Discussion Paper Series – CRC TR 224

Discussion Paper No. 650
Project B 07

The Healthcare Costs of Air Pollution in France

Julia Mink[1]

February 2025

[1]Institute for Applied Microeconomics, University of Bonn, jmink@uni-bonn.de

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Collaborative Research Center Transregio 224 - www.crctr224.de
Rheinische Friedrich-Wilhelms-Universität Bonn - Universität Mannheim

The healthcare costs of air pollution in France

Julia Mink

[a] *University of Bonn, Institute for Applied Microeconomics, Department of Economics, Lennéstrasse 43, Bonn, 53113 NRW, Germany*

Abstract

In this paper, I estimate the causal impact of short-term exposure to nitrogen dioxide (NO_2), ground-level ozone (O_3), and particulate matter (PM) on healthcare costs in France. I construct a large-scale dataset by linking administrative healthcare expenditures for a nationally representative sample with high-resolution air pollution and meteorological data. To address endogeneity concerns related to economic activity, I implement an instrumental variable (IV) strategy that exploits weekly variations in altitude atmospheric conditions—such as thermal inversions, wind speed, and the height of the planetary boundary layer—that predict local pollutant concentrations yet are unlikely to affect healthcare utilization except through pollution. My findings reveal that air pollution, even at concentrations below current European air quality standards, imposes annual healthcare costs that exceed earlier estimates by a factor of ten. Heterogeneity analyses show that pollution affects multiple medical specialties, including cardiology, pulmonology, and ophthalmology, while placebo specialties, such as trauma surgery, exhibit no significant effects. Contrary to prior work focusing on children and the elderly, I find that adverse health outcomes extend across all age groups, demonstrating broader population vulnerability. Moreover, marginal effects prove larger at lower pollution levels, implying a concave dose-response function that underscores the potential for substantial cost savings from even modest pollution abatement in relatively clean areas. These results suggest that earlier cost-benefit analyses likely undervalue the societal gains from stricter environmental regulation.

Keywords: Air pollution, healthcare cost, instrumental variables

---

1. Introduction

Air pollution is widely recognized as the most significant environmental risk to public health in Europe. Particulate matter (PM), nitrogen dioxide (NO_2), and ground-level ozone (O_3) are of particular concern (EEA, 2020). Exposure to PM alone is estimated to contribute to approximately 400,000 premature deaths each year in Europe, while NO_2 and O_3 exposures account for around 70,000 and 15,000 additional premature deaths, respectively, in 2017 (Maguire et al., 2020). Even in France, where pollution levels tend to be relatively low, air pollution remains among the top ten risk factors for both mortality and disease burden (Institute for Health Metrics and Evaluation, 2020). Numerous studies have linked exposure to air pollution with adverse health outcomes, including Chronic Obstructive Pulmonary Disease (COPD), respiratory distress (e.g., coughing, wheezing, and shortness of breath), asthma, and increases in hospital admissions for respiratory and cardiovascular diseases. For a review, see Manisalidis et al. (2020).

In response, air quality standards and target concentrations have been established for several pollutants, yet the stringency of these limits remains controversial. At present, the EU air quality guidelines are being revised to align more closely with the World Health Organization's (WHO) stricter guidelines, which themselves were updated only recently in 2021 (European Commission, 2016). Accurate estimates of the health impacts of air pollution are critical for policy makers seeking to determine the optimal level of environmental regulation. This is especially relevant when pollution levels are already relatively low, and the marginal benefits of additional pollution reduction may not obviously outweigh the costs. In this study, I quantify the healthcare costs imposed on the French healthcare system by acute (short-term) exposure to air pollution — a context in which average pollution levels remain far below current EU air quality standards.

Estimating the causal impact of air pollution on healthcare costs is complicated by the non-randomness of exposure. Individuals tend to select their residential locations based on preferences and characteristics associated with health status, which may also correlate with pollution levels, thus biasing the estimates. To address this concern, many previous studies have used quasi-experimental designs. However, data or methodological limitations often confine these studies to specific geographic areas or time periods, or narrow populations and health outcomes. Moreover, much of the existing literature focuses on mortality, a relatively rare event at moderate pollution levels; in contrast, healthcare costs may accrue even at lower concentrations. To overcome these limitations, this study combines comprehensive French administrative data on healthcare expenditure with high-resolution geospatial data on air pollution and meteorological conditions.

Specifically, I examine the causal effects of exposure to NO_2, O_3, and PM on healthcare costs from 2015 to 2018. These costs encompass all types of health services and related costs such as pharmaceutical expenditures and all types of medical specialties. To identify causal effects, I use a location fixed-effects model exploiting weekly variation in air pollution concentrations at the French zip code level. While location fixed effects account for time-invariant, location-specific population characteristics, they cannot fully address potential endogeneity arising from correlations between pollution and economic activity. Therefore, I also implement an instrumental variables (IV) approach, where altitude weather conditions serve as instruments for local pollution concentrations. After flexibly controlling for time and location fixed effects, and ground-level weather, I assume that changes in air pollution induced by altitude weather conditions are orthogonal to any changes in healthcare use or costs. These high-dimensional altitude weather instruments strongly predict air pollution variability across France and allow simultaneous identification of effects for multiple pollutants. I conduct robustness checks using alternative specifications: different sets of instruments, varying fixed-effects structures, additional ground-level weather controls, lagged pollution effects, and extended windows to capture pollution build-up over prior weeks. Furthermore, I investigate heterogeneity in the estimated effects by medical specialty, patient characteristics (e.g., age, chronic disease status, insurance status), and location characteristics (e.g., average income, population size, baseline pollution levels).

The results reveal that exposure to pollution concentrations predominantly below current European air quality standards imposes healthcare costs on the French healthcare system of several billion euros per year. Specifically, air pollution increases annual healthcare spending by at least €12.8 billion, or about 0.5% of France's 2019 GDP and 6.2% of total 2019 healthcare spending.[1] These estimates are approximately ten times larger than those reported in previous studies, suggesting that earlier assessments have substantially underestimated the health-related costs of air pollution. Although my estimates exceed prior findings, they should still be viewed as a lower bound, as they capture only short-term effects and do not include costs related to mortality or loss of productivity. Additional analyses confirm that pollution affects healthcare costs through the expected clinical pathways. Significant effects emerge in cardiology and vascular medicine, pulmonology, otolaryngology (O.R.L.), ophthalmology, gynaecology, and family practice. By contrast, no effects are detected in placebo specialties such as trauma and plastic surgery. Age-based analyses show that pollution affects healthcare costs across all age groups, diverging from prior studies that primarily highlight impacts on children and the elderly. This discrepancy may arise because most previous work focuses on mortality, while the present study investigates milder (but more widespread) health effects. These findings suggest that health problems extend beyond groups typically considered at highest risk.

[1] In 2019, France's GDP was €2,425.7 billion (INSEE, 2020), and aggregate healthcare expenditure was €208 billion (DREES, 2020).

Regarding socioeconomic heterogeneity, the analyses do not reveal any clear pattern of differential impacts by zip-code average income or by low-income status as proxied by enrolment status in the state funded complementary insurance plan available to low-income individuals. These findings are consistent with structural features of the French healthcare system, which ensures affordability through high reimbursement rates, subsidized or fully covered complementary insurance, and a nationally regulated fee schedule for most medical services. Finally, although total expenditures associated with air pollution are higher in urban areas (where populations are larger and baseline pollution is somewhat elevated), the marginal impact of pollution is greater in less populated regions with relatively low baseline concentrations. This concave relationship suggests that small pollution shocks may have disproportionately large health effects when levels are already low, echoing evidence from more recent studies on the non-linear effects of pollution.

This study contributes to the literature on measuring the health costs of air pollution for cost-benefit analyses and to the quasi-experimental literature in economics assessing the causal effects of air pollution on health. The evaluation of the healthcare costs caused by air pollution has so far been comparatively incomplete. Studies that seek to evaluate the health costs of air pollution for cost-benefit analysis often only include a selection of health effects and part of the population for which epidemiological evidence is most robust. Taking a policy relevant example from France, a 2015 Senate Committee of Inquiry into the economic and financial cost of air pollution (Sénat, 2015) searched for estimates of the total costs of air pollution to the French healthcare system to inform policy decisions. The result was a report on two studies that considered only asthma and cancer (Fontaine et al., 2007) or respiratory diseases and cancers, and hospitalisations for respiratory and cardiovascular causes in Rafenberg (2015). Quasi-experimental studies are similarly limited in scope (see for example Neidell (2004); Currie and Neidell (2005); Jayachandran (2009); Neidell (2009); Moretti and Neidell (2011); Currie and Walker (2011); Chen et al. (2013); Schlenker and Walker (2016); Knittel et al. (2016); Schwartz et al. (2017); Deschênes et al. (2017); Simeonova et al. (2019); Halliday et al. (2019); Colmer et al. (2021); Guidetti et al. (2024); Klauber et al. (2024)). Many focus on mortality or particular health outcomes, and most do not consider expenditures on drugs (see Rohlf et al. (2020) for a notable exception). The quasi-experimental studies that are probably the most comparable to the present study in terms of data quality and empirical strategy are Deryugina et al. (2019) and Barwick et al. (2024). However, Deryugina et al. (2019) is limited to analysing hospital costs and focuses on the impact on the elderly population in the USA while Barwick et al. (2024) is limited to using transaction data covering half of private healthcare spending in China in 2015 without containing information on the specific diagnoses or treatment associated with these transactions.

This study advances the existing literature in three key ways. First, the study provides a more accurate and comprehensive assessment of the health costs of acute (short-term) air pollution exposure by using a nationally representative sample including detailed and exhaustive information on all healthcare expenditures. While existing studies acknowledge that their estimates of short-term health costs are conservative, the degree of underestimation remains unclear. Although the current analysis still captures only a lower bound on the *total* health costs of air pollution —excluding costs from chronic exposure or mortality — it represents a substantially more complete lower-bound estimate than any previous similar studies. Second, the study analyses treatment effect heterogeneity across multiple patient and location characteristics. This approach sheds light on inequalities and non-linear effects that are difficult to capture in earlier research, which often relies on less comprehensive data or smaller geographic coverage. Third, this study considers the effects of multiple pollutants simultaneously using an instrumental variable approach based on a high-dimensional vector of instruments. Few studies have attempted to disentangle the effects of multiple pollutants (Deryugina et al., 2019; Godzinski and Castillo, 2021), as is challenging due to the high correlation between the pollutants.

From a policy perspective, these findings demonstrate that the health costs imposed by air pollution in France are far higher than suggested by earlier estimates. Because a large portion of these costs occurs even at pollutant concentrations below existing European air quality guidelines, the results imply that current standards may not adequately protect public health. Indeed, the concave dose-response pattern indicates that health benefits from pollution reduction are disproportionately large at low levels of pollution, further underscoring that even the most stringent WHO guideline values from 2021 should not be viewed as definitively "safe." These insights have implications for policy decisions concerning cost-benefit calculations, the setting of guideline values, and the dissemination of public health information. In particular, warning systems should emphasize that harmful health effects occur well below regulatory thresholds and affect not only traditionally vulnerable groups but the broader population as well.

The rest of the paper is structured as follows. Section 2 provides an overview of air pollution and air quality in France, and discusses altitude atmospheric conditions as instrumental variables for ground-level pollution. Section 3 presents the data sources, Section 4 outlines the empirical strategy, Section 5 reports the results, and Section 6 concludes.

2. Background

This section provides background information on air pollution and air quality in France and contains a discussion on the adequacy of altitude atmospheric conditions as instruments for air pollution concentrations.

2.1. Air pollution and air quality in France

In this study, I focus on the air pollutants nitrogen dioxide (NO_2), ozone (O_3), particulate matter 10 micrometers or less in diameter (PM_{10}). While there are many other potentially hazardous air pollutants, these air pollutants are considered key pollutants of major public health concern and have long been the focus of international and national air quality standards (EEA, 2013). For several decades, the European Union (EU) has had air quality standards in place for these pollutants in the ambient air quality directives. Current limit values are a yearly average of $40\mu g/m^3$ for NO_2, a maximum daily 8-hour mean of $120\mu g/m^3$ for O_3 not to be exceeded more than 25 days per year, a yearly mean of $40\mu g/m^3$ for PM_{10} and $25\mu g/m^3$ for $PM_{2.5}$. Although these values were based on the 2005 World Health Organisation (WHO) air quality guidelines, the EU air quality standards are less demanding than these guidelines and much less stringent compared to the most recent 2021 WHO guidelines. Table OA1 in the appendix presents the 2005 and 2021 WHO air quality guideline values and the current EU air quality standards applicable in France.

Air quality in France improved over the last two decades with the exception of O_3 pollution (Le Moullec and Meleux, 2019). I observe an average NO_2 concentration over the years 2015 to 2018 of $13.8\mu g/m^3$ and average PM_{10} and $PM_{2.5}$ concentrations are $16.61\mu g/m^3$ and $10.58\mu g/m^3$, respectively. Figure A1 shows distributions of daily mean and maximum hourly pollution levels relative to the current EU and the 2005 and 2021 WHO limit values. The EU limit values are respected on most days and in most zip code locations.

There are strong correlations between the air pollutants. NO_2 and particulate matter tend to be positively correlated as they share common sources. Nitrogen oxides (NO_x), which include nitrogen monoxide (NO) and NO_2, are emitted during the combustion of fuels from industrial plants and road traffic and contribute to the formation of O_3 and PM. PM is either directly emitted as primary particles or it forms in the atmosphere from emissions of certain precursor pollutants such as sulfur dioxide (SO_2), NO_x, ammonia (NH_3) and volatile organic compounds (VOCs). O_3 is not directly emitted into the atmosphere. It is a secondary pollutant formed from chemical reactions in the presence of sunlight, following emissions of precursor gases, mainly NO_x, carbon monoxide (CO), VOCs and methane (CH_4). The processes of O_3 formation and accumulation are complex. To put it simply, in the short-term,

NO_2 tends to be inversely related to O_3 because in many settings NO_2 disappears during the formation of O_3 and vice-versa(Lee et al., 2021; EPA, 2024; Clapp and Jenkin, 2001).

2.2. Altitude atmospheric conditions as instruments for air pollution

The quasi-experimental literature has many times relied on weather conditions as instrumental variables for pollution concentrations to address concerns of possible endogeneity bias related to the confounding effect of economic activity. The general assumption is that, conditional on ground-level weather conditions, altitude atmospheric conditions affect air pollutant concentrations on the ground while being unrelated to economic activity. I exploit a vector of instruments based on different altitude weather conditions, including thermal inversions, planetary boundary height, wind speed, and wind direction. Using a range of instruments allows to instrument for several pollutants simultaneously. As air pollutants are both highly correlated and considered to have independent effects on health, multi-pollutant approaches are regarded as desirable (Godzinski and Castillo, 2021; Mauderly et al., 2010; Vedal and Kaufman, 2011; Johns et al., 2012). Yet, challenges arise when implementing multi-pollutant approaches such as results of many regression models become unstable when incorporating more than one pollutant, and very often imprecise due to the correlation between the pollutants. Instrumenting several pollutants simultaneously can overcome this problem if different subsets of instruments better predict variation in some pollutants than in others, allowing to disentangle the effects the of different pollutants (Godzinski and Castillo, 2021; Deryugina et al., 2019).

Thermal inversions are a deviation from the normal monotonic relationship between air temperature and altitude. Under normal atmospheric conditions, warm air at the surface is drawn upwards due to its lower density. This atmospheric ventilation can help reduce air pollution at the surface. During a thermal inversion, a cooler air mass is trapped under a warm air mass, preventing normal atmospheric ventilation and trapping the polluted air at the surface. The large-scale movement of air masses in the atmosphere typically forms thermal inversions at their leading edge when warm air masses pass over cooler air masses. Thermal inversions also form when the sun heats the air in the higher parts of the atmosphere faster than the air on the ground. As variations in surface- and higher-level temperatures within a region are usually assumed to be exogenous, thermal inversions are assumed to be exogenous. However, inversions occur above ground level but are associated with weather that can potentially affect economic activity or health outcomes at ground level. To rule out a possible correlation between thermal inversions, economic conditions and health outcomes due to weather, I flexibly control for ground-level weather conditions in all regressions. Thermal inversions have often been used as an instrument for air pollution (see for example Arceo et al. (2016); Jans et al. (2018); Dechezleprêtre et al. (2019).

The planetary boundary layer is the part of the atmosphere that is directly and strongly influenced by the earth's surface. Pollutants are trapped in this vicinity of the Earth. The higher the planetary boundary layer, the greater the volume of air available for pollutants and the lower the concentration (Levi et al., 2020). Planetary boundary layer height (PBLH) responds to heating flux between the sun and the earth and can also change under unpredictable large-scale air movements. Similar to thermal inversions, PBLH is generally considered to vary exogenously, but it also has a seasonal nature and is partially related to ground-level weather. For the exclusion restriction to hold, seasonal and ground-level weather controls are included in all regressions. Although less often than thermal inversions or wind direction, PBLH has been used in the economic literature as instrument for air pollution (for example Godzinski and Castillo (2021); Schwartz et al. (2017, 2018)).

Wind characteristics are also directly influencing pollutant concentrations. While wind speed generates variation in pollution concentrations through the dispersion of locally produced pollutants, wind direction may affect pollution concentrations by bringing air composed of different pollutants from more or less distant sources, depending on the wind direction and the relative location of the pollution sources. Wind speed at altitude is correlated with ground level wind speed, which could affect health outcomes and thereby violate the exclusion restriction. Ground level wind speed is therefore included as control variable in all regressions. For wind direction, the exclusion restriction should apply without restriction. Changes in wind direction are likely to be exogenous to economic activity and should only affect health outcomes through their effect on pollution concentrations. However, the effect of wind speed on air pollution levels is location-dependent. The effect of this instrument must therefore be able to vary at the local level (see the discussion on the assumption of monotonicity of the instrument in section 4). Both wind direction and altitude wind speed have been used as instruments in the literature. For wind direction see for example Anderson (2015); Deryugina et al. (2019) and for altitude wind speed see for example Godzinski and Castillo (2021); Schwartz et al. (2017, 2018).

3. Data

I combine detailed administrative healthcare data for a representative sample of the French population with high-resolution geospatial data on pollutant concentrations and atmospheric conditions. I also use additional data on zip code characteristics, including income and population size, from tax and social benefit data sources.

The final dataset includes information on healthcare costs, concentrations of various air pollutants, weather conditions and location characteristics for the years 2015 to 2018 at the French zip code level, which typically represents an area of around 9 km $\times$ 9 km. Most of the data are available at a daily frequency, with the exception of tax and social benefit data, which are available at an annual frequency. For most analyses, I aggregate daily data to the week, but the results are qualitatively and quantitatively similar when analyses are performed with daily-frequency data.

