Inequality in Exposure to Air Pollution in France: Measurement and Impact of a City-Level Public Policy

Master's Thesis

Pascale Champalaune
M2 APE, Paris School of Economics

September 2020

Supervisors: Lucas Chancel & Thomas Piketty

Abstract

I combine measures of neighbourhood characteristics with high-resolution remote-sensing data to provide the first national-scale study of cross-sectional and longitudinal inequality in exposure to fine particulate matter ($PM_{2.5}$) in France. Descriptive evidence indicates that, at the national level, there is a U-shaped relationship between income and $PM_{2.5}$ exposure, which is not reflected within urban areas. Fixed-effect models confirm that on average, higher neighbourhood income is associated with lower exposure. Longitudinal inequality measures suggest that recent air quality improvements accrued predominantly to areas that had a lower initial exposure, and intermediate income. I then exploit a change in air quality schemes at the level of urban areas in an event-study framework, so as to shed light on potentially unequal benefits from the induced reduction in exposure. I find that initially lower-income areas received smaller benefits from the policy change, and quantile regression estimates suggest that exposure decreased more in less polluted areas. As some results are sensitive to formally accounting for spatial autocorrelation, this study also underlines the need to pay specific attention to this issue when measuring environmental inequality.

JEL Codes: I14, Q53, Q56, R23

Keywords: Air pollution, Environmental Economics, Inequality, Spatial autocorrelation

I thank François Fontaine, François Libois, and Louis Sirugue for their time and very helpful comments and suggestions. I am particularly grateful to Lucas Chancel and Thomas Piketty, who supervised this Master's Thesis, and to Camille Hémet, who accepted to act as referee, for their support.

Contents

1 Introduction

As public goods, environmental amenities are an integral part of the real income of households. This makes their distribution part of the general panorama of inequality. More specifically, unequal access to clean air may reinforce preexisting health inequalities, both across space and across income groups. Despite both policy-makers and the general public showing growing concern about it, the French economic literature on the issue is still in its infancy. In addition, over the last two decades, the country recorded substantial air quality improvements, driven in part by long-term technological development, and in part by dedicated policies. It is thus necessary to evaluate whether the latter benefitted all individuals to the same extent.

1.1 Main contributions

The contribution of this study to the academic literature is threefold. First, I provide the first nationwide evidence of cross-sectional and longitudinal inequality in exposure to fine particulate matter ($PM_{2.5}$). In mainland France, the associated health burden was recently estimated to 48,000 early deaths a year, which is equivalent to 2-year reduction of life expectancy at 30 on average (Medina et al., 2016). Albeit these detrimental health impacts, this pollutant has been overlooked by the French literature on environmental inequality, partly due to the fact that there were very few $PM_{2.5}$ monitors in France prior to the 2010s. I thus contribute in filling this gap by exploiting high-resolution satellite data, coupled with census block-level INSEE data. This allows the study not only to cover the whole of metropolitan France, but to do so at a fine spatial granularity. This is particularly crucial so as to limit the risk of ecological fallacy, i.e., to avoid making erroneous inference on individual correlations based on neighbourhood correlations (Banzhaf et al., 2019).

I show that, at the national level, there is a U-shaped relationship between $PM_{2.5}$ exposure and income, which is not reflected within urban areas, where only the lowest income deciles face a disproportionate burden of exposure. I tackle the omitted variable bias related to unobserved neighbourhood heterogeneity using fixed-effect models, and confirm that there is indeed a negative relationship between the two variables of interest in France. The results also suggest that the share of immigrants is positively associated with $PM_{2.5}$ concentration, hinting at an ethnic gap in exposure reminiscent of that observed in the United States. I also contribute to the literature by providing longitudinal measures of inequality, and present evidence that census blocks that benefitted from the smallest air quality improvements are those that had the highest initial exposure, and that were located either at the lower or upper end of the distribution of income. This study thus uncovers that despite the fact that $PM_{2.5}$ exposure dropped throughout the country between 2006 and 2016, there is growing inequality in its distribution.

My second contribution consists in taking advantage of a policy change that occurred at the level of urban areas, in order to investigate whether new $PM_{2.5}$ regulation played a role in this evolution. In France, since the 1996 LAURE law, part of the regulation of air quality is effected by the local authorities of the largest cities, through *Plans de Protection de l'Atmosphère*. Following the 2008 EU Directive on ambient air quality and cleaner air for Europe, urban areas had to revise their schemes so as to newly include measures aimed at reducing $PM_{2.5}$ concentration. Given that the years of implementation vary from 2012 to 2016, I make use of an event-study design in order to analyse how inequality in exposure to air pollution evolved as a consequence of this change. In line with the abovementioned findings, I demonstrate that these revised plans predominantly benefitted initially advantaged neighbourhoods, i.e., those whose income lay above the median. In addition, I provide quantile regression estimates which suggest that lower quantiles of exposure received larger air quality improvements.

As a third contribution, throughout this study, I pay specific attention to the robustness of the results to formally accounting for spatial autocorrelation, as opposed to most of the economic literature in the field of environmental inequality. Concomitantly, I provide a concise review of the methods used in Public Health studies, and argue that the spatial lag model, which seems the most widespread, relies on arbitrary parametric assumptions that may threaten the validity of the estimates. I thus rely on a non-parametric approach to model spatial autocorrelation, using a smoothing spline of census block geographic coordinates. I show that after controlling for neighbourhood location, while the pollution-income relationship remains similar, there is a stronger positive correlation between the share of immigrants and $PM_{2.5}$ exposure. Hence, the sensitivity of the results to controlling for neighbourhood location is consistent with observed segregation patterns.

1.2 Contextual elements

Ever since the publication of a seminal report that brought to light significant racial disparities in exposure to toxic industrial waste, at the disadvantage of Blacks (Chavis and Lee, 1987), and the subsequent emergence of the environmental justice movement in the 1980s, research in environmental inequality has flourished in the United States. A considerable body of literature has unambiguously shown that there is pervasive racial and ethnic inequality in exposure to various environmental risks and hazards, be it from air, water or soil pollution (see, e.g., reviews by Hajat et al., 2015; Mohai et al., 2009). Focusing on the former, it was demonstrated that ethnic minorities face a higher burden of exposure to air pollutants, from those of the Toxics Release Inventory (TRI) (Boyce et al., 2016; Downey, 2007; Zwickl et al., 2014), to particulate matter or ozone (Bell and Ebisu, 2012; Brochu et al., 2011). Lower-income individuals or neighbourhoods were also repeatedly shown to have a significantly higher likelihood of being exposed to high levels of pollution, although this discrepancy was generally found to be less acute than racial and ethnic differences (Bell and Ebisu, 2012; Boyce et al., 2016; Muller et al., 2018).

But as opposed to the US, where the notion of environmental inequality has been part of the public debate and the policy arsenal for several decades, the French literature on the issue is considerably scanter. The topic was mostly addressed by Public Health scholars, who provided cross-sectional evidence of inequality in exposure to air pollution based on the level of neighbourhood deprivation, an index which encapsulates both objective and subjective poverty measures (Pornet et al., 2012). Their results suggest that patterns of inequality may be heterogeneous across French cities. While it was found that disadvantaged neighbourhoods were relatively more exposed to air pollution in Marseille (Padilla et al., 2014), Strasbourg (Havard et al., 2009), Brittany (Bertin et al., 2015) or Dunkerque (Occelli et al., 2016), the relationship appeared to be reversed in Paris, and U-shaped in Lyon and Grenoble (Padilla et al., 2014; Morelli et al., 2016). These mixed results may be attributed to the fact that segregation levels are lower (Quillian and Lagrange, 2016) and urban design is more diverse in France than in the US, where most cities are built on a Black-centre/White-suburb structure,[1] which favours both racial and income gaps in exposure to air pollution.

One of the main limitations of most of these Public Health studies relates to their spatial scope. Indeed, the aforementioned papers focused on one administrative region (e.g., Bertin et al., 2015), specific urban areas (e.g., Padilla et al., 2014), or even one single municipality (e.g., Havard et al., 2009). This limited spatial extent has the benefit of bringing to light heterogeneous patterns of inequality in exposure to air pollution. Nonetheless, it also implies that not only does a number

[1]According to recent work, this historical structure is evolving, as urban centres attract more and more Whites and college-educated workers, which helped reduce the racial gap in fine particulate matter exposure between 2000 and 2015 (Currie et al., 2020).

urban areas and rural localities await to be studied, but national trends also remain to be identified. To my knowledge, two studies provided inceptive evidence of nationwide disparities in exposure to ambient air pollution. First, Ouidir et al. (2017) exploit the precision of the ELFE mother-child cohort data and link it to pollution and deprivation measures. They find that particulate matter (PM_{10}), nitrogen dioxide (NO_2) and $PM_{2.5}$ exposure during pregnancy is positively correlated with deprivation in urban areas, while the relationship is U-shaped in rural areas. Second, Lavaine (2015) investigates the extent to which socioeconomic characteristics and NO_2, PM_{10} and ozone (O_3) concentrations affect the total mortality rate of French *départements*, and finds that a higher NO_2 concentration has a greater impact on mortality in poorer areas than in wealthier ones. So far as I am aware, it is the only French study related to environmental inequality that aims to determine a causal impact: the low spatial resolution is concomitantly a weakness and a strength, as it may conceal large heterogeneity within *départements*, but allows to limit bias linked to residential sorting.

Indeed, the fact that substantial gaps are yet to be filled in terms of measurement and understanding of environmental inequality in France is at least partly attributable to how challenging it is to circumvent the endogeneity bias of the pollution-income relationship. As framed by Laurent (2009), the spatial dimension of social inequality already implies that it is not straightforward to discern it from local inequalities; adding environmental considerations further muddies the waters. Banzhaf et al. (2019) provide a summary of potential explanations to the disproportionate exposure of disadvantaged populations to pollution in a recent review: selective firm siting, selective neighbourhood sorting, or a market-like combination of the two. On the one hand, selective siting assumes that factories and other polluting economic activities disproportionately choose to locate within or in the vicinity of poorer areas. On the other hand, selective neighbourhood sorting may occur if better air quality translates into higher housing prices, with poorer (resp., wealthier) households self-selecting into more (resp., less) polluted areas. There is evidence favouring the siting hypothesis in the US (Mohai and Saha, 2015; Pastor et al., 2001), as well as in France in the context of incinerator building (Laurian and Funderburg, 2014) during the 1960-1990 period. One may imagine that nowadays, in part due to greater awareness of environmental issues, the sorting hypothesis may play a greater role. However, very few US studies were able to tackle this endogeneity issue, and have diverging conclusions (Gamper-Rabindran and Timmins, 2011; Mohai and Saha, 2015; Voorheis, 2017), while, to the best of my knowledge, there exists none in France. Although the study of environmental inequality is intrinsically interdisciplinary, the focus that Economics places on causal relationships gives it an advantage in disembroiling the mechanisms at stake here (Banzhaf et al., 2019). The tools developed by economists are also key to the evaluation of the causal impact of specific policies on environmental inequality patterns, which is what this study proceeds to do.

In particular, inequality in exposure to fine particulate matter ($PM_{2.5}$), and the influence that public policies have over it, are worth investigating. Indeed, the French literature has for now mostly focused on nitrogen dioxide (NO_2), which is often considered as an indicator for traffic (e.g., Bertin et al., 2015; Lavaine, 2015; Padilla et al., 2014), or particulate matter (PM_{10}) (e.g., Havard et al., 2008; Lavaine, 2015). This is partly due to data availability reasons, since $PM_{2.5}$ concentration was not regulated prior to 2009, which implies that the number of appropriate monitors likely was insufficient prior to the mid 2010s (Le Moullec, 2018). The fact that fine particulate matter was listed as a controlled pollutant quite recently does not entail that it is quite innocuous and uncommon; it is in fact harmful and ubiquitous. Indeed, it is constituted of aerosol particles whose diameter is smaller than 2.5 µm, hence their name, that are emitted through combustion, which implies that they are produced by various sources, from domestic wood burning to traffic, from industrial facilities to agriculture. Their small size implies that they penetrate and remain deep in the lungs, thus causing asthma and other cardiovascular and respiratory diseases, and making them the fifth-leading mortality risk factor worldwide (Cohen et al., 2017). It was also shown that they have detrimental impacts on short-term worker productivity (Graff Zivin and Neidell, 2012; Chang et al., 2016), even for indoor employees, as they easily enter buildings (Chang et al., 2019). Ebenstein et al. (2016) also provide evidence that early exposure to $PM_{2.5}$ is associated with lower human capital attainment and earnings.

This Master's Thesis is organised as follows. Section 2 presents the data on neighbourhood characteristics and fine particulate matter concentration. Section 3 describes the cross-sectional patterns of inequality in exposure to $PM_{2.5}$ in France, first using graphical evidence, then turning to fixed-effect models and their robustness to controlling for spatial autocorrelation. It also provides evidence on the evolution of these inequalities, using both vertical and horizontal equity measures. Section 4 focuses on the differential effects of the adoption of revised *Plans de Protection de l'Atmosphère*. After providing some elements on the regulative context, it describes the event-study and quantile regression approaches and discusses the results, as well as their sensitivity to the use of a spatial model. Section 5 concludes.

2 Data

2.1 Income and neighbourhood characteristics

I make use of IRIS-level data made publicly available by INSEE, the French National Institute for Statistics and Economic Studies. IRIS (*Ilôts Regroupés pour l'Information Statistique*, or aggregated units for statistical information) were designed by INSEE so as to prepare for the dissemination of the data collected through the 1999 Census. They were built using criteria based on both administrative boundaries and demographic characteristics, so as not only to be tractable in the long-term, but also to constitute a sufficiently homogeneous fraction of a municipality in terms of housing and land use. There are 3 categories of IRIS, starting with residential IRIS, which are home to between 1,800 and 5,000 inhabitants and make up 92% of the total number of IRIS. Business IRIS cluster at least 1,000 workers, and no less than twice as many workers as inhabitants. Finally, miscellaneous units are large and specific areas that are sparsely populated, such as parks, forests or harbours, and represent 3% of IRIS. To this day, all cities that are home to more than 10,000 inhabitants, and the majority of municipalities with 5,000 to 10,000 inhabitants are divided into IRIS. By extension, and so as to cover the entire French territory, every municipality that is not divided into IRIS is considered as an IRIS.[2] In 2016, there were on average 1,379 inhabitants per IRIS when counting all municipalities, and 2,860 when focusing on IRIS cities. Putting restricted-access databases aside, this is the finest spatial unit of observation available in French data.

The data used in this study is obtained by merging 77 year- and theme-specific IRIS-level files covering the period between 2006 and 2016. In addition to income, the main variable of interest that I extract from INSEE data, I select several neighbourhood characteristics listed in Table 6 in Appendix. These variables were first selected based on theoretical grounds, in the sense that they could be potential confounders in the pollution-income relationship that I study in Section 3. Moreover, epidemiological studies on social disparities in health and on environmental inequality generally rely on the European Deprivation Index (EDI) as a measure for the degree of neighbourhood precariousness (e.g., Morelli et al., 2016; Ouidir et al., 2017; Padilla et al., 2014). The variables

[2] Hereafter, I use the terminology of "IRIS cities" for those that are divided into IRIS and that were home to more than 10,000 inhabitants throughout the entire study period, while I call "non-IRIS municipalities" those that are either not divided into IRIS or those that are, but had less than 10,000 inhabitants at least once during the study period (for which I use weighted means of characteristics). However, the word "IRIS" or "census block" (the equivalent of IRIS in the American context) is used to define the observation level, regardless of the type of municipality it refers to.

included in this index are those that are the best predictors of both objective and subjective poverty measures within each European country. I thus add as controls those found by Pornet et al. (2012) to best mirror individual deprivation in France. In addition to the descriptive list provided in Table 6, summary statistics of all neighbourhood characteristics for year 2016 are available in Table 7 in Appendix. Note that the 2016 median of IRIS median incomes amounts to €20,252, which is approximately equal to the 2016 French median income of €20,500. However, due to the difference in observation level, the first decile of income is equal to €16,162 in 2016 in my data, as opposed to €11,040 based on individual-level data (Argouarc'h and Picard, 2018).

There are several issues arising with the use of a panel of yearly IRIS datasets:

1. Some IRIS boundaries (slightly) changed over time;
2. There was a large wave of mergers of towns during the study period;
3. Some cities were newly divided into IRIS during the period;
4. Income data at the IRIS level is provided only for municipalities with more than 10,000 inhabitants, and not for those with 5,000-10,000 inhabitants that are divided into IRIS.

These issues are dealt with step by step. Regarding the first issue, I construct a matching file of IRIS identifiers for IRIS that changed boundaries over time. This is arguably not problematic *per se*, since, according to INSEE, IRIS boundaries changed only slightly over the period. Still, part of this work has to be performed by hand, due to missing conversion tables.[3]

Second, there was a large wave of mergers of towns during the study period. A rather substantial fraction of these mergers occurred in 2015 and 2016, after the 2015-292 Law passed, which facilitated the creation of new merged cities if it would take place during this period. This is particularly problematic since data for year t is provided using the geographic breakdown of year $t+2$. In addition, other mergers occurred before 2015. In this case, it is not possible to only rely on a matching file. When towns are merged, it is the identifier of the chosen "head municipality" (*commune siège*) that becomes the identifier of the new merged town. As such, characteristics associated with the identifier before the merger are simply not comparable to characteristics associated with the same identifier after the merger. Hence, there are 2 available options: 1) using observations that were available at the initial pre-merger level for soon-to-be-merged municipalities, and considering the new merged municipality as a new entity, or 2) setting 2006 as a "fake" merger date, and thus

[3]The resulting dataset is available upon request.

computing the weighted average of the characteristics of the merged municipalities to assign it to the single identifier. The first option directly entails that these towns would be missing from all the parts of this study that exploit the longitudinal dimension of the data, and is thus discarded. I opt for the second option, although it presents two main disadvantages. First, it is not possible to compare, not even for a single year, the weighted average with the actual value for the new merged municipality, and second, this increases measurement error, especially in pollution, since the surface areas of merged towns are very large.[4] Nonetheless, this last issue is slight, since municipality merging mostly occurs in rural and rather sparsely populated areas. They do not carry very large weights in the following computations, since the latter are based on population weights.

Third, some (though few) municipalities were divided into IRIS during the period. This issue is tightly linked with the fourth one listed above in terms of consequences. This one implies that for the 915 municipalities of 5,000-10,000 inhabitants, while I observe socio-demographic characteristics at the IRIS level, I observe income at the municipality level. Hence, in both cases, I eventually observe some characteristics at the municipality level, and some at the IRIS level. In such a case, I keep all variables at the municipality level, and compute weighted means of IRIS-level variables. As a direct consequence, there is a need to assess the degree of validity of the weighted means. INSEE provides

Figure 1: Kernel density of p-value of the difference between actual value and weighted mean – 2015, 2016

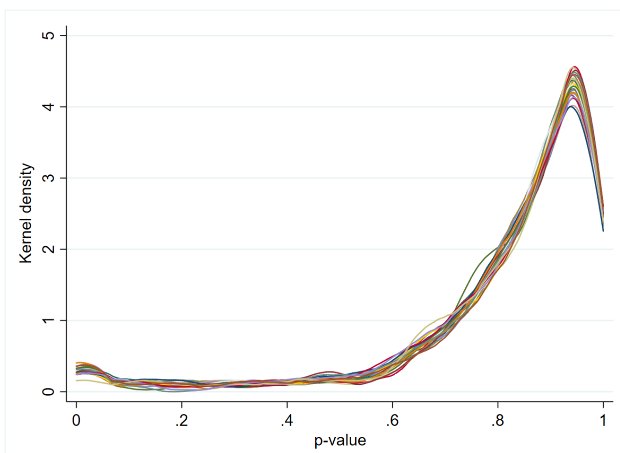

Variables tested: shares of each occupation, shares of French, immigrants and foreigners, share of unemployed, share of homeowners, tenants and subsidised housing tenants, share of each education level, share of single-parent families.

