The Cost of Air Pollution for Workers and Firms*

Marion Leroutier, Hélène Ollivier†

June 2026

Abstract

This paper shows that even moderate air pollution reduces economic activity. Using French monthly firm-level data and exploiting wind-driven variation in pollution, we find that a 10% increase in $PM_{2.5}$ exposure lowers sales by 0.23% in the same month and 0.69% in the following month. Effects vary across sectors and operate through reduced labor supply—with a 1% increase in sick leave—, reduced productivity, and demand responses. Sick leave explains about 4% of the sales decline, highlighting the importance of the other channels. Meeting WHO air quality guidelines could yield economic benefits comparable to regulation costs or benefits of reduced mortality.

Keywords: Air pollution, Firms, Absenteeism
JEL codes: Q53, H23, I10, J22

*We are grateful to Tatyana Deryugina, Gregor Singer, Geoffrey Barrows, Abhijeet Singh, Antoine Dechezlepretre, Martina Björkman-Nyqvist, Sampreet Goraya, Maiting Zhuang, Jonathan Lehne, Paul Dutronc-Postel, Xavier D'Haultfoeuille, Kirill Borusyak, Jim Sallee and Joe Shapiro for valuable feedback, as well as attendees at numerous conferences and seminars. We thank Augustin Colette and Elsa Real for sharing the INERIS air pollution data for France, Camille Regaert for help with the Hygie data, and Tom Verrier for outstanding research assistance. Authorization to use the Hygie data was granted by the CNIL in September 2019 (decision DR-2019-286). The results contain modified Copernicus Climate Change Service information from 2022. Neither the European Commission nor ECMWF is responsible for any use that may be made of the Copernicus information or data it contains. Funding from the Mistra foundation and from the French National Research Agency under grant number ANR-22-CE26-0008-01 as well as ANR-17-EURE-001 program are gratefully acknowledged.

†Leroutier: CREST, Institut Polytechnique de Paris, ENSAE. Corresponding author. Email: marion.leroutier@ensae.fr. Ollivier: Paris School of Economics, CNRS. Email: helene.ollivier@psemail.eu.

1 Introduction

Air pollution is widely recognized to have detrimental effects on human health, driving up healthcare costs (Barwick et al., 2024), increasing emergency admissions and mortality (Schlenker and Walker, 2016; Deryugina et al., 2019), and potentially impairing cognitive functions and mental health (Aguilar-Gomez et al., 2022; Bruyneel et al., 2022). These large health costs directly affect the utility of many individuals and are sufficient to justify public intervention. Yet, there might be even wider societal costs if air pollution's impacts on individuals translate into substantial economic losses for firms. Although several papers have examined how air pollution affects workers and firms using detailed data in specific settings, there is limited evidence at the scale of an entire economy.

In this paper, we fill this gap by estimating the causal effects of monthly air pollution exposure on firms' monthly sales in France, a country characterized by moderate pollution levels. The paper begins with an analytical framework that formalizes how air pollution affects private-sector firms' sales in the short run, distinguishing three channels. On the supply side, air pollution may reduce labor supply through increased absenteeism or shorter working hours. It may also lower the productivity of non-absent workers, either through mild health symptoms or through disruptions caused by absent colleagues. Finally, for firms serving local demand, pollution at the firm's location may reduce sales by deterring consumers who are themselves exposed, including through reductions in their disposable income when illness reduces their ability to work. Using administrative micro-level tax and social security data for the private sector, we estimate the overall impact on firm sales and quantify the share attributable to worker absenteeism due to sick leave; the residual captures the combined effects of reduced on-the-job productivity and lower demand.

We assemble a unique data set that combines the monthly sales of 160,000 firms, granular measures of air pollution and weather conditions at the municipality level, as well as sick leave episodes of a representative sample of private sector employees between 2009 and 2015.[1] We focus on exposure to fine particulate matter pollution ($PM_{2.5}$), a pollutant that can penetrate deep into the respiratory tract and enter the bloodstream and brain, with detrimental effects on respiratory and cardio-vascular health, as well as cognitive skills.[2] Particulate pollution can also easily penetrate indoors and affect air quality at work.

Identifying the causal effects of pollution exposure on firms' sales and workers presents three main challenges: first, air pollution is often a byproduct of economic activity; second, individual exposure is typically measured with noise; third, there may be a lag between the time of production and the time of sales.[3] To address reverse causality, measurement error, and endogeneity issues, we leverage variation in air pollution driven by changes in monthly wind directions at the municipality

[1] We define a municipality as a postcode area. An average French postcode area is thirty times smaller than an average US county.

[2] The 2.5 subscript in $PM_{2.5}$ means that these particles have a size smaller than 2.5 micrometers (μm).

[3] Two mechanisms may explain the lag between production and recorded sales. First, some production processes take time (e.g., car manufacturing, or architectural and legal services), and payment typically occurs only after the product or service has been delivered. Second, buyers may delay payment even after delivery, postponing the recording of sales. The average payment delay in the French private sector was 14 days in 2015.

level. To capture the potential lag between air pollution shocks and recorded sales, we assess how pollution exposure influences sales in both the contemporaneous and following months.

We use insights from previous work (Deryugina et al., 2019; Graff Zivin et al., 2023) to build an instrumental variable (IV) based on the changes in each municipality's monthly wind direction. After flexibly controlling for sectoral trends, weather characteristics—including temperature, rainfall and wind speed—, local seasonality and firm-year characteristics, we assume that changes in monthly exposure to cardinal wind directions at the municipality level are unrelated to changes in the sales of firms located in the same municipality, except through the influence of wind direction on air pollution. The benefit of our approach is that it neither requires identifying the sources of pollution in each municipality nor imposes the same relationship between specific wind directions and pollution in groups of municipalities. For firms that own establishments in multiple municipalities, we build an instrument for firm-level pollution exposure by computing a weighted average of predicted pollution exposure at the establishment-month level, accounting for the time-varying structure and geographic location of the intrafirm network.

Knowledge of the complete intrafirm network is crucial for estimating firm-level effects of air pollution. In our sample, firms operate an average of two establishments each year, though some manage many more. Since these establishments are not necessarily located in the same region, our measure of exposure accounts for the geographic dispersion of firms' activities. Furthermore, the intrafirm network evolves over time. By leveraging monthly variation in both sales and air pollution and including firm-year fixed effects, our estimation strategy isolates the causal effect of air pollution while holding the overall firm-year structure fixed.

Our main result is that firm-level exposure to $PM_{2.5}$ has widespread negative effects on sales. We show that a 10 percent increase in firm-level pollution exposure in month t reduces firm-level sales by 0.23 percent in the same month ($p = 0.025$) and by 0.69 percent in the following month ($p < 0.001$). The magnitude varies across sectors: manufacturing sales decrease by 0.28 percent after one month; construction sales decline by 0.36 percent after one month and 0.24 percent contemporaneously; business-to-business trade and services sales fall by 0.44 percent after one month; and business-to-consumer industries experience the largest declines— 1.17 percent after one month and 0.35 percent contemporaneously. The negative effects persist for two to three months after the pollution shock and dissipate after about five months, with no rebound. The results are similar when restricting our sample to single-establishment firms, for which pollution exposure is measured more accurately. They are not driven by restrictions imposed during air quality alerts and are robust to substituting a multi-pollutant air quality index for $PM_{2.5}$, winsorizing the outcome variable, and relying either exclusively on pollution data from monitoring stations or on satellite-based data.

