phi_{kit} \Delta A_{kit}$ as well.

To address these identification problems, we build a GMM procedure to simultaneously estimates β_0, β_1 , and β_2 . Given estimates of $\beta_0, \beta_1, \beta_2$, all other reduced-form parameters can be computed.³⁰ First, we re-write equation (1.24) as a system of equations:

$$\begin{aligned} \Delta y_{fist} &= \beta_0 R_{fist} + \beta_1 \Delta \alpha_{fist} + \eta_{ist} + \xi_{fist} \\ \eta_{ist} &= \beta_2 \sum_{k \in \Omega_{ist}^*} \phi_{kist} R_{kist} + \beta_1 \frac{\beta_2}{\beta_0} \sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta \alpha_{kist} - \frac{\beta_2}{\beta_0} \left(\Delta \lambda_{ist} - \sum_{k \in \Omega_{ist}^*} \phi_{kist} \xi_{fist} \right) + \delta_{st} + \varsigma_{ist}, \end{aligned}$$

with $\xi_{fist} \equiv \frac{\rho}{1-\rho} (A_{fist} - \overline{\Delta A_{ist}})$ and $\varsigma_{ist} \equiv \frac{\nu}{1-\nu} \overline{\Delta A_{ist}}$. Notice that the first equation is almost identical to the DiD model (1.1), except that we control for the firm-level dispersion term $\Delta \alpha_{fist}$.³¹ In the second equation that regresses industry-time fixed effects on industry-specific intensity of treatment, dispersion of activities and entry and exit controls, we purged the error term from Sato-Vartia weights using the residuals ξ_{fist} from the first equation.

Next, for any candidate vector $\beta^* = (\beta_0^*, \beta_1^*, \beta_2^*)$, we can compute

$$\widehat{\xi_{fist}} = \Delta y_{fist} - \beta_0^* R_{fist} - \beta_1^* \Delta \alpha_{fist} - \left[\frac{1}{N_{ist}^*} \sum_{k \in \Omega_{ist}^*} \Delta y_{kist} - \beta_0^* R_{kist} - \beta_1^* \Delta \alpha_{kist} \right] \quad (1.25)$$

and

$$\widehat{\varsigma_{ist}} = \left[\frac{1}{N_{ist}^*} \sum_{k \in \Omega_{ist}^*} \Delta y_{kist} - \beta_0^* R_{kist} - \beta_1^* \Delta \alpha_{kist} \right] - \beta_2^* \sum_{k \in \Omega_{ist}^*} \phi_{kist} R_{kist}$$

²⁹This problem cannot be solved simply by instrumenting $\sum_{k \in \Omega_{it}^*} \phi_{kit} R_{kit}$ with base-year weighted average regulation, because base-year weighted average regulation would correlate with $\sum_{k \in \Omega_{it}^*} \phi_{kit} \Delta A_{kit}$ as well. Rather, we would need an instrument that does not depend on the regulation.

³⁰We have $\beta_3 = \beta_1 * \beta_2 / \beta_0$, $\beta_4 = -\beta_2 / \beta_0$, $\beta_5 = -(\beta_0 + \beta_2)$, $\beta_6 = -(\beta_1 + \beta_3)$, $\beta_7 = (\beta_0 + \beta_2) / \beta_0$. Given the use of the estimates $\widehat{\xi_{fist}}$, the other two coefficients are functions of $\beta_0, \beta_1, \beta_2$, and ρ , but the term in ρ cancels out with the estimates $\widehat{\xi_{fist}}$. We have $\beta_8 = -(\beta_2 / \beta_0) * \rho / (1 - \rho)$, and $\beta_9 = (1 + \beta_2 / \beta_0) * \rho / (1 - \rho)$.

³¹Our model indicates that change in across-plant concentration $\Delta \alpha_{fit}$ should be controlled for because we expect that regulation assignment correlates with this concentration measure. When $\Delta \alpha_{fit}$ is large, economic activity across the firm becomes more dispersed over time. With greater dispersion, the labor share of regulated plants falls, so R_{fist} falls. Hence, leaving $\Delta \alpha_{fit}$ for the error term generates omitted variable bias.

$$- \beta_1^* \frac{\beta_2^*}{\beta_0^*} \sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta \alpha_{kist} + \frac{\beta_2^*}{\beta_0^*} \left(\Delta \lambda_{ist} - \sum_{k \in \Omega_{ist}^*} \phi_{kist} \widehat{\xi_{fist}} \right) - \delta_{st}, \quad (1.26)$$

where N_{ist}^* represents the number of continuing firms in industry i in post-regulation period t .³²

We then exploit the following moment conditions: (1) $E[\xi_{fist} R_{fist}] = 0$, (2) $E[\xi_{fist} \Delta \alpha_{fist}] = 0$, (3) $E[\varsigma_{ist} (\sum_{k \in \Omega_{ist}^*} \phi_{kist} R_{kist})] = 0$, and (4) $E[\varsigma_{ist} (\sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta \alpha_{kist})] = 0$. The first and third moment conditions hold under the assumption of random walk productivity growth. However, if productivity growth does not follow a random walk, and if regulation correlates with pre-period characteristics, then these conditions may not hold. To address this possibility, we control non-parametrically for observable determinants of regulation at the firm-level. In particular, we match regulated firms to unregulated firms in the same industry based on pre-period levels and growth of y_{fit} and condition on this sample when building the empirical counterpart to the first condition. To address endogeneity of regulation at the industry level, we purge ς_{ist} of industry-level average pre-period outputs and inputs before building the empirical moments. The second and fourth moment conditions follow from the exogeneity of α_{fist} .

Using (1.25) and (1.26), we build the empirical moments

$$\mathbf{\Gamma}(\beta_0^*, \beta_1^*, \beta_2^*) = \begin{pmatrix} \sum_t \sum_{f,i} \xi_{fist} \widehat{(\beta_0^*, \beta_1^*)} R_{fist} \\ \sum_t \sum_{f,i} \xi_{fist} \widehat{(\beta_0^*, \beta_1^*)} \Delta \alpha_{fist} \\ \sum_t \sum_i \varsigma_{ist} (\beta_0^*, \beta_1^*, \beta_2^*) \left(\sum_{k \in \Omega_{ist}^*} \phi_{kist} R_{kist} \right) \\ \sum_t \sum_i \varsigma_{ist} \widehat{(\beta_0^*, \beta_1^*, \beta_2^*)} \left(\sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta \alpha_{kist} \right) \end{pmatrix}$$

and choose β to minimize the usual GMM criterion function

$$\beta^{STR} = \min_{\beta_0^*, \beta_1^*, \beta_2^*} \mathbf{\Gamma}(\beta_0^*, \beta_1^*, \beta_2^*)' \mathbf{W} \mathbf{\Gamma}(\beta_0^*, \beta_1^*, \beta_2^*)$$

where the weighting matrix \mathbf{W} is estimated optimally using the two-step GMM estimator.

Finally, at the plant level, we estimate by OLS:

$$\Delta Z_{jfirst} = \beta_z R_{jfirst} + \delta_\alpha \Delta \alpha_{jfirst} + \delta_{first} + \epsilon_{jfirst}, \quad (1.27)$$

³²To control for the sector-year fixed effects, we first compute

$$\left[\frac{1}{N_{ist}} \sum_{k \in \Omega_{ist}^*} \Delta y_{kist} - \beta_0^* R_{kist} - \beta_1^* \Delta \alpha_{kist} \right] - \beta_2^* \sum_{k \in \Omega_{ist}^*} \phi_{kist} R_{kist} - \beta_1^* \frac{\beta_2^*}{\beta_0^*} \sum_{k \in \Omega_{ist}^*} \phi_{kist} \Delta \alpha_{kist} + \frac{\beta_2^*}{\beta_0^*} \left(\Delta \lambda_{ist} - \sum_{k \in \Omega_{ist}^*} \phi_{kist} \widehat{\xi_{fist}} \right),$$

then regress this quantity on a vector of sector-by-year indicator variables and extract the residual.

where δ_{fist} absorbs all determinants of emissions that are constant across plants within the firm-year. The parameter β_z represents the input price effect of the regulation on emissions ($\beta_z = -\mu_z$), holding all else constant. By exploiting within-firm variation, this equation may not suffer from the selection bias stemming from nonrandom treatment assignment, as it was the case for revenues.

1.4.2 Treatment Effects

In order to assess the effects of the regulation on average and aggregate outcomes, we must compute the counterfactual equilibrium for the case in which no firm was ever regulated and compare it to observed outcomes. This is necessary because industry- and sector-wide price indices will change with counterfactual regulation regimes. Because Sato-Vartia weights depend on both pre-period market shares as well as current market shares, we must first solve for counterfactual revenues, and then use the first order conditions to solve for counterfactual emissions at the plant level.

At the firm level, we obtain counterfactual revenues by solving a system of non-linear equations. Using equation (1.22), counterfactual revenue changes between t and t_0 for any firm with positive revenues in the post-regulation period (where the ' symbol indicates a counterfactual value) can be written as:

$$\begin{aligned}
\ln(y'_{fist}) &= \ln(y_{fist}) - \hat{\beta}_0 R_{fist} - \hat{\beta}_2 \sum_{k \in \Omega_{ist}^*} \phi_{kist} R_{kist} + (\hat{\beta}_0 + \hat{\beta}_2) \sum_{m \in \Upsilon_s} \Phi_{mst} \sum_{\ell \in \Omega_{mst}^*} \phi_{\ell mst} R_{\ell mst} \\
&+ \frac{\hat{\beta}_2}{\hat{\beta}_0} \sum_{k \in \Omega_{ist}^*} (\phi'_{kist} - \phi_{kist}) \widehat{\xi_{kist}} - \frac{\hat{\beta}_2}{\hat{\beta}_0} (\Delta \lambda'_{ist} - \Delta \lambda_{ist}) + \hat{\beta}_1 \frac{\hat{\beta}_2}{\hat{\beta}_0} \sum_{k \in \Omega_{ist}^*} (\phi'_{kist} - \phi_{kist}) \Delta \alpha_{kist} \\
&+ \frac{\hat{\beta}_0 + \hat{\beta}_2}{\hat{\beta}_0} \left(\sum_{m \in \Upsilon_s} \Phi'_{mst} \Delta \lambda'_{mst} - \sum_{m \in \Upsilon_s} \Phi_{mst} \Delta \lambda_{mst} \right) - \sum_{m \in \Upsilon_s} \left(\Phi'_{mst} - \Phi_{mst} \right) \widehat{\varsigma_{mst}} \\
&- \frac{\hat{\beta}_0 + \hat{\beta}_2}{\hat{\beta}_0} \sum_{m \in \Upsilon_s} \sum_{\ell \in \Omega_{mst}^*} (\Phi'_{mst} \phi'_{\ell mst} - \Phi_{mst} \phi_{\ell mst}) \left(\widehat{\xi_{\ell mst}} + \hat{\beta}_1 \Delta \alpha_{\ell mst} \right) \quad (1.28)
\end{aligned}$$

where ϕ'_{fist} , Φ'_{ist} and $\Delta \lambda'_{ist}$ indicate counterfactual Sato-Vartia weights and continuing-good share growth, which must be solved for endogenously. Terms $\widehat{\varsigma_{mst}}$ and $\widehat{\xi_{kist}}$ are computed according to (1.25) and (1.26), whereas $\hat{\beta}_0$, $\hat{\beta}_1$ and $\hat{\beta}_2$ result from the GMM procedure. The counterfactual equilibrium is the vector of counterfactual revenues y'_{fist} that solves the system defined by one equation (1.28) for each firm. The equilibrium can be solved either with a numerical solver or with a fixed-point algorithm.³³

³³In the fixed point algorithm, we first set $y'_{fist} = y_{fist}$, compute ϕ'_{fist} , Φ'_{ist} and λ'_{ist} under this assumption, compute the right hand side of (1.28), and then update y'_{fist} . The process then repeats until the vector y'_{fist} converges.

Once we have an estimate of y'_{fist} , counterfactual emissions at the plant level are easily computed as $Z'_{j fist} = Z_{j fist} e^{\widehat{\beta_z} R_{j fist} (y'_{fist}/y_{fist})}$ and, at the firm level, as $Z'_{fist} = \sum_j Z'_{j fist}$.

We then compute average treatment effects at the firm level and plant level. At the firm level, we define “treated” (T) firms as any firm with at least one regulated plant and “control” (C) firms as any firm without any regulated plant. At the plant level, we denote regulated plants “TT” (treated plants in treated firms), unregulated plants owned by firms with regulated plants “CT” (control plants in treated firms), and unregulated plants owned by firms without any regulated plant “CC” (control plants in control firms). Then for any group $X \in \{T, C, TT, CT, CC\}$ and outcome $v \in \{y, Z\}$, we define the average treatment effect as:

$$ATX_t^v = \frac{1}{N_t^X} \sum_{f \in \Omega_t^X} \frac{v_{fist} - v'_{fist}}{v_{fist}}, \quad (1.29)$$

where N_t^X and Ω_t^X correspond to the number and the set of studied firms or plants at time t , respectively. Firm-level treatment effects vary across firms within an industry due to the variation in the density of treatment across plants. Additionally, densities of treatment within an industry and within a sector are also heterogeneous, which leads to variation in average treatment effects across industries. At the aggregate level, we evaluate the treatment effect of the regulation on total emissions as $ATZ_t = (Z_t^{tot} - Z_t^{tot'})/Z_t^{tot}$.

1.4.3 Monte Carlo Experiments

In this section, we use Monte Carlo experiments to evaluate the performance of our estimator in a finite-sample setting and compare it to alternative estimators.

We assume the following data generating process. A fixed number of firms could potentially operate in each of two periods - a pre-regulation period and a post-regulation period. Sector-wide expenditures are drawn randomly each period, as well as firm-level entry and exit decisions. In the pre-regulation period, firms draw their productivity, number of plants, and $\alpha_{j fist}$ terms randomly. These variables completely determine marginal cost, and hence prices and revenues, given monopolistic competition. We solve for plant-level emissions using the first order conditions. In the second period, we set $R_{j fist} = 1$ for all plants with pre-period emissions above a given threshold. Productivity updates according to a first order Markov process, and marginal costs update based on regulation and updated productivity. Revenues, market shares and plant-level emissions are solved again. We simulate the model for eight different parameter combinations. We vary $\rho \in \{.8, .95\}$, the energy-efficiency effect of the regulation $\mu^e \in \{0, .8\}$, and the probability of exit $\varrho \in \{0, .1\}$. In all combinations, we set $\nu = .3$, $\gamma = .8$, $\mu^z = .2$, and $\sigma = -1$.

For each parameter combination, we simulate 100 replications. For each replication, we simulate 10 sectors, 5 industries per sector and 10 firms per industry.

Once the data is generated, we implement our estimation procedure, as well as alternative estimators that fall either in the category of DiD approaches or in the category of local approximation approaches. At the firm level, we compare our estimator to the standard DiD model, as described in equation (1.1) while making two alterations to fit our model: we include the change in dispersion forces $\Delta\alpha_{fit}$ as a control and define firm-level regulation status as the share of labor working in regulated plants. In line with common practices in the literature, we compute treatment effects ignoring the fact that CES price indices would differ in a counterfactual equilibrium, and thus build counterfactual outcomes as $\ln(v_{fist}^{DD'}) = \ln(v_{fist}) - \widehat{\beta_{v0}^{DD}} R_{fist}$. We also consider the LA model described in equation (1.2), with the same two alterations and with an additional control $\sum_{k \in \Omega_{ist_0}} \theta_{kist_0} \Delta\alpha_{kit}$ to fit our model. To compute treatment effects in this case, we use

$$\ln(v_{fist}^{LA'}) = \ln(v_{fist}) - \widehat{\beta_{v0}^{LA}} R_{fist} - \widehat{\beta_{v2}^{LA}} \sum_{k \in \Omega_{ist_0}} \theta_{kist_0} R_{kist} + (\widehat{\beta_{v0}^{LA}} + \widehat{\beta_{v2}^{LA}}) \sum_{m \in \Upsilon} \Theta_{mt_0} \sum_{k \in \Omega_{ist_0}} \theta_{kist_0} R_{kist},$$

exploiting the structural relationship identified in our model to compute the parameter in front of sector-wide spillovers. Doing this, we are getting as close as possible to our estimator while not having to solve for the counterfactual equilibrium with the system of equations (1.28). We show in Appendix A.3.2 that this amended LA model follows from taking a local approximation to the price index around the pre-regulation equilibrium.

Similarly, at the plant level, we consider a DiD model that regress plant-level emissions on plant-specific regulation status and change in labor share, $\Delta\alpha_{jfirst}$, controlling for industry-time fixed effects, δ_{ist} , in a sample that excludes unregulated plants that are owned by firms that owns at least one other regulated plant. We build counterfactual outcomes as before: $\ln(Z_{jfirst}^{DD-TT'}) = \ln(Z_{jfirst}) - \beta_{z0}^{DD-TT} R_{jfirst}$. Additionally, following previous work that investigates the role of across-plant reallocation within the firm (Bartram et al., 2022; Gibson, 2019; Soliman, 2020), we compare the evolution of plant-level outcomes for unregulated plants owned by firms that also own regulated plants to unregulated plants owned by firms that own no regulated plant, excluding directly regulated plants from the sample. We call this DiD model “DD-CT” and notice that, by only including non-directly-treated plants, the estimates capture exclusively a scale effect. Finally, even though we have not found in the literature examples of the LA model combined with a multi-plant structure, we can build the plant-level LA model based on our theoretical framework by combining the firm-level regression on revenues described in (1.2) with our plant-level

specification (1.27) to compute the following counterfactual outcomes:

$$\begin{aligned} \ln(Z_{fist}^{LA'}) &= \ln(Z_{fist}) - \widehat{\beta}_z^{LA} R_{jfist} - \widehat{\beta}_{y0}^{LA} R_{fist} - \widehat{\beta}_{y2}^{LA} \sum_{k \in \Omega_{ist_0}} \vartheta_{kist_0} R_{kist} \\ &+ (\widehat{\beta}_{y0}^{LA} + \widehat{\beta}_{y2}^{LA}) \sum_{m \in \Upsilon} \Theta_{mt_0} \sum_{k \in \Omega_{ist_0}} \vartheta_{kist_0} R_{kist}. \end{aligned}$$

Armed with our estimates, which we call structural “STR”, as well as the different versions of the DiD and LA estimators presented above, we compute firm-level ATT and ATC, plant-level ATTT, ATCT, ATCC, and aggregate effects on emissions for each simulated data set. Since we control the data generating process, we can also compute the true treatment effect of the regulation on any given firm by comparing the firm-level outcome in the regulated scenario to the unregulated scenario, keeping all exogenous parameters the same.

Figure 1.4.1 plots, for each parameter combination, the true average values across 100 replications for ATT (left) and ATC (right) for revenues (top) and emissions (bottom) on the x-axis and the average estimated value for each estimator on the y-axis for multi-plant firms. We find that (i) our procedure (indicated in red) successfully recovers unbiased estimates of the ATT and ATC, even for finite samples, whereas (ii) the DiD estimator (indicated in black) tends to overstate the effect on revenues of regulated firms in magnitude, while recovering the correct sign of the net regulation cost, and (iii) the LA estimator (indicated in blue) suffers from a small bias when the true effect is small in magnitude, but it can be severely biased in some parameterizations (for instance, when varieties are more substitutable with $\rho = .95$) and even have the wrong sign.³⁴ In the lower panels of Figure 1.4.1 we see that the biases in the DiD and the LA estimators of the average treatment effects on firm-level emissions follow a similar pattern as for revenues. However, the sign of the DiD estimate does not always correspond with the sign of the true ATT, as it does for revenues. This is because the DiD estimator in emissions conflates the scale effect with the input-price effect; and since the DiD over estimates the scale effect, the DiD estimator can yield the wrong sign. In general, however, we cannot sign the bias for the LA model, nor for the DiD model (see Appendix A.4.1).

³⁴For each parameter combination, Figure A.4.1 shows that the distribution of the biases in estimates of the ATT for our estimator is centered on zero. We also see that even when the bias in the DiD estimator is small (see the top left panel), the distribution is clearly centered away from zero. This can also be the case for the LA estimator. In Appendix A.4.2, we decompose the bias in the DiD and LA estimators into three components: specification error, weighting or spillover error, and regression error. For the multi-plant simulation results, the specification error dominates for LA, and hence dictates the sign of the bias. However, this is not true for all simulations, as indicated by the single-plant simulation results shown in the same Appendix.

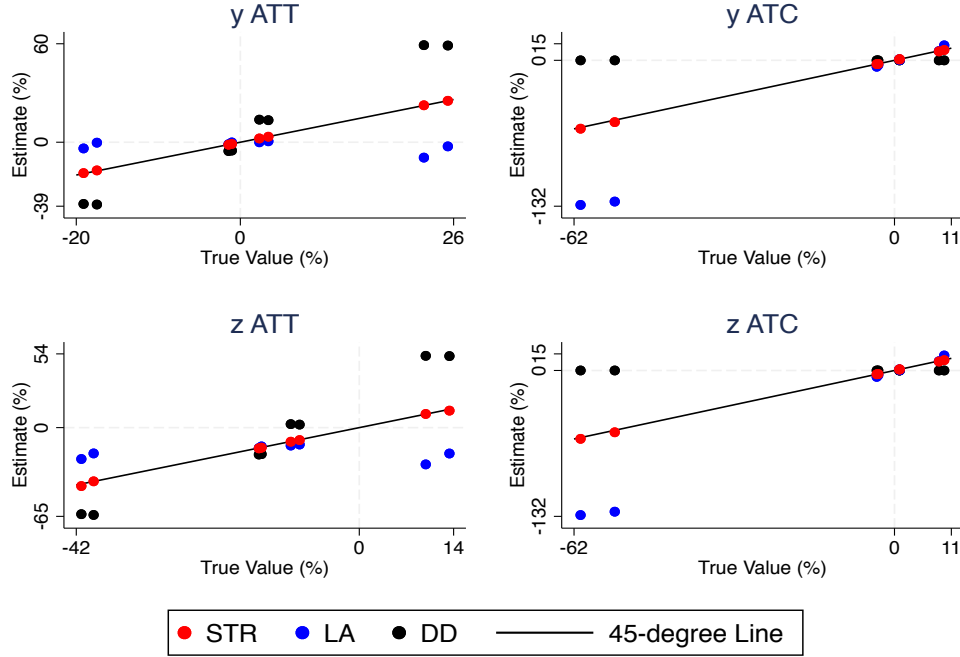


Figure 1.4.1: Average Treatment Effects at the Firm Level

Notes: Subfigures plot average estimates of treatment effects on revenues (y) and emissions (z) as a percentage of observed outcomes for each estimator across 100 replications against the average of the true metric, for each parameter combination. Left panels show the ATTs for each parameter combination, whereas right panels show the ATCs, all at the firm level. The black line corresponds to the 45-degree line.

Moving to plant-level models, Figure 1.4.2 indicates that our estimator successfully recovers unbiased estimates of treatment effects on plant-level emissions. By contrast, the DiD and LA estimators fail to recover unbiased estimates. Looking at the ATTT, the biases in the DiD and the LA follow a similar pattern as for firm-level emissions. It is thus not possible in general to sign the bias in these two estimators with respect to plant-level emissions. Looking at the ATCT, the DiD estimator that focuses on intra-firm reallocation can overestimate or underestimate these intra-firm spillovers, and even have the wrong sign.³⁵ Finally, looking at treatment effects on aggregate emissions in Figure A.4.2, we find that aggregate emissions fall across all parameter combinations since regulation always increases the cost of emissions ($\mu^z > 0$). Our estimator successfully recovers unbiased estimates of these aggregate effects whereas both the DiD and LA estimators yield biased estimates for all parameter combinations. For both estimators, the bias can

³⁵In our simulations, even when regulation lowers the unit costs of regulated firms, emissions from unregulated plants in regulated firms falls on average. This is because firms with both regulated and unregulated plants are not the biggest firms in the economy (only very large firms will have all plants regulated). Hence, while regulation lowers marginal costs for regulated firms when $\tau < 0$, firms with only some regulated plants may still lose market share to fully regulated firms.




Figure 1.4.2: Average Treatment Effect in Emissions at the Plant Level

Notes: Subfigures plot average estimates of treatment effects on emissions (z) at the plant level as a percentage of observed emissions for each estimator across 100 replications against the average of the true treatment effect, for each parameter combination. The first panel plots the ATTTs for each parameter combination, the second ATCTs, and the third ATCCs. The black line corresponds to the 45-degree line. A plant is labeled “TT” if it is regulated in period 2, “CT” if the plant is unregulated in period 2, but part of a firm with regulated plants in period 2, and “CC” if the plant is part of a firm with no regulated plants.

be either up or down, and the estimated sign can be either positive or negative.

In the model, we assume that firms do not internalize the effect of their pricing on the industry-wide price index. However, with a discrete number of firms, some of whom command nontrivial market shares, it could be that, in reality, at least some firms do internalize the effect of their pricing behavior on the CES price index. In this case, our estimator would be misspecified. To explore the potential role of misspecification bias with respect to strategic behavior, we simulate the economy assuming firms engage in Bertrand-Nash pricing. All other features of the simulations remain the same. We present results in Appendix Figures A.5.1 - A.5.4. We find that, even if firms play Bertrand-Nash, our procedure still delivers estimates of treatment effects that matches the true values quite well. This is reminiscent of the findings in Head & Mayer (2021) of the good empirical performance of the CES-MC model at estimating the effect of cost changes (international tariff changes, in their case) even though the world is more complex with firms being strategic and consumers more sophisticated.

1.5 Moving from Model to Data

In this section, we describe three empirical issues that we must address in moving from the model to the data. Additionally, we describe our matching procedure.

1.5.1 Export Markets

First, the model considers only one economy, but in reality, French firms serve not only the domestic market but also foreign markets. Our empirical strategy relies on constructing empirical counterparts to the theoretical CES price index, which requires information on all firms selling to a destination market. Focusing on the French market, we distinguish between French firm sales in France vs sales abroad. For most of the analyses below, we take domestic sales of French firms as the outcome variable. Total domestic sales represents about 75% of total French-firm revenues, depending on the year of the sample (see Figure A.2.3, left). For these flows, we can compute counterfactual revenues using our method. For comparison to other methods, we sometimes present corresponding evidence of effects on total sales.

Since our model is best suited to analyzing the domestic market, we also distinguish between total CO₂ emissions and CO₂ emissions generated in production for domestic sales. We compute this latter measure by multiplying total emissions by the domestic revenue share. Figure A.2.3 (right) shows that our measure of total emissions (inclusive of imputations described below) tracks trends in the National Emissions Inventory measure reasonably well.

1.5.2 International Imports

Second, while focusing on the French market, we need to account for foreign imports into France. While we observe the universe of French firms selling in the French market, we do not observe flows of individual foreign firms selling in France. These missing flows – which represent roughly a third of all French manufacturing consumption (see Figure A.2.2) – are problematic for building the different components of the CES price indices. To circumvent this issue, our main strategy exploits aggregate bilateral trade flows from the BACI database (Gaulier & Zignago, 2010) for which we assume that each origin-country-industry is associated with only one firm exporting to France.

It is however unlikely that the set of foreign firms shipping to France within an industry is always singleton. To investigate the role of this mis-specification error, we conduct robustness tests under alternative assumptions with respect to international trade. In Appendix Figures A.1.2-A.1.5, we repeat the analysis leaving foreign imports out of the computation

entirely. This exercise amounts to assuming Cobb-Douglas aggregation between domestic consumption and foreign consumption at the sector level so that the share of expenditures on domestic goods would be fixed. Under this assumption, industry-level CES price indices on the French market can be separated into a CES price index for domestic varieties and a CES price index for foreign varieties, and the two are completely independent. As a result, we could construct the CES price index for domestic varieties using only French firms and ignoring international imports. This is implicitly the approach taken by previous work based on local approximations (Cai & Szeidl, 2022; Muehlegger & Sweeney, 2021; Rotemberg, 2019).

1.5.3 Imputing Emissions

Third, only about 10% of all plants are included in the EACEI survey each year (see Figure A.2.1). As long as there are no systematic differences between surveyed and unsurveyed plants, conditional on observables, missing data poses no problem for the estimation and we compute our estimates of the treatment effects on the set of observed plants. However, ideally we would like to know how the regulation affected total CO₂ emissions from French manufacturing, not just the emissions from the plants observed in the EACEI. To extend the analysis to plants beyond the EACEI, we can leverage the structure of the model since emission intensity in revenues does not vary across plants within sector-year.³⁶ By taking the average of CO₂/revenue for all surveyed plants at the sector-year level, we can impute emissions for plants not included in the EACEI. See Data Appendix A.2.6 for details. This procedure is done for completeness, though we acknowledge that this extension requires potentially strong assumptions.

1.5.4 Matching

To address nonrandom treatment assignment, we match regulated firms to unregulated firms within an industry on pre-regulation characteristics. The model indicates that these characteristics should include $\ln y_{fis,2004}$ and $\alpha_{fis,2004}$, as well as pre-period average growth rate of firm-level sales. Intuitively, two firms with the same values for these characteristics should have the same productivity level, and hence should grow at the same rate in expectation following 2004, absent any effects of regulation.³⁷ To implement this matching exercise, we project an indicator variable for whether or not a firm ever

³⁶In practice, observed emission intensity in revenues does vary, but this could be due to measurement error.

³⁷Intuitively, we need to control for $\alpha_{fis,2004}$ because it has a direct effect on costs and correlates with regulation. If we omit $\alpha_{fis,2004}$ from the matching procedure, then regulated firms will systematically have higher (lower) values of $\omega_{fis,2004}$ if σ is positive (negative).

operates a regulated plant on $\ln y_{fis,2004}$, $\alpha_{fis,2004}$, and the pre-period average growth rate of firm-level sales.³⁸ We compute predicted regulation, and then identify all firms with predicted regulation scores within a given bandwidth of a regulated firm in the same industry. This procedure excludes some regulated firms from the sample for whom no unregulated firm can be found close enough in propensity score within the same industry.

By matching regulated to unregulated firms, we attempt to minimize the selection bias stemming from nonrandom regulation assignment. However, given that regulation is largely determined by pre-period size, it is important to consider where variation in firm-level regulation status comes from, conditional on size. Following Calem & Dechezlepretre (2016), most studies of the EU ETS assume that firm-level variation in regulation status comes from variation in the across-plant concentration of economic activity within the firm. Previous work has merely asserted that this variation is orthogonal to unobserved determinants of revenue growth. However, according to our model, across-plant concentration of economic activity within the firm has a direct effect on costs (c_{fist} depends on α_{fist}). Hence, the structure of our model indicates that identifying variation must come from some other measure of across-plant concentration, conditional on $\alpha_{fis,2004}$.³⁹ In appendix Table A.1.1, we show that in fact across-plant concentration as measured by $\sum_{j \in \Omega_{fis,2004}} \alpha_{jfis,2004}^2$ correlates with future regulation status, even conditional on $\alpha_{fis,2004}$. Thus, we provide empirical support for the conditional version of this widely cited identification assumption.

At the plant level, our approach exploits within-firm-year variation across plants. Hence, restricting the sample to matched regulated and unregulated plants is not necessary. We however test for robustness using a matched sample, where we match plants based on their pre-period CO₂ emission levels and growth rates.

1.6 Results

In this section, we compute the treatment effects of the EU ETS regulation on French manufacturing firms' revenues and CO₂ emissions using our estimation procedure and compare to the DiD and LA methods.

³⁸We experiment with both probit and linear probability models and find that the functional form makes little difference in terms of selecting control firms.

³⁹Ideally, one would want to know the emissions threshold used to determine regulation, and then build a measure of concentration based on this threshold. Lacking data on these thresholds, we adopt a common (though fairly arbitrary) measure of concentration and test if the measure predicts regulation, conditional on $\alpha_{fis,2004}$.

1.6.1 Estimating Reduced-form Coefficients

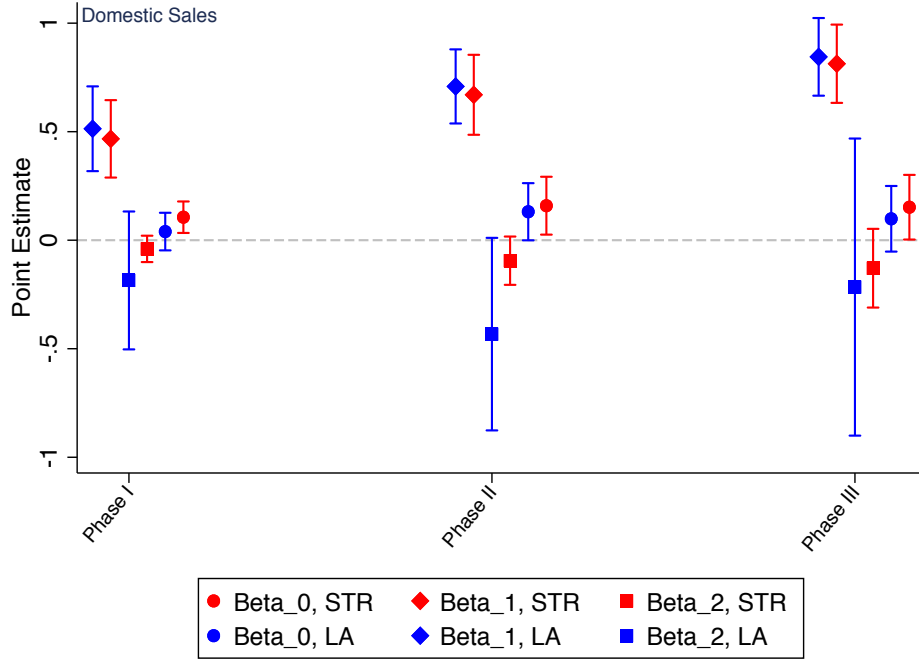


Figure 1.6.1: Reduced-Form Coefficients for the Effects on Domestic Revenues

Notes: Markers indicate estimates of β_0 , β_1 , and β_2 from our GMM estimator (in red) and from the LA model described in section 1.4.3 (in blue) for the matched sample. Bars indicate 95% confidence intervals. For each period, we estimate a pooled regression, combining all years included in the period. Top (bottom) row takes log change in domestic (total) sales as the dependent variable. Right (left) panels include all (matched) firms. The vertical line indicates the beginning of the regulation period. All regressions include industry-by-year fixed effects. Standard errors are clustered on the industry

Figure 1.6.1 presents the estimates of β_0 , β_1 , and β_2 from our GMM estimator (in red) and from the LA model described in section 1.4.3 (in blue) for the matched sample. For both estimators, the regulation variable R_{fst} is the labor share of regulated plants within the firm-year. For each phase of the regulation period, we pool all years in the period and estimate a single parameter value for each of β_0 , β_1 , and β_2 . We find that regulated firms grew faster than unregulated firms. With our estimator, we find that the point estimate for the direct effect β_0 is around 0.16, statistically significant at the 5% level for phases I and II but not for phase III, while the point estimate for the indirect effect β_2 is around -0.1 , with a 95% confidence interval that overlaps with zero in all phases. With the LA estimator, the magnitudes of the effects are similar, but the indirect effect is less precisely estimated.

Also in Figure 1.6.1, we find that β_1 is always positive, on the order of 0.5 and above, and quite precisely estimated (p-values $< .001$). The fact that we can reject $\beta_1 = 0$ indicates that the concentration of economic activity across plants within the firm indeed

affects costs, as predicted by our multi-plant framework. Moreover, with $\beta_1 > 0$, we can conclude that $\sigma < 0$, which indicates that increased concentration of economic activity within the firm increases costs: economies of scope dominate the increased cost from dispersed production.

The fact that regulated firms grew faster post-regulation is perhaps surprising, given that the regulation effectively increased the cost of energy for these firms. However, this is in fact a fairly robust finding in the EU ETS literature. For example, Löschel et al. (2018) Dechezleprêtre et al. (2018) and Calligaris et al. (2018) all find that regulated EU ETS firms grew faster post-regulation than unregulated firms. This evidence is consistent with the hypothesis that environmental regulation triggers cost-reducing technological investments (Porter & van der Linde, 1995; Ambec et al., 2013).⁴⁰

We test for parallel trends prior to 2004 by estimating the DiD model described in section 1.4.3, which controls for $\Delta\alpha_{fit}$.⁴¹ Whereas in the post-2004 years, we define R_{fist} as the labor share of regulated plants within the firm-year, in the years prior to 2004, we simply take a dummy variable for whether the firm will own at least one regulated plant post 2004. Figure A.1.1 shows that, using the full dataset (left panels), regulated firms grew faster in the pre-regulation periods and slower in the post-regulation periods than unregulated firms, conditional on industry-year fixed effects. Given the nonrandom regulation assignment, this pattern is precisely what one should expect. Regulation targeted large firms, and large firms tend to grow slower than small firms; hence, regulated firms grew slower in the post-regulation period. The fact that regulated firms grew faster in the pre-regulation period is also a function of endogenous selection: large firms grew fast prior to 2004, which is why they became regulated. These patterns indicate that productivity does not follow a random walk, and thus provide a strong motive for treating endogenous selection. Once we use the matched sample (right panels), we find that regulated and unregulated firms trended similarly prior to regulation.

In Figure 1.6.2, we present OLS estimates of equation (1.27) with the full (matched) sample indicated by black (white) diamonds. Using firm-by-year fixed effects absorbs any endogenous productivity growth, as well as the scale effect of the regulation, which explains that the estimates are similar across the two samples. The plant-level regulation status thus identifies just the price effect of the regulation. As in Figure 1.2.3, we find that emissions from regulated plants grew slower post-regulation compared to unregulated plants within a firm. The point estimates are imprecise for phases I and III, though for

⁴⁰(Calel & Dechezlepretre, 2016) and Calel (2020) find that the EU ETS induced R&D and green innovation in regulated firms, which could explain these positive effects on sales.

⁴¹As noted above, the estimates of this DiD model are closely related to our first two empirical moments in the GMM procedure.




Figure 1.6.2: Plant-level DD Estimates for Emissions

Notes: Figure presents estimates of β_z derived from OLS estimation of equation (1.27). Bars indicate 95% confidence intervals. For each period, we estimate a pooled regression, combining all years included in the period. All regressions include firm-by-year fixed effects and $\Delta\alpha_{jfst}$ as controls. Standard errors are clustered on the industry.

phase II, we can reject 0 at the 5% level.

1.6.2 Estimating Average and Aggregate Treatment Effects

Armed with our estimates of β_0 , β_1 , and β_2 , we then solve the system of equations (1.28) for counterfactual revenues under the scenario that $R_{jfst} = 0$ for all plants, and compute ATT and ATC according to (1.29). We also compute the ATT and ATC for the DiD and LA estimator following the formulas presented in section 1.4.3.

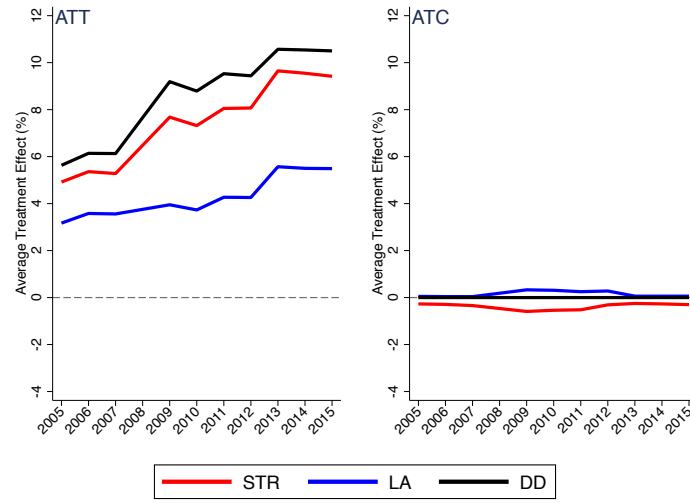
Figure 1.6.3 plots the average treatment effects by year on firm-level domestic revenues (panel a) and plant-level domestic emissions (panel b) for all three estimators. Using our estimator, we find that the EU ETS increased domestic sales for the average regulated firm between 5 and 9% annually relative to the unregulated counterfactual.⁴² The DiD estimator overstates this effect by 1-2 percentage points, and the LA estimator understates the effect by 2-4 percentage points, depending on the year. Whereas the average effect on revenues of unregulated firms is equal to zero by assumption for the DiD estimator, we

⁴²Treatment effects vary by year because of variation in values of β_0 , β_1 , and β_2 by phase, variation in industry-wide treatment density by year, as well as variation in $\Delta\omega_{fst}$, $\Delta\alpha_{fst}$, and $\Delta\lambda_{ist}$.


find that it was quite small and negative using our estimator, and slightly positive using the LA estimator.⁴³ Panel b of Figure 1.6.3 shows that, using both our estimator and the LA estimator, we find that the regulation reduced domestic emissions at regulated plants between 5 - 20% during phases I and II of the EU ETS, depending on the year and estimator. In phase III, we find a mild increase in emissions at regulated plants. By contrast, using the DiD estimator, we find little effect on emissions during phases I and II, and a fairly large increase in emissions in phase III (15%). For ATCT, we find that the DiD estimator yields a positive effect of regulation on the average unregulated plant within a firm that owns other regulated plants. These estimates are reminiscent of the “reallocation” effects estimated by Gibson (2019) and Soliman (2020). Compared to our structural estimator, we find that these DiD estimates vastly overstate the effect of regulation on these unregulated plants (i.e, an increase of 5-10% depending on the year for the DiD model compared to 0.5% for our estimator).

Finally, in Figure 1.6.4, we plot aggregate observed CO₂ emissions generated in the production of goods for the French market by year along with the counterfactual levels computed using our estimator and the LA and DiD estimators. In the top set of trajectories, we impute emission intensity for plants excluded from the EACEI surveys, while the bottom set of trajectories includes only plants observed in the EACEI survey. With imputing, we find that the regulation lowered domestic emissions between 1.5 - 3.5 mega tonnes annually through phases I and II, depending on the year, or between 4 - 10% of observed domestic emissions. Without imputing, we find that the regulation lowered domestic emissions between 1.4 - 2.7 mega tonnes annually through phases I and II, or between 5 - 12% of observed domestic emissions. In both cases, the vast majority of emission reductions came from regulated plants. In phase III, we estimate that the policy had no effect on aggregate emissions. By comparison, the LA estimator yields slightly larger treatment effects, though very similar to our preferred specification, while the DiD estimator yields zero effect during phases I and II and an effect of the opposite sign in phase III.

⁴³The fact that regulation increased domestic revenues for both regulated and unregulated firms according to the LA model is an artifact of averaging. For unregulated firms in particular, the sign of the effect of regulation varies by industry within a year. Averaged over all unregulated firms, we find slightly positive effects, but in aggregate, regulation increased domestic revenues for regulated firms as a group and reduced domestic revenues for unregulated firms as a group.



(a) ATEs on Domestic Revenues



(b) ATEs on Plant-Level Domestic Emissions

Figure 1.6.3: Average Treatment Effects on Domestic Revenues and Plant-Level Domestic Emissions

Notes: Figure presents the average treatment effects on firm-level domestic revenues (panel a) for treated firms (left) and control firms (right) and on plant-level domestic emissions (panel b) for treated plants (ATTT), control plants in treated firms (ATCT) and control plants in control firms (ATCC) as a percent of observed outcomes by year and estimator. Treatment effects on revenues are averaged over all French firms in the manufacturing sector while treatment effects on emissions are computed only for firms with observed emissions in the EACEI survey.

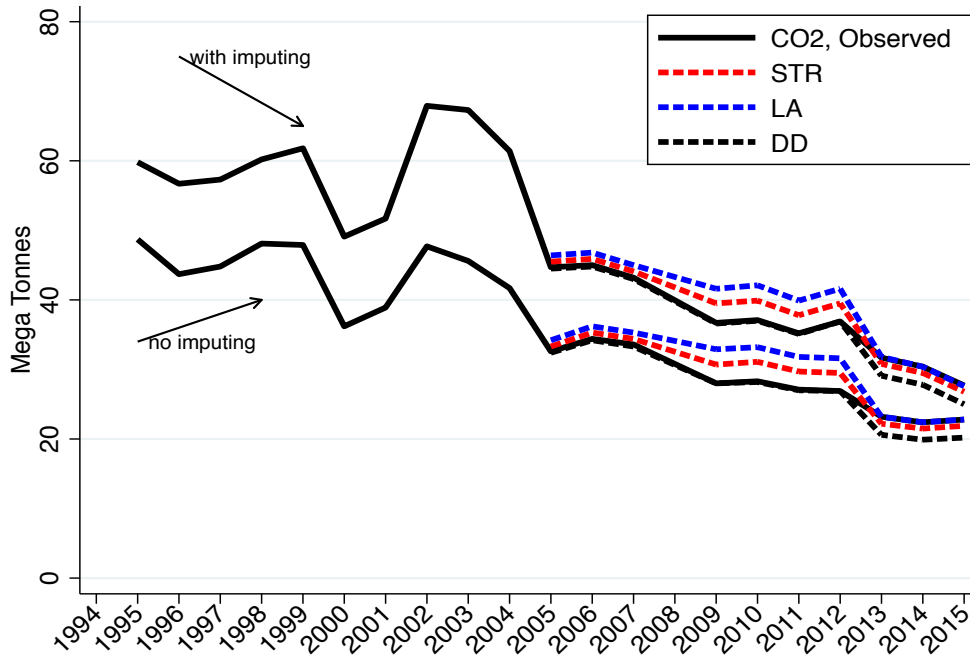


Figure 1.6.4: Treatment Effects on Aggregate Domestic Emissions

Notes: Figure presents observed total CO₂ emissions from French manufacturers generated in the production of goods for the domestic market (solid black) along with counterfactual emissions computed using STR, LA and DD methods. Bottom lines include only plants with observed emissions in the EACEI survey, while top lines impute emission intensity for plants with no emissions data.

1.7 Conclusion

The effects of industrial regulations are hard to predict for at least two reasons. First, firms may respond to regulation by investing in new technologies, which can affect costs in unanticipated ways. Second, changes in pricing induced by the regulation at regulated firms may affect the outcomes of other firms in the economy through competition on the output market. It is well known that these across-firm dependencies likely generate bias in difference-in-differences estimators of treatment effects. Controlling for industry-level regulation density goes a long way to addressing the problem, but if prices change substantially between pre-regulation and post-regulation periods, then local approximation models may lead to biased results as well.

We develop an estimation procedure to quantify the effects of industrial regulations in the presence of inter-firm and intra-firm spillovers. We extend the literature in several ways. First, we generalize the local approximation method to cases in which prices change substantially between pre and post regulation periods. Second, we incorporate multiplant production into a framework where regulation applies at the plant level. Third, we al-

low for nonrandom assignment of regulation and nonrandom walk productivity growth. Fourth, we consider heterogeneous effects of regulation by destination market.