3.1. Healthcare use and cost

I use data on healthcare costs from the French National System of Health Data (SNDS for *Système National des Données de Santé*) provided by the National Fund for Health Insurance (CNAM, 2015-2018). The French health care system is based on universal coverage by one of several health insurance plans. The SNDS database aggregates anonymous information on reimbursed claims from all these plans and is also linked to the national hospital discharge database system. The full SNDS database covers 98.8% of the French population, making it possibly the largest contiguous homogeneous benefits database in the world. The data base provides information on the nature of the medical acts, costs of treatment for all types of healthcare, including physician visits, drug purchases, and hospital care. The information is available by exact date of care.

Data on patient characteristics include patient age, sex, information on chronic health conditions, and zip code of residence. The administrative healthcare data base does not contain information about the patient's socioeconomic status. However, it does contain information about the patient's insurance coverage, including whether the individual benefits from subsidised or free supplementary health insurance which are available to individuals whose income does not exceed certain thresholds. The information on the nature of the complementary insurance is useful as a proxy measure of socioeconomic status. The income thresholds for these schemes - the Free Supplementary health insurance and the Supplementary Health Insurance with participation[2] - are around 90% and 70% of the poverty threshold, respectively.

[2] The Universal Complementary Health Coverage "*Couverture médicale universelle complémentaire*" (CMU-C), since november 2019 the Free Supplementary health insurance "*Complémentaire Santé Solidaire gratuite*" (C2SG) is a free supplementary health insurance for individuals with low income, covering healthcare expenses that are not reimbursed by the social security system, without any upfront payments. The Assistance with the Acquisition of Health Insurance "*Aide au paiement d'une Complémentaire Santé*" (ACS), since november 2019 the Supplementary Health Insurance with participation "*Complémentaire Santé Solidaire avec participation*" (C2SP) is a similar scheme, but with a small financial contribution from the

I use the general sample of beneficiaries (EGB for *"Echantillon Généraliste des Bénéficiaires"*) which is the 1/97th random permanent representative sample of the SNDS data. The EGB facilitates the implementation of longitudinal studies, as it provides panel data that makes it possible to track patients' use of healthcare. See Tuppin et al. (2010) and Bezin et al. (2017) for more information on the EGB. I use an extraction of this data with individual data aggregated to the zip code level of the patient's place of residence.

3.2. Air pollutant concentrations

I use reanalysis data on hourly concentrations of NO_2, O_3 and PM_{10} provided by the French National Institute for Industrial Environment and Risks (INERIS, 2015-2018). The data are made available in the form of high spatial resolution raster files with a cell size of approximately 4x4 km. I convert the hourly data into daily averages and overlay the raster data with a shapefile of France containing the administrative boundaries of the zip code areas to extract daily pollution levels by zip code area. For the main analyses, I then further aggregate the data to the week. Reanalysis data offers several advantages over data from measurement stations. Since the number of monitoring stations is limited and they are often only sparsely distributed in space, researchers usually have to interpolate data points for locations far away from the monitoring stations (see for example Currie and Neidell (2005); Knittel et al. (2016); Schlenker and Walker (2016)). The interpolation of pollution levels using simple distance weights, as is often done in the literature, neglects meteorological and geographical factors that influence the dispersion of pollution, which can lead to a discrepancy between the actual and assigned pollution levels, especially at locations further away from the monitoring stations. The reanalysis data from INERIS combines information from measurement stations with a climate model rather than using a statistical procedure to interpolate between observations to address this issue. For more information on the construction of the reanalysis data, see Real et al. (2022).

For sensitivity analyses I also use raw data from monitoring stations for NO_2, O_3 and PM_{10} and additionally SO_2 and CO concentrations provided by the European Environment Agency (EEA, 2015-2018). CO concentrations are recorded at 44 monitoring stations and SO_2 concentrations at 173 monitoring stations, which means that the geographical coverage is relatively sparse. The data on NO_2, O_3 and PM_{10} levels are collected at 475, 370 and 411 monitoring stations respectively, which means

individual, aimed at those whose income is slightly above the threshold for the free complementary health insurance, yet still modest (DSS, 2023).

that they are less sparse than the CO and SO_2 data, but still sparse compared to the 4x4 km grid of the INERIS reanalysis data. I convert the data from the monitoring stations into average pollutant concentrations in the zip code area by interpolating the pollution values using an inverse distance weighting, in which measurements that are geographically closer to the zip code area under consideration are weighted more strongly than measurements that are further away.

3.3. Atmospheric conditions

The data on atmospheric conditions comes from ERA5, the fifth generation of global climate and weather reanalysis produced by the Copernicus Climate Change Service (C3S) at the European Centre for Medium-Range Weather Forecasts (ECMWF). The ERA5 reanalysis combines model data with past observations from measuring stations to create a globally complete and consistent data set. The ERA5-Land hourly data provide information on atmospheric conditions at ground level, including the u and v components of wind, the height of the planetary boundary layer, temperature and precipitation with a spatial resolution of 0.1°x0.1° (ca. 9x9km). Atmospheric conditions at altitude can be retrieved for 27 pressure levels (altitude levels) and include the u- and v-component of the wind, the wind speed and the temperature with a resolution of 0.25°x0.25° (ca.22x22km). The data are freely available online at the Copernicus Climate Data Store (Muñoz Sabater (2019) for the ground-level data, Hersbach et al. (2023a) for the altitude data and Hersbach et al. (2023b) for the planetary boundary layer height). I overlay the ERA5 raster data with a shapefile of the administrative boundaries of the French zip code areas to extract the data at the zip code level. I encode the presence of a thermal inversion as a dummy variable equal to one if the temperature at the surface atmospheric layer (pressure level 1000 hPa) is lower than the temperature at the atmospheric layer just above (975 hPa). I construct the strength of the thermal inversion as the temperature difference between these two atmospheric layers. The planetary boundary layer height is provided in meters and used as a continuous variable without further transformation. Wind speed and wind direction are calculated from the u- and v-component of the wind, which are the eastward and northward component of the wind, respectively. Wind speed can be calculated as $WS = \sqrt{u^2 + v^2}$, where u and v denote the u- and v-component of the wind. I use wind speed as a continuous variable. Wind direction can be calculated as $\Phi = mod(180 + \frac{180}{\pi} atan2(v, u), 360)$ to get an answer in degrees in the range $0 \leq \Phi < 360.\Phi$ indicates the direction from which the wind is blowing. Zero means the wind is blowing from the north to the south. I use the average cardinal wind direction to construct wind direction bins. Relative humidity is calculated using the information on temperature T and dew point temperature T_d

according to the formula $RH = 100 \times \frac{e_d}{e_s}$ where $e_d = 6.11 \times 10^{\frac{7.5 \cdot T_d}{237.7+T_d}}$ (saturation vapour pressure at dew point temperature (T_d)) and $e_s = 6.11 \times 10^{\frac{7.5 \cdot T}{237.7+T}}$ (saturation vapour pressure at temperature (T)). Relative humidity is expressed in percentage.

I aggregate the hourly data to the week to construct the following ground-level meteorological variables: weekly sum of precipitation in millimetres, weekly average temperature and weekly average of daily minimum and maximum temperatures in degrees Celsius, and weekly average relative humidity in percent. I proceed in the same way to construct the following altitude atmospheric condition instrumental variables: the sum of the number of hours of thermal inversion per week, the weekly average strength of these thermal inversions in degree Celsius, the weekly average of the planetary boundary layer in meters, average wind speed at twelve pressure levels in meters per second, and wind direction bins based on weekly average cardinal wind direction.

In addition, I distinguish altitude atmospheric condition by moment of the day by constructing 4-hourly within-day averages (0 to 4 a.m., 4 to 8 a.m, 8 a.m. to 12 p.m., 12 p.m. to 4 p.m., 4 p.m. to 8 p.m and 8 p.m to 0 a.m.). I consider these six within-day averages to capture the likely differences in impact of the instruments depending on when the pollution emissions are produced. For example, a thermal inversion during the morning or evening traffic peaks should influence air pollution concentrations more than when it occurs at night (Godzinski and Castillo, 2021). I then aggregate these six within-day averages to the week, resulting in the following additional instrumental variables: the number of hours of thermal inversions per week by moment of the day, the average strength of these thermal inversions, the average of the planetary boundary height by moments of the day, the weekly average wind speed at twelve pressure levels by moment of the day, and wind direction bins based on average cardinal wind direction by moment of the day. For sensitivity analyses using daily variation, I aggregate the data in a similar way to the day.

3.4. Additional data

Data on French administrative boundaries at the zip code level are sourced from the Open Platform for French Public Data (data.gouv.fr, 2014). The data also include information on population size by zip-code.

Median household income at the zip code level, available at an annual frequency, is obtained from the FiLoSoFi database (*Fichier Localisé Social et Fiscal*). This database compiles administrative records on taxation and social benefits and is maintained by the French National Institute of Statistics and Economic Studies (INSEE, *Institut national de la statistique et des études économiques*). Aggregated data at the zip code level are publicly accessible via the INSEE website (INSEE, 2024).

Data on public holidays in France are also retrieved from the Open Platform for French Public Data (data.gouv.fr, 2024). From this dataset, I construct a control variable reflecting the weekly count of public holidays. This variable accounts for reduced economic activity during holiday periods, which is expected to influence both pollution levels and healthcare use.

3.5. Summary statistics

The final data set consists of $1,257,984$ zip code-week observations of healthcare costs across $6,048$ French zip code areas from 2015 to 2018. Missing information in the air pollution and weather databases leads to less than 3.5% missing values for any variable, with at least $1,247,428$ observations of each air pollutant concentration variable and at least $1,214,512$ observations of each weather variable. The average weekly healthcare expenditure per zip codes is €$3,609$ with a standard deviation of $7,471$. The mean concentration of NO_2 is $13.78\mu g/m^3$ (standard deviation 7.14); PM_{10} averages $16.61\mu g/m^3$ (sd 6.32) and O_3 averages $55.7\mu g/m^3$ (sd 17.49). These levels of pollutants are far below the air quality standards currently in force in France (see Table OA1 and the discussion in Section 2). Table A1 in the appendix presents summary statistics.

4. Empirical strategy

The aim of this study is to quantify the health costs caused by air pollution. Estimating causal effects presents several challenges. First, the estimates might be subject to bias due to residential sorting. People choose where they live and thus the extent of their exposure, which can lead to correlations between air pollution and personal characteristics, possibly including their health status. Not accounting for this non-random exposure may lead to biased estimates, with the direction of bias being theoretically unclear. For example, people with high socioeconomic status are on average healthier and can afford to live in areas with low air pollution, but they may also be more likely to live in highly polluted city centres because of their occupation or preferences. Second, estimates may be biased due to correlations between pollutant concentrations and health care utilization associated with economic activity. Increased economic activity is likely to be tied to increased pollution from sources such as transportation and manufacturing, and it is also likely to be tied to increased health care utilization due to work-related stress or injury or due to increased availability of health care during business hours.

4.1. Location and time fixed effects model

To address the issue of possible bias from spatial sorting, I estimate a location-fixed effect model that exploits week-to-week variation in air pollution concentrations within the same zip code area. The composition of the population in a given zip code area is plausibly stable from one week to the next, which means that the weekly variation in air pollution concentration within a zip code area is exogenous to the average location-specific population characteristics. I estimate the following model

$$H_{wp} = \sum_{x} \beta_x P_{wpx} + \alpha_p + \alpha_{m/mdep} + \alpha_y + \gamma X_{wp} + \epsilon_{wp}, \tag{1}$$

where H_{wp} denotes healthcare costs incurred in week w in zip code area p, α_p are zip code area fixed effects, P_{wpx} is the pollution concentration of pollutant x in week w in zip code area p, $\alpha_{m/mdep}$ and α_y are month or month-by-department[3] and year fixed effects, X_{pw} stands for a vector of time-varying location characteristics, and ϵ_{xdp} denotes the error term.

I construct the health care cost variable as the weekly sum of expenditure of all medical specialties and all types of health care, including physician visits, drug purchases, and hospital care. This is in contrast to existing studies, which focus on a limited number of health problems or on specific types of healthcare such as hospital admissions.

The inclusion of time fixed effects allow to flexibly control for seasonality in air pollution and healthcare use. The month-by-department fixed effects control for any seasonal correlation between pollution and healthcare use that could vary across the 96 French departments. The vector of time-varying location characteristics X_{pw} includes control variables for ground-level weather conditions. In the preferred model specification this includes indicator variables for weekly mean temperatures and wind speed and weekly sum of precipitation falling into 10 bins by decile. In sensitivity analyses, I also include indicator variables for weekly average relative humidity and weekly average of daily minimum and maximum temperature. I estimate alternative specifications with different time fixed effects and location covariates to demonstrate the robustness of the results. (see section 5.1). Controlling for these ground-level weather conditions is important. Prior research has shown that extreme weather events affect health outcomes (see for example Deschenes and Moretti (2009)) while humidity is known to favour the spread of respiratory infections (Wu et al., 2016).

[3] The department (*département* in French) is one of the three levels of administrative divisions of France, below regions and above communes. There are 96 departments in metropolitan France.

The vector of time-varying location characteristics also always includes a variable for the sum of holidays days in a given week, to account for the drop in economic activity during the holidays that is expected to affect air pollution concentration and healthcare utilization.

To minimise concerns of auto-correlation, I estimate specifications in which I control for lags of the weather and pollution variables. I also investigate the possibility that increased air pollution leads to an anticipation of some healthcare costs that would have been incurred anyway by conducting sensitivity analyses in which I consider effects over longer time windows of several weeks. For example, I estimate the effect of pollution on week w on the healthcare costs across week w to $w + 3$. To ensure that the estimates do not capture the effects of pollution or weather conditions over the following two or three weeks, I include two to three leads of the pollution and weather variables. If there is some short-term displacement of healthcare costs, then the estimates could decrease when looking at longer time windows. Otherwise, estimates should remain unchanged or increase in case pollution has some lagged effects that are not captured when looking at a one-week time window. The results are generally robust to different lag and lead structures.

While I use weekly variation of pollutant concentrations and healthcare spending in the main analysis, I also run sensitivity analyses where I exploit the daily frequency of the data. I estimate the effect of pollution on day d and healthcare spending on that same day or on the following days using again different lag and lead structures. I also reproduce the model specification used in Deryugina et al. (2019) that estimates the effect of pollution on day d on the healthcare costs across day d to $d + 3$. To ensure that the estimates do not capture the effects of pollution or weather conditions over the following two or three days, I include two to three leads of the pollution and weather variables. To minimise concerns of auto-correlation, I also estimate specifications in which I control for lags of the weather and pollution variables.

In my empirical strategy, I assume that the zip code area of residence corresponds to the location of exposure to air pollution. However, people are also exposed to air pollution at their place of work, place of leisure or while commuting. If this leads to a large measurement error in pollution exposure, my estimates could suffer from bias. I check whether the results are robust to conducting the analysis at a higher level of spatial aggregation by running the analyses at the employment zone level. The employment zone (*"zone d'emploi"* in French) is a division of the French territory into 306 geographical areas within which most of the working population resides and works. For a map showing the boundaries of the employment zones, see Figure OA1 in the Online Appendix.

Standard errors are clustered at the zip code area level. The results are robust to clustering at more aggregate geographical levels (see section 5.1 on sensitivity analyses at the employment zone level).

4.2. Instrumenting air pollution using altitude atmospheric conditions

Models with location fixed effects can account for bias by capturing cross-sectional and time-invariant location-specific population characteristics, but they cannot fully address endogeneity bias due to correlations between pollution concentrations and economic activity. Controlling for location and time fixed effects means that the remaining variation in air pollution comes from any non-seasonal events affecting local air quality, such as local traffic restrictions or economic activity. However, traffic congestion or economic activity are potentially associated with stressful conditions that could be related to healthcare use. To avoid this kind of endogeneity bias, I instrument for changes in air pollution concentrations using altitude weather conditions.

A valid instrumental variable approach requires that the instrument is relevant, i.e. that it is sufficiently correlated with the endogenous variable of interest, and that the exclusion restriction is met, i.e. that the instrument is not correlated with unobserved determinants of the outcome of interest. As for the first condition, atmospheric conditions are known to affect air pollution concentrations. See Section 2.2 for a discussion of the relationship between air pollution and weather conditions at altitude. The results of the first-stage regressions confirm that weather conditions at altitude are strong predictors of air pollution concentrations (see section 5). Regarding the exclusion restriction, the identifying assumption in the present application is that, after flexibly controlling for various time and location fixed effects and ground-level weather conditions, the variation in pollution due to changes in weather conditions at altitude is not associated with changes in health care utilisation or costs, except through the effect on air pollution. It is plausible that this assumption holds. Although weather conditions at ground level can directly influence individual behaviour and health outcomes, atmospheric conditions at altitude are unlikely to directly influence health. Atmospheric conditions at altitude are not associated with economic activity, which means that the IV approach allows me to estimate the impact of air pollution on health costs without inadvertently capturing correlations due to economic activity.

The first stage specification is as follows:

$$P_{wpx} = \sum_{k} \beta_k IV_{wpk} + \alpha_p + \alpha_{m/mdep} + \alpha_y + \delta X_{wp} + \epsilon_{wpx} \tag{2}$$

where P_{xdp} denotes the mean concentration of pollutant x in week w in zip code area p and IV_{wpk} is atmospheric conditions k in week w and location p. The control variables and the fixed effects are the same as in equation 1.

The vector of altitude atmospheric conditions includes the number of hours of thermal inversion per week and the number of hours of thermal inversions per week by moment of the day, the average strength of these thermal inversions, the average of the planetary boundary layer taken over the entire week and over the moments of the day, the weekly average wind speed at twelve pressure levels, and wind direction bins based on average cardinal wind speed. Depending on the first stage specification, I also add interactions of the instruments with the location indicators to capture potential geographical variations of the atmospheric phenomena. See section 2.2 for a description of the altitude weather phenomena that serve as the basis for the construction of the instruments, as well as a discussion of exogeneity and the exclusion restriction. See section 3.3 for a detailed description of the construction of the instrumental variables.

Using a high-dimensional vector of instruments has the advantage of allowing to instrument for several pollutants simultaneously. This is interesting because air pollutants are both highly correlated and estimated to have independent effects on health. Instrumenting for several pollutants simultaneously can overcome problems that arise when implementing multi-pollutant approaches such as results becoming unstable and imprecise when incorporating more than one pollutant due to the correlation between the pollutants. If different subsets of instruments better predict variation in some pollutants than in others, then using a high-dimensional vector of instrument should allow to disentangle the effects the of different pollutants (Deryugina et al., 2019; Godzinski and Castillo, 2021). This is plausible because the different pollutants are not perfectly transported together, can be generated by sources in different locations, and are affected differently by atmospheric conditions at altitude. Using a set of instruments that includes frequently occurring events such as changes in planetary boundary height and changes in wind direction is also useful because Bagilet and Zabrocki (2021) show that an IV strategy with low-frequency events as instruments, such as using an indicator variable for the presence of a thermal inversion, can lead to inflated estimates due to low statistical power. At the very least, instrumenting for several pollutants is interesting for sensitivity analyses and validating the results of previous studies. With the exception of a few recent studies (Deryugina et al., 2019; Godzinski and Suarez Castillo, 2019), most of the existing literature is based on single-pollutant models.

