ECONtribute
MARKETS & PUBLIC POLICY

ECONtribute
Discussion Paper No. 320

Putting a Price Tag on Air Pollution: The Social Healthcare Costs of Air Pollution in France

Julia Mink

June 2024

www.econtribute.de

UNIVERSITÄT BONN
Universität zu Köln

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1-390838866 is gratefully acknowledged.

Putting a price tag on air pollution: the social healthcare costs of air pollution in France

Julia Mink*

June 23, 2024

Abstract

This study quantifies the financial burden of acute air pollution on the French healthcare system. By combining comprehensive French administrative health data for a nationally representative sample with high-resolution geospatial data on air pollution and meteorological conditions, the healthcare costs of air pollution exposure are estimated more accurately and comprehensively than in the previous literature. I use an instrumental variable approach exploiting weekly variations in local concentrations of nitrogen dioxide, ground-level ozone and particulate matter induced by variations in altitude weather conditions. I find that air pollution causes healthcare costs to the French healthcare system in the order of several billion per year, even though air pollutant concentrations are mostly below the current European air quality standards considered safe for human health. My cost estimates are about 10 times higher than those estimated in previous studies, suggesting that the health costs of air pollution have been severely underestimated. While air pollution has a large effects on overall spending in more polluted and populated urban areas due to the higher number of affected people, the marginal effects appear to be greater in low-pollution and less populated areas. Reducing population exposure even at low air pollution concentrations should therefore be an important public health goal. Even the most stringent 2021 WHO guideline values should not be considered safe for human health.

Key words: Air pollution, healthcare cost, instrumental variable approach
JEL codes: I12, J14, Q51, Q53

*Email: julia.mink@uni-bonn.de Orcid ID: 0000-0001-9157-8227 Affiliation: University of Bonn, Institute for Applied Microeconomics, Department of Economics, Lennéstraße 43, 53113 Bonn
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 have no conflict of interest to declare.

1 Introduction

Air pollution remains the biggest environmental risk to the health of Europeans (EEA, 2020). It is not only considered one of the most important risk factors for death, but also one of the main causes of the global burden of disease. Exposure to air pollution has well-documented adverse effects on human health, such as the increased risk of cardiovascular and respiratory disease and cancer (WHO, 2017). As a response, 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. The EU air quality guidelines are currently being revised to bring them more in line with the stricter World Health Organisation (WHO) guidelines, which themselves were recently updated in 2021 and are now much stricter than their 2005 version (European Commission, 2016). Accurately quantifying the effects of air pollution exposure is essential to help policy makers determine the optimal level of environmental policy. This is particularly relevant in cases where pollution levels are already relatively low and the benefits from further pollution reduction may not offset the costs anymore. In this study, I quantify the healthcare costs caused to the French health system by acute exposure to air pollution. This is in a context where air pollution is on average far below the current EU air quality standards.

Estimating the causal impact of air pollution on healthcare costs is difficult due to the potential endogeneity of exposure and the need for high quality data. People sort spatially according to preferences and characteristics that may correlate with their health status and pollution exposure. Failure to take account of this non-random exposure might result in biased estimates. To address this problem, many studies have relied on quasi-experimental designs that exploit a plausibly exogenous source of pollution variation to estimate the causal effects of air pollution on health. However, due to the requirements of the quasi-experimental method and data limitations, these studies tend to be restricted to relatively narrow geographical areas and time periods, look at only a certain part of the population or investigate the effects of air pollution on a limited range of health conditions. Much of the existing literature focuses on mortality, a rather extreme event that is less likely to occur at moderate levels of pollution, in contrast to health costs, which may represent a significant financial burden for households and insurers even at low levels of pollution. In this study, I overcome these limitations by using comprehensive French administrative data on healthcare in combination with high-resolution geospatial data on air pollution and meteorological conditions. Exploiting high-quality data allows me to comprehensively quantify health effects in a representative sample of the French population and to use an instrumental variable approach based on a high-dimensional vector of instruments to simultaneously account for the effects of multiple pollutants. To my knowledge is the first quasi-experimental study to comprehensively quantify the healthcare costs caused by exposure to moderate levels of air pollution in a nationwide representative sample.

More specifically, this study examines the causal effects of exposure to nitrogen dioxide (NO2), ground-level ozone (O3) and particulate matter (PM 10) on health care costs, covering all types of health care, including physician visits, drug purchases, and hospital treatments, as well as all types of medical specialties. I estimate a location fixed-effects model that exploits the weekly variation in air pollution concentrations at the level of the French postcode area which typically represents an area of about 9 km 9 km. I flexibly account for the seasonality in air pollution and health care utilisation through the inclusion of month or month-by-department and year fixed effects. I also control for the influence of ground-level weather conditions by including a vector of indicator variables for temperature, precipitation and wind speed, categorised into 10 deciles, and various possible interactions of these weather indicator variables. While the location fixed effects can account for bias by capturing cross-sectional and time-invariant location-specific population characteristics, they cannot fully address endogeneity bias due to correlations between pollution concentrations and economic activity. For example, pollution could reflect increased economic activity and traffic congestion, which can be associated with stressful conditions affecting health and healthcare expenditure. I therefore adopt an instrumental variable (IV) approach where I instrument for air pollution concentrations using altitude weather conditions. Altitude weather conditions are good instruments because they are highly predictive of air pollution concentrations. They shift air pollution levels frequently and at all locations across France, which means that I do not need to limit the analyses to relatively narrow geographical areas and time periods. Using such a high-dimensional vector of instruments also has the advantage of enabling multi-pollutant models to be examined, which is important given that atmospheric pollutants are highly correlated with each other. The identifying assumption is that variation in pollution due to changes in altitude weather conditions is unrelated to changes in healthcare use or costs except through the influence on air pollution. After flexibly controlling for various time and location fixed effects and ground-level weather conditions, this assumption should hold. To demonstrate the robustness of my results, I estimate alternative specifications with different instruments, varying time fixed effects and ground-level weather controls and estimate specifications in which I control for lags of the weather and pollution variables. I also allow for the possibility that pollution build-up over the past weeks may impact health outcomes by allowing for longer effect windows. Besides estimating the overall healthcare cost of exposure to air pollution, I also examine which types of health problems in particular are affected by running separate regressions for 10 categories of medical specialities, considering both medical specialities that should be affected by air pollution (such as cardiology and vascular medicine or pulmonology) and medical specialities that should not be affected (such as plastic surgery or trauma surgery), which serve as placebo. Finally, I study heterogeneity in the effects of air pollution exposure by patient characteristics including age, chronic disease status and insurance status and location characteristics including the zip code area average income, population size and average pollution concentrations.

I find that exposure to pollutant 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. Air pollution causes additional healthcare expenditure of at least €12.8 billion year, which corresponds to 0.5% of France's GDP in 2019 and 6.2% of France's total healthcare expenditure in 2019[1]. My cost estimates are around 10 times higher than those of previous studies, suggesting that the health costs of air pollution have been severely underestimated and that policy-makers have based their decisions on incomplete information. Note that while these health costs are much higher than estimates from previous studies, they should still be considered a lower bound for the total health costs of air pollution. My estimates reflect the short-term effects of air pollution and ignore the effects of chronic exposure. Mortality costs or costs associated with sick leave are also not included here. The results from heterogeneity analyses by medical specialty are consistent with the findings from the economic and epidemiological literature. I detect no effects of pollution in the placebo medical specialties and find effects in the specialties cardiology and vascular medicine, pulmonology, otolaryngology (O.R.L.), ophthalmology, and gynaecology. Family practice shows the strongest response, which is consistent with the fact that the family practitioner is the first point of contact with the healthcare system. Heterogeneity analyses by age show that air pollution generates health costs in all age groups, suggesting that adverse health effects also occur in parts of the population that are less frequently considered. Previous studies on the health effects of air pollution have often focused on the young or elderly population, as these populations are generally considered to be the most vulnerable. A possible explanation for the fact that I find effects for all age groups is that I am looking at health costs, which include the costs of treating also milder health effects that are likely to occur in all age groups, whereas many existing studies focus on mortality, which is a rather extreme result that is likely to affect only the most vulnerable populations. Heterogeneity analyses based on location characteristics reveal geographic disparities in the effects of air pollution. Interestingly, impacts do not appear to different across locations based on average income, but I do find differential effects depending on the locality's average level of pollution and population size. While air pollution in more populated urban areas affects a larger number of people and therefore 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.

[1]In 2019, the GDP of France reached €2,425.7 billion (INSEE, 2020) and aggregate healthcare spending was €208 billion (DREES, 2020).

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 health costs caused by air pollution has so far been very 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, as they tend to focus on a relatively narrow geographical area or on events limited in time, often cover only a specific part of the population and/or investigate the effects of air pollution on a limited range of health problems (see for example Ransom and Iii (1995); Pope III and Dockery (1999); Friedman et al. (2001); Chay and Greenstone (2003); Neidell (2004); Currie and Neidell (2005); Jayachandran (2009); Neidell (2009); Moretti and Neidell (2011); Currie and Walker (2011); Chen et al. (2013); Anderson (2015); Schlenker and Walker (2015); Knittel et al. (2016); Schwartz et al. (2017); Deschênes et al. (2017); Deryugina et al. (2019); Simeonova et al. (2019); Halliday et al. (2019)). For example, the quasi-experimental study that is probably the most comparable to the present study in terms of data quality and empirical strategy is Deryugina et al. (2019), which is limited to analysing hospital costs and focuses on the impact on the elderly population in the USA. Much of this research focuses on mortality costs, while health costs are generally not quantified, as detailed information on health expenditure is rarely available.

This study goes beyond the previous literature in three important ways. First, it quantifies the health costs of exposure to air pollution more accurately and comprehensively than previous studies by examining the effects in a nationally representative sample considering information on the exact costs for all types of health care. While existing studies clearly indicate that their estimates of health costs are conservative, the extent to which total effects have been underestimated is not known. Second, the study analyses treatment effect heterogeneity as a function of several patient and location characteristics, providing information about inequalities and non-linear effects in a way that previous studies that do not draw on data from a representative sample covering a large geographical area cannot. Third, the 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), which is complicated by the high correlation between the pollutants.