We use our estimation procedure to study the effects of the EU ETS on firm-level and aggregate outcomes for French manufacturers. We find that despite fears of loss of competitiveness, the EU ETS did not cause regulated French manufacturing firms to lose market share. On the contrary, we find that the average regulated firm increased revenues earned on the French market by 5 to 9% annually as a result of regulation, consistent with a technology adoption mechanism à la Porter & van der Linde (1995). The standard DD estimator yields a higher estimate (by 1-2 percentage points), while a LA estimator yields a lower value (by 2-4 percentage points). Overall, we find that the EU ETS reduced aggregate emissions between 4 - 10% annually during phases I and II and had zero effect in phase III. The LA estimator yields a similar result, while the DD estimator predicts the opposite sign of the effect.

Our estimation procedure could be used to study other types of plant-specific shocks, such as place-based labor regulations, or regional productivity shocks. In our context, we were guided by the reduced-form evidence of across-firm and intra-firm spillovers, which led us to focus on competition forces and intra-firm reallocation forces. In a different context, researchers could potentially discover that local labor market reallocation forces are at play, for instance. The model could therefore be enriched to accommodate other spillover mechanisms.

Chapter 2

An Optimal Distribution of Polluting Activities Across Space

2.1 Introduction

It is now widely acknowledged that atmospheric pollution causes substantial damage to human health.¹ Yet, most people live in large cities that also concentrate polluting industries. Recognizing this, governments have adopted air quality regulations that are more stringent in large and populated cities. However, more stringent regulations can hurt local firms, make them less productive, and, in turn, affect local workers' income.² This local tradeoff between income and pollution raises the question of whether environmental regulations should be different across heterogeneous locations. This question requires a theory of the location choice of workers and heterogeneous firms that endogenizes local air pollution as an externality from production. In this paper, I develop such a theory and explore what would be an optimal distribution of polluting activities across space in the context of France.

I first use this framework to find the set of policies that implements an optimal allocation of workers, firms, and pollution across cities. Within this framework, location choices made by individual firms and workers depend on (1) the level of amenities offered by cities, (2) agglomeration externalities that provide a productivity advantage to firms located in more populated areas, and (3) congestion forces. Local air pollution acts as a congestion force by decreasing city-specific amenities. Workers respond to bad air quality by moving away from it. Firms, by contrast, only internalize the negative externalities associated with their polluting emissions insofar as they have to pay a local city-specific tax on emissions. They do not internalize their impact on local labor supply that responds to the decrease in amenities associated with bad air quality. Under these assumptions, I show that the largest cities, endowed with the highest level of exogenous amenities and productivities, are too small compared to the optimum when the emission tax is uniform across space.

¹Health damage have been measured even when pollution levels fall below regulatory standards (Graff Zivin & Neidell, 2012, 2013; Deryugina et al., 2019).

²This is the “jobs versus the environment” argument (Morgenstern et al., 2002).

An optimal set of city-specific pollution taxes should include higher taxes in large cities to make industrial production cleaner, air quality better, and these cities more attractive to workers.³ Second, I explore quantitatively what would be this optimal set of city-specific pollution taxes in the context of France and compare it to the estimated current regulations on industrial emissions. I find that current regulations are differentiated across space in the right direction – they are more stringent in larger cities – but welfare could increase by raising pollution taxes in larger cities even more.

I consider a unique country with a discrete number of cities that differ in their idiosyncratic and exogenous endowments in local amenities and industrial labor productivity. Workers derive utility from local amenities and from the consumption of an industrial tradable composite good. Following Rosen (1979)-Roback (1982), I assume that a fixed population of homogenous industrial workers choose to locate in different cities in equilibrium. In each city, continuums of heterogenous firms from distinct industrial sectors produce differentiated varieties that aggregate into the composite good. As in Copeland & Taylor (2004), firms use labor and polluting emissions as substitutable inputs for production.⁴ Production costs thus depend on local wages, local labor productivity, and local emission taxes. Furthermore, firms face iceberg trade costs when exporting their varieties to other cities. As is standard in economic geography models, the model includes two externalities related to the allocation of workers across space. First, local labor productivity is endogenous and reflects agglomeration economies: when the number of worker increases in a city, so does labor productivity. Second, local welfare is influenced by general congestion effects: as local population increases, congestion appears on the housing market and in transportation, which reduces each worker’s welfare.⁵ Furthermore, workers are assumed to move freely across cities.⁶

The novelty of the paper rests on the assumption that the total quantity of industrial emissions per city negatively affects city-specific amenity level, using a constant elasticity environmental damage function. Doing so, I consider polluting emissions as a congestion force that affects workers’ location decision, and firms do not internalize this effect.⁷ In

³To be able to compare sets of city-specific pollution taxes, I keep the average tax level across cities constant across alternative sets. In doing so, I focus only on the effects of spatially differentiated emission taxes, not on the welfare impacts of raising tax levels.

⁴Pollution can equivalently be modelled as an input or a by-product of industrial production in the standard Copeland & Taylor (2004)’s setup.

⁵This congestion effect potentially encompasses polluting emissions related to residential energy consumption and transportation. Even though, I do not explicitly model these pollution sources, they affect welfare negatively. I assume that these sectors are not covered by the emission tax on industrial pollution.

⁶The free mobility assumption leads to an equilibrium that is informative in the long run. In Appendix B.4, I investigate the alternative assumption of no across-city migration, which is more informative in the short run. Using the simplified setup of Section 2.2, I find that the main result holds when the central planner maximizes the average per capita welfare of workers.

⁷This is an effect that has been identified empirically (Chay & Greenstone, 2005; Davis, 2011; Khanna

the model, I consider for simplicity a representative pollutant that serves as an indicator of the combination of harmful pollutants emitted by industrial activities (PM10, PM2.5 and SO2). I further assume that this pollutant is essentially local in the sense that emissions only affect local air quality within the city limits.⁸ I show that a central planner taking into account the trade-off between agglomeration externalities and congestion forces decides to differentiate local emission taxes by setting the relative level of tax in larger cities higher.

I apply this spatial equilibrium framework to the specific case of France. This empirical setting is interesting for three reasons. First, a large share of harmful pollutant emissions result from industrial activities in France.⁹ Second, local air concentrations of pollutants in several French cities often reach levels well above World Health Organization (WHO)’s guidelines.¹⁰ Third, air quality regulations implemented in the last decades in France tend to be more stringent in larger cities, which are also more polluted and concentrate more industrial polluting activities. Figure 2.1.1 illustrates salient stylized facts that I incorporated into the model. The upper panel shows that the larger the number of workers within a French commuting zone, the higher the mean PM 2.5 concentration in 2012. The middle panel shows that the most populated commuting zones are also the most polluting in terms of PM 2.5 emitted by industrial activities. The lower panel shows that the local marginal cost of emitting PM2.5 for industrial firms tends to increase with the size of the local labor market.¹¹

Using an extensive set of city-level, firm-level, and plant-level data, I first estimate the parameters of the model. I estimate key elasticities that determine the strength of pollution externalities, agglomeration economies, and general congestion effects. I also provide novel estimates of sector-specific elasticities governing the substitution of emissions and labor in industrial production, using an instrumental variable approach. Second, I retrieve the distributions of amenities, productivities, and emission taxes across cities from the observed endogenous distributions of populations, wages, and emissions. Doing so, I obtain the set of current city-specific marginal costs of emitting atmospheric pollutants

et al., 2021), but not yet included in a quantitative setup.

⁸In reality, air pollutants can potentially travel across space depending on climatic conditions (especially, wind), some pollutants more than others. Since I consider commuting zones instead of cities in the empirical application, the assumption is however quite realistic since most pollutants do not travel very far. Furthermore, introducing an atmospheric circulation extension is outside of this paper’s scope.

⁹In 2012, industrial activities (fossil fuel combustion, chemical reactions, waste treatment or other manufacturing processes) emitted around 40% of PM2.5 and SO2 national emissions and around 25% of PM10 national emissions (see Figure B.1.1 and Tables B.1.1, B.1.2, and B.1.3 for details).

¹⁰Figure B.1.2 shows that, in 2012, both mean and maximum air concentration measurements from ground monitors over the country for these substances have remained above WHO guidelines.

¹¹Assuming that homogeneous firms in different cities have access to the same Cobb-Douglas production function between labor and PM2.5 emissions (assumptions used in the illustrative model in Section 2.2), I compute the marginal cost of PM2.5 emissions as the ratio of total wage payments over city-level PM2.5 emissions.

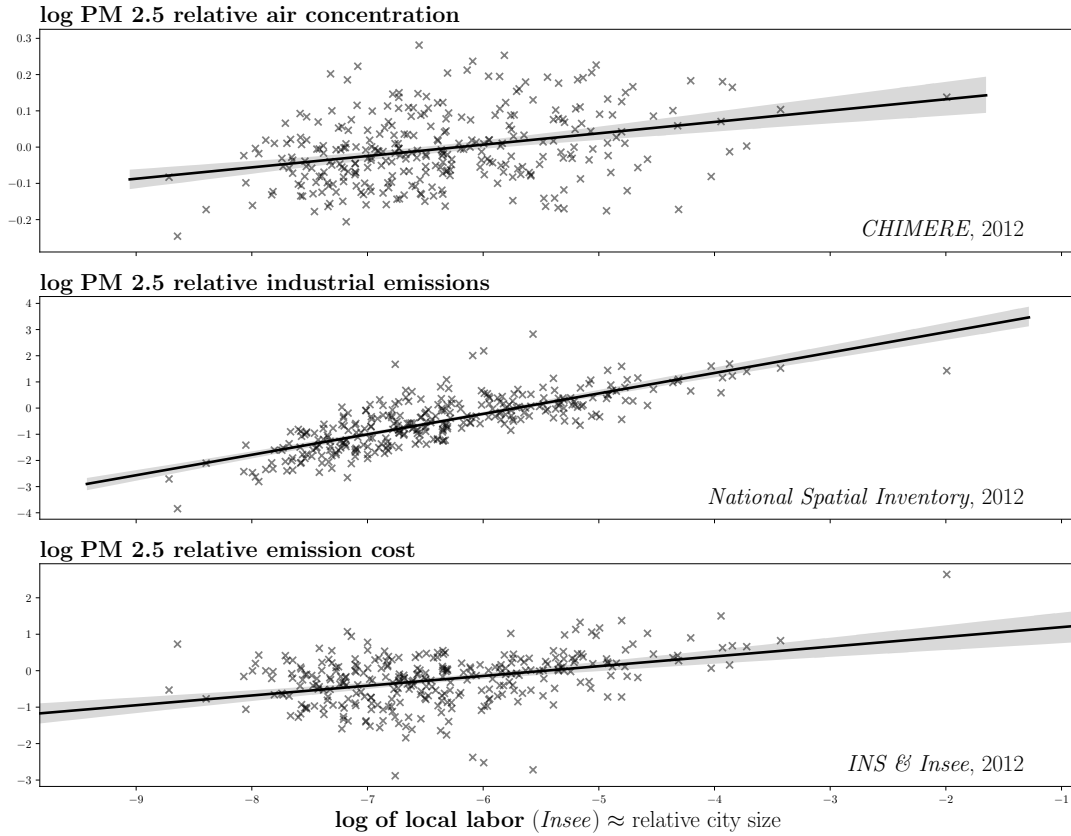


Figure 2.1.1: PM2.5 concentrations, emissions and marginal costs across commuting zones
Note: The upper panel plots the 2012 mean air concentration in PM2.5 at the commuting zone level as a function of the local number of workers. The middle panel plots the 2012 PM2.5 emissions from industrial activities at the commuting zone level as a function of the local number of workers. The lower panel plots the implicit marginal cost of emitting PM2.5 in the industry as a function of the local number of workers. This cost is the ratio of total wage payments in the commuting zone over the quantity of PM2.5 emitted by the industry.
City-level PM2.5 air concentration is from Chimere. PM 2.5 industrial emissions are from the National Spatialized Inventory. Labor data comes from the Insee and is a count of the number of workers employed in the employment area. All values are for 2012.

for industrial firms implied by current French air pollution regulations. I find that these costs are higher in large cities than in small cities. Specifically, these costs are higher in more productive cities, but not necessarily in cities endowed with more amenities.

I then characterize the welfare implications of such spatial distribution of marginal costs of emissions. I show that more stringent emission regulations in large cities enabled them to reduce pollution emissions and become larger than what they would have been under a counterfactual uniform regulation. This implies higher levels of welfare for workers living in all cities. To go further, I then identify the spatial distribution of emission taxes that would maximize workers' welfare. It would be optimal to raise emission taxes in cities with high amenities, leading to further concentration of workers in these cities, and thus exacerbating the uneven distribution of activities across space.

This paper relates to several literatures. First, it contributes to the recent quantitative

spatial economics literature on the distribution of economic activities across space (Allen & Arkolakis, 2014; Redding & Rossi-Hansberg, 2017; Redding, 2020; Allen et al., 2020). Extending the framework of Rosen (1979)-Roback (1982), this literature can be decomposed into two main strands. The first one assumes that workers have heterogeneous preferences over amenities and investigates mainly the location decision of workers across cities (Moretti, 2013; Diamond, 2016; Almagro & Domínguez-Iino, 2021). An influential paper in this literature is Diamond (2016) that endogenizes local amenities to explain the spatial distribution of skilled and unskilled workers across U.S. cities. Diamond assumes that these amenities are positively influenced by the fraction of skilled over unskilled workers. In my paper, I also introduce endogenous amenities affected by congestion forces, based on total population levels, and by the pollution externality. The second strand considers homogenous workers to explore the location decision of firms across cities and the optimal size of cities (Henderson, 1974; Eeckhout & Guner, 2015; Borck & Tabuchi, 2018; Gaubert, 2018; Albouy et al., 2019). Focusing on the supply side of the economy, namely polluting industries, I make the same assumption. In this strand, Allen & Arkolakis (2014) and Allen et al. (2020) develop a spatial quantitative model where characteristics of heterogeneous locations determine the equilibrium distribution of economic activities across space. They assume endogenous amenities and agglomeration economies.¹² Even though my model is an extension of these two papers, I depart from them by focusing on a particular source of congestion externalities, namely local air pollution, and investigate its effect on the spatial distribution of economic activities.

My paper is connected to a second literature that investigates the distribution of pollution across cities (Glaeser & Kahn, 2010; Carozzi & Roth, 2019; Colmer et al., 2020a; Borck & Schrauth, 2021; Eeckhout & Hedtrich, 2021) and its congestion effects (Drut & Mahieux, 2015; Leturque & Sanch-Maritan, 2019; Hanlon, 2019). Carozzi & Roth (2019) and Borck & Schrauth (2021) both show that denser cities are also more polluted (respectively for the US and Germany). Similarly, Colmer et al. (2020a) uncover large disparities in the spatial distribution of PM2.5 concentration across census tracts in the US. In particular they find that increasing local population leads to higher levels of ambient PM2.5 pollution. They also identify a decrease of pollution in locations where regulation became more stringent (namely census tracts in non-attainment with air quality standards under the Clean Air

¹²In this respect, I include the “productivity advantage of large cities” that have been identified for France by Combes et al. (2012) and more recently by Gaubert (2018). However, I do not make any assumption on the source of these agglomeration economies. Notably, the model does not include any sorting mechanism that may partly explain higher productivities observed in larger cities (Baldwin & Okubo, 2005; Combes et al., 2008). Another source of agglomeration comes from the costly trade assumption. Recently Bartelme (2018) showed that trade costs between US cities explained a large fraction of the spatial distribution of economic output. At the same time, my model embodies general counteracting congestion forces which limit the scale of economic concentration.

Act). In my paper, I provide a framework explaining such spatial distribution of pollution: my setup encompasses the effect of local population, local income, and local regulations on local emissions. Moreover, with a focus on France, Drut & Mahieux (2015) and Leturque & Sanch-Maritan (2019) exposed that agglomeration gains were dampened by higher levels of local air pollution. My framework includes this effect by assuming that local pollution actually acts as a congestion force. This mechanism is supported by findings from Hanlon (2019) who showed for the UK that local pollution industrial emissions reduced long-run city employment and population growth. Damage from pollution also have been found to be heterogenous across space depending on local characteristics (Aldeco et al., 2019; Deryugina et al., 2021; Desmet et al., 2021; Alvarez & Rossi-Hansberg, 2021). The assumption of constant elasticity of damage from pollution implicitly acknowledges this heterogeneity to the extent that large cities concentrate most of the welfare loss due to pollution.

Finally, my paper contributes to the literature on optimal place-based policies chosen by a central planner. In particular, Fajgelbaum & Gaubert (2020) investigate how income transfers across cities could maximize welfare by correcting spatial externalities, assuming free mobility of workers across cities. Similarly, Suarez-Serrato & Zidar (2016) consider a tax rate on firms' profits that varies across cities. In contrast to these papers, I assume an input-specific policy instrument that induces reallocation across inputs, therefore it affects the endogenous amenity level per city through emissions. This place-based literature also contains studies including endogenous pollution (Lange & Quaas, 2007; Kyriakopoulou & Xepapadeas, 2013; Yamada, 2020; Pflüger, 2021). A paper closely related to mine is Yamada (2020) that also considers atmospheric pollutants as a congestion force and shows that imposing more stringent air quality regulations in a specific set of large Chinese cities could lead to welfare gains. In contrast to this paper, I solve for the optimal distribution of local emission regulations instead of comparing ad-hoc distributions of policies inspired by planned policy projects.

The remainder of the paper is organized as follows. In the next section, I provide a simplified framework to illustrate how introducing a local pollution externality into the standard spatial equilibrium model generates new insights. In the third section I provide some context on air quality regulations in France. A more realistic general model is detailed in the fourth section. The fifth section presents the estimation of the model's parameters. Finally, the last section contains results from numerical welfare optimization problems and discusses the results.

2.2 Optimal Environmental Policy in a Simple Spatial Model

Following the Rosen (1979)-Roback (1982) framework, I consider a given set of C cities in which a fixed population of workers can live. Throughout the paper, I assume homogenous workers.¹³ The per capita utility of the representative worker in city j is given by:

$$u_j = a_j Z_j^{-\gamma} L_j^{-\delta} c_j, \quad (2.1)$$

where a_j is the local idiosyncratic endowment in amenities, L_j is the local population of workers, Z_j is the quantity of atmospheric pollutants emitted by the local industry, and c_j is the local per capita consumption of a tradable good. This utility function captures the fact that workers value the consumption of the tradable good, which depends on their income, as well as other characteristics of the locations they live in. In this paper, I focus on the local detrimental welfare effect of industrial emissions of atmospheric pollutants, as industrial activities are major emitters of several harmful pollutants.¹⁴ In equation (2.1), a positive γ implies that local industrial emissions Z_j negatively affect the local air quality and therefore welfare (through health damages). For simplicity, I abstract from spatial pollution spillovers and assume that local emissions only affect local welfare. Workers also bear various agglomeration costs, which are accounted for by the general congestion term $L_j^{-\delta}$. This captures local externalities from agglomeration, such as commuting costs and housing prices.¹⁵ This general congestion term also captures the detrimental effect on welfare of atmospheric pollution emitted by non-industrial activities (transport, residential heating, or energy production). Furthermore, each city is exogenously endowed with a fixed level of amenities a_j that explains the location choices of local workers once spatial differences in consumption and local externalities are accounted for. For instance, amenities include geographical environmental attributes, available space, institutional and social installations or any local external factor that explains why some locations are more attractive than others. Workers benefit from agglomeration economies. I assume that labor productivity in location j is equal to $b_j L_j^\nu$. Elasticity ν governs the strength of

¹³As shown in Allen & Arkolakis (2014), this choice for workers' preferences is isomorphic to a model where heterogeneous workers have idiosyncratic utility shocks, according to a Fréchet distribution, that are independent and identically distributed across locations and individuals. Such heterogeneity across workers is a dispersion force and is captured by the parameter δ in my model.

¹⁴In 2012, the industry emitted 40%, 42% and 23% of total PM2.5, SO2 and PM10 emissions, see Appendix B.1.

¹⁵Allen & Arkolakis (2014) show that constant elasticity congestion costs are isomorphic to the Helpman (1998)-Redding (2016) setup where workers spend a fixed share of their income on a non-tradable good, which can include housing.

agglomeration economies and b_j allows for idiosyncratic labor productivity differences between cities. Idiosyncratic productivity can vary across cities due to a wide array of local characteristics (better local institutions or more efficient local transport networks). Among other things, agglomeration economies can arise from local knowledge spillovers, labor markets pooling, or local economies of scale.¹⁶

In each city, identical firms produce a homogenous tradable good using labor and emissions of a representative atmospheric pollutant with a Cobb-Douglas production function defined by an expenditure share α , with $1 > \alpha > 0$. Following Copeland & Taylor (2004), it is equivalent to assuming that emissions are a by-product of production rather than an input in production. Then, firms can divert a fraction of their labor force to abate emissions. The efficiency of abatement is governed by α : the lower it is, the more efficient the abatement technology is to reduce firms' emission intensity. In location j , a unit of labor costs the local wage w_j and a unit of emissions costs the local emission tax t_j . This emission tax is a policy instrument set by a central planner. It represents the local pollution regulations (including technology standards, emission limits, or emergency responses). I assume that proceeds from the emission tax are locally redistributed to workers. These combined assumptions lead to the equilibrium distribution of industrial emissions across cities (see Appendix B.3 for computing details):

$$Z_j = \alpha \frac{w_j}{t_j} L_j, \quad (2.2)$$

which illustrates that industries emit more pollution in cities that are larger, and where wages are higher. Conversely, higher emission taxes reduce local emissions.

I assume that competition is perfect and trade is costless, which implies the spatial equalization of the output price and leads to the equilibrium distribution of local wages. Setting the average wage as the numeraire would add general equilibrium effects to the equilibrium: when regulation changes in a given city there would be spillovers on wages in other cities. I abstract from such effects to focus only on the interplay between agglomeration, congestion, and industrial pollution externalities. Therefore, I normalize the output price to 1. As a result, per capita consumption of the tradable good is equal to the local wage: $c_j = w_j$. Output price normalization also pins down local wages in location j :

$$w_j = b_j L_j^\nu t_j^{-\frac{\alpha}{1-\alpha}}, \quad (2.3)$$

which implies that wages are higher in more productive cities as well as in larger cities

¹⁶Duranton & Puga (2004) provides a large range of models that deliver this constant elasticity function. Combes & Gobillon (2015) provides a survey of the empirical literature that documents agglomeration economies.

because of agglomeration economies. It also implies that raising local emission taxes has a negative effect on local wages.

Substituting for (2.2) and (2.3) in (2.1), the indirect utility in location j can be written as:

$$u_j = a_j b_j^{1-\gamma} L_j^{-\theta} t_j^{\frac{\gamma-\alpha}{1-\alpha}}, \quad \text{with } \theta = \delta + \gamma - \nu(1 - \gamma). \quad (2.4)$$

The exponent θ on the population term in equation (2.4) reflects the fact that, since $\gamma > 0$, the emission externality both reinforces congestion and weakens agglomeration economies (when cities grow, agglomeration economies increase local wages and firms become more pollution intensive all else held equal).

Finally, I assume free mobility of workers across locations.¹⁷ As a result, utility is equalized across space in equilibrium and is equal to a level denoted \bar{u} . Indeed, if workers could reach a higher level of welfare in a given city, they would move there until the marginal gain of moving was fully compensated by the marginal cost of congestion in this city. Combining this assumption with (2.4), we obtain, for two cities i and j :

$$\frac{L_j}{L_i} = \left(\frac{a_j}{a_i} \right)^{\frac{1}{\theta}} \left(\frac{b_j}{b_i} \right)^{\frac{1-\gamma}{\theta}} \left(\frac{t_j}{t_i} \right)^{\frac{1}{\theta} \frac{\gamma-\alpha}{1-\alpha}}, \quad (2.5)$$

which implies that the distribution of workers across cities is a function of the distributions of amenities, productivities, and emission taxes. Depending on the relative strengths of the three externalities, local populations may be positively or negatively, correlated with these local characteristics. Equation (2.5) also illustrates that there exists a unique spatial equilibrium if and only if $\theta \neq 0$. This condition can be expressed as a condition on the elasticity of the pollution externality: $\gamma \neq -\frac{\delta-\nu}{1+\nu}$. Assuming that pollution is a congestion force (i.e. $\gamma > 0$), this condition ensures that the pollution externality does not exactly offset the combined effect of general congestions effect and agglomeration economies. When congestion effects strictly outweigh agglomeration economies, this condition always holds.

Without any loss of generality, I normalize the total population to 1, so that $\sum_{j \in C} L_j = 1$, and compute the common level of welfare \bar{u} reached in equilibrium:

$$\bar{u} = \left[\sum_{j \in C} a_j^{\frac{1}{\theta}} b_j^{\frac{1-\gamma}{\theta}} t_j^{\frac{1}{\theta} \frac{\gamma-\alpha}{1-\alpha}} \right]^{\theta} \quad (2.6)$$

¹⁷This assumption is informative of long term equilibria, where within-country migration costs can be considered as low. In Appendix B.4, I consider for robustness the alternative assumption, where local populations are fixed and do not adjust in equilibrium.

The optimal policy for a central planner is to maximize \bar{u} by adjusting the set of local emission taxes across cities. In Appendix B.3, I show that $1 \geq \gamma \geq \alpha$ and $1 \geq \frac{1}{\theta} \frac{\gamma - \alpha}{1 - \alpha}$ are necessary and sufficient conditions for \bar{u} to be concave and the optimization problem to have a unique solution. First, $1 \geq \gamma$ means that the positive direct effect of productivity on wages outweighs its indirect negative effect that causes firms to become more pollution intensive when wages increase. Second, $\alpha < \gamma$ means that the negative effect on wages of raising the local emission tax is weaker than the positive effect it has on air quality. Third, $1 \geq \frac{1}{\theta} \frac{\gamma - \alpha}{1 - \alpha}$ ensures that the whole population is not concentrated in a unique city. The welfare function specified by equation (2.6) is homogeneous of degree $\frac{\gamma - \alpha}{1 - \alpha}$ with respect to the set of emission taxes. This implies that, if I multiply all city level emission taxes by a common factor, workers' welfare \bar{u} is multiplied by a power $\frac{\gamma - \alpha}{1 - \alpha}$ of this factor. Therefore, I normalize the average emission tax \bar{t} to 1 and focus on the distribution of taxes across cities.

Proposition. *Consider a set of cities C , with exogenous amenities $\{a_j\}_{j \in C}$ and productivities $\{b_j\}_{j \in C}$, with relative populations specified in equation (2.5). If $1 \geq \gamma > \alpha$ and $1 > \frac{1}{\theta} \frac{\gamma - \alpha}{1 - \alpha}$, there is a unique set of emission taxes $\{t_j^*\}$ that maximizes the welfare function specified in equation (2.6) under the constraint that $\bar{t} = 1$ and it is defined by:*

$$\frac{t_j^*}{t_i^*} = \left[\left(\frac{a_j}{a_i} \right)^{\frac{1}{\theta}} \left(\frac{b_j}{b_i} \right)^{\frac{1 - \gamma}{\theta}} \right]^{\frac{1}{1 - \frac{1}{\theta} \frac{\gamma - \alpha}{1 - \alpha}}}. \quad (2.7)$$

Equation (2.7) indicates that the relationship between optimal emission taxes and local characteristics depends on the sign of θ . In particular, $\theta > 0$ means that congestion forces outweigh agglomeration economies and more workers locate in cities with a higher level of amenities and productivity (see equation (2.5)). In this case, equation (2.7) indicates that, to maximize workers' welfare, the central planner imposes higher emission taxes in cities with a higher level of amenities and productivity.

An important implication of my model is that emission taxes should be heterogeneous across space. However, national governments usually enforce regulations that are spatially uniform. Therefore, it is interesting to compare the implications of enforcing the set of emission taxes described by (2.7) to the implications of enforcing a uniform set of emission taxes (each equal to 1). Denoting L_j^* and L_j^u the populations of city j under the optimal and uniform sets of emission taxes, equations (2.5) and (2.7) imply:

$$\frac{t_j^*}{t_i^*} = \frac{L_j^*}{L_i^*} = \left(\frac{L_j^u}{L_i^u} \right)^{\frac{1}{1 - \frac{1}{\theta} \frac{\gamma - \alpha}{1 - \alpha}}}. \quad (2.8)$$

which means that if an optimal set of emission taxes exists, taxes should be higher in cities that are larger when the tax is uniform. Moreover, the sufficient condition for an optimal set of taxes to exist implies that $\frac{1}{1-\frac{1}{\theta}\frac{\gamma-\alpha}{1-\alpha}} > 1$. As a result, equation (2.8) indicates that implementing the optimal set of emission taxes reinforces the spatial concentration pattern observed under uniform taxes. When congestion forces outweigh agglomeration forces and higher labor productivity translates in higher income, it is optimal that more workers locate in cities with high amenities and productivity endowments (to compensate the costs of agglomeration). However, firms do not fully internalize the detrimental impact of their emissions on local welfare and pollute too much in cities that could attract even more workers if they were less polluted. This is why the optimal set of emission taxes imposes a relatively higher marginal cost of polluting in these cities, compared to cities where it is not optimal to concentrate workers. As a result, firms in these cities are less emission intensive, which lowers the level of emissions and makes these cities more attractive to workers.

The set of emission taxes described by equation (2.7) corrects all externalities, even the ones that are not caused by industrial pollution emissions. To analyze the case where the central planner corrects for these “non-emission” externalities using external policy instruments, one can derive the optimal set of emission taxes when the emission externality is the only externality. Assuming $\delta = \nu = 0$ in equation (2.7), the optimal set of emission taxes still imposes higher emission costs in the same cities as in the case where all externalities are corrected by the set of emission taxes.

Finally, equation (2.7) reveals that if there were no industrial emissions externality, that is to say if $\gamma = 0$, the optimal set of emission taxes would still be non-uniform, to correct for congestion and agglomeration externalities. In particular, when $\gamma = 0$, equation (2.7) leads back to the standard result arguing that large cities (with the highest productivities and amenities) should be even larger (Eeckhout & Guner, 2015; Gaubert, 2018; Albouy et al., 2019). Moreover, the average tax normalization imposes that increasing emission taxes in large cities implies decreasing them in smaller cities. Equation (2.3) shows that it is equivalent to decreasing wages in large cities and increasing them in small cities. This is similar to the system of optimal income transfers identified by Fajgelbaum & Gaubert (2020) when congestion costs outweigh agglomeration economies.

2.3 General Spatial Model

In this section I extend the model to (i) an industry composed of several polluting sectors, (ii) continuums of heterogeneous firms that compete monopolistically over differentiated varieties, (iii) costly trade between cities, and (iv) general equilibrium effects from local

changes in emission taxes.

2.3.1 Setup

The per capita utility of the representative worker in city j follows equation (2.1). The industry is now composed of S distinct sectors and c_j is the industrial composite:

$$c_j = \prod_{s \in S} \left(\sum_{i \in C} \int_{\omega \in \Omega_{is}} c_{ijs}(\omega)^{\frac{\sigma_s - 1}{\sigma_s}} d\omega \right)^{\frac{\sigma_s}{\sigma_s - 1} \beta_s}, \quad (2.9)$$

with $c_{ijs}(\omega)$ the quantity of variety ω produced in city i and consumed in city j , Ω_{is} the continuum of varieties produced in sector s in city i , σ_s the elasticity of substitution between varieties in sector $s \in S$, and β_s the share of income spent on varieties from sector s by each worker (with $\sum_{s \in S} \beta_s = 1$). These parameters are assumed to be the same across cities.

The aggregated price index of the industrial good in city j , P_j is such that:

$$P_j = \prod_{s \in S} P_{js}^{\beta_s}, \text{ and } P_{js} = \left(\sum_{i \in C} \int_{\omega \in \Omega_{is}} p_{ijs}(\omega)^{1 - \sigma_s} d\omega \right)^{\frac{1}{1 - \sigma_s}}, \quad (2.10)$$

with P_{js} city and sector specific price indices, and $p_{ijs}(\omega)$ the unit price of variety ω produced in city i in sector s and consumed in city j . I assume that there are no friction on local labor markets, so that that wages are equal across sectors. As a result, per capita consumption of the tradable good is given by $c_j = \frac{w_j}{P_j}$ and workers' indirect utility is:

$$u_j = a_j Z_j^{-\gamma} L_j^{-\delta} \frac{w_j}{P_j}. \quad (2.11)$$

In each city j there is an infinite supply of entrepreneurs in each sector s that can choose to pay a fixed sector specific entry cost f_s^e to draw a productivity ϕ from a Pareto distribution G_{js} :

$$G_{js}(\phi) = 1 - \left(\frac{\phi}{b_{js}} \right)^{-\theta_s}, \quad (2.12)$$

with $b_{js} = b_j b_s L_j^{\nu_s}$. These assumptions extend section 2.2's assumptions of idiosyncratic labor productivity and agglomeration economies to the case of multiple sectors and heterogeneous firms. In particular, the elasticities of agglomeration economies, $\{\nu_s\}_{s \in S}$, are sector specific. In addition, idiosyncratic sector and city specific labor productivities $\{b_{js}\}_{(j,s) \in C \times S}$ are multiplicatively separable between idiosyncratic city specific productivities $\{b_i\}_{i \in C}$, that are common across sectors, and sector specific productivities $\{b_s\}_{s \in S}$,

that are common across cities.

Each firm produces a specific variety of the industrial good using labor following the function:

$$q_{ijs}(\phi) = (1 - a(\phi)) \phi l_{ijs}(\phi) \quad (2.13)$$

where $q_{ijs}(\phi)$ is the quantity produced by a firm of productivity ϕ in city i and sold on city j 's market, and $l_{ijs}(\phi)$ is the quantity of labor used for production. Variable $a(\phi)$ corresponds to the fraction of the firm's labor force used to abate emissions caused by production. Production releases pollution as a by-product according to:

$$z_{ijs}(\phi) = (1 - a(\phi))^{\frac{1}{\alpha_s}} \phi l_{ijs}(\phi) \quad (2.14)$$

where $z_{ijs}(\phi)$ is the quantity of the representative pollutant emitted by a firm of productivity ϕ in city i for the production sold in city j . $\{\alpha_s\}_{s \in S}$ are the sector specific expenditure shares on emissions. Heterogeneity in these parameters across cities accounts for the fact that some industrial sectors may be more emission intensive than others. I assume that, in city j , emissions are taxed by the central planner at a rate t_j , which is the same across sectors. Local proceeds from this tax are redistributed to local workers.

These production and pollution functions closely follows Copeland & Taylor (2004) and Shapiro & Walker (2018) and are standard in the literature. Combining equations (2.13) and (2.14) implies the following:

$$q_{ijs}(\phi) = z_{ijs}(\phi)^{\alpha_s} (\phi l_{ijs}(\phi))^{1-\alpha_s}, \quad (2.15)$$

which is a Cobb-Douglas function that combines two inputs: emissions and labor. As a result, total industrial emissions of the representative pollutant, in city i , are given by:

$$Z_i = \sum_{s \in S} \sum_{j \in C} \int_{\phi} z_{ijs}(\phi) dG_{is}(\phi) \quad (2.16)$$

Trade is costly between cities. To export to city j , firms in city i have to pay an origin-destination specific iceberg cost τ_{ij} . As a result, any firm sells to all cities.¹⁸ Combining the assumption of monopolistic competition with equations (2.9) and (2.15) implies that the unit price of a variety produced in city i in sector s by a firm with productivity ϕ and

¹⁸The model could be extended to accommodate for origin-destination fixed trade costs as in Shapiro & Walker (2018). Productivity distributions of firms would then be left-truncated by endogenous zero-profit productivity cutoffs. In this case, only the most productive firms in a given city would be selling goods in all other cities.

delivered in city j follows:

$$p_{ijs}(\phi) = \frac{\sigma_s}{\sigma_s - 1} \frac{\tau_{ij} c_{is}}{\phi^{1-\alpha_s}}, \text{ with } c_{is} = \kappa_{\alpha_s} t_i^{\alpha_s} w_i^{1-\alpha_s}, \quad (2.17)$$

and $\kappa_{\alpha_s} = \alpha_s^{-\alpha_s} (1 - \alpha_s)^{-1+\alpha_s}$. Revenues and profits of a firm with productivity ϕ in city i respectively follow:

$$r_{is}(\phi) = \sum_{j \in C} p_{ijs}(\phi) q_{ijs}(\phi), \text{ and } \pi_{is}(\phi) = r_{is}(\phi) / \sigma_s. \quad (2.18)$$

I assume that entrepreneurs enter production until their expected profits equal the fixed entry cost of drawing a productivity from the local Pareto productivity distribution. I express fixed entry costs in the aggregate factor price, as in Bernard et al. (2007). Thus, for each pair of city j and sector s , I have the free entry condition:

$$\int_{\phi} \pi_{js}(\phi) dG_{js}(\phi) = c_{js} f_s^e. \quad (2.19)$$

In equilibrium, goods markets clear for each city and sector pair so that workers' expenditures are equal to firms' sales. These goods markets clearing conditions can be written for each pair of city j and sector s as:

$$P_{js}^{1-\sigma_s} = K_s \sum_{i \in C} \tau_{ij}^{1-\sigma_s} \tilde{M}_{is} (t_i^{\alpha_s} w_i^{1-\alpha_s})^{-\sigma_s} b_i^{(1-\alpha_s)(\sigma_s-1)} L_i^{\nu_s(1-\alpha_s)(\sigma_s-1)}, \quad (2.20)$$

with $\tilde{M}_{is} = \frac{\sigma_s}{\beta_s} f_s^e M_{is} c_{is}$, and M_{is} the mass of firms that produce in sector s and city i in equilibrium.¹⁹

The local labor markets also clear in equilibrium, so that the sum of employment over all sectors is equal to the local population. For each city j , it can be written as:

$$w_j L_j = \sum_{s \in S} \beta_s \tilde{M}_{js}, \quad (2.21)$$

which can be combined with equations (2.19) and (2.20) to show that, for any distribution

¹⁹ K_s is a sector-specific scaling constant given by $K_s = \frac{\beta_s \theta_s}{\theta_s - (1-\alpha_s)(\sigma_s-1)} \frac{(\sigma_s-1)^{\sigma_s-1}}{\sigma_s^{\sigma_s}} \kappa_{\alpha_s}^{-\sigma_s} \bar{b}_s^{(1-\alpha_s)(\sigma_s-1)}$.

of population $\{L_i\}_{i \in C}$ across cities, wages in city i follow:

$$w_i L_i = \sum_{s \in S} \beta_s t_i^{\alpha_s \sigma_s} w_i^{(1-\alpha_s)\sigma_s} b_i^{-(1-\alpha_s)(\sigma_s-1)} L_i^{-\nu_s(1-\alpha_s)(\sigma_s-1)} \times \sum_{j \in C} \mu_{ijs} \frac{w_j L_j}{\sum_{k \in C} \mu_{kjs} t_k^{\alpha_s \sigma_s} w_k^{(1-\alpha_s)\sigma_s} b_k^{-(1-\alpha_s)(\sigma_s-1)} L_k^{-\nu_s(1-\alpha_s)(\sigma_s-1)}} \quad (2.22)$$

, where, for each sector s , $\mu_s = \{\mu_{ijs}\}_{(i,j) \in C^2}$ are the terms of the inverse of the $\tau^{1-\sigma_s}$, with τ the iceberg trade costs matrix. Note that, as long as for any (i, j) , $\tau_{ii} < \tau_{ij}$, and for any s , $\sigma_s > 1$, $\tau^{1-\sigma_s}$ is strictly diagonally dominant. Therefore, it can be inverted under these conditions. Equation (2.22) illustrates the general equilibrium mechanisms that are included in this section: when the emission tax changes in city j , it affects wages in city j .

As in section 2.2, I assume free mobility and workers' utility maximization. As a result, welfare is the same across cities and equal to \bar{u} . From equation (2.1), free mobility implies that, in each city j :

$$a_j Z_j^{-\gamma} L_j^{-\delta} \frac{w_j}{P_j} = \bar{u}. \quad (2.23)$$

Finally, I assume that the total population of workers across cities is fixed and normalize it to 1 without loss of generality, so that:

$$\sum_{j \in C} L_j = 1, \quad (2.24)$$

which can be combined to equation (2.23) to show that:

$$L_j = \frac{\left(a_j Z_j^{-\gamma} \frac{w_j}{P_j}\right)^{1/\delta}}{\sum_{i \in C} \left(a_i Z_i^{-\gamma} \frac{w_i}{P_i}\right)^{1/\delta}}. \quad (2.25)$$

Given equation (2.16), for any set of emission taxes and masses of firms in each sector, emissions of pollution in city j follow:

$$Z_j = \sum_{s \in S} \beta_s \alpha_s \frac{\tilde{M}_{js}}{t_j}, \quad (2.26)$$

which, combined to equations (2.21), (2.22), and (2.25), implies that is an homogeneous function of degree -1 with respect to the average level of emission taxes across cities. Similarly, wages and population are homogeneous functions of degree zero with respect to the average level of emission taxes across cities. Combining equations (2.10) and (2.10)

implies that local price indices are homogeneous functions of degree $\sum_{s \in S} \beta_s \alpha_s \frac{\sigma_s}{\sigma_s - 1}$ with respect to the average level of emission taxes across cities. As a result, equation (2.23) implies that workers' welfare is a homogeneous function of degree $\gamma - \sum_{s \in S} \beta_s \alpha_s \frac{\sigma_s}{\sigma_s - 1}$ with respect to the average level of emission taxes across cities. Consequently, when investigating the central planner's welfare maximization problem, I focus on the distribution of emission taxes across cities and keep the average emission tax to 1.

2.3.2 Equilibrium and Optimal Set of Emission Taxes

An equilibrium is defined as a set of populations $\{L_j\}_{j \in C}$, wages $\{w_j\}_{j \in C}$, masses of industrial firms $\{M_{js}\}_{(j,s) \in C \times S}$, price indices $\{P_{js}\}_{(j,s) \in C \times S}$, and welfare \bar{u} , that solves the set of free entry conditions (2.19), goods market clearing conditions (2.20), local labor markets clearing conditions (2.21), free mobility conditions (2.23), and the national population clearing condition (2.24). The equilibrium is identified, to the extent that the number of equations is equal to the number of endogenous variables. As in section 2.2, I define the optimal set of emission taxes as the set of taxes that maximizes workers' welfare \bar{u} , under the constraint that the average tax remains equal to 1. In this extended version of my model, there are no closed-form solutions, neither for the equilibrium distributions of the endogenous variables, nor for the set of emission taxes that solves the maximization problem. Therefore, I rely on numerical methods to solve the model.

2.4 Estimation of the Model

To perform a quantitative policy analysis, I take the model to the data. In particular, I consider the set of French commuting zones as the set of cities in which workers live.²⁰ This set of commuting zones constitutes a partition of the full French metropolitan territory and each zone is defined statistically to be an area where local inhabitants both work and live. Using this definition of cities, I estimate the parameters of the model presented in section 2.3. First, I estimate equation (2.11) to retrieve elasticities γ and δ . Second, I use some of the model's predictions to estimate sector-specific expenditure shares $\{\beta_s\}_{s \in S}$, elasticities of substitution $\{\sigma_s\}_{s \in S}$, Pareto shape parameters $\{\theta_s\}_{s \in S}$, elasticities of agglomeration economies $\{\nu_s\}_{s \in S}$, and emission elasticities $\{\alpha_s\}_{s \in S}$. Third, I compute a matrix of road travel time between French commuting zones and use it as a proxy for the trade cost matrix. Finally, equipped with this estimated model, I retrieve the model's primitives – the level of amenities, productivities, and emission taxes – from data on the distribution of workers, wages, and industrial PM2.5 emissions across cities in 2012.

²⁰I use the 2010 definition of the French “*zones d'emplois*”.

2.4.1 Data

In this subsection, I describe the datasets used in the quantitative exercise. Firm-level datasets come from confidential French administrative data. I use annual balance sheets and income statements for the universe of French firms from 1994 to 2016 as reported in the FICUS databases for 1994-2007 and in the FARE databases for subsequent years. A firm is identified by a stable administrative code called SIREN. The main variables of interest are total sales, average employment (number of workers and total wages paid), location, and the main sector of activity.

Plant-level information on energy consumption comes from the EACEI (*Enquête Annuelle sur la Consommation d'Énergie dans l'Industrie*) surveys, which are available from 1994 to 2016. These surveys include all energy-related expenditures, with details on energy types and fuels (quantity consumed and expenditures) at the plant level. The types of energies reported are electricity (consumed and self-generated), steam, natural gas and other types of gas, coal, lignite, coke, propane and butane, domestic and heavy fuels, oil, and other types of petroleum products. The surveys also provide the plant-level number of employees. The surveys cover all large plants (over 20 employees) in the industrial sectors – with the exception of the power sector – and a subset of smaller plants (between 10 and 19 employees) that is randomly selected each year. On average, between 8,000 and 11,000 plants are included in the annual survey. I use these surveys to compute atmospheric emissions at the plant-level. The Ominea database (CITEPA, 2020) provides emission factors that associate to each fuel the corresponding amount of pollutants that are emitted. For each plant, I multiply the amount of each fuel consumed by the corresponding emission factor and sum across fuels to compute the total quantity of pollutant emitted. To ensure the validity of this approach, I compute correlation between these constructed emission values and actual values declared by large plants under the European directive and publicly available in the European Pollutant Release & Transfer Register (E-PRTR) between 2003 and 2016.²¹ Table B.5.1 in Appendix B.5.1 displays Pearson correlation coefficients along with statistical significance. It appears that constructed emissions correlate well with actual emissions of these large plants.

To estimate γ , the elasticity of environmental damage from pollution, I use geographic emissions from the Emissions Database for Global Atmospheric Research (EDGAR) from the European Commission's Joint Research Center (JRC). EDGAR provides annual grid-ded emissions of atmospheric pollutants disaggregated across polluting sectors at 0.1x0.1

²¹Although actual plant-level measures of atmospheric pollutants emissions are publicly available at the plant level in the European Pollutant Release and Transfer Register (E-PRTR), it only includes the largest plants resulting in a much restricted sample of industrial plants than the EACEI sample. Using the E-PRTR data would not allow me to estimate sector-specific emission elasticities.

degree resolution. Data is publicly available on the JRC’s dedicated website ²². Information on the methodology can be found in Crippa et al. (2018). The benefit from using this dataset rather than EACEI computed emission is that EDGAR aggregated values provide an exhaustive picture of all emissions (while EACEI surveys only account for the largest plants). I use EDGAR’s data from 2000 to 2015. I spatially aggregate it at the commuting zone level by intersecting the EDGAR grid with official ESRI shapefiles for French municipalities aggregated in commuting zones.