We next turn to the mechanisms behind the pollution-induced decline in sales, asking whether it can be attributed for the most part to labor supply adjustments due to sick leave, the channel we measure most precisely in our data. In France, sick leave requires a medical certificate issued by a general practitioner starting on the first day of absence. This makes our measure of sickness-driven absenteeism very accurate. Yet, we do not observe potential reductions in working hours by employees who continue to work while feeling unwell or who miss work to take care of sick dependents.[4] Therefore these adjustments are part of the productivity channel.

We find that air pollution reduces labor supply via an increase in sick leave. Our estimates imply that a 10 percent increase in monthly $PM_{2.5}$ exposure in the municipality of work increases sick leave episodes by 1 percent within the month of exposure ($p = 0.015$), with no persistence beyond that. The effect is also heterogeneous across economic sectors: it is strong in manufacturing ($p < 0.01$), whereas we cannot rule out a null effect in the other sectors. We quantify that absenteeism accounts for only about 4% of the overall pollution-induced sales decline, and at most 22% in the most affected sector—manufacturing. This limited contribution indicates that the other two channels—productivity and demand reductions—are important drivers of the observed sales losses.

Direct measures of physical output per hour worked are available only for a narrow set of occupations—e.g., agricultural workers, pear packers, call-center operators, or garment workers (Graff Zivin and Neidell 2012; Chang et al. 2016, 2019; Adhvaryu et al. 2022)—where micro-level data can deliver precise estimates of the productivity channel. This precision, however, comes at the cost of representativeness: such measures cannot be deployed across the entire private sector of an economy. A natural alternative in economy-wide settings is to proxy productivity with sales (or value added) per worker. In our context, however, this proxy conflates the three channels we seek to disentangle: its denominator counts workers on the payroll rather than workers actually present, so it cannot separate on-the-job productivity losses from absenteeism; and because its numerator is a market outcome, it absorbs any movement in sales driven by demand shocks. Consistent with this concern, our firm-quarterly estimates show that the response of sales per worker to pollution shocks closely mirrors the response of sales themselves.

To further isolate the productivity channel, we exploit cross-industry heterogeneity in sales responses in manufacturing, a sector where the demand channel is likely to be small. Our fixed-effects structure already absorbs national sectoral demand (through industry-by-month-by-year fixed effects) and local seasonal demand (through quarter-by-county fixed effects), so any residual demand variation must take the form of local, monthly fluctuations around the seasonal average. In manufacturing, customers are typically geographically dispersed, which further mutes this residual demand channel. To verify that demand plays no remaining role, we exploit within-sector heterogeneity in inventory holdings. Firms in industries that typically operate with large inventories can buffer temporary supply-side shocks by drawing down stocks, but inventories do not offer a comparable cushion against demand shocks. We find no significant sales response among high-inventory firms; the manufacturing sales decline is concentrated in low-inventory industries. Critically, both groups experience similar increases in sick leave following a pollution shock and are comparable in size, ruling out the compositional explanation. This evidence is consistent with the pollution-induced sales decline in manufacturing being mostly supply-driven, decomposing into 22% from absenteeism and 78% largely attributable to on-the-job productivity losses.

[4] In France, private-sector employees may take 3–5 days per year to care for a sick child, with the full 5 days applicable only when the child is under 1 year old; these days are unpaid and not captured by our dataset. This channel has, however, been documented in other contexts, such as in Lima, Peru (Aragón et al., 2017).

By contrast, for business-to-consumer industries, consumers are often located in geographic proximity to the establishments they buy from and may therefore share the same pollution shocks. To highlight the demand-side shock in these industries and the behavioral response of consumers, we show that there is a meaningful ranking in sales responses across subsectors, with the largest sales drop for discretionary goods and services, contrasting with a small to non-significant decline in essential goods purchases. This evidence suggests that budget-constrained consumers cut back more on discretionary purchases than on essential items.

Our findings of pollution-induced sales reductions are economically significant. Meeting the WHO guideline of limiting daily $PM_{2.5}$ exposure to 15 µg/m^3 would require a 25% reduction in pollution levels in our sample. Our estimates suggest that such a pollution decrease would have prevented between €25 and €38 billion in annual sales losses from 2009 to 2015, depending on whether dynamic effects after one month are included. With an average value-added-to-sales ratio of 27%, this corresponds to between €6.9 and €10.3 billion in annual foregone value added. We show that these economy-wide costs of air pollution approximately rival the widely emphasized associated mortality costs and may also be equivalent to the associated cost of reducing $PM_{2.5}$ emissions.

To the best of our knowledge, this is the first nationwide study establishing a causal link between air pollution and firm-level sales for the entire private sector. One related paper by Fu et al. (2021) examines the effect of air pollution on annual output among large Chinese manufacturing firms in a high-pollution setting. Average pollution levels in France are four to five times lower than in China or India, similar to those in Europe and fifty percent above those in the US.[5] We highlight the significant economic cost of air pollution in high-income countries, with all sectors incurring sales losses, not just manufacturing. In the European context, Dechezleprêtre et al. (2020) document that a 10% increase in regional $PM_{2.5}$ in Europe decreases regional deflated GDP per capita by 0.8% on the same year, using aggregate data.[6] While we also find substantial economic costs, our granular data capture firms' monthly sales response and our identification relies on within-firm-year variation. Annual output measures may obscure short-term shocks if firms smooth their responses over time. In our context, we find no rebound effect five months after the shock.

By exploring the channels underlying these sales reductions, our paper is connected to the literature that finds that air pollution reduces labor supply, either in the immediate days following exposure (Hanna and Oliva, 2015; Aragón et al., 2017; Holub et al., 2021; Hoffmann and Rud, 2024) or in the subsequent months (Borgschulte et al., 2024). Borgschulte et al. (2024) show that severe wildfire-related pollution shocks in the US reduce average earnings in affected counties, in part because workers exit the labor force. This reflects an extreme labor supply adjustment, likely activated only when more flexible margins, such as reducing hours, taking short-term sick leave, or

[5] In 2015, population-weighted $PM_{2.5}$ exposure was 13 µg/m^3 in France, 8 µg/m^3 in the US, 11 µg/m^{3}in Spain and the UK, 13 µg/m^3 in Germany, and 17 µg/m^3 in Italy (World Health Organization, 2016).