[4]For instance, 4 of the 10 largest municipalities in metropolitan France in terms of surface area are newly merged municipalities. All of them are located in Maine-et-Loire, where municipality merging was particularly frequent during the 2010-2016 period.

municipality-level information for the characteristics that were selected only for years 2015 and 2016, which is the reason why one cannot simply replace the IRIS-level observations by the municipality-level observations. Still, this allows me to assess how close the weighted means I computed are from the true value for these specific years. I present the results of this exercise in Figure 1, which plots the density of the *p*-value of the difference between the INSEE municipality-level value and the weighted mean I computed for each municipality in question. For any variable tested, the weighted means are not significantly different from the actual value for most observations.

2.2 Exposure to fine particulate matter

I use $PM_{2.5}$ concentration data from the Atmospheric Composition Analysis Group (ACAG) in Dalhousie University, Canada. The researchers exploited both remote-sensing sources and ground-level monitor data gathered by the European Environment Agency (EEA) in order to deduce the spatial distribution of fine particulate matter throughout the whole of Europe. They make use of GEOS-Chem, a recent chemical transport model that was jointly developed by the ACAG and Harvard University researchers.[5] I observe $PM_{2.5}$ concentration for years 2001 to 2016. The yearly files are made available to the public in raster format in grids of .01 $\times$ .01 degrees, i.e., approximately 1 km^2 at the equator. Note that the degree of precision of this data appears to be primarily adapted to large-scale studies like this one.

I infer pollution concentration at the IRIS level by computing the weighted average of the level of pollution associated with each raster grid that (at least partly) overlaps with each IRIS, using the Zonal Statistics plugin of QGIS. This implies that exposure is defined as the mean of the values of the raster grids at least part of which are inside the boundaries of each IRIS, weighted by the proportion of the area of the grid present within the area of the IRIS. I also extract the number of grid values used to compute the mean pollutant concentration for each IRIS. The latter measure is important because it is reflective of the level of measurement error of the pollution exposure variable. Indeed, as exposed in Section 2.1, not all French municipalities are divided into IRIS, and the average surface area of a French *commune* is 14.88 km^2 (14.83 in my data). This is small compared to most European countries, but remains an issue in this context: on average, 17.71 grid points were used to compute average concentration for non-IRIS municipalities, while only 2.15

[5] Further information relative to the GEOS-Chem transport model is available on the dedicated website: `http://acmg.seas.harvard.edu/geos/`.

were used in the case of IRIS municipalities. Therefore, the independent variable is subject to measurement error, which reduces the power of the statistical tests conducted in following sections, or, in other words, increases the risk of Type-II error.

I define average exposure as the population-weighted mean $PM_{2.5}$ concentration within a given IRIS. IRIS-level data, including population data, are not available for years prior to 2006. These data constraints imply that this analysis focuses on the 2006-2016 time period. According to the data, average exposure to fine particulate matter is equal to 9.54 $\mu g/m^3$ in 2016, which complies with both the European limit value of 25 $\mu g/m^3$ and the World Health Organisation (WHO) guideline of 10 $\mu g/m^3$, on average. However, the distribution of $PM_{2.5}$ concentration ranges from 4 to 15.1 $\mu g/m^3$ at the IRIS level (see the first row of Appendix Table 7).

The Atmospheric Composition Analysis Group also provides estimates of ground-level $PM_{2.5}$ concentration at a larger scale of .1 × .1 degrees, that cover almost the whole world surface (Van Donkelaar et al., 2016). These were used in the context of the Global Burden of Disease Study of 2015 (Cohen et al., 2017), and are channelled on OECD's website (OECD, 2018), making its higher-definition counterpart likely reliable. Nonetheless, I proceed to evaluate the consistency of the ACAG data using the few national-scale sources of information on fine particulate matter concentration. I compare the map of average of 2007-2008 exposure provided in a study by researchers of Santé publique France[6] with the one that I obtain using ACAG data in Figure 17 in Appendix. The two maps are very similar, which confirms the validity of the data used in this study; the minor differences may be attributed to the difference in air transport models (Gazel'Air for Medina et al., while the ACAG uses GEOS-Chem) and to the threshold effect of the caption, which is not continuous, but uses a 5-$\mu g/m^3$ increment. I also look at the consistency of ACAG data with the evolution of average $PM_{2.5}$ concentration provided in the 2017 Report on Air Quality.[7] Figure 18 in Appendix thus displays the 2009 base-100 index of $PM_{2.5}$ concentration in mainland France. While the index values at the end of the period are similar, there is up to a 7-point discrepancy between the two estimates in the 2011-2013 period. This might again be explained by the fact that Le Moullec uses another air transport model, or by a significant initial divergence between the two estimates, but the author only provides indices.

[6]See Figure 6 in Medina et al. (2016).
[7]See Figure 3 in Le Moullec (2018).

3 Descriptive evidence

3.1 Patterns of inequality

Starting with macro-scale patterns of inequality in exposure to fine particulate matter, Figure 2 depicts $PM_{2.5}$ concentration by decile alongside a map of median income by decile and by *département*. The latter administrative division roughly corresponds to the county-level of the United States or the United Kingdom. Comparing these two maps allows to confirm that, in France, even focusing on the regional scale and avoiding delving into urban area-specific heterogeneity, the relationship between $PM_{2.5}$ exposure and income is not at all monotonic. Île-de-France (excluding Seine-Saint-Denis) and the former Rhône-Alpes region combine some of the highest levels of both fine particulate matter exposure and income, while, unsurprisingly, the association is reversed for the rural areas of central France. On the other hand, the northern *départements* of Nord, Pas-de-Calais, Aisne and Ardennes belong to the highest decile of exposure and the two lowest deciles of income. On the other side of both spectra, the Atlantic South-West and Brittany compound relatively high income and low exposure. These differences cannot simply be explained by the presence of large cities, which couple comparatively higher pollution and higher income, as the situation in the north of France strongly contradicts this argument.

Figure 2: Exposure to $PM_{2.5}$ and median income by decile – 2016

(a) $PM_{2.5}$ (b) Median income

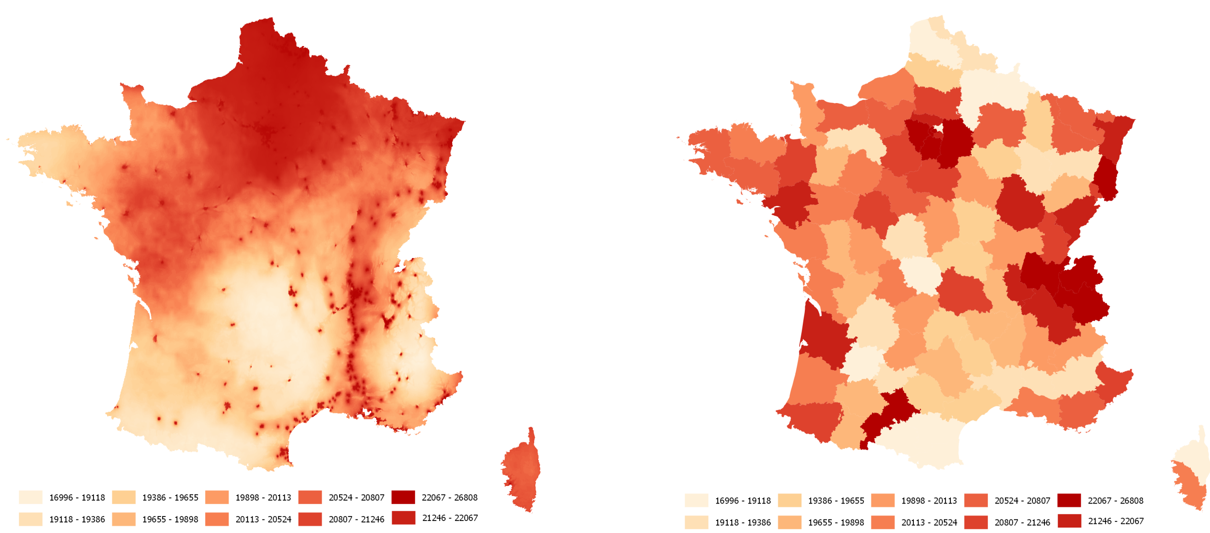

Sources: Atmospheric Composition Analysis Group (left) and INSEE (right).

I now turn to neighbourhood-level information, and split the data into two groups: the top 50% of the distribution (i.e., those that could be roughly defined as middle- and upper-class IRIS), and the bottom 50% of the initial income distribution. Recall that the median of census block-level income is roughly equal to the median of individual-level income (Argouarc'h and Picard, 2018). Average exposure is roughly equal between the two groups in 2016: the bottom 50% have a mean exposure of 9.49 $\mu g/m^3$, while that of the top 50% is equal to 9.58 $\mu g/m^3$. Throughout the study period, the gap between the exposure of the top 50% and the bottom 50% of neighbourhood income was equal to .3 $\mu g/m^3$ on average, only reaching .5 $\mu g/m^3$ in 2008-2009. This is further discussed in Section 3.4, which is dedicated to the evolution of (inequality in) $PM_{2.5}$ exposure.

This small difference in average exposure between the two groups masks specific patterns, as shown in Figure 3. The latter displays the distribution of exposure to $PM_{2.5}$ *within* the two halves of the national distribution of IRIS (median) income. Given that the number of observations is identical between the two groups, I can directly interpret gaps in the probability density functions as an occurrence of over- or under-representation. It seems that one can split these two distributions into three parts. At "very low" levels of exposure, between 4 and 7 $\mu g/m^3$, the bottom 50% of income is over-represented, in all likelihood due to the fact that rural areas usually combine low income and low pollution levels. At middle-range levels (7-10 $\mu g/m^3$), the top 50% of income is over-represented. Finally, focusing on levels of exposure above the WHO guideline of 10 $\mu g/m^3$, the top and bottom 50% are quite similarly represented, with a slightly higher proportion of top-50% neighbourhoods. This pattern is likely attributable to the characteristics of French urban areas:

Figure 3: Distribution of exposure to $PM_{2.5}$, top and bottom 50% of income – 2016

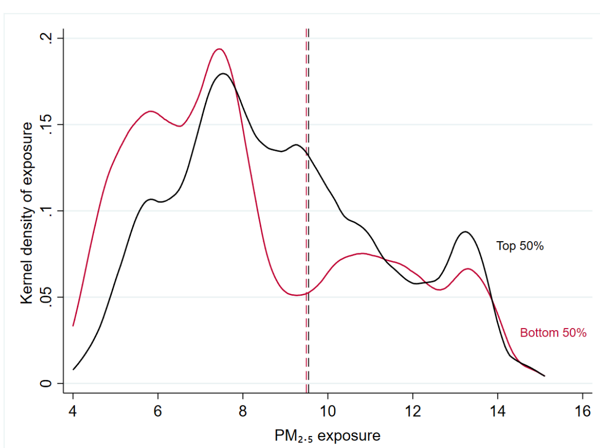

Note: The dashed vertical lines represent the mean value of exposure for each group. The median level of income is equal to €20,252.

albeit different levels of residential segregation, statistically more polluted city centres and inner suburbs usually compound both high- and very low-income neighbourhoods, and peri-urban areas are increasingly well-off (Aerts et al., 2015; Floch, 2014, 2017).

Figure 4 relies on a finer group definition: Figure 4a compares the distribution of exposure of the top 10% to that of neighbourhoods between the 8$^{\text{th}}$ and the 9$^{\text{th}}$ decile of income, while Figure 4b compares the exposure of the bottom 10% to that of neighbourhoods between the 1$^{\text{st}}$ and the 2$^{\text{nd}}$ decile of income. It appears that the likelihood of experiencing very high levels of exposure is indeed considerably higher for the top 10% IRIS than for any of the three other groups. On the other hand, a large fraction of neighbourhoods whose income lies between the 8$^{\text{th}}$ and the 9$^{\text{th}}$ decile are exposed to $PM_{2.5}$ levels that are below the national average of 9.54 µg/m^3. The patterns of distribution of exposure are very different in Figure 4b, which suggests that the neighbourhoods located below the 1$^{\text{st}}$ decile of IRIS-level income (i.e., whose median income is below €16,162) are substantially more likely to be exposed to high $PM_{2.5}$ levels: 60% of these low-income neighbourhoods are above the WHO standard, while, conversely, 80% of the neighbourhoods of the next 10% comply with this standard (see the CDFs in Appendix Figure 19). This is consistent with the fact that, in France in 2012, 65% of individuals living below the poverty line lived in large urban centres (INSEE's *grands pôles urbains*), and 20% in the Paris urban area (Aerts et al., 2015).

Figure 4: Distribution of exposure to $PM_{2.5}$ – 2016

(a) 8-9$^{\text{th}}$ decile and top 10% (b) Bottom 10% and 1-2$^{\text{nd}}$ decile

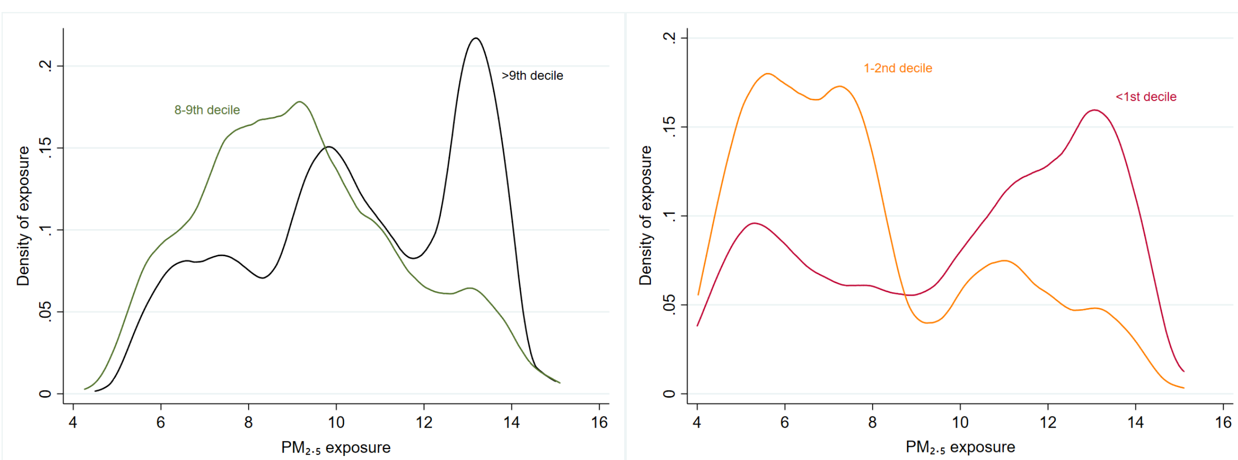

Note: The 8$^{\text{th}}$ decile of income is equal to €23,627, and the 9$^{\text{th}}$ decile is equal to €26,286 (left). The 1$^{\text{st}}$ decile of income is equal to €16,162, and the 2$^{\text{nd}}$ decile of income is equal to €17,738 (right).

Figure 5a displays the unconditional average exposure to $PM_{2.5}$ as a function of a neighbourhood's decile of income, for all even-numbered years of the study period, and rather corroborates previous evidence. As a consequence of the aforementioned phenomena relating income and location in France, there appears to be a U-shaped relationship between unconditional average exposure and income. However, the average exposure of the top 10% of IRIS income is lower than that of those located at the 9th decile of income, and similar to that of the bottom 10% of IRIS-level income. Conditioning $PM_{2.5}$ exposure on urban area yields a completely different story: *within* urban areas, neighbourhoods with very low income are substantially more polluted than their wealthier counterparts. Average exposure roughly decreases with income up to the 7th decile, and then slightly rises with income at the right end of the distribution. Notice also that conditionally on urban area, the exposure of neighbourhoods whose income lies between the 4th and the 9th decile is lower than average, while the top 10% have average exposure.

Figure 5: Average exposure to $PM_{2.5}$ based on (national-level) income decile

(a) Unconditional

(b) Conditional on urban area

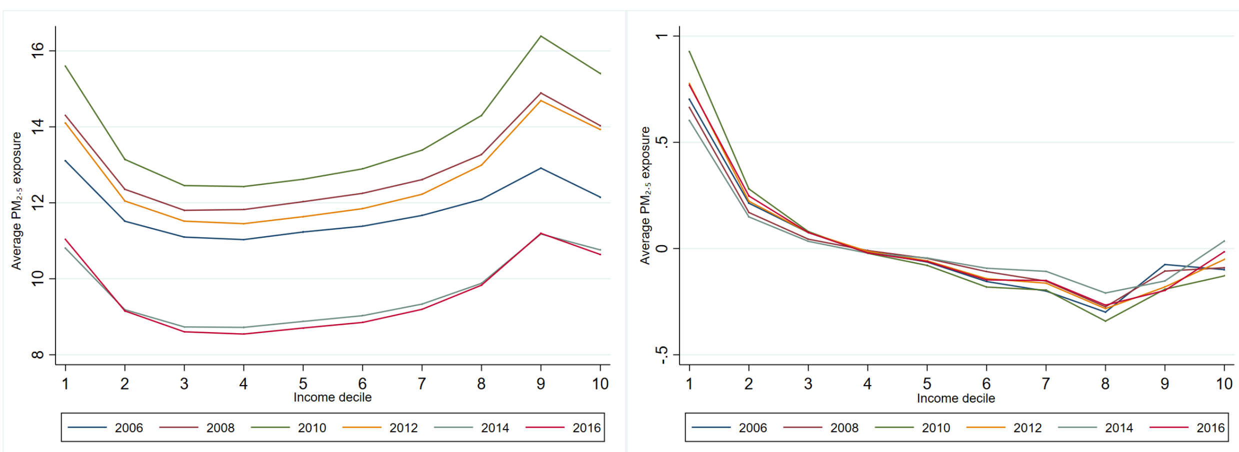

3.2 Fixed-effect models

I provide a more formal analysis of the correlation between fine particulate matter exposure and neighbourhood income by running IRIS-level fixed-effect models. By doing so, I exploit the variation in income and $PM_{2.5}$ concentration *within* census blocks, and thus difference out the potentially confounding effect of unobserved census block-level characteristics which could influence the location decision of individuals. Indeed, while I control for a series of neighbourhood socio-demographic characteristics (listed in Appendix Table 6), it remains that I do not observe the quality of the amenities present within or in the vicinity of neighbourhoods. These are particularly important in this study due to the fact that households likely substitute away air quality for other amenities, such as cultural facilities, schools or proximity to large business districts. Although not discussed, this omitted variable bias is present in any of the abovementioned cross-section analyses (e.g., Ouidir et al., 2017; Zwickl et al., 2014) whose model identification relies on variations across study areas. The corresponding equation is the following:

$$\ln(\mathrm{PM}_{it}) - \overline{\ln(\mathrm{PM}_i)} = \alpha + \beta_{INC}(\ln(\mathrm{INC}_{it}) - \overline{\ln(\mathrm{INC}_i)}) + \beta_X(X_{it} - \overline{X}_i) + \lambda_t + (\varepsilon_{it} - \overline{\varepsilon}_i) \quad (1)$$

$PM_{2.5}$ exposure is right-skewed, so I take its natural logarithm in all regressions. $\ln(\mathrm{PM}_{it})$ corresponds to the (log) average exposure in IRIS i, during year t. $\ln(\mathrm{INC}_{it})$ is the median income of IRIS i in year t. X_{it} includes the covariates taken from the list displayed in Table 6 in Appendix. I include the share of inhabitants that undertook higher education, since more educated individuals may attach a higher value to air quality than less educated individuals, and thus hold a different view of the trade-off between amenities and air pollution, which influences their residential choice. For similar reasons, I select the shares of population by occupation, divided in 8 categories (see Appendix Table 6). Moreover, homeowners and subsidised housing tenants, i.e., those who live in HLM (*Habitation à Loyer Modéré*), may have more stringent residential mobility constraints than tenants of privately owned dwellings, which would prevent them from leaving polluted areas. I also control for the share of dwellings that are not equipped with electric heating, since domestic wood burning is one of the main sources of $PM_{2.5}$ in France (Citepa, 2018), and similarly for the share of households that own a car. The fractions of immigrants, of unemployed individuals, and single-parent households are also used as measures of deprivation (Pornet et al., 2012). Finally, I cluster standard errors at the employment-zone level in order to account for autocorrelation within employment zones, and weight regressions by population.