To retrieve the model’s primitives, I use a 2012 cross-section dataset of wages, labor, and industrial PM2.5 emissions at the commuting zone level. For wages and labor, I use data from the *Insee*’s website. For PM2.5 industrial emissions, I use the National Spatial Inventory (INS) (2020). The INS is a publicly available dataset reporting emissions of around 40 pollutants from natural and anthropogenic sources.²³ The data is available at the municipality level, which is a more disaggregated spatial unit than EDGAR’s grid cells. However, data is only available for 2004, 2007, and 2012. Therefore, I use the INS only to retrieve the model’s primitives, which are identified through cross-sectional variation, and rely on EDGAR for the panel estimation. Precisely, I use PM2.5 emissions from the manufacturing industry and from production processes. These activities correspond to codes 3 and 4 in the Selected Nomenclature for Air Pollution.

2.4.2 Estimation of Congestion and Pollution Damage Elasticities

Equation (2.11) can be re-written as a log-linear relationship between wages, utility, amenities, price indices, local labor supply and emissions:

$$\log w_{it} = \log u_{it} - \log a_{it} + \log P_{it} + \delta \log L_{it} + \gamma \log Z_{it}. \quad (2.27)$$

Introducing a time variable t , city i endowments in amenities and productivity may vary over time (for instance, due to enhancements to local transport networks or to the creation of museums, parks, etc.).

Furthermore, the free migration (no migration frictions) assumption is likely not verified when using annual data. Therefore, in reality, welfare is potentially different across cities. In practice, I only observe wages, populations, and emissions. Thus, I run the following equation:

$$\log w_{it} = \delta \log L_{it} + \gamma \log Z_{it} + \chi_i + \mu_t + \epsilon_{it}. \quad (2.28)$$

Year fixed effects μ_t eliminate annual shocks that are common to all cities. City fixed

²²https://edgar.jrc.ec.europa.eu/index.php/dataset_ap50

²³<http://emissions-air.developpement-durable.gouv.fr/index.html>

effects χ_i absorb city-specific fixed characteristics. The error term ϵ_{it} corresponds to city-year specific shocks, including shocks to local amenities a_{it} , local welfare u_{it} , and local price indices P_{it} . Equation (2.20) implies that local price indices are functions of the set of local productivities, of the trade cost matrix and of the other endogenous variables. Hence, city-specific shocks to local productivities or to the trade cost matrix induce shocks on local price indices. Wages, populations and emissions are also functions of these shocks. As a result, a standard OLS estimation strategy would produce biased estimates. Indeed, any positive shock to local amenities in a given city positively affects local population and emissions. Therefore, the OLS identification assumption is unlikely to hold. To overcome this identification issue, I build shift-share instrumental variables for both local labor supply and local emissions.

To build my instruments, I follow the recent literature on shift-share instruments that has extended the Bartik (1991) approach (Card, 2001; Autor et al., 2013; Nunn & Qian, 2014; Bartelme, 2018; Bombardini & Li, 2020; Barrows & Ollivier, 2021). The idea is to approximate city-specific growth rates in populations and emissions using national growth rates of these variables in disaggregated industries and interact them with city-specific shares in an initial period. Thus, I build two instruments as:

$$\tilde{L}_{it} = \sum_s L_{it_0s} \frac{L_{st}}{L_{st_0}}, \text{ with } s \in \Omega_L, \quad (2.29)$$

and

$$\tilde{Z}_{it} = \sum_s Z_{it_0s} \frac{Z_{st}}{Z_{st_0}}, \text{ with } s \in \Omega_Z. \quad (2.30)$$

Variables \tilde{L}_{it} and \tilde{Z}_{it} are the excluded instruments of a 2SLS strategy where I instrument variables L_{it} and Z_{it} in equation (2.28). These instruments are independent of city-year specific unobservable shocks to local amenities, productivity or trade costs. They are only functions of year-specific national shocks and city-specific initial sector shares in employment and emissions. National shocks, common across cities, and initial sector shares, constant across time, do not threaten identification when year and city fixed effects are included. As thoroughly detailed in Goldsmith-Pinkham et al. (2020), a recent extensive analysis of shift-share instruments, in order for these two instruments to be exogenous and the identification assumption to be respected, my strategy implicitly assumes that unobserved shocks on productivity, amenities or trade costs are uncorrelated with initial industry shares.²⁴

My estimation strategy builds on annual industrial emissions data, aggregated at the

²⁴An alternative identification assumption is that the common shocks are exogenous (Adão et al., 2019; Borusyak et al., 2021).

commuting zone level, and on firm and plant-level data on wages, employment and activity codes also aggregated at the commuting zone level. Each firm is identified by a unique 9-digits identification code (SIREN) and each plant is identified by a 15 digits identification code (SIRET) with the 9 first digits corresponding to the SIREN of the firm to which the plant belongs. Based on these two codes I merge firm and plant-level panels. Thus, for each plant I have a value of the mean wage paid at the firm level. Based on plant-level zip codes, I compute average wages across local plants at the commuting zone level. This results in a panel of local wages across commuting zones from 2000 to 2015. The plant-level panel also includes the number of workers employed on average over the year as well as an activity code (in the French Activity Nomenclature). I sum across local plants for each sector to build a panel of total sectoral employment at the commuting zone level from 2000 to 2015. The right hand side of the estimating equation only depends on the total local employment across sectors, but the instrument defined in equation (2.29) builds on the sector local and national disaggregation across sectors. The French Activity Nomenclature distinguishes between hundreds of very precisely defined industries. I aggregate codes at a more aggregated level which I call industries and end up with 16 different categories so that Ω_L is composed of agriculture, extraction activities, manufacturing activities, the energy sector, waste management, construction, trade, transport, hotels and restauration, telecommunications, finance, real-estate, public administrations, teaching sector, health, arts and other activities.

I build a panel of pollutants emissions across commuting zones from 2000 to 2015 disaggregated across polluting sectors using the geographic emissions dataset from EDGAR. Data is available as annual $0.1^{\circ} \times 0.1^{\circ}$ degree-grid sets for each sectors. Each data point corresponds to the quantity of pollutant emitted within the grid cell annually per unit of area. Based on geographic coordinates, I attribute each cell from the EDGAR grid to the corresponding French commuting zones (based on publicly available geographic information on municipalities and compositions of commuting zones). When a cell overlays several commuting zones, I attribute emissions based on the surface share of the grid overlaying each commuting zone. Finally, I sum emissions over all grid cells within each commuting zone and I obtain a panel of sectoral emissions across years and commuting zones. I build such panel for particulate matters PM10 and PM2.5, nitrous gases NOx and ozone precursors CO and COVNM. Emissions are disaggregated across 16 polluting sectors: power industry, oil refineries and transformation industry, combustion for manufacturing, energy for buildings, fuel exploitation, non-metallic minerals production, chemical processes, iron and steel production, non-ferrous metals production, non energy use of fuels, solvents and products use, food and paper, manure management, agriculture (3 distinct activities), waste management and disposal (3 distinct activities) and fossil

fuel fires. These are polluting activities defined by the Selected Nomenclature for Air Pollution recommended in the emission register guidebooks implemented by the IPCC. I exclude some polluting activities that my framework does not include (aviation, road transportation and shipping).

Table 2.4.1 reports the first stage results. Table 2.4.2 displays the outcome of δ and γ estimations. I observe that the OLS coefficients, without and with city and year fixed-effects, underestimate the values of elasticities δ and γ . Results from the preferred strategy, reported in columns 3 to 7 of Figure 2.4.2 are all positive and statistically significant (except for the elasticity of environmental damages from emissions of NMVOC, that stands for Non Methanic Volatile Organic Components). Estimated values of δ have the same interpretation in columns 3 to 7, and correspond to the strength of congestion effects. Estimated values of γ correspond to the elasticity of environmental damage from alternative atmospheric pollutants. The effects of PM2.5 and PM10 on welfare are the stronger compared to other pollutants, and very close to each other. This is in line with evidence on the negative effect of particulate matter on human health and real estate prices. Another harmful pollutant, NOx (nitrogen dioxide), has also been found to have a negative effect on health. However, this pollutant is not primarily emitted by industrial activities (rather by the transport sector). This explains the lower value of γ in column 5. Finally, the main health effect from NMVOC and CO is due to their chemical transformation in ozone. The fact that NMVOC and CO do not directly affect health may explain the lower values obtained for γ in columns 6 and 7. To conduct my quantitative exercise, I focus on PM2.5 as the representative pollutant and use estimates from column 3 in my numerical analysis.

	(1)	(2)	(3)	(4)	(5)	(6)
	<i>Labor</i>	<i>PM2.5</i>	<i>PM10</i>	<i>NOx</i>	<i>NMVOC</i>	<i>CO</i>
Instrument	0.181*** (0.015)	0.889*** (0.028)	0.868*** (0.026)	1.211*** (0.023)	0.960*** (0.035)	1.126*** (0.035)
Observations	5136	4864	4864	4864	4864	4864
City & Year FE	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>

Standard errors in parentheses

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

Table 2.4.1: First stage results for IV estimation of δ and γ

Note: Column 1 reports the first-stage regression coefficient for local levels of employment (number of workers). Columns 2 to 6 report first-stage regression coefficients for local emissions of a set of pollutants.

2.4.3 Estimation of Sector-Specific Parameters

To calibrate parameters $\{\beta_s, \sigma_s, \theta_s\}_{s \in S}$, I closely follow the approach of Shapiro & Walker (2018). I compute Cobb-Douglas parameters $\{\beta_s\}_{s \in S}$ using the model's prediction that these parameters are the national shares of revenues of each sectors. I sum firm-level

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
δ	-0.01 (0.02)	0.09 (0.05)	0.91** (0.35)	0.89** (0.35)	1.00*** (0.38)	0.88** (0.34)	0.89** (0.35)
γ	0.05** (0.02)	-0.06 (0.06)	0.49*** (0.17)	0.49*** (0.17)	0.08* (0.05)	0.08 (0.24)	0.31* (0.16)
Observations	4864	4864	4864	4864	4864	4864	4864
City & Year FE	No	Yes	Yes	Yes	Yes	Yes	Yes
IV	No	No	Yes	Yes	Yes	Yes	Yes
Pollutant	PM2.5	PM2.5	PM2.5	PM10	NOx	NM VOC	CO

Standard errors in parentheses

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

Table 2.4.2: Estimation of δ and γ

Note: This table presents estimated coefficients δ and γ from equation (2.28). Column 1 reports coefficients estimated through a standard OLS approach, using emissions of PM2.5. Column 2 reports coefficients estimated using city and year fixed-effects, using emissions of PM2.5. Column 3 reports coefficients estimated using city and year fixed-effects and instrumenting local labor and emissions of PM2.5. Columns 4 to 7 report coefficients using emissions of other pollutants instead of PM2.5.

revenues and compute these parameters. Elasticities of substitution are recovered using the prediction that, for each sector, the ratio of total payment to labor on value-added is equal to $(1 - \alpha_s) \frac{\sigma_s - 1}{\sigma_s}$. To estimate the Pareto shape parameters, I use the prediction that the distribution of firm-level sales is Pareto with shape parameter $\theta_s/(\sigma_s - 1)$ and I estimate the following equation:

$$\log \Pr(x > X_{is}) = \gamma_{0s} + \gamma_{1s} \log(X_{is}) + \epsilon_{is} \quad (2.31)$$

where X_{ics} represents sales made by firm i from sector s in city c and $\gamma_{1s} = -\frac{\theta_s}{\sigma_s - 1}$. Because Pareto distribution better fits the right part of the productivity distribution, I restrict the sample to the upper decile of the firms sample.²⁵

I estimate sector-specific emission elasticities $\{\alpha_s\}_{s \in S}$ at the plant level. Assuming that for a plant i , installed in city c , and producing goods from sector s , with a productivity ϕ_{ics} , production follows $q_{ics} = z_{ics}^{\alpha_s} l_{ics}^{1-\alpha_s} \phi_{ics}^{1-\alpha_s}$ and that demand is $q_{ics} = k_{cs} p_{ics}^{-\sigma_s}$, where p_{ics} is the price charged, the relationship between plant-level employment and emissions follows:

$$l_{ics} = \tilde{k}_{cs} z_{ics}^{\frac{\alpha_s(\sigma_s - 1)}{1 + \alpha_s(\sigma_s - 1)}} \phi_{ics}^{\frac{(1 - \alpha_s)(\sigma_s - 1)}{1 + \alpha_s(\sigma_s - 1)}}, \quad (2.32)$$

where $\tilde{k}_{cs} = \left(\left[(1 - \alpha_s) \frac{\sigma_s - 1}{\sigma_s} \right] w_c^{-\sigma_s} k_{cs} \right)^{\frac{1}{1 + \alpha_s(\sigma_s - 1)}}$, which, in log, allows us to recover emission elasticities using the following empirical specification:

$$\log l_{icst} = \beta_{0s} + \beta_{1s} \log(z_{icst}) + \epsilon_{icst}. \quad (2.33)$$

Appendix (B.5.2) provides the algebra corresponding to the above equation. Using a

²⁵This is an approach developed in Hsieh & Ossa (2016) and Antràs et al. (2017).

standard OLS strategy to estimate equation (2.33) would yield biased estimates. Indeed, unobserved plant specific productivity shocks are correlated with the levels of employment and emissions. This is the “transmission bias” identified in the literature on production functions estimation. To circumvent this problem, I instrument plant-level emissions z_{ics} with exogenous fuel-specific energy price variation. Building on Sato et al. (2019), I use a fixed-weight energy price index that measures the plant-specific exposure to variation in fuel prices based on each plant energy consumption across fuel types in the first period when it is observed. I define this instrument as:

$$FEPI_{ist} = \sum_{f \in \Omega_{fuels}} \omega_{f,ist_0} p_{f,st}, \quad (2.34)$$

where $FEPI_{ist}$ is the plant-specific energy price index built from plant i share of energy expenditures in fuels $f \in \Omega_{fuels}$ (coal, natural gas, electricity, etc.) in period t_0 and $p_{f,st}$ the specific fuel price common to all plants in sector s in period t . An advantage of this emission elasticity estimation strategy is that it only requires to observe plant-level inputs. Conveniently, the EACEI surveys combine plant-level labor and energy consumption across fuels. Note that both the endogenous variable and the instrument are constructed from the EACEI surveys but are not aggregated across fuels using the same weights: the instrument is the sum of consumption weighted by fuel prices and the emissions are the sum of consumption weighted by emission factors that are also fuel specific. I estimate equation (2.33) for each sector using PM2.5 emissions. Table 2.4.4 displays all sector-specific coefficients that I use in my numerical analysis.

To the best of my knowledge, the only other recent work estimating industrial emission intensities is Shapiro & Walker (2018). However, they pool observations from different sectors and only estimate one coefficient. To allow the comparison with their results, Table 2.4.3 reports the results that I obtain for different pollutants when industrial sectors are pooled. The first-stage using the fixed-weight energy price index yields significant estimates that are intuitive: on average increasing polluting fuels prices by one percent causes a one percent decrease in atmospheric pollutant emissions. Across pollutants, the second stage yields similar estimates for the pooled sectors emissions intensity, approximately 5%. They use plant level emission data combined with information on emissions abatement costs from PACE surveys in the United States. Their pooled estimates are lower than mine (1.1% for PM2.5, 1.1% for PM10, .1% for NOx, .08% for CO and .8% for COVNM). The fact that their sample includes a wider range of economic activities that are less pollution intensive may explain this difference.

Finally, I estimate sector-specific elasticities of agglomeration economies $\{\nu_s\}_{s \in S}$ using a

	(1) log (<i>emiPM25</i>)	(2) log (<i>PM10</i>)	(3) log (<i>NOx</i>)	(4) log (<i>CO</i>)	(5) log (<i>COVNM</i>)
<i>First Stage:</i>					
log <i>FEPI</i>	-1.00*** (0.012)	-1.00*** (0.012)	-0.98*** (0.010)	-1.07*** (0.009)	-1.06*** (0.009)
<i>Second Stage:</i>					
$\frac{\alpha(\sigma-1)}{1+\alpha(\sigma-1)}$	0.049*** (0.004)	0.049*** (0.004)	0.051*** (0.004)	0.046*** (0.003)	0.046*** (0.003)
Observations	223,401	223,401	223,402	223,406	223,406
Year, Region & Industry FE	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>	<i>Yes</i>

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.4.3: Two-stages calibration of emission elasticities α (with pooled industrial sectors)

regression of city-specific mean firm productivity with city size. To estimate productivity, I regress the logarithm of firm-level value added on city-sector specific fixed-effects. Given equations (2.12) and (2.18), the residual corresponds to firm specific TFP (up to a proportionality factor $(1 - \alpha_s)(\sigma_s - 1)$). Finally, I compute for each sector the average of this residual within each city and regress it on the logarithm of local population to obtain the set $\{\nu_s\}_{s \in S}$.

	Sales share (β) (1)	Elasticity of substitution (σ) (2)	Pareto shape parameter (θ) (3)	Elasticity of Agglo. Economies (ν) (4)	Pollution elasticity (α) (5)
Sectors					
Automobile & transport	.02	2.27	1.85 (.19)	.05 (.006)	.036 (.001)
Chemicals	.05	3.48	1.33 (.13)	.02 (.005)	.080 (.012)
Communications & Electronics	.01	3.45	2.85 (.16)	.00 (.004)	.025 (.010)
Electrical Equipment	.01	3.91	3.59 (.31)	.00 (.004)	.079 (.010)
Extraction	.01	2.22	1.58 (.08)	.02 (.006)	.126 (.015)
Food, beverages & Tobacco	.05	3.82	2.89 (.02)	-.01 (.001)	.030 (.003)
Machinery & Equipment	.01	3.10	3.88 (.08)	.01 (.002)	.015 (.005)
Metal	.02	3.06	2.29 (.02)	.00 (.002)	.084 (.006)
Rubber & Plastic	.02	2.92	2.66 (.15)	.00 (.002)	.120 (.011)
Textile & Apparel	.01	2.99	2.21 (.04)	.02 (.002)	.089 (.006)
Wood & Paper	.01	2.90	2.42 (.03)	.01 (.002)	.035 (.005)
Other Manufacturing	.02	2.47	1.49 (.01)	.00 (.002)	.035 (.008)
Non manufacturing	.75	2.69	1.48 (.00)	.00 (.000)	.021 (.016)
<i>Pooled (except NM)</i>				.01 (.001)	

Table 2.4.4: Estimated parameters

2.4.4 Trade Costs Matrix

One option to retrieve bilateral trade costs would be to estimate a gravity model, however I do not observe bilateral trade flows between French cities. The second option is to use

a proxy for trade costs. I follow Yamada (2020) and use travel time between cities. In particular, I rely on Baum-Snow et al. (2020) and I assume the following concave relationship between iceberg bilateral trade costs and bilateral travel times:

$$\tau_{ij} = 1 + \rho \times (\text{hours of travel time}_{ij})^\xi. \quad (2.35)$$

I compute travel time using the Open Source Routing Machine (OSRM) API Python client *osrm-py*. The OSRM is a C++ routing engine for shortest paths in road networks building in the road network data of the project OpenStreetMap. Figure 2.4.1 displays the average travel time by road across all commuting zone. As expected, central areas are better connected to the rest of the country.




Figure 2.4.1: Average travel time across all commuting zone

Note: On this map I plot the average travel time from each commuting zone toward all the other commuting zones. Travel time is in hours. Travel times were computed using the Open Source Routing Machine (OSRM) API Python client *osrm-py*

2.4.5 Recovering Local Characteristics

Once all parameters are estimated, I numerically solve the non-linear system of equilibrium conditions defined in section 2.3 by equations (2.22), (2.25), and (2.26). To do so, I use data on the distribution of workers, wages, and industrial PM2.5 emissions across cities in 2012 (see subsection 2.4.1 for a description of the data). This procedure allows me to retrieve the set of idiosyncratic characteristics – amenities and productivities – as well as the set of emission taxes across cities. Without loss of generality, I set the mean emission tax, the mean productivity, and the sum of amenities across cities to be equal to one. In practice, I invert the equilibrium conditions using the Levenberg–Marquardt

algorithm as implemented in the *scipy* library in Python 3. This algorithm is an interpolation between the standard Gauss–Newton algorithm and the method of gradient descent that is more robust to the choice of initial values. Results from this step are described in the next section. Appendix B.6.1 provides further details on the sources of data for observed equilibrium distributions of wages, populations, and emissions across French cities. Appendix B.6.2 provides descriptive statistics on wages, populations, and emissions across French cities. Finally, appendix (B.6.3) provides descriptive statistics on the computed exogenous distributions of amenities, productivities, and emission taxes across French cities.

2.5 Spatially Heterogeneous Emission Regulations

2.5.1 Evidence of Spatially Heterogeneous Emission Regulations

In Figure 2.5.1, I plot the city-level emission taxes obtained through the numerical solution of equilibrium conditions described in subsection 2.4.5. More precisely, panel (a) plots local emission taxes as a function of population in 2012, panel (b) plots local emission taxes as a function of amenities, and panel (c) plots local emission taxes as a function of productivities. Throughout this section, I plot the logarithm of all variables. This is done for visual clarity, as most distributions are very skewed. I emphasize that these local emission taxes are not actual policy instruments implemented in France. As in Shapiro & Walker (2018), a way to see these values is as follows: if all French emission regulations were to be replaced by local emission taxes, Figure 2.5.1 would be the set of taxes that leads to the distribution of PM2.5 emissions that we observe across French cities. Equivalently, these local emission taxes can be seen as measures of how stringent are emission regulations in each city relative the others.

Figure 2.5.1 uncovers a large heterogeneity of emission regulations across cities. The positive correlation between local taxes and populations also reveals that existing regulations impose more stringent emissions regulations in larger cities. This observation is supported by anecdotal evidence presented in Appendix B.2. Indeed, several French regulations, that seek to enhance local air quality, impose higher emission marginal costs in larger cities. The simple model from section 2.2 predicts that an optimal set of emission taxes was positively correlated with both amenities and productivities.²⁶ Panel (c) from Figure 2.5.1 displays a strong and positive correlation between local levels of productivities and emission taxes that must not be too far from the optimum. However, panel (b) from

²⁶In section 2.2, the optimal set of emission taxes is positively correlated with both amenities and productivities when $\theta > 0$, which is the case given values computed in section 2.4.

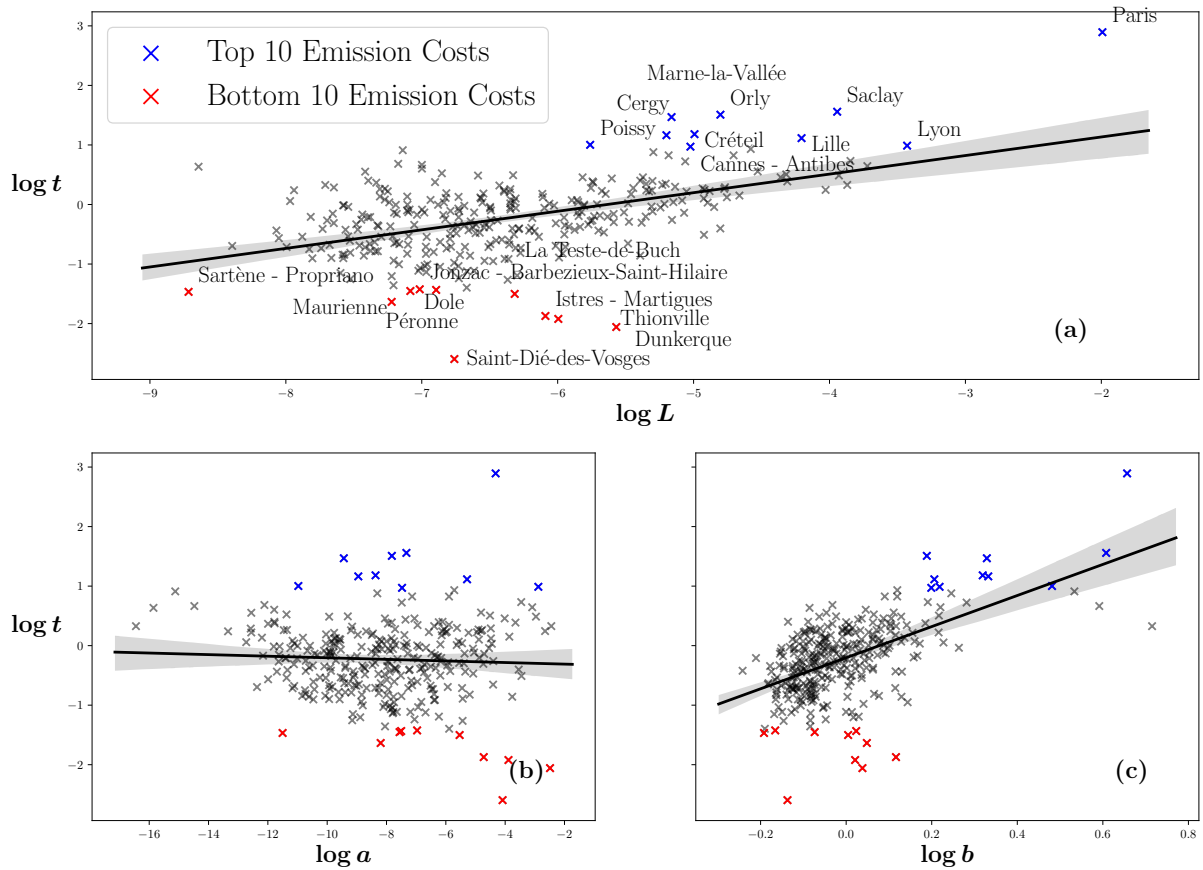


Figure 2.5.1: Log-linear relationship between local relative emission tax and local relative population, productivity and amenity

Note: On this figure I plot the current distribution of emission taxes as a function of current populations on panel (a), of local amenities on panel (b), and of local productivity on panel (c). The distributions of local relative emission taxes, productivities and amenities are estimated by inverting the observed equilibrium using one tradable sector and data for 2012.

Figure 2.5.1 shows that the current distribution of emission taxes does not display such positive correlation between emission taxes and local amenities. Therefore, the current set of emission taxes across French cities seems to potentially be off from what the model indicates to be the optimal distribution.

Figure 2.5.1 also highlights in blue the top ten cities where emission regulations are the most stringent and in red the bottom ten cities where they are the less stringent. As expected, cities with the highest levels of emission taxes are also the most productive and cities with the lowest levels of emission costs are the least productive. Yet, among cities with the highest levels of emission taxes, some have very high levels of amenities, whereas others have very low levels of amenities.

2.5.2 Welfare Impacts of Spatially Heterogeneous Emission Regulations

Results from the previous subsection highlight the spatial heterogeneity of local emission taxes across French cities. In this subsection, I analyze how this heterogeneity affects the distribution of emissions, population, and wages. To that end, I compare the observed equilibrium with a counterfactual equilibrium where the level of emission taxes is uniform across all cities. The differences between the two equilibrium can help understanding the consequences of the spatial heterogeneity of emission taxes. Using the values of idiosyncratic amenities and productivities obtained previously, I compute a new equilibrium by solving equilibrium conditions (2.22) and (2.25) using a uniform emission tax across cities. This uniform emission tax value is the mean of all emission taxes obtained in the previous computation, which is normalized to one. The counterfactual equilibrium is solved numerically using the Levenberg–Marquardt algorithm. Figure 2.5.2 displays the results from this exercise. For each endogenous variable of interest in the model, I plot a map showing the variation in level of moving from a uniform emission tax to the current distribution of emission taxes (that is heterogeneous across cities).

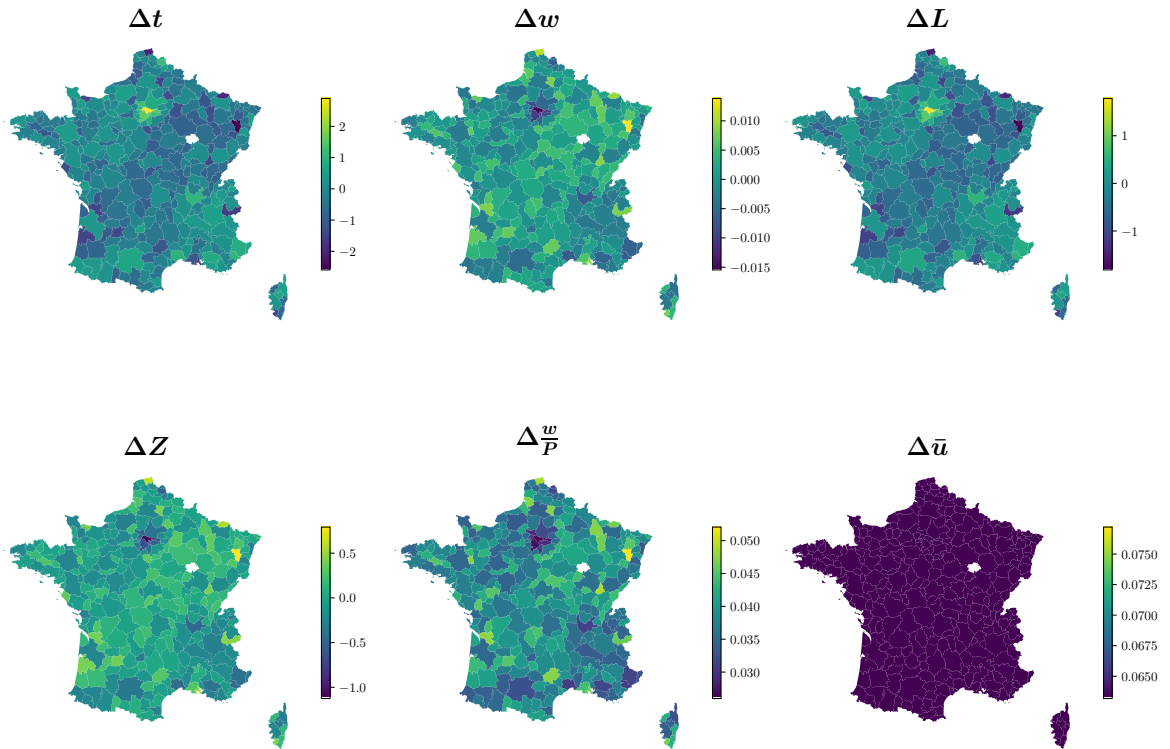


Figure 2.5.2: Spatial reallocations and welfare effect of imposing spatially heterogeneous emission taxes

Note: On this figure I map the changes in emission taxes, wages, industrial workers' populations, emissions, real wages, and welfare due to a move from a spatially uniform emission tax to the current set of emission tax.

Because emission taxes are actually higher in larger cities, such as Paris, Lyon, Bordeaux, or Toulouse, the map of changes in t reveals an increase in these cities and a decrease in less populated cities. The map of ΔZ reveals a decrease in the largest cities and an increase in the smallest cities. Hence, when the central planner implements more stringent air quality policies in largest cities relative to smaller cities, emissions are relocated from the most populated areas to the least populated areas. Figure 2.5.2 also plots the changes in the level of population (ΔL) to illustrate the reallocation of population across cities due to the move from uniform to heterogeneous emission taxes. Under the free migration assumption, workers react to changes in local air quality (and prices which are affected by the change in local emission taxes) and move away from more polluted areas. The map of ΔL shows that there is a reallocation of workers from small cities, where emissions have increased, to larger cities, where emissions have decreased. As larger cities become cleaner, more people have an incentive to live there.

Figure 2.5.3 plots the log-difference in local populations between the current equilibrium and the counterfactual one with uniform emission taxes, as a function of initial population (in panel (a)), of amenities (in panel (b)), and of productivities (in panel (c)). For a given city, a positive number indicates that the local population is higher under the current distribution of emission taxes than under the counterfactual. Panel (a) reveals that the largest cities have larger populations because of the spatial heterogeneity in regulations. Panels (b) and (c) show that the top ten cities for which population increases the most are also highly productive cities but not cities with particularly high amenities. Figure 2.5.4 is similar to Figure 2.5.3, except that it shows changes in emissions instead of populations. Panel (a) shows that the largest cities emit less pollution because of the spatial heterogeneity in regulations. Panels (b) and (c) illustrate the fact that more stringent regulations in more productive cities lead to lower levels of emissions there, compared to the counterfactual. There are no clear correlation with the distribution of local amenities. Panels (a) from Figures (2.5.3) and (2.5.4) suggest that more stringent regulations in larger cities make them cleaner (less emissions) and in turn larger (higher population) than what they would be under a uniform regulation.

2.5.3 Optimal Mean-Preserving Set of Heterogeneous Regulations

The model can be used to assess which set of emission taxes maximizes welfare. To do so, I numerically solve the \bar{u} maximization problem under the constraint that the mean emission tax remains equal to one. This optimization exercise yields a new spatial equilibrium that can be compared to the current equilibrium. With this numerical solution, I present

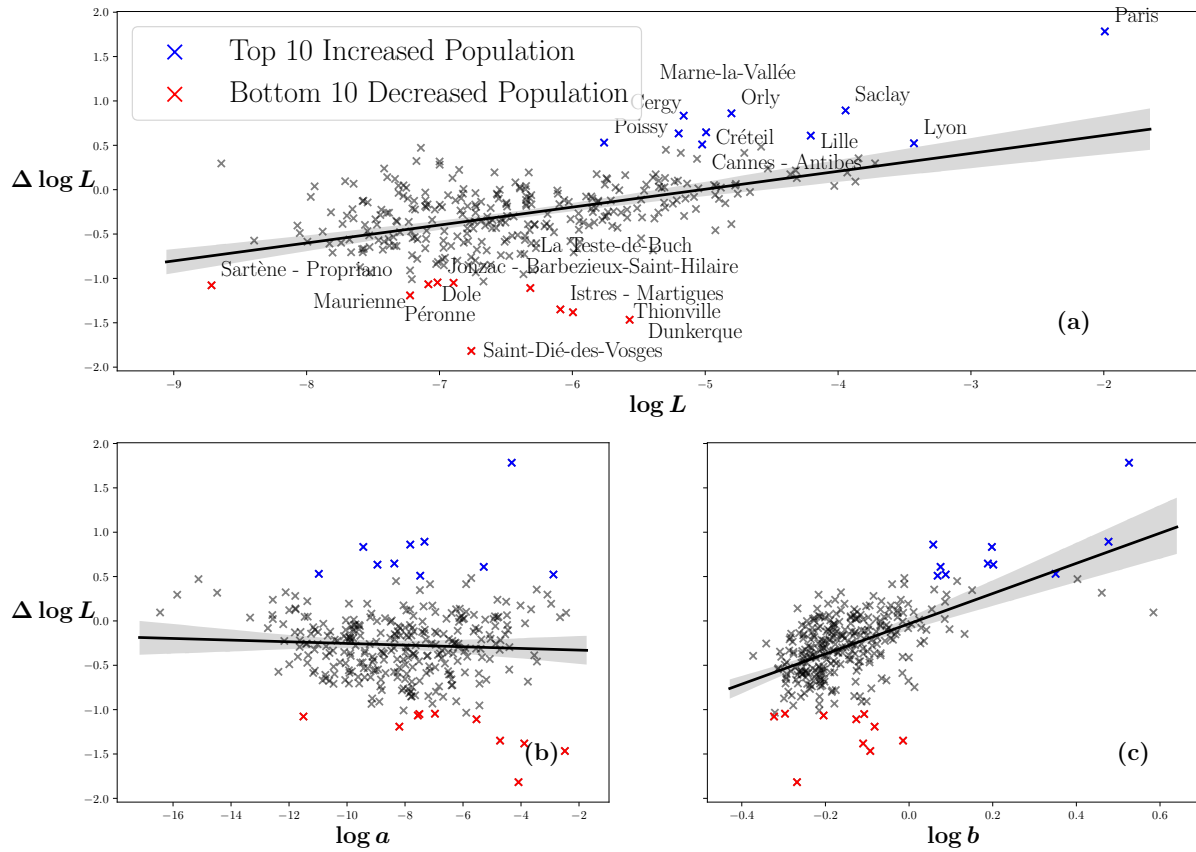


Figure 2.5.3: Effects of spatial heterogeneity emission stringency on local populations: under the current policy, large cities are larger than what they would be if they faced a uniform emission tax

Note: On this figure I plot the relative change in industrial workers' population at the commuting zone level as a result of a move from a uniform to the current distribution of emission taxes. This change is plotted as a function of current populations on panel (a), of local amenities on panel (b), and of local productivity on panel (c).

normative and empirical results on the gap between the current distribution of relative stringencies across French cities and the distribution that would maximize workers' welfare. In practice, I identify for each city an optimal level of regulation that I compare to the current level of regulation.

Figure 2.5.5 displays the results from the maximization exercise. It plots the optimal distribution of the logarithm of emission taxes as a function of the logarithm of: current city population (panel a), amenities (panel b), and productivities (panel c). The current distributions from Figure 2.5.1 are depicted in black and the optimal distributions are depicted in red. Figure 2.5.5 reveals that increasing the relative emission regulations in larger cities would achieve welfare gains. Figure 2.5.5 shows that the distributions in red display less variation on the left part than on the right part. This is because the optimal set of emission taxes results from a constrained numerical optimization algorithm: small cities are so much smaller than large cities that the constraint on the mean emission tax implies that optimal emission taxes are put to the minimum level possible in small cities.

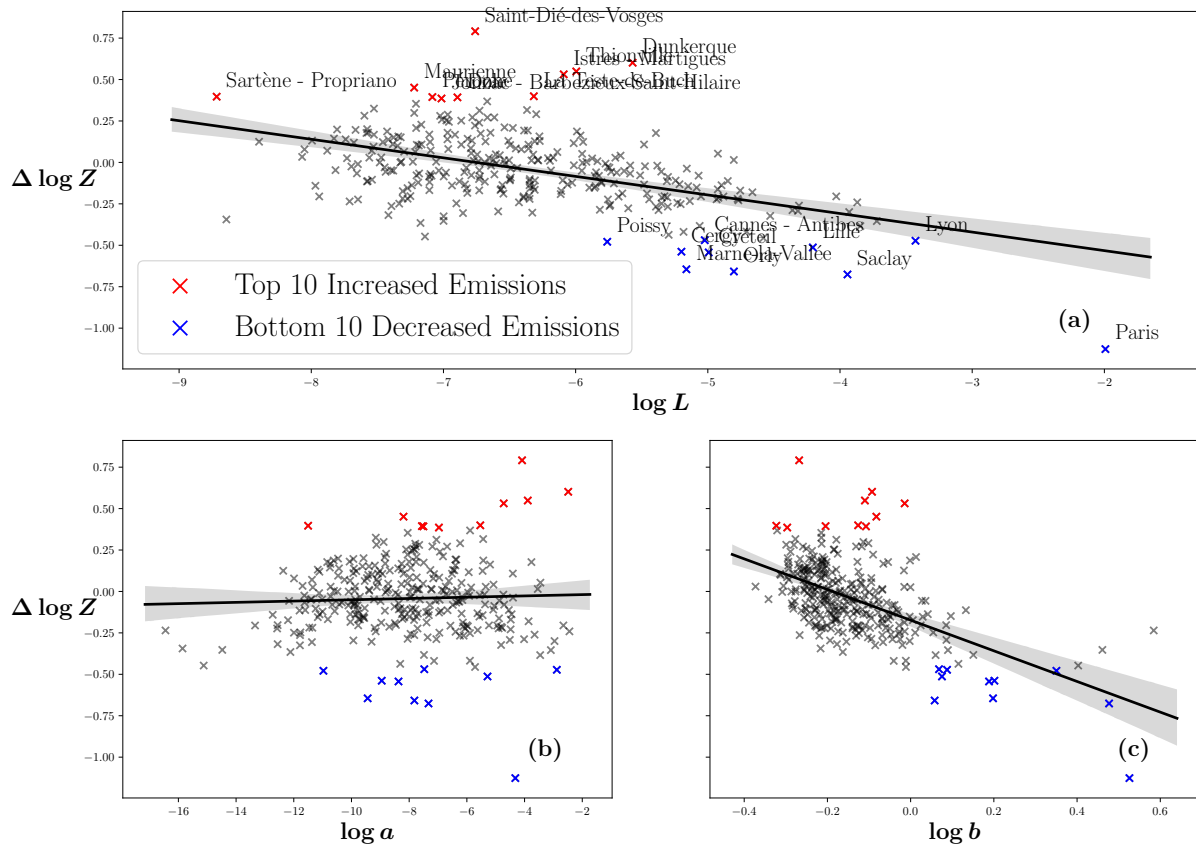


Figure 2.5.4: Effects of spatial heterogeneity emission stringency on local emissions: under the current policy, large cities emit less pollution than what they would if they faced a uniform emission tax

Note: On this figure I plot the relative change in industrial emissions at the commuting zone level as a result of a move from a uniform to the current distribution of emission taxes. This change is plotted as a function of current emissions on panel (a), of local amenities on panel (b), and of local productivity on panel (c).

These points are corner solutions. However, because distributions of cities' characteristics (amenities and productivities) are so skewed, only the largest cities matter for aggregate welfare.

As discussed in section 2.2 and illustrated in panel (b) from Figure 2.5.5, increasing the relative stringency of regulations in cities with higher amenities would be welfare-improving. This is because cities with higher amenities are generally too small because workers and firms do not internalize the impact of their location choices on local air quality. Panel (a) shows the reallocation necessary to move from the current distribution of relative emission taxes to the optimal one. These taxes need to be raised in the currently largest cities and decreased in the smallest cities. Panel (c) from Figure 2.5.5 indicates that there is already a strong positive correlation between the spatial distribution of emission taxes and the distribution of idiosyncratic productivities. Solving the optimization problem confirms that such correlation is optimal. Indeed, panel (c) shows that the coefficient of the log-linear relationship between the local idiosyncratic levels of productivities and the

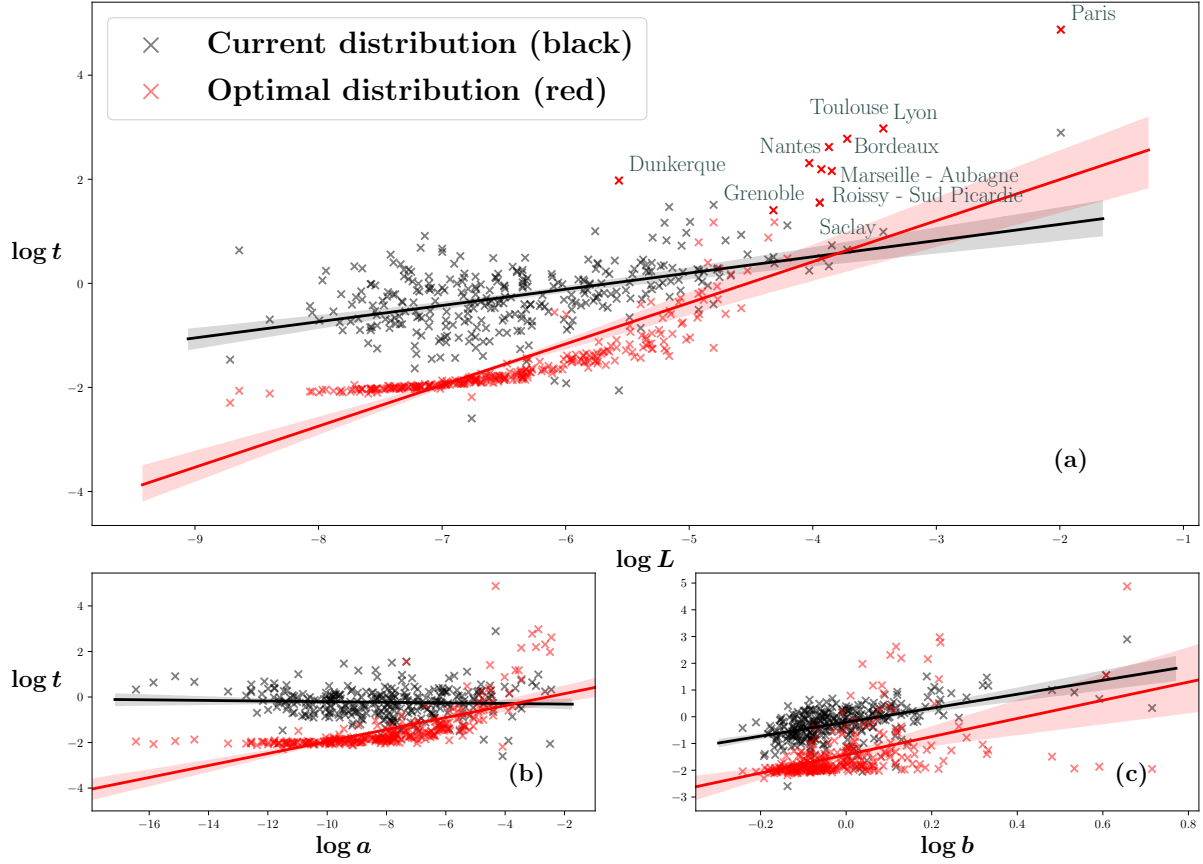


Figure 2.5.5: Comparison between current distribution of relative emission taxes (black) and the mean-preserving optimal distribution (red): emission taxes are lower in smaller cities and increased in larger cities

Note: On this figure I plot the current distribution of emission taxes in black and the optimal distribution of emission taxes as functions of current populations on panel (a), of local amenities on panel (b), and of local productivity on panel (c). The current distribution is depicted in black and the optimal distribution is depicted in red.

relative emission taxes is currently not too far from the optimal one.

Reallocating Population Across Space

Figure 2.5.6 displays the population reallocation effect that would follow from the adoption of the optimal distribution of relative emission taxes. In particular, it highlights in blue the cities that would see their population increase. Only a small subset of cities (around 15 out of 300 cities) would grow larger and the rest would become smaller. Specifically, panel (a) shows that population concentration is reinforced: the largest cities would grow even larger. Panel (b) indicates that this reallocation mainly comes from a reallocation of workers toward cities with higher amenities.