[6] Aggregate data can be prone to measurement error. In many countries, establishment-level output data is not collected. As a result, European regional GDP likely suffer from measurement error as proxies (e.g., number of workers) are used to allocate value added per industry at the regional level.

switching jobs, are unavailable. By contrast, we focus on temporary absenteeism, a more frequent adjustment margin that reveals how pollution disrupts day-to-day work continuity and firm-level productivity. Holub et al. (2021) also show that exposure to PM_{10} increases sick leave in Spain. Drawing on our matched employer-employee data, we demonstrate that the cost of pollution in terms of foregone sales far exceeds the cost attributable to pollution-induced lost work days alone.

We also relate to the literature examining how air pollution affects workers' productivity (along with the studies cited above: Graff Zivin and Neidell 2012; Lichter et al. 2017; He et al. 2019; Dong et al. 2019; Meyer and Pagel 2024). This literature is largely based on the specific settings of one or two firms, where workers are paid by the hour or productivity is easy to observe.[7] By contrast, our data cover the entire private sector of a high-income country, encompassing numerous sectors and job types for which direct measures of physical output per worker are typically unavailable. Thus, we contribute by showing that the economic effects of air pollution are pervasive across the private sector, and not confined to specific occupations where productivity measures are readily available.

In the literature on the economic effects of environmental shocks, demand-side mechanisms have received less attention than supply-side mechanisms. We thus add to a handful of papers studying how consumption behaviors change with temperature shocks (Lee and Zheng, 2025) or salient air pollution shocks (Barwick et al., 2024). In China, Barwick et al. (2024) find a negative impact of daily $PM_{2.5}$ on necessities and supermarket spending within two weeks of exposure, which can be rationalized with reductions in disposable income and short-term avoidance behaviors.[8] Using firms' recorded sales instead of credit card transaction data, we uncover demand-side responses to less salient shocks.

The remainder of the paper is organized as follows. Section 2 provides a brief background on the French context and presents our analytical framework. Section 3 presents the data. Section 4 describes our empirical strategy. Section 5 presents the main results, while section 6 discusses the channels. Section 7 puts the results in perspective, and section 8 concludes.

2 Background and Framework

2.1 The Effects of Air Pollution in the French Context

Particulate matter with a diameter below 2.5 micrometers ($PM_{2.5}$) enters the lungs and can pass into the bloodstream, resulting in significant health problems such as increased mortality and cardiovascular diseases (World Health Organization, 2016; European Environment Agency, 2020).[9] $PM_{2.5}$ also readily penetrates indoors (Chang et al., 2016; Krebs et al., 2021), thereby being likely to

[7] This point was highlighted in a review by Aguilar-Gomez et al. (2022).

[8] Because spending, like sales, is a market outcome and local consumers and retail workers face identical pollution shocks, credit card transaction data cannot separate the role of demand and supply shocks.

[9] $PM_{2.5}$ is related to other air pollutants. In particular, it is by definition included in PM_{10}, but it is deadlier because smaller-sized particles penetrate deeper into the respiratory system. $PM_{2.5}$ can be either directly emitted as "primary" particles, for which the main contributors are the residential and tertiary sector (52%), transportation (20%), manufacturing (18%) and agriculture (11%) (CITEPA, 2021) or formed in the atmosphere as "secondary" particles from the chemical reactions of gaseous pollutants, including SO_2 and NO_2.

affect individuals in their working environment. Exposure to fine particulate matter can temporarily affect cognitive functions: mounting toxicological evidence suggests that it can enter the brain and increase neuro-inflammation and oxidative stress in the central nervous system (Calderón-Garcidueñas et al., 2008). Furthermore, $PM_{2.5}$ can travel far (hundreds of kilometers) and remain in the atmosphere for a long period of time (US EPA, 2018).

Despite these detrimental health effects, air pollution appears to be a low-salience issue in France over our study period. In fact, monitoring and regulation primarily focused on PM_{10} until 2009, with $PM_{2.5}$ only gradually incorporated thereafter. There exists no maximum 24-hour concentration threshold for $PM_{2.5}$, and the annual threshold of 25 µg/m^3, defined by a European Union directive and valid until 2030, never triggers air quality alerts. Although France, like many high-income countries, has experienced significant improvements in air quality over the past decades (Champalaune, 2020; Sicard et al., 2021; Currie et al., 2023), pollution levels regularly exceed public health recommendations. In our sample, daily exposure to $PM_{2.5}$ exceeds the WHO recommended threshold of 15 µg/m^3 on 37% of worker-days. Yet, episodes where pollution levels are high enough to be visible to the naked eye are far less frequent than in heavily polluted Asian cities. Moreover, our study period and context do not include peak pollution events caused by wildfire smoke.

The low saliency of air pollution shocks in France, coupled with moderate pollution levels, has ambiguous effects on firms' economic outcomes. On the one hand, moderate pollution may lead to fewer workers experiencing severe health issues or reduced productivity, and fewer consumers avoiding shopping due to health concerns or mild symptoms. This would suggest minimal impact on output. On the other hand, reduced visibility of pollution shocks could hinder managers' ability to mitigate potential declines in productivity or labor supply, as documented in e.g, Adhvaryu et al. (2022).

Moreover, labor market and social security institutions likely influence how workers, firms, and consumers react to air pollution shocks. Workers' ability to take sick leave varies across countries, sectors, and firms, influenced by differing levels of job protection. In France, all private sector employees are eligible for sickness allowances as long as they provide a medical certificate signed by a general practitioner and have worked at least 150 hours in the past three months. The replacement rate for sick leave varies based on the duration of the leave and can differ across firms due to collective agreements, as well as among workers within the same firm depending on their wage level and type of contract. Survey data indicate that 60 percent of private-sector employees receive full wage replacement from the first day of leave (Pollak, 2015). For the remaining 40 percent, the first three days of leave are not compensated, the replacement rate from day 4 to 7 is 50% of the daily wage, and it increases to 90% of the daily wage from day 8 onward.

Given incomplete replacement rate during sick leave and associated healthcare costs, pollution can trigger a significant loss in disposable income for exposed workers-consumers. We calculate that for a 3-day sick leave episode, income loss amounts to roughly 10% of median monthly wage for the 40% of workers without employer supplements.[10] Such episode is also associated with out-of-pocket

[10] Median wage for this group is taken from Pollak (2015).

healthcare expenditures worth 9% of median monthly wage, based on median daily out-of-pocket expenditures in our sample. These reductions in disposable income may only become apparent to consumers at the end of the month, when they receive their pay and any reimbursements from social security, which may result in delayed consumer responses to air pollution shocks.

2.2 Analytical Framework

In this section, we present a stylized model that connects individual exposure to air pollution with firms' sales. Building on existing literature, we incorporate two supply-side mechanisms—a decline in labor supply and reduced productivity—as well as a demand-side mechanism, which captures income losses and behavioral changes among local consumers.