I take advantage of the panel structure of my data to deal with the omitted variable bias related to time-invariant unobserved IRIS-level characteristics, and I add year fixed effects in order to tackle the impact of year-specific shocks that would impact both income and $PM_{2.5}$ concentration, such as a shock in economic activity. However, it is likely that there are remaining biases, meaning that the models do not identify a causal impact. First, the models rely on the assumption that there exists no time-varying unobserved heterogeneity across IRIS throughout the 2006-2016 time period, since it was computationally inaccessible to introduce more than 400,000 time-varying fixed effects. Second, unlike an individual-level fixed-effect model, the models that I estimate do not fully tackle the self-selection issue that is inherent to the study of environmental inequality. Indeed, they do not resolve the potential reverse causality of the pollution-income relationship,[8] which opposes the neighbourhood sorting hypothesis to the firm siting hypothesis. As a consequence, the test that I provide here boils down to looking at whether any of these two hypotheses, or a combination of the two, may apply to France.

Table 1 shows the results associated with equation (1). The crude model with only IRIS fixed effects shows a strong negative correlation between IRIS median income and $PM_{2.5}$ exposure, and the following ones as well, though it is less pronounced. In the preferred specification (Column (3)), in which I control for year fixed effects as well as neighbourhood characteristics, the estimates suggest that over the 2006-2016 period, a 1% positive deviation from a neighbourhood's mean income is associated with a decrease of .18% of its mean $PM_{2.5}$ concentration, *ceteris paribus*. While these models do not identify a causal impact, but a correlation net of the impact of certain variables, they support the idea that, at the national scale, a higher neighbourhood income is associated with a lower exposure to $PM_{2.5}$. Therefore, these results do not contradict the Tiebout (1956) sorting mechanism evoked in Banzhaf et al. (2019). As a clean environment can be considered as a luxury good, higher air quality may prop up rents and housing prices.[9] As a direct consequence, there may be a sorting phenomenon of households across income groups, even if households do not necessarily choose to migrate due to higher air pollution (Banzhaf and McCormick, 2012). These results are also theoretically consistent with a theory of disproportionate siting by firms. However, this mechanism likely plays a less significant role in the case of fine particulate matter as compared to other pollutants, since a bit less of a quarter of emissions are emitted by the manufacturing sector, while two thirds are from residential sources and transportation (Citepa, 2018). Finally, this result may also be attributed to a combination of both processes of sorting and siting, through Coasian bargaining (Banzhaf et al., 2019). In any case, I only provide the first piece of evidence of

[8] Lagged independent variables are not particularly helpful in this context, given the degree of inertia of the variable of interest between 2 subsequent years of observation at the IRIS level.

[9] To my knowledge, only two studies investigated this in France. Lavaine (2019) found evidence of a significant impact of air quality on housing prices in the highly polluted Dunkirk metropolitan area, while Le Boennec and Salladarré (2017) did only for specific types of households in the less polluted Nantes region. It may be hypothesised that the intensity of sorting dynamics could vary depending on the overall pollution level of an urban area.

Table 1: Fixed effect models – Partial results

Variable	(1)	(2)	(3)
(Log) income	-1.274*** (.137)	-.212*** (.037)	-.182*** (.025)
% Immigrants			.143*** (.025)
% Higher education			.141*** (.039)
% White-collar			.258** (.105)
% Inactive excl. retired			.113** (.054)
% Unemployed			-.119 (.105)
% Social housing			.051 (.054)
Intercept	15.062*** (1.357)	4.542*** (.368)	4.061*** (.344)
Year fixed effects		X	X
R^2 within	0.11	0.78	0.79
R^2 between	0.02	0.01	0.21
R^2 overall	0.01	0.18	0.37
# Observations	453,386	453,386	453,306
# Groups	42,832	42,832	42,825

Standard errors clustered at the employment-zone level in parentheses. *: 10% level, **: 5% level, ***: 1% level.

this national-level correlation: further research based on individual-level data would allow to better understand the mechanisms behind this inequality.

Finally, paying specific attention to other neighbourhood characteristics, the results also suggest that there is a positive correlation between the share of immigrants and exposure to fine particulate matter, even after controlling for income. This cannot be attributed to the fact that a great fraction of immigrants are gathered in specific neighbourhoods of large metropolitan areas, since I control for IRIS-level unobserved heterogeneity. Therefore, there might also be an ethnic gap in exposure to fine particulate matter in France, as observed in the US (Currie et al., 2020; Kravitz-Wirtz et al., 2016).

The positive link between $PM_{2.5}$ concentration and the shares of highly educated individuals, of white-collar workers, and inactive individuals – the latter being mostly students, since the category excludes retirees – is consistent with the fact that these populations tend to concentrate in urban centres. On the other hand, there is no significant relationship between unemployment share and $PM_{2.5}$ level, likely due to the fact that unemployment is not only high in polluted former industrial areas, but also in cleaner rural ones. The same observation applies to the share of subsidised housing, for which I find no significant result. This may also stem from the fact that although there are significant variations in the share of social housing across IRIS, there may not be large enough variations within IRIS for any significant difference to be identified. The remainder of the estimates are provided in Appendix Table 8, and the coefficients all appear consistent with theory.

3.3 Robustness to spatial autocorrelation

Tobler's (1970) first law of geography states that "everything is related to everything else, but near things are more related than distant things". Indeed, spatial autocorrelation is likely to be of particular concern when one examines the relationship between income and access to clean air. Air pollution levels are spatially correlated by construction, and income does not verify complete spatial randomness (CSR) either: in France, residential income segregation is less pronounced than in the US (Quillian and Lagrange, 2016), but tends to increase (Beaubrun-Diant and Maury, 2020). In Section 3.2, standard errors are clustered at the employment-zone level, which means that correlation within employment zones is allowed for. This implies that spatial correlation is modelled discretely, since it is designed to follow employment zone boundaries. Pollution, however, is distributed continuously across space. As a consequence, pollution levels on either side of an employment zone boundary are thus as likely to be correlated as pollution levels of two IRIS located within the same zone. This argument also holds for other neighbourhood characteristics. Hence, the residuals of equation (1) are most likely not independently distributed. Moreover, failing to take spatial autocorrelation into account amounts to have an artificially lower variance in observations, and thus artificially lower standard errors, thus inflating the risk of Type-I error.

To my knowledge, despite the importance of the issue in this setting, environmental inequality studies performed by economists generally lack concern for spatial autocorrelation. In a related study, Lavaine (2015) does use Driscoll-Kraay standard errors, which are robust to cross-sectional correlation, in some specifications. Nonetheless, a number of studies do not mention it (e.g., Currie et al., 2020; Muller et al., 2018; Voorheis, 2017; Zwickl et al., 2014). Concern for spatial autocorrelation is however more common within the Public Health literature on the topic: some highlight the consequences of failing to account for it (e.g., Havard et al., 2009), and among those mentioned in Section 1.2, a large part uses specifically designed spatial models.

Some adopt spatial lag models, a specific version of Spatial Autoregressive (SAR) models (Havard et al., 2009; Verbeek, 2019). Spatial lag models treat spatial dependence between observations as substance, as opposed to a disturbance. They assume that the value taken by the dependent variable in each zone both affects and is affected by the values taken by the dependent variable in the neighbouring zones, which is exactly what one may expect when the outcome variable is pollution exposure. With W a weighting matrix, a basic spatial lag model may be thus written $Y = X\beta + \rho WY + \varepsilon$. The choice of the form of the spatial weighting matrix W is however arbitrary, as it amounts to making parametric assumptions about the behaviour of spatial autocorrelation. Finally, misspecifying the spatial weight matrix can introduce large biases in the final estimates (Anselin, 2002; Lam and Souza, 2014).

An appealing solution is thus to opt for a non-parametric approach to account for spatial autocorrelation. Specifically, I use a Generalised Additive Model (GAM). These are Generalised Linear Models (GLM) to which one adds a smoothing function of at least some covariates. As such, while they were not designed as spatial models to begin with, they allow to add a non-parametric function of geographic coordinates to account for neighbourhood location as a possible predictor of $PM_{2.5}$. Unlike spatial lag models, GAM thus do not impose parametric assumptions on the form of spatial autocorrelation. Hence, they were used to model spatial autocorrelation in studies of interregional knowledge spillovers (Guastella and Van Oort, 2015), interregional risk sharing (Basile and Girardi, 2010) or hedonic house pricing (Helbich et al., 2014; Von Graevenitz and Panduro, 2015). They have also been used in few environmental justice studies (Brochu et al., 2011; Padilla et al., 2014).

Thus, I add the latitude y_i and longitude x_i of the centroid of IRIS i as a smoothed term $s(x_i, y_i)$ to equation (1).[10] $s(\cdot)$ is a thin plate regression spline, which does not require to specify knots. The most widespread approach used to estimate GAM is the backfitting algorithm of Hastie and Tibshirani (1990) but, in practice, it implies that one must select the degree of smoothness of the term $s(x_i, y_i)$. Selecting the degree of smoothness amounts to selecting the span size, i.e., the

[10] Modelling spatial autocorrelation this way is thus akin to tackling the omitted variable bias arising from the fact that although location likely has an impact on both income and pollution level, equation (1) did not control for it.

Table 2: Fixed-effect generalised additive models – Partial results

Variable	Dependent variable: log(PM)		
(Log) income	-1.274*** (.005)	-0.157*** (0.003)	-.125*** (.003)
% Immigrants			.230*** (.009)
% College-educated			.219*** (.006)
% White-collar			.231*** (.009)
% Inactive excl. retired			.081*** (.006)
% Unemployed			-.121*** (.009)
% Social housing			.063*** (.004)
Intercept	.001 (.002)	-0.001*** (.000)	.030*** (.000)
Year fixed effects		X	X
Other neigh. charac.			X
p-value $s(x_i, y_i)$	0.128	0.000	0.000
Adjusted R^2	0.11	0.78	0.78
# Observations	453,386	453,386	453,306
# Groups	42,832	42,832	42,825

Standard errors clustered at the employment-zone level in parentheses. *: 10% level, **: 5% level, ***: 1% level.

number of IRIS in the vicinity of IRIS i whose outcomes are likely correlated with these of IRIS i, and thus boils down to a bias-efficiency trade-off. In order not to resort to arbitrary choices, I rely on more recent advances in research on GAM, which led to the development of a new estimation technique that singles out the degree of smoothness in an automatic and integrated fashion. Hence, the optimal smoothing parameter is obtained *via* generalised cross-validation (Wood, 2017).

The results that mirror those of Table 1 are displayed in Table 2, and the full set of estimates is provided in Appendix Table 9. The approximate p-value of the smoothing term $s(x_i, y_i)$ is highly significant, which confirms that latitude and longitude are indeed predictors of IRIS pollution level.

Another striking element is the fact that, for any specification, the intercept is substantially smaller than before, and the adjusted R^2 substantially higher than before, since neighbourhood location likely explains a great fraction of $PM_{2.5}$ variation. Although omitting to allow for spatial autocorrelation theoretically shrinks standard errors, significance levels are much higher than in previous models. It appears that as I net spatial autocorrelation out, the pollution-income relationship weakens, although the estimates are not significantly different from before: a 1% increase from the IRIS mean median income is associated with a .125% reduction in mean $PM_{2.5}$ exposure. This supports the idea that the spatial dimension does not attenuate nor invalidate the relationship between income and environmental hazard in the case of $PM_{2.5}$ exposure.

Notwithstanding, the positive relationship between $PM_{2.5}$ concentration and the share of immigrants is significantly stronger than under the previous specification at the 5% level, with an estimate of .230 as opposed to .143. Equation (1) considered IRIS as independent of each other, but diagnosed a positive correlation between $PM_{2.5}$ concentration and share of immigrants within census blocks. There is a positive spatial autocorrelation of both pollution and the share of immigrants across these blocks, which attenuates with distance. Consequently, controlling for the longitude and latitude of a census block that has a high (resp., low) pollution level accounts for the fact that its neighbours also have a high (resp., low) pollution level, and thus a likely high (resp., low) share of immigrants. In other words, allowing for spatial autocorrelation places IRIS back into their relative location, thus creating "clusters" that combine high values of $PM_{2.5}$ and share of immigrants together, medium values together, and low values together, which implies that the correlation is indeed higher than what simple FE models estimated. On the other hand, as aforementioned, the estimate associated with income is lower in absolute terms than previously evaluated, with an estimated coefficient of -.125 against -.182. These two findings can thus seem conflicting at first sight, but can be reconciled, since they imply that, for a given degree of spatial smoothing, income is distributed more uniformly across space than the share of immigrants. Indeed, in France, ethnic segregation measures provided by, e.g., Gobillon and Selod (2007), Préteceille (2011) or Safi (2009), are higher than income segregation measures (Floch, 2017; Quillian and Lagrange, 2016).

The estimates associated with the shares of college-educated individuals, white-collar workers and inactive inhabitants, and those only presented in Appendix Table 9, are qualitatively and quantitatively similar to the previous ones, making them robust to accounting for spatial autocorrelation. Finally, the estimates for the unemployment share and the social housing share are now statistically significant, which may also be attributed to segregation.

3.4 Evolution in exposure and longitudinal inequality

3.4.1 Graphical evidence

Figure 6 displays the evolution of average exposure to fine particulate matter in metropolitan France throughout the study period. According to this figure, average exposure increased from 11.85 µg/m^3 in 2006 to around 14 µg/m^3 for 2009-2011, before falling to 9.54 µg/m^3 in 2016. Hence, for years 2014 and 2016, according to matched ACAG and INSEE data, exposure seems to be line with the World Health Organisation (WHO) guideline of 10 µg/m^3, *on average*. Average exposure fell by 4.63 µg/m^3, or 32.7%, between the peak year 2009 and 2016, a decrease similar to the United States' in the 2000-2014 period (Currie et al., 2020; Voorheis, 2017).

Figure 6: Evolution of average exposure to $PM_{2.5}$ – 2006-2016

However, this national average masks some considerable spatial disparities in the evolution of fine particulate matter concentration. Figure 7 shows two years of this concentration for metropolitan France, by decile of the total distribution across these two years. In 2006, the least exposed areas were Massif Central, the Pyrénées, the southern part of the Bay of Biscay and the tip of Brittany, and remain so. More specifically, $PM_{2.5}$ concentration was already very low in Massif Central, and did not significantly decrease, while it did in the last three regions. On the other hand, the northern part of the country, the Paris region and the Rhône Valley, which were located in the top 30% of the distribution in 2006, stayed at this position in 2016. This implies that while pollution was already high in these areas, and particularly much higher than the WHO guideline, they remained at the top of the cross-year pollution distribution, while the rest of the country moved down. Taken to-

Figure 7: Exposure to $PM_{2.5}$ by decile

(a) 2006 (b) 2016

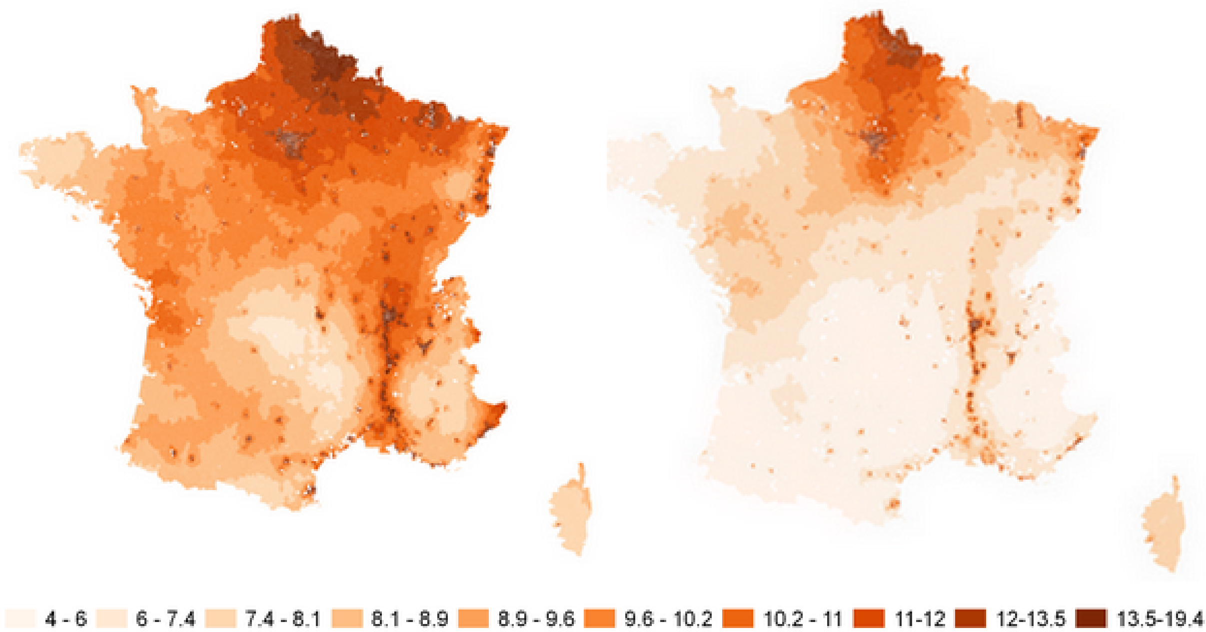

gether with the map of Figure 2b, this does not allow to draw clear-cut conclusions on longitudinal patterns of inequality: while the low-income northern *départements* remained polluted, high-income ones like Rhône or Seine-et-Marne also did.

As a first piece of evidence on longitudinal environmental inequality, I switch back to the income group definitions of Section 3.1. Figure 8a distinguishes between the top and the bottom 50% of the income distribution, and shows that the top 50% of income appears to be consistently more exposed to $PM_{2.5}$ than the bottom 50%, experiencing a larger increase in exposure during the 2006-2011 period, and a rather similar overall decrease up to 2016. The fact that urban areas concentrate both a high level of $PM_{2.5}$ concentration, a high proportion of individuals, and a relatively high level of income surely can explain a large part of the gap that we observe in this figure. However, patterns are rather different when looking at the 4 group definitions also used in Section 3.1, namely the bottom and top 10% of income, and areas located between the 1^{st} and the 2^{nd} decile, and the 8^{th} and the 9^{th} decile. The gap between the latter remains quite constant along the years. However, while they had roughly equal average levels of exposure during the 2008-2013 period, a gap in exposure

Figure 8: Evolution of average exposure to $PM_{2.5}$ in different income groups – 2006-2016

(a) Top 50% *vs* Bottom 50% (b) Two highest and lowest deciles

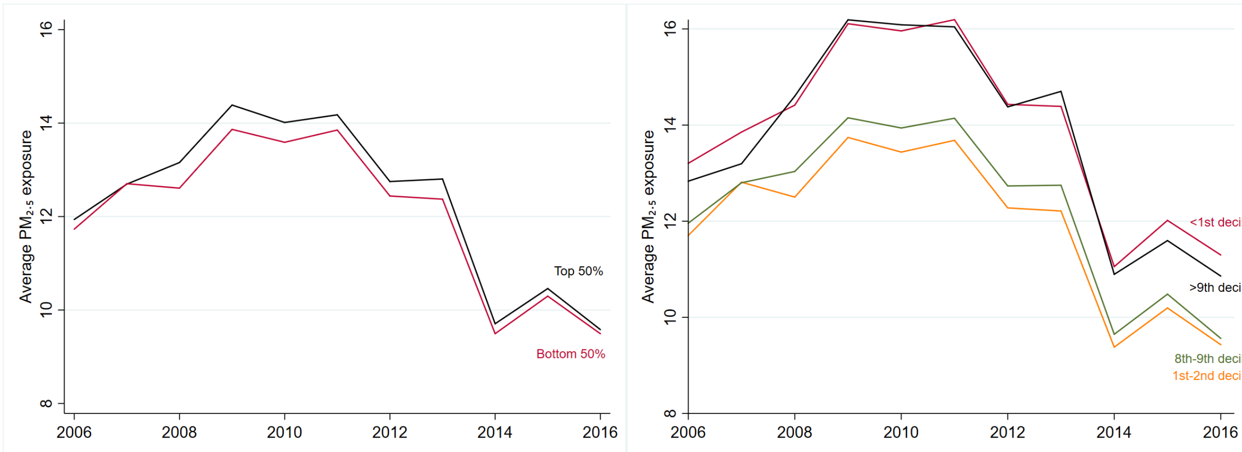

may be arising between the neighbourhoods at the top and those at the bottom 10% of income, at the favour of the top 10%. This will be a trend to pay attention to in upcoming years.