Figure 2.5.7 reveals the concentration effect of adopting the optimal set of emission taxes. I ranked cities according to their current population and then plot in black the cumulative population starting from the smallest city. The red line depicts the same cumulative sum

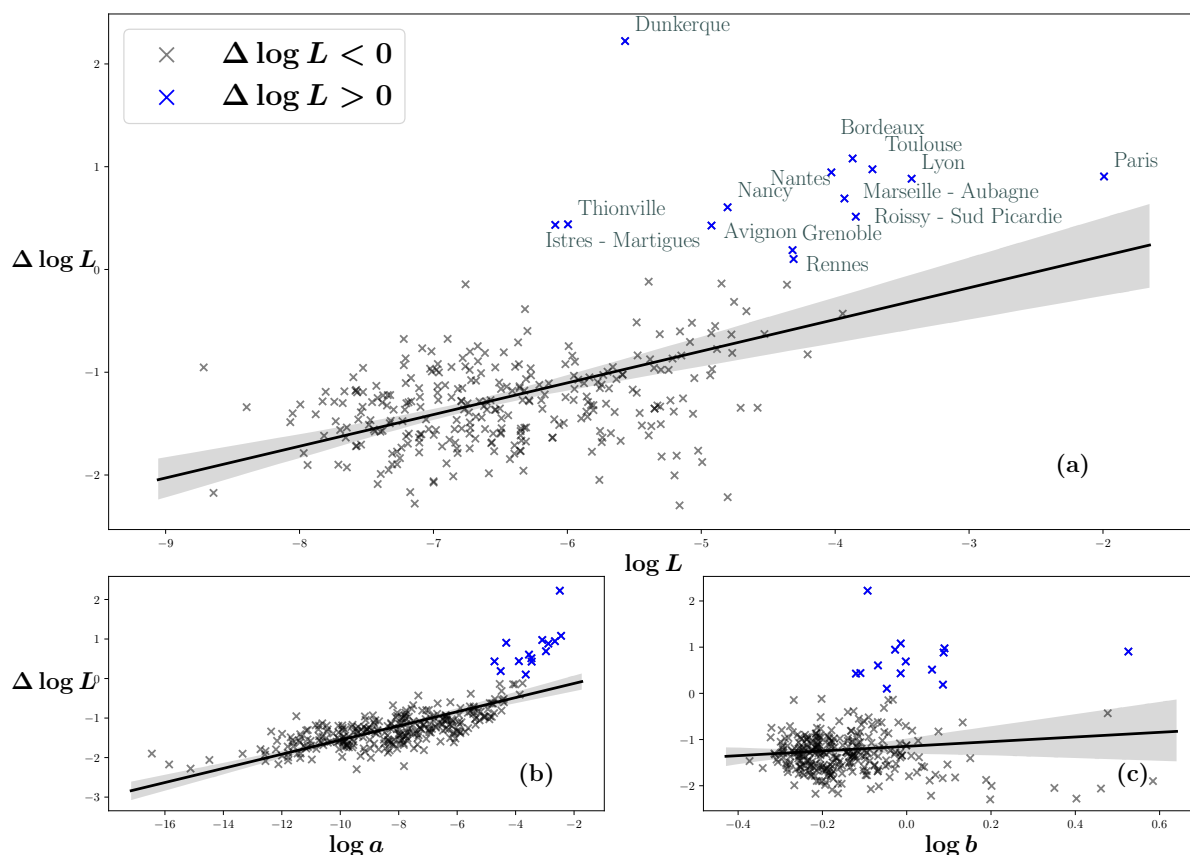


Figure 2.5.6: Population reallocation when moving from the current to the optimal distribution of relative emission taxes

Note: On this figure I plot the relative change in industrial workers' population at the commuting zone level as a result of a move from the current to the optimal distribution of emission taxes. This change is plotted as a function of current populations on panel (a), of local amenities on panel (b), and of local productivity on panel (c). In blue are the cities that would see their population increase as a result of this emission taxes distribution shift.

using what local populations would be under the optimal set of emission taxes. Since the red line is more skewed to the right than the black line, workers become more concentrated in the largest cities under the optimal distribution of emissions costs than under the current set of policies.

Reallocating Emissions Across Space

Figure 2.5.8 displays the pollution reallocation effect of moving to the optimal set of emission taxes. In particular, it highlights in green the cities that would see their local emissions decrease. In blue are the cities for which emission would decrease and population would increase. As expected, it is optimal to decrease emissions in the currently most polluted cities (panel (a)), corresponding to a reallocation of emissions from cities with high amenities to cities with low amenities (panel (b)).

Figure 2.5.9 compares the cumulative distributions of emissions under the current (in

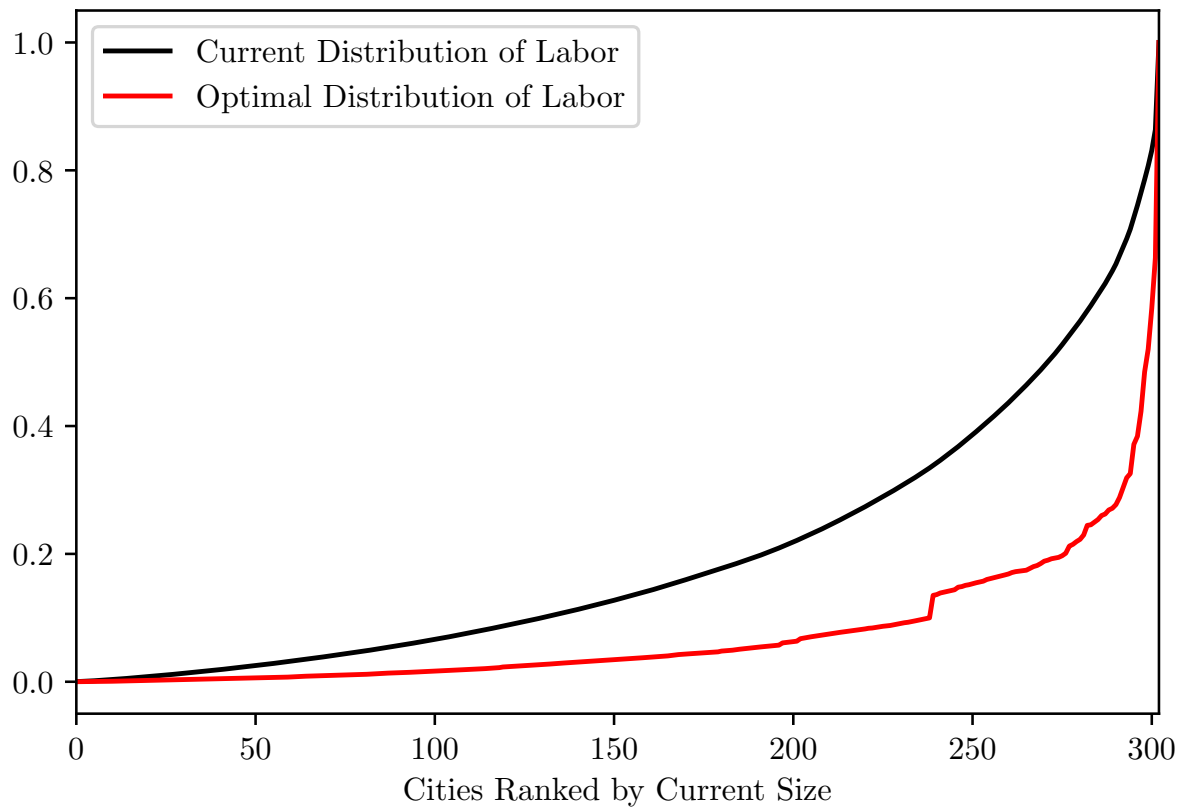


Figure 2.5.7: Cumulative distributions of population under the current and the optimal distributions of emission taxes

Note: On this graph, I plot the cumulative sum of industrial workers' population across French commuting zones as a function of their respective ranks by industrial workers' population. I depicted this distribution in black for the current distribution and in red for the optimal distribution. A 45 degree line who mean that emissions are evenly spread across cities. A curve skewed to the right indicates spatial concentration of workers in a few places.

black) and the optimal (in red) distributions of emission taxes across cities. Since the red line is less skewed to the right than the black line, the optimal reallocation of emissions across cities reduces spatial concentration of emissions.

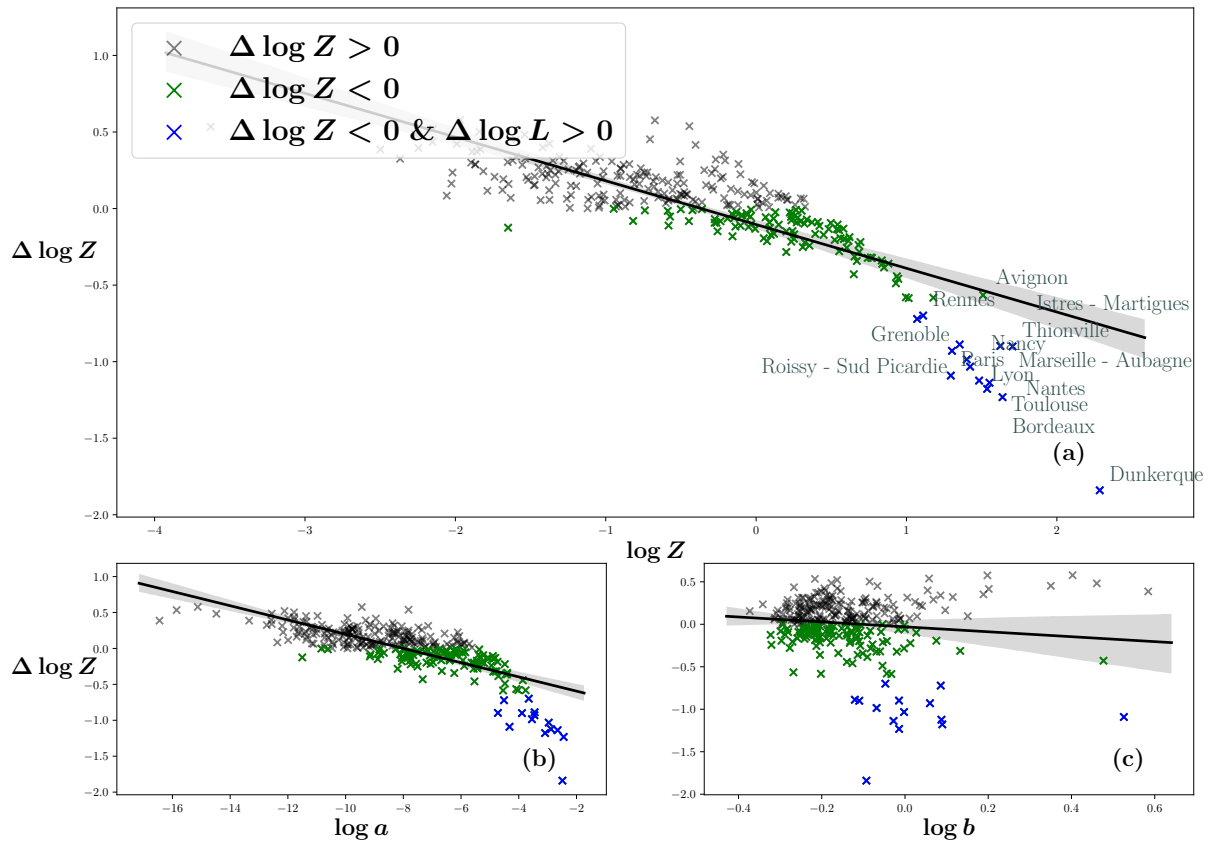


Figure 2.5.8: Emissions reallocation when moving from the current to the optimal distribution of relative emission taxes

Note: On this figure I plot the relative change in emissions at the commuting zone level as a result of a move from the current to the optimal distribution of emission taxes. This change is plotted as a function of current emissions on panel (a), of local amenities on panel (b), and of local productivity on panel (c). In blue are the cities that would see their population increase as a result of this emission taxes distribution shift. In green are the cities that would see their emissions decrease.

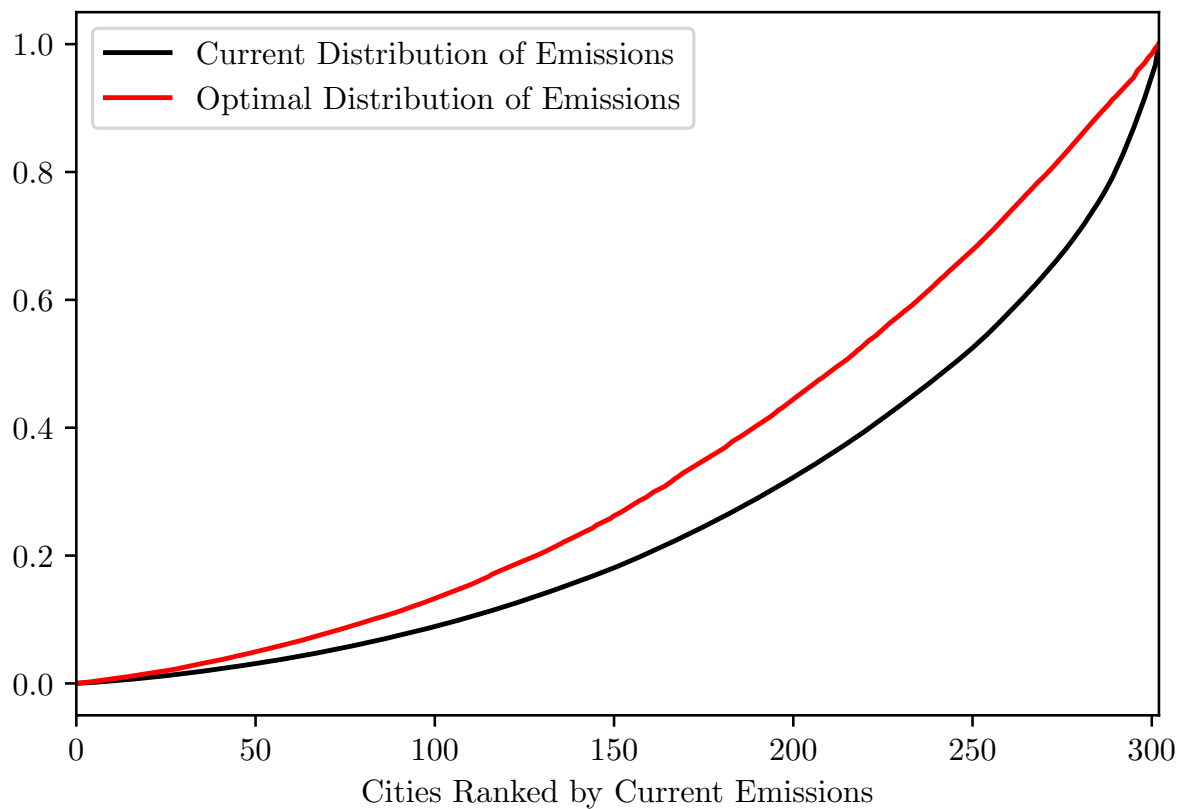


Figure 2.5.9: Cumulative distributions of pollution emissions under the current and the optimal distributions of emission taxes

Note: On this graph, I plot the cumulative sum of industrial emissions across French commuting zones as a function of their respective ranks by quantity of emission. I depicted this distribution in black for the current distribution and in red for the optimal distribution. A 45 degree line who mean that emissions are evenly spread across cities. A curve skewed to the right indicates spatial concentration of emissions in a few places.

2.6 Conclusion

In this paper, I analyze how the spatial distribution of local air quality policies can improve welfare. I incorporate in a simple spatial model endogenous industrial emissions of atmospheric pollutants. As pollution has a negative effect on welfare and workers move away from polluted areas, I show that the central planner should adopt higher emission taxes in cities that have higher levels of amenities and productivity. I build a general version of the model that includes heterogeneous industrial sectors and I estimate the model's parameters using an extensive set of data on French firms and cities. Then, I retrieve the model's primitives – local amenities, productivities, and emission taxes – from data on the distribution of workers, wages, and industrial PM2.5 emissions across cities. In the French context, I show that current regulations correspond to a set of emission taxes that are higher in the largest cities. Finally, I numerically solve the central planner's welfare maximization problem. I find that further increasing emission taxes in the largest French cities could achieve welfare gains.

This paper sheds light on a spatial tradeoff between clean air and productivity. My results indicate that policy makers should not implement environmental policies uniformly across locations. Rather, they should be more stringent in large cities. This has important implications in terms of environmental regulations design. The optimal set of relative emission taxes that I identify in this paper is independent of the average level of the tax. This means that it is applicable to any new air quality regulation: for any given objective of pollution reduction, local regulatory burdens should follow the distribution described in this paper. Future work should incorporate additional factors such as inequalities in exposure to atmospheric pollution (Agyeman et al., 2016; Banzhaf et al., 2019; Colmer et al., 2020a). Extending this work to include distributional effects of air quality regulations and their impacts on inequalities in exposure to pollution (Currie et al., 2020; Shapiro & Walker, 2021) could further help investigate optimal place-based environmental policies.

Chapter 3

The Effects of Air Pollution on Exports: Evidence from PM2.5 & French Firms

3.1 Introduction

Local air pollution has adverse effects on both health and non-health outcomes of individuals (Aguilar-Gomez et al., 2022). In particular, a large literature has documented pollution’s detrimental effects on workers’ productivity (Graff Zivin & Neidell, 2012; Chang et al., 2016; He et al., 2019; Adhvaryu et al., 2022), employees’ cognitive abilities (Allen et al., 2016; Heyes et al., 2016; Archsmith et al., 2018; Chang et al., 2019), and labor force participation (Hanna & Oliva, 2015; Holub et al., 2021). However, little is known about the aggregate consequences of pollution shocks for the economic performance of firms employing workers exposed to these shocks and only a nascent literature connects pollution-induced productivity shocks to economic outcomes (Dechezleprêtre et al., 2019; Fu et al., 2021; Khanna et al., 2021). If pollution adversely affects labor productivity and labor supply, standard microeconomic theory suggests that it should increase firms’ production costs, lower their output, and affect their revenues. The identification of such effects could reveal that the economic benefits of enhancing air quality are larger than previously estimated.

In this paper, we provide causal estimates of air pollution’s effects on firm-level exports sales and quantities. Our results indicate that higher PM2.5 concentrations have a negative effect on both the total value and the quantity of exports. Focusing on the universe of French exporting firms, between 2009 and 2015, we find that an 10% increase in the average PM2.5 concentrations faced by plants owned by an exporting firm during the three months preceding the export leads to a 1.5% decrease in the total value and a 1.4% decrease in the quantity exported.¹ These estimates are robust to a set of tests designed to investigate the main limitations of our strategy. An heterogeneity analysis across low and high labor intensity firms supports the prediction that effects of pollution on exports are consequences of negative effects on the labor force.

¹Over our period of analysis, the average monthly PM2.5 atmospheric concentration ranged between 12 and 20 $\mu g/m^3$ depending on the locations.

Our analysis builds on exhaustive and very disaggregated temporal and geographic information on firm-level trade data and local air quality. We observe French firm-level monthly exports and combine them with information on the geographic coordinates of plants owned by exporting firms as well as local environmental variables: PM2.5 atmospheric concentrations, wind direction, precipitations, and temperatures.² In comparison to previous literature that relies on pollution measurements at the monitoring station level, we use daily pollutant concentrations from a chemistry-transport model that provides local estimates of pollutant concentration over the whole French territory and not only over locations where measurement stations are potentially endogenously located. Furthermore, our panel of exporting firms, provided by French customs, is exhaustive which means that we observe the universe of exporting French firms, and not only a sub-population as is frequent in the literature. Another advantage of our trade data is its availability at a monthly resolution. Indeed, most of the literature on worker-level impacts of pollution finds that such effects happens over the course of several days or weeks.

Omitted variable bias and reverse causality are the two main obstacles to estimating pollution's causal impact on economic activity. In our context, omitted variable bias arises from local factors that directly affect both air pollution concentrations and economic activity. In particular, a large literature has shed light on the effect of temperature on health, labor supply, and productivity (Dell et al., 2014; Graff Zivin & Neidell, 2014; Addoum et al., 2020; Aguilar-Gomez et al., 2021). But temperature also affects pollution concentrations and omitting its effect is likely to bias estimates. For instance, because very high temperatures cause higher concentrations of air pollutants, omitting the effect of temperature would lead to an overestimation of the negative effect of pollution on exports. Hence, we combine pollution data with data on temperatures and precipitations to control for these confounding factors. Reverse causality is due to air pollution's endogeneity: a positive productivity shock in a given location potentially induces higher levels of production, commuting, or consumption, thus emitting more pollutants and increasing their atmospheric concentration. It ensues a positive correlation between local productivity and pollution that does not reflect the causal effect of an exogenous pollution shock on economic activity. To overcome this issue, we exploit exogenous variations in monthly wind direction in areas where plants owned by exporting firms are located. We find that wind direction is a strong predictor of PM2.5 concentrations and argue that, after controlling for seasonality effects, yearly trends, and weather variables, local changes in wind direction only affect economic activities through their effect on local PM2.5 concentrations. In practice, we

²Because most production activities are indoor (Vette et al., 2001), we focus on PM2.5 pollution, which is composed of particles small enough to enter buildings.

estimate for each location a model where monthly PM2.5 concentrations are predicted by wind direction, temperatures, precipitations, calendar month (to control for seasonality), and year (to include time trends). Based on the estimated model, we then compute predicted values for local PM2.5 concentrations based on wind direction only. Finally, we use these predicted PM2.5 values as the excluded instrument of a 2SLS estimation strategy, where we instrument firm-level PM2.5 concentrations averaged across plants.³ Conditional on controlling for relevant weather variables to solve the omitted variable bias and satisfy our instrument’s exclusion restriction, our strategy allows us to recover unbiased causal estimates of the effects of PM2.5 pollution. This strategy is inspired by the recent literature that has used wind direction as an exogenous predictor of local pollution concentration (Hansen-Lewis, 2018; Deryugina et al., 2019). We adapt the strategy to multi-plant firms that are exposed, through their set of plants, to pollution shocks that potentially happen in different regions. In particular, Deryugina et al. (2019) highlights that such strategy does not require a full description of the geographical layout of each locations where the plants are located, or to identify the sources of PM2.5 pollution (in comparison to quasi-experimental settings). This allows for a broader geographic scale of analysis, over a long time period.

This paper contributes to the literature that has identified the negative effect of air pollution on manufacturing productivity. Aguilar-Gomez et al. (2022) reviews the scientific background explaining how air pollution may have multiple health effects on individuals. Exposition to pollution can affect the functioning of the heart, lungs, and brain, in the form of both short-term impairments and latent impacts. Individuals may also suffer from non-health impacts from pollution, resulting from the interactions between physiological effects of pollution and behavioral responses from these individuals. Among these non-health effects are the decreased workers’ productivity and lower labor force participation. Most analysis are conducted at the worker level. Using data from the California air quality monitoring network, Graff Zivin & Neidell (2012) showed that a 10 parts per billion (ppb) decrease in ozone concentrations increases agricultural workers’ productivity by 5.5 percent on the same day.⁴ Chang et al. (2016) finds a negative effect of outdoor PM2.5 pollution on indoor workers at a pear-packing factory: when pollution is above $15 \mu g/m^3$, an increase by $10 \mu g/m^3$ reduces workers’ productivity by approximately 6% of average hourly earnings. Focusing on call centers in China, Chang et al. (2019) also finds that PM2.5 pollution decreases the productivity of workers whose occupation requires cognitive rather than physical faculties. Whereas these investigations are based on day-to-day

³Fontagné et al. (2018) is an other example of paper that aggregates plant-level shocks to the firm-level. They use plant-level shocks on electricity price to instrument export prices and estimate trade elasticities.

⁴A 10 ppb increases in ozone concentration corresponds to a $20 \mu g/m^3$ increase.

analysis and identify immediate effects of pollution on workers' productivity, longer periods of exposure also affect productivity. He et al. (2019) find that manufacturing workers in China suffered from prolonged exposure to PM2.5 pollution: a $10 \mu g/m^3$ increase in PM2.5 concentration over 25 days reduces daily output by 1%. Regarding labor force participation, Hanna & Oliva (2015) show that a decrease in local SO2 pollution in Mexico City following the closure of a polluting plant led to an increase in the local labor participation. In reverse, Holub et al. (2021) finds that PM10 shocks in Spain due to dust clouds from Sahara led to more workers taking sick days, and therefore to a lower labor force participation.⁵

If labor supply and workers' productivity are affected by air pollution then it is likely that it also has an impact on the activity of firms where they work. To the best of our knowledge Dechezleprêtre et al. (2019) and Fu et al. (2021) are the only other papers that estimate pollution's impact on aggregate economic outcomes. Focusing on NUTS-3 EU regions, Dechezleprêtre et al. (2019) combines annual GDP data with daily satellite data on local atmospheric PM2.5 concentrations. They show that a $1 \mu g/m^3$ increase in PM2.5 concentration leads to a .8% drop in GDP during the same year. Fu et al. (2021) focus on China and use firm-level annual output data to show that a $1 \mu g/m^3$ decrease in PM2.5 leads to a .82% increase in productivity. Both papers rely on annual data, potentially missing short-term adaptation mechanisms, such as temporal reallocation of production. In particular, firms likely engage in adaptation strategies. For instance, Adhvaryu et al. (2022) recently showed that firms' managers can react to worker-level productivity shocks caused by pollution and reallocate workers across tasks to mitigate the impact at the firm level. In particular, Aguilar-Gomez et al. (2022) argue that adaptation is an important mechanism through which pollution affects economic agents. Even if they do not observe pollution itself, agents may react to its negative consequences. As a result, even without monitoring pollution, firms may react to it because they monitor their short-term productivity. For instance, if a local pollution shock leads to health issues or productivity losses among workers in a given month, local firms may observe the lower labor supply, or the reduced productivity. In reaction, they may increase their efforts to recruit productive workers. Aguilar-Gomez et al. (2022) also highlight the fact that the timing of non-health pollution effects may also be different. A pollution shock happening at a specific time may have consequences on firm-level productivity in a longer term. This can be caused by delayed impact of pollution exposure on workers, but also because of dynamic effects due to human capital accumulation. To that extent, the investigation of both short and longer-term impact of pollution shocks using temporally

⁵On a longer term, Isen et al. (2017) shows that early childhood exposition to air pollution causes lower earnings and labor force participation at age 30.

disaggregated data is important.

Our paper also contributes to the literature on exports and exogenous supply-side shocks related to climatic events. Works within this literature have mainly build on annual and country-level data. For instance, Jones & Olken (2010) shows that higher temperatures lowers exports from poor countries, but do not find any effect on the exports from rich countries. More recently, as higher temporal resolution trade data has become available, more precise investigations have been conducted. For instance, Karlsson (2021) uses US state-level export values and shows that very high or very low temperatures have a negative effect on the value of US exports. Using the annual version of French customs trade data, Colmer et al. (2019) estimate that domestic exposure to temperature shocks leads to lower economic performances. Nedoncelle (2021) uses the same annual trade data to show that temperature shocks in destination countries lead to drops in exports by large and multi-destination French exporters. To our knowledge, only one paper builds on monthly and firm-level trade data to investigate the effect of climate shocks on exports: focusing on extreme heat days in China, Li et al. (2021) find that high temperature shocks have a persistent and negative effect on exports. In the case of France, firm-level monthly trade data has not yet been exploited to investigate climate or pollution shocks. However, some works exploit exogenous shocks and assess their impacts on French firms exports. Malgouyres & Mayer (2018) consider the effect of tax credits affecting French firms' labor costs on their exports. Lafrogne-Joussier et al. (2022) identify which French firms were impacted by China's lock-down in the early stages of the Covid crisis.

The remainder of the paper is organized as follows. In the next section, we describe the environmental, plant-geolocation, and trade data. The third section presents our empirical strategy. In the fourth section, we describe the results and conclude in the fifth section.

3.2 Data

3.2.1 Atmospheric Pollution in France

We observe local pollution concentrations from the *CHIMERE* chemistry-transport model (Real et al., 2021). This model provides daily concentrations of PM10, PM2.5, NO2 and O3 on a grid covering the French metropolitan territories between 2009 and 2015. The grid resolution is four by four kilometers on average meaning that pollution data is available for several thousands of grid points across France.⁶ This data is the output of a simulation of a chemistry transport model that uses information on meteorological

⁶To be precise, *CHIMERE* data is available on a grid of 0.03 degrees by 0.06 degrees cells.

variables and pollutant emissions from various sources. PM2.5 refers to particles smaller than 2.5 micrometer that are in suspension in the air. Because of their very small size, such particles cause adverse health outcomes, such as asthma, lung cancer, and premature death. We choose to focus on PM2.5 pollution because their small size also means that they are able to penetrate most buildings, including those where production is happening (Vette et al., 2001). The left panel of Figure 3.2.1 illustrates the spatial distribution of the monthly PM2.5 air concentrations in $\mu g/m^3$ averaged over the period 2009-2015. PM2.5 pollution is, on average, unevenly spread across space. The right panel of Figure 3.2.1 displays the spatial distribution of monthly PM2.5 concentrations standard deviations between 2009 and 2015. Considering that local average concentrations in PM2.5 range between 12 and 20 $\mu g/m^3$, the 4 to 10 $\mu g/m^3$ range of local standard deviation indicate a high level of temporal variation in local levels of PM2.5 concentration. These local variations across time are the source of variation that our estimation strategy exploits.




Figure 3.2.1: Atmospheric PM2.5 Data : Maps of Average Values and Standard Deviations
Note: Monthly PM2.5 concentrations over the 84 months included in our sample (in $\mu g/m^3$).

Figure 3.2.2 plots the evolution of the mean, maximum, and minimum concentrations across the country. In our estimation strategy, we include calendar-month fixed-effects that control for the seasonality that appears on Figure 3.2.2.

A number of studies focusing on air pollution build on monitoring stations. These stations physically measure concentrations of pollutant in the air at the hourly level. In France, this network provides publicly available data. However, it is spatially scarce and does not necessarily correspond to the spatial distribution of manufacturing plants.⁷ For that reason, it misses a significant amount of the spatial variation that one finds in satellite based estimates of air pollution (Fowlie et al., 2019).

⁷Stations are usually located near highways.

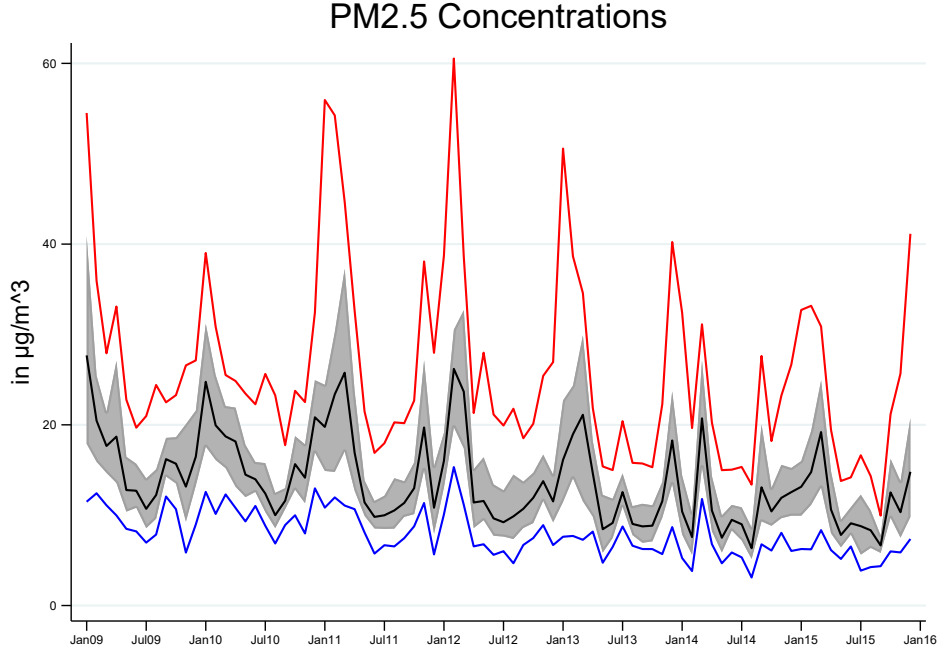


Figure 3.2.2: Atmospheric PM2.5 Data : Monthly Average, Minimum, and Maximum Local PM2.5 Concentrations

Note: The black line represent monthly PM2.5 concentrations average across *CHIMERE* grid cells. The shaded-in-grey area contains 90% of *CHIMERE* grid cells. The red and blue lines respectively represent the monthly maximum and minimum PM2.5 concentration across *CHIMERE* grid cells.

3.2.2 Wind Direction and Meteorological Controls

Our strategy also exploits data on local precipitations, temperatures, and wind direction. For precipitations and temperatures, we use the *Copernicus* satellite data.⁸ In particular, we retrieve daily measures between 2009 and 2015 over 6,453 grid cells for the French metropolitan territory. Each cell measures around 11 by 11 kilometers. We retrieve data on wind direction from the *MERRA* database, that has a resolution of 50 by 50 kilometers on average (the grid is composed of around 250 cells for the French metropolitan territory).⁹ We define wind direction as the direction in which the wind is blowing ("toward"). We define four main wind directions based on compass angles (each encompassing a 90 degree angle): northward, eastward, southward, and westward. For each *MERRA* hourly observation we convert the longitudinal and latitudinal wind speeds in a set of four dummies for whether the wind blows in one of these four directions. We then aggregate observations at the monthly level by summing the number of hours during which the wind blows in each direction. Figure 3.2.3 illustrates the dominant winds measured in each *MERRA* grid cell by plotting the average monthly percentage of hours when wind blows in each direction.

⁸Data is available at <https://surfobs.climate.copernicus.eu>.

⁹Data is available at <https://gmao.gsfc.nasa.gov/reanalysis/Merra/>.

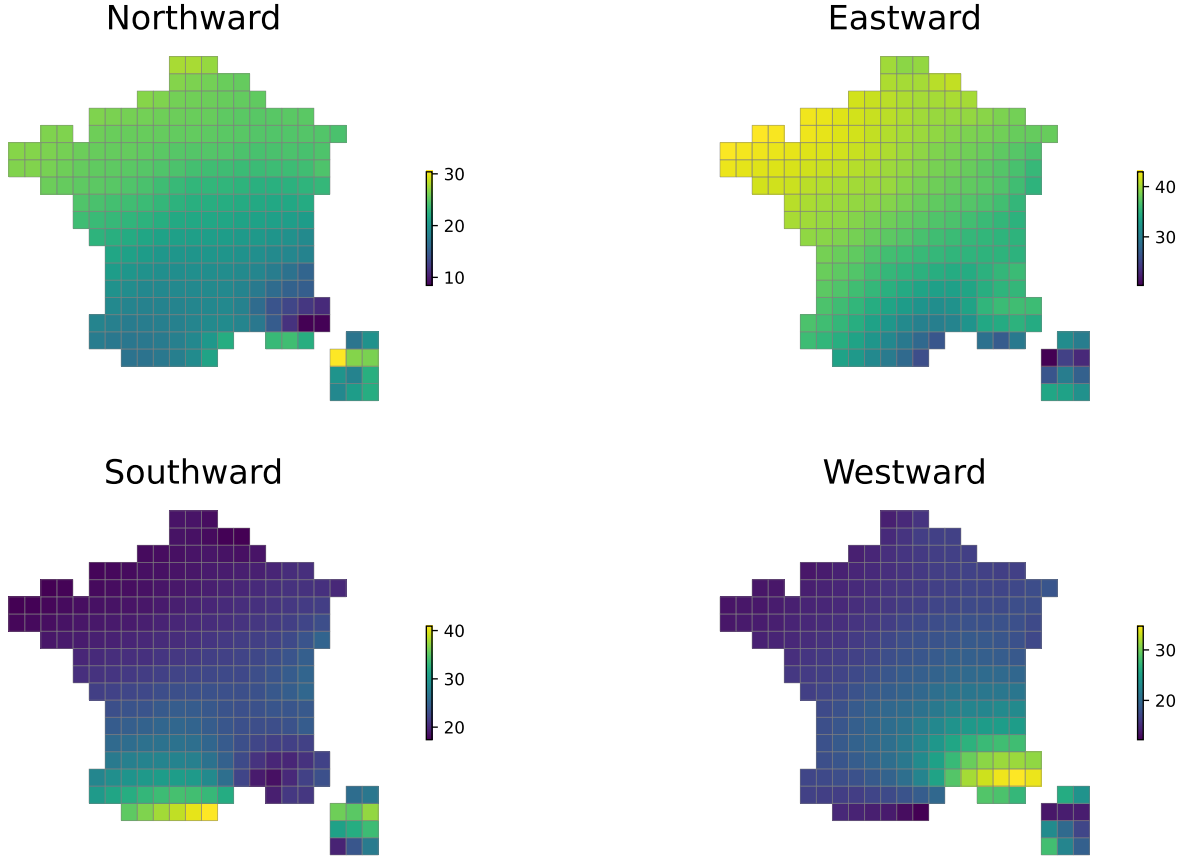


Figure 3.2.3: Wind Instrument : Dominant Winds

Note: These maps plot the average percentage of hours per month during which wind blows in each direction.

We combined PM2.5 concentrations data with temperatures, precipitations, and wind direction by overlapping *CHIMERE*, *Copernicus*, and *MERRA* grids. Both *MERRA* and *Copernicus* grids are less spatially disaggregated than the *CHIMERE* grid. As illustrated by Figure 3.2.4, for each grid, we compute the surface of each *CHIMERE* cell that intersects with each cell from the other two grids. Then we convert *Copernicus* and *MERRA* data to distribute the values across cells from the *CHIMERE* grid. We build a panel of PM2.5, temperatures, precipitations, and wind direction using average values at the monthly level, between 2009 and 2015 for each *CHIMERE* grid cell.

3.2.3 Trade Data

We exploit data on firm-level exports provided by the French customs office. The advantage of such trade data is its high level of temporal disaggregation that allows for causal inference and capturing firm level adaptation mechanisms. In particular, this level of temporal disaggregation is usually not available for complete firm or plant-level production data. This is why we focus on the effect of pollution on exports in this paper. The

trade data is composed of firm level declarations of exports of goods to foreign markets (outside of France within and outside the European Union). Trade transactions between French firms and foreign countries are subjects to the *Déclarations d'Echanges de Biens mensuelles* (DEB), a monthly trading declaration, for transactions with Member States of the European Union, and to the *Document d'Accompagnement Unique* (DAU), the unique accompaniment file, for transaction with countries outside of the EU. In both systems, firms have to declare each month transactions they have made with foreign countries. They may submit several declarations in a month, or file only one that contains all transactions that have happened during the month. For transactions within the EU, DEB follows a simplified procedure with less information when the value of the transaction is below 460 euros. In practice, only transactions with small values (inferior to 200 euros) are not declared. Before 2010, all observations below 1000 kilograms or 1000 euros were not included in the database. This represented around 0.5% of the value of exchanges but as much as 50% of the number of transactions which was misleading in some sectors such as wine production or automobiles parts (CASD, 2019). In 2010, 90% of these data were integrated to the base. Since 2011, all data is included.¹⁰

In practice, we have one observation per firm, good (defined in the Combined Nomenclature at the eight-digits level), destination, and month of sample. Each observation indicates the total value of the export in euros, its weight in kilograms, and, if applicable, the number of units exported. This last variable is only available for a few observations, we follow Fontagné et al. (2018) and only use quantities measured in weight. We construct our firm-level export panel according to Bergounhon et al. (2018)'s guidelines. To reduce the computational load of our regressions, we aggregate exports across destinations and products at the firm level. We provide a robustness test of our estimates using an alternative version of the trade panel aggregated at the four-digits activity level. We use data from 2009 to 2015, which represents 84 months in total. Figure 3.2.5 describes the monthly values of French exports, and the number of exporting firms over this period. We match these exporting firms to the set of plants they own and for which we have geographic coordinates. However, we don't have this information for all French plants, so the match is not perfect. As a result, some exporting firms are dropped out of our estimation sample. The full lines in Figure 3.2.5 represent the data contained in the raw customs datasets (all exports that are registered) and the dashed lines represent the values for the set of exporting firms matched to plants with geographic coordinates. Figure 3.2.5 shows that almost all exports are matched with at least one geo-localized plant. This is confirmed by the last line of Table (3.2.1) which shows that exporting firms that do not match with any geo-localized plant only represent 6% of total exports over the period.

¹⁰Weapons and all war materials exports are not included.

Our analysis relies on the temporal disaggregation of the trade data. To identify causal impact of pollution, we must use exogenous variations in local exposure to PM2.5. In particular, we control for yearly trends in pollution. This is one of the reason why having access to sub-year data is important. However, when we observe an export made by a firm, we cannot know when the exported goods were produced. If pollution affects export through its negative effect on workers and production, then we need to make assumption on the timing of production and exports. In our main strategy, presented in the next section, we chose to estimate the effect of firm-level exposure to pollution averaged over the three months that preceded the export. This implies that we assume that when firms produce a good to be sold on a foreign market, they ship it sometime during the next three months. To support this assumption, we provide some descriptive statistics on the timing of exports at the firm level. These statistics are presented in appendix C.0.1. Table ?? shows that around 20% of exporting firms are shipping goods at one month intervals. Moreover, these one-month-intervals exports represent around 98% of export values over the period of analysis. Although we don't have data on stocked good, it is likely that firms which export every month will also produce every month. Therefore, focusing on the effect of pollution over the three months before the export seems a valid strategy.

3.2.4 Plants Geo-localization

We match exporting firms with their set of plants using the direct correspondence between exporters' identification codes, the SIREN, with the first nine digits of plants' identification codes, the SIRET. We implement this step using the *SIRENE* database, which is the list of the universe of French firms and plants in France.¹¹

Information on the location of plants is built from the address contained in the *SIRENE* database using the geocoding engine *addok* developed by Etalab (the public administration in charge of open data). We use the list of plant coordinates built by Christian Quest using two instances of *addok*, the *Base Adresse Nationale* (BAN) and the *Base Adresses Nationale Ouverte* (BANO) from OpenStreetMap France.¹² For each plant, the address provided in the *SIRENE* database corresponds to a set of geographic coordinates. On Figure 3.2.6, we illustrate this with a map of the local density of plants belonging to exporting firms. We represent a grid, where each cell color corresponds to a quartile in the distribution of number of plants across all cells from this grid.

Multi-plants firms create an empirical challenge for our estimation. Exports are only

¹¹The database is available on <https://www.data.gouv.fr/fr/datasets/base-sirene-des-entreprises-et-de-leurs-etablissements-siren-siret/>.

¹²Data is available at https://data.cquest.org/geo_sirene/ and scripts on the corresponding Github <https://github.com/cquest/geocodage-spd>.

Activities	Number of Firms	Share of Exports	Average Number of Plants	Share of Single-plant Firms
Agriculture	15,150	.5	1.1	91.5
Manufacturing	36,375	57.6	1.6	74.8
Retail	71,797	28.6	2.4	75.4
Other Activities	45,634	7	3.7	74.6
No Plants	49,470	6.3		

Table 3.2.1: Descriptive Statistics - Trade & Plant Data

Note: This table provide some description of the match between the trade panel and the population of plants for which we have geographic coordinates. The statistics are distributed across sectors (as defined in the *SIRENE* database). The last line corresponds to exporting firms for which we don't have any information on the coordinates of their plants. The first column gives the number of exporting firms in the trade panel. The second column gives displays the share of total export values that these firms represent. The third column gives the average number of plants with coordinates for each exporting firm. The fourth column contains the share of exporting firms that only own one plant (single-plant firms).

measured at the firm-level. However, pollution shocks vary across locations, thus across plants, within firm. Therefore, we cannot measure the exact exposition to PM2.5 pollution of multi-plants firms exports. We choose to aggregate plant-level exposure at the firm-level using a weighted average of PM2.5 concentrations across plants within a given firm. The *SIRENE* database provides information on employment at the plant level so we choose to use plant-level labor to weight exposure to PM2.5 concentrations.

Therefore, a firm i with employment L_i , composed of plants $j \in [1, N_i]$ with employments $\{L_j\}_{j \in [1, N_i]}$ (with $L_i = \sum_{j=1}^{N_i} L_j$), each located in *CHIMERE* grid cells c_j , is exposed to an average concentration of PM2.5 pollution in month t that is given by:

$$PM2.5_{it} = \frac{1}{L_i} \sum_{j=1}^{N_i} L_j PM2.5_{c_j t}, \quad (3.1)$$

with $PM2.5_{c_j t}$ the average PM2.5 concentration during month t in *CHIMERE* grid cells c_j . In the last section of the paper, we test the robustness of our estimates to restricting the sample to firms that only own a single plant. The last column of Table (3.2.1) indicates that such firms represent a large share of the population of exporting firms.

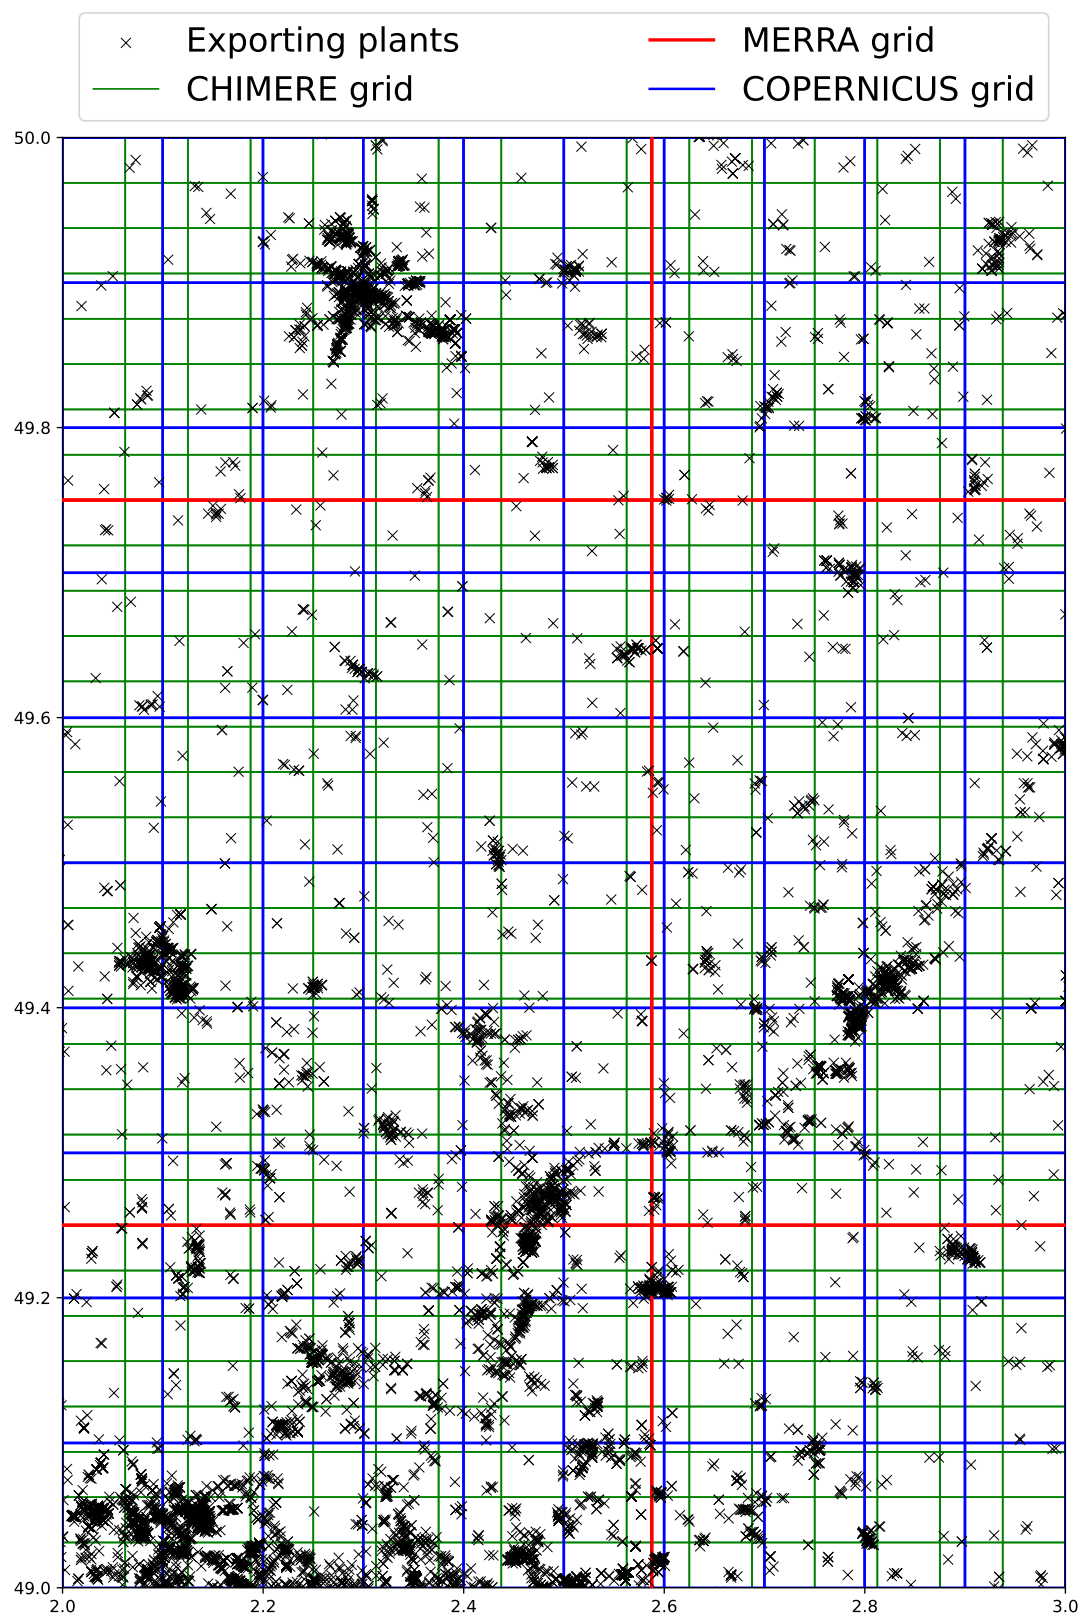


Figure 3.2.4: Illustration: Combination of Exporting Plants Coordinates, CHIMERE, MERRA, and COPERNICUS grids.