The results of this study provide information that is highly relevant to environmental policy. My cost estimates show that the health costs of air pollution to the French health system have been severely underestimated. Previous policy decisions were 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. Health costs are caused by exposure to air pollution levels that are mostly far below current European air quality guideline values, suggesting that the current guideline values, which are supposedly safe limits for human health, are set too high. Even the strictest WHO 2021 guideline values should not be considered safe for human health. In fact, the relationship between air pollution and health costs appears to be concave, with the marginal effects being greater at low pollution concentrations. This means that reducing population exposure even at low air pollution concentrations should be an important public health goal. In addition to the implications for cost-benefit calculations and the setting of air quality guideline values, the results of this study are also relevant for public health communication. Warnings are usually aimed at the population on days with peak air pollution levels and at target groups that are considered particularly vulnerable, such as the elderly, children and people with chronic health conditions. The results of this study suggest that health messages need to be updated to inform the public that health effects occur even at low levels of air pollution and affect people of all ages.

The rest of the paper is organised as follows. Section 2 provides a brief background on air pollution and air quality in France, the health effects of air pollution and a discussion of altitude atmospheric conditions as instruments for ground-level air pollution concentration. Section 3 describes the data, section 4 describes the empirical strategy, section 5 presents results and section 6 provides a conclusion.

2 Background

This section provides background information on air pollution and air quality in Frace, the health effects of air pollution and contains a discussion on the adequacy of altitude atmospheric conditions as instruments for ground-level air pollution concentrations.

2.1 Air pollution and air quality in France

Air pollution is the contamination of the environment by any chemical, physical or biological agent that modifies the natural characteristics of the atmosphere. In this study, I focus on the air pollutants nitrogen dioxide (NO2), ozone (O3), particulate matter 10 micrometers or less in diameter (PM10). 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 $40g/m^3$ for NO2, a maximum daily 8-hour mean of $120g/m^3$ for O3 not to be exceeded more than 25 days per year, a yearly mean of $40g/m^3$ for PM10 and $25g/m^3$ for PM2.5. Although these values were based on the 2005 Wold 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 A1 in the appendix presents the 2005 and 2021 WHO air quality guideline values and the current EU air quality standards applicable in France.

Following the implementation for several years of strategies and action plans in various sectors of activity, air quality in France improved over the last two decades with the exception of O3 pollution. Exceedances of regulatory air quality standards still persist, but they are fewer than in the past and affect fewer areas (Le Moullec and Meleux, 2019). Air pollution concentrations in France during the study period are mostly well below the current EU air quality limit values. I observe an average NO2 concentration over the years 2015 to 2018 of 13.8 $g/m3$ and average PM10 and PM2.5 concentrations are 16.61 $g/m3$ and 10.58 $g/m3$, 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 postcode locations. The study therefore focuses on the effects of pollutant concentrations that are considered safe according to the applicable EU limit values.

There are strong correlations between the air pollutants as many share common sources and some pollutants are precursor pollutants that are transformed into different secondary pollutants through chemical reactions. Nitrogen dioxide (NO2) and particulate matter (PM) tend to be positively correlated as they share common sources. Nitrogen oxides (NOx), which include nitrogen monoxide (NO) and NO2, are emitted during the combustion of fuels from industrial plants and road traffic and contribute to the formation of ozone (O3) and PM. PM is a mixture of solid and liquid aerosol particles covering a wide range of sizes and chemical compositions. PM is either directly emitted as primary particles or it forms in the atmosphere from emissions of certain precursor pollutants such as sulfur dioxide (SO2), NOx, ammonia (NH3) and volatile organic compounds (VOCs). O3 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 NOx, carbon monoxide (CO), volatile organic compounds (VOCs) and methane (CH4). The processes of O3 formation and accumulation are complex. To put it simply, NO2 and oxygen (O2) react with each other, resulting in NO and O3. Being an equilibrium reaction, the reaction also works in the other direction whereby ozone gets degraded again (EPA; Clapp and Jenkin, 2001). In the short-term, NO2 tends to be inversely related to O3 because in many settings NO2 disappears during the formation of O3 and vice-versa(Lee et al., 2021).

2.2 Health effects of air pollution

Although air pollution emissions have declined in the last decades, air pollution remains the single largest environmental risk to the health of Europeans, with particulate matter (PM), nitrogen dioxide (NO2) and ground-level ozone (O3) being the pollutants of greatest concern (EEA, 2020). Exposure to PM has been estimated to be responsible for around 400,000 premature deaths in Europe every year whereas exposure to NO2 and O3 were responsible for around 70,000 and 15,000 premature deaths in 2017, respectively (Maguire et al., 2020). Air pollution is one of the leading risk factors for death and also one of the main contributors to global disease burden. In low-income countries, air pollution is often very near the top of the list or the leading risk factor. In France, despite being a country with relatively low pollution levels, air pollution is among the top 10 risk factors for both mortality and disease burden (Institute for Health Metrics and Evaluation, 2020). Exposure to air pollution has been linked to various adverse health outcomes. Short-term exposure to air pollution is associated with Chronic Obstructive Pulmonary Disease (COPD), cough, shortness of breath, wheezing, asthma, respiratory disease, and high rates of hospitalisation for respiratory and cardio-vascular disease. Young children, the elderly, and people with lung disease have been shown to be especially vulnerable to air pollution. The health of susceptible and sensitive individuals can be impacted even on low air pollution days. For a review, see for example Manisalidis et al. (2020)). Most of the evidence for the health effects of air pollution comes from observational studies that provide correlational evidence. However, the correlative evidence has been largely confirmed by the results of a growing quasi-experimental literature showing effects on both mortality and morbidity (see for example Currie and Walker (2011); Schlenker and Walker (2015); Schwartz et al. (2017); Deryugina et al. (2019); Godzinski and Castillo (2021)).

2.3 Altitude atmospheric conditions and ground-level air pollution levels

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. The reasons for using a high-dimensional vector of instruments are two-fold. First, 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). Second, using a range of instruments including not only infrequently occurring events such as thermal inversions but also frequently occurring events such as changes in planetary boundary height and changes in wind direction is useful because Bagilet and Zabrocki (2021) show that an instrumental variable strategy with low frequency events as instruments may lead to inflated estimates due to low statistical power when estimating acute health effects. The weather phenomena that serve as the basis for the construction of the instruments are presented below, and exogeneity and the exclusion restriction are briefly discussed.

Thermal inversions

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 or when the surface cools overnight. This phenomena is well documented in the scientific literature (Wallace and Kanaroglou, 2009; Gramsch et al., 2014). As variations in surface- and higher-level temperatures within a region are usually assumed to be exogenous, thermal inversions are assumed to be exogenous. For the exclusion restriction to hold, thermal inversions should influence health expenditures only through their effects on pollutant concentrations. 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).

Height of the planetary boundary layer

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). The planetary boundary layer height is also closely related to thermal inversions. Planetary boundary layer height (PBLH) responds to heating flux between the sun and the earth. PBLH can also change under unpredictable large-scale air movements and responds to subsidence where air sinks down in an area of high pressure, bringing the top of the layer downward. However, whereas thermal inversions may or may not happen, encoded in a dummy variable, the height of the planetary boundary layer is always be defined and is a continuous variable. 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

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 can be used as an instrument for pollutant concentrations, but only conditional on controlling for wind speed at ground level. Because 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, it should be noted that 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 data on healthcare 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, income distribution 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 level of the French zip code, which typically represent an area of about 9 km 9 km. The raw data are available at daily frequency, with the exception of the tax and social benefit data, which are available at an annual frequency.

3.1 Administrative healthcare data

I use data on healthcare costs come from the French National System of Health Data (SNDS for *Système National des Données de Santé*). 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, i.e. more than 66 million people, from birth or immigration to death or emigration, making it possibly the largest contiguous homogeneous benefits database in the world. The data provide information on the nature of the medical acts and associated 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 and also includes codes for the classification of medical acts into medical specialities. Available data on patient characteristics include patient age, sex, information on chronic health conditions, and zip code of residence.

I use data from 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 conduct of longitudinal studies as it permits tracing back patients healthcare use history. See Tuppin et al. (2010) and Bezin et al. (2017) for more information on the EGB. I aggregate the individual-level data at the level of the zip code of the patient's place of residence. For the analyses of effect heterogeneity according to patient characteristics, I additionally divide the observations into groups according to sex, age and insurance status.

3.2 Air pollutant concentrations

For the air pollution measures, I use reanalysis data on hourly concentrations of NO2, O3 and PM10 provided by the French National Institute for Industrial Environment and Risks (INERIS for *"Institut national de l'environnement industriel et des risques"*). 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. 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 (2015)). 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. INERIS is recognised in France and internationally for its models predicting atmospheric pollution levels on different time and space scales. The institute has mapped air quality over the entire French metropolitan area and Corsica with its CHIMERE chemical transport model since 2000. For detailed information on the construction of the reanalysis data, see Real et al. (2022).

For sensitivity analyses I also use data from monitoring stations for NO2, O3 and PM10 and additionally SO2 and CO concentrations provided by the European Environment Agency (EEA)[2]. CO concentrations are recorded at 44 monitoring stations and SO2 concentrations at 173 monitoring stations, which means that the geographical coverage is relatively sparse. The data on NO2, O3 and PM10 levels are collected at 475, 370 and 411 monitoring stations respectively, which means that they are less sparse than the CO and SO2 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 areas 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.

[2] The EEA data can be downloaded here

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). The 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[3].

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. The u-component is the eastward component of the wind and the v-component is the northward component of the wind. It is expressed as the horizontal speed of air moving towards the east and the north, respectively. Wind speed can be calculated as $WS = \sqrt{u^2 + v^2}$, where u and v are the u- and v-component of the wind, respectively. 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 $0\Phi < 360$[4]. As expressed here, Φ indicates the direction from which the wind is blowing. Zero means the wind is blowing from the north to the south. Higher angles correspond to clockwise cardinal directions, so 180 means the wind is blowing from the south to the north. I use the average cardinal wind direction to construct wind direction bins.

I convert the hourly data into daily averages and 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 try 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

[3]The ground-level data can be downloaded here, the altitude data here and the planetary boundary layer height can be downloaded here.