Moreover, these results are obtained on the basis of varying ranks in the income distribution. A small counterfactual exercise can highlight the potential role of mobility. Indeed, one can compute the average exposure of neighbourhoods in 2016 using their 2006 rank, and compare it to the actual average. With this definition, the average of exposure of the bottom 10% of 2016 in 2016 is 11.30 $\mu g/m^3$, but that of the bottom 10% of 2006 in 2016 would have been 11.08 $\mu g/m^3$, holding the ranks fixed. On the other end of the spectrum, the 2016 top 10%'s average exposure is equal to 10.86 $\mu g/m^3$, while it should have been 11.01 $\mu g/m^3$ holding the ranks fixed. As such, the bottom 10% of 2016 is more exposed than the 2006 bottom 10% would have been, and the reverse holds for the top 10%. This implies that neighbourhoods that are "new" to the bottom 10% are more polluted than those that "left" the bottom 10%, and that neighbourhoods that are "new" to the top 10% are less polluted than those that "left" the top 10%. Such a fact is consistent with relative mobility patterns that would occur due to the neighbourhood sorting mechanisms, where higher-income (resp., lower-income) individuals would self-select into cleaner (resp., more polluted) neighbourhoods, be it due to pollution-related out-migration or to market forces (Banzhaf and McCormick, 2012; Banzhaf et al., 2019). Individual data would be needed to formally test these hypotheses.

3.4.2 Pollution-reduction profiles

In order to study longitudinal inequality in exposure to fine particulate matter, I proceed to compute pollution-reduction profiles, following Voorheis (2017). Voorheis argues that, although there is a sizeable body of literature on the cross-sectional measurement of environmental inequality in the US, much less is known about longitudinal environmental inequality. As such, he adapts a method used in the literature on intra-generational mobility and first proposed by Jenkins and Van Kerm (2006), called income-reduction profiles, to pollution exposure. The resulting pollution-reduction profiles (PRP) allow to capture both vertical equity concerns (i.e., how $PM_{2.5}$ exposure varies across initial ranks of exposure) and horizontal equity concerns (i.e., how $PM_{2.5}$ exposure varies across initial ranks of income). Both types of PRP are also computed both in absolute terms, using the difference between exposure in year $t + x$ and exposure in year t, and in relative terms, using the difference between the logarithm of exposure in year $t + x$ and that of year t. As shown in Figure 6, average exposure increased between 2006 and 2011, and decreased ever since. This framework thus allows to visualise the distributional impacts of the 2006-2011 air quality deterioration and the 2011-2016 (and overall) air quality improvement. All PRP are obtained by fitting a thin plate regression spline.

Figure 9 (resp., Figure 10) shows the vertical (resp., horizontal) equity profiles, with absolute changes in exposure on the left-hand side and relative changes in exposure on the right-hand side. I begin by looking at the change in exposure as a function of a census block's initial rank in the distribution of exposure. Between 2006 and 2016, regardless of their initial level of exposure to $PM_{2.5}$, on average, IRIS benefitted from overall quite comparable decreases in absolute terms, between -1.7 $\mu g/m^3$ and -2.8 $\mu g/m^3$. However, this implies that, in relative terms, census blocks that were initially less exposed to $PM_{2.5}$ benefitted from larger improvements than those who were initially more exposed. Distinguishing between the two phases of the study period, the graph suggests that initially disadvantaged census blocks experienced a higher increase in exposure between 2006 and 2011, and a smaller relative decrease in exposure between 2011 and 2016, compared to their initially less exposed counterparts. In particular, neighbourhoods at the 10^{th} percentile of initial exposure received a 40% decrease in exposure, while those at the 90^{th} percentile received a 14% decrease. In short, the vertical equity measures suggest that both the rise and the fall in $PM_{2.5}$ concentration have been regressive, with larger benefits accruing to initially less exposed areas.

Figure 9: Pollution-reduction profiles – Vertical equity

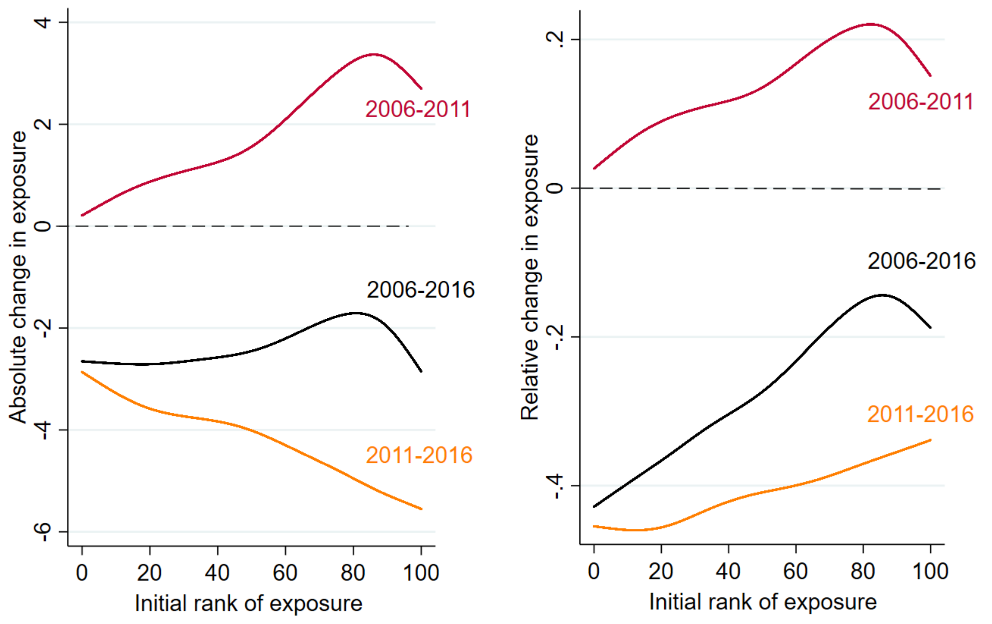

Note: The initial rank of exposure is the rank in the distribution of IRIS exposure to $PM_{2.5}$ in 2006.

This is consistent with what Figure 6 shows at a larger geographical scale: while $PM_{2.5}$ concentration decreased throughout metropolitan France, regions that were initially the most polluted, such as Hauts-de-France and Île-de-France, did not experience larger relative decreases than others. This may partly be attributed to relatively smaller decreases in emissions due to less effective policies, but another potential explanation for this may be that fine particulate matter concentration did not decrease more in initially more polluted areas because a potentially significant part of these particles is imported. Indeed, $PM_{2.5}$ can travel long distances, which implies that a potentially non-negligible part of neighbourhoods' observed concentration is not due to domestic emissions. For instance, in Île-de-France in 2010, it was estimated that 39% to 68% of their observed quantity was produced outside the region (Airparif, 2011).

Turning to the horizontal equity measures of Figure 10, the first finding is that, similarly to vertical equity measures, absolute changes in exposure over the whole study period were rather uniform across the initial income distribution, with a 2 $\mu g/m^3$ decline up to the 8^{th} decile, and an average decrease of 1.8 $\mu g/m^3$ for the top 10% of income. However, in relative terms, the bottom decile of initial income, i.e. neighbourhoods with a median income of €14,300 in 2006,[11] received a 20%

[11] All values are given in 2016 constant euros.

decrease in exposure, while those at the 4th decile (€17,360), received the largest relative decrease of 28%. Neighbourhoods of the top 10% of income (whose 2006 income is above €24,000) benefitted from the smallest relative improvement, with a 17% average decrease in $PM_{2.5}$ concentration. To summarise, in relative terms, it appears that the pollution-reduction profile is U-shaped, with air quality improvements accruing to a greater extent to neighbourhoods located in the middle 60%, while those at the top and the bottom of the initial income distribution experienced significantly smaller relative decreases. Splitting the study period into the two phases studied above, it appears that the bottom 20% and the top 20% of income not only experienced the largest (relative or absolute) increase in exposure between 2006 and 2011, but also slightly smaller relative decreases after 2011. Taken together with the evidence in Figure 9 and Section 3.1, these patterns are consistent with an overall improvement of fine particulate matter pollution throughout the country, which, however, favoured municipalities and neighbourhoods that combine intermediate income and comparatively low pollution levels. This implies that although overall $PM_{2.5}$ exposure underwent a substantial drop during the study period, inequality in exposure intensified.

Figure 10: Pollution-reduction profiles – Horizontal equity

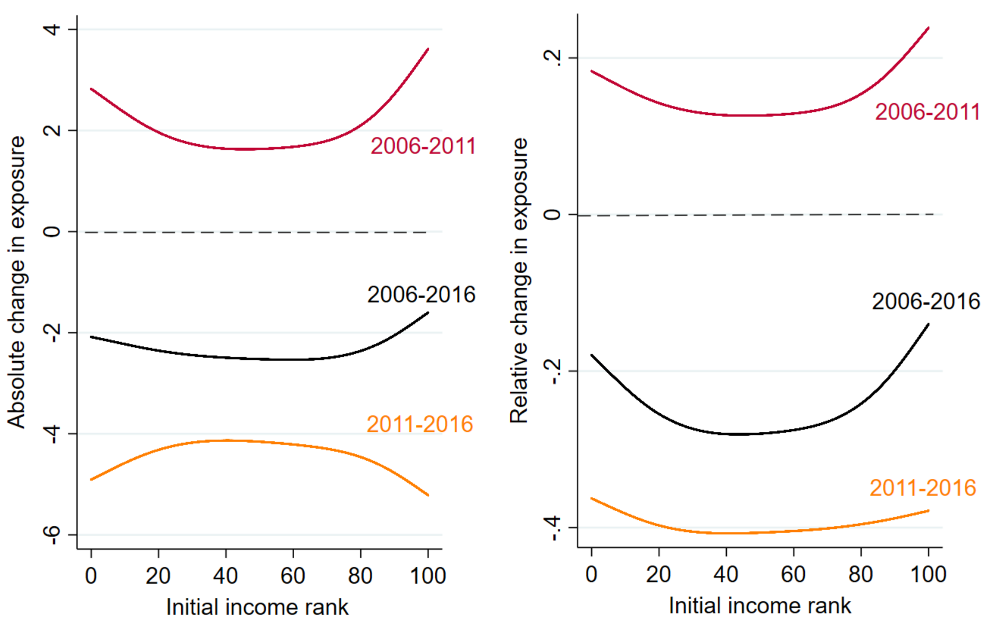

The initial income rank is the rank in the distribution of IRIS median incomes in 2006.

4 Role of Plans de Protection de l'Atmosphère

In France, part of the regulation of air quality occurs at the level of urban areas, through mandatory *Plans de Protection de l'Atmosphère*. As a consequence of a 2008 EU Directive that incorporated fine particulate matter as a newly regulated pollutant, urban areas were required to revise their schemes, so as to include measures aimed at reducing $PM_{2.5}$ concentration. Using an event-study design, this section investigates whether this change in policy helped reduce inequality in exposure to my pollutant of interest.

4.1 Context

The EU Directive 2008/50/EC on air quality, also named Directive on ambient air quality and cleaner air for Europe, includes 4 elements. First, it merged the majority of existing legislation on air quality into a single directive,[12] without any change in existing objectives. Second, it allowed for time extensions for compliance to EU standards regarding the concentration of particulate matter (PM_{10}), benzene, and nitrogen dioxide (NO_2), up to 2015. Third, it gives the opportunity for Member States to deduct emissions caused by natural sources, such as those emitted through forest fires, when assessing compliance to EU limit values of regulated pollutants. The fourth and final element of the Directive is of particular interest: it established that annual concentration in $PM_{2.5}$ has to be lower than 25 $\mu g/m^3$ by the 1st of January, 2015. In terms of exposure, the Commission chose to refer to a three-year annual average exposure (AEI, for Average Exposure Indicator), which must be lower than 20 $\mu g/m^3$. In France, the AEI is computed using monitor data of 49 urban areas. This became legally binding in 2015, i.e., starting for years 2013-2015. The Directive was translated into French law, and thus integrated into the *Code de l'Environnement*, by decree, on the 21st of October, 2010.[13]

The LAURE (Law on Air and Rational Use of Energy) of 1996 already compelled urban areas of more than 250,000 inhabitants to implement an Atmosphere Protection Plan (PPA, for *Plan de Protection de l'Atmosphère*), which, among other requirements, has to comprise a precise agenda of measures taken by local authorities so as to meet air quality standards.[14] The 2008 EU Directive

[12] The Directive 2008/50/EC merged all existing legislation on outdoor air quality, apart from the Fourth Daughter Directive 2004/107/EC, which regulates the concentration of metals, such as mercury and nickel, in ambient air.

[13] Said decree is available online at `https://www.legifrance.gouv.fr/affichTexte.do?cidTexte=JORFTEXT000022941254&categorieLien=id`.

[14] In addition to the action plan, the elements that a PPA must contain are: a) an inventory of emissions of atmospheric pollutants b) an evaluation of air quality c) a description of the sanitary impacts of air pollution d) an evaluation of the measures taken, in the form of scenarios.

Figure 11: Evolution of average $PM_{2.5}$ exposure – Whole PPA sample and selected PPA zones

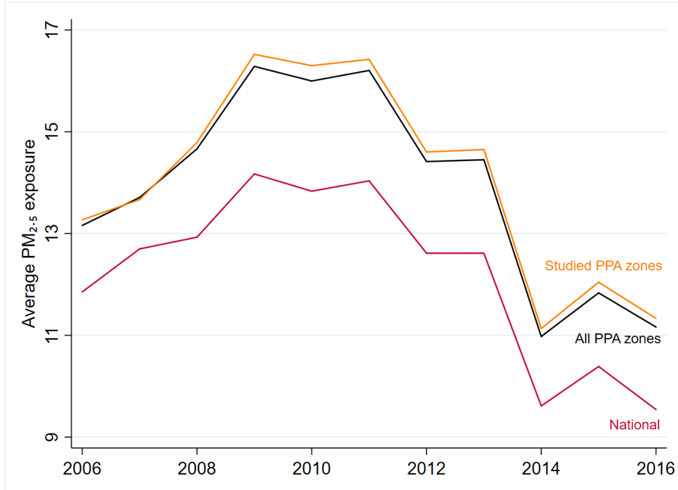

led the French government to modify the list of pollutants regulated within the PPA framework: fine particulate matter ($PM_{2.5}$), which was not included, became so. Hence, although measures that aimed at decreasing concentration in other air pollutants like PM_{10} which were implemented beforehand likely already helped mitigate $PM_{2.5}$ concentration, the date of implementation of post-2010 PPAs may constitute a shock in the concentration of fine particulate matter.

Since the 2010 decree, Atmosphere Protection Plans may be implemented if at least one of these conditions is fulfilled: a) the zone comprises an urban area of more than 250,000 inhabitants, b) limit or target concentrations of at least one pollutant are exceeded within the zone c) there is a risk that limit or target concentrations will be exceeded within the zone. To this day, 5,883 French municipalities belong to a zone covered by a *Plan de Protection de l'Atmosphère*. Using IRIS-level data, this represents 13,553 observations each year, but sample restrictions described in Section 4.2.2 drive the yearly number of observations down to 12,853. However, the selected areas represent 27.7 million inhabitants in 2016, i.e., 43% of the metropolitan French population. By construction, as shown in Figure 11, urban areas that implemented PPAs are more exposed to fine particulate matter than the national average. Notwithstanding, similarly to the average national reduction of 33% in $PM_{2.5}$ exposure between 2009 and 2016, average exposure decreased by 31% on average in all PPA zones, as well as in selected areas, although the latter are on average very slightly more exposed to $PM_{2.5}$ than the whole sample (see Figure 11). Figure 12 gives a sense of the spatial extent of these policies. Some extend up to entire administrative regions, like Île-de-France, or former Haute-Normandie and Nord-Pas-de-Calais.

Figure 12: Zones covered by a *Plan de Protection de l'Atmosphère*

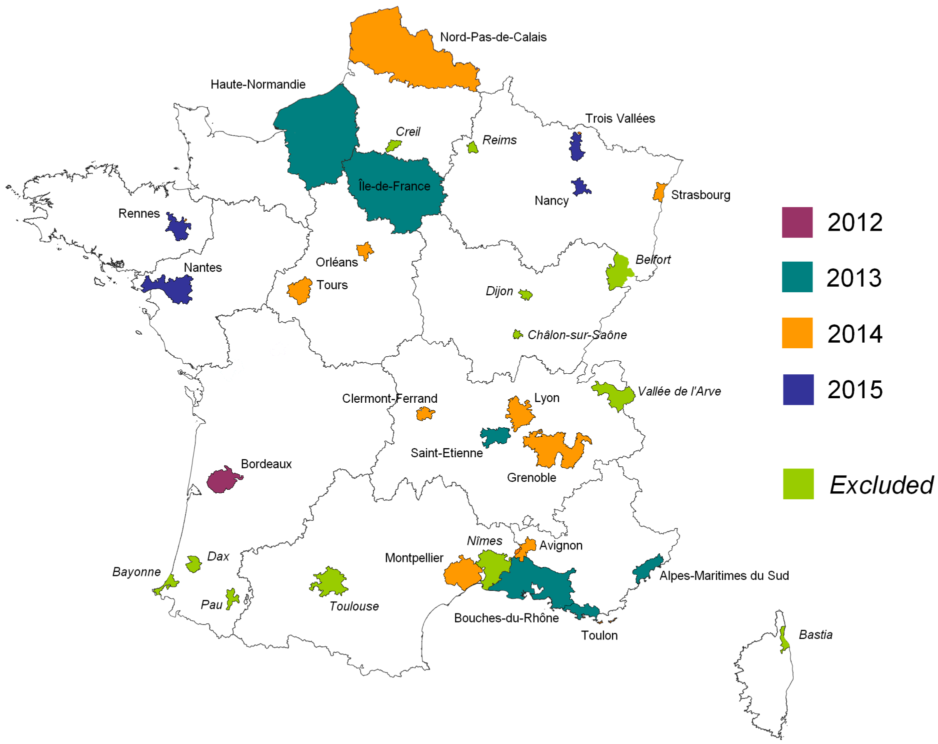

In order to comply with the law, urban areas are required to evaluate their Atmosphere Protection Plan every 5 years, and to revise it if need be. The time needed to evaluate the existing plan, to prepare new measures in conjunction with stakeholders, and to adopt the revised PPA sometimes implies that a longer time period elapses between the adoption of a scheme and its revision. In the dataset, the revised plans were passed 7.3 years after the former one was adopted, on average. By construction, this excludes urban areas that were not covered by a PPA before the beginning of the study period, displayed in italics in Figure 12 and Appendix Table 10.