Note: This map illustrates, for a fraction of the French territory corresponding to the North of the Amiens' region, the process of superposition of the CHIMERE pollution grid, the MERRA and COPERNICUS weather data grids, and the coordinates of plants belonging to exporting firms.



Figure 3.2.5: Descriptive Statistics: Evolution of Monthly Exports and Number of Exporters
Note: The upper graph illustrates the evolution of the total value of monthly exports by French firms in billion euros. The lower graph illustrates the number of exporting firms each month in thousands of firms. Full lines correspond to the raw data contained in the customs file and the dashed lines ultimately plots the exports that are included in our estimations.




Figure 3.2.6: Plants Localization Data: Map of Plants Density
Note: This figure maps the local density of plants (in number) belonging to exporting firms across the CHIMERE grid. Based on the geographic coordinates of each plant, we determine in which of the cell it is located. This grid corresponds to the grid on which PM2.5 data is available.

3.3 Empirical Strategy

Our objective is to estimate the effect of exposure to PM2.5 pollution on firm-level exports, net of any confounding factors. In our baseline strategy, we estimate:

$$y_{it} = \alpha_i + \beta_t + \gamma \ln \overline{PM2.5}_{i,t} + X_{it} + \epsilon_{it}, \quad (3.2)$$

where y_{it} is either the logarithm of total value or quantity of exports made by firm i in month t , or the export status of firm i in month t (a dummy for whether the firm exported). The first two export variables are informative of the intensive margin effect of pollution and the third one reflects the extensive margin effect of pollution. Indeed, depending on the size of the effect of pollution shocks on firm productivity, the firm may either export less goods, which would lead to a negative effect on sales and quantities, or not export at all following the shock. In particular, if there are fixed shipping costs, firms may not be able to ship quantities of goods that are under some threshold. Equation (3.2) includes the effect of the logarithm of PM2.5 pollution averaged over the three months preceding the month when the export is registered. This allows to capture variation in the time gap between the production of the exported good and the date of export. Atmospheric pollution highly depends on weather variables that also affect economic activities, such as precipitations and temperatures. Equation (3.2) includes such weather control variables, denoted as X_{it} (and also averaged over the three preceding months).

Due to reverse causality, a simple OLS approach is likely to yield biased estimates. Indeed, if pollution may affect economics activities (the effect that we want to capture), economic activities also emit pollution. For instance, higher atmospheric concentrations of pollutants may be due to polluting manufacturing processes, energy production, or workers' commuting. As a result, there may be a positive correlation between y_{it} and $\ln \overline{PM2.5}_{i,t}$ that does not represent the causal effect of PM2.5 pollution on exports. This may lead to biased estimates in equation (3.2). To overcome this issue, we use an instrumental variable approach that relies on exogenous temporal variations in wind direction at the local level. This approach builds on previous work by Hansen-Lewis (2018) and Deryugina et al. (2019).¹³ The assumption behind this strategy is that, after controlling for confounding weather variables, wind direction only affects economic activities through its effect on PM2.5 concentrations.

The first step of our strategy is to estimate the following model for each *MERRA* grid

¹³Hering & Poncet (2014), Broner et al. (2012), and Bondy et al. (2020) are also examples that build on exogenous wind direction variation for causal inference.

cell:

$$\begin{aligned} \ln PM2.5_{cym} = & \mu_c + \nu_y + \rho_m + \sum_{k \in \{E, S, W\}} \beta^k WDir_{cym}^k \\ & + \alpha^R \ln Rain_{cym} + \alpha^H \ln Temp_{cym} + \epsilon_{cym}, \end{aligned} \quad (3.3)$$

where $PM2.5_{cym}$ is the average PM2.5 concentration measured in *CHIMERE* grid cell c , in year y and calendar month m . Thus, ym identifies a unique month during our period of analysis. *Rain* measures total precipitations during the month and *Temp* is the average monthly temperature. Variables $WDir^E$, $WDir^S$, and $WDir^W$ measure the number of hours during the month when wind was blowing eastward, southward, and westward. We omit the number of hours during which wind blows northward, which is the reference. As a result, coefficients β^E , β^S , and β^W represent the effect on the monthly average PM2.5 concentrations of having one hour blowing toward the East, the South, and the West, rather than toward the North. Equation (3.3) also includes *CHIMERE* grid cell fixed-effects, μ_c , that account for time-invariant local characteristics that influence local concentration of PM2.5. Geographic attributes, such as being located on a coast, a plain, or mountains, are controlled for by this term. We also include year fixed effects, ν_y , to capture overall trends and calendar month fixed effects, ρ_m , to account for seasonality. The three maps starting from the left of Figure 3.3.1 display the estimates of these coefficients. Grid cells are left blank when the local β coefficients are not significant at the 10% level.¹⁴ The fourth map, on the right of Figure 3.3.1, displays the R^2 coefficient of each regression. Overall, Figure 3.3.1 shows that wind direction is a good predictor of local PM2.5 pollution, even after controlling for confounding weather factors, seasonality effects, and time trends. For instance, Figure 3.3.1 shows that when the wind is blowing eastward, pollution is, on average, lower than when it is blowing northward: all β^E coefficients are negative. Considering that eastward winds come from the ocean, it is intuitive that it brings down the level of pollutant concentration relative to wind blowing toward the North.

The second step is to use the estimated model defined by equation (3.3) to compute predicted monthly pollution values for each *CHIMERE* grid cell, based on local wind direction:

$$\widehat{PM2.5}_{ct} = \exp \left(\sum_{k \in \{E, S, W\}} \hat{\beta}^k WDir_{ct}^k \right). \quad (3.4)$$

Using an exponential transformation we compute the corresponding values of predicted PM2.5 pollution, $\widehat{PM2.5}_{it}$, at the firm and month level using the same weighted average

¹⁴These cells are nonetheless kept in the panel.




Figure 3.3.1: Wind Instrument : Coefficients from the First-Stage

Note: Each map represents the distribution of estimated β coefficients from the first-stage regressions. The β coefficients measure the impact on local PM2.5 of one more hour blowing in each three direction rather than northward.

across plants as for observed pollution, and use the logarithm of $\widehat{PM2.5}_{it}$, also averaged over the preceding months, to instrument the actual firm-level PM2.5 concentration in equation (3.2). In the results section, we report first-stage F-statistics, which always indicate strong first-stages.

This approach, where $\widehat{PM2.5}_{it}$ is the excluded instrument, differs from the standard approach by Deryugina et al. (2019) where variables $WDir^E$, $WDir^S$, and $WDir^W$ are the excluded instruments. In practice, we run Deryugina et al. (2019)'s first-stage separately and independently for each location. This solves the issue created by multi-plants firms that are exposed to pollution shocks in distinct regions where wind direction potentially affects PM2.5 concentrations differently. However, in both approach, causal inference relies on the same assumptions that variations in local wind direction are exogenous and only affect exports through their effect on pollution, once meteorological factors are controlled for.

3.4 Results

3.4.1 Baseline Results

Table 3.4.1 presents the results of our baseline strategy. It reports estimates of the effects on exports of firm-level exposition to PM2.5 pollution averaged over the three months preceding the export, $\log \overline{PM2.5}_{t \in [-2,0]}$. In all regressions, we also control for weather variables averaged over the same time span as pollution. Columns (1), (3), and (5) report OLS estimates of the effect of firm-level PM2.5 shocks on exports total value, export quantity, and export status. Columns (2), (4), and (6) report estimates obtained through our 2SLS approach where we instrument actual firm-level PM2.5 exposition using predicted firm-level PM2.5 exposition.

Columns (2) shows that there is a causal and negative impact of PM2.5 concentrations

on the value of firm-level exports. A 10% increase in the average concentration of PM2.5 pollution in areas where exporting plants are located, over the three months preceding the export, leads to a 1.5% decrease in the total value of exports. Estimates of the effect of PM2.5 pollution on quantities exported, from column (4), reveal that the negative effect on total value is driven by a negative effect of pollution on quantities exported. In particular, a 10% increase in the average concentration of PM2.5 pollution in areas where exporting plants are located leads to a 1.4% decrease in the quantity exported. Estimates of the effect of PM2.5 pollution on the export status of exporting firms, in columns (6), is negative but not significant which implies that an increase in PM2.5 pollution does not affect the contemporaneous probability of exporting. OLS estimates, from columns (1), (3), and (5), of the contemporaneous effect of PM2.5 are consistently superior to 2SLS estimates. These biases can be explained by reversed causality: higher economics activities cause higher exports but also lead to more atmospheric pollution due to polluting activities. This weakens the negative correlation between pollution and exports.

	Value		Quantity		Export Status	
	OLS (1)	2SLS (2)	OLS (3)	2SLS (4)	OLS (5)	2SLS (6)
$\log \overline{PM2.5}_{t \in [-2,0]}$	-0.082*** (0.011)	-0.154** (0.068)	-0.068*** (0.012)	-0.141* (0.076)	-0.005*** (0.001)	-0.003 (0.007)
Obs.	2,797,719	2,797,719	2,766,587	2,766,587	9,969,875	9,969,875
FS F-Stat.		2,937		2,915		11,108

* p<0.1, ** p<0.05, *** p<0.01 - Standard errors, clustered by firm-activity, are reported in parenthesis. PM2.5 concentrations are measured in $\mu g/m^3$. The dependent variable is: for columns (1) and (2) the logarithm of export values in euros, for columns (3) and (4) the logarithm of the export quantities in kilograms, and for columns (5) and (6) a dummy variables indicating whether the firm exported in this given month. All regressions control for weather variables temperature and precipitation.

Table 3.4.1: Effects of PM2.5 Pollution on Firm-level Exports: Three Months Before the Export

Table C.0.2 in Appendix C.0.2 displays the result of estimating the baseline equation using a weighted regression where we use export values as weights. Point estimates from columns (2) and (4) confirm results from Table 3.4.1.

3.4.2 Longer-term effects of PM2.5

Because we do not know the timing between production and export, baseline results rely on an average exposition to pollution over three months. Table 3.4.2 displays the result of estimating equation (3.2) using values of exposition to pollution (and to control weather variables) over the six months that precede the export. If the time gap between production and export is larger than three months then the potential effects will be captured by this alternative approach.

Estimates from columns (2), (4), and (6) from table 3.4.2 provide a similar picture as those from table 3.4.1. In particular, firm-level exposition to PM2.5 pollution averaged over six months has a string and negative effect on export values (column (2): a 10% increase in average PM2.5 concentrations lead to a 2.4% decrease in the value of exports). The estimate for the effect of pollution on quantities is no longer significant (column (4)), but its scale is similar to its counterpart from 3.4.1. The effect of pollution on the probability of exporting (column (6)) is still not significant, but the point estimate is still negative and larger.

	Value		Quantity		Export Status	
	OLS (1)	2SLS (2)	OLS (3)	2SLS (4)	OLS (5)	2SLS (6)
$\log \overline{PM2.5}_{t \in [-5,0]}$	-0.092*** (0.015)	-0.238** (0.118)	-0.086*** (0.018)	-0.175 (0.132)	-0.009*** (0.002)	-0.018 (0.013)
Obs.	2,703,974	2,703,974	2,679,872	2,679,872	9,622,375	9,622,375
FS F-Stat.		2,273		2,251		7,732

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ - Standard errors, clustered by firm-activity, are reported in parenthesis. PM2.5 concentrations are measured in $\mu g/m^3$. The dependent variable is: for columns (1) and (2) the logarithm of export values in euros, for columns (3) and (4) the logarithm of the export quantities in kilograms, and for columns (5) and (6) a dummy variables indicating whether the firm exported in this given month. All regressions control for weather variables temperature and precipitation.

Table 3.4.2: Effects of PM2.5 Pollution on Firm-level Exports: Six Months Before the Export

One can try to disentangle the “short-term” effect of pollution, over the three months before the export, from its “long-term” effect, the three months before that. Table 3.4.3 presents estimates of the effect of pollution averages over the fifth, fourth, and third months before the export to investigate such longer terms effects. All 2SLS estimates, in columns (2), (4), and (6), are non significant, which implies that pollution mainly has short terms effects on exports. This may have two explanations: first, the literature on the effect of pollution shocks on workers points toward very short term effect (several days, weeks at most), and second, table ?? in appendix C.0.1 shows that most exports in our panel happen at one or two months intervals which suggest a similar timing between production and export.

	Value		Quantity		Export Status	
	OLS (1)	2SLS (2)	OLS (3)	2SLS (4)	OLS (5)	2SLS (6)
$\log PM2.5_{t \in [-5, -3]}$	-0.036*** (0.011)	-0.029 (0.063)	-0.032*** (0.012)	-0.016 (0.070)	-0.006*** (0.001)	-0.004 (0.007)
Obs.	2,698,780	2,698,780	2,674,723	2,674,723	9,604,444	9,604,444
FS F-Stat.		3,476		3,446		13,217

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ - Standard errors, clustered by firm-activity, are reported in parenthesis. PM2.5 concentrations are measured in $\mu g/m^3$. The dependent variable is: for columns (1) and (2) the logarithm of export values in euros, for columns (3) and (4) the logarithm of the export quantities in kilograms, and for columns (5) and (6) a dummy variables indicating whether the firm exported in this given month. All regressions control for weather variables temperature and precipitation.

Table 3.4.3: Effects of PM2.5 Pollution on Firm-level Exports: : Lagged Three-months Period

3.4.3 Heterogeneity Analysis

Using the eight-digits product code available for each transaction in the trade data, we aggregated exports at the firm-level across two-digits activities. Table C.0.3 in Appendix C.0.3 displays the effects of PM2.5 pollution on export total values, quantities, and export status for each activity. This heterogeneity analysis reveals that aggregated results on value and quantities from table 3.4.1 are mostly driven by negative impacts in Beverages, Other Manufacturing, and Wearing Apparels. Contrasting with the non-significant aggregated results on export status, table C.0.3 reveals negative effects of pollution on the probability to export of firms in the Beverages, Fabricated Metal Products, Machinery, and Other Manufacturing sectors. This heterogeneity of the impact of pollution on exports across activities may be explained by a number of factors such as labor intensity, whether the work is indoor or outdoor, air conditioning, or the degree of insulation of the buildings. These characteristics may differ across activities. For instance, production of electronic equipment is, in some cases (*e.g.* like microchips production), conducted in an airtight environment. In such cases, workers may not be very exposed to outdoor pollution. This could explain why the analysis does not detect any effect for these activities.

Whereas information on the type of environment under which production happens is hard to gather, we can investigate the effect of labor intensity. If the impact on exports of outdoor pollution goes through its impact on workers' productivity, the effect should be larger on firms that are more labor-intensive. We combine firm-level employment data with average shipment size (in quantity) to compute a labor intensity ratio for each exporting firm within our panel (total number of employees across all plants over the average export weight). Then, using the median value of this ratio, we split the population of firms into low and high labor intensity sub-populations: $LI \in \{Low, High\}$. We estimate our baseline strategy adding an interaction term between pollution and a

dummy for whether the exporting firm is among the 50% most labor-intensive firms.¹⁵ Table 3.4.4 reports the results. Estimates from the first row corresponds to estimates from Table 3.4.1. Estimates from the second row correspond to the additional effect of being a labor-intensive firm when pollution increases during the three months preceding an export. The estimates are all negative and significant at the 1% level. This means that the negative effect of pollution on export values and quantities is stronger for more labor-intensive firms. This result supports the story of pollution affecting export through its effect on workers and production. In the baseline approach, pollution has no effect on the probability of exporting. Estimates from columns (6) in Table 3.4.4 confirms this for less labor-intensive firms but indicates that more labor-intensive firms are less likely to export following an increase in pollution.

	Value		Quantity		Export Status	
	OLS (1)	2SLS (2)	OLS (3)	2SLS (4)	OLS (5)	2SLS (6)
$\log PM2.5_{t \in [-2,0]}$	-0.089*** (0.011)	-0.149** (0.068)	-0.075*** (0.013)	-0.135* (0.076)	-0.002 (0.002)	0.004 (0.008)
$\times \mathbb{I}(LI = High)$	0.046*** (0.011)	-0.086*** (0.028)	0.046*** (0.011)	-0.103*** (0.031)	-0.005*** (0.001)	-0.017*** (0.002)
Obs.	2,797,290	2,797,290	2,766,587	2,766,587	9,922,320	9,922,320
FS F-Stat.		1,362		1,349		5,517

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ - Standard errors, clustered by firm-activity, are reported in parenthesis. PM2.5 concentrations are measured in $\mu g/m^3$. The dependent variable is: for columns (1) and (2) the logarithm of export values in euros, for columns (3) and (4) the logarithm of the export quantities in kilograms, and for columns (5) and (6) a dummy variables indicating whether the firm exported in this given month. All regressions control for weather variables temperature and precipitation.

Table 3.4.4: Effects of PM2.5 Pollution on Firm-level Exports: Heterogeneity Analysis - Labor Intensity

3.4.4 Robustness Tests

Restricting to Single-plant Firms: In multi-plants firms, each plant may specialize in the production of a specific set of goods. As a result, the goods exported in t may be associated with only one location. The average pollution shock across plants is thus a noisy signal of the pollution shock experienced at the plant where the exported good was produced. This limitation could also bias the estimates if the participation of some plants to a given export is correlated with the exposition of this plant to PM2.5 pollution. For instance, if a plant is exposed to high PM2.5 concentrations during a given month and reduces its output, then the firm-level export will be composed of goods produced

¹⁵The excluded instruments are the predicted pollution as in the baseline strategy and the predicted pollution interacted with the labor-intensity dummy.

by other plants owned by the same firm. To test the implications of this issue we run a robustness test including only single-plant firms. Table 3.4.5 displays the estimates of the effect of PM2.5 on exports for this subsample of exports. Columns (2) and (4) of Table 3.4.5 confirm our main findings for the negative effect on export total value and quantity of PM2.5 pollution.

	Value		Quantity		Export Status	
	OLS (1)	2SLS (2)	OLS (3)	2SLS (4)	OLS (5)	2SLS (6)
$\log \overline{PM2.5}_{t \in [-2,0]}$	-0.077*** (0.013)	-0.186** (0.081)	-0.058*** (0.015)	-0.164* (0.091)	-0.007*** (0.002)	-0.009 (0.008)
Obs.	1,765,448	1,765,448	1,743,679	1,743,679	7,045,560	7,045,560
FS F-Stat.		2,450		2,423		9,060

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ - Standard errors, clustered by firm-activity, are reported in parenthesis. PM2.5 concentrations are measured in $\mu g/m^3$. The dependent variable is: for columns (1) and (2) the logarithm of export values in euros, for columns (3) and (4) the logarithm of the export quantities in kilograms, and for columns (5) and (6) a dummy variables indicating whether the firm exported in this given month. All regressions control for weather variables temperature and precipitation.

Table 3.4.5: Effects of PM2.5 Pollution on Firm-level Exports: Restricting to Single-Plant Firms

Controlling for the effects of other harmful atmospheric pollutants: The atmospheric concentration of PM2.5 is often correlated with that of other pollutants. Some of those are also harmful to human health and as such may also have an effect on firm-level activities and exports. A concern regarding the interpretation of the estimates from Table 3.4.1 is that they may in fact measure the aggregated effect of several other pollutants. This is the case even in the 2SLS approach as co-pollutants are also likely to be affected by wind direction. In Table 3.4.6, we control for the contemporaneous firm-level concentrations of two other pollutants: O3 and NO2. These co-pollutants are also instrumented using their predicted values based on wind direction. Estimates from the first row of columns (2), (4), and (6) of Table 3.4.6 show that the estimates of Table 3.4.1 are robust to the inclusion of other atmospheric pollutants.

Placebo test: Table 3.4.7 presents the result of a placebo test. We estimate the effect of pollution averaged over the three months after the export has been made on the value, quantity, and export status. The fact that the estimates on value and quantities are not significant is reassuring on the fact that baseline results are not statistical artifacts.

	Value		Quantity		Export Status	
	OLS (1)	2SLS (2)	OLS (3)	2SLS (4)	OLS (5)	2SLS (6)
$\log \overline{PM2.5}_{t \in [-2,0]}$	-0.064*** (0.012)	-0.438*** (0.096)	-0.062*** (0.014)	-0.522*** (0.107)	-0.014*** (0.002)	-0.001 (0.011)
$\log \overline{O3}_{t \in [-2,0]}$	0.012 (0.010)	0.137 (0.091)	0.006 (0.011)	0.251** (0.106)	-0.010*** (0.001)	-0.023** (0.011)
$\log \overline{NO2}_{t \in [-2,0]}$	-0.046*** (0.017)	0.593*** (0.157)	-0.015 (0.019)	0.746*** (0.177)	0.022*** (0.002)	0.007 (0.019)
Obs.	2,797,719	2,797,719	2,766,587	2,766,587	9,969,875	9,969,875
FS F-Stat.		350		347		1,092

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ - Standard errors, clustered by firm-activity, are reported in parenthesis. PM2.5 concentrations are measured in $\mu g/m^3$. The dependent variable is: for columns (1) and (2) the logarithm of export values in euros, for columns (3) and (4) the logarithm of the export quantities in kilograms, and for columns (5) and (6) a dummy variables indicating whether the firm exported in this given month. All regressions control for weather variables temperature and precipitation.

Table 3.4.6: Effects of PM2.5 Pollution on Firm-level Exports: Controlling for Co-pollutants

	Value		Quantity		Export Status	
	OLS (1)	2SLS (2)	OLS (3)	2SLS (4)	OLS (5)	2SLS (6)
$\log \overline{PM2.5}_{t \in [+1,+3]}$	-0.001 (0.010)	0.065 (0.069)	0.012 (0.012)	0.094 (0.079)	-0.000 (0.001)	0.041*** (0.007)
Obs.	2,746,100	2,746,100	2,710,881	2,710,881	9,849,372	9,849,372
FS F-Stat.		2,846		2,829		11,015

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ - Standard errors, clustered by firm-activity, are reported in parenthesis. PM2.5 concentrations are measured in $\mu g/m^3$. The dependent variable is: for columns (1) and (2) the logarithm of export values in euros, for columns (3) and (4) the logarithm of the export quantities in kilograms, and for columns (5) and (6) a dummy variables indicating whether the firm exported in this given month. All regressions control for weather variables temperature and precipitation.

Table 3.4.7: Effects of Future PM2.5 Pollution on Firm-level Exports: a Placebo Test

3.5 Conclusion

Air pollution has large adverse effects on human health. Understanding whether such effects are combined with other negative economic outcomes is essential to motivate the introduction of policies that enhance local air quality. In this paper, we use very disaggregated data, both spatially and temporally, on firm-level exports and PM2.5 atmospheric concentrations to investigate whether this type of atmospheric pollutant affects firm-level exports.

Our analysis indicates that firms export less in terms of both value and quantity when their plants are exposed to higher PM2.5 concentrations during the three months before the export. To provide causal estimates and overcome reverse causality biases, we rely

on an instrumental variable approach where we predict local PM2.5 concentrations using local temporal variation in wind direction.

After extending the strategy to longer-term effects of PM2.5 over six months before the export, we find that the negative impact of pollution is mostly driven by effects of higher PM2.5 concentrations during the three months before the export. We investigate heterogeneity in the effects of pollution between low and high labor intensity firms and observe that the negative effect of pollution is stronger for more labor-intensive firms. In comparison to the rest of the population, these more labor-intensive are also less likely to export after facing higher PM2.5 pollution.

Bibliography

- Acemoglu, D. & Autor, D. (2012). Skills, tasks and technologies.
- Adao, R., Arkolakis, C., & Esposito, F. (2019). *General equilibrium effects in space: Theory and measurement*. Technical report, National Bureau of Economic Research.
- Adão, R., Kolesár, M., & Morales, E. (2019). Shift-Share Designs: Theory and Inference. *The Quarterly Journal of Economics*, 134(4), 1949–2010.
- Addoum, J. M., Ng, D. T., & Ortiz-Bobea, A. (2020). Temperature Shocks and Establishment Sales. *The Review of Financial Studies*, 33(3), 1331–1366.
- Ademe (2021). Transition(s) 2050. choisir maintenant. agir pour le climat.
- Adhvaryu, A., Kala, N., & Nyshadham, A. (2022). Management and Shocks to Worker Productivity. *Journal of Political Economy*, 130(1), 1–47.
- AEE (2013). Améliorer la qualité de l’air en Europe. *Signaux de l’AEE*.
- AEE (2020). L’environnement en EuropeÉtat et perspectives 2020. *Rapport de l’Agence européenne de l’environnement*.
- Aguilar-Gomez, S., Dwyer, H., Zivin, J. S. G., & Neidell, M. (2022). This is Air: The "Non-Health" Effects of Air Pollution.
- Aguilar-Gomez, S., Gutierrez, E., Heres, D., Jaume, D., & Tobal, M. (2021). Extreme Temperatures and Loan Defaults: Evidence from Bank-Firm Level Data in Mexico. *SSRN Electronic Journal*.
- Agyeman, J., Schlosberg, D., Craven, L., & Matthews, C. (2016). Trends and Directions in Environmental Justice: From Inequity to Everyday Life, Community, and Just Sustainabilities. *Annual Review of Environment and Resources*, 41(1), 321–340.
- Albouy, D., Behrens, K., Robert-Nicoud, F., & Seegert, N. (2019). The Optimal Distribution of Population across Cities. *Journal of Urban Economics*, 110, 102–113.
- Aldeco, L., Barrage, L., & Turner, M. A. (2019). Equilibrium Particulate Exposure. *IGC Working Paper C-89451-NOC-1*.
- Alfaro, Conconi, Fadinger, & Newman (2016). Do prices determine vertical integration?

- Allen, J. G., MacNaughton, P., Satish, U., Santanam, S., Vallarino, J., & Spengler, J. D. (2016). Associations of Cognitive Function Scores with Carbon Dioxide, Ventilation, and Volatile Organic Compound Exposures in Office Workers: A Controlled Exposure Study of Green and Conventional Office Environments. *Environmental Health Perspectives*, 124(6), 805–812.
- Allen, T. & Arkolakis, C. (2014). Trade and the Topography of the Spatial Economy. *Quarterly Journal of Economics*, 129(3), 1085–1140.
- Allen, T., Arkolakis, C., & Takahashi, Y. (2020). Universal Gravity. *Journal of Political Economy*, 128(2), 393–433.
- Almagro, M. & Domínguez-Iino, T. (2021). Location Sorting and Endogenous Amenities: Evidence from Amsterdam. *mimeo*.
- Almond, D. & Currie, J. (2011). Killing Me Softly: The Fetal Origins Hypothesis. *Journal of Economic Perspectives*, 25(3), 153–172.
- Alvarez, J. L. C. & Rossi-Hansberg, E. (2021). The Economic Geography of Global Warming. Working Paper 28466, National Bureau of Economic Research.
- Ambec, S., Cohen, M. A., Elgie, S., & Lanoie, P. (2013). The porter hypothesis at 20: Can environmental regulation enhance innovation and competitiveness? *Review of Environmental Economics and Policy*, 7(1), 2–22.
- Antràs, P., Fort, T. C., & Tintelnot, F. (2017). The Margins of Global Sourcing: Theory and Evidence from US Firms. *American Economic Review*, 107(9), 2514–2564.
- Archsmith, J., Heyes, A., & Saberian, S. (2018). Air Quality and Error Quantity: Pollution and Performance in a High-Skilled, Quality-Focused Occupation. *Journal of the Association of Environmental and Resource Economists*, 5(4), 827–863.
- Atalay, Hortacsu, & Syverson (2014). Vertical integration and input flows.
- Atkeson, A. & Burstein, A. (2008). Pricing-to-market, trade costs, and international relative prices. *American Economic Review*, 98(5), 1998–2031.
- Autor, D. H., Dorn, D., & Hanson, G. H. (2013). The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review*, 103(6), 2121–2168.
- Baier, S. L. & Bergstrand, J. H. (2009). Bonus vetus ols: A simple method for approximating international trade-cost effects using the gravity equation. *Journal of International Economics*, 77(1), 77–85.

- Baldwin, R. E. & Okubo, T. (2005). Heterogeneous firms, agglomeration and economic geography: spatial selection and sorting. *Journal of Economic Geography*, 6(3), 323–346.
- Bale, T. L., Baram, T. Z., Brown, A. S., Goldstein, J. M., Insel, T. R., McCarthy, M. M., Nemeroff, C. B., Reyes, T. M., Simerly, R. B., Susser, E. S., & Nestler, E. J. (2010). Early Life Programming and Neurodevelopmental Disorders. *Biological Psychiatry*, 68(4), 314–319.
- Banzhaf, S., Ma, L., & Timmins, C. (2019). Environmental Justice: The Economics of Race, Place, and Pollution. *Journal of Economic Perspectives*, 33(1), 185–208.
- Barrows, G. & Ollivier, H. (2021). Foreign demand, developing country exports, and CO2 emissions: Firm-level evidence from India. *Journal of Development Economics*, 149, 102587.
- Bartelme, D. (2018). Trade Costs and Economic Geography: Evidence from the U.S. *mimeo*.
- Bartik, T. J. (1991). Who Benefits from State and Local Economic Development Policies? *Upjohn Institute for Employment Research*.
- Bartram, S. M., Hou, K., & Kim, S. (2022). Real effects of climate policy: Financial constraints and spillovers. *Journal of Financial Economics*, 143(2), 668–696.
- Baum-Snow, N., Henderson, J. V., Turner, M. A., Zhang, Q., & Brandt, L. (2020). Does investment in national highways help or hurt hinterland city growth? *Journal of Urban Economics*, 115, 103124.
- Bergounhon, F., Lenoir, C., & Mejean, I. (2018). A guideline to French firm-level trade data. <http://www.isabellemejean.com/FrenchCustomsData.html>.
- Bergquist, L., Faber, B., Fally, T., Hoelzlein, M., Miguel, E., & Rodriguez-Clare, A. (2019). Scaling agricultural policy interventions: Theory and evidence from uganda. *Unpublished manuscript, University of California at berkeley*.
- Bernard, A. B., Redding, S. J., & Schott, P. K. (2007). Comparative Advantage and Heterogeneous Firms. *Review of Economic Studies*, 74(1), 31–66.
- Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, (pp. 841–890).

- Bombardini, M. & Li, B. (2020). Trade, pollution and mortality in China. *Journal of International Economics*, 125, 103321.
- Bondy, M., Roth, S., & Sager, L. (2020). Crime Is in the Air: The Contemporaneous Relationship between Air Pollution and Crime. *Journal of the Association of Environmental and Resource Economists*, 7(3), 555–585.
- Bonilla, J., Coria, J., & Sterner, T. (2017). Technical Synergies and Trade-Offs Between Abatement of Global and Local Air Pollution. *Environmental and Resource Economics*, 70(1), 191–221.
- Borck, R. & Schrauth, P. (2021). Population density and urban air quality. *Regional Science and Urban Economics*, 86, 103596.
- Borck, R. & Tabuchi, T. (2018). Pollution and city size: can cities be too small? *Journal of Economic Geography*, 19(5), 995–1020.
- Borusyak, K., Hull, P., & Jaravel, X. (2021). Quasi-Experimental Shift-Share Research Designs. *The Review of Economic Studies*.
- Bougon, P. & Lavergne, R. (2019). Des outils réglementaires et fiscaux pour réduire les émissions de polluants atmosphériques de l'industrie. *Annales des Mines - Responsabilité et environnement*, N 96(4), 81.
- Branger, F., Quirion, P., & Chevallier, J. (2017). Carbon Leakage and Competitiveness of Cement and Steel Industries Under the EU ETS: Much Ado About Nothing. *The Energy Journal*, 37(3).
- Bretschger, L. & Pittel, K. (2020). Twenty Key Challenges in Environmental and Resource Economics. *Environmental and Resource Economics*, 77(4), 725–750.
- Broner, F., Bustos, P., & Carvalho, V. (2012). Sources of Comparative Advantage in Polluting Industries.
- Brook, R. D., Rajagopalan, S., Pope, C. A., Brook, J. R., Bhatnagar, A., Diez-Roux, A. V., Holguin, F., Hong, Y., Luepker, R. V., Mittleman, M. A., Peters, A., Siscovick, D., Smith, S. C., Whitsel, L., & Kaufman, J. D. (2010). Particulate Matter Air Pollution and Cardiovascular Disease. *Circulation*, 121(21), 2331–2378.
- Cai, J. & Szeidl, A. (2022). *Indirect Effects of Access to Finance*. Technical report, National Bureau of Economic Research.

- Calel, R. (2020). Adopt or Innovate: Understanding technological responses to cap-and-trade. *American Economic Journal: Economic Policy*, 12(3), 170–201.
- Calel, R. & Dechezlepretre, A. (2016). Environmental policy and directed technological change: Evidence from the european carbon market. *Review of Economics and Statistics*, 98(1), 173–191.
- Calligaris, S., D’Arcangelo, F. M., & Pavan, G. (2018). *The impact of the european carbon market on firm productivity: Evidence from italian manufacturing firms*. Technical report, Working Paper.
- Card, D. (2001). Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration. *Journal of Labor Economics*, 19(1), 22–64.
- Card, D. & Krueger, A. B. (2000). Minimum Wages and Employment: a Case Study of the Fast-Food Industry in New Jersey and Pennsylvania: Reply. *American Economic Review*, 90(5), 1397–1420.
- Carozzi, F. & Roth, S. (2019). Dirty Density: Air Quality and the Density of American Cities . *CEP Discussion Papers. Centre for Economic Performance, LSE, London, UK*.
- CASD, D. (2019). *DONNEES STATISTIQUES DU COMMERCE EXTERIEUR: Généralités sur les Fichiers*. Technical report, RÉSERVES GÉNÉRALES DE DIFFUSION STATISTIQUES – Comité du secret.
- Chang, T., Zivin, J. G., Gross, T., & Neidell, M. (2016). Particulate Pollution and the Productivity of Pear Packers. *American Economic Journal: Economic Policy*, 8(3), 141–169.
- Chang, T. Y., Zivin, J. G., Gross, T., & Neidell, M. (2019). The Effect of Pollution on Worker Productivity: Evidence from Call Center Workers in China. *American Economic Journal: Applied Economics*, 11(1), 151–172.
- Chay, K. Y. & Greenstone, M. (2005). Does Air Quality Matter? Evidence from the Housing Market. *Journal of Political Economy*, 113(2), 376–424.
- CITEPA (2020). Rapport OMINEA – 17ème édition. <https://www.citepa.org/fr/omine/>.
- CITEPA (2021). Gaz à effet de serre et polluants atmosphériques: Bilan des émissions en France de 1990 à 2020.
- CJUE (2022). Arrêt de la cour du 28 avril 2022.

- Coase, R. (1969). The problem of social cost", *journal of law and economics*, october, 1969.
- Colmer, J., Hardman, I., Shimshack, J., & Voorheis, J. (2020a). Disparities in PM 2.5 air pollution in the United States. *Science*, 369(6503), 575–578.
- Colmer, J., Martin, R., Muuls, M., & Wagner, U. (2019). Temperature and Production in a Globalized World: Evidence from French Manufacturing Firms. *mimeo*.
- Colmer, J. & Voorheis, J. (2020). The Grandkids Aren't Alright: The Intergenerational Effects of Prenatal Pollution Exposure. *Working Paper CES-20-36, US Census Bureau, Center for Economic Studies*.
- Colmer, Jonathan and Martin, R., Muûls, M., & Wagner, U. J. (2020b). Does Pricing Carbon Mitigate Climate Change? Firm-Level Evidence from European Union Emissions Trading Scheme. *CRC TR 224*, Discussion Paper No. 232, Available at SSRN: <https://ssrn.com/abstract=3725482> or <http://dx.doi.org/10.2139/ssrn.3725482>.
- Combes, P.-P., Duranton, G., & Gobillon, L. (2008). Spatial wage disparities: Sorting matters! *Journal of Urban Economics*, 63(2), 723–742.
- Combes, P.-P., Duranton, G., Gobillon, L., Puga, D., & Roux, S. (2012). The Productivity Advantages of Large Cities: Distinguishing Agglomeration From Firm Selection. *Econometrica*, 80(6), 2543–2594.
- Combes, P.-P. & Gobillon, L. (2015). Chapter 5 - The Empirics of Agglomeration Economies. volume 5 of *Handbook of Regional and Urban Economics* (pp. 247–348). Elsevier.
- Copeland, B. R. & Taylor, M. S. (2004). Trade, Growth, and the Environment. *Journal of Economic Literature*.
- Costa, L. G., Cole, T. B., Dao, K., Chang, Y.-C., & Garrick, J. M. (2019). Developmental impact of air pollution on brain function. *Neurochemistry International*, 131, 104580.
- Crippa, M., Guizzardi, D., Muntean, M., Schaaf, E., Dentener, F., van Aardenne, J. A., Monni, S., Doering, U., Olivier, J. G. J., Pagliari, V., & Janssens-Maenhout, G. (2018). Gridded emissions of air pollutants for the period 1970–2012 within EDGAR v4.3.2. *Earth System Science Data*, 10(4), 1987–2013.
- Currie, J., Voorheis, J., & Walker, R. (2020). What Caused Racial Disparities in Particulate Exposure to Fall? New Evidence from the Clean Air Act and Satellite-Based Measures of Air Quality. *Working Paper CES-20-02*.

- Davis, L. W. (2011). The Effect of Power Plants on Local Housing Values and Rents. *Review of Economics and Statistics*, 93(4), 1391–1402.
- De Chaisemartin, C. & d’Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9), 2964–96.
- De Loecker, J. & Syverson, C. (2021). An industrial organization perspective on productivity. In *Handbook of Industrial Organization*, volume 4 (pp. 141–223). Elsevier.
- Dechezleprêtre, A., Nachtigall, D., & Venmans, F. (2018). The joint impact of the european union emissions trading system on carbon emissions and economic performance. *OECD Economics Department Working Papers*, 1515.
- Dechezleprêtre, A., Rivers, N., & Stadler, B. (2019). The economic cost of air pollution Evidence from Europe. *OECD Working Paper*, 1584.
- Dell, M., Jones, B. F., & Olken, B. A. (2014). What do we learn from the weather? the new climate-economy literature. *Journal of Economic Literature*, 52(3), 740–798.
- Demailly, D. & Quirion, P. (2006). CO2 abatement, competitiveness and leakage in the European cement industry under the EU ETS: grandfathering versus output-based allocation. *Climate Policy*, 6(1), 93–113.
- Deryugina, T., Heutel, G., Miller, N. H., Molitor, D., & Reif, J. (2019). The Mortality and Medical Costs of Air Pollution: Evidence from Changes in Wind Direction. *American Economic Review*, 109(12), 4178–4219.
- Deryugina, T., Miller, N., Molitor, D., & Reif, J. (2021). Geographic and Socioeconomic Heterogeneity in the Benefits of Reducing Air Pollution in the United States. *Environmental and Energy Policy and the Economy*, 2, 157–189.
- Desmet, K., Kopp, R. E., Kulp, S. A., Nagy, D. K., Oppenheimer, M., Rossi-Hansberg, E., & Strauss, B. H. (2021). Evaluating the Economic Cost of Coastal Flooding. *American Economic Journal: Macroeconomics*, 13(2), 444–486.
- DGEC (2016). Panorama Energie-Climat. *DICOM-DGEC/CHE/16188*.
- Di, Q., Wang, Y., Zanobetti, A., Wang, Y., Koutrakis, P., Choirat, C., Dominici, F., & Schwartz, J. D. (2017). Air Pollution and Mortality in the Medicare Population. *New England Journal of Medicine*, 376(26), 2513–2522.
- Diamond, R. (2016). The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980-2000. *American Economic Review*, 106(3), 479–524.

- Drut, M. & Mahieux, A. (2015). Correcting agglomeration economies: How air pollution matters. *HAL Working Papers hal-01007019*.
- Duranton, G. & Puga, D. (2004). Chapter 48 - Micro-Foundations of Urban Agglomeration Economies. volume 4 of *Cities and Geography* (pp. 2063–2117). Elsevier.
- Eaton, J., Kortum, S., & Sotelo, S. (2013). International Trade: Linking Micro and Macro. *Advances in Economics and Econometrics: Tenth World Congress, II*(Applied Economics), 329–370.
- Eeckhout, J. & Guner, N. (2015). Optimal Spatial Taxation: Are Big Cities too Small? *mimeo*.
- Eeckhout, J. & Hedtrich, C. (2021). Green Urbanization. *mimeo*.
- Ellerman, A. D., Marcantonini, C., & Zaklan, A. (2015). The european union emissions trading system: Ten years and counting. *Review of Environmental Economics and Policy*, 10(1), 89–107.
- European Environment Agency (2021). Health impacts of air pollution in Europe. *Briefing no. 19/2021*.
- Fabra, N. & Reguant, M. (2014). Pass-through of emissions costs in electricity markets. *American Economic Review*, 104(9), 2872–2899.
- Fajgelbaum, P. D. & Gaubert, C. (2020). Optimal Spatial Policies, Geography, and Sorting. *The Quarterly Journal of Economics*, 135(2), 959–1036.
- Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *The American Economic Review*, (pp. 157–177).
- Fontagné, L., Martin, P., & Orefice, G. (2018). The international elasticity puzzle is worse than you think. *Journal of International Economics*, 115, 115–129.
- Fowlie, M., Reguant, M., & Ryan, S. P. (2016). Market-based emissions regulation and industry dynamics. *Journal of Political Economy*, 124(1), 249–302.
- Fowlie, M., Rubin, E., & Walker, R. (2019). Bringing Satellite-Based Air Quality Estimates Down to Earth. *AEA Papers and Proceedings*, 109, 283–288.
- Franklin, S., Imbert, C., Abebe, G., & Mejia-Mantilla, C. (2021). Urban public works in spatial equilibrium: Experimental evidence from ethiopia.

- Fu, S., Viard, V. B., & Zhang, P. (2021). Air Pollution and Manufacturing Firm Productivity: Nationwide Estimates for China. *The Economic Journal*, forthcoming.
- Ganapati, S., Shapiro, J. S., & Walker, R. (2020). Energy Cost Pass-Through in US Manufacturing: Estimates and Implications for Carbon Taxes. *American Economic Journal: Applied Economics*, 12(2), 303–42.
- Gaubert, C. (2018). Firm Sorting and Agglomeration. *American Economic Review*, 108(11), 3117–3153.
- Gaulier, G. & Zignago, S. (2010). BACI: International Trade Database at the Product-Level. The 1994-2007 Version. *CEPII research center Working Papers*, 23.
- Gibson, M. (2019). Regulation-induced pollution substitution. *Review of Economics and Statistics*, 101(5), 827–840.
- Giroud, X. & Mueller, H. M. (2019). Firms’ Internal Networks and Local Economic Shocks. *American Economic Review*, 109(10), 3617–49.
- Glaeser, E. L. & Kahn, M. E. (2010). The greenness of cities: Carbon dioxide emissions and urban development. *Journal of Urban Economics*, 67(3), 404–418.
- Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2020). Bartik Instruments: What, When, Why, and How. *American Economic Review*, 110(8), 2586–2624.
- Graff Zivin, J. & Neidell, M. (2012). The Impact of Pollution on Worker Productivity. *American Economic Review*, 102(7), 3652–3673.
- Graff Zivin, J. & Neidell, M. (2013). Environment, Health, and Human Capital. *Journal of Economic Literature*, 51(3), 689–730.
- Graff Zivin, J. & Neidell, M. (2014). Temperature and the Allocation of Time: Implications for Climate Change. *Journal of Labor Economics*, 32(1), 1–26.
- Greenstone, M. (2002). The impacts of environmental regulations on industrial activity: Evidence from the 1970 and 1977 clean air act amendments and the census of manufactures. *Journal of political economy*, 110(6), 1175–1219.
- Hanlon, W. W. (2019). Coal Smoke, City Growth, and the Costs of the Industrial Revolution. *The Economic Journal*, 130(626), 462–488.
- Hanna, R. & Oliva, P. (2015). The effect of pollution on labor supply: Evidence from a natural experiment in Mexico City. *Journal of Public Economics*, 122, 68–79.