[4]See how to calculate wind speed and wind direction from u and v components of the wind from the ECMWF Q&A here. Depending on the software, the arguments entering atan2 might have to be inverted.

influence air pollution concentrations more than when it occurs at night (Godzinski and Castillo, 2021). From the daily and the six within-day averages, I construct weekly sums and averages for the analyses at weekly frequency resulting in the following variables: 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.

3.4 Additional data and summary statistics

The information on median household income at zip code level with annual frequency comes from the social and fiscal localised database FiLoSoFi (for *Fichier Localis Social et Fiscal* in French). These data are based on administrative information relating to taxes and social benefits and are provided by the French National Institute of Statistics and Economic Studies (INSEE for Institut national de la statistique et des études économiques in French). Aggregated data at postcode area level, either for all households or by household category, are publicly available on the INSEE website [5].

Data on holidays in France, to be used as control variables, are obtained from the Open platform for French public data[6]

The final dataset contains $1,257,984$ postcode-week observations on healthcare costs, air pollutant concentrations and weather conditions for the $6,048$ French zip code areas for the years 2015 to 2018. Table A2 in the appendix presents summary statistics.

Weekly average healthcare expenditure is 3,609 € with a standard deviation of 7471. Mean daily concentration of NO2 is 13.78 (standard deviation 7.14); concentration of PM 10 is 16.61 (sd 6.32); concentrations of O3 is 55.7 (sd 17.49) micro-grams per cubic meter. These pollutant levels are far below the air quality standards currently in force in France (see Table A1 and the discussion in section 2.1).

4 Empirical strategy

4.1 Location and time fixed effects model

The aim of this study is to quantify the health costs caused by air pollution. Estimating causal effects is a challenge because exposure to air pollution is not random. People choose where they live and thus the extent of their exposure, which can lead to correlations between air pollution and personal characteristics, possibly

[5] The data form the social and fiscal localised database can be downloaded here
[6] The data on public holidays can be downloaded here.

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.

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

To quantify the impact of air pollution on health care costs as comprehensively as possible, I construct the health care cost variable to include the costs 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 95 French departments. The vector of time-varying location characteristics X_{pw} includes indicator variables for holidays, indicator variables for daily mean temperatures and daily precipitation falling into 10 bins by decile and different possible interactions of these weather indicator variables. I estimate alternative specifications with different time fixed effects and location covariates to demonstrate the robustness of the results (see section 5.1).

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 (see section 5.1).

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 attenuation bias (biased towards zero). 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 A2 in the appendix.

Standard errors are clustered at the postcode 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 for example, 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 2.3 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. While weather conditions at ground level may 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 postcode 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 2.3 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). 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 postcode 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. 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 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 (see also the section 2.2). 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)). 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.[7].

To identify the populations at particular risk, I study heterogeneity in the effects of air pollution exposure by patient characteristics including age, chronic disease status and insurance status and location characteristics including the zip code area average income, population size and average pollution concentrations. Analysing the heterogeneity of effects according to the average concentration of pollutants at a given location is also interesting in order to find out whether the health effects of air pollution are non-linear. Identifying the most vulnerable population groups and knowing whether the impact of air pollution on health

[7]Except for planned operations, the choice of placebo categories is not obvious. Some studies have shown that air pollution causes an inflammatory response in the body that could lead to problems in all organs, which means that the digestive system could potentially be affected, but I don't expect this effect to be strong for acute exposure to air pollution. Accidents could be caused by poor visibility due to smog or pollution-induced aggressive behaviour Chan et al. (2022), but again I don't expect this effect to be strong except maybe on some days of large spikes in air pollution

is linear is important for setting policy priorities.

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 sizes and policy discussion.

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 postcode 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 postcode 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 NO2 leads to an average €17.23 of additional healthcare expenditure per postcode area during the same week. This corresponds to a 0.48% increase relative to the average weekly postcode area healthcare spending. The €17.23 of additional healthcare expenditure per postcode area per week in a sample of 1/97 of the French population corresponds to €525,620,310 of additional overall healthcare spending in France per year or 0.02% of France's GDP in 2019.[8]. For a discussion regarding the effect size, see section 5.3.

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

[8]The calculation €17.23 · 97 · 52 · 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 postcode areas. The GDP of France in 2019 was 2,332 billion.

Table 1: Impact of average weekly NO2, O3 and PM10 pollutant concentrations on weekly healthcare expenditure

	Weekly healthcare expenditure			
	FE		FE-IV	
NO2	44.33***	43.36***	18.42***	17.23***
	(2.692)	(2.420)	(3.820)	(3.719)
O3	4.189***	4.837***	6.282***	3.275***
	(0.383)	(0.387)	(0.773)	(0.662)
PM10	-12.06***	-13.34***	12.37***	3.540
	(0.981)	(0.996)	(2.815)	(2.843)
Lag NO2		8.947***		-3.423
		(2.119)		(4.062)
Lag O3		-0.175		6.497***
		(0.364)		(0.795)
Lag PM10		-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 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 A3 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.

are strong predictors of pollution concentration. [9] Table A3 in the appendix shows the corresponding first stage regressions. 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 A4 in the 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 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 A5 in the 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 A5 in the 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 A5 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. The fact that the results are stable when the instruments are interacted with location fixed effects is reassuring, as it provides evidence that the monotonicity assumption holds. The monotonicity assumption requires that air pollutant concentrations are always (weakly) positively or always (weakly) negatively correlated with an instrument in all postcode areas. The monotonicity of the instrument effect is necessary to be able to interpret the IV estimates as the local average treatment effect (LATE) (Imbens and Angrist, 1994). 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. Given that the results are similar in models with and without location FE interactions, this suggests that the monotonicity of instrument effects

[9] Shown in the Table are the Kleibergen-Paap Wald rk F statistics which are robust when the i.i.d. assumption is dropped using clustered standard errors as are used here with the cluster at the postcode area level. The degrees of freedom adjustment for the rk statistic is $(N-L)/L1$, as with the Cragg-Donald F statistic, except in the cluster-robust case, when the adjustment is $N/(N-1)*(N_clust-1)/N_clust$, following the standard small-sample adjustment for cluster-robust. Kleibergen-Paap statistic are the Stock-Yogo critical values for the Cragg-Donald i.i.d. case.

across location holds.

The results are also robust to changing the way the weather control variables are included as can been seen in Table A6 in the appendix. Column 1 shows the results using weather fixed effects where the variables have been partitioned into 5 bins by quintiles of their values and Column 2 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 3. Column 4 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 95 French Departments. In the main analyses, the standard errors are clustered at the level of the postcode 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.

Considering multiple pollutants simultaneously

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 A7 in the 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 A8 in the 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 A9 in the appendix.

It is important to estimate the effect of multiple pollutants simultaneously. Including only one of the pollutants at a time or including either NO2 or PM together with O3 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 A10 in the 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 NO2 and PM are larger while the coefficient on O3 is smaller. The direction of the bias is consistent with the correlations between the pollutants. NO2 and PM are positively correlated because NO2 is a precursor to PM, meaning that an omission of one of these pollutants leads the coefficient on the included pollutant to overstate its effect. O3 is mostly inversely related to NO2 and PM because the pollutants O3 and NO2 (generally NOx) are linked through equilibrium reactions (see section 2). When O3 increases, NO2 and PM tends to increase and the health effects of an increase from O3 are therefore attenuated by the health benefits from the increased PM and NO2 when these pollutants are not included in the regression. The results are robust to instrumenting only one pollutant at a time while including the others as controls as shown in Table A11 in the appendix. As NO2 is a precursor to PM it may not be very meaningful to separate the effect of the two pollutants as some of the health effects of NO2 are potentially mediated through the health effects of PM. However, the FE-IV model including all three pollutants produces the most conservative estimates for the effects of NO2 and PM10 and effects for O3 that are similar compared to the two-pollutant model, making it my preferred model specification.

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 (SO2) 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 that is comparable to the INERIS reanalysis data I have for NO2, O3 and PM pollution. 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 A12 in the appendix shows that the results for the effect of NO2, O3 and PM pollution are robust to including SO2 and CO pollution as control variables (Columns 1 and 2). The results are also similar when SO2 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 NO2 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 using the same vector of altitude atmospheric conditions. As an additional sensitivity analysis, I use the EEA measuring station data on NO2, O3 and

PM 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 A13 in the appendix).

Considering the timing of the effects

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. For example, a person with chronic asthma might have an attack triggered by exposure to increased air pollution that would have been triggered at a later date anyway without the exposure. 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 NO2 pollution on health expenditure in the same week remains unchanged, but the size of the coefficient for the effect of O3 pollution is reduced by half and the effect of PM10 is no longer statistically distinguishable from zero. The coefficients for the lags of O3 and PM10 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 5.3. 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 A14 in the 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. 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 the one week lag of air pollution (week -1) on healthcare spending during the following 3 weeks (week 1 to 3) shown in Column 4. I find that the estimates increase with the length of the time window, with one exception for the coefficient for the lag of NO2 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.

Additional sensitivity analyses

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 A15 in the appendix. Controlling for two days lag of pollutant concentrations, I find that an increase of daily average NO2 pollution by 1 $\mu g/m^3$ leads to an increase in daily postcode area healthcare costs of €4.95 while an increase in daily average O3 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 NO2 and O3 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 NO2 pollution and O3 pollution, respectively. Similar to the results from the weekly frequency analyses, the effect for PM10 is not statistically significant when pollution lags are included.

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 A2 in the appendix. Table A16 in the appendix shows that the results are qualitatively similar when the analysis is carried out at the employment zone level, albeit less statistically significant. 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. Some of the results at the employment zone level are even quantitatively close to the results at the postcode area level. Column 3 indicates that an increase of one unit in the weekly average NO2 and O3 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 NO2 and O3 pollution, respectively, estimated using the weekly frequency data.