4.2 Methodological framework

4.2.1 Event-study design

Given that the years of implementation of the revised *Plans de Protection de l'Atmosphère* vary, I evaluate the impact of the policy change described above on IRIS-level exposure to fine particulate matter within an event-study framework. Prior to investigating potential unequal benefits from new PPA adoption, the baseline equation (2) focuses on the evolution of exposure to $PM_{2.5}$ in the years to and from the event.

$$\ln(PM_{izct}) = \alpha + \sum_{j=-5}^{4} \mu_j \, \mathbb{1}\{t = j\} + X'_{ict}\eta + \gamma_c + \lambda_z + \varepsilon_{izct} \tag{2}$$

$\ln(PM_{izct})$ is the natural logarithm of particulate matter exposure for census block i, located in PPA zone z, during calendar year c, and year relative to PPA implementation t. Coefficients μ_j are the coefficients of interest of this baseline equation, and capture the change in exposure prior to and after the event.

The matrix X'_{ict} contains a series of census block characteristics, listed in Table 6 in Appendix, as well as census block initial pollution level, since the pollution path of an IRIS likely is dependent on prior levels. Initial pollution level is defined as $PM_{2.5}$ concentration in 2009, the year prior to the signing of the decree that enacted the change in policy. It is defined based on a calendar year ($c = 2009$) and not based on a relative year t. At first sight, it would be sensible to assume that pollution levels follow a (partly) auto-regressive process. Hence, looking at the case where initial pollution level would be defined as that of $t = -2$, I focus on the Bouches-du-Rhône and Montpellier PPA zones, which are geographically close to each other. They adopted their new plans in 2013 and 2014, respectively, meaning that their initial pollution level would be set as being that of 2011 and 2012. But if there was a common year-specific shock to pollution levels (due to similar meteorological conditions, or a generalised lockdown, for instance) in 2012 that would affect both zones, then, one would not be controlling for comparable initial $PM_{2.5}$ concentrations. Therefore, in order to avoid this, initial $PM_{2.5}$ concentration is defined as that of 2009.

Regressions also include calendar-year and time-varying PPA-zone fixed effects. Calendar year fixed effects allow to non-parametrically control for time trends, while year-PPA zone fixed effects allow to take into account potential unobserved heterogeneity in the measures enacted across zones. Given that variables are at the IRIS level, as opposed to the individual level, all regressions are weighted using analytic weights that correspond to the population in each IRIS. Standard errors are clustered at the PPA zone level to account for within-zone autocorrelation.

I first incorporate income in the study by assessing the relationship between the median income of a neighbourhood and the degree to which it benefitted from the implementation of a revised PPA without discretising the variable. A significant effect of income would provide inceptive evidence of the fact that it is a modification factor of the impact of the adoption of a revised PPA on $PM_{2.5}$ exposure. Denoting inc_i the median income of an IRIS i, I estimate:

$$\ln(PM_{izct}) = \alpha + \sum_{j=-5}^{4} \mathbb{1}\{t=j\}(\mu_j + \beta_j \ln(\text{inc}_i)) + X'_{ict}\eta + \gamma_c + \lambda_z + \varepsilon_{izct} \tag{3}$$

Then, I evaluate the potential differentiated impact of the implementation of revised PPAs between the different groups of income by replacing the (logarithm of) neighbourhood median income by a dummy $\mathbb{1}\{\text{ADV}_i\}$. The discrete equivalent of equation (3) writes as follows:

$$\ln(PM_{izct}) = \alpha + \sum_{j=-5}^{4} \mathbb{1}\{t=j\}\,(\mu_j + \beta_j\,\mathbb{1}\{\text{ADV}_i\}) + X'_{ic}\eta + \gamma_c + \lambda_z + \varepsilon_{izct} \tag{4}$$

In the main specifications, $\mathbb{1}\{\text{ADV}_i\}$ is a dummy equal to one when census block i is considered as advantaged at the beginning of the period, and 0 otherwise. This dummy can take two different forms: it can refer to the initial location in the national distribution of income or to the initial location in the PPA zone distribution of income. Specifically, $\mathbb{1}\{\text{ADV}_i\}$ takes the value 1 when the census block was located in the top 50% of the national (/PPA-zone) distribution of census block median incomes at the beginning of the study period, and 0 when it belonged to the bottom 50%. I also provide estimations that define ADV_i as a 4-category variable representing the 4 quartiles of the national or PPA-zone income distribution, taking the first quartile as a reference. Coefficients β_j, which are associated with the interaction term between belonging to an initially advantaged group and the year relative to PPA implementation, are the coefficients of interest.

4.2.2 Assumptions and sample selection

In an event-study design, the following assumptions need to be fulfilled. First, the outcome should not have diverged in $t \geq 0$ absent the event. In other words, as long as there are no systematic changes between neighbourhoods nor PPA zones over time, except for the treatment, the coefficient of interest can be interpreted as causal. This could arise from different patterns of in- and out-migration, for instance if certain neighbourhoods become increasingly attractive to a certain type of population relative to others, since I do not observe individual movements. Note that focusing on initial median income instead of a varying version of the variable (e.g., $\ln(\text{inc}_{ict})$) allows to circumvent part of this potential bias: if income would vary across (relative) years, it could be due to these mobility patterns. In addition, I use a rich set of controls for IRIS characteristics in X_{ict}, as well as a control for time-varying PPA zone fixed effects, which account for remaining heterogeneity between urban areas.

Second, the outcome in the reference category of the relative time variable should be unaffected by the event. The chosen reference category is thus $t = -1$, for both relative year fixed effects (μ_t) and relative year interacted with the income dummy (β_t). Indeed, some schemes were adopted early in $t = 0$, meaning that the event may have an effect at $t = 0$, the first, even though non-full, year of implementation. Hence, μ_j coefficients measure the impact of the revised plan relative to the year just before the adoption of the plan. However, there exists another threat to identification. 12 urban areas were not yet covered by an Atmosphere Protection Plan prior to the 2010 decree, but adopted a plan for the first time during the study period. They did so on the basis of the fact that they showed an exceedance of air quality standards for at least one of the regulated pollutants. All these exceedances do not necessarily involve fine particulate matter. However, $PM_{2.5}$ are emitted through various sources that also diffuse other regulated pollutants. For instance, traffic is a source of both $PM_{2.5}$ and nitrogen dioxide (NO_2). Hence, the positive correlation of the concentration in other pollutants with that in $PM_{2.5}$ implies that including zones that adopted a PPA on the sole basis of them exceeding air quality standards is a threat to identification, since the timing of the event would be correlated with the outcome variable. I thus exclude urban areas that were not yet covered by an Atmosphere Protection Plan prior to 2010 from the analysis. By doing so, the second assumption of the event-study design, which states that $PM_{2.5}$ concentration in $t = -1$ should not be affected by the adoption of a PPA in $t = 0$, is likely to hold.

Remaining sample and timeframe restrictions are the following. First, given that in order to comply with the law, urban areas are required to revise their Atmosphere Protection Plan every 5 years, I reduce the number of relative years prior to the new PPA implementation to 5. I study the 4 years that follow the adoption of the new plan due to data restrictions: the study period extends to 2016, but new PPAs were adopted only starting in 2012. Most of the plans were revised in 2013 and 2014; for instance, the new scheme for Île-de-France was adopted in 2013. PPAs adopted in 2016 were dropped from the analysis. This implies that I do not study the impact of the measure for the Toulouse urban area.[15] A summary of adoption dates and of the number of observations per urban area, which distinguishes between those that revised their plan and those that did not have one prior to 2010, is made available in Table 10 in Appendix. The ensuing number of observations per relative year t is available in Table 11 in Appendix.

[15] The Creil and Nîmes urban areas also adopted their plan in 2016, but they were already discarded from the analysis due to the fact that they did not have any PPA before the 2010 decree passed.

4.2.3 Quantile regression

The pollution-reduction profiles computed in Section 3.4 hint at the fact that the decrease in $PM_{2.5}$ concentration that occurred between 2011 and 2016 was not uniform over the whole pollution distribution, and that this change was progressive, in the sense that higher (absolute) gains in air quality accrued to initially most exposed individuals. Although these computations do not comprise only studied PPA zones, the latter represent 43% of the 2016 population. As such, it is worth examining how the adoption of new Atmosphere Protection Plans affected different quantiles of the distribution of fine particulate matter exposure. This will allow to understand whether the adoption of new PPAs contributed to the observed difference in air quality improvements across the pollution distribution.

Moreover, the comparison between the cumulative distribution functions of $PM_{2.5}$ exposure before and after the event, depicted in Figure 20 in Appendix, is an indication that there may be some heterogeneity in the change in exposure resulting from PPA adoption, depending on the census block's location in the $PM_{2.5}$ distribution. Indeed, had the effects been unconditionally homogeneous, the cumulative distribution of pollution at $t = 1$ would have been a perfect translation of the cumulative distribution at $t = -1$, which is not the case. Hence, in addition to the question being worth investigating in itself, the data calls for a quantile regression approach.

One of the main advantages of quantile regression is that it relies on weaker assumptions than standard OLS does, as error terms need not be identically distributed. Nonetheless, they are still assumed to be independent from one another. Hence, PPA zone fixed effects are, for now, assumed to drain all the autocorrelation out of the model. There is an additional requirement that could potentially not be fulfilled when applying quantile regression techniques to this study's data. Indeed, the dependent variable should be very continuous, as quantile regression performs less well when there are many ties at some values of Y. For instance, this may happen if two IRIS are located within one single raster grid of 1 km^2, thus taking the same value of $PM_{2.5}$ concentration. For instance, in Paris and its inner suburbs, the average surface area of an IRIS is .3 km^2, which indeed creates ties at some values of the dependent variable, but this issue remains marginal.

Defining the τ-th conditional quantile function $Q_{Y|X}(\tau) = X'\beta_\tau$, the coefficient of interest β_τ must be interpreted as the average marginal effect of X on the conditional quantile of Y. As census blocks likely have a different position in the pollution distribution depending on the value of X, β_τ cannot generally be interpreted as the effect of a small variation in X for *census blocks* at the τ-th quantile of $Y|X$. However, if census blocks have the same ranking in the distribution of $Y(x)$ whatever x, that is, if the rank invariance assumption holds, then the latter interpretation is valid. This is likely not to hold unconditionally on urban area, since census blocks that have similar socioeconomic characteristics may be located in either the very polluted North of France, or in the less polluted South-West. This assumption would thus theoretically be more likely to hold thanks to the use of PPA-zone fixed effects. Nonetheless, there likely is some heterogeneity in the patterns relating income to $PM_{2.5}$ exposure within PPA zones as well. As such, the rank similarity condition does not hold conditionally on socioeconomic characteristics, but likely does conditionally on socioeconomic characteristics, PPA-zone fixed effects, and neighbourhood-level initial pollution level. Nevertheless, formally testing for rank similarity in an event-study design is not straightforward. Indeed, existing methods proposed by, e.g., Frandsen and Lefgren (2018) or Dong and Shen (2018) are applicable to a program evaluation framework, but cannot be adapted to an event study without a control group. As such, I cannot and do not formally test this assumption.

4.3 Results

4.3.1 Indirect assumption tests

Prior to interpreting the results, I look at whether the assumptions formulated in Section 4.2.2 and 4.2.3 are likely to hold. Indeed, although they are untestable by construction, the event-study design provides an indirect test of these assumptions, by focusing on coefficients μ_t for $t < 0$. Them being statistically significant would imply that exposure to $PM_{2.5}$ in PPA zones was not stable during the years leading up to the adoption of a new PPA. This would imply that the timing of the event would likely be endogenous, even conditionally on the control variables, which would hamper identification. This test is also necessary for β_t coefficients for $t < 0$. Indeed, if initially richer neighbourhoods were to trend differently from initially disadvantaged neighbourhoods prior to the event, I would not be able to infer a causal impact. This test is presented for both baseline event-study coefficients μ_j and income-related β_j in Table 12 in Appendix. It should be remembered that the reference category is $t = -1$, hence the absence of an associated coefficient in Table 12. The check can also be performed based on a visual inspection of the pre-trends in Figures 13 and 14. None of the pre-trend μ_j coefficients are statistically different from zero, which supports the validity of the baseline event-study design. This also holds for β_j coefficients, apart from $\beta_{t=-5}$,

Figure 13: Effect of PPA adoption on average exposure

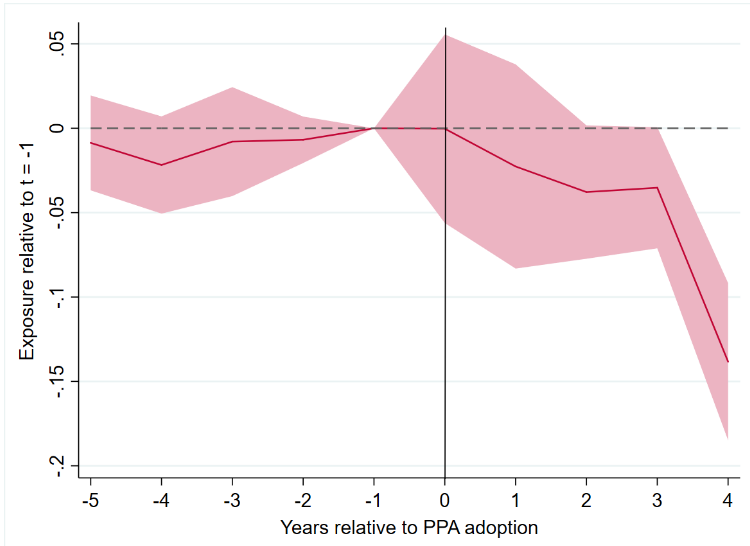

Note: Estimates of coefficients μ_j in equation (2). The shaded area corresponds to the 95% confidence interval.

which is statistically significantly positive at the 10% level. Given that this coefficient accounts for a differential trend 5 years prior to the event, this does not seem alarming.

4.3.2 Baseline results

I start by examining the effect of the adoption of revised PPAs on average exposure to fine particulate matter. Figure 13 shows the estimates of the event-study coefficients μ_j in equation (2), and Table 3 presents the detailed coefficients, with all coefficients multiplied by 100 for the sake of readability. The implementation of a new Atmosphere Protection Plan does not have any significant effect on exposure during the year it is implemented, nor the following one. Everything else equal, two years after the adoption of a new PPA, $PM_{2.5}$ concentration decreased by $(e^{-.0378} - 1 =)$ 3.71% on average relative to the year before the signing, an effect significant at the 10% level. This effect is stable through the third year following the event, with an average decrease of $(e^{-.352} - 1 =)$ 3.46%, also significant at the 10% level, and is similar when controlling for initial (log) income. Average exposure throughout metropolitan France decreased by 33% between 2009 and 2016, and by 31% on average in the urban areas I study here. Hence, the identified effects remain significantly smaller than the average observed decrease in $PM_{2.5}$ exposure, suggesting that even if a different year of reference is used for every different year of implementation, the adoption of revised *Plans de Protection de l'Atmosphère* did not drive much of the improvement in air quality.

Finally, relative year $t = 4$ is associated with a much greater decrease in exposure (-13%), but this effect only concerns the Bordeaux PPA zone, since it is the only studied urban area that adopted its plan in 2012, thus allowing to observe the outcome variable 4 years as of the event. This brings to light a weakness of this study, namely the lack of post-event years. Indeed, the ACAG has not yet disseminated $PM_{2.5}$ data for years after 2016, which restricts the study period, and thus, the number of exploitable relative-year timepoints.

I start by looking at the results associated with equation (3), shown in Column (1) of Table 3. Note that again, in Table 3, all coefficients and standard errors were multiplied by 100 for the sake of readability of coefficients in Columns (2) and (3). Similarly to the estimates of coefficient $\mu_{t=1}$, the first post-event estimate $\hat{\beta}_{t=1}$ is not statistically different from zero, suggesting an absence of differential effect between initially advantaged and initially disadvantaged neighbourhoods in the first year after the adoption of the new scheme. Starting two years after, it appears that income indeed alters the impact of the event on exposure. For instance, on average at $t = 2$, in addition to the 3.7% average decrease in exposure, an additional 10% in initial ($t = -1$) income allows to benefit from a .37% lower exposure on average. However, at $t = 3$, this negative effect is a bit smaller. Taken together with the fact that the absolute value of $\hat{\mu}_{t=3}$ is greater than that of $\hat{\mu}_{t=2}$, though not significantly, this alludes that, three years after the event, after accruing chiefly to higher income areas, air quality improvements may start benefitting more deprived areas as well. This cannot be attributed to mobility between neighbourhoods, which could induce variations in income, since this value is fixed. Again, observing later timepoints would prove handy to verify this finding.

Figure 14: Relative effect of PPA adoption on exposure of initially higher income areas (β_j)

(a) National distribution

(b) PPA zone distribution

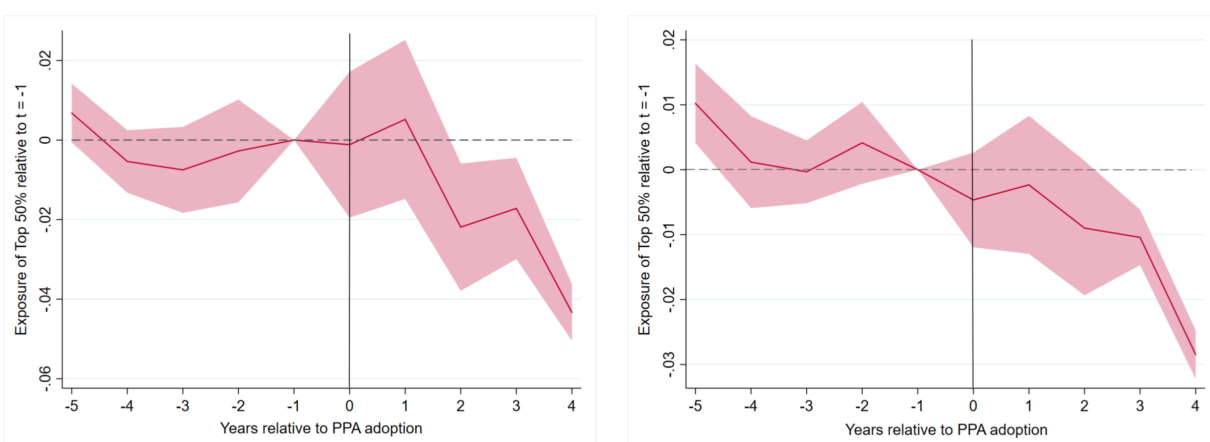

Note: Estimates of coefficients β_j in equation (4). The shaded area corresponds to the 95% confidence interval.

The results for the estimation of equation (4) taking $\mathbb{1}\{\text{ADV}_i\}$ as reflecting initial location in the national distribution of income are displayed in Column (3) of Table 3 and in Figure 14a. Again, there is no significant difference in exposure between initially advantaged and disadvantaged census blocks at $t = 1$. However, starting from the second year after the event, it appears that neighbourhoods that were located in the top 50% of the national distribution of income prior to the event benefitted more from the implementation of a new air quality scheme than those that were located in the bottom 50%. For instance, 2 years after the implementation of a new PPA, living in an initially advantaged census block in terms of income implies that one's exposure would be 2.2% lower than if she lived in an initially disadvantaged census block, *ceteris paribus*. Hence, it appears that neighbourhoods that were initially more advantaged in terms of income obtained larger benefits from the measures taken so as to reduce $PM_{2.5}$ exposure. One could argue that this result may be driven by the fact that there are initial differences between PPA zones, because some of them may comprise more initially disadvantaged neighbourhoods than others, and because the implementation of a revised air quality plan may be more or less effective depending on PPA zones, but this confounding effect should be taken into account by the year-PPA zone fixed effect. As the effect is smaller at $t = 3$, one may reiterate the hypothesis that this differential impact may dissipate with time. Moreover, it appears that, as I control for the indicator $\mathbb{1}\{\text{ADV}_i\}$, the overall impact on exposure, captured by the estimates in the upper part of the table, becomes smaller in magnitude and insignificant. Hence, it seems that *only* neighbourhoods located within the top 50% of the national income distribution significantly benefitted from the adoption of a revised PPA.