- Hansen-Lewis, J. (2018). Does Air Pollution Lower Productivity? Evidence from Manufacturing in India. *mimeo*.
- Harrigan, J., Reshef, A., & Toubal, F. (2018). *Techies, trade, and skill-biased productivity*. Technical report, National Bureau of Economic Research.
- Hauvuy, E. & Riedinger, N. (2005). Une estimation du coût d’abattement de la pollution atmosphérique pour les entreprises françaises. *Économie & prévision*, 168(2), 63.
- HCC (2020). Maîtriser l’empreinte carbone de la France. *Haut Conseil pour le Climat: Réponse à la saisine du gouvernement*.
- He, J., Liu, H., & Salvo, A. (2019). Severe Air Pollution and Labor Productivity: Evidence from Industrial Towns in China. *American Economic Journal: Applied Economics*, 11(1), 173–201.
- Head, K. & Mayer, T. (2021). *Poor Substitutes? Counterfactual methods in IO and Trade compared*. Technical report, CEPII research center.
- Helpman, E. (1998). The Size of Regions. *Topics in Public Economics. Theoretical and Applied Analysis*, (pp. 33–54).
- Helpman, E., Melitz, M., & Rubinstein, Y. (2008). Estimating trade flows: trading partners and trading volumes. *Quarterly Journal of Economics*, 123(2), 441–487.
- Henderson, J. V. (1974). The Sizes and Types of Cities. *The American Economic Review*.
- Hering, L. & Poncet, S. (2014). Environmental policy and exports: Evidence from Chinese cities. *Journal of Environmental Economics and Management*, 68(2), 296–318.
- Heyes, A., Neidell, M., & Saberian, S. (2016). The Effect of Air Pollution on Investor Behavior: Evidence from the S&P 500. *NBER Working Paper 22753*.
- Hintermann, B. (2017). Market power in emission permit markets: Theory and evidence from the EU ETS. *Environmental and Resource Economics*, 66(1), 89–112.
- Holub, F., Hospido, L., & Wagner, U. J. (2021). Urban Air Pollution and Sick Leaves: Evidence from Social Security Data. *SSRN Electronic Journal*.
- Hsieh, C.-T. & Ossa, R. (2016). A global view of productivity growth in China. *Journal of International Economics*, 102, 209–224.
- IGF (2018). Évaluation de l’impact environnemental et économique de la Taxe générale sur les activités polluantes (TGAP) sur les émissions de polluants atmosphériques. *Rapport IGF n 2018-M-027-03*.

- INS (2020). Inventaire national spatialisé des émissions de polluants dans l'air. <http://emissions-air.developpement-durable.gouv.fr>.
- IPCC (2021a). Climate Change 2021: The Physical Science Basis - Summary for Policy-makers.
- IPCC (2021b). COMMUNIQUÉ DE PRESSE - Changement climatique généralisé et rapide, d'intensité croissante – GIEC.
- IPCC (2022). Climate Change 2022: Mitigation of Climate Change - Summary for Policymakers.
- Isen, A., Rossin-Slater, M., & Walker, W. R. (2017). Every Breath You Take—Every Dollar You'll Make: The Long-Term Consequences of the Clean Air Act of 1970. *Journal of Political Economy*, 125(3), 848–902.
- Jaraite, J. & Di Maria, C. (2018). Did the eu ets make a difference? an empirical assessment using lithuanian firm-level data. *The Energy Journal*.
- Joltreau, E. & Sommerfeld, K. (2018). Why Does Emissions Trading Under the EU ETS not Affect Firms' Competitiveness? Empirical Findings from the Literature. *Climate Policy*, 19(4), 453–471.
- Jones, B. F. & Olken, B. A. (2010). Climate Shocks and Exports. *American Economic Review*, 100(2), 454–459.
- Karlsson, J. (2021). Temperature and Exports: Evidence from the United States. *Environmental and Resource Economics*, 80(2), 311–337.
- Khanna, G., Liang, W., Mobarak, A. M., & Song, R. (2021). The Productivity Consequences of Pollution-Induced Migration in China. *National Bureau of Economic Research - Working Paper 28401*.
- Kyriakopoulou, E. & Xepapadeas, A. (2013). Environmental policy, first nature advantage and the emergence of economic clusters. *Regional Science and Urban Economics*, 43(1), 101–116.
- Lafrogne-Joussier, R., Martin, J., & Mejean, I. (2022). Supply shocks in supply chains: Evidence from the early lockdown in China. *mimeo*.
- Lange, A. & Quaas, M. F. (2007). Economic Geography and the Effect of Environmental Pollution on Agglomeration. *The B.E. Journal of Economic Analysis & Policy*, 7(1).

- Leturque, E. & Sanch-Maritan, M. (2019). Do agglomeration economies are lower for polluting sectors? *HAL Working Papers hal-02347595*.
- Li, C., Cong, J., & Yin, L. (2021). Extreme heat and exports: Evidence from Chinese exporters. *China Economic Review*, 66, 101593.
- Löschel, A., Lutz, B. J., & Managi, S. (2018). The impacts of the eu ets on efficiency and economic performance – an empirical analyses for german manufacturing firms. *Resource and Energy Economics*.
- Malgouyres, C. & Mayer, T. (2018). Exports and labor costs: evidence from a French policy. *Review of World Economics*, 154(3), 429–454.
- Marin, G. & Vona, F. (2021). The Impact of Energy Prices on Socioeconomic and Environmental Performance: Evidence from French Manufacturing Establishments, 1997-2015. *European Economic Review*, 135.
- Martin, L. A., Nataraj, S., & Harrison, A. E. (2017). In with the big, out with the small: Removing small-scale reservations in india. *American Economic Review*, 107(2), 354–86.
- Martin, R., Muûls, M., & Wagner, U. J. (2015). 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*, 0(0), 1–21.
- Massard-Guilbaud, G. (1999). La régulation des nuisances industrielles urbaines. *Vingtième Siècle. Revue d'histoire*, (64), 53–65.
- Melitz, M. J. & Ottaviano, G. I. (2008). Market Size, Trade, and Productivity. *Review of Economic Studies*, 75, 295–316.
- Millock, K. & Nauges, C. (2006). Ex Post Evaluation of an Earmarked Tax on Air Pollution. *Land Economics*, 82(1), 68–84.
- Moretti, E. (2013). Real Wage Inequality. *American Economic Journal: Applied Economics*, 5(1), 65–103.
- Morgenstern, R. D., Pizer, W. A., & Shih, J.-S. (2002). Jobs Versus the Environment: An Industry-Level Perspective. *Journal of Environmental Economics and Management*, 43(3), 412–436.
- Muehlegger, E. & Sweeney, R. L. (2021). Pass-through of own and rival cost shocks: Evidence from the us fracking boom. *Review of Economics and Statistics*.

- Nedoncelle, C. (2021). Temperatures, Heterogeneous Exporters and Aggregate Trade. *mimeo*.
- Neumark, D. & Simpson, H. (2015). Place-based policies. In *Handbook of regional and urban economics*, volume 5 (pp. 1197–1287). Elsevier.
- Newell, R. G. & Stavins, R. N. (2003). Cost Heterogeneity and the Potential Savings from Market-Based Policies. *Journal of Regulatory Economics*, 23(1), 43–59.
- Nunn, N. & Qian, N. (2014). US Food Aid and Civil Conflict. *American Economic Review*, 104(6), 1630–1666.
- Petrack, S. & Wagner, U. J. (2014). The impact of carbon trading on industry: Evidence from german manufacturing firms. *Kiel Institute Working Papers*, 1912.
- Pflüger, M. (2021). City Size, Pollution and Emission Policies. *Journal of Urban Economics*, 126, 103391.
- Pigou, A. C. (1920). The Economics of Welfare. *Macmillan and Co., London*.
- Porter, M. E. & van der Linde, C. (1995). Toward a new conception of the environment-competitiveness relationship. *Journal of Economic Perspectives*, 9(4), 97–118.
- Real, E., Couvidat, F., Ung, A., Malherbe, L., Raux, B., & Colette, A. (2021). Historical reconstruction of background air pollution over France for 2000–2015. *Earth Syst. Sci. Data Discuss.* [preprint].
- Redding, S. (2020). Trade and Geography. *NBER Working Paper 27821*.
- Redding, S. J. (2016). Goods trade, factor mobility and welfare. *Journal of International Economics*, 101, 148–167.
- Redding, S. J. & Rossi-Hansberg, E. (2017). Quantitative Spatial Economics. *Annual Review of Economics*, 9(1), 21–58.
- Ribière, G. (2010). La gestion des aménités et le développement durable. *Revue française d'administration publique*, 134(2), 373.
- Roback, J. (1982). Wages, Rents, and the Quality of Life. *Journal of Political Economy*, 90(6), 1257–1278.
- Rosen, S. (1979). *Wage-Based Indexes of Urban Quality of Life*, (pp. 74–104). Johns Hopkins University Press.

- Rotemberg, M. (2019). Equilibrium effects of firm subsidies. *American Economic Review*, 109(10), 3475–3513.
- 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*.
- Sato, K. (1976). The ideal log-change index number. *The Review of Economics and Statistics*, (pp. 223–228).
- Sato, M., Singer, G., Dussaux, D., & Lovo, S. (2019). International and sectoral variation in industrial energy prices 1995–2015. *Energy Economics*, 78, 235–258.
- Shapiro, J. S. (2022). Pollution Trends and US Environmental Policy: Lessons from the Past Half Century. *Review of Environmental Economics and Policy*, 16(1), 42–61.
- Shapiro, J. S. & Walker, R. (2018). Why Is Pollution from US Manufacturing Declining? The Roles of Environmental Regulation, Productivity, and Trade. *American Economic Review*, 108(12), 3814–3854.
- Shapiro, J. S. & Walker, R. (2021). Where Is Pollution Moving? Environmental Markets and Environmental Justice. *AEA Papers and Proceedings*, 111, 410–414.
- Sicard, P., Agathokleous, E., Marco, A. D., Paoletti, E., & Calatayud, V. (2021). Urban population exposure to air pollution in Europe over the last decades. *Environmental Sciences Europe*, 33(1).
- Smagghue, G. (2019). Estimating the cost of regulation thresholds : A structural approach. *mimeo*.
- Soliman, F. (2020). Intrafirm leakage.
- Sraer, D. & Thesmar, D. (2020). How to use natural experiments to estimate misallocation. *SSRN available at ssrn.com/abstract=3099007*.
- Suarez-Serrato, J. C. & Zidar, O. (2016). Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms. *American Economic Review*, 106(9), 2582–2624.
- Thomas, A. (2021). ‘Heart of steel’: how trade unions lobby the European Union over emissions trading. *Environmental Politics*, 30(7), 1217–1236.
- UNEP (2021). The Heat Is On - A world of climate promises not yet delivered - Executive Summary. *Emissions Gap Report*.

- US EPA (2022). Air Quality - National Summary.
- Vartia, Y. O. (1976). Ideal log-change index numbers. *scandinavian Journal of statistics*, (pp. 121–126).
- Vette, A. F., Rea, A. W., Lawless, P. A., Rodes, C. E., Evans, G., Highsmith, V. R., & Sheldon, L. (2001). Characterization of Indoor-Outdoor Aerosol Concentration Relationships during the Fresno PM Exposure Studies. *Aerosol Science and Technology*, 34(1), 118–126.
- Viehmann, A., Hertel, S., Fuks, K., Eisele, L., Moebus, S., Möhlenkamp, S., Nonnemacher, M., Jakobs, H., Erbel, R., Jöckel, K.-H., & Hoffmann, B. (2015). Long-term residential exposure to urban air pollution, and repeated measures of systemic blood markers of inflammation and coagulation. *Occupational and Environmental Medicine*, 72(9), 656–663.
- Waidyatillake, N. T., Campbell, P. T., Vicendese, D., Dharmage, S. C., Curto, A., & Stevenson, M. (2021). Particulate Matter and Premature Mortality: A Bayesian Meta-Analysis. *International Journal of Environmental Research and Public Health*, 18(14), 7655.
- WHO (2014). 7 million premature deaths annually linked to air pollution.
- WHO (2016). WHO releases country estimates on air pollution exposure and health impact. <https://www.who.int/news/item/27-09-2016-who-releases-country-estimates-on-air-pollution-exposure-and-health-impact>.
- WHO (2021). Lignes directrices OMS relatives à la qualité de l'air: Particules (PM_{2,5} et PM₁₀), ozone, dioxyde d'azote, dioxyde de soufre et monoxyde de carbone - Résumé d'orientation.
- Yamada, E. (2020). A Spatial Equilibrium Analysis of Air Pollution in China. *JICA-RI Working Paper*.

Appendix A

Appendix to Chapter 1

A.1 Further Empirical Results

	Probit		OLS			
	everR (1)	everR (2)	everR (3)	everR (4)	maxR (5)	maxR (6)
$\ln y_{2004}$	0.614*** (0.030)	0.610*** (0.030)	.0041*** (0.00009)	.0039*** (0.00009)	.0035*** (0.00008)	.0024*** (0.00008)
$grwth_y$	0.220*** (0.075)	0.222*** (0.074)	.0035*** (0.00024)	.0034*** (0.00024)	.0029*** (0.00021)	.0028*** (0.00021)
$\alpha_{f,2004}$		0.282 (0.075)		0.0412*** (0.248)		0.0154*** (0.0019)
$\sum_{j \in \Omega_{f,2004}} \alpha_{jf,2004}^2$	0.336** (0.155)	0.899* (0.521)	-0.0167*** (0.0011)	0.0554*** (0.0039)	-0.0042*** (0.0009)	0.0227*** (0.0035)
# obs	25,871	25,871	139,950	139,950	139,950	139,950

Notes: Table reports probit (columns 1 and 2), and OLS (columns 3-6) estimation of correlations between firm-level characteristics in 2004 and future regulation. Outcome variable in columns 1-4 is indicator variable for whether a firm ever operates a regulated plant. Outcome variable in columns 5-6 is the maximum value of $R_{fist} \in [0, 1]$ observed for the firm. All regressions include industry fixed effects. Columns 1-2 include only industries with some ever-regulated firms. Columns 3-6 include all industries. Standard errors are clustered at the industry level.

Table A.1.1: Predicting Regulation Status

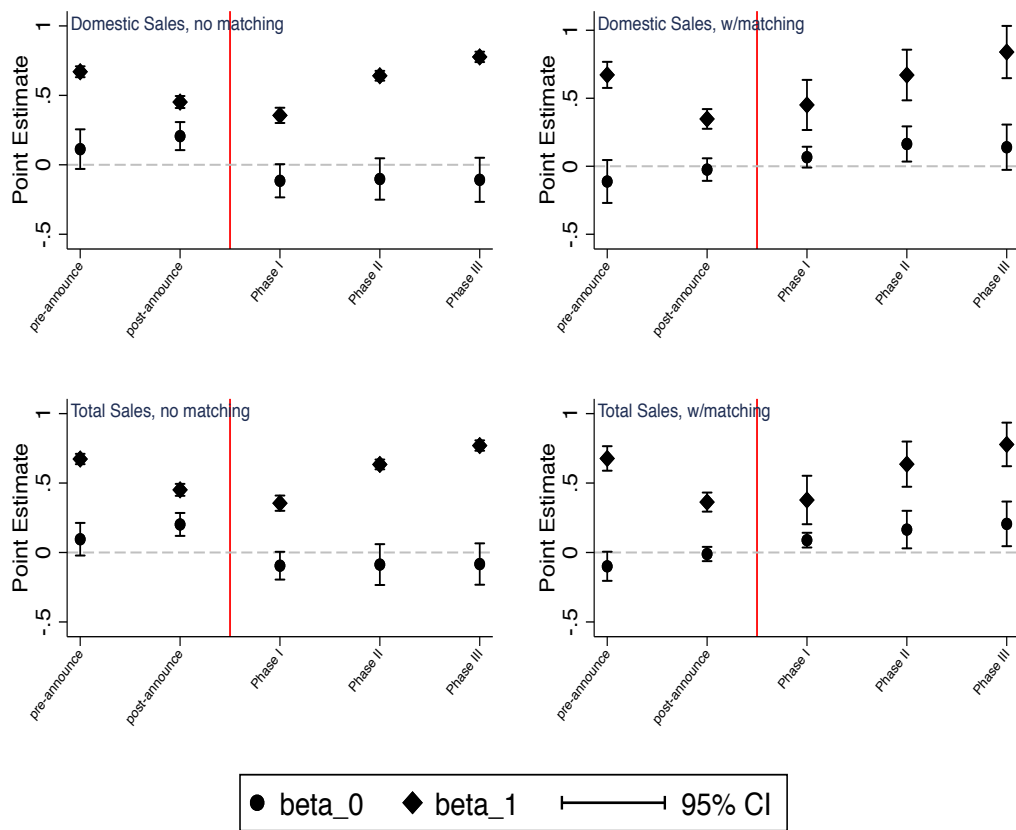


Figure A.1.1: Firm-level DD Estimates for Revenues

Notes: Markers indicate estimates of β_0 and β_1 derived from OLS estimation of DiD model described in section 1.4.3, which controls for $\Delta\alpha_{fit}$. Bars indicate 95% confidence intervals. For each period, we estimate a pooled regression, combining all years included in the period. Top (bottom) row takes log change in domestic (total) sales as the dependent variable. Right (left) panels include all (matched) firms. The vertical line indicates the beginning of the regulation period. All regressions include industry-by-year fixed effects. Standard errors are clustered on the industry

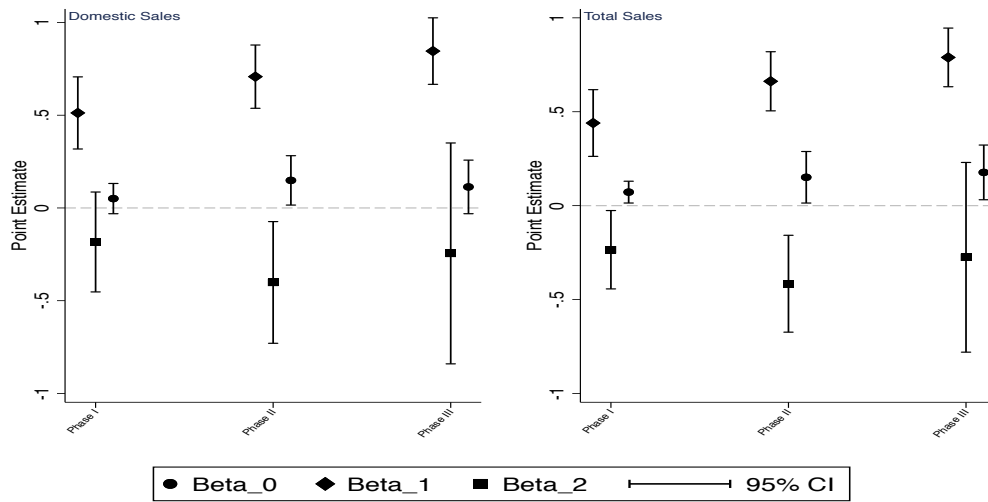


Figure A.1.2: Point Estimates, LA Excluding Foreign Imports

Notes: Markers indicate point estimates resulting from OLS estimation of equation 1.2, excluding foreign imports. Bars indicate 95% confidence intervals. For each period, we estimate a pooled regression, combining all years included in the period. All regressions include sector-by-year fixed effects. Left (right) panel takes log change in domestic (total) sales as the dependent variable. Standard errors are clustered on the industry.

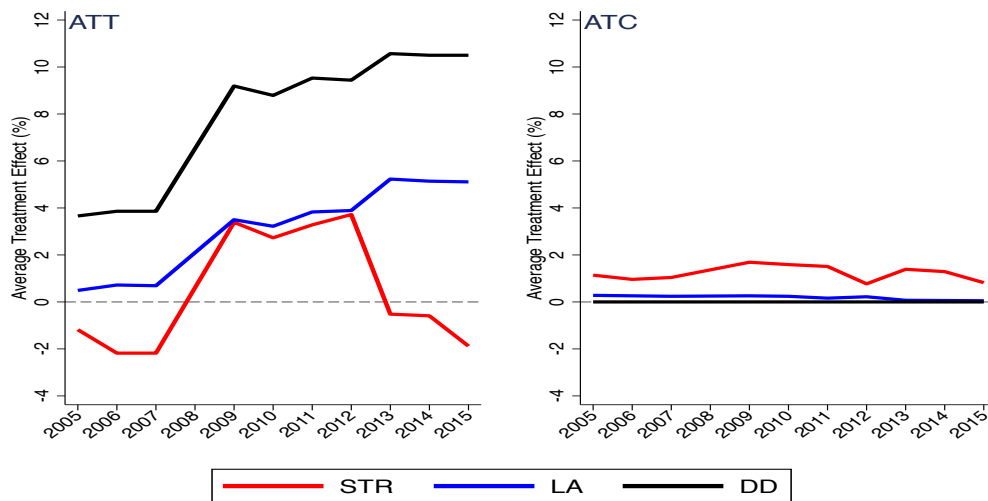


Figure A.1.3: ATEs on Revenues, Excluding Foreign Imports

Notes: Figure presents the average treatment effects on firm-level revenues for treated firms (left) and control firms (right) as a percent of observed firm-level revenues by year and estimator, excluding foreign imports. Treatment effects are averaged over all French firms in the manufacturing sector.

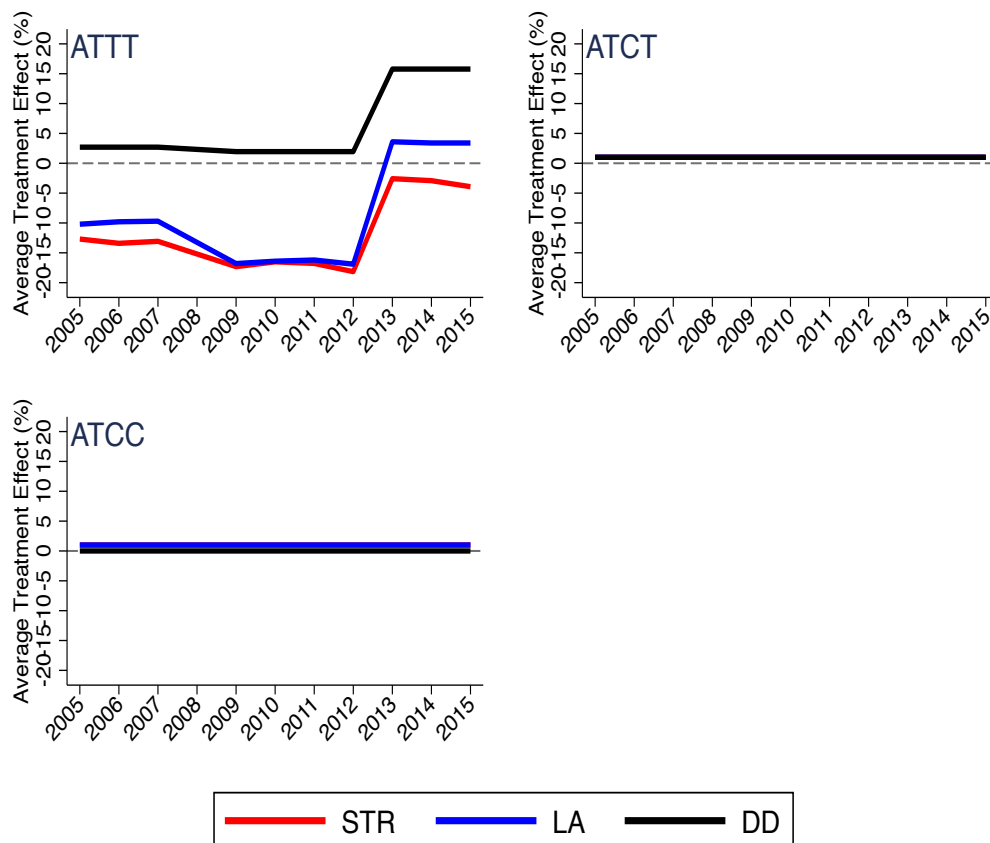


Figure A.1.4: ATEs on Emissions, Plant-level Excluding Foreign Imports

Notes: Figure presents the average treatment effects on plant-level emissions for TT plants (top left), CT plants (top right) and CC plants (bottom right) generated in the production of goods for the domestic market as a percent of observed plant-level emissions generated in the production of goods for the domestic market, excluding foreign imports. Average treatment effects are computed only for firms with observed emissions in the EACEI survey.

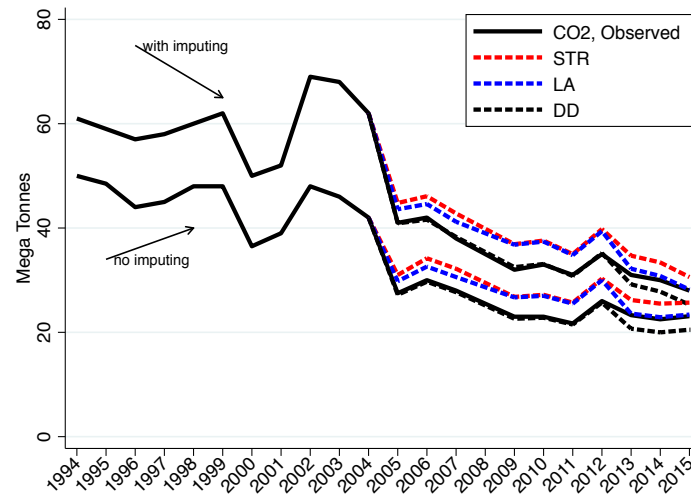


Figure A.1.5: Aggregate Emissions, Excluding Foreign Imports

Notes: Figure presents observed total CO₂ emissions from French manufacturers generated in the production of goods for the domestic market (solid black) along with counterfactual emissions computed using STR, LA and DD methods, excluding foreign imports. Bottom lines include only plants with observed emissions in the EACEI survey, while top lines impute emission intensity for plants with no emissions data.

A.2 Data Appendix

A.2.1 FICUS-FARE

Balance sheet information for the universe of French firms was retrieved from the FICUS (Annual structural statistics of companies from the SUSE scheme, 1994-2007) and FARE (Annual structural statistics of companies from the ESANE scheme, 2008-2016). These firm-level data originate from a fiscal source, since firms need to declare their profits to the tax authorities. Firm-level data from FICUS/FARE were linked to other data sets based on the unique identifier of French firms (SIREN).

We identify industries and sectors based on the activity code APE (*Activité Principale de l'Etablissement*), from the national activity nomenclature, namely the *Nomenclature des Activités Françaises* (NAF). These 4-digit codes correspond to an industry, whereas the first two digits of the codes are common within a sector. The NAF classification was revised in 2003 (from the “NAF93” to the “NAFRev.1”) and again in 2008 (to the “NAFRev.2”). If the conversion from the NAF93 to the NAFRev1 is straightforward, the revision in 2008 deeply modified the within-sector industry decomposition and resulted in a many-to-many mapping between NAFRev.1 and NAFRev.2. Since our analysis rests on a structure of industries and sectors, we must assign one activity code to each firm over time. We consider alternately the two revisions – NAFRev.1 and NAFRev.2. If a firm is active both before and after 2008, it is assigned two stable activity codes, one for each revision. Indeed, we assign to firms that switch industry codes their modal code within a revision. If a firm is not active in one of the period, it is either assigned a 1-to-1 match in nomenclatures (if available) or a code in the NAF classification that is observed most frequently for firms with the same industry code.

A.2.2 EACEI

The information on energy use comes from the EACEI (*Enquête sur les consommations d'énergie dans l'industrie*) survey. This is a survey of manufacturing establishments (identified by a unique identifier called SIRET, whose first 9 digits identify the SIREN of the firm) that provides information on the consumption (quantity and value) of energy, broken down by energy types: electricity (bought and self-generated), steam, natural gas, other types of gas, coal, lignite, coke, propane/butane, heavy fuel oil, heating oil and other petroleum products. All establishments with more than 250 employees receive the survey each year, whereas only a sample of establishments with 20 or more employees (stratified by industries, number of employees, and region) receive it. The response rate is nearly 90% (see Marin & Vona 2021).

We compute CO2 emissions at the plant level by combining the quantity of energy consumed with an energy-specific conversion factor that indicates the amount of CO2 released in the atmosphere when the type of energy considered is used. We thus focus on energy-related CO2 emissions released by each plant during production through the combustion of fossil energy or through the indirect emissions related to electricity bought from the grid or steam bought externally. We thus ignore any emissions released in the upstream stages of extraction, transformation and transportation of energy. We use CO2 emission factors for the different energy types from ADEME’s “Base Carbone”.¹ Whereas the conversion factors are constant over time for most energy types, given the proportionality between CO2 emissions and the carbon content of fossil fuels, the amount of CO2 emissions associated with the use of electricity from the grid is time varying and depends on the energy mix of the electricity sector in France. Table A.2.1 reports these conversion factors.

Fuel name	Unit	Conversion Factor
Coal	t	3.07
Lignite	t	1.72
Coke of coal	t	3.03
Petroleum coke	t	3.1
Gas from the grid	MWh	0.169
Non-natural Gas	MWh	0.469
Butane/Propane	t	2.965
Heavy fuel oil	t	3.14
Heating oil	L	0.00268
Steam	t	0.113
Electricity		
2012	MWh	0.0552
2013	MWh	0.0526
2014	MWh	0.0492
2015	MWh	0.0408
2016	MWh	0.0379

Notes: These conversion factors have been computed by the authors using ADEME’s “Base Carbone”. The units for fuels are in tons (t), megawatt hours (MWh) or liters (L). The conversion is toward tons of CO2.

Table A.2.1: CO2 Conversion Factors

To clean the data, we first identify whether there are within-plant observations that are 50 times smaller or larger than others over time. If so, we drop the entire history of plant CO2 emissions. Next, we identify within-plant year-to-year changes that are equivalent

¹We use CO2 conversion factors from the document published by ADEME entitled “Base Carbone: Documentation des Facteurs d’Emissions de la Base Carbone” in its version 11.2.0, available from March 2015.

to being multiplied by 10 or divided by 10. If so, we again drop the entire plant-level history.




Figure A.2.1: Share of Plants Covered by the EACEI Survey

Notes: Figure plots the shares of plants, labor share, and emissions shares of plants included in the EACEI survey by year. Emissions share is computed by imputing emissions for plants excluded from the EACEI survey.

A.2.3 EUTL

To assess the treatment status of each firm and plant, we use the European Union Transaction Log (EUTL), which is the central reporting tool for the EU ETS. Transactions are reported through the EUTL with a delay of three years. The data can be downloaded from the EUTL webpage (<https://ec.europa.eu/clima/ets/>). We downloaded two registries, one for 2014 and the other for 2018.

Within the EUTL, an installation is a regulated entity. It faces the obligation to surrender allowances at least equal to its verified emissions of the previous year to the regulatory authority. Installations either receive these allowances for free or buy them on the allowances market. To be able to receive, transfer, and surrender allowances, each installation must be represented by an operator holding account (OHA). For each account, the EUTL provides a primary contact for the account holder with an address. Each installation receives a unique identifier and a registry, which corresponds to its country of location. In some cases, the EUTL will provide information on companies related to the installation, that is on the SIREN in the French context. SIREN identifiers are usually reported for manufacturing industries, but not for the power sector and for heating units. To illustrate,

from the registry of 2014, we obtain 1256 account holders with SIREN information and 130 without. To complement the original file, we use the Ownership links and enhanced EUTL dataset from Jaraite, Jong, Kazukauskas, Zaklan and Zeitlberger (2016).² From the registry of 2018, we obtain 995 account holders with SIREN and 320 without, some of them being identical to the ones from the registry of 2014.

To assign a treatment status to each plant, we must match the EUTL installations with the list of SIRET owned by a firm (SIREN) in all manufacturing sectors. We first retrieve the list of plants and their locations (city code, address) from the “Base des Etablissements” from the French statistical agency INSEE. We then merge the installations from EUTL with this list using SIREN and city codes, retrieving roughly 600 unique matches and leaving 600 unmatched installations. Next, for the installations that match with several SIRET, we use the street names and a web scraping python code to select among these different SIRET.

A.2.4 Imports and Exports

We use the Customs data for the export sales by firm-product-destination for each year of the sample, and aggregate at the firm level. We use BACI from CEPII for aggregate trade flows and select all flows toward France.




Figure A.2.2: Sales on French Market by Firm Origin

Notes: This figure plots the total sales in France across all manufacturing sectors for French firms vs foreign firms. Foreign sales are computed from BACI.

²The dataset can be downloaded from <https://cadmus.eui.eu/handle/1814/64596>.

A.2.5 Variable Definitions

Revenue is total sales, including exports. In FICUS, this is CATOTAL, whereas in FARE this is REDI_R310. Domestic sales are the sales realized on the domestic market (CAFRANC in FICUS, REDI_R420 in FARE). We compute them as the difference between total sales and export sales from the French Customs dataset.

Employment is the full-time equivalent of the number of directly employed workers by the plant averaged over the year. We use the variable reported in the Stock d'Etablissements.

Energy consumption per energy type is defined as the quantity of energy used by a plant. In EACEI, the variable is CONS_UP for the surveys from 1994 to 1999, and CSUP for the surveys from 2000 onward.

A.2.6 Imputing Emissions

There are two empirical matters with respect to emissions. First, we don't observe all plants in the energy survey. Hence, we do not have emissions information for all plants. Second, we do not observe emissions broken down by destination. Our empirical strategy requires to compute counterfactual revenues on the domestic market. Since we cannot compute counterfactual sales in foreign markets, we cannot compute counterfactual emissions for the foreign market.

First, to address missing plants. Recall that emissions at the plant-level are given by

$$z_{jfst} = \underbrace{\frac{\rho\gamma\kappa_t}{w_t^E}}_{\equiv x_t} e^{\widehat{\beta}_{10}R_{jfst}} \alpha_{jfst} y_{fst}$$

where we have used out estimate $\widehat{\beta}_{10}$ in place of the underlying structural parameter $-\mu_z$. Note that the term $x_t = \frac{\rho\gamma\kappa_t}{w_t^E}$ does not vary by firm within a year. Hence, we can estimate x_t by taking averages over the year

$$\widehat{x}_t = \frac{1}{N_t} \sum_j \frac{z_{jfst}}{\alpha_{jfst} * y_{fst}} e^{-\widehat{\beta}_{10}R_{jfst}}$$

Then for any plant whose emissions we don't observe, we can compute

$$\widehat{z_{jfst}} = \widehat{x}_t e^{\widehat{\beta}_{10}R_{jfst}} \alpha_{jfst} y_{fst}$$

Then, to compute emissions related to production for the domestic market, observe that

$$z_{jfst} = x_t e^{\widehat{\beta_{10} R_{jfst}}} \alpha_{jfst} * (y_{jfst}^{dom} + y_{jfst}^{for})$$

where y_{jfst}^{dom} indicates domestic sales and y_{jfst}^{for} indicates foreign sales.

Then we define plant-level emissions generating in production for the domestic market

$$\widehat{z_{jfst}^{dom}} = \frac{z_{jfst}}{y_{jfst}} * y_{jfst}^{dom}$$

Similarly, counterfactual plant-level emissions generating in production for the domestic market can be computed as

$$\left(\widehat{z_{jfst}^{dom}}\right)' = \frac{z_{jfst}}{y_{jfst}} * e^{-\widehat{\beta_{10} R_{jfst}}} * (y_{jfst}^{dom})'$$




Figure A.2.3: Total vs Domestic Sales and Emissions for French Firms

Notes: Figure plots the total revenues vs revenues on the French market (domestic) for French firms (left) and total CO₂ emissions vs CO₂ generated in production for domestic sales for French firms (right). CO₂ emissions are imputed for plants not included in the EACEI survey. The red line plots the total CO₂ emissions from all manufacturing sectors reported in the National Emissions Inventory of France.

Figure A.2.4 plots total CO₂ emissions (solid) vs CO₂ generated in production for domestic sales (dashed) with imputing emissions (right) and without imputing emissions (left). Our goal will be to construct counterfactual emissions series for the two dashed lines for the scenario in which there had been no EU ETS.




Figure A.2.4: Emissions, with and without Imputations

Notes: Figure plots .

A.3 Theoretical Appendix

A.3.1 Proof of the exact CES price index change

In this section, we derive the change in CES price indices between t and t_0 as described in (1.15) with the definition of Sato-Vartia weights. The proof builds upon Feenstra (1994).

Using the expressions for price indices (1.5) and consumers' optimal expenditures (1.6), we can show that, for any firm f in the set Ω_{ist} of firms from industry i and sector s :

$$\frac{y_{fist}}{\sum_{\ell \in \Omega_{ist}} y_{list}} = p_{fist}^{\frac{\rho}{\rho-1}} P_{ist}^{-\frac{\rho}{\rho-1}},$$

where the LHS represents the market share of firm f within its industry at time t . For any pair of two periods t_0 and t , we define the set of continuing firms as $\Omega_{ist}^* \equiv \Omega_{ist} \cap \Omega_{ist_0}$.

We can write:

$$\underbrace{\frac{y_{fist}}{\sum_{\ell \in \Omega_{ist}^*} y_{list}}}_{\equiv \vartheta_{fist}} \underbrace{\frac{\sum_{\ell \in \Omega_{ist}^*} y_{list}}{\sum_{\ell \in \Omega_{ist}} y_{list}}}_{\equiv \lambda_{ist}} = p_{fist}^{\frac{\rho}{\rho-1}} P_{ist}^{-\frac{\rho}{\rho-1}}.$$

In log differences between periods t_0 and t , we obtain

$$\log \frac{\vartheta_{fist}}{\vartheta_{fist_0}} + \log \frac{\lambda_{ist}}{\lambda_{ist_0}} = \frac{\rho}{\rho-1} \left(\log \frac{p_{fist}}{p_{fist_0}} - \log \frac{P_{ist}}{P_{ist_0}} \right),$$

which we rearrange and multiply on both sides by $\vartheta_{fist} - \vartheta_{fist_0}$ to get:

$$\frac{\rho - 1}{\rho} (\vartheta_{fist} - \vartheta_{fist_0}) = \frac{\vartheta_{fist} - \vartheta_{fist_0}}{\log \frac{\vartheta_{fist}}{\vartheta_{fist_0}}} \log \frac{p_{fist}}{p_{fist_0}} - \frac{\vartheta_{fist} - \vartheta_{fist_0}}{\log \frac{\vartheta_{fist}}{\vartheta_{fist_0}}} \log \frac{P_{ist}}{P_{ist_0}} + \frac{1 - \rho}{\rho} \frac{\vartheta_{fist} - \vartheta_{fist_0}}{\log \frac{\vartheta_{fist}}{\vartheta_{fist_0}}} \log \frac{\lambda_{ist}}{\lambda_{ist_0}}.$$

Summing over firms from the continuing set Ω_{ist}^* , we obtain:

$$\left(\sum_{\ell \in \Omega_{ist}^*} \frac{\vartheta_{\ell ist} - \vartheta_{\ell ist_0}}{\log \frac{\vartheta_{\ell ist}}{\vartheta_{\ell ist_0}}} \right) \log \frac{P_{ist}}{P_{ist_0}} = \sum_{\ell \in \Omega_{ist}^*} \frac{\vartheta_{fist} - \vartheta_{fist_0}}{\log \frac{\vartheta_{fist}}{\vartheta_{fist_0}}} \log \frac{p_{fist}}{p_{fist_0}} + \frac{1 - \rho}{\rho} \left(\sum_{\ell \in \Omega_{ist}^*} \frac{\vartheta_{\ell ist} - \vartheta_{\ell ist_0}}{\log \frac{\vartheta_{\ell ist}}{\vartheta_{\ell ist_0}}} \right) \log \frac{\lambda_{ist}}{\lambda_{ist_0}},$$

which yields (1.14) and the definition of firm-specific Sato-Vartia weights:

$$\phi_{fist} \equiv \frac{\vartheta_{fist} - \vartheta_{fist_0}}{\log \vartheta_{fist} - \log \vartheta_{fist_0}} \bigg/ \sum_{\ell \in \Omega_{ist}^*} \frac{\vartheta_{\ell ist} - \vartheta_{\ell ist_0}}{\log \vartheta_{\ell ist} - \log \vartheta_{\ell ist_0}}.$$

Similarly, for any industry i in the set Υ_s of industries from sector s , we get:

$$\Theta_{ist} = P_{ist}^{\frac{\nu}{\nu-1}} \Psi_{st}^{-\frac{\nu}{\nu-1}}.$$

Taking the log differences between t_0 and t yields:

$$\frac{\log \frac{P_{ist}}{P_{ist_0}} - \log \frac{\Psi_{st}}{\Psi_{st_0}}}{\log \frac{\Theta_{ist}}{\Theta_{ist_0}}} = \frac{1 - \nu}{\nu}.$$

Multiplying each side of the equation by $\Theta_{ist} - \Theta_{ist_0}$ gives:

$$\frac{\Theta_{ist} - \Theta_{ist_0}}{\log \frac{\Theta_{ist}}{\Theta_{ist_0}}} \log \frac{P_{ist}}{P_{ist_0}} - \frac{\Theta_{ist} - \Theta_{ist_0}}{\log \frac{\Theta_{ist}}{\Theta_{ist_0}}} \log \frac{\Psi_{st}}{\Psi_{st_0}} = \frac{1 - \nu}{\nu} (\Theta_{ist} - \Theta_{ist_0}).$$

Summing over industries that belong to the stable set Υ_s from sector s , we obtain:

$$\sum_{i \in \Upsilon_s} \frac{\Theta_{ist} - \Theta_{ist_0}}{\log \frac{\Theta_{ist}}{\Theta_{ist_0}}} \log \frac{P_{ist}}{P_{ist_0}} = \left(\sum_{h \in \Upsilon_s} \frac{\Theta_{hst} - \Theta_{hst_0}}{\log \frac{\Theta_{hst}}{\Theta_{hst_0}}} \right) \log \frac{\Psi_{st}}{\Psi_{st_0}}.$$

which yields (1.15) and the definition of industry-specific Sato-Vartia weights:

$$\Phi_{ist} \equiv \frac{\Theta_{ist} - \Theta_{ist_0}}{\log \Theta_{ist} - \log \Theta_{ist_0}} \bigg/ \sum_{h \in \Upsilon_s} \frac{\Theta_{hst} - \Theta_{hst_0}}{\log \Theta_{hst} - \log \Theta_{hst_0}}.$$

A.3.2 Deriving the Local Approximation Estimation Equation

In what follows we show that the approach described in equation (1.2) corresponds to a local approximation of changes around an equilibrium set at time t_0 .

For any vector of infinitesimal changes in firm-level prices relative to the equilibrium obtained at time t_0 , noted $\{dp_{fist_0}\}_{f \in \Omega_{ist_0}}$, the change in the industry price index P_{ist} is given by:

$$dP_{ist_0} = \sum_{f \in \Omega_{ist_0}} \frac{\partial P_{ist}}{\partial p_{fist} |_{t=t_0}} dp_{fist_0}. \quad (\text{A.1})$$

Given (1.5), we have that, at any period t :

$$\begin{aligned} \frac{\partial P_{ist}}{\partial p_{fist}} &= \left(\sum_{f \in \Omega_{ist}} (p_{fist})^{\frac{\rho}{\rho-1}} \right)^{\frac{\rho-1}{\rho}-1} (p_{fist})^{\frac{\rho}{\rho-1}-1} \\ \Rightarrow \frac{\partial P_{ist}}{\partial p_{fist}} &= \left(\frac{p_{fist}}{P_{ist}} \right)^{\frac{1}{\rho-1}} \\ \Rightarrow \frac{\partial P_{ist}}{\partial p_{fist}} &= \theta_{fist}, \end{aligned}$$

which we can plug in (A.1) to obtain:

$$dP_{ist_0} = \sum_{f \in \Omega_{ist_0}} \theta_{fist_0} dp_{fist_0}. \quad (\text{A.2})$$

As a result, if we approximate small variation in firms' prices with infinitesimal changes around t_0 equilibrium, we obtain $\Delta P_{ist_0} \approx \sum_{f \in \Omega_{ist_0}} \theta_{fist_0} \Delta p_{fist_0}$ for the change in industry i 's price index. We could similarly build $\Delta \Psi_{st_0} \approx \sum_{i \in \Upsilon_{st_0}} \Theta_{fist_0} \Delta P_{ist_0}$ for the change in sector s 's price index, but controlling for sector-time fixed effects absorbs these spillovers.

A.4 Comparing to other models

A.4.1 Analytical Results on Bias in DiD and LA Models

In this section, we derive analytically the biases associated with estimating the average treatment effects using the DiD model and the LA model described in section 1.4.3.

Firm-level treatment effects are defined as $(y_{fit} - y'_{fit})/y_{fit}$. Focusing on directly treated

firms and taking averages over all regulated firms, we have

$$ATT_t^y = \frac{1}{N_t^T} \sum_s \sum_{i \in \Upsilon_s} \sum_{f \in \Omega_{it}^T} \frac{y_{fit} - y'_{fit}}{y_{fit}} = \frac{1}{N_t^T} \sum_s \sum_{i \in \Upsilon_s} \sum_{f \in \Omega_{it}^T} \left(1 - \exp^{-\ln(y_{fit}/y'_{fit})} \right),$$

where N^T is the total number of treated firms in the economy. We define the estimation bias for ATTs as the difference between the estimate of ATTs under a specific model and the true ATTs: $Bias_t^{y,model} \equiv E[\widehat{ATT_t^{y,model}} - ATT_t^{y,true}]$. Similarly, the industry-level equivalent is $Bias_{ist}^{y,model} \equiv E[\widehat{ATT_{ist}^{y,model}} - ATT_{ist}^{y,true}]$, where treatment effects are averaged over firms within industry i only.

First, the DiD model yields the following bias in the estimation of the ATTs on revenues:

$$Bias_t^{yDD} = \frac{1}{N^T} E_\omega \left[\sum_s \sum_{i \in \Upsilon_s} \sum_{f \in \Omega_{it}^T} \left(\exp^{-\ln(y_{fit}/y'_{fit})} - \exp^{-\widehat{\beta}_{y0}^{DD} R_{fit}} \right) \right].$$

Denoting by

$$\text{Reg-Error}_{it}^{yDD} = \ln \left(\frac{\sum_{f \in \Omega_{it}^*} \exp^{-\widehat{\beta}_{y0}^{DD} R_{fit}}}{\sum_{f \in \Omega_{it}^*} \exp^{-\beta_0 R_{fit}}} \right),$$

$$\text{Spillover}_{it} = \frac{(\nu - \rho) \tau}{(\nu - 1)(1 - \rho)} \sum_{k \in \Omega_{it}^*} \phi_{kit} R_{kit} + \frac{\nu \tau}{1 - \nu} \sum_{m \in \Upsilon} \Phi_{mt} \sum_{\ell \in \Omega_{mt}^*} \phi_{\ell mt} R_{\ell mt},$$

and

$$\begin{aligned} \text{Spec-Error}_{it} &= \frac{(\nu - \rho) \sigma}{(\nu - 1)(1 - \rho)} \sum_{k \in \Omega_{it}^*} (\phi_{kit} - \phi'_{kit}) \Delta \alpha_{kit} + \frac{\nu \sigma}{1 - \nu} \sum_{m \in \Upsilon} \sum_{\ell \in \Omega_{mt}^*} (\Phi_{mt} \phi_{\ell mt} - \Phi'_{mt} \phi'_{\ell mt}) \Delta \alpha_{\ell mt} \\ &+ \frac{(\nu - \rho)}{(1 - \nu)(1 - \rho)} \sum_{k \in \Omega_{it}^*} (\phi_{kit} - \phi'_{kit}) \Delta A_{kit} + \frac{\nu}{\nu - 1} \sum_{m \in \Upsilon} \sum_{\ell \in \Omega_{mt}^*} (\Phi_{mt} \phi_{\ell mt} - \Phi'_{mt} \phi'_{\ell mt}) \Delta A_{\ell mt} \\ &+ \frac{(\nu - \rho)}{\rho(\nu - 1)} (\Delta \lambda_{it} - \Delta \lambda'_{it}) + \frac{\nu(1 - \rho)}{\rho(1 - \nu)} \sum_{m \in \Upsilon} (\Phi_{mt} \Delta \lambda_{mt} - \Phi'_{mt} \Delta \lambda'_{mt}), \end{aligned}$$

we have

$$\text{sign}(Bias_{ist}^{yDD}) = -\text{sign} \left(E \left[\text{Reg-Error}_{it}^{yDD} + \text{Spillover}_{it} + \text{Spec-Error}_{it} \right] \right).$$

If $\text{sign}(Bias_{ist}^{yDD})$ is independent from the intensities of treatment, then it would not vary by industry and we can sign $Bias_t^{yDD}$. Otherwise, we cannot.