5.2 Heterogeneity analyses

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. 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 NO2 and O3 pollution while pulmonology is affected by PM10 pollution. For O.R.L., effects are found for the lags of O3 and PM10 exposure, while there are effects of NO2 and the lags of O3 and PM10 pollution on expenditures for ophthalmology. For gynaecology, I find effects for the lag of O3 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. It is reassuring to see that 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. Even more reassuring is that the placebo categories do not appear to be affected, as none of the estimates are statistically significantly different from zero. This suggests that the estimates from the IV-FE model are not simply picking up some spurious correlation between air pollution concentrations and health expenditure, which could be due, for example, to fluctuations in economic activity. In contrast, many of the estimates from the simple location FE model shown in Table A17 in the Appendix are negative, highlighting again the difficulty of estimating the effects of multiple correlated pollutants simultaneously without using instruments (see again the discussion in section 5.1). 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 some sort of spurious correlation between healthcare expenditure and pollution.

Results by patient and location characteristics

To identify the populations at particular risk, I study heterogeneity in the effects of air pollution exposure by patient characteristics including age, chronic disease status and insurance status. 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. Table A18 in the appendix shows the FE-IV model results for regressions run separately for observations divided into age groups. The estimated level effect is higher for older individuals of 40 years and above, but the effect relative to the age group's average expenditure is more similar across age groups. 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. A potential explanation for this is that many of the previous studies focus on the effects of mortality, which is a rather extreme outcome likely to affect the only the most vulnerable populations. I look at overall healthcare costs, which include the costs of treating milder health consequences that are likely to occur in all age groups. Heterogeneity analyses by chronic disease status and socioeconomic status as indicated by enrolment status in the state funded complementary insurance plan available to low-income individuals (CMUC for *Couverture médicale universelle complémentaire*) produces unstable results and no clear pattern.

Table 2: Impact of average weekly NO2, O3 and PM10 pollutant concentrations on weekly healthcare expenditure - regressions run separately by medical specialty

	Family practice	Cardio-vascular	Pulmo.	O.R.L.	Ophthalmo.
NO2	4.956***	0.466*	0.0177	0.0236	1.108***
	(1.492)	(0.223)	(0.179)	(0.084)	(0.228)
O3	0.927***	0.0401	0.0127	0.00108	0.107*
	(0.235)	(0.040)	(0.035)	(0.017)	(0.042)
PM10	-1.180	-0.0541	0.180	-0.0468	-0.336*
	(1.143)	(0.159)	(0.139)	(0.062)	(0.170)
Lag NO2	2.513	-0.00495	-0.300	0.0897	0.192
	(1.297)	(0.225)	(0.202)	(0.084)	(0.240)
Lag O3	1.217***	0.178***	0.0273	0.0476**	0.206***
	(0.264)	(0.041)	(0.031)	(0.017)	(0.044)
Lag PM10	3.329***	0.239	0.260*	0.140**	0.318*
	(0.835)	(0.142)	(0.126)	(0.052)	(0.149)
	Gynaeco.	Nephro.	Gastro-hepato.	Trauma surg.	Plastic surg.
NO2	0.102	0.0517	-0.513	-0.107	-0.0235
	(0.147)	(0.082)	(0.345)	(0.218)	(0.108)
O3	0.00422	0.0130	0.0480	0.0276	0.0306
	(0.029)	(0.017)	(0.084)	(0.040)	(0.021)
PM10	0.170	-0.0335	0.370	0.172	0.129
	(0.111)	(0.060)	(0.278)	(0.159)	(0.080)
Lag NO2	0.0581	0.0115	-0.285	0.327	-0.111
	(0.160)	(0.091)	(0.410)	(0.222)	(0.106)
Lag O3	0.0644*	0.0138	0.0281	0.0756	-0.0109
	(0.031)	(0.017)	(0.074)	(0.041)	(0.022)
Lag PM10	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 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. 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. The specialties of gastro-hepatology, nephrology, trauma surgery, and plastic surgery are considered as 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.

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 postcode area characteristics. Tables A19 and A20 in the appendix shows the results of the regressions for observations categorised into groups below and above the median in terms of postcode 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 A19 in the appendix). Higher expenditure in areas where the number of people affected by air pollution is higher is not surprising. Regressions weighted by population size yield estimates about seven times larger than the unweighted main regression specification. See Table A21 in the appendix.

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 postcode average household income. As can be seen in Panel A of Table A20, 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 the effects are not stronger in places with more low-income populations because the French healthcare system is universal and lower-income individuals respond in the same way as higher-income individuals when seeking medical care. In addition, I do not find stronger effects for people with chronic health conditions which means that I also should not expect to find stronger effects in poorer areas, even if people living there have poorer health on average. However, it is also possible that there are differences between low- and high-income locations but that there are not detectable in the data. It is likely that there is significant income heterogeneity within a particular postcode area that is unobserved here and that is relevant for differences in the health effects. Second, I find differences in the effects by postcode average NO2 concentration and population size. The results in panels B and C of Table A20 indicate that the effects of pollution on per capita healthcare expenditure are stronger in locations with NO2 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. The relationship between pollution exposure and health outcomes is often considered in the epidemiological literature in the form of a concentration-response functions. These function are used to quantify the health effects of air pollution and to predict the health benefits of reductions in air pollutant concentrations. Uncertainties remain regarding the shape of the concentration-response function with previous studies calling for more research in this area for more informed and effective air pollution abatement policies Pope III et al. (2015). Some studies have indeed suggested that the concentration-response function for particulate matter pollution might be supra-linear. For example, Miller et al. (2021) and Henderson et al. (2024) show that small air PM2.5 pollution shocks have proportionally larger mortality effects than large air pollution shocks. 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 A22 in the 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 O3 and PM pollution. This is consistent again with a concave relationship between air pollution exposure and health effects or concave concentration-response function.

5.3 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 NO2 concentrations leads to an average €17.23 of additional healthcare expenditure per postcode area during the same week. Similarly, a on 1 $\mu g/m^3$ increase in weekly average O3 levels leads to an average €3.28 of additional healthcare expenditure. These €17.23 and €3.28 of additional healthcare expenditure per postcode 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.[10]. These are changes in annual health expenditures for a $1\mu g/m^3$ change in air pollution concentration, which does not provide an intuitive understanding of the effect size. To better understand the magnitude of the effect, consider instead the total effect of air pollution 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 NO2 and €5,573,344,618 for O3 pollution, resulting in an overall effect of €12,816,392,490[11]. 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[12].

[10]To illustrate, consider the calculation for NO2: €17.23 · 97 · 52 · 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 postcode areas to obtain the effect for France.

[11]$13.78 \cdot €525,620,310 = €7,243,047,872$ for NO2 and $55.7 \cdot €100,060,047 = €5,573,344,618$ for O3 pollution.

[12]In 2019, the GDP of France reached €2,425.7 billion (INSEE, 2020) and aggregate healthcare spending was €208 billion (DREES, 2020).

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 cited in the report 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. Although the study sought to assess the cost of air pollution to the French healthcare system as comprehensively as possible, it 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 evaluation of the health costs generated by air pollution has been very partial, both in the quasi-experimental literature and in the epidemiological literature. 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. For example, the Environmental Benefits Mapping and Analysis ProgramCommunity 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 (for example Birnbaum 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 is on mortality costs, an estimate of hospital costs is also provided. They estimate that a decrease in average PM 2.5 concentrations of $4.9 \mu g/m^3$ in the US between 1999 and 2013 saved hospital costs of USD 1.5 billion per year. As PM2.5 is nested withing PM10 with a correlation between both larger than 0.9 in my data, I consider that my results for the effect of PM10 should be comparable to the effects of

PM2.5 in Deryugina et al. (2019). Considering €12.3 of additional weekly healthcare costs at the postcode area level for a one unit increase in weekly average PM10 concentrations from the two pollutant FE-IV model in Table A10, scaling it to a yearly estimate for France and multiplying it by 4.9 for the change considered in Deryugina et al. (2019) yields €1,920,817,342. Scaling the cost for the French population of roughly 67 million to the population of 55 million people aged over 65 in the US yields €1,575,070,220 or USD 1,696,744,394. This estimate of overall costs for the French healthcare system is very close to the USD 1.5 billion estimate for the United States from Deryugina et al. (2019), which only considers hospital costs. While it is difficult to compare healthcare costs between the United States and France, and between the US elderly and the general French population, the comparison suggests that my estimates are not unrealistic. Healthcare costs are on average much higher in the US than in France[13], which could explain why my estimate of total healthcare costs in France is not much higher than the estimate of only the hospital costs in the US.

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 (NEC, 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 NO2 but also the cost of reducing other pollutants. Compliance with the NEC Directive requires France to reduce nitrogen oxides (NOx, composed of both NO2 and NO) by 50% compared to 2005 values, to be achieved from 2030. In 2005, annual NO2 concentrations in France were 17.5 $\mu g/m^3$ (INERIS, 2024), which means that France should reduce NO2 by 8.75 $\mu g/m^3$ from its 2005 levels 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 NO2 concentration. According to my estimate, this 3.26 $\mu g/m^3$ of annual NO2 concentration should lead to savings of €1.7 billion in annual health expenditure. The health cost savings from complying with the NO2 pollution reductions alone 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 NO2 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 estimate for France (home to 13% of the total EU population) for compliance regarding NO2 standards disregarding reductions of other air pollutant concentrations already corresponds to 70% of the health costs considered in Amann et al. (2017), suggesting that the health costs considered in Amann et al. (2017) are largely underestimated.

[13]The United States has the most expensive healthcare system of any country. For example, 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. Information availablehere.

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. First, the estimates only reflect the short-term effects of air pollution, not the effects of chronic exposure, as my identification strategy is based on short-term fluctuations in air pollution concentrations. This may mean that a significant portion of health costs are not considered. One example are treatment costs related to cancers. Air pollution may be linked to 0.5-1% of all cancer cases in Europe (Couespel and Price, 2020) and to over 7% of lung cancer cases (Kulhanova et al., 2018). Second, the cost estimates do not include the costs of 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). I cannot account for this kind of avoidance behaviour as it remains unobserved.Third, 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. Many model specifications produce larger cost estimates, such as the population size-weighted regressions.