Next, I reshape the variable $\mathbb{1}\{\text{ADV}_i\}$ and define it based on the neighbourhood's location in the distribution of income of its PPA zone. The corresponding results are provided in Column (3) of Table 3 and Figure 14b. The estimates of β_j are qualitatively similar to those obtained using other definitions of income, in the sense that there is a differential effect only 2 years after the event, which only seems to benefit neighbourhoods located above the median of the income distribution. In other words, one recovers "national" trends even within urban areas. However, the estimates are smaller in magnitude: at $t = 2$, on average, an initially advantaged neighbourhood has a $PM_{2.5}$ exposure that is .9% lower than an initially disadvantaged neighbourhood, *ceteris paribus*. Contrarily to the results of previous specifications, this effect does not seem to be attenuated at $t = 3$. This differential effect in favour of higher-income neighbourhoods could be attributed to the fact that census blocks located in historical centres of cities or urban areas, which may be wealthier

Table 3: Event study results – Change in exposure relative to $t = -1$

Variable	(1) Baseline	(2) Continuous	(3) 50/50 national	(4) 50/50 PPA zone
$t = 0$	-.026	-15.260	-.009	.204
	(2.851)	(9.116)	(2.766)	(2.861)
$t = 1$	-2.262	-9.116	-2.671	-2.149
	(3.086)	(12.669)	(3.457)	(3.305)
$t = 2$	-3.781*	-3.681***	-2.574	-3.340
	(2.015)	(1.220)	(2.266)	(2.167)
$t = 3$	-3.523*	-3.804***	-2.368	-3.012
	(1.831)	(.787)	(1.960)	(1.960)
$t = 4$	-13.828***	-53.110***	-10.648***	-.881
	(2.376)	(4.945)	(2.427)	(.679)
$(t = 0) \times$ income var.		.169	-.116	-.465
		(.167)	(.936)	(.371)
$(t = 1) \times$ income var.		.067	-.519	-.234
		(.106)	(1.020)	(.543)
$(t = 2) \times$ income var.		-.374***	-2.189**	-1.043***
		(.108)	(.816)	(.218)
$(t = 3) \times$ income var.		-.264***	-1.718**	-.899*
		(.063)	(.649)	(.428)
$(t = 4) \times$ income var.		-.517***	-4.333***	-2.847***
		(.038)	(.362)	(.190)
Intercept	30.783***	21.489***	29.754***	29.836***
	(4.562)	(7.275)	(4.516)	(4.223)
Initial pollution	X	X	X	X
Neigh. charac.	X	X	X	X
PPA zone FE	X	X	X	X
Year FE	X	X	X	X
# Obs.	102,704	102,704	102,704	102,704
# Groups	12,853	12,853	12,853	12,853
# Clusters	20	20	20	20
R^2	0.92	0.93	0.93	0.93

Coefficients multiplied by 100 for the sake of readability. Standard errors clustered at the PPA zone level in parentheses. *: 10% level, **: 5% level, ***: 1% level. Column (1) gives the results of equation (2), Column (2) the results of equation (3), and Column (3) and (4) the results of equation (4).

on average, were targeted by public authorities to be cleaned up. It may be worth obtaining more information on the exact measures undertaken by urban areas in the context of these revised air quality schemes in order to further analyse the reasons behind this difference.

Again, the results obtained for $t = 4$ are to be taken with caution, since they only regard the Bordeaux PPA area, but they are so specific they are worth analysing. As one can easily see in the graphical representation of Figure 14, the confidence interval for this specific timepoint is very small, which brings more confidence to the result. It appears that in this area, when I control for location in the national distribution of income, the overall effect at $t = 4$ is significant (-10%), with a greater reduction accruing to higher-income IRIS (additional -4% compared the lower-income IRIS). But on the other hand, the pattern is very different when controlling for location within the PPA zone income distribution: four years after the implementation, in the Bordeaux area, the estimates suggest that the effect of the PPA adoption is solely received by neighbourhoods located above the zone's median income (-3% on average, *ceteris paribus*).

In order to provide a more precise picture of the differential effects from the change in policy, I look at what takes place *within* the two halves of the distribution of PPA-zone income. Given that certain areas under study only comprise about 100 observations (see Appendix Table 10), I proceed to use quartile-based bins. As shown in Figure 21 in Appendix, similarly to the upper panel of Column (4) in Table 3, there is no significant impact of the event on average exposure at $t \geq 0$, abstracting from $t = 4$, when using this definition as a control for income. Selected results are presented in Table 4, using the neighbourhoods located in the first 25% of the PPA-zone distribution of income as a reference. Hence, there is no significant fall in $PM_{2.5}$ for neighbourhoods located below the 1st income quartile. At $t = 0$ and $t = 1$, the estimates associated with the 2nd and 3rd quarter of income are statistically significantly different from zero, while those of the upper quarter are not. This suggests that only the middle 50% of the PPA-zone neighbourhood income distribution received any improvement in air quality during the first 2 years of implementation, which likely explains the insignificance of the corresponding estimates when using the former 50/50 splitting definition. After 2 years, the relative impact of the event gets larger for the second quarter, but is still insignificant for the top quarter, and after 3 years, all income groups seem to have benefitted, to different extents, from the adoption of a revised air quality plan, apart from the bottom 25% of income. Indeed, 3 years after the event, all other things held equal, the bottom quarter does not benefit from any reduction, the second quarter a 1.2% one, the third quarter a 1.3% one, and the

Table 4: Change in exposure relative to $t = -1$ and PPA-zone 1st quartile of income

Variable	(1) 2nd quarter	(2) 3rd quarter	(3) 4th quarter
$(t = 0) \times$ income group	-.009*	-.006*	-.007
	(.005)	(.004)	(.009)
$(t = 1) \times$ income group	-.008**	-.011*	-.005
	(.004)	(.006)	(.009)
$(t = 2) \times$ income group	-.012**	-.013**	-.009
	(.005)	(.005)	(.007)
$(t = 3) \times$ income group	-.011***	-.016***	-.011*
	(.039)	(.036)	(.064)
$(t = 4) \times$ income group	-.037***	-.040***	-.054***
	(.002)	(.002)	(.003)

Standard errors clustered at the PPA zone level displayed in parentheses. *: 10% level, **: 5% level, ***: 1% level. Results based on the formulation of equation (4) defining ADV_i as a 4-category variable.

top quarter a 1.1% one. This is consistent with the previous finding that the top 50% of PPA-zone income did receive a 1% decrease in $PM_{2.5}$ exposure at $t = 3$. However, it is not shared equally within this group, nor is it within the bottom 50%. All in all, the $PM_{2.5}$ concentration improvements arising from the adoption of revised PPAs seem to follow a similar trend as the nationwide evolution depicted by the horizontal-equity pollution-reduction profiles of Section 3.4.2: larger benefits seem to accrue to the middle of the income distribution while the two ends obtain little to no benefit. Given that, on average, the middle of the income distribution is also initially less polluted, results would suggest that the PPA policy reinforced preexisting inequality in exposure to $PM_{2.5}$. One must recall that the magnitude of the effects is small compared to the overall decrease in exposure experienced during the study period, meaning that the scope of this regressive impact was limited.

4.3.3 Quantile regression results

I also explore how the adoption of revised *Plans de Protection de l'Atmosphère* affects different quantiles of the distribution of fine particulate matter. Figure 15 depicts the effect three years after the event, by vigintile of initial exposure, while Appendix Figure 22 presents the equivalent for $t = 2$. Under the assumption that, conditionally on socioeconomic characteristics, PPA-zone fixed effects and neighbourhood-level initial pollution level, the rank invariance condition holds, one can

Figure 15: Quantile event-study effects – Focus on $t = 3$

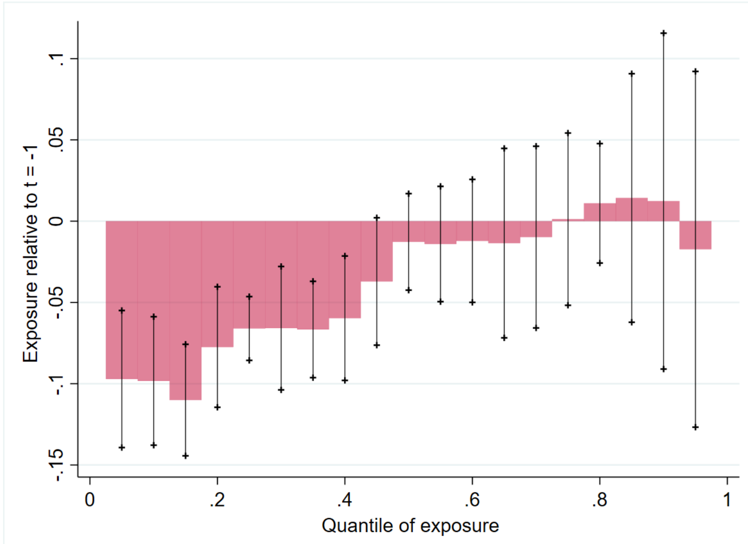

interpret the coefficients as follows: everything else equal, PPA adoption leads to a decline of 10% in exposure on average, for neighbourhoods located at the first to third vigintile of exposure. Given that this assumption is not formally tested for, however, it may be argued that these effects should be interpreted as marginal effects on the conditional quantile of exposure, instead of neighbourhoods located at this quantile. In any case, there is a significant benefit from the adoption of a revised PPA only up to the 4^{th} decile of initial exposure. From the latter to the 7^{th} decile, the effect is also negative, but not significantly. And finally, there might even be an economically positive effect of the policy change at the upper vigintiles, although it is not statistically significantly different from zero. The confidence intervals are substantially larger on the right-hand side of the graph than on the left-hand side, suggesting that there is greater heterogeneity in the effects of PPA implementation at higher quantiles of exposure.

Finally, the degree to which income alters the effect of the new policy may vary depending on the initial level of exposure. For instance, one could hypothesise that in a large PPA zone, conditionally on wanting to target poorer neighbourhoods, policy-makers may be more likely to specifically target more polluted areas than less polluted ones. I thus investigate the benefits for neighbourhoods whose median income lay either at the bottom or the top 50% of the national and PPA-zone income distributions at different quantiles of initial pollution level. Note that unlike previously, the splitting based on the position in the PPA-zone income distribution may arguably not be as informative, because the support of the distribution of $PM_{2.5}$ within PPA zones may be narrower than that of the entire sample, especially those of limited spatial extent, e.g., these of Orléans or Clermont-Ferrand.

Figure 16 presents the quantile regression results three years after the event, distinguishing between neighbourhoods that are initially located above *vs* below the median of the national (Figure 16a) and PPA-zone (Figure 16b) income distribution. The equivalent for $t = 2$ is shown in Figure 23 in Appendix. It appears that at the lower quantiles of the pollution distribution, neighbourhoods that were initially below the median income benefitted from substantially lower air quality improvements than their top 50% counterparts. This difference is reversed and very small when looking at the estimates within the PPA-zone income distribution. At upper quantiles of the pollution distribution, the results are coherent with those of Figure 15, in two respects. First, the effects are not statistically different from zero, again hinting at heterogeneity, this time even within income groups. Second, abstracting from significance, the estimates suggest that while census blocks located at the top 50% of the PPA-zone distribution received a reduction in exposure, though to a lower extent than those at lower quantiles of exposure, those at the bottom saw theirs slightly rise. For the sake of visual clarity, the confidence intervals are omitted, but either way, at any vigintile, they overlap, meaning that I cannot reject the null hypothesis that there is no income-related difference in benefits from the adoption of PPAs. This might be due to a lack of power arising from the fact that the specification may be too demanding compared to the number of available observations.

Figure 16: Quantile event-study effects – Focus on $t = 3$, based on initial income

(a) National distribution (b) PPA-zone distribution

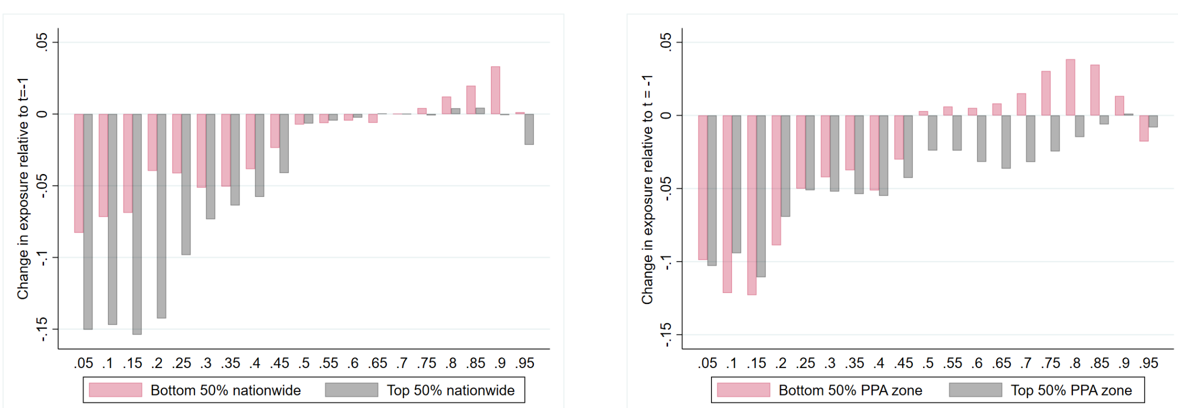

Note: Confidence intervals were omitted for the sake of visual clarity. Estimates are not statistically significant above the 4$^{\text{th}}$ decile, and are not significantly different from each other within the same vigintile when I distinguish between income groups.

4.4 Robustness to spatial autocorrelation

Clustering standard errors at the PPA zone level strips the effect of autocorrelation within PPA zones away, but, as in Section 3.2, this likely is not enough for the model errors to be independently distributed. Indeed, the same argument as in Section 3.3 holds: clustering errors at the PPA zone level amounts to assume that there are discrete shifts in the distribution of residuals at the cluster zone boundary, while they are probably continuously distributed across space. In other words, PPA zones that are contiguous, such as those of Bouches-de-Rhône, Avignon and Toulon, which are likely to exhibit correlated pollution levels, are not considered as neighbouring in the models of the previous section, due to the fact that clusters are not defined continuously, but based on administrative divisions. The argument of continuousness also pertains with regard to non-neighbouring PPA zones. For instance, since fine particulate matter can travel long distances, especially when it is carried by speedy winds, the effectiveness of the measures implemented within the PPA framework in Grenoble might influence pollution levels in Bouches-du-Rhône. Moreover, computing Moran's I test-statistic[16] confirmed that there is remaining positive spatial autocorrelation in the residuals, meaning that high (resp., low) values of residuals are located close to high (resp., low) values of residuals in the 2 areas that encompass the PPA zones of Avignon, Bouches-du-Rhône and Toulon, and those of Haute-Normandie and Île-de-France.[17] Test results are provided in Appendix Table 13. This is coherent with the fact that fine particulate matter is itself positively spatially autocorrelated. Regarding the modelling approach, the same reasoning as in Section 3.3 applies: spatial lag models involve making arbitrary and potentially consequential parametric choices on the form of spatial autocorrelation, and modelling it in a smoother fashion seems relevant. Hence, I opt for a Generalised Additive Model, meaning that I add a smoothing spline of longitude x_i and latitude y_i of the centroid of IRIS i as a potential predictor of (log) $\mathrm{PM}_{2.5}$ concentration. Equation (4) thus rewrites:

$$\ln(PM_{izct}) = \alpha + \sum_{j=-5}^{4} \mathbb{1}\{t=j\}\left(\mu_j + \beta_j\, \mathbb{1}\{\mathrm{ADV}_i\}\right) + X'_{ic}\eta + \gamma_c + \lambda_z + s(x_i, y_i) + \varepsilon_{izct} \qquad (5)$$

[16]Specifically, Moran's I (1950) test-statistic ranges from -1 (indicating perfect negative autocorrelation) to 1 (indicating perfect positive correlation. A value of 0 means that there is no autocorrelation across residual values of the estimated model.

[17]In order to conduct the test, I use a spatial weight matrix built using inverse-distance weighting. Missing values are automatically created since only certain parts of the French country are studied. To my knowledge, the only way to handle missing values in this context is to interpolate values over neighbours. However, given the spatial scale of missing observations in this context (see Figure 12), this approach is not adapted. Hence, I only run the tests on 2 sets of contiguous PPA zones, and separately for each aggregate.

I keep PPA-zone fixed effects so as to account for potential heterogeneity in the measures implemented as part of each specific *Plan de Protection de l'Atmosphère*. $\mathbb{1}\{\text{ADV}_i\}$ is equal to 1 when the census block belongs to the top 50% of the national or PPA-zone income distribution. The model is estimated in the same fashion as in Section 3.3, using generalised cross-validation to automatically select the degree of smoothness of the spline.

Table 5 presents the corresponding results, with coefficients multiplied by 100 for ease of reading. The results are rather different from those obtained in the previous subsection, in which I concluded that only the top 50% of the national or PPA-zone distributions of income perceived air quality improvements from the adoption of revised PPAs. Indeed, once I allow for spatial autocorrelation, the average effect μ_j is not only significantly negative, but it is so as soon as one year after the event. Moreover, the magnitude of the coefficient is non-negligible, which comes to slightly contradict the previous conclusion that the revision of PPAs played a minor role in the decrease in $PM_{2.5}$ concentration. Indeed, according to these results, 2 years (resp., 3 years) after the revision of a PPA, there was an average decrease of 6% to 7% (resp., 5%) in exposure, *ceteris paribus*. This fall was even greater in the Bordeaux area (-15% on average). Turning to estimates of β_j, however, the results are similar to those of the lower panel of Table 3, be it in terms of significance or in terms of magnitude. They imply that, regardless of whether one looks at the income distribution at the national level or within PPA zones, initially advantaged neighbourhoods perceived additional gains from this policy change of around 2%, i.e., 25% to 40% higher benefits than the initially disadvantaged neighbourhoods. The discrepancy observed for μ_j is thus absent for β_j. This may stem from reasons akin to those of Section 3.3. The $PM_{2.5}$ concentration of a census block influences, and is influenced by, that of nearby census blocks. Therefore, by assuming that observations are randomly distributed across space, one may fail to detect variations that are statistically significant at higher aggregation levels than the high-resolution IRIS. In other words, allowing for spatial autocorrelation rightly "clusters" census blocks that have comparable outcomes, and allows not to underestimate the impact encapsulated by μ_j. However, the results for β_j are similar to previous ones. This can at least partly be attributed to residential segregation patterns: neighbourhoods whose income is located above the median are very likely to be located close to each other, while those whose income is located below the median are very likely to be located further away, but close to each other as well. Hence, in the baseline equation (4), the dummy variable $\mathbb{1}\{\text{ADV}_i\}$ in itself already captured part of the spatial autocorrelation.