Demand. We consider an economy in which a representative consumer divides expenditures between a set of differentiated products available in different industries, denoted by $i \in \{1, ..., \mathcal{I}\}$. The utility function takes the following form:

$$U_t = \prod_{i=1}^{\mathcal{I}} \left[\left(\sum_{f \in \Omega_{it}} X_{fit}^{\rho_i} e^{u_{fit}} \right)^{1/\rho_i} \right]^{\nu_i}, \tag{1}$$

where X_{fit} denote the consumption at time t of variety f in industry i and u_{fit} is an *ex post* variety-specific demand shock (realized at the point of sales).[11] The utility function has two tiers. The top tier aggregates consumption in a Cobb-Douglas function across industries, which implies that expenditures on each industry i, Y_{it}, are determined as fixed shares of total expenditures, Y_t: $Y_{it} = \nu_i Y_t$. The second tier aggregates consumption via a Constant Elasticity of Substitution (CES) function across the set of varieties available in each industry i at time t, Ω_{it}. We assume varieties are imperfect substitutes within an industry and ρ_i is the parameter that governs the substitutability of varieties in industry i, with $0 < \rho_i < 1$.

On the demand side, two variables may be influenced by air pollution shocks. First, the ex-post variety-specific demand shock, $u_{fit}(c)$, depends on the level of air pollution exposure, c. Health effects from exposure or avoidance behaviors may lead consumers to alter their spending behavior, such as by staying home and postponing purchases. The sign of the derivative $u'_{fit}(c)$ is ambiguous, however, since consumers may decide to buy more or less of each variety—e.g., higher healthcare expenditures reduce the disposable income for other expenditures whereas staying home may induce a higher demand for food delivery services.[12]

Second, in a developed country context with established sick leave rights and provisions, consumers' income in the period following pollution exposure may be impacted. Therefore, income is given by $Y_t(c) \equiv (1 - \zeta \bar{a}_t(c)) w L_t$, where ζ represents the income loss due to partial sick leave compensation (with $\zeta = 0$ indicating full compensation), $\bar{a}_t(c)$ denotes the average worker absence rate

[11] For simplicity, we assume that $E[e^{u_{fit}}] = 1$ for all firms.

[12] Note that food delivery services are uncommon in France over our study period, accounting only for 0.6% of restaurant sales in 2012 (Source: INSEE)

across firms, w represents the wage rate, and L_t denotes the contractual number of hours worked per employee.[13] We expect $\bar{a}_t'(c) \geq 0$ as higher pollution concentrations likely worsen health effects.

The representative consumer's objective is to maximize her utility (1) given her budget constraint. The CES structure yields an expression for expenditures y_{fit} on each variety f at time t that depends on air pollution exposure, c, through at least the demand-side mechanism:

$$y_{fit}(c) = (p_{fit})^{\frac{\rho_i}{\rho_i - 1}} (P_{it})^{\frac{\rho_i}{1-\rho_i}} e^{\frac{u_{fit}(c)}{1-\rho_i}} \nu_i Y_t(c), \tag{2}$$

where p_{fit} is the price of variety f at time t and P_{it} corresponds to the CES price index for industry i, which is defined in the usual way: $P_{it} = \left[\sum_{f \in \Omega_{it}} (p_{fit})^{\frac{\rho_i}{\rho_i - 1}} e^{\frac{u_{fit}(c)}{1-\rho_i}} \right]^{\frac{\rho_i - 1}{\rho_i}}$.

Production. On the supply side, air pollution exposure influences output through two mechanisms that concur in reducing effective labor, which is the only factor of production. First, workers exposed to pollution shocks may be less productive due to health symptoms, cognitive effects, and work disruptions. Second, some workers may decide to take a sick leave. We assume that each firm produces a single differentiated variety, allowing f to represent both varieties and firms interchangeably. As a result, the production technology for output Q is[14]

$$Q_{fit} = L_{fit}^A(c) \exp(\omega_{fit}) = \lambda_{fit}(c)[1 - a_{fit}(c)]^\theta L_{fit} \exp(\omega_{fit}), \tag{3}$$

where L_{fit}^A denotes effective labor, L_{fit} denotes the number of workers employed at time t, and ω_{fit} is a Hicks-neutral productivity shock that is exogenous to air pollution exposure. Effective labor, L_{fit}^A, responds to air pollution exposure, c, through firm f's marginal productivity of workers at time t without absenteeism, $\lambda_{fit}(c)$, and through firm f's average worker absence rate at time t, $a_{fit}(c)$, combined with a parameter reflecting the attendance impact on marginal productivity, θ. Both mechanisms worsen with higher air pollution levels: $\lambda_{fit}'(c) \leq 0$ and $a_{fit}'(c) \geq 0$.

While the number of workers employed by firm f at time t may not adjust to short-term fluctuations in air pollution, a firm whose employees take leaves of absence experiences a change in the marginal cost of labor. Indeed, we express the firm-specific marginal cost of labor as $w[1 - \eta a_{fit}(c)]$, which depends on the wage rate w, the average worker absence rate $a_{fit}(c)$, and a parameter η that denotes the social security system's contribution to employees' sick leave benefits (with $\eta = 1$ if the social security system covers all sick leave benefits, and $\eta = 0$ if the firms fully compensate their absent workers).

Each firm faces a residual demand curve with constant elasticity $\sigma_i = 1/(1 - \rho_i)$ within industry i and thus chooses the same profit maximizing markup equal to $1/\rho_i$. This yields the pricing rule

$$p_{fit} = \frac{w[1 - \eta a_{fit}(c)] e^{-\omega_{fit}}}{\rho_i \lambda_{fit}(c)[1 - a_{fit}(c)]^\theta}. \tag{4}$$

[13] In a context where wages are flexibly adjusted based on output per hour, air pollution exposure could affect a third variable, the wage rate $w(c)$. In France, such adjustments are unlikely because low-skilled workers are typically paid a regulated minimum wage, and high-skilled workers often negotiate their wages on a long-term basis.

[14] The production function is similar to the one-worker-type production function in Zhang et al. (2017).

Effects of Pollution Shocks on Firms' Sales. Combining (2) with (4) yields the following expression for firm f's sales at time t:

$$y_{fit} = \left(\frac{w[1-\eta a_{fit}(c)]e^{-\omega_{fit}}}{\rho_i \lambda_{fit}(c)[1-a_{fit}(c)]^{\theta}} \right)^{\frac{\rho_i}{\rho_i - 1}} (P_{it})^{\frac{\rho_i}{1-\rho_i}} e^{\frac{u_{fit}(c)}{1-\rho_i}} \nu_i Y_t(c), \tag{5}$$

Taking logs, assuming that the average absence rate is quite small[15] (hence, $\log(1-x) \approx -x$) and reorganizing terms yields

$$\log y_{fit} = \underbrace{\frac{\rho_i}{1-\rho_i} \log \lambda_{fit}(c)}_{\text{Productivity effect}} + \underbrace{\frac{\rho_i(\eta-\theta)}{1-\rho_i} a_{fit}(c)}_{\text{Absenteeism effect}} + \underbrace{\frac{u_{fit}(c)}{1-\rho_i} + \log Y_t(c)}_{\text{Demand effect}} + \delta_{it} + \epsilon_{fit}, \tag{6}$$

with $\delta_{it} \equiv \frac{\rho_i}{1-\rho_i} \log P_{it} + \frac{\rho_i}{\rho_i - 1} \log \left(\frac{w}{\rho_i} \right) + \log \nu_i$ and $\epsilon_{fit} \equiv \frac{\rho_i}{1-\rho_i} \omega_{fit}$. Equation (6) summarizes the three mechanisms through which air pollution affects firms' sales. First, air pollution may decrease the marginal productivity of workers, resulting in sales losses. Second, the labor effectively supplied by employees may decrease with air pollution exposure, especially if they take sick leaves. This mechanism also lowers sales if and only if $\eta < \theta$.[16] Third, firms' sales may fluctuate following an air pollution shock due to consumer behavior changes and the income losses associated with low replacement rates during sick leave.