The results of this study provide highly relevant information for public policy decisions. Although the cost estimates still represent a lower limit for the total health costs, they 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. Subsequent cost-benefit studies should include these considerable costs. The health costs estimated here are caused by air pollution levels that are mostly far below the current European air quality guideline values (see Figure A1 in the Appendix), indicating that the current guideline values, which are supposedly safe limits for human health, 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 NO2 from an annual average of 40 $\mu g/m^3$ to 10 $\mu g/m^3$, for PM10 from 40 $\mu g/m^3$ to 15 $\mu g/m^3$ and for PM2.5 from 25 $\mu g/m^3$ to 5 $\mu g/m^3$. However, this study provides evidence that there are likely significant health benefits from reducing pollutant levels even further below the current WHO guideline values. While air pollution has a large effect on overall spending in more populated urban areas due to the higher number of affected people, the marginal effects appear to be 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 NO2 and 15 $\mu g/m^3$ annual mean for PM10 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 A23 in the appendix. Even the most stringent 2021 WHO guideline values should not be considered safe for human health. In addition to a revision of the guideline values for air quality, this fact also has implications for public health communication. Usually, warning messages are addressed to the population on days with peak levels of air pollution. Further communication should inform the population about the quality of air pollution even on days with low pollution levels and point out that pollution can have a significant impact on health even at low concentrations. Public health warnings have also typically targeted populations considered to be particularly vulnerable, such as the elderly, children and pregnant women or people with chronic health conditions, as previous studies found that mortality impacts were concentrated in these groups. The present study provides evidence of health effects in all age groups, which may also justify updating public health messages to inform the population that health effects can occur even in apparently healthy adults. Finally, while I find that the health costs of short-term air pollution exposure alone are large enough to motivate further reductions in air pollution concentrations, the effects of chronic exposure to air pollution may be even more important in terms of overall public health significance(Pope III et al., 2009). I therefore recommend further research into the long-term effects.

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. Consistent with evidence from the economic and epidemiological literature, I find that pollution affects spending in the specialties of family medicine, cardiology and vascular medicine, pulmonary medicine, otolaryngology, ophthalmology and gynaecology, while no effects are found in medical specialties that should not be affected, such as plastic surgery and trauma surgery. 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 less frequently studied. While air pollution in more populated urban areas affects a larger number of people and therefore 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. The health costs of air pollution for the French health system have been greatly underestimated. 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, which are supposed to be safe limits for human health, 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. The results of this study also suggest that public health messages need to be updated to inform the population that health effects occur at low levels of air pollution and affect people of all ages, as opposed to the current warnings, which are usually issued on peak air pollution days and target groups that are considered particularly vulnerable, such as the elderly, children and pregnant women.

References

(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`.

Ahrens, A., Hansen, C. B., and Schaffer, M. E. (2020). lassopack: Model selection and prediction with regularized regression in stata. *The Stata Journal*, 20(1):176–235.

Amann, M., Holland, M., Maas, R., Vandyck, T., and 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*, pages 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., and 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(591):257–280.

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

Belloni, A., Chen, D., Chernozhukov, V., and Hansen, C. (2012). Sparse models and methods for optimal instruments with an application to eminent domain. *Econometrica*, 80(6):2369–2429.

Bezin, J., Duong, M., Lassalle, R., Droz, C., Pariente, A., Blin, P., and Moore, N. (2017). The national healthcare system claims databases in france, sniiram and egb: powerful tools for pharmacoepidemiology. *Pharmacoepidemiology and drug safety*, 26(8):954–962.

Birnbaum, H. G., Carley, C. D., Desai, U., Ou, S., and 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(12):2113–2119.

Chan, T.-C., Pai, C.-W., Wu, C.-C., Hsu, J. C., Chen, R.-J., Chiu, W.-T., and Lam, C. (2022). Association of air pollution and weather factors with traffic injury severity: a study in taiwan. *International journal of environmental research and public health*, 19(12):7442.

Chay, K. Y. and Greenstone, M. (2003). The impact of air pollution on infant mortality: evidence from geographic variation in pollution shocks induced by a recession. *The quarterly journal of economics*, 118(3):1121–1167.

Chen, Y., Ebenstein, A., Greenstone, M., and Li, H. (2013). Evidence on the impact of sustained exposure to air pollution on life expectancy from chinas huai river policy. *Proceedings of the National Academy of Sciences*, 110(32):12936–12941.

Clapp, L. J. and 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(36):6391–6405.

Couespel, N. and Price, R. (2020). Strengthening europe in the fight against cancer. *European Union: Brussels, Belgium*.

Currie, J. and Neidell, M. (2005). Air pollution and infant health: what can we learn from california's recent experience? *The Quarterly Journal of Economics*, 120(3):1003–1030.

Currie, J. and Walker, R. (2011). Traffic congestion and infant health: Evidence from e-zpass. *American Economic Journal: Applied Economics*, 3(1):65–90.

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

Deryugina, T., Heutel, G., Miller, N. H., Molitor, D., and Reif, J. (2019). The mortality and medical costs of air pollution: Evidence from changes in wind direction. *American Economic Review*, 109(12):4178–4219.

Deschênes, O., Greenstone, M., and Shapiro, J. S. (2017). Defensive investments and the demand for air quality: Evidence from the nox budget program. *American Economic Review*, 107(10):2958–89.

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`.

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 (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. What is ozone. `https://www.epa.gov/ozone-pollution-and-your-patients-health/what-ozone`. Accessed: 2021-10-21.

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`.

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

Friedman, M. S., Powell, K. E., Hutwagner, L., Graham, L. M., and Teague, W. G. (2001). Impact of changes in transportation and commuting behaviors during the 1996 summer olympic games in atlanta on air quality and childhood asthma. *Jama*, 285(7):897–905.

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

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

Gramsch, E., Cáceres, D., Oyola, P., Reyes, F., Vásquez, Y., Rubio, M., and Sánchez, G. (2014). Influence of surface and subsidence thermal inversion on pm2. 5 and black carbon concentration. *Atmospheric Environment*, 98:290–298.

Halliday, T. J., Lynham, J., and De Paula, Á. (2019). Vog: Using volcanic eruptions to estimate the health costs of particulates. *The Economic Journal*, 129(620):1782–1816.

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

Imbens, G. and Angrist, J. (1994). Identification and estimation of local average treatment effects. *Econometrica*, 62(2):467–475.

INERIS (2024). La qualité de lair en france métropolitaine cartographiée de 2000 aujourdhui par lineris. `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`.

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

Ito, K. and Zhang, S. (2020). Willingness to pay for clean air: Evidence from air purifier markets in china. *Journal of Political Economy*, 128(5):1627–1672.

Jans, J., Johansson, P., and 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 indonesias wildfires. *Journal of Human resources*, 44(4):916–954.

Johns, D. O., Stanek, L. W., Walker, K., Benromdhane, S., Hubbell, B., Ross, M., Devlin, R. B., Costa, D. L., and Greenbaum, D. S. (2012). Practical advancement of multipollutant scientific and risk assessment approaches for ambient air pollution. *Environmental health perspectives*, 120(9):1238–1242.

Knittel, C. R., Miller, D. L., and Sanders, N. J. (2016). Caution, drivers! children present: Traffic, pollution, and infant health. *Review of Economics and Statistics*, 98(2):350–366.

Kulhanova, I., Morelli, X., Le Tertre, A., Loomis, D., Charbotel, B., Medina, S., Ormsby, J.-N., Lepeule, J., Slama, R., and Soerjomataram, I. (2018). The fraction of lung cancer incidence attributable to fine particulate air pollution in france: impact of spatial resolution of air pollution models. *Environment international*, 121:1079–1086.

Le Moullec, A. and Meleux, F. (2019). Bilan de la qualité de lair 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., and Kim, C.-H. (2021). Ozone continues to increase in east asia despite decreasing no2: Causes and abatements. *Remote Sensing*, 13(11):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(2):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., and Viaud, V. (2020). *The European Environment-state and outlook 2020*. European Environment Agency.

Manisalidis, I., Stavropoulou, E., Stavropoulos, A., and 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., and Wyzga, R. E. (2010). Is the air pollution health research community prepared to support a multipollutant air quality management framework? *Inhalation toxicology*, 22(sup1):1–19.

Miller, N. H., Molitor, D., and Zou, E. Y. (2021). A causal concentration-response function for air pollution: Evidence from wildfire smoke *.

Moretti, E. and Neidell, M. (2011). Pollution, health, and avoidance behavior evidence from the ports of los angeles. *Journal of human Resources*, 46(1):154–175.

Neidell, M. (2004). Air pollution, health, and socio-economic status: the effect of outdoor air quality on childhood asthma. *Journal of health economics*, 23(6):1209–1236.

Neidell, M. (2009). Information, avoidance behavior, and health the effect of ozone on asthma hospitalizations. *Journal of Human resources*, 44(2):450–478.

Pope III, C. A., Cropper, M., Coggins, J., and Cohen, A. (2015). Health benefits of air pollution abatement policy: Role of the shape of the concentration–response function. *Journal of the Air & Waste Management Association*, 65(5):516–522.

Pope III, C. A. and Dockery, D. W. (1999). Epidemiology of particle effects. In *Air pollution and health*, pages 673–705. Elsevier.

Pope III, C. A., Ezzati, M., and Dockery, D. W. (2009). Fine-particulate air pollution and life expectancy in the united states. *New England Journal of Medicine*, 360(4):376–386.

Rafenberg, C. (2015). Estimation des coûts pour le système de soins français de cinq maladies respiratoires et des hospitalisations attribuables à la pollution de lair. *Etud Doc*, 122:36.

Ransom, M. R. and Iii, C. A. P. (1995). External health costs of a steel mill. *Contemporary economic policy*, 13(2):86–97.

Real, E., Couvidat, F., Ung, A., Malherbe, L., Raux, B., Gressent, A., and Colette, A. (2022). Historical reconstruction of background air pollution over france for 2000–2015. *Earth System Science Data*, 14(5):2419–2443.

Schlenker, W. and Walker, W. R. (2015). Airports, air pollution, and contemporaneous health. *The Review of Economic Studies*, 83(2):768–809.

Schwartz, J., Bind, M.-A., and Koutrakis, P. (2017). Estimating causal effects of local air pollution on daily deaths: effect of low levels. *Environmental health perspectives*, 125(1):23–29.

Schwartz, J., Fong, K., and 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(8):087004.