Given that $E_\omega [\widehat{\beta_{y0}^{DD}}] = \beta_0$, the regression error $\text{Reg-Error}_{it}^{yDD}$ should be in expectation zero. Assuming that $0 < \nu < \rho < 1$, we have $\text{Spillover}_{it} > 0 \iff \tau > 0$. The specification error, Spec-Error_{it} , tends toward zero when market shares under the counterfactual unregulated regime remain almost identical to the observed market shares under the regulated regime (hence, $\phi'_{fist} \approx \phi_{fist}$, $\Phi'_{ist} \approx \Phi_{ist}$ and $\Delta\lambda'_{it} \approx \Delta\lambda_{it}$). Under this condition, the DiD model yields an upward bias in magnitude due to the omission of the spillovers. By contrast, if counterfactual market shares differ widely from observed market shares, we cannot identify the sign of the specification error.

The DiD model for firm-level emissions is only correctly specified for single-plant firms or if emissions from all plants within a firm have been aggregated as it relies on variation across firms. If it is the case, we have

$$\text{sign}(\text{Bias}_{ist}^{ZDD}) = -\text{sign}(E[\text{Reg-Error}_{it}^{zDD} + \text{Spillover}_{it} + \text{Spec-Error}_{it}]),$$

where

$$\text{Reg-Error}_{it}^{zDD} = \ln \left(\frac{\sum_{f \in \Omega_{it}^*} \exp^{-\widehat{\beta_z^{DD}} R_{fit}}}{\sum_{f \in \Omega_{it}^*} \exp^{-(\beta_0 - \mu^z) R_{fit}}} \right).$$

Given that $E_\omega [\widehat{\beta_z^{DD}}] = \beta_0 - \mu^z$, it yields the same sign of bias as for the ATTs on revenues. When $\tau > 0$, emissions decrease and the DiD model yields an upward bias in magnitude when the specification error tends toward zero. However, when $\tau < 0$, the combination of a scale effect (β_0) and a technique effect ($-\mu^z$) implies that the ATTs on firm-level emissions could be either positive or negative. Suppose emissions decrease despite the increase in revenues, then $\text{Bias}_{ist}^{ZDD} > 0$ since $\text{Spillover}_{ist} < 0$, abstracting from the specification error. In this case, the bias can push $\widehat{ATT}_{ist}^{ZDD}$ so much upward relative to the true ATT that it could reverse the sign of it. We summarize these results in the following Proposition A.4.1:

Proposition. *i/ The ATTs on revenues predicted by the DiD model provide information on the sign of the net regulation cost: $\text{sign}(\widehat{\beta_{y0}^{DD}}) = -\text{sign}(\tau)$.*

ii/ If market shares under the counterfactual unregulated regime remain almost identical to the observed market shares under the regulated regime (hence, $\phi'_{fist} \approx \phi_{fist}$, $\Phi'_{ist} \approx \Phi_{ist}$ and $\Delta\lambda'_{it} \approx \Delta\lambda_{it}$), the ATTs on revenues predicted by the DiD model are biased up in magnitude.

iii/ The ATTs on firm-level emissions predicted by the DiD model is not informative about the sign of the true effect of the regulation on these emissions.

Proposition A.4.1 states that the DiD model reveals the correct sign of the net regulation cost, but it overestimates the ATTs on revenues when we abstract from large changes in market shares between the regulated and unregulated regimes. If market shares vary widely between the regulated and unregulated regimes, we can no longer sign the bias on the ATTs analytically. The ATTs on firm-level emissions predicted by the DiD model may not reveal the right sign of the true ATTs. Indeed, the estimate of β_{z0}^{DD} combines a scale effect (a change in sales) with a technique effect (a change in emission intensity) and suffers from a bias due to the omission of within-industry spillovers. If the bias is large, pushing the scale effect upward, then the sign of the ATTs predicted by the DiD model may be wrong. We consider below alternative plant-level models that take into account the existence of multiplant firms.

Second, the LA model yields the following bias in the estimation of the ATTs on revenues:

$$Bias_t^{yLA} = \frac{1}{NT} E \left[\sum_{f \in \Omega_t^T} \left(\exp^{-\ln(y_{fit}/y'_{fit})} - \exp^{-\widehat{\beta}_{y0}^{LA} R_{fit} - \widehat{\beta}_{y2}^{LA} \sum_{k \in \Omega_{it}^*} \phi_{kit} R_{kit} - \widehat{\beta}_{y5}^{LA} \sum_{m \in \Upsilon} \Phi_{mt} \sum_{\ell \in \Omega_{mt}^*} \phi_{\ell mt} R_{\ell mt}} \right) \right].$$

We decompose the bias into three industry-specific components. First, the specification error, $Spec-Error_{it}$, defined above for the DiD model is also present in the LA model. Second, the regression error is defined as:

$$\begin{aligned} Reg-Error_{ist}^{yLA} &= \ln \left(\frac{\sum_{f \in \Omega_{it}^*} \exp^{-\widehat{\beta}_{y0}^{LA} R_{fit}}}{\sum_{f \in \Omega_{it}^*} \exp^{-\beta_0 R_{fit}}} \right) + \left(\widehat{\beta}_{y2}^{LA} - \beta_2 \right) \sum_{k \in \Omega_{it}^*} \phi_{kit} R_{kit} \\ &+ \left(\widehat{\beta}_{y5}^{LA} - \beta_5 \right) \sum_{m \in \Upsilon} \Phi_{mt} \sum_{\ell \in \Omega_{mt}^*} \phi_{\ell mt} R_{\ell mt}. \end{aligned}$$

Third, the weighting error, $Weighting-Error_{ist}$, which corresponds to the error coming from using pre-period market shares instead of Sato-Vartia weights is defined as follows:

$$\begin{aligned} Weighting-Error_{ist} &= \widehat{\beta}_{y2}^{LA} \left(\sum_{k \in \Omega_{it}} \theta_{kit} R_{kit} - \sum_{k \in \Omega_{it}^*} \phi_{kit} R_{kit} \right) \\ &+ \widehat{\beta}_{y5}^{LA} \left(\sum_{m \in \Upsilon} \Theta_{mt} \sum_{\ell \in \Omega_{mt}} \theta_{\ell mt} R_{\ell mt} - \sum_{m \in \Upsilon} \Phi_{mt} \sum_{\ell \in \Omega_{mt}^*} \phi_{\ell mt} R_{\ell mt} \right). \end{aligned}$$

We thus obtain:

$$sign(Bias_{ist}^{LA}) = sign \left(Reg-Error_{ist}^{yLA} + Weighting-Error_{ist} + Spec-Error_{it} \right).$$

As above, the specification error tends toward zero when market shares under the counter-

factual unregulated regime remain almost identical to the observed market shares under the regulated regime. Otherwise, its sign is ambiguous. We cannot determine the sign of the weighting error, which depends on the difference between pre-period revenue shares and Sato-Vartia weights, themselves being functions of the difference between current and pre-period revenue shares on a continuing good set. We cannot sign the regression error in general. Yet, as shown in section A.3.2, pre-period market shares are a correct approximation of Sato-Vartia weights for infinitesimally small changes around the pre-period equilibrium. As a result, under this condition, the LA model is unbiased. These results are summarized in the following Proposition A.4.1:

Proposition. *i/ The ATTs predicted by the LA model are unbiased if and only if firm-level price changes relative to the pre-period t_0 equilibrium are infinitesimally small.*

ii/ Otherwise, the ATTs predicted by the LA model are biased, and the sign of the bias is ambiguous.

A.4.2 Monte Carlo Experiments

We use Monte Carlo experiments to evaluate the performance of our estimator in a finite-sample setting and compare to the DiD and LA estimators. The data generating process is described in section 1.4.3.

For each parameter combination, Figure A.4.1 reports the distribution of the biases in the estimates of the ATTs for our estimator (in red), for the DiD estimator (in black), and for the LA estimator (in blue). For our estimator, we find that the distribution of biases is centered on zero and quite precise. By contrast, we also see that even when the bias in the DiD estimator is small (for example, the top left panel in Figure A.4.1), the distribution is clearly centered away from zero. Specifically, we find that the DiD estimator is biased up in magnitude on average for all parameter combinations (i.e., the bias is on average negative when $\mu_t^E = 0$ and positive when $\mu_t^E = 0.8$). The bias in the LA estimator can be small, especially with low substitutability ($\rho = .8$). However, when varieties are more substitutable ($\rho = .95$), the LA estimator is severely biased. For the case with $\rho = .95$ and $\mu^e = .8$, the LA estimator yields an ATT that has the wrong sign, on average. Looking at Figure A.4.1), we see that for these cases, the entire distribution of estimates for the LA estimator is shifted away from zero.

We explore the biases on ATTs for aggregate emissions in Figure A.4.2. Across all parameter combinations, we find that aggregate emissions fall. This is because we assume that regulation increases the cost of emissions ($\mu^z > 0$). Hence, even though the regulation reallocates resources across firms, the direct effect of the regulation on input costs dominates, and total emissions fall. As we can see in Figure A.4.2, our estimator successfully

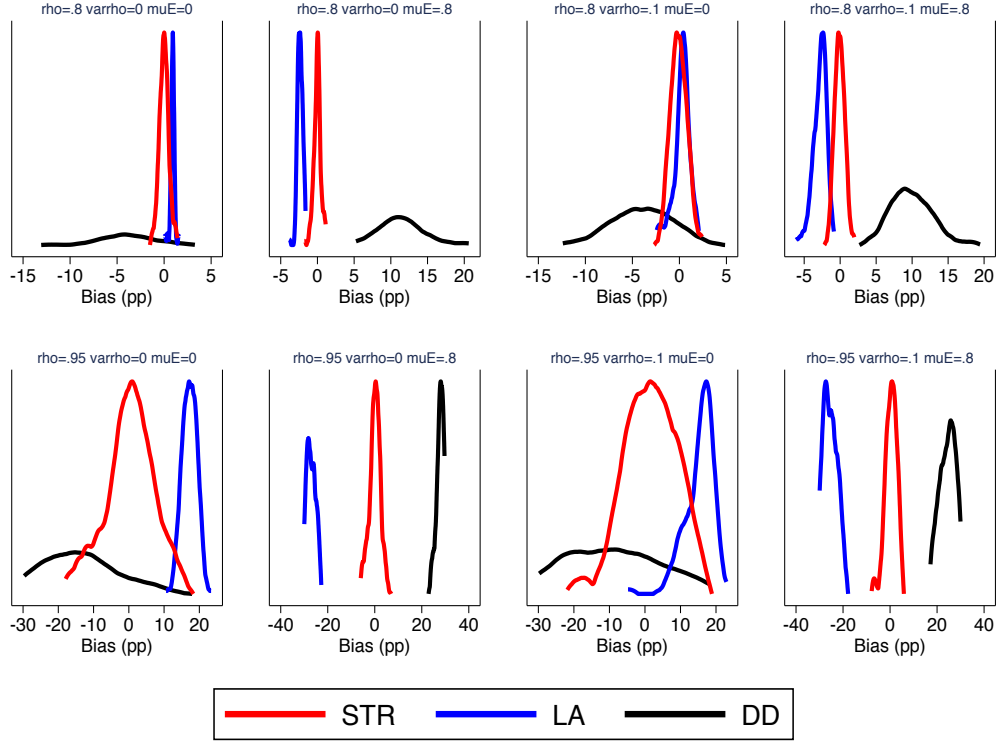


Figure A.4.1: Distributions of ATT for revenues at the firm level, Multi-Plant Firms

Notes: Each subfigure plots the distribution of bias in the ATTs for revenues for each estimator across 100 replications for a given parameter combination. The bias is computed as $\widehat{ATT} - ATT^{true}$ and denominated in percentage points (pp).

recovers unbiased estimates of these aggregate effects. By contrast, both the DiD and LA estimators yield biased estimates for all parameter combinations. For both estimators, the bias can be either up or down, and the estimated sign can be either positive or negative.

In Appendix section A.4.1, we decompose the bias in the DiD estimator into three components: regression error, spillover, and specification error. We also decompose the bias in the LA estimator into three components: specification error, weighting error, and regression error. Only the specification error is the same for both estimators. Figure A.4.3 reproduces Figure 1.4.1 for these two estimators and explores the sources of their biases by reporting the average magnitude of each of the three components of the bias, for each parameter combination. For this set of simulation results, the spillover error dominates for the DiD estimator whereas the specification error dominates for the LA estimator, and hence each dictates the sign of the bias. However, this is not true for all simulations. For instance, Figures A.4.4-A.4.5 present results for an economy with all single-plant firms and random walk productivity growth. In this set of simulation results, Figure A.4.6 shows that, in some cases, the weighting error can dominate, thereby revealing that it is




Figure A.4.2: Effects on Aggregate Emissions, Multi-Plant Firms

Notes: This figure plots the average aggregate effect on total emissions relative to the counterfactual unregulated equilibrium for each estimator across 100 replications against the average of the true metric, for each parameter combination. The black line corresponds to the 45-degree line.

not possible to sign the bias in the LA estimator in general.

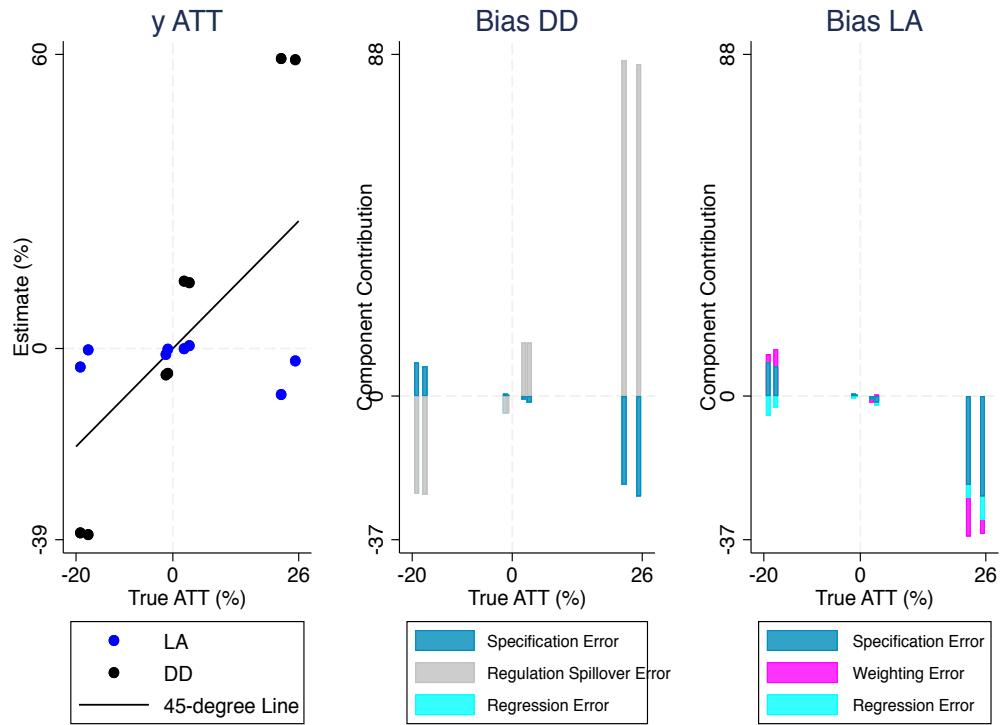


Figure A.4.3: Sources of Bias, Multi-Plant Firms

Notes: Left panel reproduces the top left panel of Figure 1.4.1, leaving out the STR estimator. Middle panel plots the contributions of different sources of bias to the overall bias in the DD estimator for different parameter combinations. The terms are defined in appendix A.4.1. The average of the true ATT on revenues by parameter combination is measured on the x-axis. Right panel plots the contributions of different sources of bias to the overall bias in the LA estimator. In the data generating process, firms own multiple plants and productivity follows an AR(1) process.

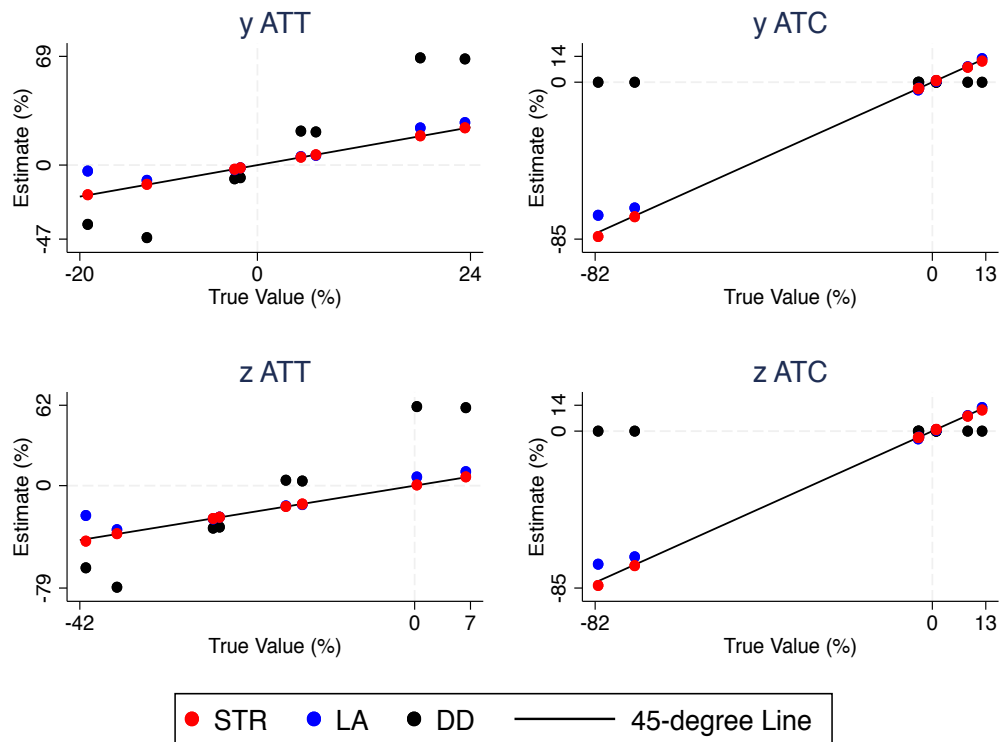


Figure A.4.4: ATE for revenues and emissions at the firm level, Single-Plant Firms

Notes: Subfigures plot average estimates of treatment effects on revenues (y) and emissions (z) as a percentage of observed outcomes for each estimator across 100 replications against the average of the true metric, for each parameter combination. Data generating process assumes single-plant firms and random walk productivity growth. Left panels show the ATTs for each parameter combination, whereas right panels show the ATCs, all at the firm level. The black line corresponds to the 45-degree line.




Figure A.4.5: Effects on Aggregate Emissions, Single-Plant Firms

Notes: This figure plots the average aggregate effect on total emissions relative to the counterfactual unregulated equilibrium for each estimator across 100 replications against the average of the true metric, for each parameter combination. The black line corresponds to the 45-degree line. Data generating process assumes single-plant firms and random walk productivity growth.

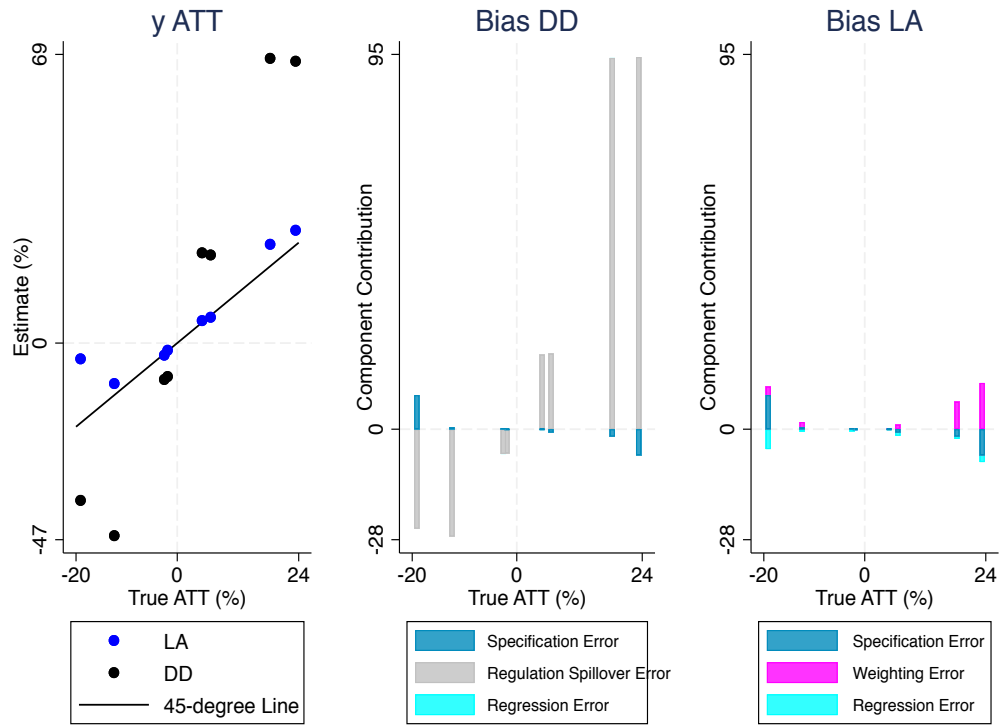


Figure A.4.6: Sources of Bias, Single-Plant Firms

Notes: Left panel reproduces the top left panel of Figure A.4.4, leaving out the STR estimator. Middle panel plots the contributions of different sources of bias to the overall bias in the DD estimator for different parameter combinations. The terms are defined in appendix A.4.1. The average of the true ATT on revenues by parameter combination is measured on the x-axis. Right panel plots the contributions of different sources of bias to the overall bias in the LA estimator. In the data generating process, firms own a single plants and productivity follows a random walk.

A.5 Extension to Oligopoly

A.5.1 Model with Bertrand-Nash competition

In this section, we discuss a modification to our baseline model allowing for oligopolistic behaviors in the manufacturing sectors. The demand function, the production function, and the emissions function are unchanged relative to the baseline model. The only modeling difference is that firms play a Bertrand-Nash equilibrium in prices. As a result, firms take into account the impact of their pricing decision on price indices.

Given demand function (1.6) and unit cost function (1.21), profit-maximization yields the following pricing rule, when firms internalize the effect of their pricing decision on the CES price index:

$$p_{fit}^O = c_{fit} \left[\frac{(\nu - 1) - (\nu - \rho) \theta_{fit} + (1 - \rho) \nu \theta_{fit} \Theta_{it}}{\rho (\nu - 1) - (\nu - \rho) \theta_{fit} + (1 - \rho) \nu \theta_{fit} \Theta_{it}} \right], \quad (\text{A.3})$$

where $\theta_{fit} = (p_{fit}/P_{it})^{\frac{-\rho}{1-\rho}}$ describes the response of the within-industry price index to firm f 's price change, and $\Theta_{it} = (P_{it}/\Psi_t)^{\frac{-\nu}{1-\nu}}$ describes the response of the sector price index to industry i 's price index change. The pricing rule thus varies with firm f 's market share in industry i and industry i 's market share in sector s at time t . We denote by φ_{fit} the term in brackets in (A.3), which describes firm f 's markup in this oligopoly setting. Hence, $p_{fit}^O = c_{fit} \varphi_{fit}$.

When firms have market power, given the pricing rule (A.3), changes in firm f 's price between post-regulation period t and pre-regulation base year t_0 can be described by:

$$\Delta p_{fit}^O = \Delta c_{fit} + \Delta \varphi_{ist}. \quad (\text{A.4})$$

Hence, it varies with its own cost changes and its markup changes. Combining (1.6), (1.10) and (A.4) allows us to compute the changes in revenues between t and t_0 :

$$\begin{aligned} \Delta y_{fit}^O &= \frac{\rho\tau}{\rho-1} R_{fit} + \frac{\rho}{\rho-1} \Delta \varphi_{fit} + \frac{\rho\sigma}{\rho-1} \Delta \alpha_{fit} + \frac{(\nu-\rho)\tau}{(\nu-1)(1-\rho)} \sum_{k \in \Omega_{it}^*} \phi_{kit} R_{kit} + \frac{(\nu-\rho)}{\rho(\nu-1)} \Delta \lambda_{it} \\ &+ \frac{(\nu-\rho)\sigma}{(\nu-1)(1-\rho)} \sum_{k \in \Omega_{it}^*} \phi_{kit} \Delta \alpha_{kit} + \frac{\nu\sigma}{1-\nu} \sum_{m \in \Upsilon} \Phi_{mt} \sum_{\ell \in \Omega_{mt}^*} \phi_{\ell mt} \Delta \alpha_{\ell mt} \\ &+ \frac{\nu\tau}{1-\nu} \sum_{m \in \Upsilon} \Phi_{mt} \sum_{\ell \in \Omega_{mt}^*} \phi_{\ell mt} R_{\ell mt} + \frac{\nu(1-\rho)}{\rho(1-\nu)} \sum_{m \in \Upsilon} \Phi_{mt} \Delta \lambda_{mt} + \Delta Y_t \\ &+ \frac{\nu-\rho}{(\nu-1)(1-\rho)} \sum_{k \in \Omega_{it}^*} \phi_{kit} \Delta \varphi_{kit} + \frac{\nu}{1-\nu} \sum_{m \in \Upsilon} \Phi_{mt} \sum_{\ell \in \Omega_{mt}^*} \phi_{\ell mt} \Delta \varphi_{\ell mt} \\ &+ \frac{\rho}{\rho-1} \Delta A_{fit} + \frac{(\nu-\rho)}{(1-\nu)(\rho-1)} \sum_{k \in \Omega_{it}^*} \phi_{kit} \Delta A_{kit} + \frac{\nu}{1-\nu} \sum_{m \in \Upsilon} \Phi_{mt} \sum_{\ell \in \Omega_{mt}^*} \phi_{\ell mt} \Delta A_{\ell mt}. \end{aligned} \quad (\text{A.5})$$

Compared to (1.16), equation (A.5) includes additive terms related to firms' markup changes. The expression of markup φ_{fit} itself depends on parameters ρ and ν , as well as on market share information. As a result, the effects of being treated and of the spillovers are no longer homogeneous across firms in the same industry: these impacts will depend on each firm's market share and markup. We can no longer apply the same empirical strategy to estimate treatment effects on revenues as in the baseline model.

The change in emissions from pre-period t_0 to period t depends on the changes in revenues, which differ depending on the competitive environment. We obtain

$$\Delta Z_{fit}^O = -\mu^z R_{fit} + \Delta \kappa_t - \Delta w_{Et} + \Delta y_{fit}^O, \quad (\text{A.6})$$

where Δy_{fit}^O corresponds to (A.5). To estimate the treatment effects on emissions, the same empirical strategy as in the baseline model can apply using plant-level regressions.

To estimate the treatment effect on revenues under oligopoly, we can rewrite the log change in revenues described in (A.5) as:

$$\Delta y_{fit} = \frac{\rho\tau}{\rho-1} R_{fit} + \frac{\rho\sigma}{\rho-1} \Delta \alpha_{fit} + \frac{\rho}{\rho-1} \ln \left[\frac{\frac{(\nu-1) - (\nu-\rho)\theta_{fit} + (1-\rho)\nu\theta_{fit}\Theta_{it}}{\rho(\nu-1) - (\nu-\rho)\theta_{fit} + (1-\rho)\nu\theta_{fit}\Theta_{it}}}{\frac{(\nu-1) - (\nu-\rho)\theta_{fit_0} + (1-\rho)\nu\theta_{fit_0}\Theta_{it_0}}{\rho(\nu-1) - (\nu-\rho)\theta_{fit_0} + (1-\rho)\nu\theta_{fit_0}\Theta_{it_0}}} \right] + \delta_{ist} + \epsilon_{fit} \quad (\text{A.7})$$

where all terms that are the same within an industry i in sector s are collected into an industry-time fixed effect δ_{ist} . Endogeneity concerns emerge from the fact that θ_{fit} and Θ_{it} correlate with firm-level productivity shocks, $\Delta \alpha_{fit}$. In the oligopoly context, we could estimate parameters ρ , ν , σ and τ directly by nonlinear GMM with the following instruments: $\Lambda_{fist} = \sum_{k \in \Omega_{ist,-f}} \theta_{kist_0} R_{kist}$ for θ_{fit} and $\Gamma_{fist} = \sum_{m \in \Upsilon_s} \Theta_{mt} \sum_{\ell \in \Omega_{mt,-f}^*} \theta_{\ell mst_0} R_{\ell mst}$ for Θ_{it} . Doing so, we would use other firms' treatment intensity within an industry i and within a sector s , excluding firm f , as an instrument for firm f 's market share within industry i and for industry i 's market share within sector s , respectively. These instruments do not respond to firm f 's productivity shocks (exclusion restriction) and predict firm f 's market share in an industry or a sector.

A.5.2 Monte-Carlo experiments for Oligopoly

In this section, we explore the potential role of misspecification bias with respect to strategic behavior. Hence, we simulate the economy assuming all firms engage in Bertrand-Nash pricing. All other features of the simulations remain the same. We present results in Figures A.5.1 - A.5.4. We find that, even if firms play Bertrand-Nash, our procedure still delivers estimates of treatment effects that matches the true values quite well. We can see minor differences relative to Figures 1.4.1-1.4.2-A.4.2, but misspecifying the nature of

competition does not appear to generate a noticeable bias.

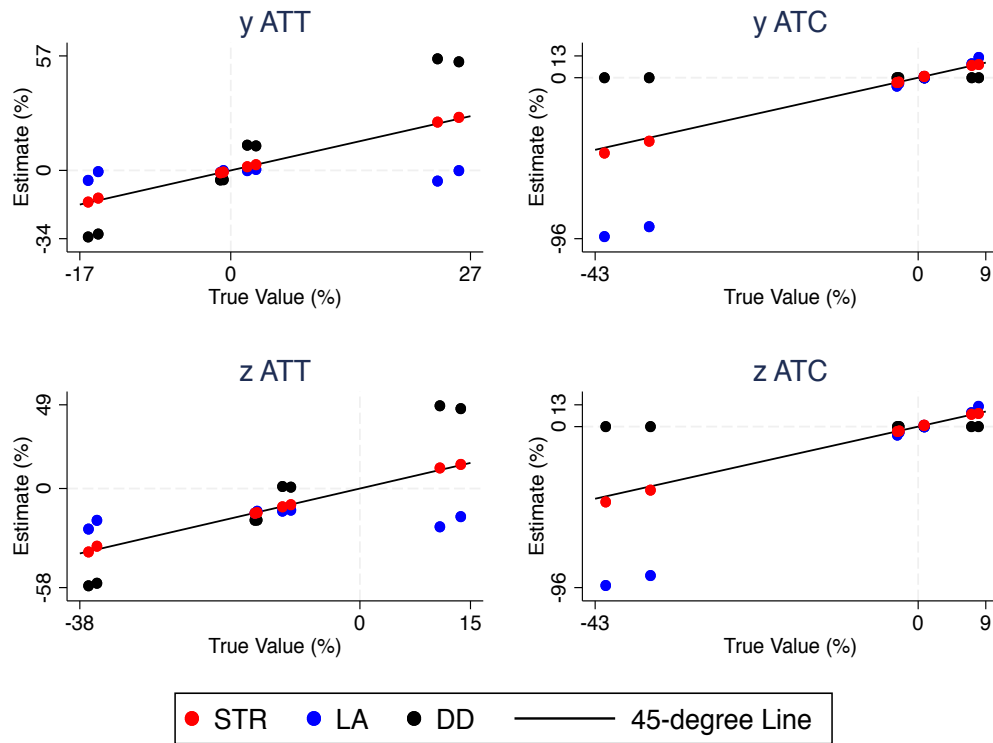


Figure A.5.1: ATE for revenues and emissions at the firm level, MP Oligopoly

Notes: Subfigures plot average estimates of treatment effects on revenues (y) and emissions (z) as a percentage of observed outcomes for each estimator across 100 replications against the average of the true metric, for each parameter combination. Data generating process assumes firms engage in Bertrand-Nash competition on prices. Firms own multiple plants and productivity follows an AR(1) process. Left panels show the ATTs for each parameter combination, whereas right panels show the ATCs, all at the firm level. The black line corresponds to the 45-degree line.

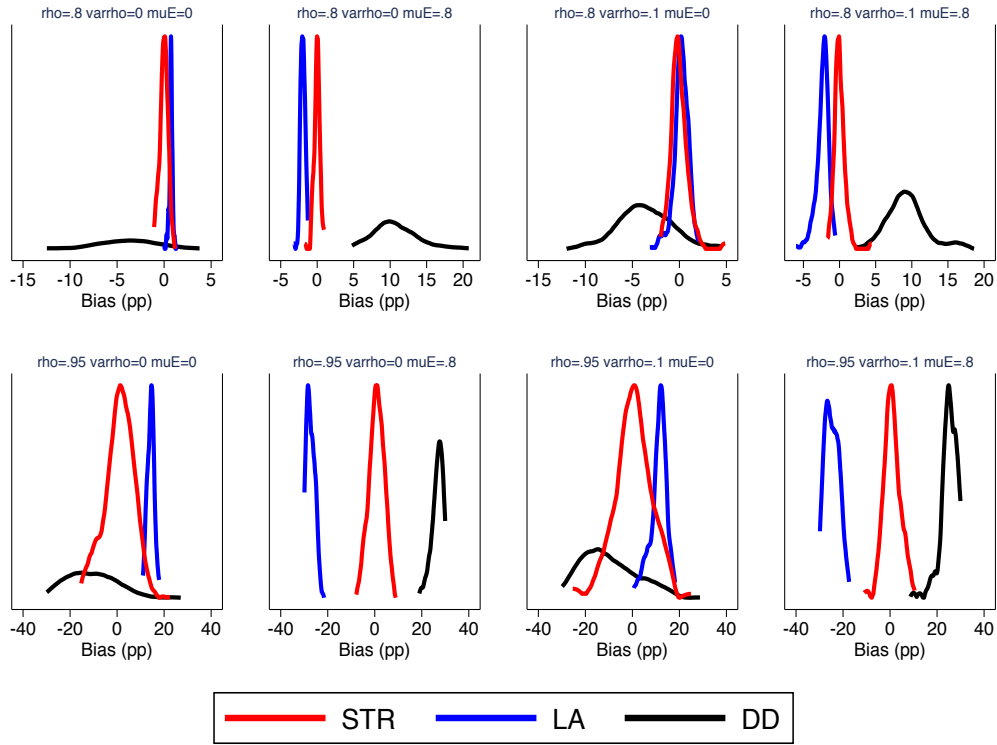


Figure A.5.2: Distributions of ATEs for revenues at the firm level, MP Oligopoly

Notes: Each subfigure plots the distribution of bias in the ATTs for revenues for each estimator across 100 replications for a given parameter combination. The bias is computed as $\widehat{ATT} - ATT^{true}$ and denominated in percentage points (pp). Data generating process assumes firms engage in Bertrand-Nash competition on prices. Firms own multiple plants and productivity follows an AR(1) process.

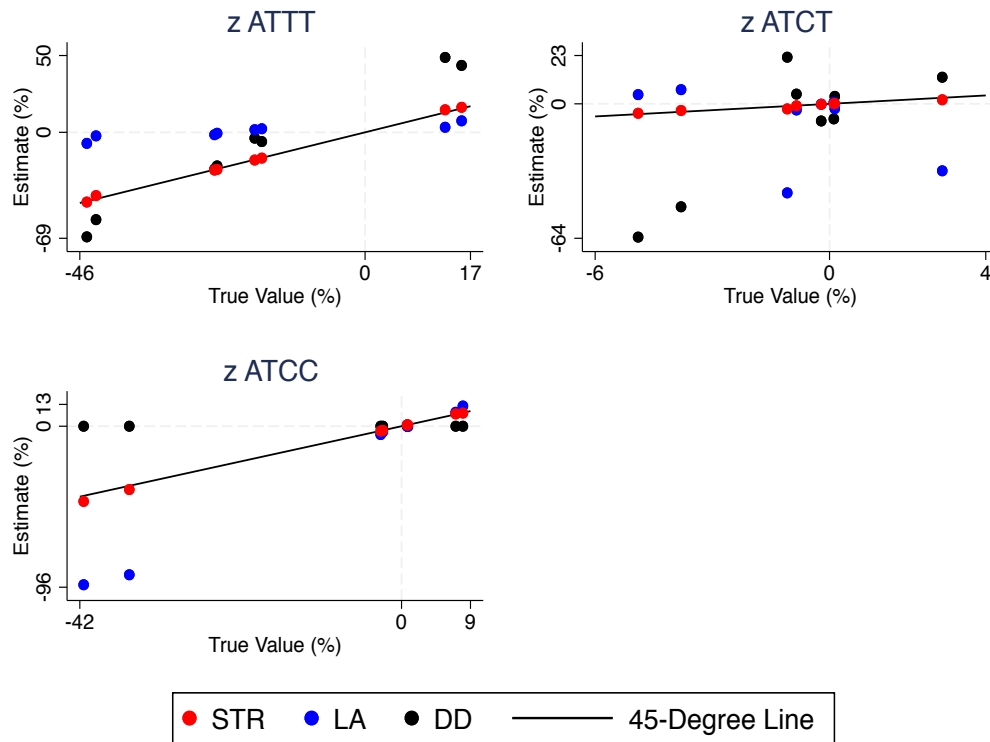


Figure A.5.3: ATE for emissions at the plant level, MP Oligopoly

Notes: Subfigures plot average estimates of treatment effects on emissions (z) at the plant level as a percentage of observed emissions for each estimator across 100 replications against the average of the true treatment effect, for each parameter combination. The first panel plots the ATTTs for each parameter combination, the second ATCTs, and the third ATCCs. The black line corresponds to the 45-degree line. A plant is labeled “TT” if it is regulated in period 2, “CT” if the plant is unregulated in period 2, but part of a firm with regulated plants in period 2, and “CC” if the plant is part of a firm with no regulated plants. Data generating process assumes firms engage in Bertrand-Nash competition on prices. Firms own multiple plants and productivity follows an AR(1) process.




Figure A.5.4: Effects on Aggregate Emissions, Oligopoly

Notes: This figure plots the average aggregate effect on total emissions relative to the counterfactual unregulated equilibrium for each estimator across 100 replications against the average of the true metric, for each parameter combination. The black line corresponds to the 45-degree line. Data generating process assumes firms engage in Bertrand-Nash competition on prices. Firms own multiple plants and productivity follows an AR(1) process.

Appendix B


Appendix to Chapter 2

B.1 Descriptive Statistics on Pollution from the Industry

Over the last decades, several regulations have been implemented in France to reduce atmospheric pollution. As a result, national levels of emissions have decreased in the past. Figure (B.1.1), constructed from the European Emissions Database for Global Atmospheric Research (EDGAR), shows this decrease for particulate matters 10 and 2.5 as well as for sulfur dioxide. However, observed levels of pollution in France are still well above World Health Organization (WHO) guidelines in several areas over the country. Using data on air concentration in PM10, PM2.5 and SO2 from ground monitors installed in diverse locations, I illustrate this situation on Figure (B.1.2) and show that a substantial fraction of monitors measures concentrations above the 24-hour and annual guidelines from the WHO guidelines. There are many sources for this pollution. Figure (B.1.1) also illustrates the distribution of the pollution sources. One can observe that for the three pollutants highlighted, PM10, PM2.5 and SO2, industrial activities represent a large fraction of the emissions. Tables (B.1.1), (B.1.2) and (B.1.3) present the respective contributions of the main polluting activities: in 2012, industrial activities contributed to respectively 23%, 40% and 42% of PM10, PM2.5 and SO2 national emissions.

% Share from:	1990	2000	2012
Industry	19.4	21.4	23.0
Energy	6.7	4.8	3.8
Transport	12.2	15.4	9.4
Residential	40.1	34.6	39.1
Agriculture	17.2	22.5	24.6
Waste	4.4	1.3	0.1

Table B.1.1: Distribution of PM10 Emissions Across Activities (data from EDGAR timeseries): in 2012, 23% of PM10 emissions result from industrial activities.



Data are from the EDGAR timeseries (see Crippa et al. (2019)). Industry sector includes IPPC codes 1A2 & 2. Units are Gg for all pollutants.

Figure B.1.1: 1990-2012 Evolution of PM10, PM2.5 and SO2 emissions by Activities (data from EDGAR timeseries): for these pollutants, a substantial fraction of national emissions still result from industrial activities.

% Share from:	1990	2000	2012
Industry	30.5	28.0	39.6
Energy	10.8	7.5	5.7
Transport	37.6	46.3	33.2
Residential	12.0	8.3	10.6
Agriculture	5.5	9.0	10.8
Waste	3.6	1.0	0.1

Table B.1.2: Distribution of PM2.5 Emissions Across Activities (data from EDGAR timeseries): in 2012, almost 40% of PM2.5 emissions result from industrial activities.

% Share from:	1990	2000	2012
Industry	32.3	28.8	41.7
Energy	43.0	47.2	31.9
Transport	2.7	3.9	5.5
Residential	21.3	19.6	20.3
Agriculture	0.2	0.2	0.5
Waste	0.5	0.4	0.2

Table B.1.3: Distribution of SO2 Emissions Across Activities (data from EDGAR timeseries): in 2012, almost 42% of SO2 emissions result from industrial activities.




Figure B.1.2: Annual and 24-hour Mean Concentration Measurement from French Ground Monitors in 2012 (data from the EEA AirBase V8): for PM10, PM2.5 and SO2, a substantial fraction of ground monitors measures concentrations above WHO guidelines.

B.2 Air Quality Regulations in France

Air pollution regulation in France is based on European standards. Limits for maximum air concentration of various pollutants and reduction objectives are set at the EU level (directives 2004/107 and 2008/50/CE) to be enforced in each Member States. In turn, they must implement specific action plans to remain within these limits and reach the air quality objectives. In the case of France, these actions are taken both by central and local authorities.

The French government regulates pollutant emissions through tax instruments and regulatory standards (Hauvuy & Riedinger, 2005; Bougon & Lavergne, 2019). The main tax instrument is the General Tax on Polluting Activities (*TGAP*), which is a national tax on the quantity of pollutant emitted per year. In 2016, more than a thousand installations were subjects to this tax (IGF, 2018). However, the air pollution component of the *TGAP* is limited,¹ and does not constitute the most efficient instrument enforced to limit air pollution (Millock & Nauges, 2006; IGF, 2018). Frameworks regulating industrial emissions are mainly set at a national level. However, their implementation is usually enforced by local authorities. For instance, the main regulatory framework for polluting activities is the *ICPE* regulation, standing for Plants Classified for the Protection of the Environment. The *ICPE* is a set of norms governing polluting plants activities in relation with their impact on the environment. Notably, it requires that an opening permit is delivered by the authority of the departmental prefects for any polluting plant under the condition of implementing specific technology standards and after conducting a local survey of the local population.² Moreover, some of these classified plants may fall under the EU Directive on Industrial Emissions (IED).³ This is the case if they are running polluting installations with capacities above thresholds set in the IED. The main obligation enforced under this regulation is the Best Available Technique (BAT): authorizations for industrial installations are conditional to the use of the least pollution intensive techniques. In that respect, the current national regulation focuses mainly on the largest plants, but any smaller polluting plant not covered by the *ICPE* regulation is regulated at the municipality level.

Furthermore, since the LAURE law, adopted in 1996, other local authorities can use specific measures to improve local air quality and reach air quality national targets. For

¹In 2016, €50 millions were collected compared to the €3.8 billions collected for the carbon component of the energy consumption tax in 2016 (DGEC, 2016)

²These prefects are the regional authorities for the “*départements*” which are NUTS 3 geographic units (larger than commuting zones).

³Since 2010, the Industrial Emissions Directive has replaced the Integrated Pollution Prevention and Control Directive. The guidelines are similar and aim at preventing air, water, and soil pollution.

instance, starting in the early 2000's, several "Atmospheric Protection Plans" (*PPA*) have been implemented in different areas. Within their respective application zone, many of these plans adopted differentiated measures for more versus less densely populated areas. Some of these plans introduced more stringent environmental standards for the manufacturing industries, especially in agglomerations above a certain threshold of inhabitants. For instance, the three consecutive *PPAs* for the Paris region mandated lower industrial NOx and PM emission caps relative to the national caps. These plans also implement emergency responses when air concentration of certain pollutant exceed national and European standards. Figure (B.2.1) shows the distribution of existent *PPA* across commuting zones⁴ and highlights the positive relationship between the size of commuting zones and their probability to lie in the perimeter of an *PPA*.

Finally, both national and local regulations mandate that when air quality reaches regulatory thresholds, emergency actions must be launched by the departmental prefect. Figure (B.2.2) illustrates the positive relationship between the size of these areas and the number of emergency actions cases accounted for since 2017.

⁴*PPA* are adopted at a geographic level potentially lower than the commuting zone, which is a statistical construction. The map in Figure (B.2.1) displays commuting zone where at least one *PPA* is implemented in its sub-areas.