We draw two main implications from this model. First, sales will decrease with high pollution levels either if all three channels move together or if the productivity and absenteeism effects dominate an opposite demand effect. One of our main objectives in the empirical analysis is thus to evaluate the reduced-form effect of air pollution on firms' sales.

Second, equation (6) shows that the strength of all three channels depends on the elasticity of substitution across varieties within an industry. Industries with highly elastic demand—and therefore low profit margins—are particularly exposed to both supply- and demand-side shocks. Because firms in these sectors cannot raise prices without losing customers, they operate with thin margins and therefore likely suffer disproportionately large sales losses when pollution shocks occur. Using demand elasticity estimates from Harrigan et al. (2024) for France, we distinguish two groups. Manufacturing, construction, and administrative and support activities exhibit relatively low demand elasticities, with average $\sigma_i = 3.89$, 2.67, and 3.34, respectively. In contrast, wholesale trade, retail, and hospitality have much higher elasticities—8.93, 6.03, and 5.52. For illustration, a given log-point decline in worker productivity would reduce sales about 1.74 times more in retail than in manufacturing, purely due to the higher demand elasticity.

[15] In our data, the average share of workers starting a sick leave each month is 23 per 1,000 workers, so the absence rate is very small.

[16] In France, publicly-funded sickness allowances cover only a small share of wages: for instance, for a 5-day SLE with full replacement rate, we calculate that $\eta = 0.2$. Zhang et al. (2017) obtain an estimate of θ equal to 0.46 on Canadian private sector employees, which is higher than η.

3 Data

We combine value added tax records available for essentially all French firms above a reporting threshold, a representative panel dataset of private sector employees affiliated to France's universal sickness-leave insurance, a quarterly panel of the number of employees per firm, and nationwide gridded reanalysis pollution and weather data. Our main analysis is based on two monthly panels over the 2009-2015 period, one at the firm level and one at the establishment level.

Pollution. We use gridded reanalysis air pollution data produced by the French National Institute for Industrial Environment and Risks (INERIS), available at the 4 km by 4 km scale. We obtain hourly concentrations for $PM_{2.5}$, PM_{10}, ozone and nitrogen dioxide. The dataset, described in Real et al. (2022), results from a spatial interpolation of measurements of air pollution concentrations from monitoring stations that is disciplined by the modeled concentrations obtained with a chemistry-transport model built for France named CHIMERE.[17] The resulting dataset is better suited to capture the average pollution exposure of local residents than monitor readings. Monitors are sparse, so their readings may not take into account all polluting sources.[18] On the other hand, the CHIMERE chemistry transport model uses meteorological variables as inputs, including wind patterns, that may produce a mechanical correlation between the modeled $PM_{2.5}$ and our wind-based instrumental variables. In section 5.3, we show that our results are robust to using $PM_{2.5}$ exposure based on a simpler spatial interpolation of monitor readings. Figure A.1 shows the spatial distribution of annual exposure at different points in time and the significant reduction in average $PM_{2.5}$ concentration over the period.

Weather. We use gridded reanalysis weather data from the Copernicus Climate Change Service (ERA5 dataset, from Hersbach et al. 2018). We obtain hourly precipitations, surface temperature, wind direction, and wind speed at the 0.25° x 0.25° resolution (approximately 28 km by 28 km). We compute monthly averages for daily maximum temperature and hourly wind speed, and sum hourly precipitation over each month. For wind direction, we compute for each month the share of hours when the wind blows from each of four cardinal directions: North (below 45° or above 315°), East (between 45° and 135°), South (between 135° and 225°) and West (between 225° and 315°).

Firm-level sales. We use monthly sales data at the firm level from Value Added Tax (VAT) records collected by the French administration. The tax administration imposes monthly declarations to firms with annual sales above certain industry-specific thresholds, while small-sized firms are allowed to report either monthly or quarterly.[19] Firms filling monthly VAT declarations account for 66% of all French firms, but generate 91% of total sales (France Stratégie and Inspection

[17] Pollution concentrations are obtained by co-kriging measurements from background monitoring stations onto a 4km×4km grid, using outputs from the CHIMERE chemistry-transport model as auxiliary variables. The procedure follows Real et al. (2022), who validate the resulting dataset through leave-one-out cross-validation and show that it provides reliable estimates of background pollution levels across most pollutants.

[18] Over the study period, there are between 62 and 105 background monitoring stations for $PM_{2.5}$.

[19] The threshold is €818,000 of annual sales for manufacturing and hospitality, and €247,000 of annual sales for the other sectors.

générale des Finances, 2021). Our dataset therefore provides broad coverage of the French private sector.

We restrict our sample to firms that operate in one of four broad economic sectors: manufacturing, construction, business-to-business trade and services (communication and IT services, wholesale trade, professional services, and cleaning services), and business-to-consumer retail and services (groceries and supermarkets, restaurants, hairdressers, clothing stores, furniture stores, and car sales and repair). This restriction excludes farming, the financial sector, the not-for-profit sector, and two industries for which the timing and location of sales are often disconnected from the timing and location of consumption: hotels and transportation. We use the 2-digit level of the European Union industry classification to identify 88 industries grouped into the four main sectors. Our final dataset includes firms with at least one employee covered in the sick leave dataset described below. It comprises 158,223 firms totaling €1,900 billion sales in 2013 which represents 56% of all sales in the four broad sectors defined above.[20]

Sales are reported at the firm level. For the 64% of firms with a single establishment, we assign monthly pollution and weather exposure from the nearest grid cell to the establishment's municipality in the pollution and weather data. The remaining 36% percent of firms own more than one establishment and jointly represent 75% of total sales in our sample. To build firm-level pollution and weather exposure for them, we leverage exhaustive matched employer-employee data that provide for each firm the number and location of all its establishments and the number of workers employed in each establishment each year.[21] We build a weighted-average firm-level exposure to pollution and weather characteristics, with weights equal to the annual number of workers employed in each establishment owned by the firm.