Table 5: Event study results – Generalised additive models

Variable	(1) Baseline	(2) 50/50 national	(3) 50/50 PPA zone
$t = 0$	-3.191	-3.943	-3.451
	(3.682)	(2.847)	(2.850)
$t = 1$	-7.311***	-7.261**	-7.189**
	(1.944)	(2.174)	(2.057)
$t = 2$	-6.637***	-6.074***	-6.182***
	(1.789)	(1.934)	(1.890)
$t = 3$	-5.921***	-5.970***	-5.812***
	(1.660)	(1.813)	(1.763)
$t = 4$	-17.614***	-14.441***	-16.199***
	(4.956)	(.627)	(.553)
$(t = 0) \times$ income var.		.120	-.059
		(.128)	(.120)
$(t = 1) \times$ income var.		-.276**	-.258**
		(.128)	(.123)
$(t = 2) \times$ income var.		-2.680***	-1.087***
		(.131)	(.126)
$(t = 3) \times$ income var.		-1.797***	-.848***
		(.148)	(.140)
$(t = 4) \times$ income var.		-4.468***	-3.043***
		(.539)	(.476)
Intercept	-9.384***	-8.798**	-7.998**
	(3.552)	(.892)	(.897)
Initial pollution	X	X	X
Neigh. charac.	X	X	X
PPA zone FE	X	X	X
Year FE	X	X	X
p-value $s(x_i, y_i)$	$<2 \times 10^{-16}$	$<2 \times 10^{-16}$	$<2 \times 10^{-16}$
# Obs.	102,704	102,704	102,704
# Groups	12,853	12,853	12,853
R^2	0.96	0.97	0.97

Coefficients and standard errors multiplied by 100 for the sake of readability. Standard errors displayed in parentheses. *: 10% level, **: 5% level, ***: 1% level. Column (1) gives the results of equation (2), and Column (2) and (3) the results of the two forms of equation (4).

5 Conclusion and discussion

Throughout this study, the matching of IRIS-level information on income and other neighbourhood characteristics with high-resolution satellite-based data on ground-level fine particulate matter ($PM_{2.5}$) concentration allowed to draw a picture of current national-scale patterns of inequality in exposure to this harmful pollutant. According to population-weighted estimates, average $PM_{2.5}$ exposure in metropolitan France amounts to 9.54 $\mu g/m^3$ in 2016, which is lower than the WHO guideline. Nonetheless, this average masks pronounced heterogeneity, both across space and across income groups. While no clear trend seems to emerge across *départements*, descriptive evidence suggests that, at the national level, the unconditional relationship between $PM_{2.5}$ exposure and neighbourhood-level income decile follows a U-shape. Indeed, in 2016, only those located at the lower and upper tail of the distribution of income are exposed to average concentrations that are above the WHO standard of 10 $\mu g/m^3$. Specifically, 60% of neighbourhoods of the bottom 10% of income are exposed to $PM_{2.5}$ levels that exceed the WHO standard, while, conversely, 80% of the neighbourhoods of the next 10% comply with this standard. Turning to the top 20% of the distribution, the highest levels of exposure are also experienced by neighbourhoods whose income lays above the 9^{th} decile of income. These patterns are coherent with the fact that despite different levels of residential segregation (Quillian and Lagrange, 2016), statistically more polluted city centres and inner suburbs usually compound both high- and very low-income neighbourhoods, while less polluted peri-urban areas are increasingly well-off (Aerts et al., 2015; Floch, 2014, 2017).

However, these findings are not reflected within urban areas, as, on average, exposure appears to decrease as a neighbourhood gets higher up the distribution of income. Fixed-effect models then allow to exploit the panel structure of the data to get rid of the confounding impact of unobserved neighbourhood-level heterogeneity, and confirm that, over the 2006-2016 period in France, higher neighbourhood income is indeed associated with lower $PM_{2.5}$ exposure. This result is robust to formally accounting for spatial autocorrelation using a smoothing spline of census block coordinates. I also find a significantly positive relationship between the share of immigrants and fine particulate matter concentration, hinting at the fact that there may be an ethnic gap in pollution exposure in France as well. This positive link is stronger when accounting for spatial autocorrelation, which not only confirms the need to pay great attention to this issue, but also brings to light one of the consequences of the racial and ethnic segregation that pertains to large French cities (Préteceille, 2011; Safi, 2009). These findings echo those of the American literature, which unambiguously shows that there is a pervasive income and racial gap in exposure to various pollutants (see the reviews of Banzhaf et al., 2019; Mohai et al., 2009). They can be explained either by selective neighbourhood sorting (Tiebout, 1956), selective facility siting, or a combination of both. One limitation of this study is that IRIS-level data does not allow to study relative mobility patterns so as to test the sorting hypothesis. Indeed, aggregate changes in the number of inhabitants of a census block are very often uncorrelated with changes in pollution levels, even in the presence of Tiebout sorting, be it due to general equilibrium effects, stickiness in neighbourhood composition or multiplier effects (Banzhaf et al., 2019). Depro et al. (2015) show that individual sorting is not even identified from non-micro data, since one does not observe substitution behaviours. Consequently, further research relying on individual-level data will provide a better understanding of the mechanisms underlying the phenomena that this study brought to light.

Throughout the study period, average exposure to fine particulate matter decreased by 33%. Following Voorheis (2017), pollution-reduction profiles allow to depict that this improvement in air quality was regressive in terms of vertical equity, meaning that initially less exposed neighbourhoods benefitted from larger relative gains in air quality than their more exposed counterparts. Moreover, neighbourhoods located in the middle 80% of the initial distribution of income, i.e., whose median income ranged from €14,300 to €24,000 in 2006 (in 2016 euros), received higher air quality improvements than the bottom 10% and top 10%, which, as previously shown, are those that are initially most exposed to pollution. Hence, as deduced from IRIS-level data, although the overall exposure fell during the study period, inequality in $PM_{2.5}$ exposure intensified. This calls for follow-up studies in later years, as well as an adapted policy response in order to halt this trend.

This is of particular importance since it appears that a policy implemented during the study period contributed to these unequal changes. Indeed, following the 2008 EU Directive for ambient air quality and cleaner air for Europe, a number of urban areas that concentrate 43% of the metropolitan French population revised their existing air quality schemes, called *Plans de Protection de l'Atmosphère* (PPA). The adjusted plans had to newly include measures aimed at reducing fine particulate matter exposure, and were adopted between 2012 and 2016. The event-study design allows me to estimate the causal impact of this new policy change, since necessary assumptions are very likely to hold, as confirmed by an examination of the trend in exposure in the years leading up to the event. Results are particularly subject to residual autocorrelation, since both the significance and the magnitude of the baseline effect change substantially before and after taking spatial autocorrelation into account. According to preferred specifications, which control for neighbourhood location as a potential predictor of $PM_{2.5}$ concentration, the latter dropped by 5% to 7% in the years following the adoption of revised PPAs, which indicates that the measure played a role in the overall air quality improvements over the period. Results suggest that, regardless of whether one looks at the income distribution at the national level or within PPA zones, initially advantaged neighbourhoods perceived additional gains from this policy change, estimated around 25% to 40% higher than the average. Quantile regression estimates suggest that only lower quantiles of exposure benefitted from this decrease, while the effect may have been quite heterogeneous at upper quantiles. Hence, in the same vein as the country-level trends inferred using pollution-reduction profiles, it appears that the air quality improvements attributable to the revision of Atmosphere Protection Plans reduced overall exposure, but likely have exacerbated inequality in exposure. Nonetheless, I must emphasise that in order to verify this hypothesis, it would be worth conducting this analysis using data on a larger time span. Indeed, while the revision of PPAs occurred between 2012 and 2016, data availability restricts the study period up to 2016. Observing later years would also allow to examine potential indirect impacts of these policies, e.g., in the form of in-migration to neighbourhoods that were cleaned, which would be in line with the neighbourhood sorting hypothesis. Finally, it would be worth drawing a clear typology of the measures implemented as part of these air quality schemes, so as to evaluate their relative effectiveness, and thus provide recommendations to local governmental bodies.

References

Airparif, the Observatory of Air Quality in the Paris region (2011). Source apportionment of airborne particles in the Île-de-France region. [Available online at: https://www.airparif.asso.fr/_pdf/publications/rapport-particules-anglais-120829.pdf].

Aerts, A.-T., S. Chirazi, and L. Cros (2015). Une pauvreté très présente dans les villes-centres des grands pôles urbains. *INSEE Première 1952*, 1–4.

Anselin, L. (2002). Under the hood: Issues in the specification and interpretation of spatial regression models. *Agricultural Economics 27*(3), 247–267.

Argouarc'h, J. and S. Picard (2018). Les niveaux de vie en 2016: La prime d'activité soutient l'évolution du niveau de vie des plus modestes. *INSEE Première 1752*, 1–4.

Banzhaf, H. and E. McCormick (2012). Moving beyond cleanup: Identifying the crucibles of environmental gentrification. In *The Political Economy of Environmental Justice*, pp. 23–51. Stanford University Press.

Banzhaf, S., L. Ma, and C. Timmins (2019). Environmental justice: The economics of race, place, and pollution. *Journal of Economic Perspectives 33*(1), 185–208.

Basile, R. and A. Girardi (2010). Specialization and risk sharing in European regions. *Journal of Economic Geography 10*(5), 645–659.

Beaubrun-Diant, K. and T.-P. Maury (2020). Income Segregation and Social Housing in France. HAL Working Paper, No. 02526776.

Bell, M. L. and K. Ebisu (2012). Environmental inequality in exposures to airborne particulate matter components in the United States. *Environmental Health Perspectives 120*(12), 1699–1704.

Bertin, M., C. Chevrier, T. Serrano, C. Monfort, F. Rouget, S. Cordier, and J.-F. Viel (2015). Association between prenatal exposure to traffic-related air pollution and preterm birth in the PELAGIE mother–child cohort, Brittany, France. Does the urban–rural context matter? *Environmental Research 142*, 17–24.

Boyce, J. K., K. Zwickl, and M. Ash (2016). Measuring environmental inequality. *Ecological Economics 124*, 114–123.

Brochu, P. J., J. D. Yanosky, C. J. Paciorek, J. Schwartz, J. T. Chen, R. F. Herrick, and H. H. Suh (2011). Particulate air pollution and socioeconomic position in rural and urban areas of the Northeastern United States. *American Journal of Public Health 101*(S1), S224–S230.

Chang, T. Y., J. Graff Zivin, T. Gross, and M. Neidell (2016). Particulate pollution and the productivity of pear packers. *American Economic Journal: Economic Policy 8*(3), 141–69.

Chang, T. Y., J. Graff Zivin, T. Gross, and M. Neidell (2019). The effect of pollution on worker productivity: Evidence from call center workers in China. *American Economic Journal: Applied Economics 11*(1), 151–72.

Chavis, B. F. and C. Lee (1987). *Toxic wastes and race in the United States: A national report on the racial and socio-economic characteristics of communities with hazardous waste sites*. Commission for Racial Justice – United Church of Christ.

Citepa (2018). Gaz à effet de serre et polluants atmosphériques: Bilan des émissions en France de 1990 à 2017.

Cohen, A. J., M. Brauer, R. Burnett, H. R. Anderson, J. Frostad, K. Estep, K. Balakrishnan, B. Brunekreef, L. Dandona, R. Dandona, et al. (2017). Estimates and 25-year trends of the global burden of disease attributable to ambient air pollution: An analysis of data from the Global Burden of Diseases Study 2015. *The Lancet 389*(10082), 1907–1918.

Currie, J., J. Voorheis, and R. Walker (2020). What caused racial disparities in particulate exposure to fall? New evidence from the Clean Air Act and satellite-based measures of air quality. NBER Working Paper, No. 26659.

Depro, B., C. Timmins, and M. O'Neil (2015). White flight and coming to the nuisance: Can residential mobility explain environmental injustice? *Journal of the Association of Environmental and Resource Economists 2*(3), 439–468.

Dong, Y. and S. Shen (2018). Testing for rank invariance or similarity in program evaluation. *Review of Economics and Statistics 100*(1), 78–85.

Downey, L. (2007). US metropolitan-area variation in environmental inequality outcomes. *Urban Studies 44*(5-6), 953–977.

Ebenstein, A., V. Lavy, and S. Roth (2016). The long-run economic consequences of high-stakes examinations: Evidence from transitory variation in pollution. *American Economic Journal: Applied Economics 8*(4), 36–65.

Floch, J.-M. (2014). Des revenus élevés et en plus forte hausse dans les couronnes des grandes aires urbaines. In *France, portrait social*, pp. 69–81. INSEE.

Floch, J.-M. (2017). Standards of living and segregation in twelve French metropolises. *Économie et Statistique 497*(1), 73–96.

Frandsen, B. R. and L. J. Lefgren (2018). Testing rank similarity. *Review of Economics and Statistics 100*(1), 86–91.

Gamper-Rabindran, S. and C. Timmins (2011). Hazardous waste cleanup, neighborhood gentrification, and environmental justice: Evidence from restricted access census block data. *American Economic Review 101*(3), 620–24.

Gobillon, L. and H. Selod (2007). Les déterminants locaux du chômage en région parisienne. *Economie et Prévision 4*, 19–38.

Graff Zivin, J. and M. Neidell (2012). The impact of pollution on worker productivity. *American Economic Review 102*(7), 3652–73.

Guastella, G. and F. G. Van Oort (2015). Regional heterogeneity and interregional research spillovers in European innovation: Modelling and policy implications. *Regional Studies 49*(11), 1772–1787.

Hajat, A., C. Hsia, and M. S. O'Neill (2015). Socioeconomic disparities and air pollution exposure: A global review. *Current Environmental Health Reports 2*(4), 440–450.

Hastie, T. J. and R. J. Tibshirani (1990). *Generalized additive models.* CRC Press.

Havard, S., S. Deguen, D. Zmirou-Navier, C. Schillinger, and D. Bard (2009). Traffic-related air pollution and socioeconomic status: A spatial autocorrelation study to assess environmental equity on a small-area scale. *Epidemiology 20*(2), 223–230.

Havard, S., G. Pédrono, S. Deguen, C. Schillinger, C. Segala, E. Rivière, D. Arveiler, and D. Bard (2008). Air pollution and myocardial infarction. A small-area case crossover in Strasbourg, France: Influence of individual and area characteristics. *Epidemiology 19*(6), S197.

Helbich, M., W. Brunauer, E. Vaz, and P. Nijkamp (2014). Spatial heterogeneity in hedonic house price models: The case of Austria. *Urban Studies 51*(2), 390–411.

Jenkins, S. P. and P. Van Kerm (2006). Trends in income inequality, pro-poor income growth, and income mobility. *Oxford Economic Papers 58*(3), 531–548.

Kravitz-Wirtz, N., K. Crowder, A. Hajat, and V. Sass (2016). The long-term dynamics of racial/ethnic inequality in neighborhood air pollution exposure, 1990-2009. *Du Bois Review: Social Science Research on Race 13*(2), 237–259.

Lam, C. and P. C. Souza (2014). Regularization for spatial panel time series using the adaptive lasso. LSE STICERD - Econometrics Paper Series.

Laurent, É. (2009). Écologie et inégalités. *Revue de l'OFCE 2*, 33–57.

Laurian, L. and R. Funderburg (2014). Environmental justice in France? A spatio-temporal analysis of incinerator location. *Journal of Environmental Planning and Management 57*(3), 424–446.

Lavaine, E. (2015). An econometric analysis of atmospheric pollution, environmental disparities and mortality rates. *Environmental and Resource Economics 60*(2), 215–242.

Lavaine, E. (2019). Environmental risk and differentiated housing values: Evidence from the north of France. *Journal of Housing Economics 44*, 74–87.

Le Boennec, R. and F. Salladarré (2017). The impact of air pollution and noise on the real estate market. The case of the 2013 European Green Capital: Nantes, France. *Ecological Economics 138*, 82–89.

Le Moullec, A. (2018). Bilan de la qualité de l'air extérieur en France en 2017. Technical report, Ministère de la Transition Écologique et Solidaire. [Available online at: `https://www.statistiques.developpement-durable.gouv.fr/sites/default/files/2018-10/datalab-45-bilan-qualite-air-exterieur-france-2017-octobre2018.pdf`].

Medina, S., M. Pascal, and C. Tillier (2016). Impacts de l'exposition chronique aux particules fines sur la mortalité en France continentale et analyse des gains en santé de plusieurs scénarios de réduction de la pollution atmosphérique. Santé publique France.

Mohai, P., D. Pellow, and J. T. Roberts (2009). Environmental justice. *Annual Review of Environment and Resources 34*, 405–430.

Mohai, P. and R. Saha (2015). Which came first, people or pollution? A review of theory and evidence from longitudinal environmental justice studies. *Environmental Research Letters 10*(12), 125011.

Moran, P. A. (1950). Notes on continuous stochastic phenomena. *Biometrika 37*(1/2), 17–23.

Morelli, X., C. Rieux, J. Cyrys, B. Forsberg, and R. Slama (2016). Air pollution, health and social deprivation: A fine-scale risk assessment. *Environmental Research 147*, 59–70.

Muller, N. Z., P. H. Matthews, and V. Wiltshire-Gordon (2018). The distribution of income is worse than you think: Including pollution impacts into measures of income inequality. *PloS One 13*(3).

Occelli, F., R. Bavdek, A. Deram, A.-P. Hellequin, M.-A. Cuny, I. Zwarterook, and D. Cuny (2016). Using lichen biomonitoring to assess environmental justice at a neighbourhood level in an industrial area of Northern France. *Ecological Indicators 60*, 781–788.

OECD (2018). Exposure to $PM_{2.5}$ fine particles – Countries and regions. [Available online at: `https://www.oecd-ilibrary.org/content/data/ba4edead-fr`].

Ouidir, M., J. Lepeule, V. Siroux, L. Malherbe, F. Meleux, E. Rivière, L. Launay, C. Zaros, M. Cheminat, M.-A. Charles, et al. (2017). Is atmospheric pollution exposure during pregnancy associated with individual and contextual characteristics? A nationwide study in France. *Journal of Epidemiology & Community Health 71*(10), 1026–1036.

Padilla, C. M., W. Kihal-Talantikite, V. M. Vieira, P. Rossello, G. Le Nir, D. Zmirou-Navier, and S. Deguen (2014). Air quality and social deprivation in 4 French metropolitan areas: A localized spatio-temporal environmental inequality analysis. *Environmental Research 134*, 315–324.

Pastor, M., J. Sadd, and J. Hipp (2001). Which came first? Toxic facilities, minority move-in, and environmental justice. *Journal of Urban Affairs 23*(1), 1–21.

Pornet, C., C. Delpierre, O. Dejardin, P. Grosclaude, L. Launay, L. Guittet, T. Lang, and G. Launoy (2012). Construction of an adaptable European transnational ecological deprivation index: The French version. *Journal of Epidemiology & Community Health 66*(11), 982–989.

Préteceille, E. (2011). Has ethno-racial segregation increased in the greater Paris metropolitan area? *Revue française de sociologie 52*(5), 31–62.

Quillian, L. and H. Lagrange (2016). Socioeconomic segregation in large cities in France and the United States. *Demography 53*(4), 1051–1084.

Safi, M. (2009). The spatial dimension of integration: Developments in segregation of immigrant populations in France from 1968 to 1999. *Revue française de sociologie 50*(3), 521–552.

Tiebout, C. M. (1956). A pure theory of local expenditures. *Journal of Political Economy 64*(5), 416–424.

Tobler, W. R. (1970). A computer movie simulating urban growth in the Detroit region. *Economic Geography 46*(1), 234–240.

Van Donkelaar, A., R. V. Martin, M. Brauer, N. C. Hsu, R. A. Kahn, R. C. Levy, A. Lyapustin, A. M. Sayer, and D. M. Winker (2016). Global estimates of fine particulate matter using a combined geophysical-statistical method with information from satellites, models, and monitors. *Environmental Science & Technology 50*(7), 3762–3772.

Verbeek, T. (2019). Unequal residential exposure to air pollution and noise: A geospatial environmental justice analysis for Ghent, Belgium. *SSM - Population Health 7*.

Von Graevenitz, K. and T. E. Panduro (2015). An alternative to the standard spatial econometric approaches in hedonic house price models. *Land Economics 91*(2), 386–409.

