The Cost of Air Pollution for Workers and Firms: Evidence from Sickness Leave Episodes*

Work in progress, please ask before sharing or citing

Marion Leroutier[†], Hélène Ollivier[‡]

This version: May 2022

Abstract

Poor air quality is known to be bad for health. How do air pollution's health effects on workers translate into economic costs? And do firms adapt via their payroll decisions? We answer these questions by combining French administrative data on sickness leave episodes and workers' flows with fine-grained pollution and weather data. We exploit short-term variations in wind direction as an instrument for exposure to particulate matter ($PM_{2.5}$) pollution. A one standard deviation increase in weekly $PM_{2.5}$ increases the share of workers starting a sickness leave that week by 4.5%. Respecting the pollution thresholds recommended by the World Health Organization - which implies decreasing pollution by 18% on average over the study period - would have avoided 3 million days of sickness leave every year. This would have saved an annual €88 million in publicly-funded benefits, €110 million in employer-funded benefits, and €391 million in foregone production valued at the wage level. According to preliminary analyses at the establishment level, a one standard deviation increase in monthly $PM_{2.5}$ decreases new workers' inflow rate in the current and following two months by 19%. The effect is driven by a decrease in both permanent hires and in the transfer of workers from other establishments of the same firm.

Keywords: air pollution, instrumental variable, absenteeism, labour market
JEL codes: Q53, I1, J22

*We thank Augustin Colette and Elsa Real from INERIS for sharing the air pollution data for France with us, Camille Regaert from IRDES for answering our questions on the Hygie data, and Tom Verrier for outstanding research assistance.

[†]Misum, Stockholm School of Economics. Email: marion.leroutier@hhs.se

[‡]Paris School of Economics, CNRS.

1 Introduction

It is widely acknowledged that air pollution has detrimental effects on human health.[1] Even in Europe, where air pollution has been regulated for several decades, an annual 307,000 premature deaths are attributed to $PM_{2.5}$ pollution (European Environment Agency, 2020). Many papers have documented a plausibly causal relationship between exposure to air pollution and a variety of health outcomes, such as emergency admissions (Schlenker and Walker, 2016; Deryugina et al., 2019), medical expenditures (Deschênes et al., 2017), sickness leave (Holub et al., 2021), or mortality (Deryugina et al., 2019). Furthermore, people breathing polluted air may suffer from a wide range of diseases, which may impair their productivity (Graff Zivin and Neidell, 2012; Chang et al., 2016; Lichter et al., 2017; Meyer and Pagel, 2017; He et al., 2019; Chang et al., 2019; Fu et al., 2021; Adhvaryu et al., 2022), reduce their labour supply (Aragón et al., 2017), and lead to aggregate economic costs (Dechezleprêtre et al., 2019).[2]

In this paper, we jointly examine how air pollution affects the incidence of sickness leave among French workers, and the economic and employment outcomes of firms employing them. To the best of our knowledge, this paper is the first to jointly consider worker-level and firm-level responses to a pollution shock in the context of a developed country with institutionalized sickness leave. Recent papers focusing on developing countries or settings where workers are paid by the hour found that a pollution-induced negative health shock impacts workers' productivity primarily via a decrease in output per hour, rather than a decrease in the number of hours worked (Chang et al., 2016; Adhvaryu et al., 2022). In a developed country with institutionalised sickness leave, being absent from work is not very costly, so we expect that the channel will differ. Pollution shocks will likely increase workers' propensity for calling in sick, while potentially also affecting the productivity of

[1]According to the World Health Organization (WHO) (WHO, 2014), air pollution is the world's largest single environmental health risk. Exposure to fine particulate matter ($PM_{2.5}$), for instance, causes approximately 4.2 million premature deaths every year globally.

[2]For instance, Dechezleprêtre et al. (2019) found that a 10% increase in $PM_{2.5}$ pollution reduces real GDP by 0.8% in Europe.

those workers not calling in sick. Firms are likely to see their output and profits decrease due to the pollution-induced productivity shock. Firms may also try to adapt to the absenteeims and the production shock by adjusting their employment decisions in the short term.

We pair detailed data on sickness leave episodes from a random sample of French private sector employees with data on payroll decisions for a subset of establishment employing them, and granular measures of air pollution exposure at the location of the establishment where workers are employed. We plan to add monthly data on sales from mandatory VAT declarations, observed at the firm level, to capture the effect of pollution on sales.

We begin by showing that fairly large air pollution shocks in fine particulate matter ($PM_{2.5}$), exceeding WHO recommended thresholds, are quite common in our setting, and that workers take sickness leaves in response to these shocks. We rely on a panel of roughly 400,000 employed individuals followed over the 2009-2015 period, with records of all sickness leave episodes at the daily level. Workers are geolocated based on their workplace's post-code and attributed a certain level of air pollution exposure from their workplace location, using gridded reanalysis pollution data. To recover plausibly causal estimates, we adopt an instrumental variable approach exploiting within-grid-cell changes in wind direction, similar to Deryugina et al. (2019).

We find that a one standard deviation change in weekly exposure to $PM_{2.5}$ increases the share of workers starting a sickness leave that week by 4.5%. We calculate the benefits in terms of avoided sickness leave payment if France had respected the daily exposure threshold recommended by the World Health Organization between 2009 and 2015 - this corresponds to a decrease by 18% of observed pollution levels. Based on our estimates, respecting the thresholds would have avoided at least 2.9 million days of sickness leave every year, saving €88 million in publicly-funded sickness benefits (1.3% of total sickness leave payments over the period) and €110 million in employer-funded sickness benefits.

Crucial to our research question is the ability to link workers taking sick leaves to their employer's characteristics and decisions. In our matched dataset, we identify the exact establishment where individuals work. Thus, we conduct an analysis at the establishment level identifying firms' response to the pollution shock in terms of direct effect on sales, and payroll and hiring vs firing decisions. The analysis on sales is ongoing. Regarding payroll adjustments, we find that a one standard deviation change in monthly exposure to $PM_{2.5}$ decreases the inflow rate of new workers by 19% in the current and following two months, while there is not significant effect on the outflow rate. The effect is driven by a decrease in permanent hires and transfers of workers from other establishments. One interpretation is that establishments faced with a productivity shock freeze new hires rather than adjusting pollution-induced absenteeism with new workers.

The primary contribution of this study is to the literature examining the effect of air pollution on economic activity. A growing literature examines how pollution impacts workers, in terms of health-related absenteeism (Holub et al., 2021), productivity (Graff Zivin and Neidell, 2012; Chang et al., 2016; Lichter et al., 2017; Meyer and Pagel, 2017; He et al., 2019; Chang et al., 2019; Adhvaryu et al., 2022), decision-making (Meyer and Pagel, 2017; Dong et al., 2019; Aguilar-Gomez et al., 2022) and labour supply (Aragón et al., 2017). Some of these papers consider firm-level consequences of such worker-level effects, but they are often based on a small and non-representative sample. To our knowledge, only one paper, by Fu et al. (2021), examines firm-level productivity effects of pollution for a large representative sample of firms. But it is based on a high-pollution context (China) and only examines manufacturing firms; firms' margins of adjustment and the estimated elasticity of output to pollution are probably different in a developed, low-pollution country.

Furthermore, our paper contributes to the burgeoning economic literature that uses micro-level data to estimate the cost of air pollution. Such cost estimates are essential for policy appraisals and cost-benefit analysis more broadly. In many developed countries where air pollution levels have decreased in the past two decades, it is not necessarily clear that further emission reductions would bring net benefits. Accurately estimating how much air pollution costs to society today is all the more needed. Recent plausibly causal estimates are almost exclusively based on the health costs for the individuals. For example, Deryugina et al. (2019) estimate the benefits from pollution reductions induced by the US Clean Air Act to around $24 billion annually, based on the monetary value of the reduced number of years lost for the elderly. Based on social security data, (Mink, 2022) estimate that meeting European standards for nitrogen dioxide (NO_2) would save an annual 5.2 billion in healthcare costs in France. Observing sickness leave and its effects on firms allows us to estimate an additional component of the benefits from avoided pollution, namely the avoided detrimental effects on workers and firms.

Closer to this aim, Holub et al. (2021) examine the benefits of recent improvements in air quality in Spain in terms of avoided production loss, also using sickness leave data. Their analysis is based on PM_{10}, while we focus on $PM_{2.5}$, a pollutant with more severe health effects and more likely to affect individuals in their working environment Krebs et al. (2021). Holub et al. (2021) proxy the cost of a day of sickness leave with the daily wage. In contrast, we are able to decompose the cost in several components: the cost borne by taxpayers in terms of publicly funded sickness leave benefit spending; that borne by firms in terms of privately-funded sickness leave benefit spending; and the cost borne by firms in terms of foregone sales (this third analysis is ongoing). Foregone sales are likely to differ from the foregone production proxied with the daily wage for two reasons: first, within an establishment, the workers who do not enter sickness leave are presumably also affected by pollution, both in a direct and indirect way: even if they do not take sickness leave, their health and cognitive performances are probably also directly affected by pollution. At the same time, their productivity and labour supply may be indirectly affected by the absenteeism of their co-workers via spillover effects or conscious task reallocation - as evidenced in Adhvaryu et al. (2022). Second, as discussed, firms may cushion the effect of absenteeism on production via their payroll decisions.

The paper is organised as follows. Section 2 presents the French setting with air pollution regulation and institutionalised sickness leave. Section 3 presents the data, section 4 describes the empirical strategy, section 5 presents the main results, and section 6 some robustness checks.

2 Background

2.1 Particulate matter air pollution in France

Five pollutants enter the air quality index that is used to monitor local air quality on a daily basis in France:[3] particulate matter with a diameter below 2.5 micrometers ($PM_{2.5}$), particulate matter with a diameter below 10 micrometers (PM_{10}), nitrogen dioxide (NO_2), sulphur dioxide (SO_2), and ozone (O_3). While each of these pollutants has some detrimental effects on human health, fine particulate matter ($PM_{2.5}$) is the one driving the most significant health problems (European Environment Agency, 2020). According to epidemiological evidence, short- and long-term exposure to $PM_{2.5}$ is associated with increased mortality and cardiovascular diseases, even at low levels of exposure. While $PM_{2.5}$ is included in PM_{10}, the former is deadlier because smaller-sized particles penetrate deeper into the respiratory system. Furthermore, $PM_{2.5}$ can easily penetrate indoors and affect indoor exposure to pollution (Krebs et al., 2021), while PM_{10} cannot (Thatcher and Layton, 1995; Vette et al., 2001). There is growing evidence that health effects from $PM_{2.5}$ manifest even after very short-term exposure, such as the hourly level (World Health Organization, 2016). Focusing on these short-term effects, Deryugina et al. (2019) found that in the US, a 1 µg/m^3 increase in $PM_{2.5}$ exposure for one day causes 0.69 additional deaths per million elderly individuals over the three following days.

$PM_{2.5}$ can be emitted directly or formed in the atmosphere. "Primary" particles are directly emitted from different sources: in France, in 2015 52% of such $PM_{2.5}$ emissions come from the residential and tertiary sector (in particular from heating and wood burning), 20% from transport, 18% from manufacturing and 11% from agriculture (CITEPA, 2021).

[3]See https://atmo-france.org/lindice-atmo/ for more information.

"Secondary" particles form in the atmosphere and result from the chemical reactions of gaseous pollutants, including SO_2 and NO_2. $PM_{2.5}$ can travel far (hundreds of kilometres) and remain in the atmosphere for a long period of time (US EPA, 2018).

While the severe effects of $PM_{2.5}$ justify our focus on this specific pollutant, the effect of other pollutants are worth mentioning, especially since pollutants are often correlated (see section 4 for a discussion of these correlations in the context of our study). Short-term exposure to NO_2 causes respiratory issues, and long-term exposure is associated with respiratory diseases such as asthma and a potential increase in lung cancer. Short-term exposure to SO_2 also causes respiratory issues, but evidence on the effects of long-term exposure are more mixed. For ozone, evidence on the effect of short-term exposure is mixed - potentially due to the confounding effect of other pollutants in a multi-pollutant context. Long-term exposure has been linked with cardiovascular, reproductive and developmental effects, and to an increased risk of cancer and total mortality (World Health Organization, 2016).

In France, air quality is regulated via command-and-control regulation taking the form of maximum concentration thresholds, both at the annual and 24-hour level. The legal thresholds are defined at the European level and transposed into French law.[4]

Table 1 shows that the European standards were higher than the WHO recommendations for annual exposure to $PM_{2.5}$. The average annual exposure in France was below the European standards but more than twice higher than the WHO recommended thresholds in 2015. It was also 50% higher than exposure in the US, comparable to exposure in Germany, and much lower than exposure in India and China.

Figure 1 shows aggregate trends in $PM_{2.5}$ concentrations between 2009 and 2015, using the pollution data presented in the next section. While concentrations have been decreasing over time, they remained higher than the WHO threshold of 5 µg/m^3 in almost all grid cells

[4]The French government must comply with these thresholds or risk incurring sanctions: in 2020, France has been referred to the Court of Justice of the European Union for exceeding the daily thresholds for particulate matter PM_{10} (European Commission, 2020)

Table 1: PM2.5: legal standards and population-weighted exposure in 2015

		Annual (µg/m^3)	**24-hour** (µg/m^3)
WHO recommendation		5	15
European standard		25	-
2015 exposure: France		12.5	
	Germany	12.7	
	United States	8.2	
	India	67.2	
	China	50.3	

Notes: source for WHO standards: https://apps.who.int/iris/handle/10665/345329 ; source for European standards: https://www.eea.europa.eu/themes/air/air-quality-concentrations/air-quality-standards; source for population-weighted exposure in 2015: https://www.who.int/data/gho/data/themes/air-pollution/modelled-exposure-of-pm-air-pollution-exposure

in 2015. Concentrations are especially high in areas such as the Centre-North (Paris area and further North), the East (Alsace region) and the South-East (near the Alps and Rhone region). Figures A.1 to A.3 show similar figures for the other pollutants.

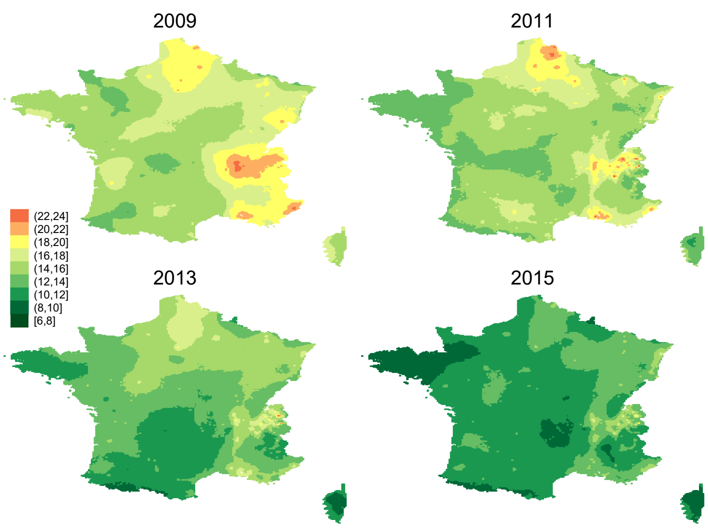

Figure 1: Average annual concentrations of $PM_{2.5}$ (µg/m^3)

2.2 Sickness leave in France

In France, any worker who does not show up to work must provide a justification for her absence. An absence caused by a sickness must be justified by a medical certificate, which must indicate the length of the leave and can be renewed if necessary. All private sector employees are eligible to publicly funded sickness leave benefits (hereafter SLB) starting from the fourth day of the sickness leave episode (hereafter SLE), as long as (i) they provide a medical certificate; (ii) they have worked at least 150 hours in the past three months. The daily SLB amounts to 50% of the daily gross wage, with a salary cap fixed at 1.8 times the daily equivalent of the minimum wage. In 2015, the implied daily SLB cap was €43. The SLB can be received for up to 365 days (and up to 3 years in the case of a long-standing disease).