Sick leave episodes. We obtain data on sick leave episodes (SLE) for a representative sample of private sector employees (Hygie dataset). This dataset reports for each worker the exact start date and duration of each SLE that occurred between 2009 and 2015, as well as individual characteristics (gender, age, wage, annual medical expenditures). Our measure of absenteeism is an indicator for an individual starting a SLE in a given month. We focus on SLEs that last less than three months, capturing 93% of the spells.[22]

We restrict our dataset to workers with an establishment-level identifier (see Appendix C for more details). This enables us to attribute pollution and weather exposure to each employee based on the municipality of their workplace, absent information on their municipality of residence.

[20] We restrict the sample to firms for which we can estimate both sales and absenteeism responses to compare their magnitudes. Because Hygie is a representative sample of private sector employees across industries, this restriction is unlikely to induce a bias in our sample, except for possible underrepresentation of industries with high prevalence of self-employed relative to salaried workers. Our final dataset covers 73% of nationwide manufacturing sales, 43% of construction sales, 55% of sales in the business-to-consumer goods and services sector, and 49% of the sales in the business-to-business trade and services sector.

[21] We build exposure to pollution and weather shocks using the exact network of establishments within each firm for each year. It is necessary to update the network each year because of substantial entry and exit of establishments within firms. From 2009 to 2015, we observe 67,878 entries or exits of plants. The churn rate, i.e., net plant turnover relative to the number of plants per firm-year, is 2.34%.

[22] In our data, the average sick leave episode lasts 29 days whereas the median duration is only 9 days. Figure C.2 shows how a small number of SLEs lasting more than 3 months influences this average duration.

Our analysis of exhaustive matched employer-employee data shows that the distributions of $PM_{2.5}$ exposure at workplace and residential municipalities are nearly identical (see Figure A.3).[23] Panel (a) in Figure A.2 shows the resulting average monthly exposure of workers over the period. Daily exposure to $PM_{2.5}$ exceeds the WHO recommended threshold of 15 µg/m^3 (red line) on 37% of worker-days. Panel (b) illustrates the substantial variation in monthly exposure to $PM_{2.5}$ within a quarter-year. We aggregate sick leave data at the establishment-month level.

Establishment-level employment. To construct establishment-level employment, we use the Labor Movements survey (MMO, Mouvements de Main d'Oeuvre), an administrative dataset that records worker flows at least at quarterly level for French private-sector, non-agricultural establishments. All establishments with more than 50 employees are included, as well as a random sample of those with fewer than 50 employees, based on a survey. The survey response rate dropped in 2015, so we restrict the data to 2009-2014. We use the headcount of workers at the beginning of each quarter, available at firm level only for a sample of single-establishment firms, as the denominator of our sales-per-worker measure.

Descriptive statistics. Panel (a) of Table 1 shows that firms in our sample employ on average 60 workers and report €1.32 million in monthly sales, while the medians—15 workers and €145,372—indicate substantial skewness. The sectoral composition is balanced, with 20% of firms in manufacturing, 16% in construction, 31% in business-to-business trade and services, and 33% in business-to-consumer industries.

Panel (b) reports descriptive statistics for workers with sick leave information employed in firms with monthly VAT records. This sample—used for our absenteeism analysis—includes roughly 400,000 workers in 353,155 establishments between 2009 and 2015. Workers are on average 40 years old, earn €28,542 annually, and incur €442 in annual medical expenditures, including €140 out-of-pocket. Each month, about 23 per 1,000 workers begin a SLE lasting less than three months. Appendix Table A.1 compares this sample with the initial representative worker sample. Because firms with monthly VAT filings are typically larger, workers in our main sample have somewhat higher earnings. However, both samples exhibit similar demographics, sick leave rates, and pollution exposure.

[23] In 2015, 27% lived and worked in the same municipality, and the 2017 median commuting distance was 9.2 km (INSEE, 2021). Figure A.3 shows nearly identical workplace and residence exposure distributions (overall and by income quintile) based on the 2009 exhaustive matched employer–employee data.

Table 1: Summary Statistics, 2009-2015

	Mean	Sd	Count
Panel a: Firms' characteristics			
Single-establishment	0.64	0.48	9,832,620
Number of workers	59.68	482.76	9,832,620
Monthly sales (k€)	1,316.30	18,153.87	9,831,760
Share in: Manufacturing	0.20	0.40	9,832,620
Construction	0.16	0.37	9,832,620
Business-to-business trade and services	0.31	0.46	9,832,620
Business-to-consumer retail and services	0.33	0.47	9,832,620
Monthly exposure to $PM_{2.5}$ (µg/m^3)	15.17	6.22	9,832,620
Panel b: Workers' characteristics (aggregated at establishment level)			
Age	40.19	8.74	8,233,440
Annual wage (euros €)	28,541.97	20,576.10	8,233,440
Annual medical expenditures (€)	442.02	809.78	8,233,440
Annual out-of-the-pocket medical expenditures (€)	139.88	172.21	8,233,440
Works in a single-establishment firm	0.40	0.49	8,239,344
Nb workers falling sick per month, per 1,000 workers	24.70	113.44	8,239,344
incl: for <93 days	23.00	109.24	8,239,344
Nb of associated sick days per 1,000 workers	758.91	9,404.01	8,239,344
incl: for <93 days	363.52	2,655.22	8,239,344
Share in: Manufacturing	0.28	0.45	8,239,344
Construction	0.12	0.32	8,239,344
Business-to-business trade and services	0.33	0.47	8,239,344
Business-to-consumer retail and services	0.27	0.42	8,239,344
Monthly exposure to $PM_{2.5}$ (µg/m^3)	15.34	6.33	8,239,344

Notes: Data from panel a) is based on the firm-level dataset. Data from panel b) is based on the establishment-level dataset, with weights corresponding to the number of workers for whom we observe sick leave within each establishment.

4 Empirical Strategy

Our objective is to identify the short-term causal effect of $PM_{2.5}$ on firms' sales and on their employees' absenteeism due to sick leave. Our main identification challenge is that there may be unobserved determinants of both local air pollution and firms' sales and worker absenteeism. These determinants include time-invariant characteristics, such as local economic conditions, and time-varying factors, such as weather conditions or demand seasonality. To address these concerns, our econometric specification combines a rich set of fixed effects with instrumental variables.

4.1 Effect on sales: firm-level econometric specification

We model the relationship between firms' sales and pollution exposure using the following equation:

$$Y_{fiyt} = \sum_{\tau=-1}^{T} \beta_\tau PM_{fyt-\tau} + \mathbf{W'_{fyt-1}}\gamma_0 + \mathbf{W'_{fyt}}\gamma_1 + \mathbf{W'_{fyt+1}}\gamma_2 + \nu_{fy} + \theta_{iyt} + \delta_{cq} + \epsilon_{fiyt}, \tag{7}$$

where the unit of observation is a firm f producing in industry i in month t in year y. The outcome Y_{fiyt} is the logarithm of the average sales recorded by firm f for months t and $t+1$ in year y. This aggregation smooths sector-specific and idiosyncratic variability in the lag between the economic transaction and its recording as sales.[24] $PM_{fyt-\tau}$ is either contemporaneous ($\tau = 0$), leaded ($\tau = -1$) or lagged ($\tau > 0$) pollution exposure at firm level, and T is the number of lagged pollution variables.