Voorheis, J. (2017). Longitudinal Environmental Inequality and Environmental Gentrification: Who Gains From Cleaner Air? CARRA Working Paper, US Census Bureau, No. 2017-04.

Wood, S. N. (2017). *Generalized additive models: An introduction with R, Second Edition*. CRC Press.

Zwickl, K., M. Ash, and J. K. Boyce (2014). Regional variation in environmental inequality: Industrial air toxics exposure in US cities. *Ecological Economics 107*, 494–509.

Appendix

Figure 17: Exposure to $PM_{2.5}$ – ACAG and Santé publique France data

(a) ACAG (2008) (b) SPF (2007-2008)

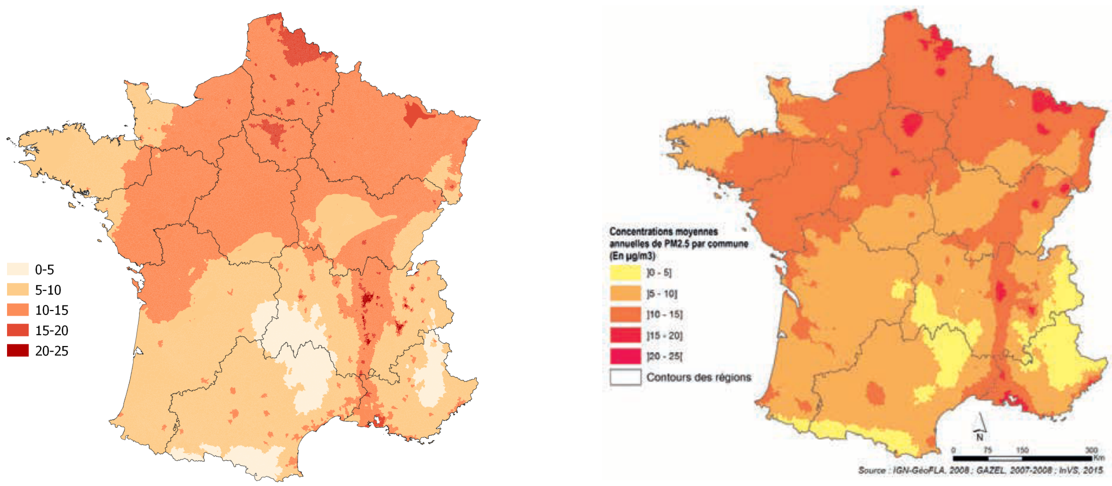

Sources: Atmospheric Composition Analysis Group and Santé publique France (Medina et al., 2016).

Figure 18: Evolution of exposure to $PM_{2.5}$ – ACAG and SDES data

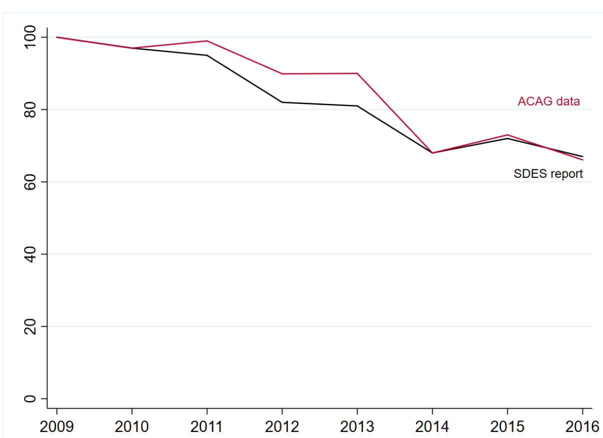

Sources: Atmospheric Composition Analysis Group and Le Moullec (2018).

Table 6: List of selected neighbourhood characteristics

Category	Variable
Unemployment	Share of unemployed in population aged 15-64
Origin	Share of immigrants
Education	Share of population with at least some higher education
Occupation	Share of population aged 15+ by occupation: 1) Farmers 2) Craftsmen, tradesmen, business managers (excl. licensed professionals) 3) White-collar jobs: executives and other intellectual professions 4) Intermediary: Technicians, foremen, school teachers, nurses, intermediary public servant occupations, and other intermediary occupations 5) Employees: Lower civil servant positions, policemen, military, intermediary administrative positions, service workers, and other employees 6) Blue-collar jobs: Industrial and craft workers, agricultural workers, drivers, and other blue-collar jobs 7) Retired (excluded to avoid multicollinearity) 8) Others without occupation: unemployed who never worked, students
Housing	Share of population that are homeowners Share of population that live in subsidised housing (HLM) Share of dwellings without electric heating
Households	Share of population that live in single-parent households Share of households that do not own a car
Location	Indicator variable of urban/rural

Table 7: Summary statistics of all variables – Whole sample, 2016

Variable	Min.	1^{st} quartile	Mean	Median	3^{rd} quartile	Max.	Std dev.
$PM_{2.5}$	4	6.53	8.59	7.97	10.58	15.1	2.64
Median income	2,210	18,218	20,897	20,252	22,819	70,674	5,108
% Unemployed	0	.06	.09	.08	.11	.38	.04
% Immigrants	0	.02	.07	.04	.08	.60	.07
% College-educated	0	.18	.26	.23	.31	.84	.12
% Farmers	0	0	.02	.004	.03	.47	.04
% Crafts-/tradesmen	0	.02	.04	.03	.05	.32	.03
% White-collar	0	.03	.07	.06	.10	.49	.07
% Intermediary	0	.10	.14	.14	.17	.52	.06
% Employees	0	.13	.16	.16	.19	.61	.05
% Blue-collar	0	.09	.14	.14	.18	.55	.07
% Retired	0	.22	.29	.28	.35	.88	.10
% Other inactive	0	.10	.14	.13	.17	.71	.07
% Homeowners	0	.63	.71	.78	.85	1	.20
% Social housing	0	0	.09	.02	.09	1	.17
% Trad. heating	0	.06	.23	.23	.35	.94	.17
% Single-parent	0	.05	.09	.09	.13	.5	.06
% No car	0	.05	.13	.08	.15	.91	.13

Figure 19: Empirical cumulative density of $PM_{2.5}$ exposure for selected income groups – 2016

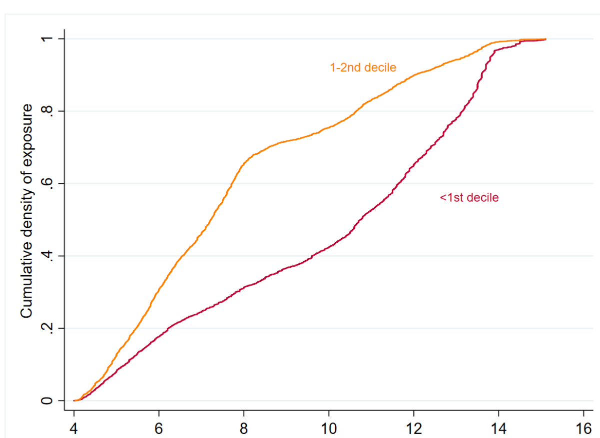

Table 8: Fixed-effect models – Full set of results

Variable	Dependent variable: log(PM)		
(Log) income	-1.274***	-.212***	-.182***
	(.137)	(.037)	(.030)
% Immigrants			.143**
			(.055)
% Unemployed			-.119
			(.105)
% College-educated			.141**
			(.039)
% No car			.337***
			(.105)
% No electric heating			-.124***
			(.042)
% Farmers			.383***
			(.105)
% Crafts-/tradesmen			.042
			(.045)
% White-collar			.258**
			(.105)
% Intermediary prof.			-.003
			(.019)
% Employees			.048**
			(.024)
% Blue-collar			.069
			(.053)
% Inactive excl. retired			.113**
			(.052)
% Single-parent			.014
			(.009)
% Social housing			.051
			(.053)
% Homeowners			.076
			(.051)
Intercept	15.062***	4.542***	4.061***
	(1.357)	(.368)	(.344)
Year fixed effects		X	X
R^2 within	0.11	0.78	0.74
R^2 between	0.02	0.01	0.22
R^2 overall	0.01	0.18	0.34
Observations	453,386	453,386	411,458
Groups	42,832	42,832	42,790

Standard errors clustered at the employment-zone level in parentheses. *: 10% level, **: 5% level, ***: 1% level.

Table 9: Fixed-effect generalised additive models – Full set of results

Variable	Dependent variable: log(PM)		
(Log) income	-.860***	-.191***	-.160***
	(.122)	(.031)	(.026)
% Immigrants			.132**
			(.051)
% Unemployed			-.090
			(.116)
% College-educated			.128**
			(.053)
% No car			.334***
			(.114)
% No electric heating			-.145***
			(.053)
% Farmers			.347***
			(.113)
% Crafts-/tradesmen			.016
			(.041)
% White-collar			.242**
			(.113)
% Intermediary prof.			-.009
			(.017)
% Employees			.028
			(.021)
% Blue-collar			.061
			(.057)
% Inactive excl. retired			.114**
			(.054)
% Single-parent			.022**
			(.010)
% Social housing			.051
			(.059)
% Homeowners			.089*
			(.052)
Intercept	10.992***	4.334***	3.830***
	(1.209)	(.312)	(.311)
Year fixed effects		X	X
R^2 within	0.05	0.73	0.74
R^2 between	0.03	0.01	0.22
R^2 overall	0.01	0.13	0.34
Observations	411,531	411,531	411,458
Groups	42,797	42,797	42,790

Standard errors clustered at the employment-zone level in parentheses. *: 10% level, **: 5% level, ***: 1% level.

Table 10: Years of adoption of analysed *Plans de Protection de l'Atmosphère*

Year	Zone	Number of IRIS	Number of municipalities
2012	*Bayonne*	*66*	*18*
	Bordeaux	277	52
	Dax	*31*	*20*
	Pau	*60*	*22*
	Vallée de l'Arve	*58*	*41*
		277 (*492*)	153
2013	Alpes-Maritimes du Sud	350	52
	Belfort (incl. Montbéliard)	*244*	*199*
	Bouches-du-Rhône	705	112
	Haute-Normandie	1,585	1,291
	Saint-Étienne	170	55
	Toulon	206	26
	Île-de-France	4,883	1,267
		7,899 (*8,143*)	3,002
2014	Avignon	86	22
	Clermont-Ferrand	88	22
	Dijon	*90*	*15*
	Grenoble	403	272
	Lyon	508	115
	Montpellier	211	114
	Nord-Pas-de-Calais	2,300	1,538
	Orléans	97	22
	Strasbourg	163	28
	Tours	138	40
		3,994 (*4,084*)	2,188
2015	*Bastia*	*26*	*12*
	Chalon-sur-Saône	*29*	*10*
	Nancy	106	38
	Nantes	260	58
	Reims	*96*	*16*
	Rennes	142	43
	Trois Vallées (incl. Metz, Thionville)	175	67
		683 (*834*)	244
		12,853 (13,553)	**5,234** (5,587)

PPA zones whose names are in italics are those that newly adopted an Atmosphere Protection Plan during the study period and are excluded from the study in Section 4.3. Totals in italics and parentheses include these urban areas.

Table 11: Number of observations by relative year of PPA adoption

Relative year	-5	-4	-3	-2	-1	0	1	2	3	4
# IRIS	13,360	13,360	13,360	13,360	13,360	13,360	13,360	12,526	8,442	492
# Obs. used	12,853	12,853	12,853	12,853	12,853	12,853	12,853	12,170	8,176	277

Figure 20: Cumulative distribution of $\ln(PM\mathrm{PM}_{2.5})$ at $t = -1$ and $t = 1$

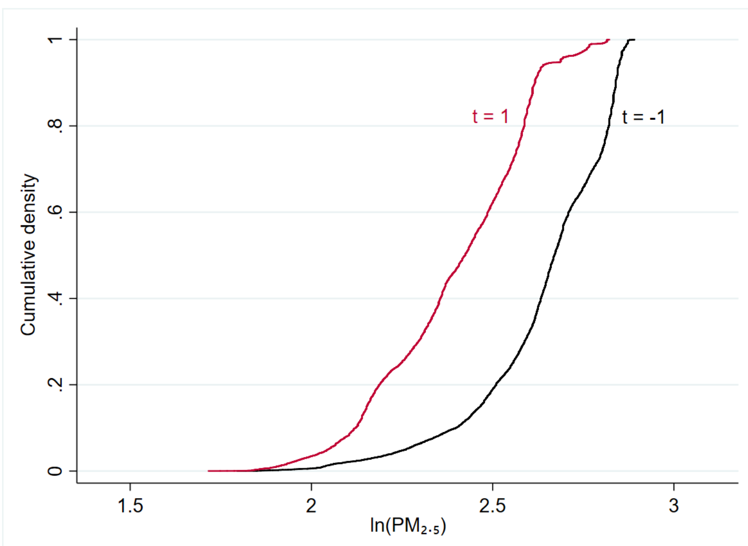

Table 12: Event study results – Pre-trends of μ_j and β_j

Variable	(1) Baseline	(2) Continuous	(3) 50/50 national	(4) 50/50 PPA zone
$t = -5$	-.867	-6.477	-1.154	-1.352
	(1.434)	(3.847)	(1.494)	(1.413)
$t = -4$	-2.176	3.181	-1.817	-2.221
	(1.491)	(7.035)	(1.388)	(1.336)
$t = -3$	-.790	8.068	-.285	-.765
	(1.647)	(9.032)	(1.765)	(1.603)
$t = -2$	-.682	7.592	-.472	-.881
	(.702)	(5.598)	(.766)	(.679)
$(t = -5) \times$ income var.$_i$		.144	.681*	1.022**
		(.108)	(.376)	(.312)
$(t = -4) \times$ income var.$_i$		-.108	-.538	.117
		(.157)	(.401)	(.362)
$(t = -3) \times$ income var.$_i$		-.148	-.749	-.032
		(.148)	(.551)	(.247)
$(t = -2) \times$ income var.$_i$		-.119	-.273	.414
		(.082)	(.659)	(.321)
Initial pollution	X	X	X	X
Neigh. charac.	X	X	X	X
PPA zone FE	X	X	X	X
Year FE	X	X	X	X
# Obs.	102,704	102,704	102,704	102,704
# Groups	12,853	12,853	12,853	12,853
# Clusters	20	20	20	20
R^2	0.96	0.97	0.97	0.97

Coefficients multiplied by 100 for the sake of readability. Standard errors clustered at the PPA zone level displayed in parentheses. *: 10% level, **: 5% level, ***: 1% level. Column (1) gives the results of equation (2), Column (2) the results of equation (3), and Column (3) and (4) the results of equation (4).

Figure 21: Relative effect of PPA adoption on average exposure – ADV_i in 4 categories

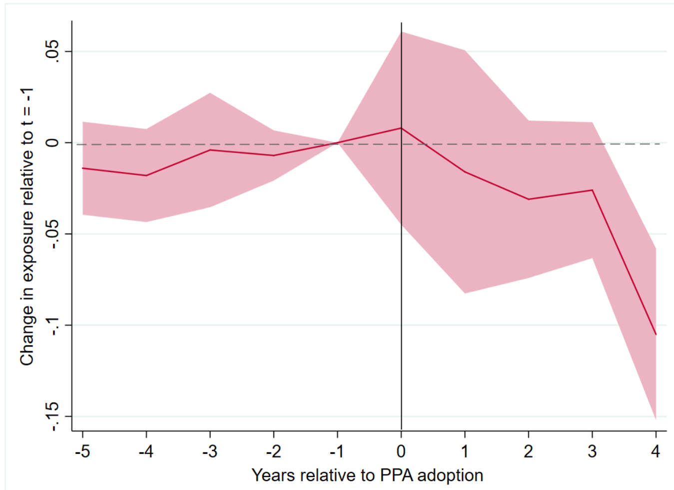

Note: Estimates of coefficient μ_j in equation (4) taking ADV_i as a 4-category variable. The shaded area corresponds to the 95% confidence interval.

Figure 22: Quantile event-study effects – Focus on $t = 2$

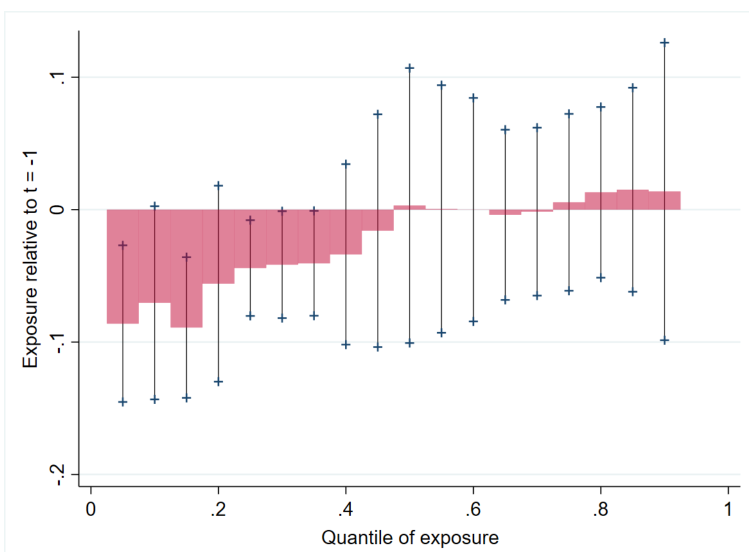

Figure 23: Quantile event-study effects – Focus on $t = 2$, based on initial income

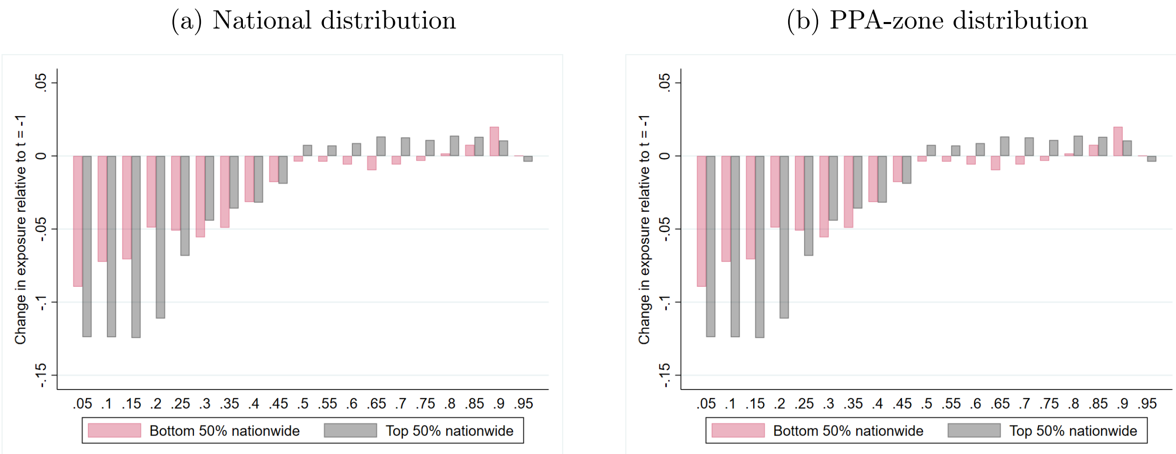

Note: Confidence intervals were omitted for the sake of visual clarity. Estimates are not statistically significant above the 4th decile, and are not significantly different from each other within the same vigintile when I distinguish between income groups.

Table 13: Event-study – Results of Moran's I test for residual spatial autocorrelation

	Baseline		50/50 national		50/50 PPA zone	
	(A)	(B)	(A)	(B)	(A)	(B)
Observed	.179	.215	.172	.207	.156	.198
Expected	-.012	-.012	-.012	-.004	-.004	-.004
p-value	<.001	<.001	<.001	<.001	<.001	<.001

Note: (A) columns give results for the area that encompasses the PPA zones of Avignon, Bouches-du-Rhône and Toulon, and (B) columns for the area that encompasses the PPA zones of Haute-Normandie and Île-de-France.