Sénat (2015). Rapport fait au nom de la commission denqute sur le cot économique et financier de la pollution de lair. `http://www.senat.fr/rap/r14-610-1/r14-610-11.pdf`.

Simeonova, E., Currie, J., Nilsson, P., and Walker, R. (2019). Congestion pricing, air pollution, and childrens health. *Journal of Human Resources*, pages 0218–9363R2.

Tuppin, P., De Roquefeuil, L., Weill, A., Ricordeau, P., and Merlière, Y. (2010). French national health insurance information system and the permanent beneficiaries sample. *Revue d'epidemiologie et de sante publique*, 58(4):286–290.

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

Wallace, J. and Kanaroglou, P. (2009). The effect of temperature inversions on ground-level nitrogen dioxide (no2) and fine particulate matter (pm2. 5) using temperature profiles from the atmospheric infrared sounder (airs). *Science of the Total Environment*, 407(18):5085–5095.

WHO (2017). Preventing noncommunicable diseases (ncds) by reducing environmental risk factors. Technical report, World Health Organization.

Zivin, J. G. and Neidell, M. (2009). Days of haze: Environmental information disclosure and intertemporal avoidance behavior. *Journal of Environmental Economics and Management*, 58(2):119–128.

Appendix

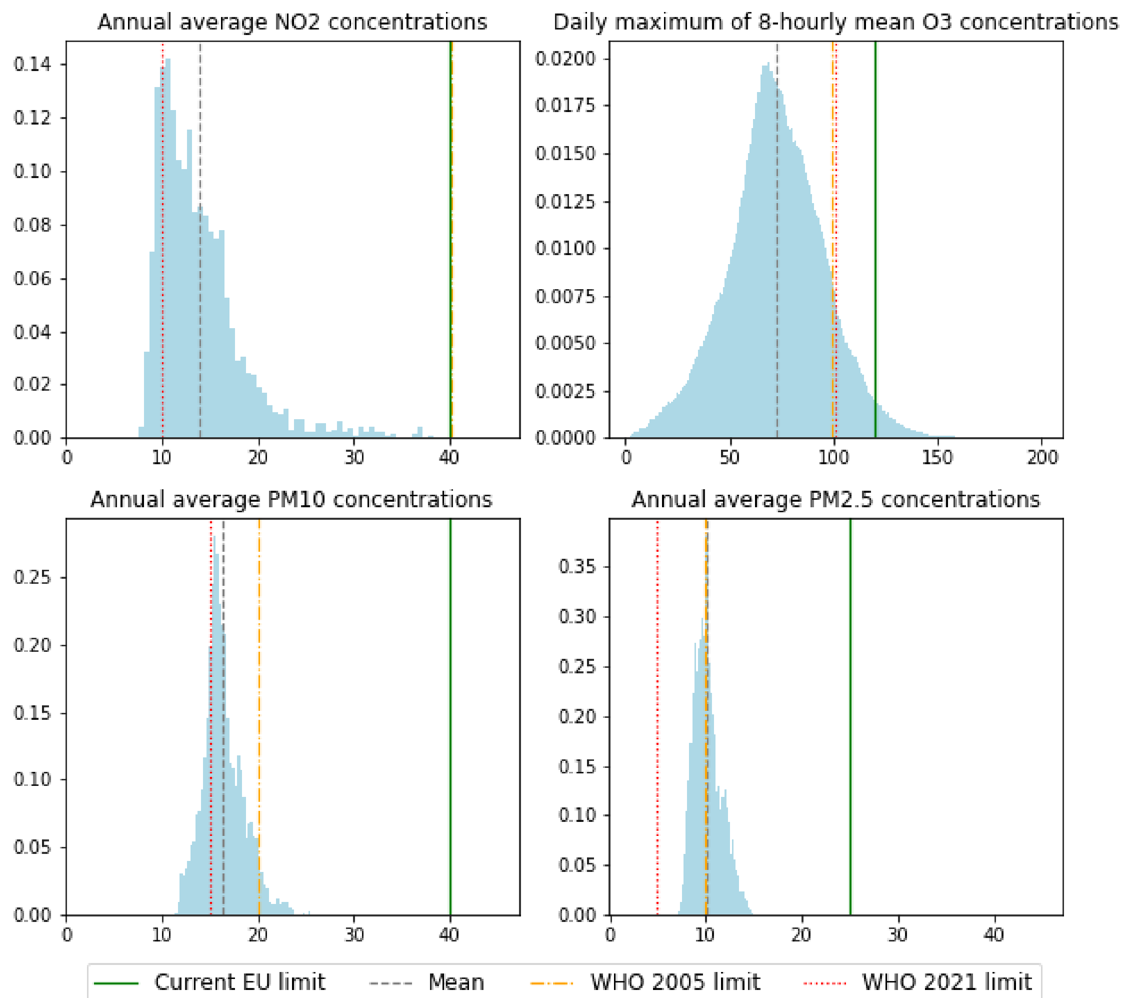

Figure A1: Level of pollutants relative to the limit values presented in Table A1. The Figure shows the distribution of pollution concentrations in 2017 across postcode areas in light blue, the current limit values in France/the EU (solid green), the average value across postcodes (grey, dashed) and 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.

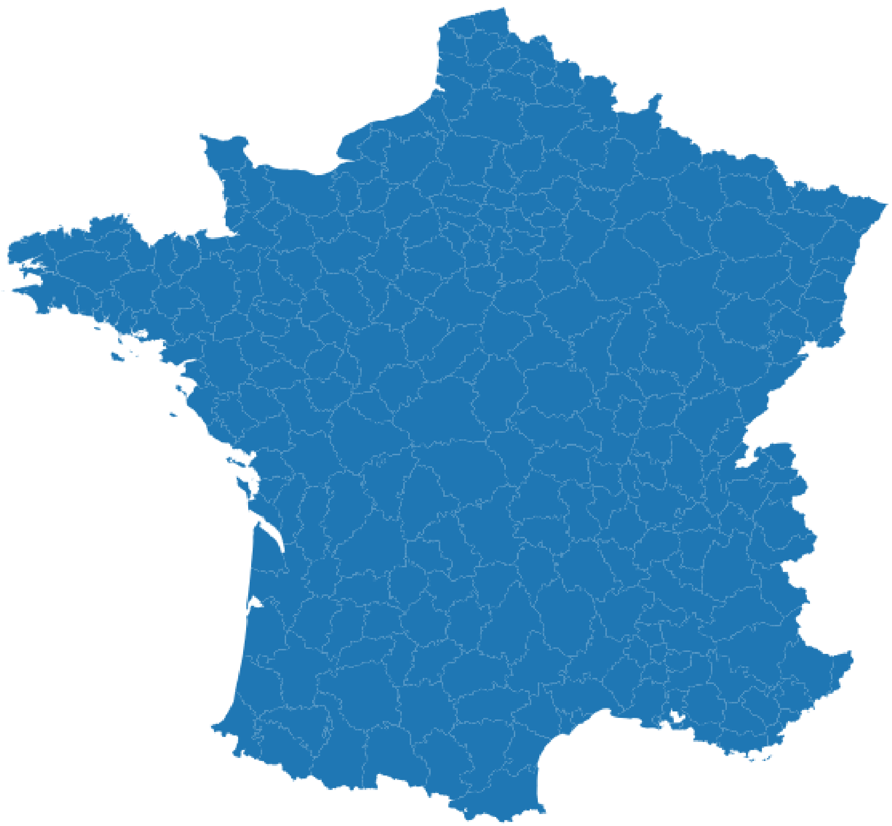

Figure A2: Division of France into employment zones. The employment zone (*"zone d'emploi"* in French) is a division of the French territory into 306 geographical areas within which most of the working population resides and works.

Table A1: Summary of the main WHO and EU Air Quality Standard values

Pollutant	Averaging time	WHO 2005 Guidelines	WHO 2021 Guidelines	EU/France current limit values
Nitrogen dioxide (NO2)	Annual	40	10	40
Nitrogen dioxide (NO2)	24-hour	N/A	25	N/A
Ozone (O3)	8-hour	100	100	120
Ozone (O3)	Peak season[a]	N/A	60	N/A
Particles $\varnothing \leq 10\mu m$ (PM10)	Annual	20	15	40
Particles $\varnothing \leq 10\mu m$ (PM10)	24-hour	50	45	50
Particles $\varnothing \leq 2.5\mu m$ (PM2.5)	Annual	10	5	25
Particles $\varnothing \leq 2.5\mu m$ (PM2.5)	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 O3 concentration in the six consecutive months with the highest six-month running- average O3 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 A2: Summary statistics - pooled postcode-week observations

Variable	Mean	Std. Dev.	Min.	Max.	N
Weekly sum of healthcare spending in €					
Total spending	3608.63	7471.21	0	370436.63	1257984
Family practice	1212.08	2502.24	0	75846.8	1257984
Cardiology and vascular medicine	50.9	184.84	0	37266.45	1257984
Pulmonology	22.76	142.19	0	15688.04	1257984
Otorhinolaryngology	19.34	74.56	0	10203.16	1257984
Ophtalmology	82.36	228.2	0	9585.98	1257984
Gynecology	43.19	150.61	0	9300.96	1257984
Nephrology	11.48	79.62	0	11234.26	1257984
Gastroenterology and hepatology	32.37	305.67	0	26562.06	1257984
Trauma surgery	36.01	164.33	0	14695.7	1257984
Plastic surgery	5.22	76.64	0	6468.49	1257984
Weekly average pollution concentrations in $\mu g/m^3$					
NO2	13.78	7.14	2.68	76.84	1247428
O3	55.7	17.49	0.4	119.09	1247428
PM10	16.61	6.32	4.28	92.76	1247428
Meteorological conditions					
Temperature (mean, C)	12.39	6.74	-17.2	31.3	1257984
Precipitation (sum, mm)	2.17	4.92	0	123	1257984
Ground-level wind speed (mean, m/s)	3.05	1.68	0	29.6	1257984
Thermal inversions (sum of occurrence)	0.36	0.94	0	7	1257984
Temperature difference (mean, C)	-1.2	0.47	-3.56	3.06	1257984
Planetary boundary height (mean, m)	539.85	186.81	12.05	1590.81	1257984