We include a rich set of time-varying controls $\mathbf{W_{fyt}}$ for weather conditions and holiday. Specifically, we generate bins for monthly averages of daily maximum temperatures, wind speed and precipitation in each location, and include the set of indicators for all possible interactions of these weather parameters as controls.[25] When firms own multiple establishments, we build these weather controls based on weighted averages of the values taken at each establishment. To account for the lower economic activity and pollution levels during holiday periods, we also include the monthly count of days of school holiday in each location.[26] Given that our outcome is averaged at t and $t+1$, our regressions include time-varying controls at t and $t+1$.

We also control for firm-by-year fixed effects (ν_{fy}), industry-by-month-by-year fixed effects (θ_{iyt}), and quarter-by-county seasonality (δ_{cq}).[27] Firm-by-year fixed effects isolate variation in pollution exposure around the mean exposure of a firm at the annual level, thereby absorbing any annually-invariant firm characteristics while also controlling for annual shocks jointly affecting exposure to pollution and sales. Such shocks include any productivity shock or any change in the number or location of establishments belonging to a firm, which we only observe at the annual level. Industry-by-month-by-year fixed effects capture monthly shocks that are common across all firms in the same industry. Quarter-by-county fixed effects capture seasonality in pollution, and also in wind patterns for the instrumented version, that are specific to a county and may be correlated with local seasonal fluctuations in economic activity. It captures for instance the seasonal demand variation in ski or sea resort areas.

The key parameters of interest are β_τ, which capture the contemporaneous and delayed effects of monthly air pollution exposure on firms' sales. In our main specification, we include one lag

[24] Sales and VAT must be declared to the tax administration in the month of delivery for domestic goods, in the month of payment for domestic services, and one month after delivery for exported goods and services.

[25] For temperature we define 12 bins spanning 3°C each, except for the first bin including all negative temperatures, and for the twelfth bin including all temperatures above 33°C. For wind speed and precipitation, we compute indicators for each quintile of these variables.

[26] Beside the July-August and Christmas school breaks, which occur at the same time for all schools in France, the two-week school breaks in the Fall, Winter, and Spring are staggered by region.

[27] We use "county" to denote a French *département*, the second smallest administrative subdivision after municipality, on average 1.8 times larger than a US county. There are 96 *départements* in mainland France.

of pollution exposure ($T = 1$) and interpret β_1 as our main coefficient of interest, but still report β_0. The reasons for focusing on this delayed effect vary by sector. First, in some industries, there is a lag between production and delivery of the finished product. These lead times depend on product complexity and supply-chain structure, and can be substantial: for example, lead times for steel, concrete, and prefabricated components in France are estimated at 6 to 12 weeks.[28] Second, sales may be recorded either at delivery—as in manufacturing, wholesale trade and retail—or at payment, as in construction, business-to-business or consumer services. When consumers delay payment, a gap arises between the time of purchase and the time of recorded sales, and these delays can be substantial.[29]

We also explore the dynamic effects of air pollution on sales up to five months after exposure by considering five lags in (7). To reduce the noise due to the serial correlation in wind direction and pollution exposure over time, we use a polynomial distributed lag (PDL) specification (Schwartz, 2000; He et al., 2019; Barwick et al., 2024) and impose a smooth polynomial function on the lag structure to discipline the coefficients. We assume a cubic polynomial functional form on the coefficients β_l, for $l \in \{0, .., 5\}$: $\beta_l = \sum_{k=0}^{3} \gamma_k l^k$. For example, $\beta_0 = \gamma_0$, $\beta_1 = \gamma_0 + \gamma_1 + \gamma_2 + \gamma_3$, and $\beta_2 = \gamma_0 + 2\gamma_1 + 4\gamma_2 + 8\gamma_3$ for the first parameters. Using these relationships, we rewrite the regression equation as a function of γ_ks and estimate the coefficients γ_0, γ_1, γ_2, and γ_3 by OLS and two-stage-least-squares for the instrumented version. Combining these point estimates and associated standard errors, we recover the point estimates β_ls and their associated standard errors.

4.2 Effect on sick leave: establishment-level econometric specification

We model the relationship between contemporaneous pollution exposure and worker absenteeism at the establishment level using the following equation:

$$Y_{eiyt} = \beta^A PM_{gyt} + \mathbf{W'_{gyt}}\gamma + \nu_e + \theta_{iyt} + \delta_{cq} + \epsilon_{eiyt}, \tag{8}$$

where the dependent variable Y_{eiyt} is the sick leave outcome measured in month t in year y in establishment e operating in industry i. Unlike sales, we observe worker absenteeism at the establishment level, even within multi-establishment firms. The parameter of interest is β^A, the coefficient on contemporaneous monthly $PM_{2.5}$ exposure for establishment e located in municipality g. Pollution exposure and control variables $\mathbf{W_{gyt}}$ are defined as before, and computed at the municipality level. As in (7), we control for industry-by-month-by-year (θ_{iyt}) and quarter-by-county (δ_{cq}) fixed effects. Additionally, we control for establishment fixed effects, ν_e, which isolate monthly variation in pollution exposure within an establishment and absorbs any time-invariant establishment-specific characteristic. To assess the robustness of our specification, we explore the

[28] These estimates are components of the HCOB France Manufacturing Purchasing Managers' Index (PMI), compiled by S&P Global. See https://www.spglobal.com/market-intelligence/en/solutions/products/pmi [Online; accessed 11 January 2025].

[29] For instance, mean payment delays in 2015 amount to 11 days in construction, 17 days in business support activities, and 15 days in business-to-consumer services on average (Prost and Villetelle, 2018).

dynamic effects of pollution exposure on worker absenteeism using five lags, as we do for sales.

4.3 Wind direction instruments

Despite the inclusion of high-dimensional fixed effects, OLS estimates of equation (7) remain vulnerable to bias from reverse causality, measurement error in pollution exposure, and omitted variables. Higher sales can mechanically raise air pollution through higher production and higher road traffic, creating reverse causality. Measurement error also arises when individuals' pollution exposure is assigned based on the workplace municipality only. Under classical (mean zero and i.i.d) measurement error, this induces attenuation bias, potentially amplified by fixed effects (Griliches and Hausman, 1986). Unobserved local shocks that jointly affect pollution and economic outcomes are also a concern. For instance, a positive shock to local demand outside seasonal patterns—e.g., due to sports or cultural events—could raise retail and service-sector sales while increasing car traffic and related transport $PM_2.5$.