I test the sensitivity of the results to different first-stage specifications including different combinations of instruments. I also test whether the results are robust when only one of the pollutants is instrumented while the others are included as controls.

To see if different sets of instruments indeed better predict a certain pollutant, I apply the IV LASSO approach proposed by Belloni et al. (2012) and implemented in *ivlasso* (Ahrens et al., 2020) to select pollutant-specific vectors of instruments. I then compare the model fit when the pollutant-specific instruments predict the pollutant for which they were selected with the model fit when these instruments are used to predict the concentrations of the other pollutants using the Bayes Information Criterion (BIC) and the Akaike Information Criterion (AIC) for model selection.

IV estimates a weighted average of the individual causal effects, also called the local average treatment effect (LATE). The term local emphasises that it is the weighted average that places the most weight on those entities whose treatment probability is most influenced by the instrumental variable. Interpreting IV estimates as a LATE requires imposing a monotonicity assumption (Imbens and Angrist, 1994; Mogstad et al., 2021). In the present application it means that I need to assume that air pollutant concentrations are always (weakly) positively or always (weakly) negatively correlate with a certain instrument in all zip code areas. The monotonicity assumption would be violated if the direction of the instrument-pollution relationship differs across zip code areas. To test whether the monotonicity assumption holds, I interact the instruments with location fixed effects which relaxes the assumption that the effect has to be monotonous across locations. The results remain robust to this approach.

4.3. Heterogeneity analyses

Besides estimating the overall healthcare costs of exposure to air pollution, I am also examining which types of health problems are particularly affected. To do so, I run separate regressions for 10 different categories of medical specialities. The dependent variable in these regressions is constructed as the weekly sum of expenditure on medical care in the respective medical specialty, including the costs of physician fees and all related expenditures such as drugs and exams. While interesting in itself, this exercise also serves as a sanity check. I examine both a set of medical specialties that are expected to be affected by air pollution and - as a placebo exercise - medical specialties that are not expected to be affected. I should find that air pollution has no effect on expenditure in the placebo categories. Otherwise, it would suggest that the estimates pick up some spurious correlation between air pollution and health expenditure and that the estimates of overall health effects are likely biased. The categories that I expect to be affected are family practice (primary care physician), cardiology and vascular medicine, pulmonology, otorhinolaryngology (O.R.L.), ophthalmology, and gynaecology. Family practice has been chosen because the first point of contact with the healthcare system in France is the family doctor, unless the health problem is an emergency that needs to be treated in the hospital emergency department. In France, the healthcare system is a gate-keeping system in which people must first visit their family doctor, who then refers them to specialists. Cardiology and vascular medicine, and pulmonology have been chosen because these are medical specialties that are frequently considered in the literature and for which effects of acute (short-term) exposure to air pollution has repeatedly been shown. O.R.L. and ophthalmology were selected because the short-term effects of air pollutants are irritation of the respiratory tract and mucous membranes. Gynaecology is considered as an additional category because there is evidence of short term effects of air pollution exposure on pregnancy outcomes (see Leiser et al. (2019)). Recent epidemiological research also shows that exposure to air pollution can affect the central nervous system and lead to cognitive issues (see for example Calderón-Garcidueñas et al. (2007)). Neurology was considered as an additional medical specialty likely to be affected by air pollution, but the results are not reported here as these analyses yielded only null results. As placebo, I consider the specialties of gastro-hepatology, nephrology, trauma surgery, and plastic surgery. Problems with the digestive system should not be affected by air pollution and trauma surgery and plastic surgery should not be affected, as accidents and planned operations should not react to air pollution exposure.

To identify populations at particular risk, I study heterogeneity in the effects of air pollution exposure as a function of patient characteristics, including age, chronic disease status and insurance status, and location characteristics, including average zip code area income, population size and average pollution concentrations. Previous literature has shown that air pollution-related mortality is higher in the very young and the very old, so a similar pattern might be expected for morbidity. Previous studies have also shown that the effects of air pollution are more pronounced in chronically ill and low-income populations. Many of these studies are correlative in nature, so it will be interesting to see if I find similar results. Analysis of heterogeneity by insurance status serves as an additional analysis of differential effects by economic status. I compare the results of those eligible for subsidized or publicly funded supplementary health insurance with those who are not. The income thresholds for these plans are between 70% and 90% of the poverty line (for more information, see Section 3. I expect the effects to be greater in more populated areas because of the size of the population, but also because more populated areas are generally more polluted. As regards heterogeneity by baseline pollution level, the results are a priori ambiguous. Some studies show that marginal effects are greater at high pollution levels (convex dose-response function), while others show that marginal effects are greater at low pollution levels (concave dose-response function). For references, see the discussion in section 5.4. These heterogeneity analyses can provide information for defining policy priorities. If effects are concentrated in densely populated, highly polluted cities, and marginal effects are greatest at higher concentrations, then policy efforts should clearly focus on improving local air pollution in large cities. If the effects are still significant outside cities, and marginal effects are higher at lower concentrations, it may be unwise to focus solely on air pollution reduction in the big cities.

5. Results

This section first presents the results of the effects of air pollution exposure on aggregate health costs comparing the location-fixed effects model and the model instrumenting air pollution with atmospheric conditions including sensitivity analyses. Then the results of heterogeneity analyses are presented, where the effects are estimated separately by medical specialty and by patient and location characteristics. This is followed by a discussion of effect size and policy implications.

5.1. Effect of air pollution exposure on healthcare costs

Table 1 reports the main estimates of the relationship between weekly average air pollutant concentrations and weekly healthcare expenditure at the zip code area level. The first two columns show results for the location fixed effect model (FE) and the last two columns show results for the location fixed effect instrumental variable model (FE-IV) in which altitude atmospheric conditions are used as instruments for the air pollutant concentrations. Columns 1 and 3 present results for a model excluding any lags of the pollutants and columns 2 and 4 present the model including one week lag of pollutant concentrations. The coefficients indicate the increase in average additional healthcare spending per zip code area for a 1 $\mu g/m^3$ increase in weekly average pollutant concentrations. For example, in the FE-IV model in column 4, each 1 $\mu g/m^3$ increase in weekly average NO_2 leads to an average €17.23 of additional healthcare expenditure per zip code area during the same week. This corresponds to a 0.48% increase relative to the average weekly zip code area healthcare spending. The €17.23 of additional healthcare expenditure per zip code area per week in a sample of 1/97 of the French population is linearly extrapolated to €525,620,310 of additional overall healthcare spending in France per year or 0.02% of France's GDP in 2019[4]. For a discussion regarding the effect size, see section 5.4.

[4]The calculation €17.23 $\cdot$ 97 $\cdot$ 52 $\cdot$ 6,048 = €525,620,310 for the effect times the adjustment for the sample of the total population, times the number of weeks in a year, times the number of zip code areas. The GDP of France in 2019 was 2.43 trillion (INSEE, 2020).

Table 1: Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly healthcare expenditure

	Weekly healthcare expenditure			
	FE		FE-IV	
	(1)	(2)	(3)	(4)
NO_2	44.33***	43.36***	18.42***	17.23***
	(2.692)	(2.420)	(3.820)	(3.719)
O_3	4.189***	4.837***	6.282***	3.275***
	(0.383)	(0.387)	(0.773)	(0.662)
PM_{10}	-12.06***	-13.34***	12.37***	3.540
	(0.981)	(0.996)	(2.815)	(2.843)
Lag NO_2		8.947***		-3.423
		(2.119)		(4.062)
Lag O_3		-0.175		6.497***
		(0.364)		(0.795)
Lag PM_{10}		-1.412		18.14***
		(0.872)		(2.616)
Observations	1,209,572	1,186,311	1,209,572	1,186,311
First-stage F-stat.			2055.4	824.2

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. The table reports the main estimates of the relationship between average weekly air pollutant concentrations and weekly healthcare expenditure. The coefficients indicate the increase in average healthcare spending per zip code area per week for a 1 $\mu g/m^3$ increase in weekly average pollutant concentrations. The first two columns show results for the location fixed effect model (FE) and the last two columns show results for the location fixed effect instrumental variable model (FE-IV) in which altitude atmospheric conditions are used as instruments for the air pollutant concentrations. Columns 1 and 3 present results for a model excluding any lags of the pollutants and columns 2 and 4 present the model including one week lag of pollutant concentrations. Table OA2 in the appendix shows the corresponding first stage regressions. All regressions include month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level in parenthesis.

The instruments used in the first stage regressions corresponding to the FE-IV models in Columns 3 and 4 are the number of hours of thermal inversion per week and the number of hours of thermal inversions per week by moment of the day, the average strength of these thermal inversions by moment of the day, the average of the planetary boundary layer taken over the moments of the day, and the weekly average wind speed at twelve pressure levels. The large F-statistics shown at the bottom of Table 1 indicate that the instruments are strong predictors of pollution concentration. Table OA2 in the Online Appendix shows the corresponding first stage regressions.

Disentangling the effects of different air pollutants in a multi-pollutant model is challenging. The location fixed effects model produces negative coefficients for the effect of PM on healthcare expenditure as shown in columns 1 and 2 of Table 1. This result is counter-intuitive, as it would mean that the increase in particle pollution has protective effects on health, since it leads to a reduction in healthcare expenditure. Unexpected negative coefficients and unstable results have also been found in the previous literature when incorporating more than one pollutant due to the correlation between the pollutants. The estimation of an instrumental variable model appears to solve this problem, since it produces the expected positive signs as shown in columns 3 and 4 of Table 1). The use of a high-dimensional vector of instruments could indeed make it possible to distinguish the effects of different pollutants if different subsets of instruments are better at predicting the variation of some pollutants than others (see for example Godzinski and Castillo (2021)). When I apply an IV LASSO approach to select pollutant-specific instrument vectors for the first stage regression I find that each pollutant is indeed better predicted by a different set of instruments. Table OA3 in the Online Appendix shows the first stage regression results where each pollutant is regressed over the LASSO-selected variables. Different instruments are selected for each pollutant and when the same instruments are selected, the sign and magnitude of their effect differs by pollutant. The model fit in terms of the Bayesian Information Criterion (BIC) and the Akaike Information Criterion (AIC) is greater when the vector of instruments predicts the pollutant that it has been chosen to predict compared to when it predicts the other pollutants. See Table OA4 in the Online Appendix. The results using the LASSO-selected instruments are almost identical to the results using the FE-IV approach as can be seen in Table OA5 in the Online Appendix.

5.2. Sensitivity analyses

5.2.1. Lagged effects and effects over longer time windows

My estimation strategy relies on short term variation in atmospheric conditions and air pollutant concentrations. The estimation results therefore reflect only the short term health effects of air pollution exposure. Without capturing the effects of chronic exposure, linearly scaling the estimated effects to obtain yearly healthcare costs should yield a lower bound for the overall health effects of air pollution exposure. There are two potential problems with this interpretation as a lower bound. First, despite the inclusion of month and season fixed effects, pollutant concentrations might be auto-correlated and exposure could have some lagged effects. The estimates for the effect of air pollution concentration on health expenditure in the same week could therefore pick up the effects of the previous week and therefore be inflated. Second, the estimates may be overestimated if exposure to increased air pollution leads to an anticipation of some healthcare costs that would have been incurred anyway. To address the first issue, I estimate models that include lags of the pollutants. Column 4 in Table 1 shows results for the preferred model specification with one week lags of the pollution concentrations, using one week lags of the atmospheric conditions as instruments. The coefficient for the effect of NO_2 pollution on health expenditure in the same week remains unchanged, but the size of the coefficient for the effect of O_3 pollution is reduced by half and the effect of PM_{10} is no longer statistically distinguishable from zero. The coefficients for the lags of O_3 and PM_{10} pollution are statistically significant and of a similar order of magnitude to the coefficients for the same week effects in the model that excludes lags. This suggests that there are some lagged effects of exposure and that these effects are partially captured by the coefficients for the same week effects in models that do not consider lags. For a more conservative approach, I use the estimates of the same-week effects from the model including the lags in the calculation of the healthcare costs in section 5.4. To investigate the second issue of whether some of the estimated healthcare costs might result from a shift in spending over time rather than from additional costs arising from pollution, I conduct sensitivity analyses in which I consider effects over longer time windows of several weeks. If there is some short-term displacement of healthcare costs, then the estimates could potentially decrease when longer time windows are considered. Otherwise, the estimates should remain unchanged or increase in case pollution has some lagged effects that are not captured when considering a one week window. Table OA6 in the Online Appendix shows results for models where I estimate the effect of weekly air pollution exposure and its one week lag on healthcare expenditure over two weeks to four weeks, controlling for the appropriate number of weather and instrument leads. I find that the estimates increase with the length of the time window, with one exception for the coefficient for the lag of NO_2 pollution, where a sign reversal occurs. Overall, this suggests that pollution has some lagged effects and that the effects are not due to a displacement of expenditure over time. When I consider even longer time windows, the results become unstable, including some sign changes. However, these results are likely due to the difficulty of estimating a model with multiple pollutants, rather than evidence for expenditure displacement, as the coefficients increase monotonically when I estimate single-pollutant models.

5.2.2. Importance of considering multiple pollutants

It is important to estimate the effect of multiple pollutants simultaneously. Including only one of the pollutants at a time or including either NO_2 or PM together with O_3 in a two-pollutant model yields coefficients that are mostly of the expected positive sign in both the FE and FE-IV models, as can be seen in Table OA8 in the Online Appendix. However, excluding some of the pollutants mean that part of their effect are now captured by the coefficients on the included pollutants. The coefficients in the FE-IV model main specifications in Table1 that include all of the pollutants are indeed different from the coefficients in the one or two-pollutant models. In the one and in the two-pollutant models, the coefficients on NO_2 and PM are larger while the coefficient on O_3 is smaller. The direction of the bias is consistent with the correlations between the pollutants. NO_2 and PM are positively correlated, meaning that an omission of one of these pollutants leads the coefficient on the included pollutant to overstate its effect. O_3 is mostly inversely related to NO_2 and PM. When O_3 increases, NO_2 and PM tends to increase and the health effects of an increase from O_3 are therefore attenuated by the health benefits from the increased PM and NO_2 when these pollutants are not included in the regression. For a discussion on the correlation between the pollutants, see section 2. The results are robust to instrumenting only one pollutant at a time while including the others as controls as shown in Table OA9 in the Online Appendix.

In the preferred model specification I use PM_{10} pollutant concentrations to estimates the effect of particulate matter on healthcare costs. This choice is motivated by the fact that $PM_{2.5}$ is nested within PM_{10} and PM_{10} is therefore a broader measure of particulate matter pollution. Both pollutants are highly correlated and the results for models including $PM_{2.5}$ instead of PM_{10} pollution are quantitatively and qualitatively similar to the results from the preferred model specification using PM_{10} as can be seen in Table OA10 in the Appendix.

A remaining concern is that there are other air pollutants that impact health and are correlated with the pollutants examined in this study. While I am analysing the effects of the three pollutants that are generally considered to have the greatest impact on health, sulphur dioxide (SO_2) and carbon monoxide (CO) are two additional pollutants that are also widely considered key pollutants and the subject of regulatory measures. Unfortunately, I do not have high quality, high spatial resolution data for these pollutants. Instead, I use data from monitoring stations provided by the European Environment Agency (EEA) for sensitivity analyses. See section 3 for information about this data source. Table OA11 in the Online Appendix shows that the results for the effect of NO_2, O_3 and PM pollution are robust to including SO_2 and CO pollution as control variables (columns 1 and 2). The results are also similar when SO_2 and CO are included as additional instrumented pollutants (columns 3 and 4), except for unexpected negative coefficients on the effect of CO pollution and the lag of NO_2 pollution in the model with one week lagged effects. This is probably due to the fact that it becomes more difficult to disentangle the effects of a greater number of pollutants.As an additional sensitivity analysis, I use the EEA measuring station data on NO_2, O_3 and PM_{10} instead of the more high-resolution reanalyses data from INERIS. The results are qualitatively similar as long as a one week lag of the pollutants are included as control variables or as instrumented variables. See Table OA12 in the Online Appendix).

5.2.3. Robustness to alternative model specifications and placebo exercise

In general, the results are robust to different first-stage specifications including a different number and different combinations of instruments. The regression results shown in Table 1 are the most conservative across different first-stage specifications. The second-stage coefficients tend to be larger when I use fewer instruments. See Columns 1 and 2 of Table OA13 in the Online Appendix that show results for a regression including as instruments only the number of thermal inversions per week, their average strength, average planetary boundary height and average wind speed at the lowest altitude layer above ground-level. Adding in addition wind direction as instrument where I interact dummies for the weekly average wind direction by 90-degree intervals with location dummies similar to the IV specification used by Deryugina et al. (2019) yields results that are very similar to the results from the main specification as can be seen in Columns 3 and 4 of Table OA13 in the Online Appendix. The wind direction instrument must necessarily be interacted with the location fixed effects to account for the fact that wind direction shifts pollution concentrations differently depending on the location of pollution sources relative to the location under consideration. The results are robust to interacting the instruments with location fixed effects more in general. Columns 5 and 6 of Table OA13 show results where all instruments are interacted with location (employment zone) fixed effects. Adding interactions of the instruments with the location indicators should capture potential geographical variations of the effect of the atmospheric conditions instruments on pollutant concentrations. Similar results for models with and without location FE interactions suggests that the monotonicity of instrument effects across location holds. The monotonicity assumption requires that air pollutant concentrations are always weakly positively or always weakly negatively correlated with an instrument in all zip code areas. The interaction of instruments with location fixed effects relaxes the assumption of monotonicity between locations since only monotonicity of the instrument effect within a given location is required. The monotonicity assumption is necessary to be able to interpret the IV estimates as the local average treatment effect (LATE) (Imbens and Angrist, 1994).

The results are also robust to changing the way the weather control variables are included as can been seen in Table OA14 in the Online Appendix. Column 1 shows results for a model including humidity and minimum and maximum temperatures as additional ground-level weather controls. Column 2 shows the results using weather fixed effects where the variables have been partitioned into 5 bins by quintiles of their values and Column 3 shows results for 15 bins. The results are similar to the results from the preferred specification using 10 bins. The results are also similar when I use the non-transformed weather variables as shown in column 4. Column 5 shows that the results are also robust to using month-by-department fixed effects rather than month fixed effects to allow for different effects of seasonality in pollution and healthcare expenditure across the 96 French Departments.

I conduct a placebo exercise where I randomly reshuffle the values of the instrumental variables and use those shuffled instruments in the first stage instead of the actual instruments. As can be seen in Table OA15 in the Online Appendix, the first-stage F-statistics are very small, which provides evidence that the instruments are picking up meaningful rather than spurious variation in pollution levels. Using the shuffled instruments also leads to second-stage estimates that are statistically not significant.