In addition to the publicly funded benefit, there exists an employer-funded allowance with two components, one mandatory and one optional. The mandatory allowance is paid to workers having worked at least one year for the company from day 8 of the sickness leave. The allowance initially represents 40% of the daily gross wage, and decreases to 16% after a period of between 30 and 90 days, depending on the workers' seniority in the firm. The allowance is paid for a maximum of between 60 and 180 days, again depending on the worker's seniority.

The optional allowance is negotiated in collective agreements at the firm- or industry-level. It can cover any amount not already covered by the public or mandatory employer-funded allowances. It is paid by the employer via an insurance fund (*prévoyance entreprises*). According to survey evidence in Pollak (2015), two-thirds of private sector employees receive this optional allowance, which most of the time guarantees a 100% replacement rate: typically the allowance covers 100% of the daily wage in the first three days of leave, and the difference between the daily wage and the other benefits from day 4 on.

To sum up, for a sickness leave episode of 29 days - the average duration in our sample -, a worker working at a firm with an optional allowance receives 100% of her wage for the whole period. A worker working at a firm with no optional allowance receives nothing between day 1 and day 3, the minimum between 50% of her wage and 43€ between day 4 and day 8, and the minimum between 90% of her wage and 43€+40% of her wage from day 8. Compared to countries without sickness leave coverage, the income loss associated with being sick is low, especially for workers with generous collective agreements and optional allowance. This might give workers an incentive to take sickness leave even with mild symptoms. In our sample, 21% of workers have at least one SLE during a given year.

3 Data

3.1 Sickness leave episodes

We obtain data on sickness leave episodes (SLE) from the Hygie dataset. Hygie is a panel dataset made of a random sample of former and current private sector employees born between 1935 and 1989, having worked at least for one quarter. It combines administrative data on health from the organization managing the public health insurance (CNAM) with administrative data on employees' careers from the organization managing the public pension system (CNAV). We follow roughly 900,000 employees during the period 2009-2015.

For each individual, we know the exact start date and duration of each SLE, as well as some characteristics such as age, gender, annual wage, contract type, and annual medical expenditures. If the individual is employed, we have information on her exact workplace via an establishment-level identifier used by all French administrations and called SIRET. This has two advantages. First, it allows us to precisely geolocate the workplace and allocate pollution exposure. Second, we can use this identifier to combine the Hygie panel with other establishment-level data sets.

To build our sample of analysis, we make three restrictions. First, we only keep individuals to whom we are able to assign a place of work based on the establishment's unique identifier[5]. Second, we discard individuals whose establishment identifier corresponds to a

[5]This makes us discard two types of individuals: first, individuals with no employment history declared between 2009 and 2015, who represent 25% of the sample. Although we cannot check the exact reason (apart from those retired), these individuals are probably retired, unemployed or out of the labour force over the whole period. Two-thirds of them should be retired in 2009 given their age; second, individuals for whom we do not have an establishment identifier despite the fact that they did work and contribute to the pension system over the 2009-2015 period, who represent 6% of the sample. Two third of these

public institution such as hospital or schools, because we want to focus the analysis on private sector employees[6]. Third, we discard the few individuals who did not work enough to contribute to the public pension system for any of the years included in the period[7]. We assign each individual to the postcode of her workplace (there are around 6,000 postcodes in France). Figure A.5 shows the geographic distribution of the employees' workplaces in 2009, which is consistent with the distribution of the French population across the territory.

3.2 Workers' flows

We use administrative data on workers' flows called DMMO. All establishments having more than 50 employees must report all workers' inflows and outflows every quarter as well as the type of inflow/outflow. Total inflows are the sum of hires on a permanent contract, hires on a fixed-term contract, and transfers from another establishment of the same firm. Total outflows are the sum of in the case of outflows, whether the worker resigns, is fired, retires, or whether her fixed-term contract just comes to an end). We use the exact entry and exit dates reported for each flow to build a monthly panel. The data also includes rich information on the establishment's sector of activity as well as the number of workers at the start and end of each quarter. The process used to collect this administrative data changed in 2015, so we only include the 2009-2014 period in our analysis. Over that period, there appears to be some under-reporting of inflows and outflows for temporary contracts of less than a month. Accordingly, we may not be able to accurately capture firms' employment decisions for very short-term contracts. Flows of temporary workers employed via an external temp agency do not appear in the data, but establishment separately declare their stock of temp workers every quarter.

individuals have zero employers declared over the period. They may have switched to the public sector or to the agricultural sector or started their own business; or they may work in the domestic care sector, where there is no establishment-level identifier (since these workers work for private individuals).

[6]Some individuals working in these institutions have a private sector type of contract and are thus eligible to enter the Hygie sample.

[7]Each year these individuals worked less than 150 equivalent hours valued at the minimum wage per year, which is the minimum to contribute to public pension. With such a low labour supply, they are unlikely to experience sickness leave episodes.

We only include establishments reporting in DMMO which also appear in the final dataset on sickness leave (using the establishment identifier), having in mind to link the workers' absenteeism response to the establishment's employment decisions. We discard the establishments for which we do not observe any worker in the sickness leave dataset. We then build a monthly panel including, for each establishment, the stock of workers at the start of the month and flow of entries and exits broken down by cause. We assign each establishment to its postcode.

3.3 VAT declarations

We obtained the exhaustive set of French firms' monthly VAT records, where we observe the gross VAT paid every month and the components used to calculate it. Firms' sales is one of these components. The analysis of this dataset is ongoing.

3.4 Pollution and Weather

We use air pollution data from the French National Institute for Industrial Environment and Risks (INERIS), which provides gridded reanalysed historical pollution data for metropolitan France (Real et al., 2021). Using a kriging method that combines background measurements of air quality from monitoring stations and modeling with the chemistry-transport model CHIMERE, Real et al. (2021) produce hourly concentrations of $PM_{2.5}$, PM_{10}, NO_2, and O_3 with a spatial resolution of approximately 4 km x 4 km for the period 2000-2015. We refer to the grid scale of this pollution data as delimited by the "Chimere grid cell". There are 33,252 Chimere grid cells in metropolitan France.

Gridded reanalysed pollution data are better suited to capture the average pollution exposure for the residents in a grid cell than pollution-monitor readings. Indeed, monitors are often sparsely placed and may capture locally produced pollution that does not reflect the pollution exposure of all the residents in a given grid cell. By contrast, reanalysed data combine these monitor readings with a chemistry-transport model that takes account of all sources of pollution to give a measure of average exposure. As a result, these data probably suffer less from measurement error. It is particularly suited to examine $PM_{2.5}$ exposure as the network of $PM_{2.5}$ monitoring stations is even more sparse than for other pollutants in France: over the study period, there are between 62 and 105 background monitoring stations measuring $PM_{2.5}$ concentrations; between 173 and 251 measuring PM_{10}, between 318 and 385 measuring ozone, and between 282 and 337 measuring NO_2. The reanalysis takes into account the correlation between $PM_{2.5}$ and PM_{10} using a co-kriging method, allowing the $PM_{2.5}$ estimation to benefit from the higher density of PM_{10} monitoring stations.

We use gridded weather data derived from satellite observations, coming from two sources. We obtain daily average precipitations, minimum, maximum and average temperatures at the 11 km x 11 km scale from the E-OBS dataset of the Copernicus Climate Change Service (C3S), an information service provided by the Copernicus Earth Observation Programme of the European Union.[8] There are 6,453 Copernicus grid cells in metropolitan France. Finally, we obtain hourly wind speed and wind direction data at the 50 km x 50 km scale from MERRA (the Modern-Era Retrospective analysis for Research and Applications), a NASA research project providing satellite-derived reanalysis data.[9] There are 247 Merra grid cells in metropolitan France. We average wind speed at the daily level. For wind direction, we sum the number of hours in the day where wind blows in a given direction, with four directions: North (direction below 45°or above 315°), East (between 45° and 135°), South (between 135° and 225°) and West (between 225° and 315°). In total, we have three different grid cell sizes across the different weather and pollution data: from the smallest to the largest, Chimere (pollution), Copernicus (temperature and precipitation) and Merra

[8]We acknowledge the E-OBS dataset from the EU-FP6 project UERRA (https://www.uerra.eu) and the Copernicus Climate Change Service, and the data providers in the ECA&D project (https://www.ecad.eu). The data can be downloaded from: https://surfobs.climate.copernicus.eu/dataaccess/access_eobs.php

[9]We acknowledge using the data from Global Modeling and Assimilation Office (GMAO) (2015), MERRA-2 tavg1_2d_flx_Nx: 2d, 1-Hourly, Time-Averaged, Single-Level, Assimilation, Surface Flux Diagnostics V5.12.4, Greenbelt, MD, USA, Goddard Earth Sciences Data and Information Services Center (GES DISC), Accessed: January 2021, 10.5067/7MCPBJ41Y0K6. See https://disc.gsfc.nasa.gov/datasets/M2T1NXFLX_5.12.4/summary.

(wind direction and wind speed). Figure A.10 gives an idea of the relative size of the different grids for one of the 22 French regions.

3.5 Final datasets

Weekly worker-level dataset: In the weekly dataset, we combine the sickness leave data with the pollution and weather data based on the workplace postcode and a spatial matching between postcodes, Chimere, Copernicus and Merra grids (we assign one single Chimere and one single Merra grid to each postcode, based on the postcode's centroid). We thus allocate pollution exposure based on the location of work, rather than on the place of residence, which we do not observe. We expect the difference in exposure at workplace and at place of residence to be small because most workers live close to their workplace. We compare the distribution of pollution exposure on the workplace and on the place of residence for the population of French workers (of which our worker-level dataset is drawn) using exhaustive matched employer-employee data (called *DADS-Postes*). As shown on figure 2, the two distributions almost overlap, which indicates that our measure of workers' exposure is probably very close to their exposure at home[10].

Whether measured at the workplace or at the place of residence, any static measure of exposure introduces a measurement error in the attribution of exposure to air pollution because it neglects exposure from all other places visited during the day. As previously noted in the literature, this measurement error likely produces an attenuation bias in the estimation of pollution impacts. We discuss this issue further in section 4.

We aggregate daily observations at the weekly level to ease computational require-

[10]We also compare exposure between workplace and residence for each quintile of the wage distribution, having in mind that different income groups sort differently across space and that this may introduce a greater measurement error for some groups than others. For example, if high-income workers work predominantly in dense polluted city centres but live predominantly in green low-pollution suburban areas while low-income workers live closer to their workplace, the error in the allocation of exposure would be larger for high-income workers. Figure A.4 shows the difference in exposure by quintile of hourly wage in 2009. The distributions are again very close for all quintiles, with a slightly higher exposure at work than the place of residence for the top quintile

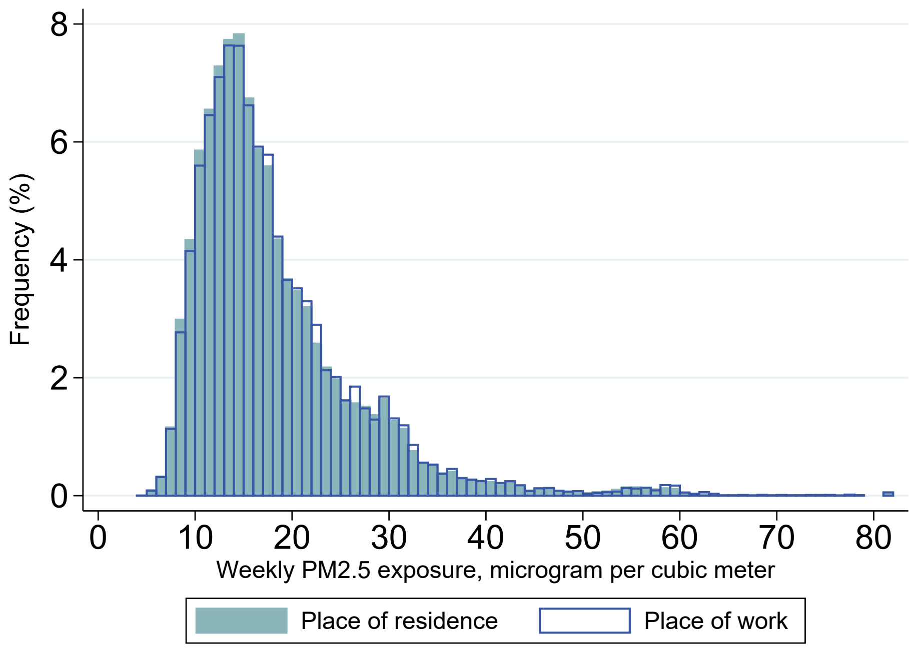

Figure 2: Distribution of weekly $PM_{2.5}$ exposure on the workplace vs the place of residence (µg/m^3), all French workers, 2009

Notes: the sample includes all private sector workers from the exhaustive matched employer-employee data of 2009, who work in a postcode where we observe at least one individual from the sickness leave dataset (N≈13,000,000). For reasons of statistical confidentiality, the few postcodes with fewer than 11 workers were dropped (<0.4% of workers).

ments.[11]. The final sample is a weekly unbalanced panel of around 450,000 individuals working in close to 430,000 private sector establishments over the period considered. The annual sample size varies from around 367,000 in 2009 to 317,000 in 2015, as the panel ages and more and more individuals reach retirement age. Table 2 shows summary statistics for the weekly sample at the individual level. Employees in our sample are 55 percent male, 40 years old on average, earn an average annual gross wage of 25,865 euros and 75 percent are full time employed.

We use the exhaustive matched employer-employee data to compare the characteristics of our sample of workers to the characteristics of the whole population of private sector employees. If we apply the same age restrictions to the exhaustive matched employer-employee data[12], we find that those workers representing the population from which our sample is drawn are 55% male, 41 on average, and earn an average annual gross wage of 26,204 euros. Thus, the average individual in our final worker sample is very close to the average private sector employee.

At the postcode level, the average $PM_{2.5}$ exposure over the study period is 15.3 µg/m^3.[13]. Figure 3 shows the distribution of weekly exposure. The recommended threshold of 15µg/m^3 for *daily* exposure set by the WHO is often exceeded, as that level is already exceeded at the *weekly* level for 39% of the worker-weeks. The threshold is exceeded at least once in every single postcode.