Table A3: First stage regression regression results, preferred FE-IV model specification

	Weekly mean NO2	Weekly mean O3	Weekly mean PM10
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***
	(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***
	(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 A4: Impact of average weekly NO2, O3 and PM10 pollutant concentrations on weekly healthcare expenditure - placebo regressions using shuffled instruments

	Weekly healthcare spending	
NO2	37.09	89.04
	(164.970)	(133.743)
O3	19.51	20.79
	(61.500)	(51.431)
PM10	-96.78	-3.562
	(116.092)	(105.872)
Lag NO2		206.0
		(135.917)
Lag O3		91.33
		(54.970)
Lag PM10		-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 A5: Impact of average weekly NO2, O3 and PM10 pollutant concentrations on weekly healthcare expenditure - robustness to different first stage specifications

	Weekly healthcare spending					
NO2	67.13*	34.96*	14.19**	17.18***	35.60***	26.71***
	(28.411)	(17.679)	(4.843)	(4.706)	(3.950)	(3.419)
O3	13.10***	5.033*	4.809***	3.223***	5.172***	3.308***
	(3.193)	(1.954)	(0.891)	(0.860)	(0.794)	(0.799)
PM10	-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 NO2		65.44**		13.65**		23.30***
		(21.345)		(4.781)		(4.180)
Lag O3		17.09***		6.963***		5.742***
		(2.458)		(0.895)		(0.756)
Lag PM10		-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 A6: Impact of average weekly NO2, O3 and PM10 pollutant concentrations on weekly healthcare expenditure - robustness to different fixed effects structures

	Sum of weekly healthcare spending			
NO2	12.00***	17.36***	19.29*	25.71***
	(3.450)	(3.901)	(7.950)	(3.692)
O3	4.440***	2.760***	-0.370	4.194***
	(0.617)	(0.671)	(0.653)	(0.667)
PM10	4.852	2.370	-5.493	-1.476
	(2.787)	(2.927)	(5.031)	(2.593)
Lag NO2	0.528	-3.675	7.638	-3.079
	(3.962)	(4.101)	(7.504)	(3.530)
Lag O3	4.328***	7.831***	7.449***	7.022***
	(0.647)	(0.853)	(0.948)	(0.736)
Lag PM10	15.78***	18.28***	12.99***	17.77***
	(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 and 2 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 3 shows results for a model including the non-transformed weather variables. Column 4 shows results for a model including month-by-department FE rather than month fixed effects.

Table A7: First stage using LASSO selected instruments

	Weekly mean NO2	Weekly mean O3	Weekly mean PM10
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***
		(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. The results (the second stage regression) using the LASSO-selected instruments are shown in Table A9. 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 A8: First stage model fit in terms of the Bayesian Information Criterion (BIC) and the Akaike Information Criterion (AIC)

	First stage regression - estimating NO2 pollution		
	Using NO2 instruments	Using O3 instruments	Using PM10 instruments
AIC	**5854194.7**	5857163.3	5864386.1
BIC	**5854879.1**	5857883.6	5865082.4
	First stage regression AIC and BIC estimating O3 pollution		
	Using NO2 instruments	Using O3 instruments	Using PM10 instruments
AIC	8074782.3	**7958247.7**	7974813.1
BIC	8075466.6	**7958968.0**	7975509.4
	First stage regression AIC and BIC estimating PM10 pollution		
	Using NO2 instruments	Using O3 instruments	Using PM10 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 A9: Impact of average weekly NO2, O3 and PM10 pollutant concentrations on weekly healthcare expenditure - FE-IV LASSO regression results

	Weekly healthcare spending	
Weekly mean NO2	20.40***	20.18***
	(3.881)	(3.750)
Weekly mean O3	6.177***	3.296***
	(0.783)	(0.666)
Weekly mean PM10	10.75***	1.519
	(2.839)	(2.842)
Lag weekly mean NO2		-6.877
		(4.134)
Lag weekly mean O3		7.033***
		(0.814)
Lag weekly mean PM10		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 A7. 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 A10: Impact of average weekly NO2, O3 and PM10 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					
Weekly mean NO2	30.33***			33.34***		44.33***
	(1.927)			(2.029)		(2.692)
Weekly mean O3		0.362		4.076***	0.754*	4.189***
		(0.353)		(0.381)	(0.355)	(0.383)
Weekly mean PM10			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					
Weekly mean NO2	22.71***			32.29***		18.42***
	(1.952)			(2.152)		(3.820)
Weekly mean O3		0.957		5.984***	5.477***	6.282***
		(0.680)		(0.756)	(0.789)	(0.773)
Weekly mean PM10			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
	Panel C: Location FE model, lags					
Weekly mean NO2	27.44***			31.12***		43.36***
	(1.689)			(1.814)		(2.420)
Lag weekly mean NO2	8.213***			8.322***		8.947***
	(1.518)			(1.580)		(2.119)
Weekly mean O3		0.939**		4.769***	1.256***	4.837***
		(0.339)		(0.386)	(0.342)	(0.387)
Lag weekly mean O3		-0.890*		-0.294	-0.760*	-0.175
		(0.351)		(0.362)	(0.355)	(0.364)
Weekly mean PM10			2.665***		3.022***	-13.34***
			(0.598)		(0.604)	(0.996)
Lag weekly mean PM10			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					
Weekly mean NO2	15.87***			17.85***		17.23***
	(1.805)			(2.140)		(3.719)
Lag weekly mean NO2	8.286***			19.98***		-3.423
	(1.873)			(2.082)		(4.062)
Weekly mean O3		-0.618		3.015***	2.345***	3.275***
		(0.557)		(0.653)	(0.625)	(0.662)
Lag weekly mean O3		4.699***		6.528***	6.890***	6.497***
		(0.688)		(0.775)	(0.804)	(0.795)
Weekly mean PM10			11.59***		12.85***	3.540
			(1.335)		(1.514)	(2.843)
Lag weekly mean PM10			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 A11: Impact of average weekly NO2, O3 and PM10 pollutant concentrations on weekly healthcare expenditure - instrumenting only one pollutant and including the others as controls

	Weekly healthcare spending		
NO2	32.69^{***}		
	(3.223)		
O3		5.482^{***}	
		(0.766)	
PM10			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 A12: Impact of average weekly NO2, O3 and PM10 pollutant concentrations on weekly healthcare expenditure - robustness to controlling and instrumenting for CO and SO2 pollution

	Weekly healthcare spending			
NO2	23.32***	11.69**	22.61***	27.20***
	(4.062)	(3.704)	(5.061)	(5.180)
O3	7.035***	3.115***	6.707***	4.208***
	(0.797)	(0.689)	(0.860)	(0.785)
PM10	12.64***	6.056*	12.42***	2.252
	(2.818)	(2.776)	(2.815)	(2.935)
Lag NO2		3.359		-16.56***
		(4.230)		(4.798)
Lag O3		7.413***		3.923***
		(0.873)		(0.923)
Lag PM10		16.60***		18.48***
		(2.594)		(2.569)
CO			-61.04	-1756.7***
			(294.946)	(257.931)
SO2			-43.46	-33.12
			(29.125)	(29.695)
Lag CO				1790.6***
				(257.966)
Lag SO2				73.94*
				(29.274)
CO and SO2	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 NO2, O3 and PM10 including in addition SO2 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 A13: Impact of average weekly NO2, O3 and PM10 pollutant concentrations on weekly healthcare expenditure - robustness to using EEA measuring station data

	Weekly healthcare spending			
NO2 (EEA data)	53.55***	50.23***	77.78***	76.44***
	(3.645)	(4.190)	(5.440)	(5.709)
O3 (EEA data)	10.60***	9.382***	12.86***	11.57***
	(1.188)	(1.204)	(1.281)	(1.246)
PM10 (EEA data)	-3.553	-4.618	-3.995	-2.048
	(3.836)	(2.823)	(3.904)	(2.856)
Lag NO2 (EEA data)		26.41***		27.99***
		(3.157)		(3.882)
Lag O3 (EEA data)		3.745**		1.269
		(1.303)		(1.342)
Lag PM10 (EEA data)		3.994		1.269
		(2.667)		(2.735)
CO			-1671.5***	-1800.9***
			(344.310)	(251.484)
SO2			-351.5***	-349.7***
			(36.966)	(40.012)
Lag CO				519.5*
				(253.911)
Lag SO2				-282.5***
				(38.559)
CO and SO2	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 NO2, O3 and PM10 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 SO2 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 A14: Impact of average weekly NO2, O3 and PM10 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 NO2	17.23***	25.63***	54.10***
	(3.719)	(5.503)	(7.602)
Weekly mean O3	3.275***	2.692**	8.057***
	(0.662)	(1.009)	(1.675)
Weekly mean PM10	3.540	6.277	23.63***
	(2.843)	(4.356)	(5.824)
Lag weekly mean NO2	-3.423	1.672	-29.71***
	(4.062)	(6.089)	(8.766)
Lag weekly mean O3	6.497***	8.627***	16.55***
	(0.795)	(1.377)	(1.628)
Lag weekly mean PM10	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 A15: Impact of average daily NO2, O3 and PM10 pollutant concentrations on daily healthcare expenditure

	Daily healthcare spending	
Daily NO2	3.488***	4.946***
	(0.488)	(0.522)
Daily O3	0.429**	0.692***
	(0.136)	(0.119)
Daily PM10	2.059**	-0.0566
	(0.676)	(0.466)
One day lag NO2		-2.036***
		(0.481)
Two day lag NO2		0.957*
		(0.450)
One day lag O3		-0.417**
		(0.133)
Two day lag O3		0.813***
		(0.150)
One day lag PM10		2.447***
		(0.530)
Two day lag PM10		-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 A16: Impact of average weekly NO2, O3 and PM10 pollutant concentrations on weekly healthcare expenditure at the level of the employment zone