In the main analyses, the standard errors are clustered at the level of the zip code area but the results are robust to clustering at the more aggregate level of the employment zone that divide the French territory into 306 zones.

5.2.4. Robustness to different levels of spatial and temporal aggregation

As an additional robustness exercise, I run regressions using the data at daily frequency to estimate the very short-term impact of an increase in pollution on a given day on the impact of health care spending on the same day. The results are shown in Table OA7 in the Online Appendix. Controlling for two days lag of pollutant concentrations, I find that an increase of daily average NO_2 pollution by 1 $\mu g/m^3$ leads to an increase in daily zip code area healthcare costs of €4.95 while an increase in daily average O_3 pollution by 1 $\mu g/m^3$ leads to an increase in healthcare spending of €0.7. Linearly scaling these effects to a week yields an effect of €34.7 and €4.9 additional healthcare costs for increases in NO_2 and O_3 pollution, respectively. While qualitatively similar, the results from models using weekly frequency data are comparatively more conservative - €17.2 and €3.3 for a one-unit increase in NO_2 pollution and O_3 pollution, respectively.

In my empirical strategy, I assume that the zip code area of residence corresponds to the location of exposure to air pollution. However, people are also exposed to air pollution at their place of work, place of leisure or while commuting. I check whether the results are robust to conducting the analysis at a higher level of spatial aggregation by running the analyses at the employment zone level. The employment zone (*"zone d'emploi"* in French) is a division of the French territory into 306 geographical areas within which most of the working population resides and works. For a map showing the boundaries of the employment zones, see Figure OA1 in the Online Appendix. Table OA16 in the Online Appendix shows that the results are qualitatively similar when the analysis is carried out at the employment zone level, albeit less statistically significant. Some of the results at the employment zone level are even quantitatively close to the results at the zip code area level. Column 3 indicates that an increase of one unit in the weekly average NO_2 and O_3 exposure increases weekly health expenditure at the employment zone level by €438.3 and €83.68 respectively. A linear scaling of these amounts to the annual costs for the entire French population results in €663,235,560 and €126,624,576. These estimates are similar to the additional healthcare costs of €525,620,310 and €99,907,517 resulting from a one-unit increase in average NO_2 and O_3 pollution, respectively, estimated using the weekly frequency data.

5.3. Effect heterogeneity

This section presents the results of heterogeneity analyses, including the results of regressions conducted separately by medical speciality, patient and location characteristics.

Results by medical speciality

I investigate which type of health condition is affected by exposure to air pollution by running separate regressions for 10 categories of medical specialities. While interesting in its own right, this exercise also serves as a sanity check. I examine both a set of medical specialities that should be affected by air pollution and - as a placebo exercise - medical specialities that should not be affected. I should find that air pollution has no effect on expenditure in the placebo categories. Otherwise, it would suggest that the estimates pick up some spurious correlation between air pollution and health expenditure and that the estimates of overall health effects are likely biased. The categories that I expect to be affected are family practice (primary care physician), cardiology and vascular medicine, pulmonology, otorhinolaryngology (O.R.L.), ophthalmology, and gynaecology. The placebo specialties are gastro-hepatology, nephrology, trauma surgery, and plastic surgery.

Table 2 shows results by medical speciality using the preferred FE-IV specification. The dependent variable in these regressions is constructed as the weekly sum of expenditure on medical treatments in the respective medical specialty, including the costs of physician fees and all related expenditures such as drugs and exams. Family practice shows the strongest response, with coefficients for all air pollutants statistically significantly different from zero. This is consistent with the fact that the family practitioner is the first point of contact with the healthcare system before orienting patients to specialist care or the only point of contact in case of minor health problems. The estimates suggest that cardio-vascular issues are affected by NO_2 and O_3 pollution while pulmonology is affected by PM_{10} pollution. For O.R.L., effects are found for the lags of O_3 and PM_{10} exposure, while there are effects of NO_2 and the lags of O_3 and PM_{10} pollution on expenditures for ophthalmology. For gynaecology, I find effects for the lag of O_3 exposure but the coefficient is only significant at the 5% level. These effects are consistent with the findings from the economic and epidemiological literature, which have shown effects of all three pollutants on health problems falling into these medical specialities. All but one of the coefficients have the expected positive sign and the only statistically significant negative coefficient is only significant at the 5% level. The placebo categories do not appear to be affected, as none of the estimates are statistically significantly different from zero. In contrast, many of the estimates from the simple location FE model shown in Table OA17 in the Online Appendix are negative, highlighting again the difficulty of estimating the effects of multiple correlated pollutants simultaneously without using instruments. The simple location FE model also yields statistically significant estimates for the placebo categories which suggests that model estimates from models that do not use instruments for pollution concentrations pick up spurious correlation.

Assuming no knowledge of which outcomes are affected and which pollutants affect the outcomes, I conduct a total of 60 hypothesis tests across 10 specialties for three pollutants and three pollutant lags. This increases the likelihood of false positives by a factor of 60 compared to analysing a single treatment. To control the Family-Wise Error Rate (FWER) — the probability of at least one false rejection

Table 2: Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly healthcare expenditure - regressions run separately by medical specialty

	Family practice	Cardio-vasc.	Pulmo.	O.R.L.	Ophthalmo.
NO_2	4.956***	0.466*	0.0177	0.0236	1.108***
	(1.492)	(0.223)	(0.179)	(0.084)	(0.228)
O_3	0.927***	0.0401	0.0127	0.00108	0.107*
	(0.235)	(0.040)	(0.035)	(0.017)	(0.042)
PM_{10}	-1.180	-0.0541	0.180	-0.0468	-0.336*
	(1.143)	(0.159)	(0.139)	(0.062)	(0.170)
Lag NO_2	2.513	-0.00495	-0.300	0.0897	0.192
	(1.297)	(0.225)	(0.202)	(0.084)	(0.240)
Lag O_3	1.217***	0.178***	0.0273	0.0476**	0.206***
	(0.264)	(0.041)	(0.031)	(0.017)	(0.044)
Lag PM_{10}	3.329***	0.239	0.260*	0.140**	0.318*
	(0.835)	(0.142)	(0.126)	(0.052)	(0.149)
	Gynaeco.	Nephro.	Gastro-hep.	Trauma surg.	Plastic surg.
NO_2	0.102	0.0517	-0.513	-0.107	-0.0235
	(0.147)	(0.082)	(0.345)	(0.218)	(0.108)
O_3	0.00422	0.0130	0.0480	0.0276	0.0306
	(0.029)	(0.017)	(0.084)	(0.040)	(0.021)
PM_{10}	0.170	-0.0335	0.370	0.172	0.129
	(0.111)	(0.060)	(0.278)	(0.159)	(0.080)
Lag NO_2	0.0581	0.0115	-0.285	0.327	-0.111
	(0.160)	(0.091)	(0.410)	(0.222)	(0.106)
Lag O_3	0.0644*	0.0138	0.0281	0.0756	-0.0109
	(0.031)	(0.017)	(0.074)	(0.041)	(0.022)
Lag PM_{10}	0.0318	0.0418	0.206	-0.0926	0.0409
	(0.094)	(0.056)	(0.286)	(0.139)	(0.068)
Observations	1186311	1186311	1186311	1186311	1186311
FS F-stat	824.2	824.2	824.2	824.2	824.2

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. The table reports results for regressions run separately by medical speciality using the preferred FE-IV specification including a one week lag of the pollutant concentrations. The dependent variables are the weekly sum of expenditure on medical treatments in the respective medical specialty. The coefficients indicate the increase in average healthcare spending per zip code area per week for a 1 $\mu g/m^3$ increase in weekly average pollutant concentrations. The medical specialties expected to be affected are family practice (primary care physician), cardiology and vascular medicine, pulmonology, otorhinolaryngology (O.R.L.), ophthalmology, and gynaecology. Gastro-hepatology, nephrology, trauma surgery, and plastic surgery are placebo categories. All regressions include month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level in parenthesis.

— I use the Holm-Bonferroni correction (Holm, 1979), which ranks p-values in ascending order $p_{(1)} \leq p_{(2)} \leq \cdots \leq p_{(60)}$ and sequentially compares the p-value ($p_{(i)}$) against the threshold ($\frac{\alpha}{m-i+1}$), stopping at the first failure to reject. Using this approach, the coefficients that remain statistically significant are NO_2 and lagged O_3 for ophthalmology, lagged O_3 for cardio-vascular disease, lagged PM_{10}, O_3, lagged O_3 and NO_2 for family practice (see Table OA18 in the Online Appendix). Alternatively, one could remain agnostic only about which pollutants affect spending within each specialty, rather than across all specialties. In that scenario, each speciality's regressions form a separate "family" of hypotheses, and one applies the correction within each specialty for 6 hypotheses (3 pollutants and their lags). This is a less conservative approach, but still reasonable as the categories are selected based on evidence from the literature. This leads to 3 statistically significant coefficients for ophthalmology (NO_2, lagged O_3, and O_3), 1 for cardio-vascular disease (lagged O_3), 2 for O.R.L.(lagged O_3 and lagged PM_{10}) and 4 for family practice (NO_2, O_3, lagged O_3 and lagged PM_{10}).

Results by patient and location characteristics

Many of the existing studies on the health effect of air pollution focus on the young or elderly populations as these populations are generally considered to be the most vulnerable (see for example Manisalidis et al. (2020) for a review). The middle-aged adult population is often omitted from these studies, which makes it impossible to compare the severity of the impact in this group compared to other age groups. I find evidence of effects across all age categories, suggesting that adverse health effects also manifest in parts of the population that are less often considered. See Table OA19 in the Online Appendix that shows the FE-IV model results for regressions run separately for observations divided into age groups. One possible explanation is that many studies conducted so far focus on mortality, an outcome that may concern only the most vulnerable populations. I look at healthcare costs, which include the costs of treating milder health outcomes that are likely to occur in all age groups. Another possibility is that the middle-aged population is often under-represented in many epidemiological studies. The effects appear to be most pronounced in the 40-60 age group, possibly because people in this age group already have some pre-existing health conditions and may be frailer than younger people, and because this age group is also likely to be more exposed to air pollution than older people as they are more likely to lead active lives and spend more time outdoors.

To determine whether there are geographic disparities in the impact of pollution on health expenditures, I run separate regressions for observations divided into groups according to their average zip code area characteristics. Tables OA20 and OA21 in the Online Appendix shows the results of the regressions for observations categorised into groups below and above the median in terms of zip code average household income, pollutant concentration and population size. Examination of the impact of pollution on healthcare expenditure in absolute terms reveals that most healthcare expenditure is incurred in the more populated areas, which are on average also higher-income and more polluted areas (see Table OA20 in the Online Appendix). Higher expenditure in areas where the number of people affected by air pollution is higher is not surprising.

A more interesting picture emerges when studying the effect on per capita healthcare expenditure. First, I find no clear evidence of differential effects of pollution on per capita healthcare spending based on zip code average household income. As can be seen in Panel A of Table OA21, the increase in weekly per capita health care spending for a one-unit increase in average pollutant concentration is similar in locations with average household income below and above the median. When the observations are categorised into groups in a different way - for example into terciles and quartiles of average income - no clear pattern emerges either. It is possible that there are differences between low- and high-income locations but that they are not detectable in the data. There might exist significant income heterogeneity within a particular zip code area that is unobserved here and that is relevant for differences in the health effects. However, analyses of effect heterogeneity by economic status as approximated by enrolment status in the state funded complementary insurance plan available to low-income individuals (see Section 3 also do not yield any clear pattern. Individuals of low socio-economic status could sort into more polluted areas but consume less healthcare services because they are liquidity constrained. However, affordability of care should not be a concern in the French context. The French healthcare system is universal, and reimbursement rates for medical expenses are generally high. In addition to compulsory basic health insurance, which reimburses 70-80% of costs, most people have supplementary insurance, often through their employer, which reimburses the remainder. For people on low incomes, the state provides subsidies or covers the full cost of supplementary health insurance. France also has nationally-regulated fee schedules for most medical services, ensuring that basic costs are relatively consistent across the country. Differences in healthcare prices depending on location should therefore not play a significant role. The most likely explanation is that, given the characteristics of the French healthcare system, exposure to air pollution generates similar healthcare costs across income groups.

Second, I find differences in the effects by zip code average NO_2 concentration and population size. The results in panels B and C of Table OA21 indicate that the effects of pollution on per capita healthcare expenditure are stronger in locations with NO_2 concentration and population size below the median. While pollution in more populated urban areas affects a greater number of people and therefore has a large effect on overall spending, pollution seems to have greater marginal health effects in relatively clean and less populated areas. These results are consistent with a concave relationship between air pollution exposure and health effects where pollution has greater marginal effects on health at low concentrations. To further examine the non-linearity of the effects, I run piece-wise linear regressions in which I interact the weekly pollutant concentration with a dummy variable that categorises that week's pollutant concentration into four categories per quartile of its value. Table OA22 in the Online Appendix shows that the effect of a one-unit increase of average weekly pollution concentration when the pollution concentration during that week belongs into the lowest quartile is greater than the effect of a one-unit increase when the pollution concentration is in the second, third or fourth quartile. The same applies to the effect of O_3 and PM pollution. This is consistent again with a concave relationship between air pollution exposure and health effects or concave concentration-response function. Greater marginal effect of pollution at low pollution levels is in contrast with some of the findings in the literature (Schlenker and Walker, 2016; Arceo et al., 2016; Dechezleprêtre et al., 2019) but some more recent studies have indeed suggested that the concentration-response function for pollution might be supra-linear. For example, Miller et al. (2021) and Henderson et al. (2024) show that small air $PM_{2.5}$ pollution shocks have proportionally larger mortality effects than large air pollution shocks.

5.4. Effect size and policy discussion

The results from the preferred FE-IV model specification indicate that a 1 $\mu g/m^3$ increase in weekly average NO_2 concentrations leads to an average €17.23 of additional healthcare expenditure per zip code area during the same week. Similarly, a on 1 $\mu g/m^3$ increase in weekly average O_3 levels leads to an average €3.28 of additional healthcare expenditure. These €17.23 and €3.28 of additional healthcare expenditure per zip code area per week in a sample of 1/97 of the French population corresponds to €525,620,310 and €100,060,047 of additional healthcare spending in France per year or 0.03% of France's GDP in 2019.[5]. These are changes in annual health expenditures for a $1\mu g/m^3$ change in air pollution concentration. To better understand the magnitude of the effect, consider the total effect of air pollution

[5]To illustrate, consider the calculation for NO_2: $€17.23 \cdot 97 \cdot 52 \cdot 6{,}048 = €525{,}620{,}310$ where the effect is adjusted for the sample size, multiplied by 52 to obtain the yearly effect and multiplied by the number of zip code areas to obtain the effect for France.

by multiplying the estimated annual health costs for a one-unit change by the 2015-2018 average air pollution concentrations in the data. The total effect of air pollution are then a yearly additional healthcare expenditure of €7,243,047,872 for NO_2 and €5,573,344,618 for O_3 pollution, resulting in an overall effect of €12,816,392,490[6]. These €12.8 billion of additional healthcare costs per year correspond to 0.5% of France's GDP in 2019 and 6.2% of France's total healthcare expenditure in 2019[7].

These cost estimates are around 10 times larger than previous estimates of the costs of air pollution to the French health system. A 2015 Senate Committee of Inquiry into the economic and financial cost of air pollution (Sénat, 2015) searched for estimates of the total costs of air pollution to the French healthcare system, resulting in a report on two studies that considered only a fraction of the total possible healthcare costs and a recommendation that more research be conducted in this area. One of the studies is a 2007 impact study conducted by the French Agency for Environmental and Occupational Health Safety (Fontaine et al., 2007) investigating the costs related to asthma and cancer and presenting an estimate of the overall cost situated between 0.3 and 1.3 billion euros. The other study dates from 2015 and was carried out by the General Commission for Sustainable Development (Rafenberg, 2015) arriving at an overall cost of between 0.9 billion euros and 1.8 billion euros per year. The study covered only the costs related to respiratory diseases (asthma, acute bronchitis, chronic bronchitis, chronic obstructive pulmonary disease, cancers), and hospitalisations for respiratory and cardiovascular issues. I am not aware of any other study that quantified healthcare costs in France more comprehensively.

In general, the assessment of the health costs of air pollution has been comparatively incomplete in both the quasi-experimental and epidemiological literature. Studies that attempt to assess the health costs of air pollution for cost-benefit analyses often include only a selection of health effects and a part of the population for which the epidemiological evidence is most robust. For example, the Environmental Benefits Mapping and Analysis Program–Community Edition (BenMAP-CE), a tool historically used by the Environmental Protection Agency (EPA) and widely employed to estimate the economic impact of the health outcomes of air pollution, considers in its default features only the costs of hospital and emergency department admissions. When an additional quantification including also ambulatory care is added, only a subset of health effects have been considered (see for example Birn-

[6] $13.78 \cdot N525,620,310 = N7,243,047,872$ for NO_2 and $55.7 \cdot N100,060,047 = N5,573,344,618$ for O_3 pollution.

[7] In 2019, the GDP of France reached €2.43 trillion (INSEE, 2020) and aggregate healthcare spending was €208 billion (DREES, 2020).

baum et al. (2020) who consider only two disease categories, respiratory and all cardiovascular disease). Quasi-experimental studies are similarly limited in scope, since they focus on relatively narrow geographical areas and time periods and/or concern only a specific part of the population and a selection of health effects, often concentrating on mortality. The quasi-experimental studies that is most comparable to the present study in terms of data quality and empirical strategy is Deryugina et al. (2019). Using wind direction as instrument for PM 2.5 pollution, the study investigates the health effects on Medicare beneficiaries in the US, i.e. people aged over 65. While the focus lies on mortality costs, the study also provide an estimate of hospital costs. A decrease in average PM 2.5 concentrations of $4.9\mu g/m^3$ in the US between 1999 and 2013 is estimated to have saved hospital costs of USD 1.5 billion per year. Considering €13.03 of additional weekly healthcare costs at the zip code area level for a one unit increase in weekly average $PM_{2.5}$ concentrations from the multi-pollutant FE-IV model in Table OA10, scaling it to a yearly estimate for France and multiplying it by 4.9 for the change considered in Deryugina et al. (2019) yields €1,947,723,733. Scaling the cost for the French population of roughly 67 million to the size of the population of 55 million people aged over 65 in the US yields €1,598,877,691 or USD 1,647,219,757. My estimate of the overall costs for the French healthcare system is not too far away from the USD 1.5 billion estimate for the United States from Deryugina et al. (2019), which only considers hospital costs. However, healthcare costs are on average much higher in the US than in France and also any other country in the world(Papanicolas et al., 2018)[8]. This comparison shows that it is important to have separate cost estimates for Europe, as the cost estimates from the USA are not applicable to other health systems.