21 percent of employees take at least one sickness leave episode within a year[14]. The average sickness leave episode lasts 29 days and costs 808 euros in public sickness leave

[11] Access to the data is obtained through the CASD (Secure Data Access Center), which provides the service of making confidential data sets available to researchers using a secured server, for which there are constraints on the size devoted to each project, thereby computational constraints.

[12] keeping only private sector employees, only workers born between 1935 and 1989 and removing those older than 71, who should be retired. Note that in the matched employer-employee data, a worker having two different employers appears twice. We aggregate wage information at the worker level, summing up the wages she receives from different employers.

[13] In 2015, it is 12.5 µg/m^3 in 2015, the same as the population-weighted annual exposure reported by the WHO for that year (see Table 1)

[14] A national survey on Working Conditions estimated that 28 percent of private sector employees in France took at least one sick leave during 2013. Source: https://www.fonction-publique.gouv.fr/files/files/statistiques/rapports_annuels/2015/RA2015_dossier_1.pdf

benefits, whereas the median duration is only 9 days and the median spending 183 euros. On average, the number of workers starting a sickness leave episode on a given week is 5.9 per 1,000 workers. The number of sick days associated with weekly entry in sickness leave is 168 per 1,000 workers, and the associated spending is 4,742 euros per 1,000 workers.

Table 2: Summary statistics at the worker (top), sickness leave episode (middle), and postcode (bottom) level, 2009-2015

	Mean	Sd	Median	Count
Share men	0.55	0.50		2,408,715
Age	40.44	10.98	40	2,408,715
Annual gross wage (€)	25,865	26,155	22,467	2,408,715
Share full-time employed	0.75	0.43		2,408,715
Share with at least one SLE in a year	0.21	0.40		2,408,715
Nb. SLE per year.worker	0.31	0.73		2,408,715
SLE duration (days)	29	69	9	753,522
SLE benefits (publicly-funded component) (€)	808	2,291	183	753,522
Weekly exposure to $PM_{2.5}$ at workplace, average 2009-2015 (µg/m^3)	15.3	8.4	13.0	1,875,939
Nb. workers starting sickness leave episode per week (per 1,000 workers)	5.9	9.7	4.1	1,875,939
Nb. sickness days starting per week (per 1,000 workers)	168	755	53	1,875,939
Sickness leave benefits starting per week (per 1,000 workers)	4,742	23,549	1,085	1,875,939

Notes: We add median values for the continuous variables only. For statistics at the postcode level (bottom), observations are weighted with the number of individuals working in that postcode.

Monthly establishment-level dataset In the monthly dataset, we combine establishment-level information on the monthly inflow and outflow of workers, broken down by category, with weather and pollution information averaged at the monthly level. The two main outcomes are the total monthly inflow of workers joining an establishment and the total monthly outflow of workers leaving the establishment. We also examine specific categories of inflows and outflows which we expect to be affected by a pollution shock.

If establishments experience a decrease in output and profits on a polluted month, we expect them to adjust by slowing down new hires, especially those on a permanent contract, and potentially try to increase the outflow of workers. Although firing workers is costly in France, firms can decrease payroll by i)not renewing an expiring temporary contract ii)letting

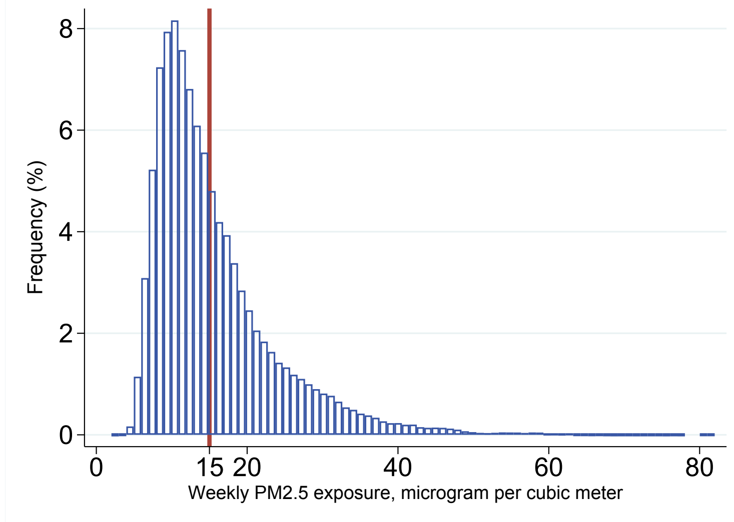

Figure 3: Workers' weekly exposure to $PM_{2.5}$ (µg/m^3), 2009-2015

Notes: The red vertical line shows the 24-hour WHO recommended threshold at 15µg/m^3.

go workers after a probation period.

If, on the other hand, establishments try to compensate worker absenteeism, one margin of adjustment would be to hire more workers on temporary contracts. Another margin could be to decrease separation from existing workers. We might then observe an increase in the inflow of workers on temporary contracts and/or a decrease in the outflow of workers. Overall, the effect of a pollution shock on labour flows is ambiguous. It probably affects as well as is affected by the elasticity of output/sales to pollution, which we will examine with the VAT data.

The final sample is a monthly unbalanced panel of around 55,000 establishments over 2009-2014. Table 3 shows summary statistics for the establishment-level sample. The average establishment has 166 employees, with monthly inflows and outflows of 7-8 workers. 70% of the establishments are active in the service sector, 25% in the manufacturing sector, 7% in the construction sector and 0.03% in the agriculture sector. Monthly exposure to pollution of their workers is close to the exposure observed for the worker-level dataset. The higher median is probably explained by the fact that the 50+ workers establishments are located in denser and more polluted areas than smaller establishments.

Table 3: Summary statistics, establishments with more than 50 workers employing at least one individual from the Hygie dataset, 2009-2014

		Mean	Sd	Median	Count
Stock of workers		166	299	97	2,509,197
Sector (%):	**Agriculture**	**0.03**			2,509,197
	Construction	**7.0**			
	Manufacturing, of which:	**25.0**			
	Manufacture of food products, beverage and tobacco	*3.6*			
	Manufacture of coke and refined petroleum products	*0.08*			
	Manufacture of computer, electronic and optical products, machinery and equipment	*4.0*			
	Manufacture of transport equipment	*1.5*			
	Manufacture of other products	*13.3*			
	Mining, quarrying, energy and water supply , waste management	*2.2*			
	Service, of which:	**70.0**			
	Wholesale and retail trade; repair of motor vehicles and motorcycles	*17.1*			
	Transportation and storage	*7.3*			
	Accommodation and food service activities	*2.5*			
	Information and Communication	*4.4*			
	Financial and insurance activities	*4.1*			
	Real estate activities	*1.1*			
	Professional, scientific and technical activities; administrative and support service activities	*12.6*			
	Public administration and defence; compulsory social security; education, health and social work	*16.5*			
	Other service activities	*2.5*			
Monthly inflow of workers, of which:		**7.7**	34.5	2	2,509,197
	hires temporary contract (%)	*58.4*			
	hires permanent contract (%)	*35.0*			
	transfers (%)	*6.6*			
Monthly outflow of workers, of which:		**7.2**	32.7	2	2,509,197
	temporary contracts coming to an end (%)	*43.2*			
	resignations (%)	*17.5*			
	lay-offs (%)	*16.0*			
	pre-retirements or retirements (%)	*7.3*			
	transfers (%)	*7.2*			
	others (%)	*4.0*			
	Mutual agreement layoff (%)	*4.8*			
Quarterly stock of workers from temp agencies		**5.4**	22.5	0	2,598,879
Monthly exposure to $PM_{2.5}$, 2009-2014 (µg/m^3)		15.6	6.4	13.9	2,598,879

4 Empirical strategy

4.1 Air Pollution Impacts on Sickness Leave Episodes

Our first objective is to estimate the impact of short-run exposure to fine particulate matter on sickness leave outcomes, net of any potentially confounding factors. We aggregate the data at the grid cell level to ease computational requirements. The grid cell level is also the scale at which we observe pollution. We focus on three sickness leave outcomes at the grid cell level: the number of workers starting a SLE, per 1,000 workers ; the associated number of sick days per 1,000 workers, and the associated sickness leave benefit spending.

We model this relationship using the following equation:

$$Y_{g,t} = \alpha + \beta PM_{2.5g,t} + W_{g,t}\gamma + h_{d,t}\delta + \nu_g + \theta_{q,y} + \epsilon_{g,t}, \tag{1}$$

where the dependent variable $Y_{g,t}$ is the sickness leave outcome measured at time t and in grid cell/postcode g . the parameter of interest is β, the coefficient on weekly $PM_{2.5g,t}$ concentration in grid cell g.

The high granularity of our data allows us to include multiple sets of high-dimensional fixed effects. We generate indicators for weekly average of the daily maximum temperatures falling into 1 of 8 bins[15]. We do the same for wind speed and rainfall, for which we compute indicators for each quintile of these variables. We then generate a set of indicators for all possible interactions of these weather controls and include it in all our regressions as $W_{g,t}$. Our estimates also include a vector of *departement*-level time-varying variables $h_{d,t}$, including school holiday dummies and the estimated number of flu cases per 100,000 individuals in the *departement* d where grid cell g falls into (where the *departement* is an administrative and jurisdictional unit).[16] The grid cell fixed effects ν_g control for cross-sectional geographic

[15]The bins span 5°C each, except for the first bin including all negative temperatures, and for the eighth bin including all temperatures above 30 °C

[16]In metropolitan France, the median size of a *departement* is 5 880 km^2, which is equivalent to 3.5 times the size of a median US country.

differences in working population and pollution at a very fine level (approximately 4 km by 4 km). Finally, quarter-by-year fixed effects $\theta_{q,y}$ control flexibly for common time-varying shocks, such as those induced by regulation. We cluster all standard errors $\epsilon_{g,t}$ at the grid cell level and weight all estimates by the number of workers present in the grid cell a given year.

Since our identification relies on within-grid cell variations, we need to have enough variation in pollution over time within a grid cell. Figure A.9 shows the week-on-week variations in $PM_{2.5}$ levels for four grid cells, two in urban areas and two in rural areas. In cities, the average week-on-week variation is around 6 µg/m^3, which represents 80% of the SD for the whole sample. Pollution levels are lower on average in rural areas but week-on-week variations are still quite high, around 4 µg/m^3 on average, which represents 50% of an SD.

OLS estimates of equation (1) are prone to bias because exposure to $PM_{2.5}$ is likely to be measured with error and not to be randomly assigned. Indeed, as mentioned above, pollution exposure based on the workplace location is a static measure that potentially suffers from a classical measurement error which gives rise to an attenuation bias. Also, individuals may take into account air pollution concentrations when choosing where to live, and this sorting behavior may induce a bias in the OLS estimates upward or downward, depending on how individuals sort.[17] There may be an upward bias if, for example, poorer individuals with a usually worse health status live (and work) in relatively more polluted areas, where housing prices are lower. There may also be a downward bias if individuals who are more vulnerable to pollution sort themselves into places that are less affected by pollution. Additionally, there may be other omitted variables influencing both the formation of pollutants and health. For example, ozone formation is more frequent under high temperatures, and high temperatures can lead to heat waves having detrimental effects on health. Finally, there

[17] This has been well-documented in the literature on residential sorting (see, eg. Banzhaf et al. 2019; Lee and Lin 2018). To the extent that pollution levels at the workplace and at the place of residence are correlated, residential sorting implies that pollution levels at the workplace is also not randomly assigned.

may be a simultaneity bias if high absenteeism levels lead to a lower economic activity and lower commuting, thereby decreasing local pollution levels.

We can partly address these concerns with the use of appropriate fixed effects and control variables. In particular, controlling for weekly variations in temperature, precipitation or wind speed at the grid cell level absorbs the joint effect of weather conditions on pollution and health. Similarly, local school holidays may influence pollution concentrations and the propensity to start a SLE. Controlling for quarter-by-year and *departement*-week fixed effects is important to capture time-specific shocks influencing both pollution and absenteeism - an example of such shock is the Covid-19 epidemic and associated lockdowns or the flu epidemic. In our main specification, we control for grid cell fixed effects to capture time-invariant area-specific characteristics, such as the demographic and socio-economic composition of an area (gender, age profile, socio-economic status); but, in a robustness test, we also control for grid-year fixed effects to allow for this composition to vary over the years.[18]

To address remaining potential biases, we rely on an instrumental variable approach exploiting week-to-week variations in wind direction at the grid cell level, in the spirit of Deryugina et al. (2019); Anderson (2020). The validity of our IV estimation rests on two conditions. First, pollution needs to be sufficiently correlated with changes in wind direction. Intuitively, pollution travels in the air and pollution particles are very light, such that wind direction is likely to influence pollution concentrations. Suppose that for each grid cell, there is one main source of pollution (e.g., a power plant or a manufacturing firm), hence only the direction of the wind that comes from that source will increase pollution concentration. Second, the exclusion restriction requires that wind directions should only affect the propensity to start a SLE via its effect on air pollution. The specification of our first stage is:

[18]Ideally, we would take advantage of the individual panel dimension and use individual-fixed effects. However, the resulting dataset is very large and all the regressions need to be run from a secure server with a limited computational capability. We have been unable to run regressions with individual fixed effects thus far.

$$PM_{2.5g,m,t} = \alpha + \sum_{j=2}^{4} \sum_{k=1}^{K} \beta_{jk} \text{WIND}_{j,k,t} \mathbb{1}(k = m) + W_{g,t}\gamma + h_{g,d,t}\delta + \nu_g + \theta_{q,y} + \epsilon_{g,m,t} \quad (2)$$

where $PM_{2.5g,m,t}$ is the average $\text{PM}_{2.5}$ concentration in grid point g located in Merra grid m at time t. The excluded instruments are $\text{WIND}_{j,k,t}\mathbb{1}(k = m)$, where each variable in the set $\text{WIND}_{j,k,t}$ corresponds to the number of hours in week t where the wind comes from direction j, with $j = 1$ being wind blowing from the South (omitted category), $j = 2$ from the West, $j = 3$ from the North, and $j = 4$ from the East.[19] The variable $\mathbb{1}(k = m)$ is an indicator for Chimere grid cell g to be included in Merra grid cell m, with K being the total number of Merra grids. The coefficient on their interactions, β_{jk}, is the parameter of interest for the first stage and is allowed to vary across Merra grid cells. The other control variables are defined as in equation (1).

For the identification of the β_{jk}, we rely on the week-to-week variation in wind direction within a Merra grid cell and a quarter-year. Figure 4 illustrates the magnitude of these short-term variations. For example, the upper left panel indicates that in Merra grid cell #1, wind blew disproportionately more from the South and the West, and less from the East, in January 2009 (red bar) compared to the average first quarter of 2009; by contrast, the lower left panel indicates that, for the same grid in a different month (June), wind blew disproportionately more from the East and less from the North and the West in 2015 compared to the average of the quarter (light blue), while variations were more limited in 2009 (red) and 2012 (dark blue). The patterns differ across Merra grid cells within a given month. This Figure suggests there is a large amount of variation in wind direction within a given Merra grid cell within a quarter.