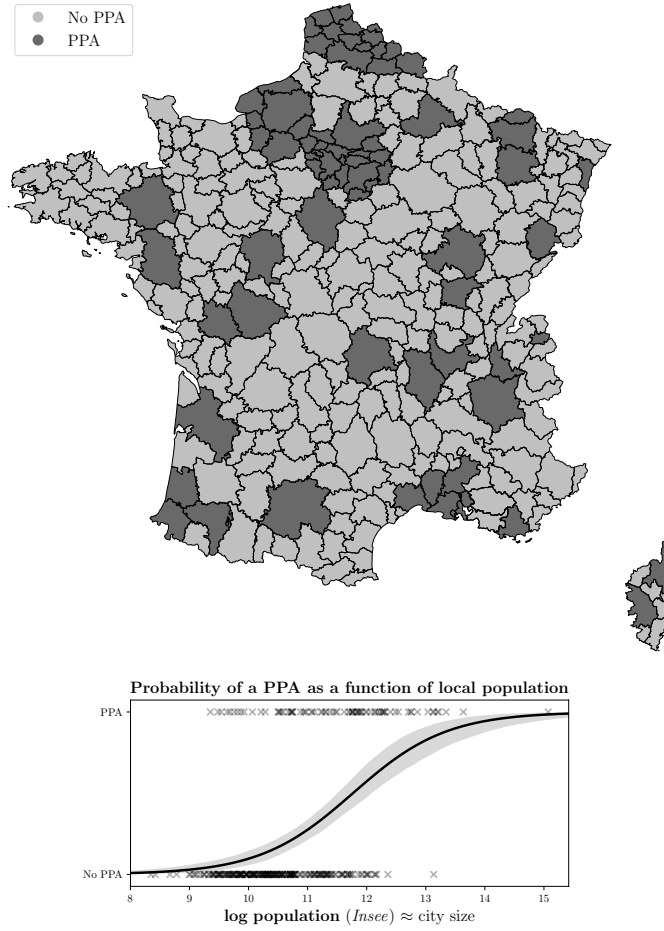


Figure B.2.1: Distribution of Atmospheric Protection Plans across commuting zones
Note: The upper map shows the set of commuting zones that contains municipalities that adopted PPAs since the LAURE law. The lower panel plots the probability that a commuting zone contains municipalities that adopted a PPA as a function of its population. The black line is the result of a logit function fitted on the data on PPA adoption across French commuting zones.
Data on PPAs is from the French Ministry to the Ecological Transition




Figure B.2.2: Number of emergency actions due to excessive measures of atmospheric pollutants
Note: This figure plots the number of daily emergency procedures at the “département” level since 2017 as a function of the “département” population.

B.3 Illustrative Framework Appendix

In each city j , the quantity of the tradable good produced is:

$$Q_j = Z_j^\alpha (\phi_j^L L_j^{prod})^{1-\alpha}, \quad (\text{B.1})$$

with L_j^{prod} the local population employed for production, ϕ_j^L the local labor productivity and Z_j the local quantity of emissions. Each input has the respective unit price w_j and t_j . Moreover, remember that due to the assumption of agglomeration economies local labor productivity ϕ_j^L is equal to $b_j L_j^\nu$. Local firms solve the cost minimization program:

$$\min_{Z_j, L_j} L_j^{prod} w_j + Z_j t_j \text{ s.t. } Q_j = Z_j^\alpha (\phi_j^L L_j^{prod})^{1-\alpha}. \quad (\text{B.2})$$

FOCs with respect to labor and emissions yield the relative labor to emissions intensity, which is city specific and depends on the local ratio of input prices:

$$\frac{Z_j}{L_j^{prod}} = \frac{\alpha}{1-\alpha} \frac{w_j}{t_j}. \quad (\text{B.3})$$

Agglomeration economies are an externality and are not taken into account by firms in their cost minimization problem. That is to say that each firm is atomistic and does not consider that fact that the higher is local employment the larger labor productivity will be. From this proportion of each input in production, one can compute the marginal production price c_j :

$$c_j = \kappa(\alpha) t_j^\alpha \left(\frac{w_j}{\phi_j^L} \right)^{1-\alpha}, \quad (\text{B.4})$$

with $\kappa(\alpha) = \frac{(1-\alpha)^{1-\alpha}}{\alpha^\alpha}$ which I drop in the following computations for simplicity and without loss of generality (α remains a constant in the equilibrium and the central planner's optimization problem).

Given the assumption of perfect competition and not mark-up with a fixed output price fixed to 1 on foreign markets, we then have in equilibrium that $c_j = 1$ which gives the expression for equilibrium wages:

$$w_j = b_j L_j^\nu t_j^{-\frac{\alpha}{1-\alpha}}. \quad (\text{B.5})$$

The total income accounts for payments for production labor and labor employed for

abatement (or). As a result one can write:

$$w_j L_j = w_j L_j^{prod} + t_j Z_j, \quad (B.6)$$

which yields the formula for local emissions as function of local population, local wages and local emission cost:

$$Z_j = \alpha \frac{w_j L_j}{t_j}. \quad (B.7)$$

Given that the output price is the same across cities and normalized to 1, each worker's budget constraint can be written as:

$$w_j \times 1 = 1 \times c_j. \quad (B.8)$$

Using equations (B.5), (B.7) and (B.8) to substitute in the representative worker's utility function we get:

$$u_j = a_j b_j^{1-\gamma} L_j^{-\theta} t_j^{\frac{\gamma-\alpha}{1-\alpha}}, \quad \text{with } \theta = \delta + \gamma - \nu(1 - \gamma). \quad (B.9)$$

Given the assumption of free migration of workers, utility is equalized across cities and I note its equilibrium level \bar{u} . The national population is fixed and normalized to 1, so that:

$$\sum_{j \in C} L_j = 1. \quad (B.10)$$

Using the fact that equation (B.9) can be rewritten as:

$$L_j = \left[\bar{u}^{-1} a_j b_j^{1-\gamma} t_j^{\frac{\gamma-\alpha}{1-\alpha}} \right]^{\frac{1}{\theta}}, \quad (B.11)$$

which illustrates that there exist a unique spatial equilibrium if and only if $\theta \neq 0$. This condition can be expressed as a condition on the elasticity of the pollution externality: $\gamma \neq -\frac{\delta-\nu}{1+\nu}$. Assuming that pollution is a congestion force (i.e. $\gamma > 0$), this condition ensures that the pollution externality does not exactly offset the combined effect of general congestions effect and agglomeration economies. When congestion effects strictly outweigh agglomeration economies, this condition always holds. We consider that this is always the case and that $\theta > 0$ (which is supported by empirical evidence and by the estimates obtained in section 2.4). I substitute in equation (B.10) and extract the equilibrium

welfare:

$$\bar{u} = \left[\sum_{j \in C} a_j^{\frac{1}{\theta}} b_j^{\frac{1-\gamma}{\theta}} t_j^{\frac{1}{\theta} \frac{\gamma-\alpha}{1-\alpha}} \right]^{\theta}. \quad (\text{B.12})$$

Finally I solve that optimization program:

$$\max_{\{t_i\}_{i \in C}} \bar{u}(t_1, \dots, t_C) \text{ s.t. } \sum_{j \in C} t_j = 1, \quad (\text{B.13})$$

by computing the first and second order conditions of the Lagrangian:

$$\mathcal{L} = \bar{u}(t_1, \dots, t_C) + \lambda \left(C - \sum_{j \in C} t_j \right), \quad (\text{B.14})$$

with λ the Lagrangian multiplier. Computing the derivatives of \bar{u} , I find:

$$\frac{t_j}{\bar{u}} \frac{\partial \bar{u}}{\partial t_j} = \frac{\gamma - \alpha}{1 - \alpha} L_j \quad (\text{B.15})$$

and

$$\frac{\partial^2 \bar{u}}{\partial t_j^2} = \frac{1}{t_j} \frac{\partial \bar{u}}{\partial t_j} \left[\left(\frac{\gamma - \alpha}{1 - \alpha} - \frac{1}{\theta} \frac{\gamma - \alpha}{1 - \alpha} \right) L_j - \left(1 - \frac{1}{\theta} \frac{\gamma - \alpha}{1 - \alpha} \right) \right]. \quad (\text{B.16})$$

Then, the first order conditions are:

$$\forall j \in C, \bar{u}^{\frac{\theta-1}{\theta}} \frac{\gamma - \alpha}{1 - \alpha} a_j^{\frac{1}{\theta}} b_j^{\frac{1-\gamma}{\theta}} t_j^{*\frac{1}{\theta} \frac{\gamma-\alpha}{1-\alpha} - 1} = \lambda, \quad (\text{B.17})$$

from which we get equation (2.7).

The concavity condition that ensures that the set of emission taxes identified in equation (2.7) is unique and corresponds to a maximization of workers' welfare is:

$$\forall j \in C, \frac{\partial^2 \bar{u}}{\partial t_j^2} < 0. \quad (\text{B.18})$$

which can be expressed as:

$$\left(\frac{\gamma - \alpha}{1 - \alpha} \right)^2 \left(1 - \frac{1}{\theta} \right) L_j - \left(\frac{\gamma - \alpha}{1 - \alpha} \right) \left(1 - \frac{1}{\theta} \frac{\gamma - \alpha}{1 - \alpha} \right) < 0. \quad (\text{B.19})$$

In particular, the left-hand side term of inequation (B.19) is a linear function of each L_j with variations that only depend on the sign of $\theta - 1$. We are looking for the set of parameters for which the welfare function is concave over any set of cities (in particular, the concavity condition must hold for city sizes infinitely close to 0 or equal to 1, if there is

a unique city). Consequently, if $\theta > 1$, we only need $\left(\frac{\gamma-\alpha}{1-\alpha}\right)^2 \left(1 - \frac{1}{\theta}\right) < \left(\frac{\gamma-\alpha}{1-\alpha}\right) \left(1 - \frac{1}{\theta} \frac{\gamma-\alpha}{1-\alpha}\right)$. Respectively, if $0 < \theta \leq 1$, we only need $0 \leq \left(\frac{\gamma-\alpha}{1-\alpha}\right) \left(1 - \frac{1}{\theta} \frac{\gamma-\alpha}{1-\alpha}\right)$.

If $\theta > 1$, the condition holds if and only if $1 > \gamma > \alpha$. Indeed, if $\theta > 1$ the LHS is positive, so the RHS must also be positive so either: (1) $\frac{\gamma-\alpha}{1-\alpha} < 0$ and $1 < \frac{1}{\theta} \frac{\gamma-\alpha}{1-\alpha}$, or (2) $\frac{\gamma-\alpha}{1-\alpha} > 0$ and $1 > \frac{1}{\theta} \frac{\gamma-\alpha}{1-\alpha}$. The first option is not possible, because if $\frac{\gamma-\alpha}{1-\alpha} < 0$ then $\frac{1}{\theta} \frac{\gamma-\alpha}{1-\alpha} < 0$ (because $\theta > 0$ for an equilibrium to exist) and cannot be higher than one. The second option is necessary for the condition to hold. As a result, $\left(\frac{\gamma-\alpha}{1-\alpha}\right)^2 \left(1 - \frac{1}{\theta}\right) < \left(\frac{\gamma-\alpha}{1-\alpha}\right) \left(1 - \frac{1}{\theta} \frac{\gamma-\alpha}{1-\alpha}\right)$ is equivalent to $1 > \gamma > \alpha$ (combined with $\theta > 1$, it also implies $1 > \frac{1}{\theta} \frac{\gamma-\alpha}{1-\alpha}$). If $0 < \theta \leq 1$, the condition holds if and only if $1 > \gamma > \alpha$ and $\theta > \frac{\gamma-\alpha}{1-\alpha}$. Indeed, $\frac{\gamma-\alpha}{1-\alpha}$ is necessarily positive because $\theta > 0$ and we cannot have $\frac{\gamma-\alpha}{1-\alpha} < 0$ and $1 < \frac{1}{\theta} \frac{\gamma-\alpha}{1-\alpha}$. As a result, when $0 < \theta \leq 1$ the condition is equivalent to $\gamma > \alpha$ and $\theta > \frac{\gamma-\alpha}{1-\alpha}$ (which implies that $\gamma < 1$). Finally, note that $\theta > \frac{\gamma-\alpha}{1-\alpha}$ is implied when $1 > \gamma > \alpha$ and $\theta > 1$. As a conclusion, there is a unique solution to (B.13) if and only if $1 > \gamma > \alpha$ and $\theta > \frac{\gamma-\alpha}{1-\alpha}$.

B.4 Relaxing the Free Migration Assumption

Assume all the hypotheses of the simplified framework. The only difference is that, now, the distribution of workers across locations is fixed, meaning that for every location j , L_j is exogenously given. This assumption means that there are not adjustment of populations to local changes of wage or air quality. This corresponds to a “very” short-term version of the model where movement frictions are infinitely high.

In this small extension I consider that the central planner still tries to maximize welfare by choosing the distribution of relative emission taxes that maximizes the weighted average welfare per capita. Indeed, without free migration, welfare is not equalized across locations anymore. Still, the central planner has to maximize an objective function. Summing the local welfare levels using local populations as weight corresponds to computing average welfare per capita. It puts the same weight on each worker welfare. In the equilibrium with free migration, it corresponds exactly to the common level of welfare reached in all locations and $\tilde{u} = \bar{u}$.

We note this alternative objective function \tilde{u} and define it as:

$$\tilde{u} = \sum_{j \in C} L_j u_j = \sum_{j \in C} L_j a_j L_j^{-\delta} Z_j^{-\gamma} c_j \quad (\text{B.20})$$

This extension constitute the alternative to the free migration model. Both are extreme and in reality one may find a situation that is between these two cases. Introducing costs of moving that may allow endogenous distribution of labor across space would require adding a dynamic side to the model. Such assumption would also call for the companion extension to heterogenous workers to encompass the empirical fact that workers with higher income are more likely to adjust to local pollution than low income workers;

Keeping the other assumptions of the simple model, we have:

$$w_j = b_j L_j^\nu t_j^{-\frac{\alpha}{1-\alpha}} \quad (\text{B.21})$$

$$Z_j = \alpha \frac{w_j L_j}{t_j} \quad (\text{B.22})$$

Replacing the \tilde{u} , I get that the weighted average per capita welfare is (without constant):

$$\tilde{u} = \sum_{j \in C} L_j^{1-\theta} a_j b_j^{1-\gamma} t_j^{\frac{\gamma-\alpha}{1-\alpha}} \quad (\text{B.23})$$

with $\theta = \gamma + \delta - \nu(1 - \gamma)$.

Assuming that the central planner solves:

$$\max_{\{t_j\}_{j \in C}} \tilde{u} \quad \text{s.t.} \quad \sum_{j \in C} t_j = 1 \quad (\text{B.24})$$

She finds the following distribution of relative emission taxes:

$$\forall (j, i) \in C^2, \quad \frac{t_j^*}{t_i^*} = \left(\frac{a_j}{a_i} \right)^{\frac{1-\alpha}{1-\gamma}} \left(\frac{b_j}{b_i} \right)^{1-\alpha} \left(\frac{L_j}{L_i} \right)^{\frac{1-\alpha}{1-\gamma}(1-\theta)} \quad (\text{B.25})$$

Using estimated elasticities this imply that optimal taxes should be higher in larger cities and cities more productive and with higher amenities which is in line with the result obtained when free migration is assumed.

B.5 Estimation Appendix

B.5.1 Building Plant-Level Pollutants Emissions from Energy Surveys

Pollutant	PCC	Nb. Obs.
SO2	.1929*	1060
PM10	.7758*	45
NOX	.8192*	1734
COVNM	.2143*	4689
CO	.7898*	313

Table B.5.1: Pearson correlation coefficients between actual emissions values from the E-PRTR and estimated values from the EACEI survey (* indicates significance at the 1% level).

B.5.2 Estimation of the Sector Specific Emission Intensities

Assuming that for a plant i , installed in city c and producing goods from sector s , that has an intrinsic productivity equal to ϕ_{ics} production follows $q_{ics} = z_{ics}^{\alpha_s} l_{ics}^{1-\alpha_s} \phi_{ics}^{1-\alpha_s}$ and that it faces demand $q_{ics} = k_{cs} p_{ics}^{-\sigma_s}$, where p_{ics} is the price charged.

As I do not observe the distribution of inputs across destination, I abstract from the fact the plants may face distinct demand functions in each destination (in the case where trade between cities is costly). This assumption is mainly due to data constraint but considering that, in practice, within-country trade costs are low, it is also acceptable. The variable k_{cs} encompasses both the size of the output market and the level of competition on this market. What is important is that all plants within the same region face the same k_{cs} .

Writing the expression for plant revenues:

$$r_{ics} = p_{ics} q_{ics} \quad (\text{B.26})$$

$$= k_{cs}^{\frac{1}{\sigma_s}} q_{ics}^{\frac{\sigma_s-1}{\sigma_s}} \quad (\text{B.27})$$

$$= k_{cs}^{\frac{1}{\sigma_s}} \left(z_{ics}^{\alpha_s} l_{ics}^{1-\alpha_s} \phi_{ics}^{1-\alpha_s} \right)^{\frac{\sigma_s-1}{\sigma_s}} \quad (\text{B.28})$$

Based on the standard results from monopolistic competition with CES demand, markups are fixed and depend on σ_s . Based on the standard results of cost minimization under Cobb-Douglas production functions, plant-level expenditure shares across inputs are fixed and given by α_s . These two results combined allow the following formulation of the relationship between revenues and employment:

$$r_{ics} = \frac{\sigma_s}{\sigma_s - 1} \frac{1}{1 - \alpha_s} w_{ics} l_{ics} \quad (\text{B.29})$$

with w_{ics} the plant-specific unit cost of labor. In the main model of this paper I make the assumption that this cost is the wage and is the same across sectors and firms within a city. However, in this appendix on the estimation of emission elasticities, I temporarily relax this assumption to show that if, in reality, there are plant specific shocks on wages (which is likely to be the case) it does not threaten my estimation strategy.

Substituting revenues in (B.26) using (B.29):

$$\frac{\sigma_s}{\sigma_s - 1} \frac{1}{1 - \alpha_s} w_{ics} l_{ics} = k_{cs}^{\frac{1}{\sigma_s}} \left(z_{ics}^{\alpha_s} l_{ics}^{1-\alpha_s} \phi_{ics}^{1-\alpha_s} \right)^{\frac{\sigma_s-1}{\sigma_s}} \quad (\text{B.30})$$

$$\Rightarrow l_{ics}^{\frac{1-(1-\alpha_s)\frac{\sigma_s-1}{\sigma_s}}{\sigma_s}} = \frac{\sigma_s - 1}{\sigma_s} (1 - \alpha_s) k_{cs}^{\frac{1}{\sigma_s}} w_{ics}^{-1} \phi_{ics}^{(1-\alpha_s)\frac{\sigma_s-1}{\sigma_s}} z_{ics}^{\alpha_s \frac{\sigma_s-1}{\sigma_s}} \quad (\text{B.31})$$

$$\Rightarrow l_{ics} = \left[\frac{\sigma_s - 1}{\sigma_s} (1 - \alpha_s) k_{cs}^{\frac{1}{\sigma_s}} \right]^{\frac{1}{1-(1-\alpha_s)\frac{\sigma_s-1}{\sigma_s}}} z_{ics}^{\frac{\alpha_s \frac{\sigma_s-1}{\sigma_s}}{1-(1-\alpha_s)\frac{\sigma_s-1}{\sigma_s}}} \left(w_{ics}^{-1} \phi_{ics}^{(1-\alpha_s)\frac{\sigma_s-1}{\sigma_s}} \right)^{\frac{1}{1-(1-\alpha_s)\frac{\sigma_s-1}{\sigma_s}}} \quad (\text{B.32})$$

which can be simplified as:

$$l_{ics} = \tilde{k}_{cs}(k_{cs}) z_{ics}^{\frac{\alpha_s(\sigma_s-1)}{1+\alpha_s(\sigma_s-1)}} \tilde{\phi}_{ics}(\phi_{ics}, w_{ics}) \quad (\text{B.33})$$

with \tilde{k}_{cs} a function only of k_{cs} , so common to all plants within a city-sector pair, and $\tilde{\phi}_{ics}$ a function of plant specific productivity and wage shocks.

Using the previous relationship, I estimate:

$$\log l_{icst} = \beta_{0s} + \beta_{1s} \log z_{icst} + \mu_t + \rho_{ics} + \epsilon_{icst} \quad (\text{B.34})$$

to retrieve $\beta_{1s} = \frac{\alpha_s(\sigma_s-1)}{1+\alpha_s(\sigma_s-1)}$. In introduce time fixed-effects, μ_t , and plant fixed effects, ρ_{ics} .

However, unobserved shocks on productivity and wages at the plant level are implicitly included in the residual ϵ_{icst} and are correlated with input use l_{icst} and z_{icst} , OLS estimates are biased. This is a sort of “transmission bias” which is frequent in the literature estimating production functions. The only difference here is that I do not estimate the production function using revenues on the left-hand side. This is due to not observing revenues at the plant level but only at the firm level. Indeed, I also only observe emissions

of atmospheric pollutants at the plant level for subsets of plants within firms. Using plant and firm level observations in the same equation would force me to drop observations from multi-plants firms. As a result, the estimation sample would be too small.

I rely on an IV approach to estimate β_{1s} . I build an instrument that uses exogenous shocks on fuel prices as source of variation of plant-level emissions. This instrument is a fixed-weight energy price index that measures the plant-specific exposure to variation in fuels prices based on each plant distribution of energy consumption across fuel types in the first period where it is observed. I build on Sato et al. (2019) to compute this instrumental variable.

$$FEPI_{ist} = \sum_{f \in \Omega_{fuels}} \omega_{f,ist_0} p_{f,st} \quad (\text{B.35})$$

where $FEPI_{ist}$ is the plant-specific energy price index build from plant i share of energy expenditures in fuels $f \in \Omega_{fuels}$ (coal, natural gas, electricity, etc.) in period t_0 and $p_{f,st}$ the specific fuel price common to all plants in sector s in period t . The exclusion restriction requires in this case that plant specific shares of expenditures across fuels are not correlated with unobserved shocks that are included in ϵ_{icst} . Given (B.33), shocks in the residual only comes from demand shocks that are common to all plants, shocks on plant specific wage and shocks on plant specific productivity. Therefore, the exclusion restriction is, in this case, likely to be verified.

B.6 Application to France Appendix

B.6.1 Sources of Equilibrium Data

Empirical Distribution of Wages Across Commuting Zones: I use the distribution of average hourly wages at the commuting zone level for 2012 from the INSEE dataset publicly available at <https://www.insee.fr/fr/statistiques/2021266>. This data is aggregated geographically by the INSEE using worker level data from the Social Data Annual Declaration (DADS - *Déclaration annuelle des données sociales*). I use the 2010 commuting zones (*Zones d'Emploi*) definition and normalize mean hourly local wages so that the average over all commuting zones is equal to one.

Empirical Distribution of Populations Across Commuting Zones: I use the distribution of working population at the “*Communes*” level for 2012 from the INSEE dataset publicly available at <https://www.insee.fr/fr/statistiques/2128672>. The “*Communes*” is one of the most disaggregated administrative geographic unit in France. There are around 36,000 of them over the territory. The 2010 commuting zone’s (*Zones d'Emploi*) definition is a partition of the set of French “*Communes*” in around 300 areas where inhabitants both work and live. I use the correspondence table between “*Communes*” and *Zones d'Emploi* from the INSEE and publicly available at <https://www.insee.fr/fr/information/2114596>. I sum working populations at the “*Communes*” level to get aggregated working populations at the *Zones d'Emploi* level. Finally I normalize local populations so that the total national working population over the whole set of commuting zones is equal to one.

Empirical Distribution of Emissions Across Commuting Zones: I use data from the National Spatialized Inventory (“*Inventaire National Spatialisé*”) which is built by the French ministry in charge of environmental issues. The inventory is publicly available at <http://emissions-air.developpement-durable.gouv.fr/>. I exported data for the 2012 platform (data is also available for 2004 and 2007) for PM10, PM2.5, SO2, NOx and COVNM. These datasets provide amounts of pollutants emitted disaggregated at the “*Communes*” level and across broad sectors of activities following the Selected Nomenclature for Air Pollution (SNAP). I aggregate emissions at the commuting zone level by using the correspondence table corresponding to the 2010 definition of commuting zones and summing emissions within commuting zones across “*Communes*”. I keep emissions from codes 3 and 4 of the SNAP which correspond to emissions from industrial combustion plants and industrial processes without combustion. These two codes broadly correspond to emissions due to the manufacturing industries. I aggregate over these two emission

sectors by summing emissions. Figure(B.6.1) displays the correlation between levels of emissions of available pollutants. One can observe that the correlation is pretty strong. I only retain the distribution of PM10 emissions as the empirical data for the distribution of my model's representative pollutant across commuting zones.

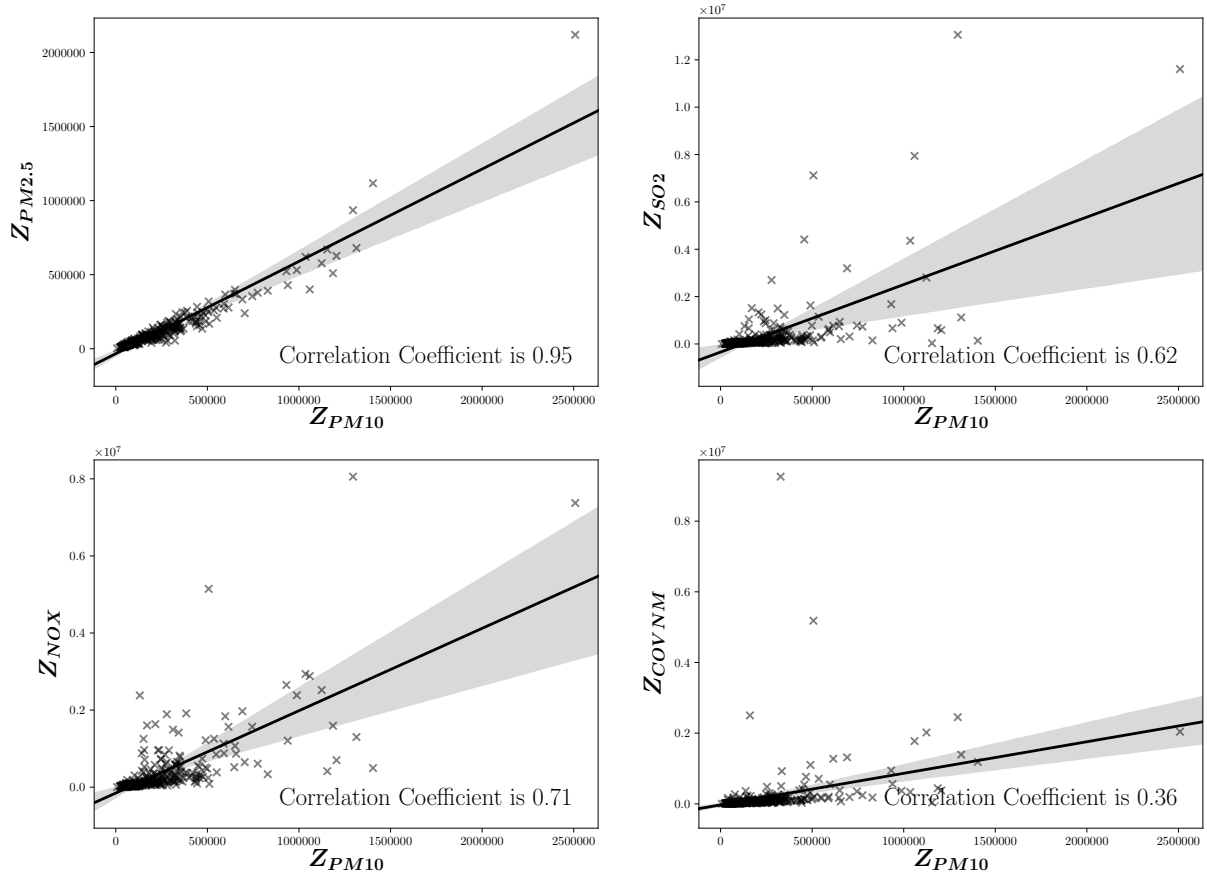


Figure B.6.1: Correlation Between Emissions of Harmful Industrial Pollutants at the Commuting Zone Level

B.6.2 Descriptive Statistics on Equilibrium Data

B.6.3 Descriptive Statistics on Computed Local Characteristics

	Wages	Populations	PM2.5 Emissions
count	303.000	303.000	303.000
mean	1.000	0.003	1.000
std	0.101	0.009	1.072
min	0.847	0.000	0.027
25%	0.938	0.001	0.352
50%	0.973	0.001	0.697
75%	1.037	0.003	1.256
max	1.585	0.136	9.830

Table B.6.1: Descriptive Statistics on the Spatial Distributions of Wages, Populations and Emissions across French Cities

Highest Populations	Highest Wages	Highest Emissions
Paris	Saint-Quentin-en-Yvelines	Dunkerque
Lyon	Paris	Thionville
Toulouse	Saclay	Bordeaux
Roissy - Sud Picardie	Rambouillet	Istres - Martigues
Bordeaux	Versailles	Nantes
Marseille - Aubagne	Poissy	Toulouse
Saclay	Cergy	Saint-Dié-des-Vosges
Nantes	Marne-la-Vallée	Lyon
Lille	Créteil	Marseille - Aubagne
Rennes	Aix-en-Provence	Nancy

Table B.6.2: Top Ten Cities with Highest Populations, Wages and Emissions

Lowest Populations	Lowest Wages	Lowest Emissions
Sartène - Propriano	Saint-Flour	Corte
Corte	Mauriac	Issoudun
Ghisonaccia - Aléria	Sartène - Propriano	Loches
Le Blanc	Le Blanc	Figeac
Ambert	Sarlat-la-Canéda	L'Aigle
Calvi - L'Île-Rousse	Brioude	Vire
Loches	Lozère	Menton - Vallée de la Roya
Issoudun	Villeneuve-sur-Lot	Ghisonaccia - Aléria
Avallon	Saint-Amand-Montrond	Le Blanc
Porto-Vecchio	Ussel	Mauriac

Table B.6.3: Bottom Ten Cities with Lowest Populations, Wages and Emissions

	Amenities	Productivities	Emission Costs
count	303.000	303.000	303.000
mean	0.003	1.000	1.000
std	0.010	0.160	1.173
min	0.000	0.785	0.075
25%	0.000	0.906	0.521
50%	0.000	0.956	0.812
75%	0.002	1.045	1.153
max	0.087	2.044	18.059

Table B.6.4: Descriptive Statistics on the Spatial Distributions of Amenities, Productivities and Emission Costs across French Cities

Highest Amenities	Highest Productivities	Highest Emission Costs
Bordeaux	Saint-Quentin-en-Yvelines	Paris
Dunkerque	Paris	Saclay
Nantes	Saclay	Orly
Lyon	Rambouillet	Marne-la-Vallée
Marseille - Aubagne	Versailles	Créteil
Toulouse	Poissy	Cergy
Avignon	Cergy	Lille
Roissy - Sud Picardie	Marne-la-Vallée	Poissy
Nancy	Créteil	Lyon
Rennes	Aix-en-Provence	Cannes - Antibes

Table B.6.5: Top Ten Cities with Highest Amenities, Productivities and Emission Costs

Lowest Amenities	Lowest Productivities	Lowest Emission Costs
Saint-Quentin-en-Yvelines	Saint-Flour	Saint-Dié-des-Vosges
Corte	Mauriac	Dunkerque
Versailles	Sartène - Propriano	Thionville
Rambouillet	Villeneuve-sur-Lot	Istres - Martigues
Issoudun	Sarlat-la-Canéda	Maurienne
Loches	Brioude	La Teste-de-Buch
L'Aigle	Lozère	Sartène - Propriano
Figeac	Le Blanc	Péronne
Étampes	Saint-Amand-Montrond	Dole
Ghisonaccia - Aléria	Bressuire	Jonzac - Barbezieux-Saint-Hilaire

Table B.6.6: Top Ten Cities with Lowest Amenities, Productivities and Emission Costs




Figure B.6.2: Histograms of Spatial Distributions of Wages, Populations and Emissions across French Cities
All variables are in log.




Figure B.6.3: Cumulative Distributions of Populations and Emissions across French Cities
Cities are ranked by size, from smallest to largest, and the normalized cumulative sum of populations and emissions is plotted.

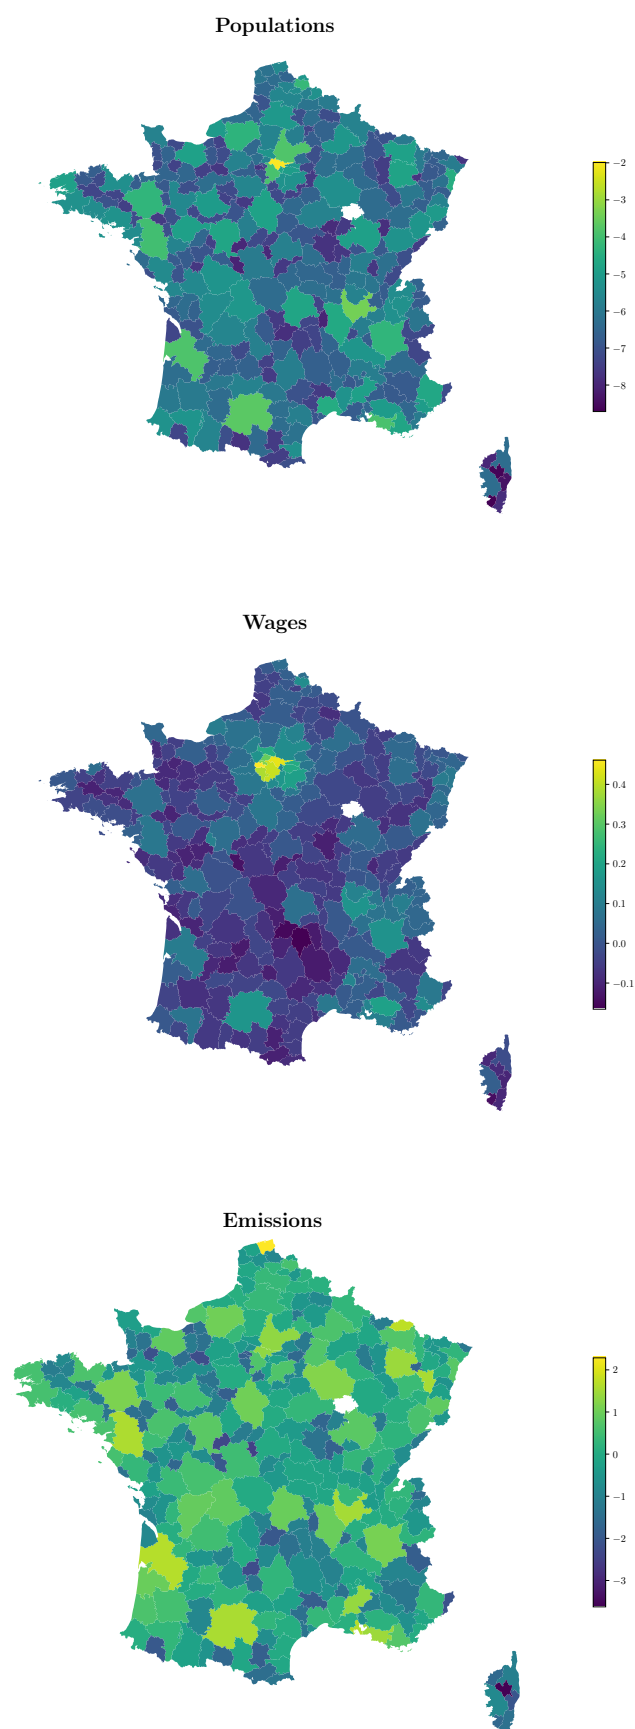


Figure B.6.4: Spatial Distributions of Wages, Populations and Emissions across French Cities


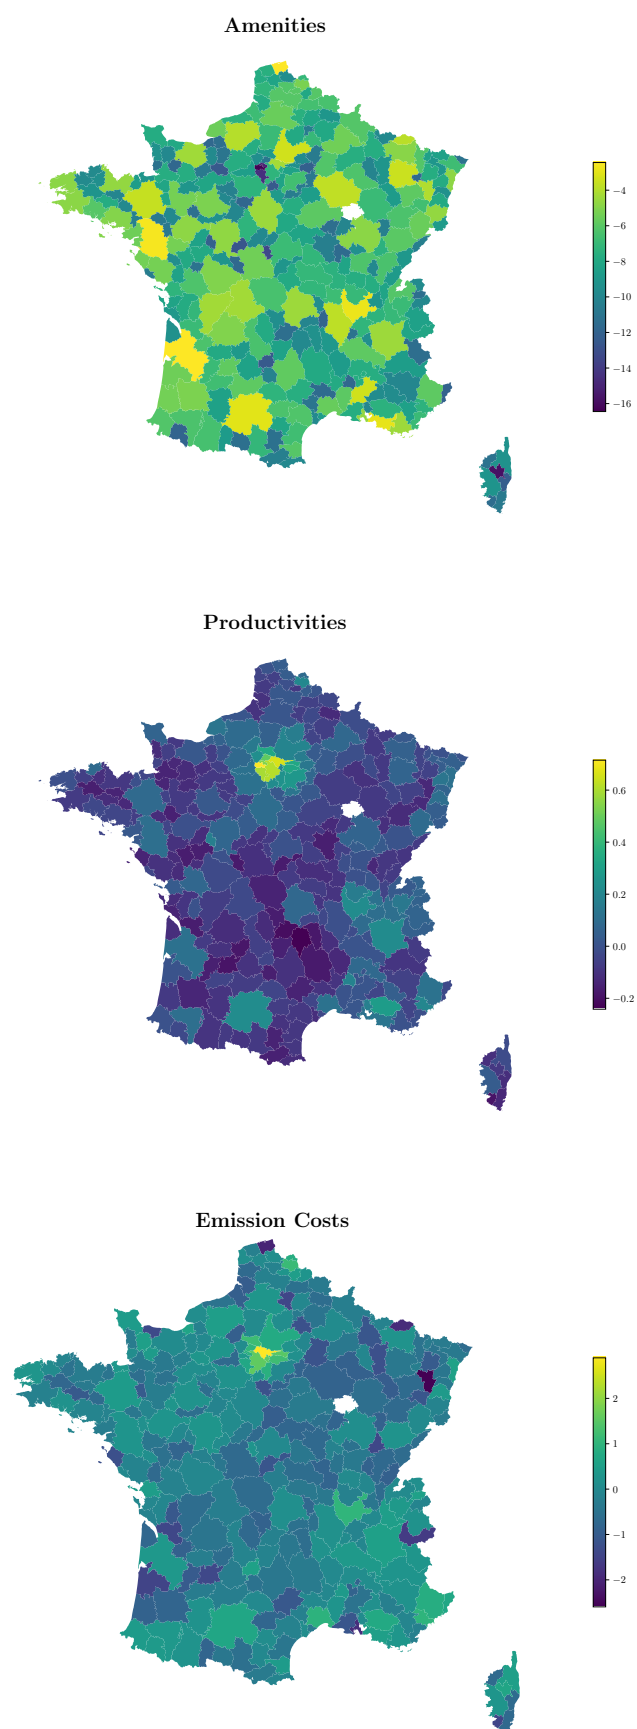


Figure B.6.5: Histograms of Spatial Distributions of Wages, Populations and Emissions across French Cities
All variables are in log.



All variables are in log.

Figure B.6.6: Spatial Distributions of Amenities, Productivities and Emission Costs across French Cities

Appendix C

Appendix to Chapter 3

C.0.1 Timing of Exports

Mean Nb. Months	Std. Deviation ∈					1 ^{ce}	2 ^{ce}	Total
	[0, 1[[1, 2[[2, 3[[3, 4[≥ 4			
∈ [1, 2[19.59	4.23	0.67	0.18	0.06	.	1.54	26.27
	<i>96.89</i>	<i>1.31</i>	<i>0.24</i>	<i>0.13</i>	<i>0.03</i>	.	<i>0.01</i>	<i>98.61</i>
∈ [2, 3[0.47	3.39	2.56	0.79	0.67	.	0.93	8.80
	<i>0.00</i>	<i>0.38</i>	<i>0.16</i>	<i>0.10</i>	<i>0.15</i>	.	<i>0.00</i>	<i>0.78</i>
∈ [3, 4[0.18	0.62	1.91	1.38	1.22	.	0.69	6.00
	<i>0.00</i>	<i>0.02</i>	<i>0.05</i>	<i>0.06</i>	<i>0.10</i>	.	<i>0.00</i>	<i>0.22</i>
∈ [4, 5[0.09	0.20	0.73	1.06	1.81	.	0.55	4.44
	<i>0.00</i>	<i>0.00</i>	<i>0.01</i>	<i>0.02</i>	<i>0.06</i>	.	<i>0.00</i>	<i>0.09</i>
∈ [5, 6[0.07	0.09	0.28	0.54	2.01	.	0.51	3.50
	<i>0.00</i>	<i>0.00</i>	<i>0.00</i>	<i>0.01</i>	<i>0.08</i>	.	<i>0.00</i>	<i>0.09</i>
∈ [6, 7[0.06	0.07	0.16	0.25	1.89	.	0.42	2.86
	<i>0.00</i>	<i>0.00</i>	<i>0.00</i>	<i>0.00</i>	<i>0.03</i>	.	<i>0.00</i>	<i>0.04</i>
≥ 7	0.27	0.21	0.40	0.43	10.80	.	5.93	18.04
	<i>0.00</i>	<i>0.00</i>	<i>0.00</i>	<i>0.00</i>	<i>0.09</i>	.	<i>0.03</i>	<i>0.12</i>
Total	20.72	8.82	6.71	4.64	18.47	30.08	10.58	100.00
	<i>96.89</i>	<i>1.70</i>	<i>0.45</i>	<i>0.32</i>	<i>0.54</i>	<i>0.05</i>	<i>0.04</i>	<i>100.00</i>

This table displays the repartition of exporting firms according to the number of months between two of their exports. For each categorie, there is the percentage of the total number of distinct firms appearing in the trade panel (217412 firms) and in italic if the fraction of total export values that this categorie represents.

Table C.0.1: Descriptive Statistics - Trade Data : Timing Exports

C.0.2 Extended Tables - Baseline Results

	Value		Quantity		Export Status	
	OLS (1)	2SLS (2)	OLS (3)	2SLS (4)	OLS (5)	2SLS (6)
$\log \overline{PM2.5}_{t \in [-2,0]}$	-0.102* (0.062)	-0.457 (0.333)	-0.043 (0.053)	-0.647** (0.321)	0.000 (.)	-0.003 (0.007)
Obs.	2,797,719	2,797,719	2,766,517	2,766,517	2,797,719	9,969,875
FS F-Stat.		80		80		11,108

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ - Standard errors, clustered by firm-activity, are reported in parenthesis. PM2.5 concentrations are measured in $\mu g/m^3$. The dependent variable is: for columns (1) and (2) the logarithm of export values in euros, for columns (3) and (4) the logarithm of the export quantities in kilograms, and for columns (5) and (6) a dummy variables indicating whether the firm exported in this given month. All regressions control for weather variables temperature and precipitation. Observations are weighted using export values.

Table C.0.2: Effects of PM2.5 Pollution on Firm-level Exports: Weighted Regression

C.0.3 Heterogeneity Across Manufacturing Activities

Activities	Value	Quantity	Export Status
Agriculture <i>196,172 obs.</i>	-.3 (.23)	-.16 (.24)	-.03 (.03)
Arts, entertainment and recreation <i>23,317 obs.</i>	.57 (1.03)	1 (1.1)	.13** (.05)
Communication <i>186,282 obs.</i>	.16 (.41)	-.39 (.4)	-.02 (.02)
Energy <i>152 obs.</i>	-5.42 (6.57)	-5.42 (6.57)	-.25 (.17)
Manufacturing - Basic Metals <i>150,226 obs.</i>	.12 (.6)	-.16 (.7)	-.03 (.04)
Manufacturing - Beverages <i>260,664 obs.</i>	-.37*** (.14)	-.43*** (.14)	-.05** (.02)
Manufacturing - Chemicals <i>372,096 obs.</i>	.13 (.17)	.11 (.18)	.01 (.02)
Manufacturing - Coke and Refined Petroleum Products <i>31,567 obs.</i>	.63 (.81)	.92 (.87)	-.03 (.06)
Manufacturing - Computers, Electronic, & Optical <i>363,359 obs.</i>	-.14 (.21)	-.22 (.23)	-.01 (.01)
Manufacturing - Electrical Equipment <i>345,965 obs.</i>	-.07 (.27)	-.05 (.3)	-.01 (.02)
Manufacturing - Fabricated Metal Products <i>452,184 obs.</i>	.67 (.46)	.39 (.51)	-.05** (.02)
Manufacturing - Food <i>247,291 obs.</i>	-.14 (.19)	-.28 (.21)	0 (.02)
Manufacturing - Furniture <i>115,387 obs.</i>	-.65 (.71)	-.63 (.68)	.03 (.02)
Manufacturing - Leather <i>145,115 obs.</i>	.09 (.38)	.22 (.41)	0 (.02)
Manufacturing - Machinery <i>558,333 obs.</i>	.03 (.28)	.09 (.33)	-.03* (.01)
Manufacturing - Motor Vehicles <i>200,772 obs.</i>	.05 (.38)	.16 (.44)	-.01 (.01)
Manufacturing - Non-metallic Mineral Products <i>179,739 obs.</i>	.43 (.41)	.05 (.41)	-.01 (.02)
Manufacturing - Other Manufacturing <i>282,914 obs.</i>	-.59** (.28)	-.53* (.31)	-.04*** (.02)
Manufacturing - Other Transport Equipment <i>74,811 obs.</i>	.37 (.51)	.55 (.56)	-.01 (.02)
Manufacturing - Paper <i>195,723 obs.</i>	-.08 (.34)	.15 (.34)	-.02 (.02)
Manufacturing - Pharmaceuticals <i>63,485 obs.</i>	.21 (.41)	.65 (.45)	-.01 (.04)
Manufacturing - Recorded Media <i>3,589 obs.</i>	.15 (2.09)	-.17 (2.67)	.08 (.1)
Manufacturing - Rubber & Plastics <i>424,311 obs.</i>	-.13 (.23)	-.1 (.25)	-.01 (.02)
Manufacturing - Textiles <i>213,930 obs.</i>	.11 (.29)	.18 (.28)	-.02 (.02)
Manufacturing - Tobacco <i>856 obs.</i>	-3.36 (3.2)	-1.41 (3.36)	-.23 (.22)
Manufacturing - Wearing Apparels <i>206,639 obs.</i>	-1.05*** (.36)	-.99*** (.36)	-.02 (.02)
Manufacturing - Wood <i>136,617 obs.</i>	-.24 (.42)	-.67 (.42)	.01 (.03)
Mining <i>39,646 obs.</i>	.24 (.66)	-.18 (.85)	.05 (.05)
Other service activities <i>4 obs.</i>	49.39*** (0)	-32.31*** (0)	-.43 (.27)
Scientific and technical activities <i>3,251 obs.</i>	-1.95 (2.84)	1.74 (2.82)	.11 (.07)
Water & waste management <i>59,117 obs.</i>	.16 (.62)	.02 (.69)	.07 (.06)

* p<0.1, ** p<0.05, *** p<0.01 - This table presents the results of estimating equation (3.2) at the activity level. We split firm-level exports according to NAF 2-digits level. As in the baseline strategy, we estimate the effect of exposure to PM2.5 concentrations averaged over the three months preceding the export. Numbers are in bold when significant. We report the number of positive exports for each activity.

Table C.0.3: Effects of Pollution on Exports Across Activities