The healthcare cost estimates presented in this study are sizeable compared to estimates of the costs of further pollution reduction. The total cost of complying with the EU National Emission Commitment (NEC) Directive (European Parliament, 2016) 2030 air pollution target values considering 2017 pollution levels has been estimated at €9.9 billion per year by Amann et al. (2017). This includes not only the cost of reducing NO_2 but also the cost of reducing other pollutants. Compliance with the NEC Directive requires France to reduce nitrogen oxides (NOx, composed of both NO_2 and NO) by 50% compared to 2005 values, to be achieved from 2030. In 2005, annual NO_2 concentrations in France were 17.5 $\mu g/m^3$ (INERIS, 2024), which means that France should reduce NO_2 by 8.75 $\mu g/m^3$ from its 2005 levels

[8]To illustrate, an appendectomy performed in the United States will cost an average of USD 33,000, or around €29,000, compared with only €600 in France. Taken as an example from the website of an international health insurance for expatriates availablehere.

until 2030. Given the 2017 average of 12.01 $\mu g/m^3$ (INERIS, 2024), this implies a further decrease of 3.26 $\mu g/m^3$ of annual NO_2 concentration. According to my estimate, this 3.26 $\mu g/m^3$ of annual NO_2 concentration should lead to savings of €1.7 billion in annual health expenditure. The health cost savings from complying with the NO_2 pollution reductions alone (disregarding other pollutants) should therefore account for 17% of the estimated total cost of complying with the NEC Directive for France. My estimate of €1.7 billion savings in annual health expenditure for compliance with the NO_2 limit values in France alone are almost as large as the estimate of €2.4 billion of annual health cost savings of full compliance with the NEC Directive for the *entire European Union (EU28)* considered in Amann et al. (2017). My health cost savings estimate for France (home to 13% of the total EU population) for compliance regarding NO_2 standards disregarding reductions of other air pollutant concentrations already corresponds to 70% of the health cost savings considered in Amann et al. (2017), suggesting that the health cost savings considered in Amann et al. (2017) are largely underestimated.

Although the health costs presented in this study are higher than those from the previous literature, they still represent a lower bound for the total health costs of air pollution. The estimates only reflect the short-term effects of air pollution, not the effects of chronic exposure. The cost estimates do not include the costs of unobserved behavioural responses. Short-run increases in air pollution have been shown to cause people to stay indoors (Neidell, 2009; Zivin and Neidell, 2009) or buy indoor air purifiers (Ito and Zhang, 2020). Mortality costs or the costs of lost productivity due to illness are not considered in this study. Finally, the choice of estimates for the cost calculation is conservative. Some model specifications produce larger health cost estimates.

The results of this study provide highly relevant information for public policy decisions. My cost estimates show that the health costs of air pollution to the French health system have been severely underestimated. Previous policy decisions were therefore based on cost-benefit calculations that did not take into account health cost savings from further reductions in air pollution in the order of several billion euros per year. The health costs estimated here are caused by air pollution levels that are mostly far below the current European air quality guideline values, indicating that the current guideline values are set too high. A review of EU air quality guidelines is currently underway. On 26 October 2022, as part of the European Green Deal, the Commission proposed to revise the Ambient Air Quality Directives to align the air quality standards more closely with the 2021 recommendations of the WHO (European Commission, 2016). This planned revision is a step in the good direction. It would signify a reduction of the limit values for NO_2, PM_{10} and $PM_{2.5}$. However, this study provides evidence that there are likely significant health benefits from reducing pollutant levels even further below the current WHO guideline values. I find that the marginal effects are greater in relatively low-pollution and less populated areas. Reducing population exposure even at low air pollution concentrations should therefore be an important public health goal. When I consider the current WHO air quality guideline values of 10 $\mu g/m^3$ annual mean for NO_2 and 15 $\mu g/m^3$ annual mean for PM_{10} pollution more specifically, I find that an increase in pollutant concentrations of one unit below the guideline values has a greater impact than an increase of one unit above the guideline values. See Table OA23 in the Online Appendix. Even the most stringent 2021 WHO guideline values should not be considered safe for human health. The results also have implications for public health communication. Usually, warning messages are addressed to the population on days with peak levels of air pollution and typically target populations considered to be particularly vulnerable, such as the elderly and children. The messages should be updated to inform the population that pollution can have a significant impact on health even at low concentrations and that health effects can occur even in apparently healthy adults.

6. Conclusion

This study quantifies the healthcare costs caused to the French healthcare system by acute exposure to air pollution. Air pollution remains the greatest environmental risk to the health of Europeans. Air quality standards and targets have been set for a number of air pollutants, but the appropriateness of these limits remains the subject of debate and the object of recent policy changes. Accurately quantifying the effects of air pollution exposure is essential to determine the optimal level of environmental policy.

I combine comprehensive French administrative health data for a nationally representative sample with high-resolution geospatial data on air pollution and meteorological conditions to estimate the health costs of air pollution more accurately and comprehensively than previous studies, which tend to be limited to relatively narrow geographical areas and time periods, look at only a specific part of the population, or examine the effects of air pollution on a limited range of health conditions. Using high-quality data from a nationally representative sample also makes it possible to analyse treatment effect heterogeneity as a function of patient and location characteristics in a way that has not been possible in previous studies based on non-representative samples. To estimate causal effects, I adopt an instrumental variable approach that exploits weekly variations in local concentrations of nitrogen dioxide, ground-level ozone and particulate matter caused by variations in altitude weather conditions. Weather conditions at altitude are good instruments because they are highly predictive of air pollution concentrations and are unlikely to be associated with changes in health care use other than through their effect on air pollution, conditional on controlling for various time and location fixed effects as well as weather at ground level.

Exposure to air pollution concentrations that are predominantly below the current European air quality standard values causes healthcare costs to the French health system in the order of several billions a year. The costs are about 10 times higher than those estimated in previous studies, suggesting that the health costs of air pollution have been severely underestimated. Air pollution causes health costs in all age groups, suggesting that adverse health effects also occur in parts of the population that were considered less vulnerable and were less frequently studied. While air pollution in more populated urban areas has a large impact on total health expenditure, pollution in relatively cleaner and less populated areas appears to have a larger marginal effect on healthcare costs. These results are consistent with a concave relationship between air pollution exposure and health outcomes, with pollution having larger marginal health effects at low concentrations.

These results are highly relevant for environmental policy. Previous policy decisions have been based on estimates that do not account for health cost savings of several billion per year and should be updated. Significant health costs are caused by air pollution levels that are below current European air quality guideline values, suggesting that the guideline values are set too high. The apparent concave relationship between air pollution and health costs means that reducing population exposure, even at low air pollution concentrations, should be an important public health objective.

References

Ahrens, A., Hansen, C.B., Schaffer, M.E., 2020. lassopack: Model selection and prediction with regularized regression in stata. The Stata Journal 20, 176–235.

Amann, M., Holland, M., Maas, R., Vandyck, T., Saveyn, B., 2017. Costs, benefits and economic impacts of the eu clean air strategy and their implications on innovation and competitiveness. IIASA Report; International Institute for Applied Systems Analysis (IIASA): Laxenburg, Austria , 1–59.

Anderson, M.L., 2015. As the wind blows: The effects of long-term exposure to air pollution on mortality .

Arceo, E., Hanna, R., Oliva, P., 2016. Does the effect of pollution on infant mortality differ between developing and developed countries? evidence from mexico city. The Economic Journal 126, 257–280.

Bagilet, V., Zabrocki, L., 2021. Inference design in studies on acute health effects of air pollution .

Barwick, P.J., Li, S., Rao, D., Zahur, N.B., 2024. The healthcare cost of air pollution: Evidence from the world's largest payment network. Review of Economics and Statistics , 1–52.

Belloni, A., Chen, D., Chernozhukov, V., Hansen, C., 2012. Sparse models and methods for optimal instruments with an application to eminent domain. Econometrica 80, 2369–2429.

Bezin, J., Duong, M., Lassalle, R., Droz, C., Pariente, A., Blin, P., Moore, N., 2017. The national healthcare system claims databases in france, sniiram and egb: powerful tools for pharmacoepidemiology. Pharmacoepidemiology and drug safety 26, 954–962.

Birnbaum, H.G., Carley, C.D., Desai, U., Ou, S., Zuckerman, P.R., 2020. Measuring the impact of air pollution on health care costs: Study examines the impact of air pollution on health care costs. Health Affairs 39, 2113–2119.

Calderón-Garcidueñas, L., Franco-Lira, M., Torres-Jardon, R., Henriquez-Roldan, C., Barragán-Mejía, G., Valencia-Salazar, G., Gonzalez-Maciel, A., Reynoso-Robles, R., Villarreal-Calderón, R., Reed, W., 2007. Pediatric respiratory and systemic effects of chronic air pollution exposure: nose, lung, heart, and brain pathology. Toxicologic pathology 35, 154–162.

Chen, Y., Ebenstein, A., Greenstone, M., Li, H., 2013. Evidence on the impact of sustained exposure to air pollution on life expectancy from china's huai river policy. Proceedings of the National Academy of Sciences 110, 12936–12941.

Clapp, L.J., Jenkin, M.E., 2001. Analysis of the relationship between ambient levels of o3, no2 and no as a function of nox in the uk. Atmospheric Environment 35, 6391–6405.

CNAM, 2015-2018. Système national des données de santé (snds). `https://www.health-data-hub.fr/`. Accessed: 2020-01-01.

Colmer, J., Lin, D., Liu, S., Shimshack, J., 2021. Why are pollution damages lower in developed countries? insights from high-income, high-particulate matter hong kong. Journal of Health Economics 79, 102511.

Currie, J., Neidell, M., 2005. Air pollution and infant health: what can we learn from california's recent experience? The Quarterly Journal of Economics 120, 1003–1030.

Currie, J., Walker, R., 2011. Traffic congestion and infant health: Evidence from e-zpass. American Economic Journal: Applied Economics 3, 65–90.

data.gouv.fr, 2014. Fond de carte des codes postaux. https://www.data.gouv.fr/fr/datasets/fond-de-carte-des-codes-postaux/#/information. Accessed: 2025-01-15.

data.gouv.fr, 2024. Jours fériés en france. https://www.data.gouv.fr/en/datasets/jours-feries-en-france/. Accessed: 2021-01-01.

Dechezleprêtre, A., Rivers, N., Stadler, B., 2019. The economic cost of air pollution: Evidence from europe .

Deryugina, T., Heutel, G., Miller, N.H., Molitor, D., Reif, J., 2019. The mortality and medical costs of air pollution: Evidence from changes in wind direction. American Economic Review 109, 4178–4219.

Deschênes, O., Greenstone, M., Shapiro, J.S., 2017. Defensive investments and the demand for air quality: Evidence from the nox budget program. American Economic Review 107, 2958–89.

Deschenes, O., Moretti, E., 2009. Extreme weather events, mortality, and migration. The Review of Economics and Statistics 91, 659–681.

DREES, 2020. Les dépenses de santé en 2019 - résultats des comptes de la santé. https://drees.solidarites-sante.gouv.fr/publications-documents-de-reference/panoramas-de-la-drees/les-depenses-de-sante-en-2019-resultats.

DSS, 2023. Rapports d'évaluation des politiques de sécurité sociale.

EEA, 2013. Every breath we take - improving air quality in europe. https://www.eea.europa.eu/publications/eea-signals-2013. Accessed: 2023-10-07.

EEA, 2015-2018. Air quality e-reporting. https://www.eea.europa.eu/en/datahub/datahubitem-view/3b390c9c-f321-490a-b25a-ae93b2ed80c1?activeAccordion=1084370%2C1084350%2C1084338%2C1084391%2C1084346. Accessed: 2024-03-01.

EEA, 2020. Air pollution: how it affects our health. https://www.eea.europa.eu/themes/air/health-impacts-of-air-pollution. Accessed: 2021-07-07.

EPA, 2024. What is ozone. https://www.epa.gov/ozone-pollution-and-your-patients-health/what-ozone. Accessed: 2024-01-18.

European Commission, 2016. Proposal for a directive of the european parliament and of the council on ambient air quality and cleaner air for europe (recast), com/2022/542 final. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM:2022:542:FIN.

European Parliament, 2016. Directive (eu) 2016/2284 of the european parliament and of the council of 14 december 2016 on the reduction of national emissions of certain atmospheric pollutants, amending directive 2003/35/ec and repealing directive 2001/81. https://eur-lex.europa.eu/eli/dir/2016/2284/oj.

Fontaine, A., Bonvalot, Y., Lim, T.A., Duée, M., Pernelet-Joly, V., Thuret, A., 2007. Impacts économiques des pathologies liées à la pollution .

Godzinski, A., Castillo, M.S., 2021. Disentangling the effects of air pollutants with many instruments. Journal of Environmental Economics and Management 109, 102489.

Godzinski, A., Suarez Castillo, M., 2019. Short-term health effects of public transport disruptions: air pollution and viral spread channels .

Guidetti, B., Pereda, P., Severnini, E.R., 2024. Health shocks under hospital capacity constraint: Evidence from air pollution in sao paulo, brazil. Technical Report. National Bureau of Economic Research.

Halliday, T.J., Lynham, J., De Paula, Á., 2019. Vog: Using volcanic eruptions to estimate the health costs of particulates. The Economic Journal 129, 1782–1816.

Henderson, S.B., Environmental, B., Nguyen, P.D., 2024. The public health paradox of wildfire smoke. Province-wide implementation of the Vancouver Chest Pain Rule , 93.

Hersbach, H., Bell, B., Berrisford, P., Biavati, G., Horányi, A., Muñoz Sabater, J., Nicolas, J., Peubey, C., Radu, R., Rozum, I., Schepers, D., Simmons, A., Soci, C., Dee, D., Thépaut, J.N., 2023a. Era5 hourly data on pressure levels from 1940 to present. DOI: 10.24381/cds.bd0915c6. Accessed: 2023-01-01.

Hersbach, H., Bell, B., Berrisford, P., Biavati, G., Horányi, A., Muñoz Sabater, J., Nicolas, J., Peubey, C., Radu, R., Rozum, I., Schepers, D., Simmons, A., Soci, C. andDee, D., Thépaut, J.N., 2023b. Era5 hourly data on single levels from 1940 to present. https://cds.climate.copernicus.eu/cdsapp#!/dataset/reanalysis-era5-single-levels?tab=overview. Accessed: 2023-01-01.

Holm, S., 1979. A simple sequentially rejective multiple test procedure. Scandinavian Journal of Statistics 6, 65–70.

Imbens, G., Angrist, J., 1994. Identification and estimation of local average treatment effects. Econometrica 62, 467–475.

INERIS, 2015-2018. Historical reconstruction of background air pollution over france. Accessed: 2020-01-01.

INERIS, 2024. La qualité de l'air en france métropolitaine cartographiée de 2000 à aujourd'hui par l'ineris. https://www.ineris.fr/fr/recherche-appui/risques-chroniques/mesure-prevision-qualite-air/qualite-air-france-metropolitaine.

INSEE, 2020. Les comptes de la nation en 2019. https://www.insee.fr/fr/statistiques/4500483.

INSEE, 2024. Revenus et pauvreté des ménages aux niveaux national et local - revenus localisés sociaux et fiscaux. https://www.data.gouv.fr/fr/datasets/revenus-et-pauvrete-des-menages-aux-niveaux-national-et-local-revenus-localises-sc Accessed: 2021-01-01.

Institute for Health Metrics and Evaluation, 2020. Global burden of disease study 2019. https://www.healthdata.org/research-analysis/gbd.

Ito, K., Zhang, S., 2020. Willingness to pay for clean air: Evidence from air purifier markets in china. Journal of Political Economy 128, 1627–1672.

Jans, J., Johansson, P., Nilsson, J.P., 2018. Economic status, air quality, and child health: Evidence from inversion episodes. Journal of health economics 61, 220–232.

Jayachandran, S., 2009. Air quality and early-life mortality evidence from indonesia's wildfires. Journal of Human resources 44, 916–954.

Johns, D.O., Stanek, L.W., Walker, K., Benromdhane, S., Hubbell, B., Ross, M., Devlin, R.B., Costa, D.L., Greenbaum, D.S., 2012. Practical advancement of multipollutant scientific and risk assessment approaches for ambient air pollution. Environmental health perspectives 120, 1238–1242.

Klauber, H., Holub, F., Koch, N., Pestel, N., Ritter, N., Rohlf, A., 2024. Killing prescriptions softly: Low emission zones and child health from birth to school. American Economic Journal: Economic Policy 16, 220–248.

Knittel, C.R., Miller, D.L., Sanders, N.J., 2016. Caution, drivers! children present: Traffic, pollution, and infant health. Review of Economics and Statistics 98, 350–366.

Le Moullec, A., Meleux, F., 2019. Bilan de la qualité de l'air extérieur en france en 2017. Rapport Commissariat Général au Développement Durable, Ministère de la Transition Ecologique et Solidaire .

Lee, H.J., Chang, L.S., Jaffe, D.A., Bak, J., Liu, X., Abad, G.G., Jo, H.Y., Jo, Y.J., Lee, J.B., Kim, C.H., 2021. Ozone continues to increase in east asia despite decreasing no2: Causes and abatements. Remote Sensing 13, 2177.

Leiser, C.L., Hanson, H.A., Sawyer, K., Steenblik, J., Al-Dulaimi, R., Madsen, T., Gibbins, K., Hotaling, J.M., Ibrahim, Y.O., VanDerslice, J.A., et al., 2019. Acute effects of air pollutants on spontaneous pregnancy loss: a case-crossover study. Fertility and sterility 111, 341–347.

Levi, Y., Dayan, U., Levy, I., Broday, D.M., et al., 2020. On the association between characteristics of the atmospheric boundary layer and air pollution concentrations. Atmospheric Research 231, 104675.

Maguire, C., Asquith, M., Lung, T., Viaud, V., 2020. The European Environment-state and outlook 2020. European Environment Agency.

Manisalidis, I., Stavropoulou, E., Stavropoulos, A., Bezirtzoglou, E., 2020. Environmental and health impacts of air pollution: A review. Frontiers in public health 8.

Mauderly, J.L., Burnett, R.T., Castillejos, M., Özkaynak, H., Samet, J.M., Stieb, D.M., Vedal, S., Wyzga, R.E., 2010. Is the air pollution health research community prepared to support a multipollutant air quality management framework? Inhalation toxicology 22, 1–19.

Miller, N.H., Molitor, D., Zou, E.Y., 2021. A causal concentration-response function for air pollution: Evidence from wildfire smoke *. URL: https://api.semanticscholar.org/CorpusID:247961998.

Mogstad, M., Torgovitsky, A., Walters, C.R., 2021. The causal interpretation of two-stage least squares with multiple instrumental variables. American Economic Review 111, 3663–3698.

Moretti, E., Neidell, M., 2011. Pollution, health, and avoidance behavior evidence from the ports of los angeles. Journal of human Resources 46, 154–175.