Weak instrument bias is not a concern in our setting. We formally show that local variations in wind direction are a strong predictor for local $\text{PM}_{2.5}$ concentrations. Figure 5

[19] We decompose the 360 degrees of a compass into 90-degree quadrants corresponding to a main direction.

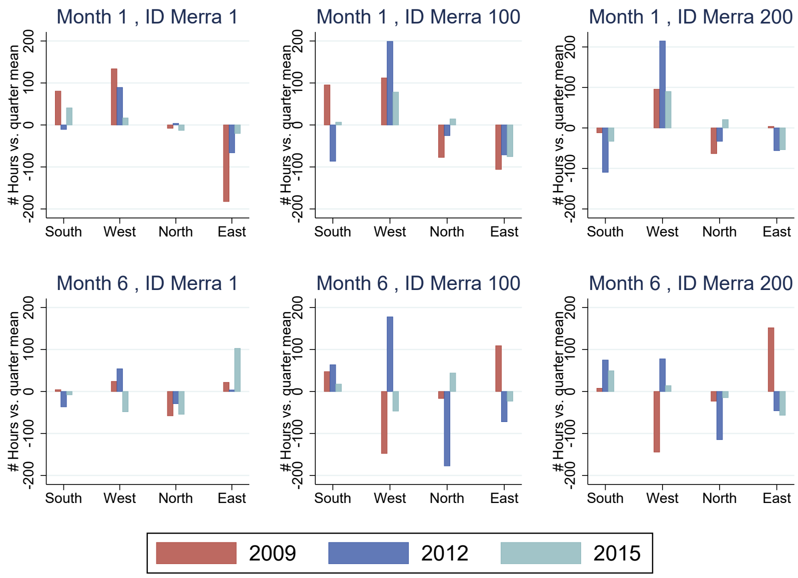

Figure 4: Variation in wind direction

Notes: Figure shows the number of hours in a month in which the wind blows from a given direction, demeaned by the average for the quarter that month is part of, for three different Merra grid cells (1, 100, and 200), two different months (January or Month 1 in the upper panel and June or Month 6 in the lower panel), and three different years (2009, 2012, 2015).

plots the magnitude of the 711 $\widehat{\beta_{jk}}$ coefficients (3 wind directions x 237 Merra grids) obtained with an OLS regression on equation (4), against their t-stat. Only few point estimates are not statistically significant at the 95% level, and most coefficients have a large t-stat. The F-statistic of the first stage is 4,077, way above the recommended threshold value of 104.7 for weak instrument detection (Lee et al., 2021).

Figure 6 shows a map of the estimated coefficients. Compared to wind from the South (the omitted category), winds blowing from the West (coming from the Atlantic ocean) significantly decrease pollution levels in all grid cells, which is quite intuitive since the ocean does not contain many sources of air pollution. By contrast, winds blowing from the East increase pollution in the North-West, and decrease pollution in the South, relatively to Southern winds. Note that wind blowing from the South may bring dust from the Sahara, thereby increasing the $PM_{2.5}$ concentrations.[20]

[20]France does not report these dust events to the EU even though there is a EU directive that allows countries to do so to disentangle external pollution from nationally produced pollution.

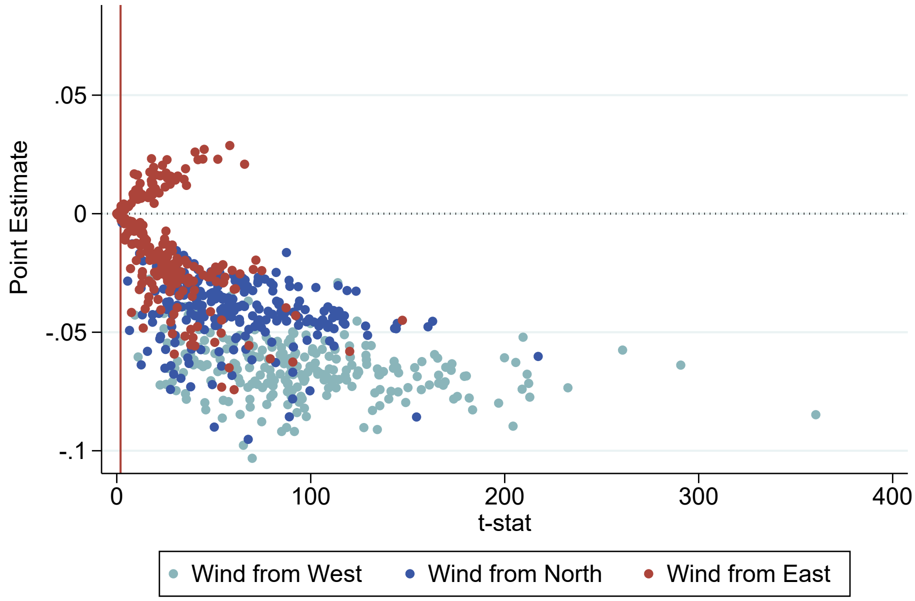

Figure 5: Point estimates and t-stat for each $\widehat{\beta_{jk}}$

Notes: The estimated coefficients $\widehat{\beta_{jk}}$ express the average increase in weekly $PM_{2.5}$ in a given Merra grid when wind blows one additional hour from direction j compared to blowing from the South. The red vertical line is set at x=1.96 and the black horizontal line at y=0.

In our context, the biggest threat to the exclusion restriction is that other pollutants that also affect health co-vary with wind direction. Of the four other regulated air pollutants (SO_2, NO_2, PM_{10} and ozone) that may impact sickness leave, SO_2 and NO_2 are primary pollutants that convert to particulate matter within two or three days; thus we cannot estimate their effect on health independently given our aggregation of the pollution data at the weekly level. PM_{10} is highly correlated with $PM_{2.5}$ (Pearson correlation coefficient: ρ=0.93) and actually includes $PM_{2.5}$. Finally, ozone is a pollutant that is typically anti-correlated with other pollutants due to how it is formed in the atmosphere: ozone results from the chemical reaction between solar radiation, nitrogen oxide and volatile organic compound (Nasa Earth Observatory, 2003). In our data, the correlation coefficient between weekly $PM_{2.5}$ and ozone is ρ=-0.3. Figures A.6-A.8 illustrate this anti-correlation by showing the seasonality of ozone, $PM_{2.5}$ and NO_2 concentrations: while $PM_{2.5}$ and NO_2 concentrations peak in winter, ozone peaks in spring and summer. In section 6.1, we check whether our results change when we

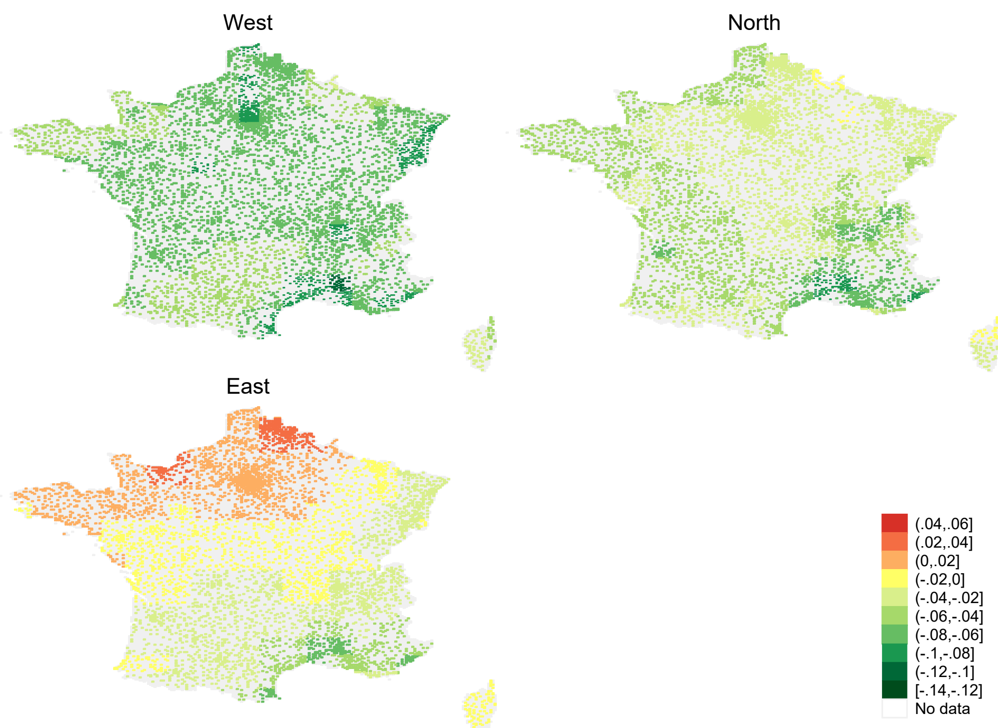

Figure 6: Map of the $\widehat{\beta_{jk}}$

Notes: The estimated coefficients $\widehat{\beta_{jk}}$ express the average increase in weekly $PM_{2.5}$ in a given Merra grid when wind blows one additional hour from direction j compared to blowing from the South.

add ozone as an additional endogenous regressor in the two-stage least square estimation. In Appendix A.1, we test whether our results are robust to specicying a Poisson model, assuming non-linear effects of pollution on health, instead of a linear model.

4.2 Plant-Level Response to Air Pollution Shocks

Our strategy to examine plant-level responses to pollution differs from the weekly analysis described above in two ways:: first, we now rely on month-to-month variation in wind direction and pollution, instead on week-to-week variations. Second, we allow pollution to affect labour flows with a lag. This seems realistic given that a hiring or separation process takes some time. In practice, we examine the relationship between pollution on month t and the average of labour flows over month t, $t+1$ and $t+2$. We also aggregate the data at the grid cell level g given computational constraints.

We model the relationship between monthly pollution $PM_{2.5g,t}$ on month t and grid cell g and workers' flows $Y_{g,t}$ in establishments located in that grid as follows:

$$Y_{g,t} = \alpha + \beta PM_{2.5g,t} + W_{g,t}\gamma + h_{d,t}\delta + \nu_g + \theta_{q,y} + \epsilon_{g,t}, \tag{3}$$

$Y_{g,t}$ represents the average of the labour flow outcome over month t, $t+1$ and $t+2$. The labour flow outcome observed at month t, for all t, is the flow measured in number of workers divided by the sock of workers in the establishment at the beginning of the month. In terms of flow types, we examine total inflows, total outflows, as well as specific categories of inflows and outflows which we expect may be affected by a pollution shock.

The vector $W_{g,t}$ is now a set of weather indicators combining all the possible interactions between bins of monthly average of daily maximum temperature and quintiles of monthly wind speed and monthly precipitation. The vector of *departement*-level time-varying variables $h_{d,t}$ now includes the share of days in the month where there are school holidays in that departement and the share of incidence of flu in the *departement* d where grid cell g falls into.

Like in the specification at the weekly level, the grid cell fixed effects ν_g control for cross-sectional geographic differences in working population and pollution, and quarter-by-year fixed effects $\theta_{q,y}$ control flexibly for common time-varying shocks. We cluster all standard errors $\epsilon_{g,t}$ at the grid cell level and weight all estimates by the number of establishments in each grid cell-month.

To avoid that the effect of pollution on labour flow in the following two months captures changes in weather, holiday and flu conditions in these two months, we also control for two leads of the instrument and two leads of the holiday and flu epidemic variables.

The specification of the first stage is now:

$$PM_{2.5g,m,t} = \alpha + \sum_{j=2}^{4} \sum_{k=1}^{K} \beta_{jk} \text{WIND}_{j,k,t} \mathbb{1}(k = m) + W_{g,t}\gamma + h_{d,t}\delta + \nu_g + \theta_{q,y} + \epsilon_{g,m,t} \quad (4)$$

where $PM_{2.5g,m,t}$ is the average $\text{PM}_{2.5}$ concentration in grid point g located in Merra grid m in month t. The excluded instruments are still $\text{WIND}_{j,k,t}\mathbb{1}(k = m)$, but now each variable in the set $\text{WIND}_{j,k,t}$ corresponds to the number of hours in a month t where the wind comes from direction j rather than a week. The variable $\mathbb{1}(k = m)$ is an indicator for Chimere grid cell g to be included in Merra grid cell m, with K being the total number of Merra grids. The β_{jk} is the parameter of interest for the first stage at the monthly level and is allowed to vary across Merra grid cells. The other control variables are defined as in equation (3).

In spite of aggregating at the monthly level, we still have a strong first stage. The F-statistic is 1,036, and 84% of the $\widehat{\beta_{jk}}$ coefficients are statistically significant at the 5% level. The estimated coefficients are also consistent across the two estimations: for a given Merra grid cell, they have the same sign in the monthly as in the weekly specification in 99% of the cases for the wind direction West, 59% of the cases for the wind direction North, and 64% of the cases for the wind direction South. The monthly-level first stage does not include 2015 while the weekly-level dataset does, so the two are not completely comparable.

5 Results

5.1 Pollution and Sickness Leave

Main results Table 4 reports the OLS and 2SLS estimates of the relationship between $PM_{2.5}$ concentration and the three sickness leave outcomes. While the OLS coefficients reported in odd-numbered columns are relatively small in magnitudes and not always statistically different from zero, the IV coefficients reported in even-numbered columns are positive and quite precisely estimated. From column (2), an increase by 1 µg/m^3 of weekly $PM_{2.5}$ concentrations is predicted to increase the number of workers starting a sickness leave that week by 0.03 per 1,000 workers, starting from a baseline average of 5.9 per 1,000 workers. This corresponds to a 0.55% increase. Put differently, a one standard deviation increase in $PM_{2.5}$ increases the share of workers starting a SLE that week by 4.6%. The results by Holub et al. (2021) – the closest paper to ours – indicate that a one standard deviation increase in PM_{10} increases the weekly absence rate by 1.4%. Since $PM_{2.5}$ is known to be more damaging for health than PM_{10}, it would be logical to obtain a larger magnitude in our case. However, we cannot directly compare the two results because Holub et al. (2021) use a different IV strategy and measure the share of sickness leaves by including all sickness leaves (even ones that started earlier) whereas we only consider the share starting sickness leave on a given week.

Columns (3) and (4) report the OLS and 2SLS estimates of the relationship between $PM_{2.5}$ concentration and the average number of sickness leave days per worker. An increase by 1 µg/m^3 in weekly PM2.5 pollution is predicted to increases the number of sickness leave days by 0.943 days per 1,000 workers. This corresponds to a 0.56% increase compared to the variable mean.