	Weekly healthcare spending at employment zone level			
NO2	520.9*	550.8	438.3***	210.4*
	(240.411)	(353.243)	(107.389)	(104.483)
O3	45.08	7.188	83.68***	81.81***
	(39.424)	(40.754)	(16.960)	(17.488)
PM10	-115.9	-526.2	19.21	25.56
	(206.098)	(467.519)	(52.511)	(59.223)
Lag NO2		983.6		737.0
		(653.242)		(385.839)
Lag O3		145.9***		143.6***
		(42.129)		(33.353)
Lag PM10		-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 A17: Impact of average weekly NO2, O3 and PM10 pollutant concentrations on weekly healthcare expenditure - by medical specialty, location FE

	Family practice	Cardio-vascular	Pulmo.	O.R.L.	Ophthalmo.
NO2	8.912***	0.749***	0.0721	0.332***	1.187***
	(0.610)	(0.104)	(0.072)	(0.040)	(0.109)
O3	1.018***	0.0419*	0.0306	0.0383***	0.122***
	(0.121)	(0.019)	(0.017)	(0.008)	(0.021)
PM10	-2.747***	-0.131*	0.0495	-0.0827***	-0.172***
	(0.350)	(0.052)	(0.038)	(0.018)	(0.049)
Lag NO2	2.722***	0.304**	-0.118	0.146***	0.376**
	(0.586)	(0.100)	(0.075)	(0.044)	(0.114)
Lag O3	-0.557***	0.0376*	-0.0218	0.00457	0.0372
	(0.133)	(0.019)	(0.019)	(0.008)	(0.024)
Lag PM10	-0.680*	-0.127**	0.0141	-0.0449*	-0.225***
	(0.302)	(0.044)	(0.034)	(0.020)	(0.049)
	Gynaeco.	Nephro.	Gastro-hepato.	Trauma surg.	Plastic surg.
NO2	0.488***	0.0413	0.394*	0.642***	0.177***
	(0.071)	(0.044)	(0.158)	(0.092)	(0.053)
O3	0.0222	0.0279***	0.104	0.0788***	0.0210*
	(0.015)	(0.008)	(0.063)	(0.021)	(0.010)
PM10	-0.0663	-0.0378	-0.0837	-0.122*	-0.0163
	(0.035)	(0.023)	(0.089)	(0.048)	(0.024)
Lag NO2	0.317***	-0.0218	0.218	0.223*	0.0641
	(0.091)	(0.044)	(0.149)	(0.091)	(0.048)
Lag O3	0.00326	-0.00501	0.00481	0.0340	-0.0165
	(0.015)	(0.009)	(0.038)	(0.020)	(0.010)
Lag PM10	-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 A18: Impact of average weekly NO2, O3 and PM10 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	
NO2	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)
O3	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)
PM10	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 NO2		1.726		-0.00790		-4.374*		4.012*		0.0855
		(0.898)		(1.144)		(1.887)		(1.732)		(1.220)
Lag O3		1.149***		-0.0286		1.503***		2.754***		0.797**
		(0.178)		(0.229)		(0.388)		(0.361)		(0.245)
Lag PM10		1.930***		1.162		8.750***		-2.741*		2.236**
		(0.571)		(0.859)		(1.251)		(1.082)		(0.738)
Observations	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 A19: Impact of average weekly NO2, O3 and PM10 pollutant concentrations on weekly healthcare expenditure in locations with below and above postcode average income, NO2 pollution concentrations and population size

Panel A: *Heterogeneity by average postcode average income*

	Dependent variable: Weekly healthcare spending			
	Below median income	Above median income	Below median income	Above median income
NO2	15.47**	22.01***	22.04***	16.12***
	(5.937)	(4.746)	(5.992)	(4.542)
O3	5.792***	7.174***	3.171**	3.378***
	(1.279)	(0.869)	(0.974)	(0.884)
PM10	10.33*	13.40***	-1.149	4.677
	(4.301)	(3.450)	(4.322)	(3.514)
Lag NO2			-7.131	0.644
			(6.431)	(4.898)
Lag O3			5.748***	7.642***
			(1.249)	(0.977)
Lag PM10			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 NO2 concentration*

	Dependent variable: Weekly healthcare spending			
	Below median pollution	Above median pollution	Below median pollution	Above median pollution
NO2	3.197	26.95***	10.09*	21.91***
	(4.130)	(5.148)	(4.225)	(4.852)
O3	2.669***	12.27***	1.423*	6.221***
	(0.632)	(1.499)	(0.631)	(1.201)
PM10	16.11***	13.13***	6.881*	1.700
	(2.954)	(3.867)	(3.062)	(3.604)
Lag NO2			-17.48***	4.989
			(4.779)	(5.380)
Lag O3			3.159***	12.61***
			(0.708)	(1.508)
Lag PM10			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

Panel C: *Heterogeneity by postcode population size*

	Dependent variable: Weekly healthcare spending			
	Below median population	Above median population	Below median population	Above median population
NO2	2.229	26.92***	4.381	29.18***
	(2.896)	(6.368)	(2.857)	(6.134)
O3	2.531***	9.723***	1.382**	4.660***
	(0.492)	(1.444)	(0.475)	(1.197)
PM10	8.032***	17.44***	3.722	0.503
	(2.010)	(4.927)	(2.059)	(4.725)
Lag NO2			-8.585**	-2.112
			(3.083)	(6.571)
Lag O3			1.725**	11.19***
			(0.559)	(1.458)
Lag PM10			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 A20: Impact of average weekly NO2, O3 and PM10 pollutant concentrations on weekly *per capita* healthcare expenditure in locations with below and above postcode average income, NO2 and population size

Panel A: *Heterogeneity by average postcode income*

	Dependent variable: Weekly per capita healthcare spending			
	Below median income	Above median income	Below median income	Above median income
NO2	0.0584	0.0750	0.171*	0.0699
	(0.070)	(0.039)	(0.070)	(0.040)
O3	0.0441***	0.0492***	0.0210	0.0230**
	(0.012)	(0.009)	(0.012)	(0.008)
PM10	0.117*	0.133***	-0.0197	0.0557
	(0.050)	(0.029)	(0.049)	(0.029)
Lag NO2			-0.166*	-0.0621
			(0.073)	(0.038)
Lag O3			0.0424**	0.0423***
			(0.013)	(0.009)
Lag PM10			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 NO2 concentration*

	Dependent variable: Weekly per capita healthcare spending			
	Below median pollution	Above median pollution	Below median pollution	Above median pollution
NO2	0.0941	0.0934**	0.194*	0.0764*
	(0.081)	(0.032)	(0.083)	(0.034)
O3	0.0357*	0.0475***	0.0143	0.0216*
	(0.017)	(0.009)	(0.014)	(0.008)
PM10	0.115*	0.0872***	-0.0331	0.0382
	(0.056)	(0.023)	(0.058)	(0.025)
Lag NO2			-0.169	-0.0496
			(0.105)	(0.031)
Lag O3			0.0321	0.0486***
			(0.017)	(0.009)
Lag PM10			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

Panel C: *Heterogeneity by postcode population size*

	Dependent variable: Weekly per capita healthcare spending			
	Below median population	Above median population	Below median population	Above median population
NO2	0.0640	0.0641*	0.144	0.0663*
	(0.082)	(0.028)	(0.082)	(0.028)
O3	0.0564**	0.0308***	0.0306*	0.00997
	(0.019)	(0.006)	(0.015)	(0.006)
PM10	0.165**	0.0868***	0.0458	0.0206
	(0.059)	(0.021)	(0.058)	(0.022)
Lag NO2			-0.298**	0.000430
			(0.096)	(0.027)
Lag O3			0.0268	0.0476***
			(0.019)	(0.006)
Lag PM10			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 A21: Impact of average weekly NO2, O3 and PM10 pollutant concentrations on weekly healthcare expenditure - population weighted regressions

	Weekly healthcare spending			
	FE		FE-IV	
NO2	136.7***	185.4***	72.87*	124.3***
	(16.092)	(25.026)	(34.575)	(36.172)
O3	14.09***	28.61***	48.89**	25.96**
	(1.996)	(6.256)	(15.576)	(8.289)
PM10	-38.57***	-61.09***	59.05	-19.17
	(4.494)	(9.340)	(34.057)	(33.881)
Lag NO2		40.66**		-37.54
		(14.492)		(42.059)
Lag O3		7.899		41.51*
		(5.687)		(17.188)
Lag PM10		-6.436		91.13**
		(6.504)		(31.207)
Observations	1209572	1145103	1167556	1145103
First-stage F-stat			560.2	260.9

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$. This table present results from FE-IV regressions weighted by population size. 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 A22: Impact of average weekly NO2, O3 and PM10 pollutant concentrations on weekly healthcare expenditure - heterogeneity by pollution quartile, piece-wise linear regression

	Weekly healthcare expenditure
NO2 x first NO2 quartile	200.6***
	(42.270)
NO2 x second NO2 quartile	102.0***
	(19.183)
NO2 x third NO2 quartile	103.6***
	(22.807)
NO2 x fourth NO2 quartile	61.74***
	(9.416)
O3 x first O3 quartile	38.78***
	(5.548)
O3 x second O3 quartile	35.49***
	(3.213)
O3 x third O3 quartile	25.99***
	(3.055)
O3 x fourth O3 quartile	23.76***
	(2.307)
PM10 x first PM10 quartile	215.2***
	(24.213)
PM10 x second PM10 quartile	120.6***
	(16.096)
PM10 x third PM10 quartile	102.2***
	(15.918)
PM10 x fourth PM10 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 A23: Impact of average weekly NO2, O3 and PM10 pollutant concentrations above and below the annual average WHO limit value on weekly healthcare expenditure

	Weekly healthcare spending	
NO2 x below WHO limit	48.53***	30.33**
	(11.972)	(11.714)
NO2 x above WHO limit	29.96***	18.80***
	(4.882)	(4.645)
PM10 x below WHO limit	84.70***	11.20
	(11.755)	(8.458)
PM10 x above WHO limit	43.08***	6.981
	(5.874)	(4.445)
Lag NO2 x below WHO limit		45.34***
		(11.832)
Lag NO2 x above WHO limit O		10.97*
		(5.159)
Lag PM10 x below WHO limit		12.58
		(8.532)
Lag PM10 x above WHO limit		13.33**
		(4.163)
O3	9.471***	3.482***
	(0.903)	(0.757)
Lag O3		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.