Muñoz Sabater, J., 2019. Era5-land hourly data from 1950 to present. DOI: 10.24381/cds.e2161bac. Accessed: 2020-01-01.

Neidell, M., 2004. Air pollution, health, and socio-economic status: the effect of outdoor air quality on childhood asthma. Journal of health economics 23, 1209–1236.

Neidell, M., 2009. Information, avoidance behavior, and health the effect of ozone on asthma hospitalizations. Journal of Human resources 44, 450–478.

Papanicolas, I., Woskie, L.R., Jha, A.K., 2018. Health care spending in the united states and other high-income countries. Jama 319, 1024–1039.

Rafenberg, C., 2015. Estimation des coûts pour le système de soins français de cinq maladies respiratoires et des hospitalisations attribuables à la pollution de l'air. Etud Doc 122, 36.

Real, E., Couvidat, F., Ung, A., Malherbe, L., Raux, B., Gressent, A., Colette, A., 2022. Historical reconstruction of background air pollution over france for 2000–2015. Earth System Science Data 14, 2419–2443.

Rohlf, A., Holub, F., Koch, N., Ritter, N., 2020. The effect of clean air on pharmaceutical expenditures. Economics Letters 192, 109221.

Schlenker, W., Walker, W.R., 2016. Airports, air pollution, and contemporaneous health. The Review of Economic Studies 83, 768–809.

Schwartz, J., Bind, M.A., Koutrakis, P., 2017. Estimating causal effects of local air pollution on daily deaths: effect of low levels. Environmental health perspectives 125, 23–29.

Schwartz, J., Fong, K., Zanobetti, A., 2018. A national multicity analysis of the causal effect of local pollution, no 2, and pm 2.5 on mortality. Environmental health perspectives 126, 087004.

Sénat, 2015. Rapport fait au nom de la commission d'enquête sur le coût économique et financier de la pollution de l'air. http://www.senat.fr/rap/r14-610-1/r14-610-11.pdf.

Simeonova, E., Currie, J., Nilsson, P., Walker, R., 2019. Congestion pricing, air pollution, and children's health. Journal of Human Resources , 0218–9363R2.

Tuppin, P., De Roquefeuil, L., Weill, A., Ricordeau, P., Merlière, Y., 2010. French national health insurance information system and the permanent beneficiaries sample. Revue d'epidemiologie et de sante publique 58, 286–290.

Vedal, S., Kaufman, J.D., 2011. What does multi-pollutant air pollution research mean?

Wu, X., Lu, Y., Zhou, S., Chen, L., Xu, B., 2016. Impact of climate change on human infectious diseases: Empirical evidence and human adaptation. Environment international 86, 14–23.

Zivin, J.G., Neidell, M., 2009. Days of haze: Environmental information disclosure and intertemporal avoidance behavior. Journal of Environmental Economics and Management 58, 119–128.

Acknowledgements

This research is supported by the Strategic Research Initiative NutriPerso from University Paris-Saclay and the project AlimaSSen [ANR-14-CE20-0003-01] from the French National Research Agency (ANR). Support by the German Research Foundation (DFG) through CRC TR 224 (Project B07) and through Germany's Excellence Strategy (EXC 2126/1-390838866) is gratefully acknowledged.
I am most grateful to the French National Institute for Industrial Environment and Risks (INERIS) for providing the data on air pollution and to the French National Fund for Health Insurance (CNAM) for access to the French administrative health data. The author has no conflict of interest to declare.

Appendix

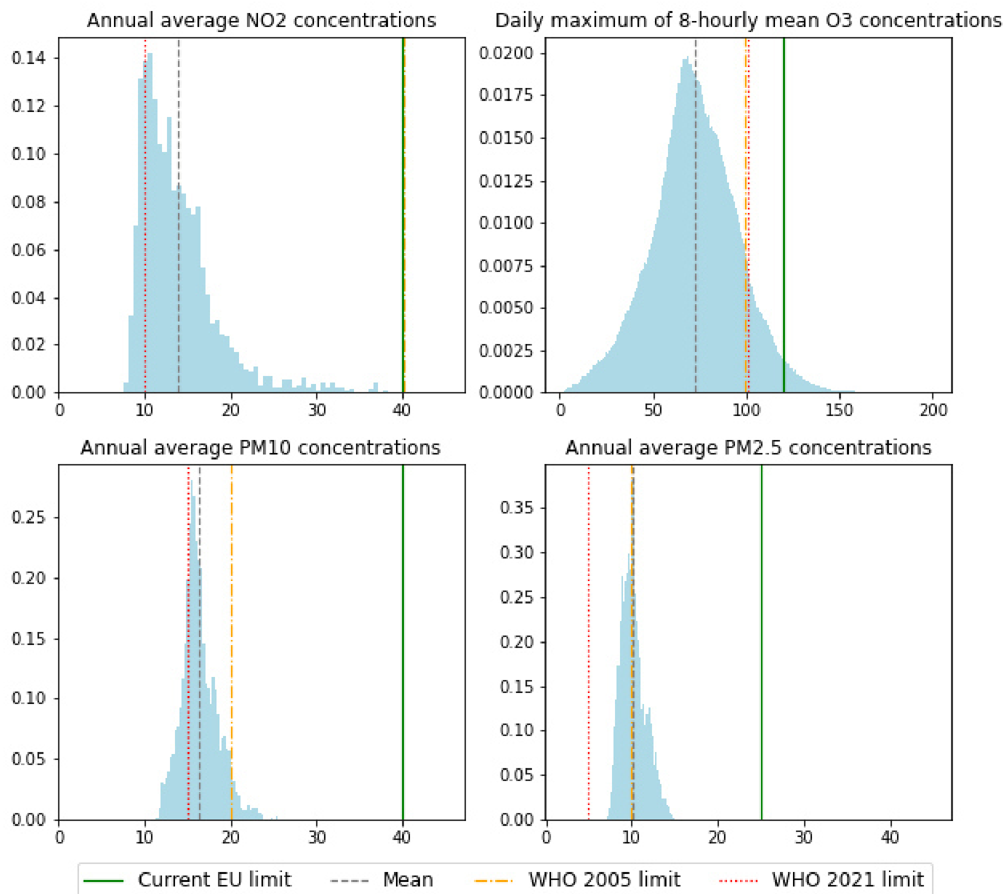

Figure A1: Level of pollutants relative to the limit values presented in Table OA1. The Figure shows the distribution of pollution concentrations in 2017 across zip code areas in light blue, the current limit values in France/the EU (solid green), the average value across zip codes (grey, dashed) and the WHO 2005 and 2021 guideline values (yellow dashdot and red dotted, respectively). Pollution levels in France are generally below the prevailing EU air quality guideline values.

Table A1: Summary statistics - pooled postcode-week observations

Variable	Mean	Std. Dev.	Min.	Max.	N
Sum of healthcare spending in €					
Total spending	3608.63	7471.21	0	370436.63	1,257,984
Family practice	1212.08	2502.24	0	75846.8	1,257,984
Cardiology and vascular medicine	50.9	184.84	0	37266.45	1,257,984
Pulmonology	22.76	142.19	0	15688.04	1,257,984
Otorhinolaryngology	19.34	74.56	0	10203.16	1,257,984
Ophtalmology	82.36	228.2	0	9585.98	1,257,984
Gynecology	43.19	150.61	0	9300.96	1,257,984
Nephrology	11.48	79.62	0	11234.26	1,257,984
Gastroenterology and hepatology	32.37	305.67	0	26562.06	1,257,984
Trauma surgery	36.01	164.33	0	14695.7	1,257,984
Plastic surgery	5.22	76.64	0	6468.49	1,257,984
Pollution concentrations in $\mu g/m^3$					
NO_2	13.78	7.14	2.68	76.84	1,247,428
O_3	55.7	17.49	0.4	119.09	1,247,428
PM_{10}	16.61	6.32	4.28	92.76	1,247,428
$PM_{2.5}$	10.57	5.75	1.97	79.67	1247428
NO_2 (data from EEA)	21.72	5.58	12.44	41.56	1,257,984
O_3 (data from EEA)	54.55	15.78	19.68	89.71	1,257,984
PM_{10} (data from EEA)	18.76	5.37	10.69	52.76	1,257,984
SO_2 (data from EEA)	2.22	0.39	1.4	3.91	1,257,984
CO (data from EEA)	0.23	0.1	0.06	0.55	1,257,984
Altitude atmospheric conditions					
Thermal inversions (sum of occurrence)	0.36	0.94	0	7	1,257,984
Temperature difference (mean, C°)	-1.2	0.47	-3.56	3.06	1,257,984
Planetary boundary height (mean, m)	539.85	186.81	12.05	1590.81	1,257,984
Wind speed at altitude ((mean, m/s)	12.52	4.52	2.89	32.6	1,257,984
Ground-level meteorological conditions					
Precipitation (sum, mm)	14.11	16.44	0	297.2	1,257,984
Wind speed (mean, m/s)	2.52	1.12	0.08	11.71	1,214,512
Temperature (mean, C°)	12.52	6.37	-12.43	32.09	1,257,984
Relative humidity (mean, in %)	76.11	9.58	33.78	97.86	1,257,984
Postcode-level characteristics					
Household income	22096.28	4050.53	7910	52670	1,251,536
Population size	10377.76	15258.49	0	236715	1,257,984
Holidays (nb. of days)	0.21	0.44	0	2	1257984

The table presents summary statistics of the pooled postcode-week observations. The information on healthcare spending come from the French National System of Health Data (SNDS) provided by the National Fund for Health Insurance (CNAM, 2015-2018). The data sources for pollutant concentrations are reanalysis data provided by the French National Institute for Industrial Environment and Risks (INERIS, 2015-2018) and measurement station data provided by the European Environment Agency (EEA, 2015-2018). Information on atmospheric and meteorological conditions come from ERA5 reanalysis data produced by the Copernicus Climate Change Service (C3S, Muñoz Sabater (2019); Hersbach et al. (2023a,b)). Information on household income comes from the FiLoSoFi social and fiscal localized database provided by the French National Institute of Statistics and Economic Studies (INSEE, INSEE (2024)). The data on population size is included in the French shapefile data available at the Open platform for French public data (data.gouv.fr, 2014). Data on holidays in France are also obtained from the Open platform (data.gouv.fr, 2024).

Online appendix

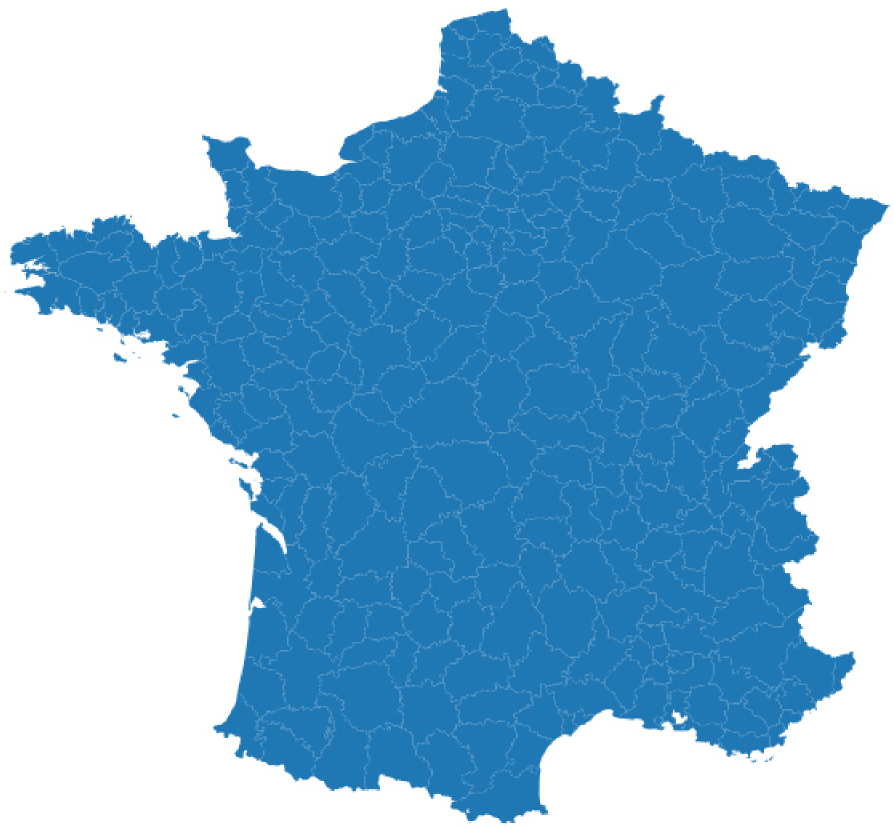

Figure OA1: Division of France into employment zones. The employment zone (*"zone d'emploi"* in French) is a division of the French territory into 306 geographical areas within which most of the working population resides and works.

Table OA1: Summary of the main WHO and EU Air Quality Standard values

Pollutant	Averaging time	WHO 2005 Guidelines	WHO 2021 Guidelines	EU/France limit values
Nitrogen dioxide (NO_2)	Annual	40	10	40
Nitrogen dioxide (NO_2)	24-hour	N/A	25	N/A
Ozone (O_3)	8-hour	100	100	120
Ozone (O_3)	Peak season[a]	N/A	60	N/A
Particles $\varnothing \leq 10\mu m$ (PM_{10})	Annual	20	15	40
Particles $\varnothing \leq 10\mu m$ (PM_{10})	24-hour	50	45	50
Particles $\varnothing \leq 2.5\mu m$ ($PM_{2.5}$)	Annual	10	5	25
Particles $\varnothing \leq 2.5\mu m$ ($PM_{2.5}$)	24-hour	25	15	N/A

The table presents a summary of the main World Health Organisation (WHO) and European Union (EU) air quality standard values. Guidelines and limit values are expressed in $\mu g/m^3$.

[a] Average of daily maximum 8-hour mean O_3 concentration in the six consecutive months with the highest six-month running- average O_3 concentration.

Sources: WHO, `https://www.who.int/news-room/feature-stories/detail/what-are-the-who-air-quality-guidelines`

Airparif, `https://www.airparif.asso.fr/la-reglementation-en-france`

Table OA2: First stage regression regression, preferred FE-IV model specification

	NO_2	O_3	PM_{10}
Thermal inversion (nb. h per week)	0.176***	0.0126	0.347***
	(0.009)	(0.021)	(0.012)
TI 0-4 h (nb. h per week)	0.0953***	0.0124**	0.189***
	(0.002)	(0.004)	(0.004)
TI 4-8 h (nb. h per week)	-0.0416***	0.119***	-0.0567***
	(0.002)	(0.005)	(0.004)
TI 8-12 h (nb. h per week)	-0.0759***	-0.425***	-0.0397***
	(0.005)	(0.010)	(0.006)
TI 12-16 h (nb. h per week)	0.201***	-0.663***	0.319***
	(0.007)	(0.019)	(0.009)
TI 16-20 h (nb. h per week)	0.0764***	-0.282***	0.155***
	(0.006)	(0.014)	(0.008)
TI 20-24 h (nb. h per week)	0.0630***	0.163***	0.0142***
	(0.002)	(0.006)	(0.004)
TI strength 0-4 h (diff degree C)	1.445***	-0.225**	0.0965*
	(0.034)	(0.072)	(0.047)
TI strength 4-8 h (diff degree C)	-0.842***	-1.500***	0.641***
	(0.032)	(0.066)	(0.036)
TI strength 8-12 h (diff degree C)	-1.222***	1.657***	-1.571***
	(0.045)	(0.104)	(0.043)
TI strength 12-16 h (diff degree C)	1.905***	-6.413***	3.612***
	(0.050)	(0.133)	(0.069)
TI strength 16-20 h (diff degree C)	-0.138*	-1.503***	-0.776***
	(0.061)	(0.153)	(0.080)
TI strength 20-24 h (diff degree C)	0.765***	0.455***	1.084***
	(0.034)	(0.073)	(0.046)
PBLH 0-4 h (m)	0.0000389	0.0114***	-0.00636***
	(0.000)	(0.000)	(0.000)
PBLH 4-8 h (m)	-0.00327***	-0.00595***	0.00115***
	(0.000)	(0.000)	(0.000)
PBLH 8-12 h (m)	-0.00284***	0.00266***	-0.00370***

(Continued on next page)

Table OA2: (continued) First stage regression, preferred FE-IV model specification

	NO_2	O_3	PM_{10}
	(0.000)	(0.000)	(0.000)
PBLH 12-16 h (m)	0.00108***	0.0192***	-0.000876***
	(0.000)	(0.000)	(0.000)
PBLH 16-20 h (m)	-0.00254***	-0.00219***	0.000179**
	(0.000)	(0.000)	(0.000)
PBLH 20-24 h (m)	-0.00420***	0.00310***	0.00263***
	(0.000)	(0.000)	(0.000)
Wind speed at 350 hPa (m/s)	0.0608***	-0.638***	-0.212***
	(0.005)	(0.013)	(0.008)
Wind speed at 400 hPa (m/s)	-0.254***	0.997***	-0.0726***
	(0.012)	(0.031)	(0.019)
Wind speed at 450 hPa (m/s)	0.182***	-1.117***	0.217***
	(0.016)	(0.041)	(0.027)
Wind speed at 500 hPa (m/s)	0.0279	1.713***	0.0238
	(0.018)	(0.055)	(0.027)
Wind speed at 550 hPa (m/s)	0.122***	-1.852***	0.150***
	(0.019)	(0.061)	(0.032)
Wind speed at 600 hPa (m/s)	-0.00984	1.520***	0.0843*
	(0.023)	(0.061)	(0.039)
Wind speed at 650 hPa (m/s)	-0.738***	-0.526***	-1.193***
	(0.025)	(0.068)	(0.045)
Wind speed at 700 hPa (m/s)	0.774***	0.168*	1.244***
	(0.025)	(0.071)	(0.045)
Wind speed at 750 hPa (m/s)	-0.166***	-1.422***	-0.651***
	(0.028)	(0.076)	(0.050)
Wind speed at 800 hPa (m/s)	-0.965***	2.390***	-0.0596
	(0.054)	(0.149)	(0.114)
Wind speed at 825 hPa (m/s)	1.285***	-3.477***	0.288*
	(0.063)	(0.180)	(0.142)
Wind speed at 850 hPa (m/s)	-0.309***	2.437***	-0.0563
	(0.028)	(0.079)	(0.064)
Constant	13.79***	65.18***	18.48***

(Continued on next page)

Table OA2: (continued) First stage regression, preferred FE-IV model specification

	NO_2	O_3	PM_{10}
	(0.127)	(0.305)	(0.169)
Observations	1209572	1209572	1209572

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. The table shows the first stage regression corresponding to the FE-IV models in Columns 3 and 4 in table 1. The instruments are the number of hours of thermal inversion (TI) per week and the number of hours of thermal inversions per week by moment of the day, the average strength of these thermal inversions in terms of weekly average temperature difference between the lowest and second lowest atmospheric layer by moment of the day, the average of the planetary boundary layer height (PBLH) taken over the moments of the day, and the weekly average wind speed at twelve pressure levels. All regressions include month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level in parenthesis.