Finally, columns (5) and (6) report the OLS and 2SLS estimates of the relationship between $PM_{2.5}$ concentration and the average amount of public sickness leave benefit per worker. We can expect the result on sickness leave benefit to differ from the result on sickness

Table 4: Pollution and sickness leave outcomes - OLS and IV

	Workers starting sickness leave		Number of sickness days		Sickness leave benefit	
	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	IV	OLS	IV	OLS	IV
$PM_{2.5}$	0.00388**	0.0321***	0.194	0.943***	7.37*	28.3***
	(0.00137)	(0.00299)	(0.100)	(0.24)	(3.30)	(7.97)
Weather controls	Yes	Yes	Yes	Yes	Yes	Yes
Holiday and flu controls	Yes	Yes	Yes	Yes	Yes	Yes
Quarter-year FE	Yes	Yes	Yes	Yes	Yes	Yes
Chimere grid cell FE	Yes	Yes	Yes	Yes	Yes	Yes
Dependant variable mean	5.9	5.9	168	168	4,742	4,742
N	1,869,578	1,869,578	1,869,578	1,869,578	1,869,578	1,869,578
1st stage F-statistic		6,538		6,538		6,538

Notes: The unit of observation is the postcode x week. All estimates are per 1,000 workers. Estimates are weighted by the number of workers in each postcode. Robust standard errors clustered at the chimere grid cell level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

leave days due to two factors: first, the first three days of leave are not compensated, so we could see a lower increase in sickness benefits than in sickness days if the pollution-induced leaves are predominantly short. Second, the amount of sickness benefit is linked to the worker's wage, as explained in section 2. If pollution-induced sickness leave spells affect more high-wage workers than the average sickness leave spell, we could see a higher increase in sickness benefits compared to the increase in sickness leave days. We find that an increase by 1 µg/m^3 in weekly PM2.5 pollution increases the average sickness leave benefit spending per 1,000 worker by 28.3€. This corresponds to a 0.6% increase compared to the variable mean, slightly higher than the percent change in terms of sickness leave days.

Benefits of reduced pollution A back-of-the-envelope calculation illustrates the policy relevance of these results. In particular, we can estimate the benefits of meeting the WHO targets in terms of avoided sickness days and avoided sickness leave benefit spending. Over our 7-year study period, the 15µg/m^3 threshold is exceeded for 39% of the worker-weeks (the part of the distribution to the right of the red line on Figure 3), corresponding to close to 52 millions worker-weeks. We compute, for each grid cell-week, (i) the decrease in $PM_{2.5}$ that is required to meet the 15 µg/m^3 threshold; (ii) the associated number of avoided sickness days; (iii) the associated avoided sickness leave benefit spending.

Regarding step (i), bringing each grid cell-week exceeding the 15µg/m^3 threshold to a PM2.5 level of 15µg/m^3 implies decreasing pollution by 7.6 on average for those observations. For the entire sample, such a decrease would imply a 18% decrease in annual pollution concentrations, compared to the levels observed over 2009-2015.

To calculate (ii) and (iii) for a given grid cell-week, we simply multiply the estimated coefficient from columns (4) and (6) of table 4, respectively, by the required decrease. This is conservative because estimates from table 4 is the average marginal effect estimated from the whole distribution of PM2.5 levels ; actually, the marginal effect of pollution on the number of sickness days is likely to be non-linear and increasing with the initial level of pollution.

We estimate that 63,426 days of sickness leave would have been avoided every year if none of the grid cell-weeks had had pollution above 15 µg/m^3. Scaling it to the actual population of 18,730,000 private sector employees, this corresponds to 2.9 million sickness days per year[21].In terms of sickness leave benefit, not exceeding this threshold would have avoided 88 million euros of spending every year, at population scale. Given that total benefits related to sickness leaves amounted to €7,091 million each year over the period (DREES, 2020), respecting the WHO standard could have saved 1.2% of that cost. This is only including the short-term effects of pollution on absenteeism. Respecting the WHO standards would probably also have long-term health benefits translating in reduced absenteeism, a benefit that we are unable to capture in the present analysis.

In addition to the publicly funded SLB, employers pay the mandatory employer-funded allowance, and the optional allowance for 2/3 of the employees. With an average SLE duration of 29 days and an average daily wage at 71€ in our sample, we calculate that meeting the WHO threshold would have saved an annual 110 million euros for firms, only

[21]Our sample has on average 344,052 individuals per year, 1.84% of the total population of private sector employees, 18,730,000 workers in 2015. (See INSEE: https://www.insee.fr/fr/statistiques/2496914).

in terms of avoided sickness benefit spending.[22]

There is an additional cost to the spending related to sickness leave insurance: that in terms of foregone production due to workers' absenteeism. In the second part of the analysis, we focus on outcomes at the establishment level. Some of these results are still preliminary, so it is useful to value this foregone production at the daily wage rate, in a first step. Valued at a daily wage of 133 euros per working day[23], the cost of failing to meet the WHO standards is 391 million euros in terms of foregone production.

5.2 Pollution and workers' flows

Main results Tables 5 and 6 show the results from running a 2SLS regression on the model described in equation 3. From column (1) of table 5, a one-unit increase in average PM2.5 levels on month t decreases workers' inflow rate by 0.21 percentage points on average over the current and following two months. This corresponds to a 3% decrease relative to the average inflow rate of 6.8 percent. Put differently, a one standard deviation increase in monthly $PM_{2.5}$ decreases the inflow rate of new workers by 19%.

The decrease is driven by a small decrease in the hiring rate of workers on permanent contracts (column (3)) and by a large decrease in the rate of transfers of workers from other establishments belonging to the same firm (column (4)). Establishments faced with a pollution-induced productivity/production shock seem to cope by reducing the most costly and long-term labour inflows. For multi-establishment firms, slowing down transfers from other establishments is another way to cope.

From table 6, we fail to detect any effect of pollution on workers' outflows. The results by category of outflow suggest rather a decrease in outflows, especially those at the discretion of the employer: we see a small significant decrease in the separation rate due to a failed

[22]the avoided 2.9 million of sick days per year correspond to $2,900,000/29 = 101,484$ avoided sickness leave episodes. The average daily wage in our sample is $25,865/365 = 71e$. Thus, the benefit of avoiding one 29-day SLE for private sector employers is approximately $(29-7)*0.4*71+(2/3)[3*71+4*0.5*71+(29-7)*0.1*71] =$€1,084. The avoided spending is $1,084 \times 101,484 = 110,008,656$

[23]This measure of wage is per working day, with 260 working days per year, while the measure of hourly wage used for calculating sickness leave benefit is per calendar day

Table 5: Pollution and workers' inflows - IV

	(1) Total	(2) Temporar contract hires	(3) Permanent contract hires	(4) Transfers
$PM_{2.5}$	-0.204**	-0.0327	-0.0159*	-0.156**
	(0.0747)	(0.0499)	(0.00686)	(0.0561)
Weather controls	Yes	Yes	Yes	Yes
Holiday and flu controls	Yes	Yes	Yes	Yes
Quarter by Year FE	Yes	Yes	Yes	Yes
Chimere grid FE	Yes	Yes	Yes	Yes
Dependent variable mean	6.8	5.3	1.0	0.5
N	255,500	255,500	255,500	255,500
1st stage F-statistic	1,036	1,036	1,036	1,036

Notes: Estimates are weighted by the number of establishments in each grid cell. Robust standard errors clustered at the chimere grid level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Outcomes and point estimates have been scaled by a factor 100 to improve readability. The outcomes are standardized by the stock of workers. For example, the inflow outcome designate the share of new workers in the plant as a share of the stock of workers in the beginning of the month.

probation period or a layoff, and a decrease in transfers to other establishments.

One interpretation could be that establishments faced with a pollution shock cope with the associated absenteeism by retaining existing workers more. Analyses by sector and taking into account the vulnerability of establishments to pollution-induced absenteeism (based on their workers' characteristics) are ongoing. They will help explain the mechanisms at play.

5.3 How are the costs distributed? Heterogeneity analysis

5.3.1 Heterogeneity of the pollution-sickness leave response

By initial health status: We expect to see a stronger effect of pollution on sickness leave for individuals with an initially poorer health status. While we do not observe workers' health status, we observe their annual healthcare costs in terms of medical visits (except hospitalisations and dentists) and medical drug purchase. Importantly, we observe both the total healthcare costs and the out-of-pocket payments, from which we can infer the

Table 6: Pollution and workers' outflows - IV

	(1) Total	(2) Resign	(3) End of contract	(4) End of probation and layoff	(5) Transfers
$PM_{2.5}$	-0.0363	-0.0011	-0.0099	-0.0032**	-0.0044**
	(0.0502)	(0.00084)	(0.0024)	(0.0083)	(0.00145)
Weather controls	Yes	Yes	Yes	Yes	Yes
Holiday and flu controls	Yes	Yes	Yes	Yes	Yes
Quarter-year FE	Yes	Yes	Yes	Yes	Yes
ID Chimere FE	Yes	Yes	Yes	Yes	Yes
Dependent variable mean	6.2	0.5	4.7	0.4	0.2
N	255,500	255,500	255,500	255,500	255,500
1st stage F-statistic	1,036	1,036	1,036	1,036	1,036

Notes: Estimates are weighted by the number of establishments in each grid cell. Robust standard errors clustered at the chimere grid level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Outcomes and point estimates have been scaled by a factor 100 to improve readability. The outcomes are standardized by the stock of workers. For example, the outflow outcome designate the share of workers leaving the establishment as a share of the stock of workers in the beginning of the month.

healthcare costs covered by the public health insurance. Total healthcare costs are probably not a great proxy for health status: all else equal, healthcare costs are likely to increase with income because the price of a medical visit is not regulated for all doctors in France. On the other hand, the amount covered by the public health insurance is regulated, capped for each medical visit and drug type, and does not vary with income level (except for rare exceptions). It is thus a better reflection of actual health status. We define quintiles of healthcare utilization based on the healthcare costs covered by the public health insurance in 2009 as a proxy for initial health status. We then build five grid cell-level datasets, each only including individuals of a given quintile. We run the 2SLS estimation from equation 1, using the number of workers starting a sickness leave per 1,000 workers as an outcome. We restrict the regression period to 2010-2015, because pollution in 2009 may affect the composition of each quintile via its effects on healthcare spending that year.

Table A.4 shows the composition of each quintile in terms of healthcare costs covered by the public health insurance in 2009, and mean $PM_{2.5}$ exposure and average gross wage over the 2010-2015 period. Healthcare costs covered by the public health insurance are almost 300 times larger for Q5 than Q1 on average. On the other hand, there is little difference in pollution exposure and annual gross wage across the different quintiles.

Table 8 reports the 2SLS estimates from equation 1 for the whole sample over 2010-2015, and by quintile of covered healthcare costs. Sickness leave incidence increases with our proxy of initial health status, which makes sense. The marginal effect of pollution on sickness leave is low and not statistically significant for the 20% of workers with the lowest healthcare utilization in 2009, while it is 50% larger than the average effect for the quintiles quintiles 3, 4 and 5. The effect for quintiles 3 and 4 is significantly different from that of quintile 1. The fact that pollution increases sickness leave incidence even for workers with a median health status implies that pollution effects go well beyond the vulnerable individuals. Taken as a percent change relative to the mean, the effect is strongest for quintile 3.

Table 7: Number of workers starting sickness leave per 1,000 workers - heterogeneity analysis by quintile of covered healthcare costs - IV regression

	(1)	(2)	(3)	(4)	(5)	(6)
	All	Q1	Q2	Q3	Q4	Q5
$PM_{2.5}$	0.0253***	0.0078	0.014*	0.037***	0.041***	0.037***
	(0.00329)	(0.00599)	(0.00651)	(0.00798)	(0.00865)	(0.0104)
Dependent variable mean	5.9	3.3	4.7	6.1	7.6	9.6
Effect relative to mean	0.43%	0.24%	0.30%	0.61%	0.54%	0.39%
N	1,601,310	1,342,109	1,326,507	1,323,835	1,315,133	1,281,716

Notes: Robust standard errors clustered at the chimere grid level. * $p < 0.05$, ** $p < 0.01$ ***$p < 0.001$. Estimates are weighted by the number of workers in each grid cell. Quintiles are based on the distribution of social security reimbursements for medical appointments and drug purchase in 2009. The regression analysis is restricted to the 2010-2015 period. All regressions include weather, holiday and flu controls, Chimere grid cell fixed effects and quarter-by-year fixed effects. Quintile categories take into account the number of individuals by grid cell, hence different quintiles have a different number of grid cells and regression by quintiles have different N.

By wage level: We may expect the pollution-sickness leave relationship to vary by wage levels via several mechanisms. First, lower-income individuals have been found to incur more damages from air pollution in a variety of contexts (Hsiang et al., 2019). These higher damages can be due to two factors: a higher exposure combined with increasing marginal damages from pollution, and/or a higher vulnerability for a given exposure. In the case of sickness leaves, another mechanism may explain heterogeneous effects along the wage dimension: facing the same health shock, high-wage and low-wage individuals may have a different propensity to actually take a sickness leave. Low-wage individuals in precarious contracts may be less willing to take a sickness leave if it endangers their job security. On the other hand, high-wage individuals have a lower replacement rate given the cap on publicly funded sickness leave benefits, so they may be less willing to enter sickness leave. Overall, the effect of wage on the marginal effect of pollution on sickness leave depends on how high-wage and low-wage individuals differ in terms of i)their exposure to pollution ii)their vulnerability to pollution iii)their propensity to take a sickness leave for a given health shock.

We test for heterogeneous effects along the wage dimension based on wage quintiles from 2009 and restricting the analysis to 2010-2015. We only consider workers who are full-time employed and have an annual wage at least equal to the minimum wage in 2009, to avoid misclassifying a high-wage part-time worker in a low-wage quintile. Like above, we run the 2SLS estimation from equation 1 on five grid cell-level datasets, each including individuals from a given wage quintile.

Table A.5 shows the composition of each quintile in terms of annual gross wage in 2009, and $PM_{2.5}$ exposure and covered healthcare costs (the same variable as the one used in the heterogeneity analysis above) averaged over the 2010-2015 period by wage quintile. Workers in the top wage quintile earn close to three times as much as the ones in the bottom quintile on average. They are also the most exposed to pollution. While this may partly reflect a greater discrepancy in our measure of exposure (as mentioned in section 3.5), this is also consistent with the top quintile's higher exposure based on the place of residence reported in table A.3 and figure A.4e. In contrast, covered healthcare costs are fairly similar across wage quintiles, suggesting that differences in initial health status may be quite low.

Table 8 reports the 2SLS estimates from equation 1 for the five different wage quintiles.

Note that individuals in the bottom quintile of the wage distribution are twice as likely to start a sickness leave on a given week than individuals in the top quintile. Given that both quintiles have similar healthcare costs from table A.5, indicative of a similar initial health status, this may reflect the fact that high-wage individuals are less willing to enter sickness leave. The marginal effect of pollution on sickness leave follows an inverted U-shaped curve with respect to the wage quintile: it is lowest for the bottom and top quintiles and highest for the middle quintile. However, the magnitudes are not statistically significant from each other. The interplay of the mechanisms mentioned above may explain the shape of the curve: the higher exposure to pollution of the top wage quintile may be compensated by a lower willingness to enter sickness leave or/and a lower vulnerability to pollution.