Table OA3: First stage using LASSO selected instruments

	NO_2	O_3	PM_{10}
Thermal inversion (nb. h per week)	0.294***		0.339***
	(0.009)		(0.012)
TI 0-4 h (nb. h per week)	0.0809***	0.00256	0.186***
	(0.002)	(0.004)	(0.003)
TI 4-8 h (nb. h per week)	0.00956***	0.160***	-0.0548***
	(0.002)	(0.005)	(0.003)
TI 8-12 h (nb. h per week)	-0.0309***	-0.607***	-0.0488***
	(0.005)	(0.010)	(0.006)
TI 12-16 h (nb. h per week)		-0.658***	0.278***
		(0.020)	(0.009)
TI 16-20 h (nb. h per week)		-0.190***	0.208***
		(0.014)	(0.006)
TI 20-24 h (nb. h per week)	0.0227***	0.157***	
	(0.002)	(0.005)	
TI strength 0-4 h (diff degree C)	1.019***	-0.419***	
	(0.021)	(0.068)	
TI strength 4-8 h (diff degree C)		-0.894***	0.767***
		(0.058)	(0.026)
TI strength 8-12 h (diff degree C)	-1.078***		-1.756***
	(0.048)		(0.043)
TI strength 12-16 h (diff degree C)	0.882***	-6.229***	3.266***
	(0.026)	(0.114)	(0.060)
TI strength 20-24 h (diff degree C)			0.793***
			(0.021)
PBLH 0-4 h (m)		0.0120***	-0.00535***
		(0.000)	(0.000)
PBLH 4-8 h (m)	-0.00369***	-0.00626***	
	(0.000)	(0.000)	
PBLH 8-12 h (m)	-0.00196***	0.00249***	-0.00317***
	(0.000)	(0.000)	(0.000)
PBLH 12-16 h (m)		0.0182***	-0.000915***

(Continued on the next page)

Table OA3: (continued) First stage using LASSO selected instruments

	NO_2	O_3	PM_{10}
		(0.000)	(0.000)
PBLH 16-20 h (m)	-0.00156***		
	(0.000)		
PBLH 20-24 h (m)	-0.00433***	0.00168***	0.00265***
	(0.000)	(0.000)	(0.000)
Wind speed at 350 hPa (m/s)	-0.0156***	-0.257***	-0.117***
	(0.001)	(0.003)	(0.002)
Wind speed at 500 hPa (m/s)		0.391***	
		(0.006)	
Wind speed at 650 hPa (m/s)	-0.126***		-0.143***
	(0.002)		(0.003)
Wind speed at 750 hPa (m/s)		-0.863***	
		(0.015)	
Wind speed at 850 hPa (m/s)	0.144***	0.879***	
	(0.006)	(0.021)	
Observations	1209572	1209572	1209572

***$p < 0.001$, **$p < 0.01$, *$p < 0.05$. The table presents the first stage regression results where each pollutant is regressed over the LASSO-selected variables. The instruments are the number of hours of thermal inversion (TI) per week and the number of hours of thermal inversions per week by moment of the day, the average strength of these thermal inversions in terms of weekly average temperature difference between the lowest and second lowest atmospheric layer by moment of the day, the average of the planetary boundary layer height (PBLH) taken over the moments of the day, and the weekly average wind speed at twelve pressure levels. All regressions include month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level in parentheses.

Table OA4: First stage model fit in terms of the Bayesian Information Criterion (BIC) and the Akaike Information Criterion (AIC)

	First stage regression - estimating NO_2 pollution		
	Use NO_2 instruments	Use O_3 instruments	Use PM_{10} instruments
AIC	**5854194.7**	5857163.3	5864386.1
BIC	**5854879.1**	5857883.6	5865082.4
	First stage regression AIC and BIC estimating O_3 pollution		
	Use NO_2 instruments	Use O_3 instruments	Use PM_{10} instruments
AIC	8074782.3	**7958247.7**	7974813.1
BIC	8075466.6	**7958968.0**	7975509.4
	First stage regression AIC and BIC estimating PM_{10} pollution		
	Use NO_2 instruments	Use O_3 instruments	Use PM_{10} instruments
AIC	6932516.4	6925064.4	**6919790.2**
BIC	6933200.8	6925784.7	**6920486.6**

The table compares the model fit in terms of the Bayesian Information Criterion (BIC) and the Akaike Information Criterion (AIC) when the pollutant-specific instruments predict the pollutant for which they were selected with the model fit when these instruments are used to predict the concentrations of the other pollutants. The terms in bold show the model fit for models where the pollutant-specific instruments predict the pollutant for which they have been selected, which corresponds to the best model fit (lowest AIC and BIC).

Table OA5: Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly healthcare expenditure - FE-IV LASSO regression results

	Weekly healthcare spending	
	(1)	(2)
Weekly mean NO_2	20.40^{***}	20.18^{***}
	(3.881)	(3.750)
Weekly mean O_3	6.177^{***}	3.296^{***}
	(0.783)	(0.666)
Weekly mean PM_{10}	10.75^{***}	1.519
	(2.839)	(2.842)
Lag weekly mean NO_2		-6.877
		(4.134)
Lag weekly mean O_3		7.033^{***}
		(0.814)
Lag weekly mean PM_{10}		23.10^{***}
		(2.724)
Observations	1,209,572	1,186,311

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. The table presents the results for the FE-IV regressions using LASSO selected pollutant-specific instrument vectors for the first stage regression. The corresponding first stage regression results are presented in table OA3. All regressions include month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level in parenthesis.

Table OA6: Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly healthcare expenditure over a time window of several weeks

	Sum of healthcare spending		
	Same week	Over 2 weeks	Over 3 weeks
Weekly mean NO_2	17.23***	25.63***	54.10***
	(3.719)	(5.503)	(7.602)
Weekly mean O_3	3.275***	2.692**	8.057***
	(0.662)	(1.009)	(1.675)
Weekly mean PM_{10}	3.540	6.277	23.63***
	(2.843)	(4.356)	(5.824)
Lag weekly mean NO_2	-3.423	1.672	-29.71***
	(4.062)	(6.089)	(8.766)
Lag weekly mean O_3	6.497***	8.627***	16.55***
	(0.795)	(1.377)	(1.628)
Lag weekly mean PM_{10}	18.14***	41.77***	69.78***
	(2.616)	(3.966)	(5.740)
Observations	1186311	1163050	1139789

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. The table shows results for the effect of weekly air pollution exposure and its one week lag on healthcare expenditure over a longer time window of two to four weeks, controlling for the appropriate number of weather and instrument leads. Column 1 shows results for the baseline model that estimates the effects of weekly average air pollution concentration and its lag on healthcare expenditure during the same week for reference. Column 2 shows results for the effects during the same week and the following week and Column 3 shows the effects for the same week and the following two weeks of healthcare expenditure. The total lag considered is a month for the effect of the one week lag of air pollution (week -1) on healthcare spending during the following three weeks (week 1 to 3) shown in Column 4. All regressions include month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level in parenthesis.

Table OA7: Impact of average daily NO_2, O_3 and PM_{10} pollutant concentrations on daily healthcare expenditure

	Daily healthcare spending	
	(1)	(2)
Daily NO_2	3.488***	4.946***
	(0.488)	(0.522)
Daily O_3	0.429**	0.692***
	(0.136)	(0.119)
Daily PM_{10}	2.059**	-0.0566
	(0.676)	(0.466)
One day lag NO_2		-2.036***
		(0.481)
Two day lag NO_2		0.957*
		(0.450)
One day lag O_3		-0.417**
		(0.133)
Two day lag O_3		0.813***
		(0.150)
One day lag PM_{10}		2.447***
		(0.530)
Two day lag PM_{10}		-1.343**
		(0.437)
Observations	8484329	8484121
First-stage F-stat	2068.6	2328.7

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. The table shows results for the effect of average daily pollutant concentrations on daily healthcare expenditure. All regressions include day-of-the-week, month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level in parenthesis.

Table OA8: Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly healthcare expenditure - single- and two-pollutant models

	Dependent variable: Sum of weekly healthcare spending					
	Panel A: Location FE model, no lags					
	(1)	(2)	(3)	(4)	(5)	(6)
Weekly mean NO_2	30.33***			33.34***		44.33***
	(1.927)			(2.029)		(2.692)
Weekly mean O_3		0.362		4.076***	0.754*	4.189***
		(0.353)		(0.381)	(0.355)	(0.383)
Weekly mean PM_{10}			4.053***		4.251***	-12.06***
			(0.570)		(0.573)	(0.981)
Observations	1209572	1209572	1209572	1209572	1209572	1209572
	Panel B: Location FE-IV model, no lags					
	(1)	(2)	(3)	(4)	(5)	(6)
Weekly mean NO_2	22.71***			32.29***		18.42***
	(1.952)			(2.152)		(3.820)
Weekly mean O_3		0.957		5.984***	5.477***	6.282***
		(0.680)		(0.756)	(0.789)	(0.773)
Weekly mean PM_{10}			16.87***		22.77***	12.37***
			(1.375)		(1.607)	(2.815)
Observations	1209572	1209572	1209572	1209572	1209572	1209572
First-stage F-stat	2648.7	5768.1	3763.8	2648.7	5768.1	2648.7

(Continued on the next page)

Table OA8: (continued) Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly healthcare expenditure - single- and two-pollutant models

	Dependent variable: Sum of weekly healthcare spending					
	Panel C: Location FE model, lags					
	(1)	(2)	(3)	(4)	(5)	(6)
Weekly mean NO_2	27.44***			31.12***		43.36***
	(1.689)			(1.814)		(2.420)
Lag weekly mean NO_2	8.213***			8.322***		8.947***
	(1.518)			(1.580)		(2.119)
Weekly mean O_3		0.939**		4.769***	1.256***	4.837***
		(0.339)		(0.386)	(0.342)	(0.387)
Lag weekly mean O_3		-0.890*		-0.294	-0.760*	-0.175
		(0.351)		(0.362)	(0.355)	(0.364)
Weekly mean PM_{10}			2.665***		3.022***	-13.34***
			(0.598)		(0.604)	(0.996)
Lag weekly mean PM_{10}			2.380***		2.237***	-1.412
			(0.562)		(0.567)	(0.872)
Observations	1186311	1186311	1186311	1186311	1186311	1186311
	Panel D: Location FE-IV model, lags					
	(1)	(2)	(3)	(4)	(5)	(6)
Weekly mean NO_2	15.87***			17.85***		17.23***
	(1.805)			(2.140)		(3.719)
Lag weekly mean NO_2	8.286***			19.98***		-3.423
	(1.873)			(2.082)		(4.062)
Weekly mean O_3		-0.618		3.015***	2.345***	3.275***
		(0.557)		(0.653)	(0.625)	(0.662)
Lag weekly mean O_3		4.699***		6.528***	6.890***	6.497***
		(0.688)		(0.775)	(0.804)	(0.795)
Weekly mean PM_{10}			11.59***		12.85***	3.540
			(1.335)		(1.514)	(2.843)
Lag weekly mean PM_{10}			8.493***		15.48***	18.14***
			(1.242)		(1.485)	(2.616)
Observations	1186311	1186311	1186311	1186311	1186311	1186311
First-stage F-stat	2063.7	6746.1	4417.5	2063.7	6746.1	2063.7

***$p < 0.001$, **$p < 0.01$, *$p < 0.05$. This table shows results for one and two-pollutant models. Panels A and C presents results for the location fixed effects (FE) model while Panels B and D present results for the location fixed effects instrumental variable model (FE-IV). Panels C and D include one week lag of the pollutants. Columns 1 to 3 show results for models including only one pollutants at a time. Columns 4 and 5 show results for two-pollutant models and column 6 shows results for the model including all three pollutants. All regressions include month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level in parenthesis.

Table OA9: Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly healthcare expenditure - instrumenting only one pollutant and including the others as controls

	Weekly healthcare spending		
NO_2	32.69^{***}		
	(3.223)		
O_3		5.482^{***}	
		(0.766)	
PM_{10}			6.853^{**}
			(2.429)
Observations	1209572	1209572	1209572
First-stage F-stat	2016.3	6165.0	2853.5

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. The table shows results for models instrumenting only one pollutant at a time while the other pollutants are included as controls (not shown in the table). All regressions include month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level in parenthesis.

Table OA10: Impact of average weekly NO_2, O_3 and $PM_{2.5}$ pollutant concentrations on weekly healthcare expenditure - model investigating the effect of $PM_{2.5}$ instead of PM_{10}

	Dependent variable: Sum of weekly healthcare spending					
	Panel A: Location FE model					
	(1)	(2)	(3)	(4)	(5)	(6)
Weekly mean $PM_{2.5}$	2.779***	1.365*	3.189***	2.209***	-16.45***	-16.31***
	(0.638)	(0.653)	(0.651)	(0.667)	(1.209)	(1.158)
Lag weekly mean $PM_{2.5}$		4.919***		3.854***		-0.567
		(0.550)		(0.573)		(1.023)
Weekly mean O_3			1.021*	1.905***	5.532***	4.303***
			(0.413)	(0.406)	(0.449)	(0.380)
Lag weekly mean O_3				-1.571***		0.887*
				(0.295)		(0.420)
Weekly mean NO_2					48.12***	42.86***
					(2.826)	(2.393)
Lag weekly mean NO_2						9.236***
						(2.247)
Observations	1209572	1186311	1209572	1186311	1209572	1186311
	Panel B: Location FE-IV model					
	(1)	(2)	(3)	(4)	(5)	(6)
Weekly mean $PM_{2.5}$	17.36***	12.09***	29.78***	17.05***	15.24***	13.03***
	(1.434)	(1.461)	(1.860)	(1.901)	(3.172)	(3.284)
Lag weekly mean $PM_{2.5}$		11.26***		23.29***		24.75***
		(1.435)		(1.801)		(2.868)
Weekly mean O_3			12.43***	7.186***	14.63***	7.868***
			(0.947)	(0.907)	(1.031)	(0.987)
Lag weekly mean O_3				12.31***		11.99***
				(1.066)		(1.151)
Weekly mean NO_2					24.84***	7.287
					(4.398)	(4.274)
Lag weekly mean NO_2						-2.208
						(4.234)
Observations	1209572	1186311	1209572	1186311	1209572	1186311

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. The table reports the main estimates of the relationship between average weekly air pollutant concentrations and weekly healthcare expenditure including $PM_{2.5}$ instead of PM_{10}. The coefficients indicate the increase in average healthcare spending per zip code area for a 1 $\mu g/m^3$ increase in weekly average pollutant concentrations. Panel A presents results for the location fixed effect model (FE) and Panel B shows results for the location fixed effect instrumental variable model (FE-IV) in which altitude atmospheric conditions are used as instruments for the air pollutant concentrations. All regressions include month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level are in parenthesis.

Table OA11: Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly healthcare expenditure - robustness to controlling and instrumenting for CO and SO_2 pollution

	Weekly healthcare spending			
	(1)	(2)	(3)	(4)
NO_2	23.32***	11.69**	22.61***	27.20***
	(4.062)	(3.704)	(5.061)	(5.180)
O_3	7.035***	3.115***	6.707***	4.208***
	(0.797)	(0.689)	(0.860)	(0.785)
PM_{10}	12.64***	6.056*	12.42***	2.252
	(2.818)	(2.776)	(2.815)	(2.935)
Lag NO_2		3.359		-16.56***
		(4.230)		(4.798)
Lag O_3		7.413***		3.923***
		(0.873)		(0.923)
Lag PM_{10}		16.60***		18.48***
		(2.594)		(2.569)
CO			-61.04	-1756.7***
			(294.946)	(257.931)
SO_2			-43.46	-33.12
			(29.125)	(29.695)
Lag CO				1790.6***
				(257.966)
Lag SO_2				73.94*
				(29.274)
CO and SO_2	controlled	controlled	instrumented	instrumented
Observations	1209572	1186311	1209572	1186311
First-stage F-stat	2106.8	1023.9	917.8	845.4

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. The table shows results for the effects of the three main pollutants NO_2, O_3 and PM_{10} including in addition SO_2 and CO pollution concentrations as control variables in columns 1 and 2 and as additional instrumented pollutants in columns 3 and 4. All regressions include month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level in parenthesis.

Table OA12: Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly healthcare expenditure - robustness to using EEA measuring station data

	Weekly healthcare spending			
	(1)	(2)	(3)	(4)
NO_2 (EEA data)	53.55***	50.23***	77.78***	76.44***
	(3.645)	(4.190)	(5.440)	(5.709)
O_3 (EEA data)	10.60***	9.382***	12.86***	11.57***
	(1.188)	(1.204)	(1.281)	(1.246)
PM_{10} (EEA data)	-3.553	-4.618	-3.995	-2.048
	(3.836)	(2.823)	(3.904)	(2.856)
Lag NO_2 (EEA data)		26.41***		27.99***
		(3.157)		(3.882)
Lag O_3 (EEA data)		3.745**		1.269
		(1.303)		(1.342)
Lag PM_{10} (EEA data)		3.994		1.269
		(2.667)		(2.735)
CO			-1671.5***	-1800.9***
			(344.310)	(251.484)
SO_2			-351.5***	-349.7***
			(36.966)	(40.012)
Lag CO				519.5*
				(253.911)
Lag SO_2				-282.5***
				(38.559)
CO and SO_2	controlled	conrolled	instrumented	instrumented
Lagged pollutants	controlled	instrumented	controlled	instrumented
Observations	1191156	1191156	1191156	1191156
First-stage F-stat	2009.5	1909.8	1559.9	1317.1

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. The table shows results of weekly average pollution concentrations on weekly healthcare expenditure for models using EEA measuring station data on NO_2, O_3 and PM_{10} instead of the reanalyses data from INERIS. Columns 1 and 2 show results including the pollutants from the main analyses while columns 3 and 4 also include SO_2 and CO pollution concentrations. A one week lag of the pollutants are included as control variables in columns 1 and 3 (coefficients not shown) or as instrumented variables in columns 2 and 4. All regressions include month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level in parenthesis.

Table OA13: Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly healthcare expenditure - robustness to different first stage specifications

	Weekly healthcare spending					
	(1)	(2)	(3)	(4)	(5)	(6)
NO_2	67.13*	34.96*	14.19**	17.18***	35.60***	26.71***
	(28.411)	(17.679)	(4.843)	(4.706)	(3.950)	(3.419)
O_3	13.10***	5.033*	4.809***	3.223***	5.172***	3.308***
	(3.193)	(1.954)	(0.891)	(0.860)	(0.794)	(0.799)
PM_{10}	-12.91	-14.28	10.42***	3.814	-3.967	-2.776
	(15.795)	(11.717)	(2.698)	(2.583)	(2.376)	(2.214)
Lag NO_2		65.44**		13.65**		23.30***
		(21.345)		(4.781)		(4.180)
Lag O_3		17.09***		6.963***		5.742***
		(2.458)		(0.895)		(0.756)
Lag PM_{10}		-6.946		2.708		-4.354
		(10.971)		(2.536)		(2.572)
Observations	1209572	1186311	1014676	995163	1014676	995163
First stage	Fewer instruments		Fewer instruments and wind speed		Instruments interacted with location FE	

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Columns 1 and 2 show results for models including as instruments only the number of thermal inversions per week, their average strength, average planetary boundary height and average wind speed at the lowest altitude layer above ground-level. Columns 3 and 4 show results for models that include in addition to the instruments in Columns 1 and 2 the weekly average wind direction by 90-degree intervals interacted with zip code fixed effects. Columns 5 and 6 show results for models where all instruments are interacted with location (employment zone) fixed effects.

Table OA14: Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly healthcare expenditure - robustness to different fixed effects structures

	Sum of weekly healthcare spending				
	(1)	(2)	(3)	(4)	(5)
NO_2	14.37***	12.00***	17.36***	19.29*	25.71***
	(4.101)	(3.450)	(3.901)	(7.950)	(3.692)
O_3	7.600***	4.440***	2.760***	-0.370	4.194***
	(0.987)	(0.617)	(0.671)	(0.653)	(0.667)
PM_{10}	6.138*	4.852	2.370	-5.493	-1.476
	(2.803)	(2.787)	(2.927)	(5.031)	(2.593)
Lag NO_2	0.594	0.528	-3.675	7.638	-3.079
	(4.153)	(3.962)	(4.101)	(7.504)	(3.530)
Lag O_3	12.59***	4.328***	7.831***	7.449***	7.022***
	(1.139)	(0.647)	(0.853)	(0.948)	(0.736)
Lag PM_{10}	23.00***	15.78***	18.28***	12.99***	17.77***
	(2.601)	(2.470)	(2.645)	(3.726)	(2.420)
Observations	1186311	1186311	995163	1186311	
First-stage F-stat	876.4	734.7	70.21	999.3	

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. Column 1 shows results for a model including humidity and minimum and maximum temperatures as additional ground-level weather controls. Column 2 and 3 show results for models including weather fixed effects variables partitioned into 5 and 15 bins instead of the 10 bins used in the main model specification. Column 4 shows results for a model including the non-transformed weather variables. Column 5 shows results for a model including month-by-department FE rather than month fixed effects.

Table OA15: Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly healthcare expenditure - placebo regressions using shuffled instruments

	Weekly healthcare spending	
	(1)	(2)
NO_2	37.09	89.04
	(164.970)	(133.743)
O_3	19.51	20.79
	(61.500)	(51.431)
PM_{10}	-96.78	-3.562
	(116.092)	(105.872)
Lag NO_2		206.0
		(135.917)
Lag O_3		91.33
		(54.970)
Lag PM_{10}		-14.70
		(89.382)
Observations	1209572	1186311
First-stage F-stat	0.833	0.619

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. The table presents results for a placebo exercise where the values of the instrumental variables are randomly reshuffled. All regressions include month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level in parenthesis.

Table OA16: Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly healthcare expenditure at the level of the employment zone

	Weekly healthcare spending at employment zone level			
	(1)	(2)	(3)	(4)
NO_2	520.9*	550.8	438.3***	210.4*
	(240.411)	(353.243)	(107.389)	(104.483)
O_3	45.08	7.188	83.68***	81.81***
	(39.424)	(40.754)	(16.960)	(17.488)
PM_{10}	-115.9	-526.2	19.21	25.56
	(206.098)	(467.519)	(52.511)	(59.223)
Lag NO_2		983.6		737.0
		(653.242)		(385.839)
Lag O_3		145.9***		143.6***
		(42.129)		(33.353)
Lag PM_{10}		-239.5		-44.42
		(244.924)		(141.435)
Observations	59696	58548	59696	58548
First-stage F-stat	320.2	262.7		

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$.

This table shows results for analyses at the employment zone level instead of the zip-code level. Columns 1 and 2 show the results for models using as instruments the vector of altitude atmospheric conditions from the main specification. Columns 3 and 4 show the results for models using fewer instruments, including the number of thermal inversions per week, their average strength, average planetary boundary height, average wind speed at the lowest altitude layer above ground-level and wind speed interacted with the employment zone location indicator variables. All regressions include month, year and employment zone fixed effects and ground-level weather controls. Robust standard errors clustered at the employment zone level in parenthesis. The employment zone (*"zone d'emploi"*) is a higher level of spatial aggregation as it divides the French territory into 306 geographical areas within which most of the working population resides and works.

Table OA17: Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly healthcare expenditure - by medical specialty, location FE

	Family practice	Cardio-vasc.	Pulmo.	O.R.L.	Ophthalmo.
NO_2	8.912***	0.749***	0.0721	0.332***	1.187***
	(0.610)	(0.104)	(0.072)	(0.040)	(0.109)
O_3	1.018***	0.0419*	0.0306	0.0383***	0.122***
	(0.121)	(0.019)	(0.017)	(0.008)	(0.021)
PM_{10}	-2.747***	-0.131*	0.0495	-0.0827***	-0.172***
	(0.350)	(0.052)	(0.038)	(0.018)	(0.049)
Lag NO_2	2.722***	0.304**	-0.118	0.146***	0.376**
	(0.586)	(0.100)	(0.075)	(0.044)	(0.114)
Lag O_3	-0.557***	0.0376*	-0.0218	0.00457	0.0372
	(0.133)	(0.019)	(0.019)	(0.008)	(0.024)
Lag PM_{10}	-0.680*	-0.127**	0.0141	-0.0449*	-0.225***
	(0.302)	(0.044)	(0.034)	(0.020)	(0.049)
	Gynaeco.	Nephro.	Gastro-hep.	Trauma surg.	Plastic surg.
NO_2	0.488***	0.0413	0.394*	0.642***	0.177***
	(0.071)	(0.044)	(0.158)	(0.092)	(0.053)
O_3	0.0222	0.0279***	0.104	0.0788***	0.0210*
	(0.015)	(0.008)	(0.063)	(0.021)	(0.010)
PM_{10}	-0.0663	-0.0378	-0.0837	-0.122*	-0.0163
	(0.035)	(0.023)	(0.089)	(0.048)	(0.024)
Lag NO_2	0.317***	-0.0218	0.218	0.223*	0.0641
	(0.091)	(0.044)	(0.149)	(0.091)	(0.048)
Lag O_3	0.00326	-0.00501	0.00481	0.0340	-0.0165
	(0.015)	(0.009)	(0.038)	(0.020)	(0.010)
Lag PM_{10}	-0.167***	-0.0160	-0.0245	-0.0629	-0.0476*
	(0.039)	(0.021)	(0.085)	(0.046)	(0.023)
Observations	1186311	1186311	1186311	1186311	1186311

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. This table presents results for regressions run separately by medical speciality using the location fixed effect (FE) model including a one week lag of the pollutant concentrations. For results using the FE-IV model, see table 2. All regressions include month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level in parenthesis.

Table OA18: Statistically significant coefficients from the heterogeneity analyses by medical specialty - Bonferroni–Holm multiple hypothesis testing correction

#	Specialty	Variable	Coefficient	SE	p-value	Holm Threshold
1	Ophthalmology	NO_2	1.108	0.228	0.0000012	0.000833
2	Ophthalmology	Lag O_3	0.206	0.044	0.0000029	0.000847
3	Cardio-vascular disease	Lag O_3	0.178	0.041	0.000014	0.000862
4	Family Practice	Lag PM_{10}	3.329	0.835	0.00007	0.000877
5	Family Practice	O_3	0.927	0.235	0.00008	0.000893
6	Family Practice	Lag O_3	1.217	0.264	0.0001	0.000909
7	Family Practice	NO_2	4.956	1.492	0.0009	0.000925
8	O.R.L.	Lag O_3	0.0476	0.017	0.0052	0.000943

The table shows the coefficients from the heterogeneity analyses by medical specialty that remain statistically significant after adjusting for multiple hypothesis testing according to Holm-Bonferroni Stepwise Adjustment. We fail to reject at step 8 as $0.0052 > 0.000943$. The first 7 tests are rejected at the 5% FWER level.

Table OA19: Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly healthcare expenditure - by age group

	Sum of weekly healthcare spending									
	Age 0-20		Age 21-40		Age 41-60		Age 61-80		Age over 80	
NO_2	2.974**	1.622	2.844**	0.140	7.062***	10.30***	2.559	3.270	1.508	2.241
	(0.962)	(0.980)	(1.057)	(1.095)	(1.867)	(1.935)	(1.651)	(1.669)	(1.173)	(1.189)
O_3	0.876***	0.611***	0.650**	0.601**	2.177***	1.275***	2.722***	1.158***	0.557*	0.285
	(0.176)	(0.168)	(0.198)	(0.198)	(0.403)	(0.364)	(0.359)	(0.311)	(0.219)	(0.216)
PM_{10}	1.313	0.389	0.431	1.443	4.705***	-0.364	-1.002	-2.128	1.506	-0.0898
	(0.696)	(0.725)	(0.817)	(0.862)	(1.371)	(1.346)	(1.191)	(1.250)	(0.819)	(0.876)
Lag NO_2		1.726		-0.00790		-4.374*		4.012*		0.0855
		(0.898)		(1.144)		(1.887)		(1.732)		(1.220)
Lag O_3		1.149***		-0.0286		1.503***		2.754***		0.797**
		(0.178)		(0.229)		(0.388)		(0.361)		(0.245)
Lag PM_{10}		1.930***		1.162		8.750***		-2.741*		2.236**
		(0.571)		(0.859)		(1.251)		(1.082)		(0.738)
Obs.	1209572	1186311	1209572	1186311	1209572	1186311	1209572	1186311	1209572	1186311
FS F-stat	2055.4	824.2	2055.4	824.2	2055.4	824.2	2055.4	824.2	2055.4	824.2

***$p < 0.001$, **$p < 0.01$, *$p < 0.05$. This table shows the FE-IV model results for regressions run separately for observations divided into age groups. All regressions include month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level in parenthesis.

Table OA20: Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly healthcare expenditure in locations with below and above postcode average income, NO_2 pollution concentrations and population size

	Dependent variable: Weekly healthcare spending			
	Panel A: Heterogeneity by average postcode average income			
	Below median income	Above median income	Below median income	Above median income
NO_2	15.47**	22.01***	22.04***	16.12***
	(5.937)	(4.746)	(5.992)	(4.542)
O_3	5.792***	7.174***	3.171**	3.378***
	(1.279)	(0.869)	(0.974)	(0.884)
PM_{10}	10.33*	13.40***	-1.149	4.677
	(4.301)	(3.450)	(4.322)	(3.514)
Lag NO_2			-7.131	0.644
			(6.431)	(4.898)
Lag O_3			5.748***	7.642***
			(1.249)	(0.977)
Lag PM_{10}			22.00***	14.41***
			(4.139)	(3.051)
Observations	607672	596076	595986	584613
First-stage F-stat	1288.0	1059.9	477.7	495.9
	Panel B: Heterogeneity by postcode average NO_2 concentration			
	Below median pollution	Above median pollution	Below median pollution	Above median pollution
NO_2	3.197	26.95***	10.09*	21.91***
	(4.130)	(5.148)	(4.225)	(4.852)
O_3	2.669***	12.27***	1.423*	6.221***
	(0.632)	(1.499)	(0.631)	(1.201)
PM_{10}	16.11***	13.13***	6.881*	1.700
	(2.954)	(3.867)	(3.062)	(3.604)
Lag NO_2			-17.48***	4.989
			(4.779)	(5.380)
Lag O_3			3.159***	12.61***
			(0.708)	(1.508)
Lag PM_{10}			20.20***	17.43***
			(2.921)	(3.590)
Observations	599092	610480	587571	598740
First-stage F-stat	1628.6	1454.4	672.7	682.3

(Continued on the next page)

Table OA20: (continued) Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly healthcare expenditure in locations with below and above postcode average income, NO_2 pollution concentrations and population size

	Dependent variable: Weekly healthcare spending			
	Panel C: Heterogeneity by postcode population size			
	Below median population	Above median population	Below median population	Above median population
NO_2	2.229	26.92***	4.381	29.18***
	(2.896)	(6.368)	(2.857)	(6.134)
O_3	2.531***	9.723***	1.382**	4.660***
	(0.492)	(1.444)	(0.475)	(1.197)
PM_{10}	8.032***	17.44***	3.722	0.503
	(2.010)	(4.927)	(2.059)	(4.725)
Lag NO_2			-8.585**	-2.112
			(3.083)	(6.571)
Lag O_3			1.725**	11.19***
			(0.559)	(1.458)
Lag PM_{10}			11.81***	24.33***
			(2.004)	(4.369)
Observations	601536	608036	589968	596343
First-stage F-stat	1136.4	1042.7	441.3	431.3

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. This table shows the results of the impact of pollution on healthcare expenditure *in absolute terms* from regressions run separately for observations categorised into groups below and above the median in terms of postcode average household income (panel A), pollutant concentration (panel B) and population size (panel C). All regressions include month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level in parenthesis.

Table OA21: Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly *per capita* healthcare expenditure in locations with below and above postcode average income, NO_2 and population size

	Dependent variable: Weekly per capita healthcare spending			
	Panel A: Heterogeneity by average postcode income			
	Below median income	Above median income	Below median income	Above median income
NO_2	0.0584	0.0750	0.171*	0.0699
	(0.070)	(0.039)	(0.070)	(0.040)
O_3	0.0441***	0.0492***	0.0210	0.0230**
	(0.012)	(0.009)	(0.012)	(0.008)
PM_{10}	0.117*	0.133***	-0.0197	0.0557
	(0.050)	(0.029)	(0.049)	(0.029)
Lag NO_2			-0.166*	-0.0621
			(0.073)	(0.038)
Lag O_3			0.0424**	0.0423***
			(0.013)	(0.009)
Lag PM_{10}			0.237***	0.158***
			(0.048)	(0.025)
Observations	607672	596076	595986	584613
First-stage F-stat	1288.0	1059.9	477.7	495.9
	Panel B: Heterogeneity by postcode average NO_2 concentration			
	Below median pollution	Above median pollution	Below median pollution	Above median pollution
NO_2	0.0941	0.0934**	0.194*	0.0764*
	(0.081)	(0.032)	(0.083)	(0.034)
O_3	0.0357*	0.0475***	0.0143	0.0216*
	(0.017)	(0.009)	(0.014)	(0.008)
PM_{10}	0.115*	0.0872***	-0.0331	0.0382
	(0.056)	(0.023)	(0.058)	(0.025)
Lag NO_2			-0.169	-0.0496
			(0.105)	(0.031)
Lag O_3			0.0321	0.0486***
			(0.017)	(0.009)
Lag PM_{10}			0.228***	0.131***
			(0.064)	(0.020)
Observations	599092	610480	587571	598740
First-stage F-stat	1628.6	1454.4	672.7	682.3

(Continued on the next page)

Table OA21: (continued) Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly *per capita* healthcare expenditure in locations with below and above postcode average income, NO_2 and population size

	Dependent variable: Weekly per capita healthcare spending			
	Panel C: Heterogeneity by postcode population size			
	Below median population	Above median population	Below median population	Above median population
NO_2	0.0640	0.0641*	0.144	0.0663*
	(0.082)	(0.028)	(0.082)	(0.028)
O_3	0.0564**	0.0308***	0.0306*	0.00997
	(0.019)	(0.006)	(0.015)	(0.006)
PM_{10}	0.165**	0.0868***	0.0458	0.0206
	(0.059)	(0.021)	(0.058)	(0.022)
Lag NO_2			-0.298**	0.000430
			(0.096)	(0.027)
Lag O_3			0.0268	0.0476***
			(0.019)	(0.006)
Lag PM_{10}			0.319***	0.0997***
			(0.063)	(0.017)
Observations	601536	608036	589968	596343
First-stage F-stat	1136.4	1042.7	441.3	431.3

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. This table shows the results of the impact of pollution on healthcare expenditure *per capita* from regressions run separately for observations categorised into groups below and above the median in terms of postcode average household income (panel A), pollutant concentration (panel B) and population size (panel C). All regressions include month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level in parenthesis.

Table OA22: Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations on weekly healthcare expenditure - heterogeneity by pollution quartile, piece-wise linear regression

	Weekly healthcare expenditure
NO_2 x first NO_2 quartile	200.6***
	(42.270)
NO_2 x second NO_2 quartile	102.0***
	(19.183)
NO_2 x third NO_2 quartile	103.6***
	(22.807)
NO_2 x fourth NO_2 quartile	61.74***
	(9.416)
O_3 x first O_3 quartile	38.78***
	(5.548)
O_3 x second O_3 quartile	35.49***
	(3.213)
O_3 x third O_3 quartile	25.99***
	(3.055)
O_3 x fourth O_3 quartile	23.76***
	(2.307)
PM_{10} x first PM_{10} quartile	215.2***
	(24.213)
PM_{10} x second PM_{10} quartile	120.6***
	(16.096)
PM_{10} x third PM_{10} quartile	102.2***
	(15.918)
PM_{10} x fourth PM_{10} quartile	84.58***
	(9.103)
Observations	1209572
First-stage F-stat	609.5

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. This table shows results for piece-wise linear regressions in which the weekly pollutant concentration are interacted with a dummy variable that categorises that week's pollutant concentration into four categories per quartile of its value. All regressions include month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level in parenthesis.

Table OA23: Impact of average weekly NO_2, O_3 and PM_{10} pollutant concentrations above and below the annual average WHO limit value on weekly healthcare expenditure

	Weekly healthcare spending	
	(1)	(2)
NO_2 x below WHO limit	48.53***	30.33**
	(11.972)	(11.714)
NO_2 x above WHO limit	29.96***	18.80***
	(4.882)	(4.645)
PM_{10} x below WHO limit	84.70***	11.20
	(11.755)	(8.458)
PM_{10} x above WHO limit	43.08***	6.981
	(5.874)	(4.445)
Lag NO_2 x below WHO limit		45.34***
		(11.832)
Lag NO_2 x above WHO limit		10.97*
		(5.159)
Lag PM_{10} x below WHO limit		12.58
		(8.532)
Lag PM_{10} x above WHO limit		13.33**
		(4.163)
O_3	9.471***	3.482***
	(0.903)	(0.757)
Lag O_3		6.894***
		(0.886)
Observations	1209572	1186311
R^2	0.015	0.021
First-stage F-stat	196.2	143.8

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. This table presents results for the effect of average weekly pollutant concentrations above and below the annual average World Health Organisation (WHO) limit value of 10 $\mu g/m^3$ and 15 $\mu g/m^3$ on weekly healthcare expenditure. All regressions include month, year and zip code fixed effects and ground-level weather controls. Robust standard errors clustered at the zip code level in parenthesis.