Table 8: Share of individuals starting sickness leave - heterogeneity analysis by wage quintile - IV regression

	(1)	(2)	(3)	(4)	(5)
	Q1	Q2	Q3	Q4	Q5
$PM_{2.5}$	0.024*	0.044***	0.044***	0.036***	0.019*
	(0.0120)	(0.0113)	(0.0102)	(0.0092)	(0.00754)
Dep. var. mean	7.9	7.4	6.8	5.9	3.8
Effect vs. mean	0.30%	0.59%	0.65%	0.61%	0.50%
N	1,239,872	1,241,642	1,168,443	1,048,142	893,582

Notes: Robust standard errors clustered at the chimere grid level. * $p < 0.05$, ** $p < 0.01$ ***$p < 0.001$. Estimates are weighted by the number of workers in each grid cell. quintiles are based on the distribution of wages in 2009. The regression analysis is restricted to the 2010-2015 period. Compared to the full sample, wage quintile categories only include individuals with a full-time job and earning at least the equivalent of a full time minimum wage in 2009. All regressions include weather, holiday and flu controls, Chimere grid cell fixed effects and quarter-by-year fixed effects. Quintile categories take into account the number of individuals by grid cell, hence different quintiles have a different number of grid cells and regression by quintiles have different N.

5.3.2 Heterogeneity of the pollution-labour flow response

Forthcoming.

6 Robustness checks

6.1 Worker-level analysis

Varying specifications: We test the robustness of our main results to a different set of weather controls and fixed effects. Columns (1)-(3) of table 9 report the results. Column (1) reports the main 2SLS estimate on the relationship between PM2.5 and the share of individuals starting a sickness leave. In Column (2), weather controls take the form of three continuous variables – one for the average daily maximum temperature, one for wind speed and one for total weekly rainfall – rather than a set of interactions, as in the main specification. In column (3), we add Chimere grid cell by year fixed effects to capture changes in the composition of Chimere grid cells over time (these changes arise both because our individual panel is not balanced and because individuals may change workplace over time). The estimates are not too far from the main coefficient estimate reported in column (1). Appendix A.2 also shows our the coefficient estimate varies when we add controls progressively.

Adding other pollutants: One concern with interpreting our estimates as the causal effects of $PM_{2.5}$ is that other pollutants, ozone in particular, may also be influenced by wind direction but are omitted in our main specification. We exploit the fact that these pollutants are not perfectly co-transported and that they can be produced by sources located in different places and carried differently by the wind. Hence, our empirical strategy allows us to instrument separately for several pollutants. As described above, two pollutants, SO_2 and NO_2 are precursors to $PM_{2.5}$ and convert to particulate matter within two to three days, whereas PM_{10} includes $PM_{2.5}$. Given our specification of weekly concentration in $PM_{2.5}$, we are not able to distinguish the independent effects of these pollutants from those of $PM_{2.5}$. The other pollutant left, ozone (O_3), is anti-correlated with $PM_{2.5}$ and may have independent effects that we can capture. In column (4) of table 9, we consider both the effects of $PM_{2.5}$

Table 9: Varying specifications - share starting sickness leave - 2SLS

	(1) Main	(2) Continuous weather	(3) ID Chimere by year FE	(4) With ozone	(5) Ozone only
$PM_{2.5}$	0.0321***	0.0489***	0.0286***	0.0370***	
	(0.00299)	(0.00331)	(0.00289)	(0.00316)	
O_3				0.0179***	-0.00113
				(0.00514)	(0.00492)
Weather	Yes	Continuous	Yes	Yes	Yes
Holiday and flu control	Yes	Yes	Yes	Yes	Yes
Quarter by Year FE	Yes	Yes	No	Yes	Yes
Chimere grid FE	Yes	Yes	No	Yes	Yes
Chimere grid x Year FE	No	No	Yes	No	No
Quarter FE	No	No	Yes	No	No
Dependent variable mean	5.9	5.9	5.9	5.9	5.9
N	1,869,578	1,869,578	1,869,578	1,869,578	1,869,578

Notes: The unit of observation is the postcode x week. All estimates are per 1,000 workers. Estimates are weighted by the number of workers in each postcode. Robust standard errors clustered at the chimere grid cell level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

and ozone in the 2SLS regression by adding ozone as an endogenous regressor. The point estimate on $PM_{2.5}$ is close to the main estimate. This suggests that our main result is not driven by the confounding role of ozone. In column (5), we run the analysis considering only ozone as an endogenous regressor. We fail to detect an impact on sickness leave incidence, which may be due to the fact that wind direction changes are not a valid instrument for ozone, or to the fact that ozone does not affect sickness leave incidence (high ozone levels are mostly found in summer, where many French workers take holidays and are thus less likely to take sickness leave).

6.2 Establishment-level analysis

This part is preliminary.

Placebo test We conduct a falsification test where we test the effect of pollution at month $t+4$ on our worker's flows outcomes at t (which is the average of flows at t, $t+1$ and $t+2$).

We should not detect any effect. To ease computational requirement, instead of using wind direction changes at $t+4$ as an instrument for pollution at $t+4$, we use the control function approach described in Appendix 1: we first estimate the first stage at the monthly level and save the residuals. We then test whether the (endogenous) measure of pollution at $t+4$, $PM_{2.5g,t+4}$, controlling for the estimated first stage residual $\widehat{\epsilon_{g,t+4}}$, has a significant impact on workers' flows. The other covariates are the same as in the main specification from equation 3.

Tables 10 and 11 show the results. Except for an imprecisely estimated effect and very low effect on transfers into the establishment and resignations, we do not detect an effect of pollution at $t+4$ on workers' flows at t.

Table 10: Placebo test - Pollution at $t+4$ and workers' outflows at t - IV

	(1) Total	(2) Temporar contract hires	(3) Permanent contract hires	(4) Transfers
$\text{PM}_{2.5t+4}$	0.000731	-0.0000656	0.000174	0.000622*
	(0.000465)	(0.000345)	(0.000114)	(0.000279)
$\widehat{\epsilon}_{t+4}$	-0.000140	0.000219	-0.000243	-0.000116
	(0.000713)	(0.000480)	(0.000150)	(0.000469)
Weather controls	Yes	Yes	Yes	Yes
Holiday and flu controls	Yes	Yes	Yes	Yes
Quarter by Year FE	Yes	Yes	Yes	Yes
Chimere grid FE	Yes	Yes	Yes	Yes
Dependent variable mean	6.8	5.3	1.0	0.5
N	246,310	246,310	246,310	246,310

Notes: Estimates are weighted by the number of establishments in each grid cell. Robust standard errors clustered at the chimere grid level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Outcomes and point estimates have been scaled by a factor 100 to improve readability. The outcomes are standardized by the stock of workers. For example, the inflow outcome designate the share of new workers in the plant as a share of the stock of workers in the beginning of the month.

Table 11: Placebo test - Pollution at $t+4$ and workers' outflows at t - IV

	(1) Total	(2) Resign	(3) End of contract	(4) End of probation and layoff	(5) Transfers
$PM_{2.5}$ $t+4$	0.000120	0.0000198**	0.0000624	0.00000947	-0.000000964
	(0.000352)	(0.00000738)	(0.000347)	(0.00000886)	(0.0000136)
$\widehat{\epsilon}$ $t+4$	0.0000166	-0.0000135	0.0000355	0.0000114	0.00000386
	(0.000487)	(0.00000915)	(0.000482)	(0.0000103)	(0.0000186)
Weather controls	Yes	Yes	Yes	Yes	Yes
Holiday and flu controls	Yes	Yes	Yes	Yes	Yes
Quarter-year FE	Yes	Yes	Yes	Yes	Yes
ID Chimere FE	Yes	Yes	Yes	Yes	Yes
Dependent variable mean	6.2	0.5	4.7	0.4	0.2
N	246,310	246,310	246,310	246,310	246,310

Notes: Estimates are weighted by the number of establishments in each grid cell. Robust standard errors clustered at the chimere grid level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Outcomes and point estimates have been scaled by a factor 100 to improve readability. The outcomes are standardized by the stock of workers. For example, the outflow outcome designate the share of workers leaving the establishment as a share of the stock of workers in the beginning of the month.

References

Adhvaryu, A., N. Kala, and A. Nyshadham (2022). Management and shocks to worker productivity. *Journal of Political Economy 130*(1).

Aguilar-Gomez, S., H. Dwyer, J. S. Graff Zivin, and M. J. Neidell (2022, March). This is Air: The "Non-Health" Effects of Air Pollution. Working Paper 29848, National Bureau of Economic Research. Series: Working Paper Series.

Anderson, M. L. (2020). As the Wind Blows: The Effects of Long-Term Exposure to Air Pollution on Mortality. *Journal of the European Economic Association 18*(4), 1886–1927.

Aragón, F. M., J. J. Miranda, and P. Oliva (2017). Particulate matter and labor supply: The role of caregiving and non-linearities. *Journal of Environmental Economics and Management 86*, 295–309.

Banzhaf, S., L. Ma, and C. Timmins (2019). Environmental Justice: The Economics of Race, Place, and Pollution. *Journal of Economic Perspectives 33*(1), 185–208.

Chang, T., J. Graff Zivin, T. Gross, and M. Neidell (2016). Particulate Pollution and the Productivity of Pear Packers. *American Economic Journal: Economic Policy 8*(3), 141–169.

Chang, T. Y., J. G. Zivin, T. Gross, and M. Neidell (2019). The Effect of Pollution on Worker Productivity: Evidence from Call Center Workers in China. *American Economic Journal: Applied Economics 11*(1), 151–172.

CITEPA (2021). Secten – le rapport de référence sur les émissions de gaz à effet de serre et de polluants atmosphériques en France.

Correia, S., P. Guimarães, and T. Zylkin (2020). Fast Poisson estimation with high-dimensional fixed effects. *The Stata Journal 20*(1), 95–115. Publisher: SAGE Publications.

Dechezleprêtre, A., N. Rivers, and B. Stadler (2019). The economic cost of air pollution: Evidence from Europe. Technical report, OECD, Paris.

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

Deschênes, O., M. Greenstone, and J. S. Shapiro (2017). Defensive Investments and the Demand for Air Quality: Evidence from the NOx Budget Program. *American Economic Review 107*(10), 2958–2989.

Dong, R., R. Fisman, Y. Wang, and N. Xu (2019, December). Air Pollution, Affect, and Forecasting Bias: Evidence from Chinese Financial Analysts. *Journal of Financial Economics 139*(3).

DREES (2020). Les dépenses de santé en 2019 - Edition 2020 - Les indemnités journalières.

European Commission (2020). Air quality: Commission decides to refer France to the Court.

European Environment Agency (2020). Air pollution: how it affects our health.

Fu, S., V. B. Viard, and P. Zhang (2021). Air Pollution and Manufacturing Firm Productivity: Nationwide Estimates for China. *The Economic Journal 131*(640), 3241–3273.

Graff Zivin, J. and M. Neidell (2012). The Impact of Pollution on Worker Productivity. *American Economic Review 102*(7), 3652–3673.

He, J., H. Liu, and A. Salvo (2019). Severe Air Pollution and Labor Productivity: Evidence from Industrial Towns in China. *American Economic Journal: Applied Economics 11*(1), 173–201.

Holub, F., L. Hospido, and U. Wagner (2021). Urban Air Pollution and Sick Leaves: Evidence from Social Security Data.

Hsiang, S., P. Oliva, and R. Walker (2019). The Distribution of Environmental Damages. *Review of Environmental Economics and Policy 13*(1), 83–103. Publisher: The University of Chicago Press.

Krebs, B., J. Burney, J. G. Zivin, and M. Neidell (2021, May). Using Crowd-Sourced Data to Assess the Temporal and Spatial Relationship between Indoor and Outdoor Particulate Matter. *Environmental Science & Technology 55*(9), 6107–6115. Publisher: American Chemical Society.

Lee, D. S., J. McCrary, M. J. Moreira, and J. R. Porter (2021). Valid t-ratio Inference for IV. Working Paper 29124, National Bureau of Economic Research. Series: Working Paper Series.

Lee, S. and J. Lin (2018). Natural Amenities, Neighbourhood Dynamics, and Persistence in the Spatial Distribution of Income. *The Review of Economic Studies 85*(1), 663–694.

Lichter, A., N. Pestel, and E. Sommer (2017). Productivity effects of air pollution: Evidence from professional soccer. *Labour Economics 48*, 54–66.

Meyer, S. and M. Pagel (2017). Fresh Air Eases Work – The Effect of Air Quality on Individual Investor Activity. Technical Report 24048, National Bureau of Economic Research, Inc. Publication Title: NBER Working Papers.

Mink, J. (2022). Putting a price tag on air pollution: the social healthcare costs of air pollution in France.

Nasa Earth Observatory (2003). Chemistry in the Sunlight. Publisher: NASA Earth Observatory.

Pollak, C. (2015). L'effet du délai de carence sur le recours aux arrêts maladie des salariés du secteur privé.

Real, E., F. Couvidat, A. Ung, L. Malherbe, B. Raux, and A. Colette (2021). Historical reconstruction of background air pollution over france for 2000–2015. *Earth System Science Data preprint, in review.*(https://doi.org/10.5194/essd-2021-182).

Schlenker, W. and W. R. Walker (2016). Airports, Air Pollution, and Contemporaneous Health. *Review of Economic Studies 83*(2), 768–809.

Thatcher, T. L. and D. W. Layton (1995). Deposition, resuspension, and penetration of particles within a residence. *Atmospheric Environment 29*(13), 1487–1497.

US EPA (2018). EPA Report on the Environment - Particulate Matter Emissions. Technical report.

Vette, A. F., A. W. Rea, P. A. Lawless, C. E. Rodes, G. Evans, V. R. Highsmith, and L. Sheldon (2001). Characterization of Indoor-Outdoor Aerosol Concentration Relationships during the Fresno PM Exposure Studies. *Aerosol Science and Technology 34*(1), 118–126. Publisher: Taylor & Francis _eprint: https://www.tandfonline.com/doi/pdf/10.1080/02786820117903.

WHO (2014). 7 million premature deaths annually linked to air pollution.

Wooldridge, J. M. (2015). Control Function Methods in Applied Econometrics. *Journal of Human Resources 50*(2), 420–445. Publisher: University of Wisconsin Press.

World Health Organization (2016). WHO Expert Consultation: Available evidence for the future update of the WHO Global Air Quality Guidelines (AQGs). pp. 50.

A Appendix

A.1 Worker-level specification - Pseudo-Poisson Maximum Likelihood estimator

Given the large number of zeros for SLEs per grid cell-week and the fact that the linear probability model is not necessarily the most accurate representation of the impacts of pollution when there are non-linearities, we alternatively use a Pseudo-Poisson Maximum Likelihood (PPML) estimator, in a robustness exercise. It rests on the following assumption for the conditional expectation of the outcome:

$$E[Y_{g,t}|PM_{2.5g,t}, W_{g,t}, h_{g,d,t}, \nu_g, \theta_{q,y}] = exp(\alpha + \beta PM_{2.5g,t} + W_{g,t}\gamma + h_{g,d,t}\delta + \nu_g + \theta_{q,y}) \quad (5)$$

As a result, we can no longer apply 2SLS to equation (1), using the wind direction x Merra grid interaction terms as excluded instruments. Instead, we adopt a control function approach recommended by Wooldridge (2015): we first run the OLS regression of the first stage and obtain OLS residuals $\widehat{\epsilon_{g,t}}$. We then run an augmented second stage where we include both the endogenous regressor $PM2.5_{g,t}$ and the residuals $\widehat{\epsilon_{g,t}}$, using a Poisson pseudomaximum likelihood (PPML) estimator. Specifically, we use the ppmlhdfe command (Correia et al., 2020) to accommodate high-dimensional fixed effects. By separately running the first and second stages, the standard errors of the second stage are underestimated (Wooldridge, 2015); hence, we must bootstrap in order to obtain the right standard errors. In this preliminary version, we present the results with non-bootstrapped standard errors.

Table A.1 reports the PPML estimates of the relationship between PM25 and the sickness leave outcomes, as described in equation 5. The difference in the estimated coefficients for PM2.5 before and after controlling for first stage residuals $\widehat{\epsilon_{g,t}}$ is similar to the difference between the OLS and IV coefficients reported in table 4: without controlling for first-stage residuals, the coefficients on the endogenous PM2.5 are close to zero and not significant,

Table A.1: Pollution and sickness leave outcomes - PPML

	Workers starting sickness leave		Number of sickness days		Sickness leave benefit	
	(1)	(2)	(3)	(4)	(5)	(6)
$PM_{2.5}$	0.000460*	0.00586***	0.00113*	0.00595***	0.00146*	0.00626***
	(0.000200)	(0.000496)	(0.000540)	(0.00144)	(0.000635)	(0.00169)
$\widehat{\epsilon}$		-0.00646***		-0.00579***		-0.00576**
		(0.000556)		(0.00154)		(0.00178)
Weather controls	Yes	Yes	Yes	Yes	Yes	Yes
Holiday and flu controls	Yes	Yes	Yes	Yes	Yes	Yes
Quarter-year FE	Yes	Yes	Yes	Yes	Yes	Yes
Chimere grid cell FE	Yes	Yes	Yes	Yes	Yes	Yes
Dependant variable mean	5.9	5.9	168	168	4,742	4,742
N	1,824,688	1,824,688	1,824,688	1,824,688	1,824,688	1,824,688

Notes: The unit of observation is the postcode x week. All estimates are per 1,000 workers. Estimates are weighted by the number of workers in each postcode. Robust standard errors clustered at the chimere grid cell level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

while they are positive and significant after controlling for the first stage residuals (even-numbered columns). The fact that the estimates associated to $\widehat{\epsilon_{g,t}}$ are statistically significant indicates that ppml estimates without the control function suffer from an endogeneity bias. Not including this correction term would result in biased estimates of the effect of pollution. The estimate from column (2) suggests that an increase by 1 µg/m^3 of weekly $PM_{2.5}$ concentrations is predicted to increase the average number of individuals starting a SLE on the same week by $(e^{0.00586} - 1) * 100 = 0.59\%$. This is close to the 2SLS estimate expressed as a percent change in the mean outcome. For now, the standard errors are underestimated given the two-step estimation procedure. However, the coefficient is relatively precisely estimated. Once standard errors are corrected, it will likely remain significantly different from zero at conventional levels of statistical significance.

Regarding the number of sickness leave days and amount of sickness leave benefit, estimates from columns (4) and (6) indicate that on average, an increase by 1 µg/m^3 of weekly $PM_{2.5}$ concentrations is predicted to increase the number of sickness leave days per worker by $(e^{0.00113} - 1) * 100 = 1.1\%$, and the amount of sickness leave benefits by

$(e^{0.00146} - 1) * 100 = 1.5\%$. These two results are larger than the IV estimates reported as a percent increase from the mean, which suggest that non-linearities may be worth taking into account for these two outcome variables.

A.2 Worker-level specification - Adding controls and fixed effects progressively:

Table A.2 shows how the OLS and IV estimates vary when controls are added progressively, for the main sickness leave outcome (number of workers entering sickness leave). Columns (1) to (4) progressively add controls when running an OLS estimation. Column (2) shows that the positive association between PM2.5 and sickness leave incidence, conditional on holiday status and flu incidence, is stable after controlling for weather characteristics. However, when the time trend is controlled for with the quarter by year fixed effects in column (3), the effect fades out. It may be that part of the effect from column (2) was driven by the seasonal effects influencing pollution (such as that the fact to use more heating in winter) but also sickness leave (even controlling for flu, there are other diseases that are more prevalent in winter). In addition, with the time controls, much of the identifying variation comes from variation in pollution *across* postcodes. Given the spatial sorting of skills across the French territory, there are probably omitted variables at the postcode level that are correlated with both pollution and sickness leave incidence (high-skill individuals are typically likely to live in dense, high-pollution areas, while being less likely to take sickness leave). In column (4), adding Chimere grid fixed effects corrects for this spatial sorting. Column (4) corresponds to the main specification and the results are the same as in column (1) of table 4.

Columns (5) to (8) similarly show the effect of progressively adding controls to the specification using local changes in wind direction as an instrument. In column (5), the first stage predicts PM2.5 levels off variations in local wind direction, controlling only for flu incidence and holidays but not other weather characteristics. From column (6), adding weather controls to the IV estimate affects the estimate much more than adding it to the OLS estimate.

Table A.2: Introducing controls progressively - number starting sickness leave per 1,000 workers

	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) 2SLS	(6) 2SLS	(7) 2SLS	(8) 2SLS
$PM_{2.5}$	0.0247***	0.0223***	-0.000919	0.00388**	0.0296***	0.00457	0.0203*	0.0321***
	(0.00315)	(0.00368)	(0.00382)	(0.00137)	(0.00592)	(0.00840)	(0.00939)	(0.00299)
Flu incidence	0.422***	0.461***	0.333***	0.345***	0.416***	0.466***	0.334***	0.349***
	(0.00896)	(0.0113)	(0.0120)	(0.00840)	(0.0109)	(0.0119)	(0.0122)	(0.00859)
Holiday dummy	-1.54***	-1.51***	-1.34***	-1.34***	-1.53***	-1.53***	-1.33***	-1.34***
	(0.0244)	(0.0255)	(0.0282)	(0.0287)	(0.0219)	(0.0238)	(0.0281)	(0.0293)
Weather controls	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Quarter by Year FE	No	No	Yes	Yes	No	No	Yes	Yes
Chimere grid FE	No	No	No	Yes	No	No	No	Yes
Dep. var. mean	5.9	5.9	5.9	5.9	5.9	5.9	5.9	5.9
N	1,875,939	1,869,578	1,869,578	1,869,578	1,875,939	1,869,578	1,869,578	1,869,578

Notes: Flu incidence emeasures the numer of estimated flu cases per week at the department level. The holiday dummy is observed at the departement level. The unit of observation is the postcode x week. All estimates are per 1,000 workers. Estimates are weighted by the number of workers in each postcode. Robust standard errors clustered at the chimere grid cell level in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Presumably, in column (5) the exclusion restriction is violated: the wind direction changes that predict pollution changes are associated with other weather changes (such as colder temperatures), and these weather changes affect both pollution and sickness leave incidence. Once weather conditions are controlled for in column (6), variations in PM2.5 predicted by variations in local wind direction do not affect sickness leave incidence. The difference between column (6) and (7) may indicate that wind patterns are strongly seasonal: if North wind is more prevailing in winter and South wind in Spring, the change in wind direction will influence sickness leave not only via its effect on pollution, but also due to seasonality. This is again a case of violation of the exclusion restriction. In column (8), finally, the grid fixed effect ensures that the instrument captures the relative change in pollution from a change in wind direction relative to wind coming from the South *in the same grid cell*, rather than relative to the effect of wind coming from the South averaged across all grid cells. This mostly improves the precision of the estimate.

A.3 Additional figures and tables

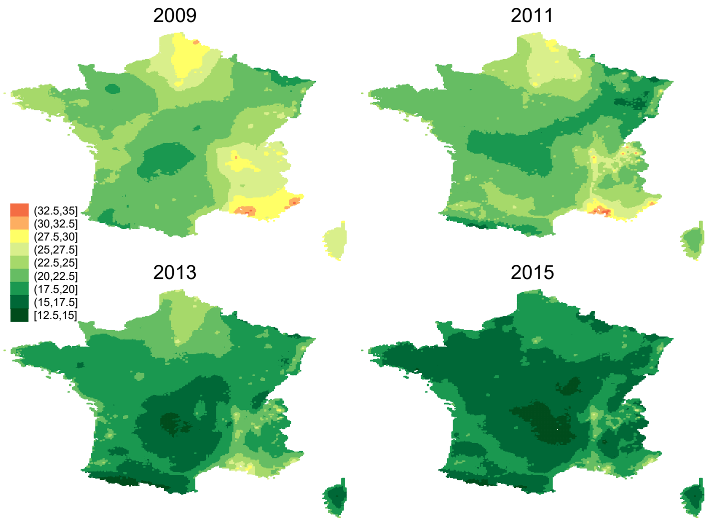

Figure A.1: Average annual concentrations of PM_{10} (µg/m^3)

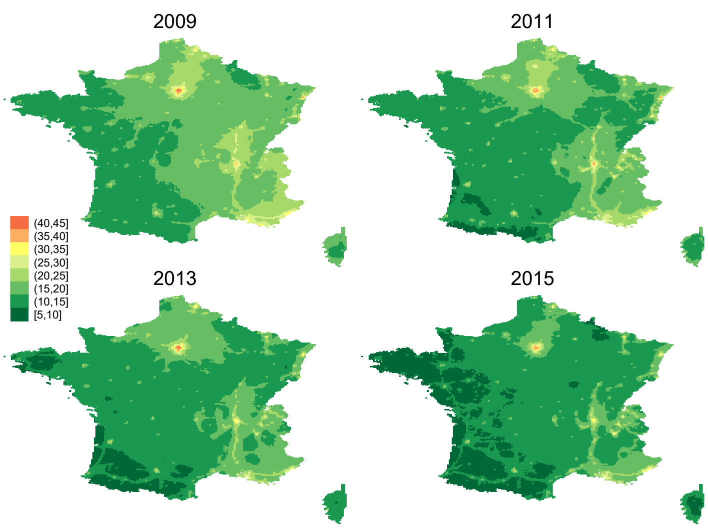

Figure A.2: Average annual concentrations of NO_2 (µg/m^3)

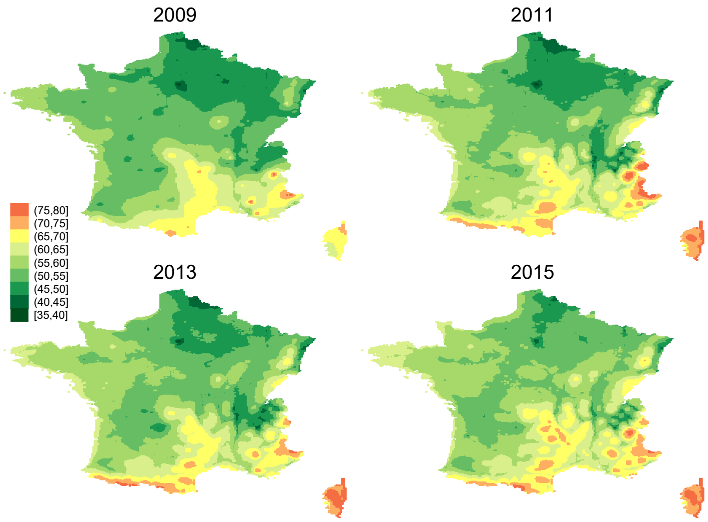

Figure A.3: Average annual concentrations of O_3 (µg/m^3)

Table A.3: Exposure at the place of work vs. place of residence by wage quintile in 2009

	Hourly gross wage (€)	$PM_{2.5}$ exposure work µg/m^3	$PM_{2.5}$ exposure residence µg/m^3	Difference exposure µg/m^3
Q1	9.4	17.7	17.6	0.1
Q2	11.4	17.7	17.5	0.2
Q3	13.8	17.8	17.6	0.2
Q4	18.1	18.0	17.7	0.3
Q5	36.0	18.3	18.0	0.3

Notes: Source: exhaustive matched employer-employee data (DADS Postes) for year 2009. Sample: individuals working in the same set of Chimere grids as the Hygie individuals. Quintile of income based on gross salary divided by number of hours worked for the main position.

Table A.4: Workers quintiles based on healthcare costs covered by the public health insurance in 2009

	Annual covered healthcare costs 2009 (€)	$PM_{2.5}$ exposure 2010-2015 (µg/m^3)	Annual gross wage 2010-2015 (€)
Q1	4	15.0	27,371
Q2	46	14.9	29,244
Q3	123	14.9	28,656
Q4	278	14.9	27,759
Q5	1,152	15.1	26,155

Notes: Covered healthcare costs is obtained by taking the difference between total healthcare costs and out-of-the-pocket spending for visits to GP, visits to specialists and drug purchase. Healthcare costs at hospitals and visits to dentists are not included because the data is not reliable for these categories. Observations are weighted by the number of workers in each grid cell.

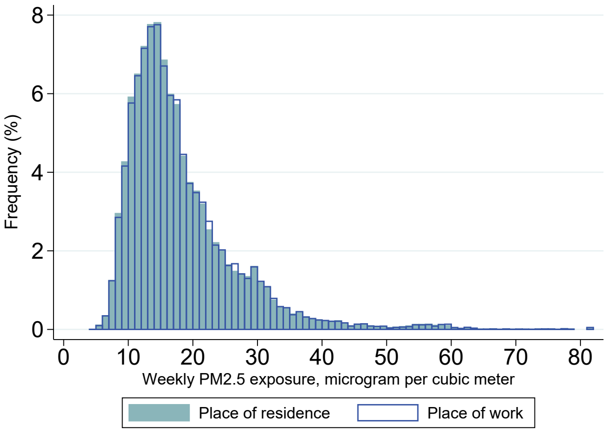

(a) Q1

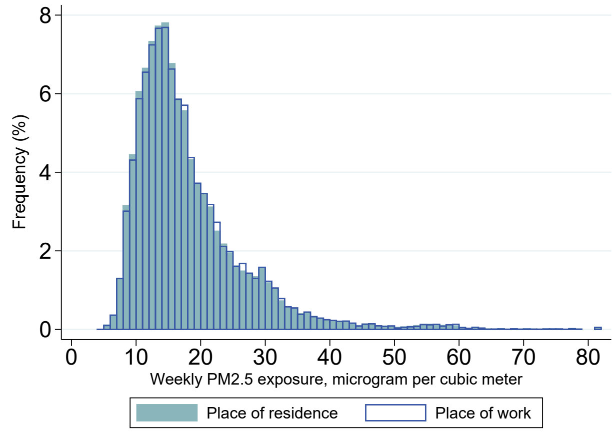

(b) Q2

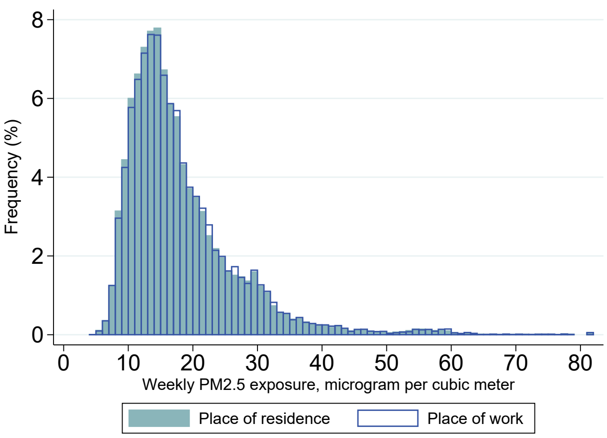

(c) Q3

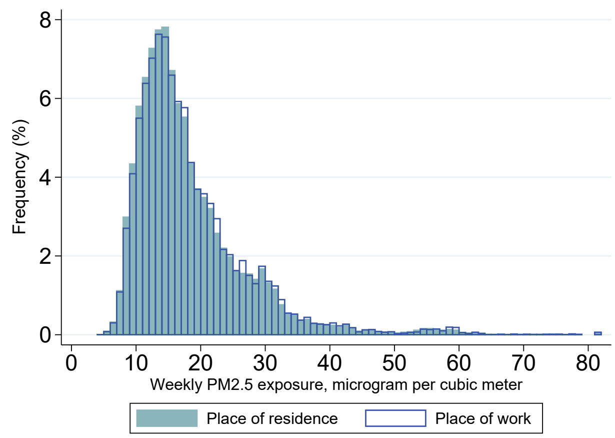

(d) Q4

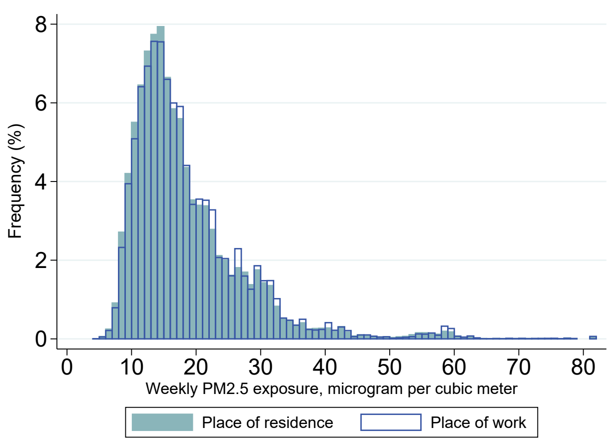

(e) Q5

Figure A.4: Distribution of weekly PM2.5 exposure at the place of residence vs. the workplace, by quintile of hourly gross wage

Note: the sample includes all private sector workers from exhaustive matched employer-employee data, who work in a postcode where we observe at least one individual from the sickness leave dataset. For reasons of statistical confidentiality, the few postcodes with fewer than 11 workers were dropped (<0.4% of workers). Each figure is based on observations from around 2.6 million workers.

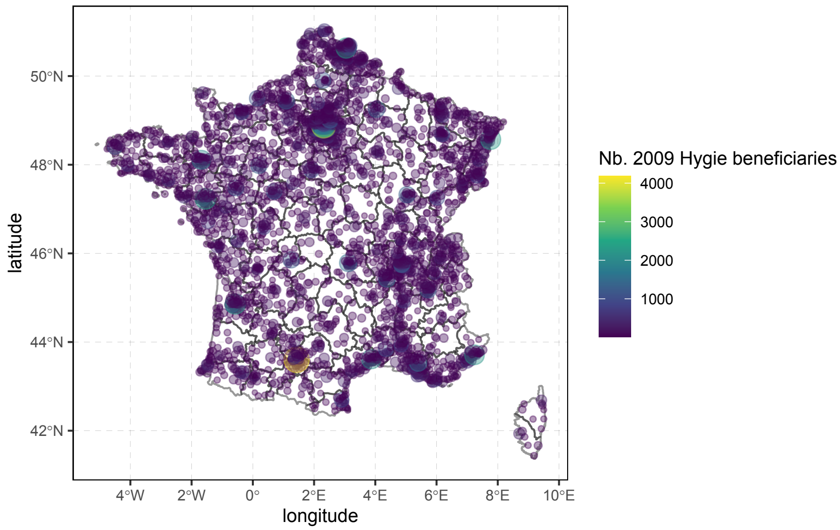

Figure A.5: Location of individuals based on workplace ZIP code

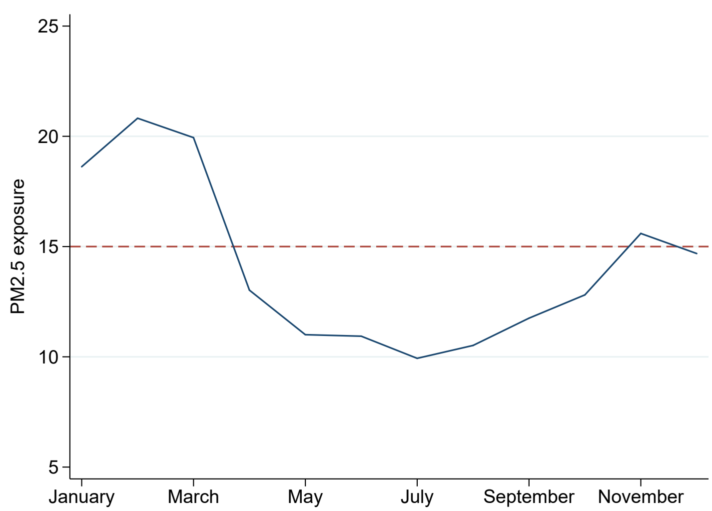

Figure A.6: Seasonality in $PM_{2.5}$ concentrations (µg/m^3)

Notes: $PM_{2.5}$ concentrations are averaged by month, across years between 2009 and 2015. The horizontal red line shows the WHO recommended threshold for daily $PM_{2.5}$ concentrations.

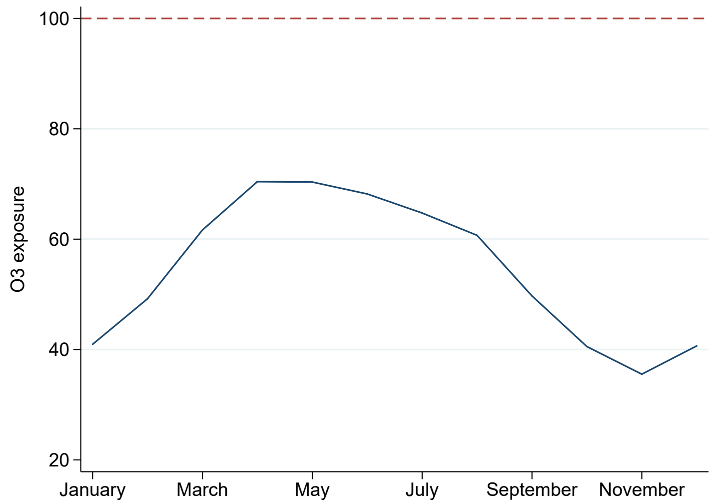

Figure A.7: Seasonality in ozone concentrations (µg/m^3)

Notes: ozone concentrations are averaged by month, across years between 2009 and 2015. The horizontal red line shows the WHO recommended threshold for daily ozone concentrations.

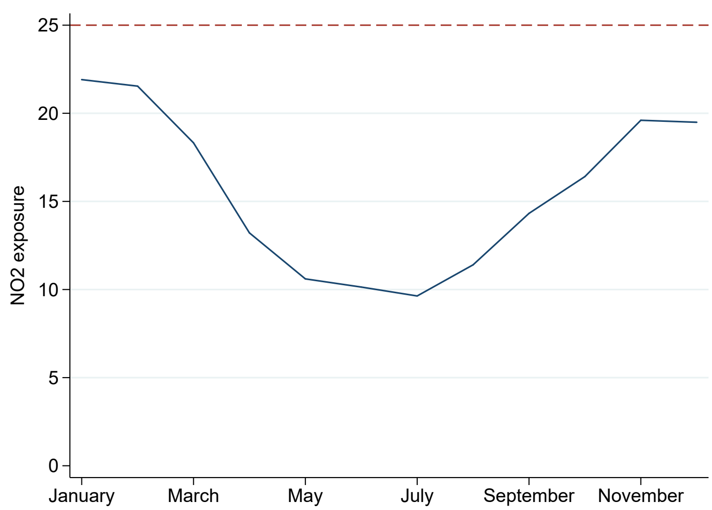

Figure A.8: Seasonality in NO_2 concentrations (µg/m^3)

Notes: NO_2 concentrations are averaged by month, across years between 2009 and 2015. The horizontal red line shows the WHO recommended threshold for daily NO_2 concentrations.

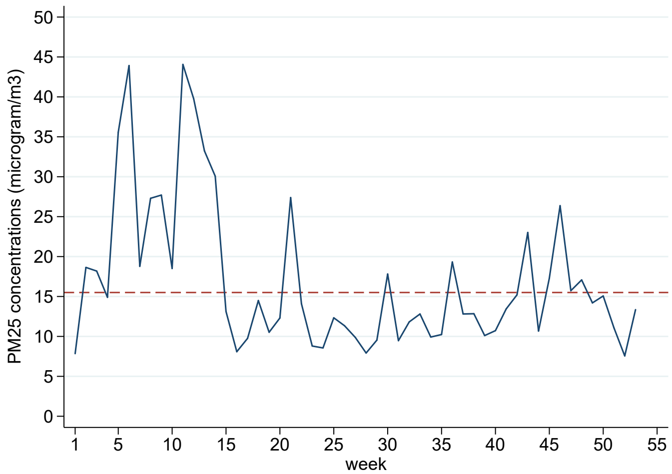

(a) Urban area - Paris

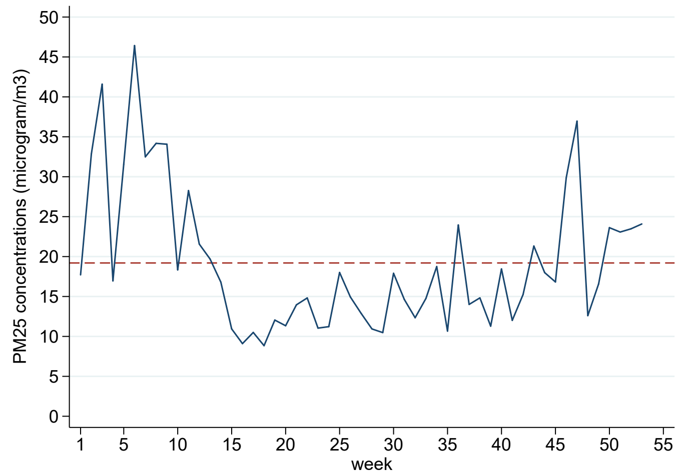

(b) Urban area - Grenoble

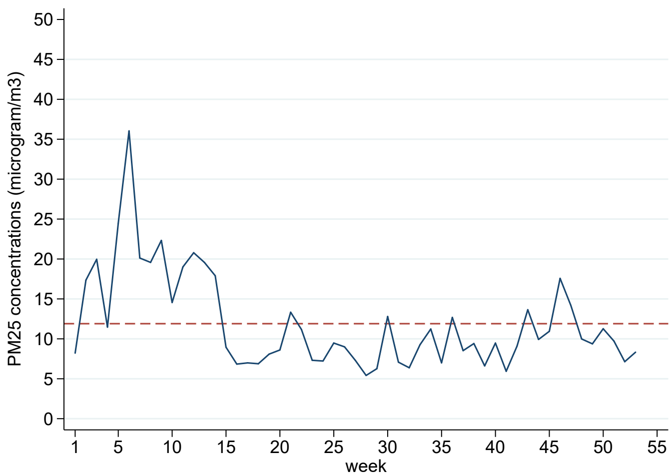

(c) Rural area - Creuse (Centre)

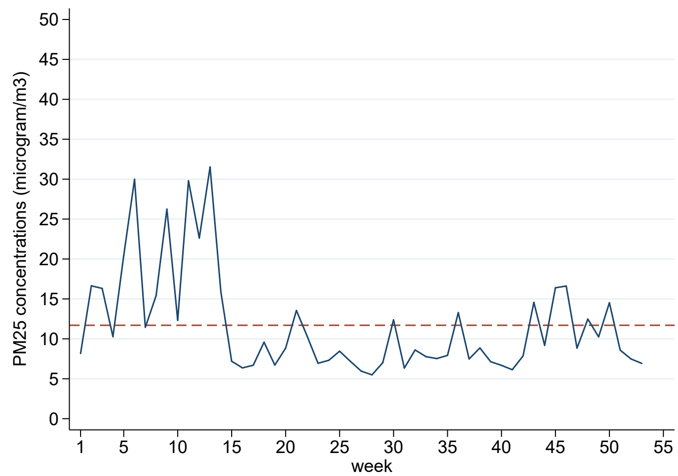

(d) Rural area - Bretagne (West)

Figure A.9: Week-on-week variations in $PM_{2.5}$ pollution for selected grid cells in 2012

Note: The red horizontal line shows the average for that grid cell in 2012.

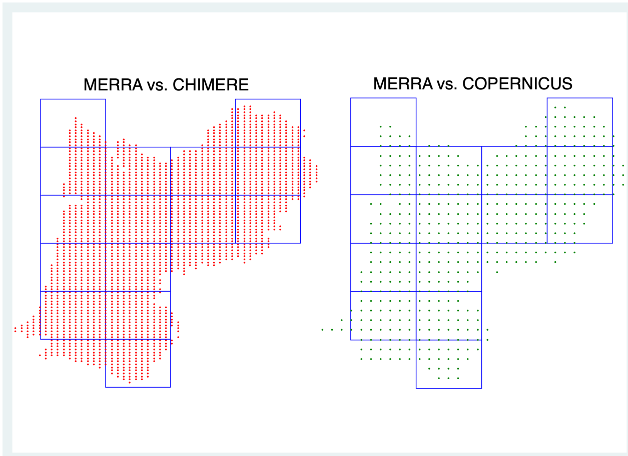

Figure A.10: Grid cell sizes, illustration with one of the 22 French regions, Aquitaine

Notes: Chimere grids are the red dots, Copernicus grids are the green dots, and Merra grids are the blue squares

Table A.5: Workers' quintiles based on annual gross wage in 2009

	Annual gross wage 2009 (€)	$PM_{2.5}$ exposure 2010-2015 (µg/m^3)	Annual covered healthcare costs 2010-2015 (€)
Q1	18,962	14.7	316
Q2	22,295	14.7	316
Q3	26,193	14.8	313
Q4	32,256	15.1	313
Q5	52,074	15.5	316

Notes: Sample: individuals working full-time in 2009 and earning at least the full-time minimum wage. Observations are weighted by the number of workers in each grid cell.