To address these remaining potential biases, we rely on an instrumental variable approach exploiting month-to-month variation in wind direction at the municipality level, in the spirit of Deryugina et al. (2019) and Graff Zivin et al. (2023). We instrument monthly pollution exposure with a combination of the share of hours in a month where wind blows from each of the four cardinal directions (South, West, East, and North) and a pollution intensity factor for each direction in each municipality. This flexible approach acknowledges that a given wind direction might affect air pollution differently in different municipalities, depending on the location of polluting sources. Following Graff Zivin et al. (2023), we compute four instruments Z^j_{gyt}, one for each wind direction $j \in \{South, West, East, North\}$ as follows:

$$\mathrm{Z}^j_{gyt} = \underbrace{WIND^j_{gyt}}_{\text{A: Time-varying}} \underbrace{\left(\frac{1}{N^j}\sum_{d\in D^j} PM_{gd} - \frac{1}{N}\sum_{d\in D} PM_{gd}\right)}_{\text{B: Time-invariant}} \tag{9}$$

where component A, $WIND^j_{gyt}$, identifies the share of hours in calendar month t in year y where the wind blows from direction j in municipality g, while component B reflects the average deviation from daily mean pollution levels on days where the wind blows from direction j in municipality g, across the entire sample period.[30] N^j and D^j are the number and set of days where the dominant wind blows from direction j, and N and D are the total number and set of days over the period of analysis.

Component B is a time-invariant component akin to a pollution intensity factor by wind direction, calculated for each municipality using the average pollution level across sample days, by wind direction and overall. Figure 1 plots this component for each wind direction across munic-

[30] A one-unit increase in Z^j_{gyt} reflects different combinations of the frequency of wind j and its influence on the pollution level in municipality g. For instance, a 10 pp increase in the share of North wind in municipality A, where North wind deviates from the mean pollution by 0.1µg/m^3, and a 20 pp decrease in the share of North wind in municipality B, where North wind deviates from the mean pollution by -0.05µg/m^3, would both result in a one-unit increase in Z^{North}_{Ayt} and Z^{North}_{Byt}.

ipalities in France. While East (West) winds increase (decrease) pollution in the vast majority of municipalities, there is still a lot of variation in the magnitude of the increase (decrease). By contrast, winds blowing from the North and the South have heterogeneous effects across regions: North (South) winds increase (decrease) pollution in the Northern half of the country, while having moderate effects in the Southern half of the country.

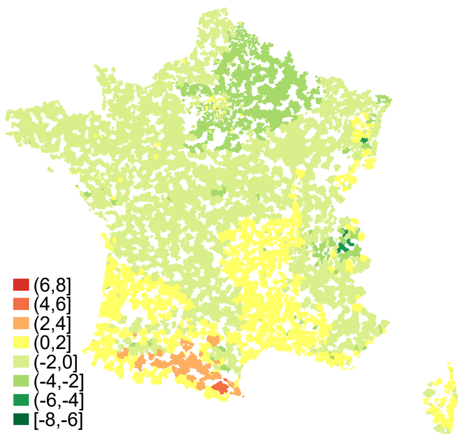

(a) South wind

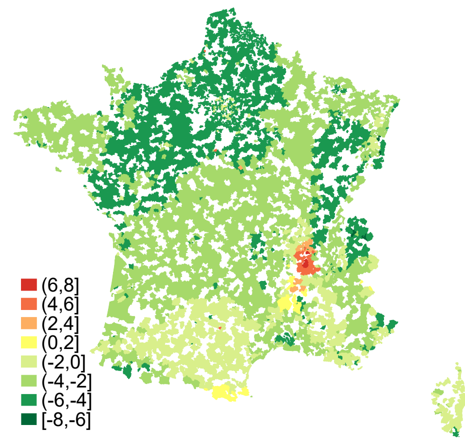

(b) West wind

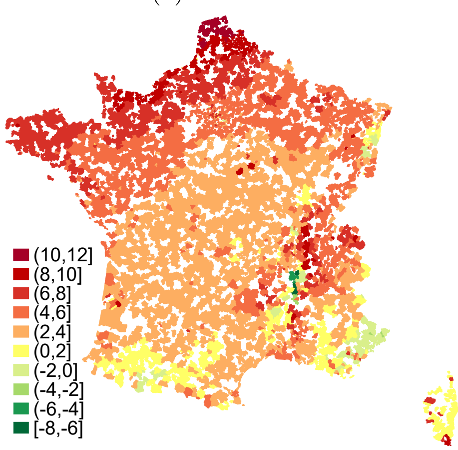

(c) East wind

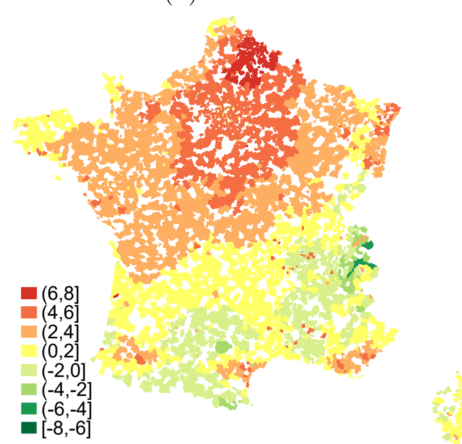

(d) North wind

Figure 1: Deviation from daily mean PM 2.5 for each wind direction

Notes: Figure shows the component B of the instrument Z_{jgyt}, which describes for each municipality the deviation from daily mean pollution levels on days where the dominant wind blows from direction j.

To the extent that component A—the month-to-month variation in wind direction—is as-good-as-random conditional on the controls and fixed effects, the potential endogeneity of component B is not a threat to identification: our design is close to a shift-share instrument with a shift-based identification (Borusyak et al., 2025), where component A is akin to the exogenous shift and component B associates different intensities to the random shifts and enhances the power of our instrument. We show in section 5.3 that our results are robust to different designs of the wind instruments.

For single-establishment firms, pollution and weather exposure can be precisely measured at a single location; hence, the specification of the first stage includes all four wind-direction instruments:

$$PM_{fiyt} = \sum_{j=1}^{4} \beta^j \mathrm{Z}_{gyt}^{j} + \mathbf{W'_{gyt}}\gamma + \nu_{fy} + \theta_{iyt} + \delta_{cq} + u_{fiyt}, \tag{10}$$

with PM_{fiyt} and weather controls $\mathbf{W'_{gyt}}$ varying at the municipality level g, β^js the parameters of interest, and fixed effects defined like in equation (7). For a given wind direction j, β^j captures the average effect of a marginal increase in the intensity of wind direction j. The identifying variation is the quasi-random change in wind direction intensity stemming from variation in component A of the wind instrument, around the mean exposure of each municipality within a year, after partialling out quarter-by-county-specific variation, industry-specific national trends, and after controlling for other weather conditions. In our specifications with different lags of the pollution variable, we instrument each pollution lag (or lead) with the wind instrument at the corresponding time.

Figure A.5 plots the distribution of the raw and residualized wind instrument variables for the subsample of single-establishment firms. To match the main specification where the coefficient of interest is for lagged pollution, we residualize the lagged wind direction instrument by regressing it on the fixed effects, the controls, the contemporaneous wind instrument and one lead of the wind instrument. There remains substantial variation in each instrument after residualizing. Figure A.6 plots similar distributions for the time-varying part of the instruments only (component A), and shows substantial remaining variation. We further illustrate the substantial variation in wind direction within a municipality within a calendar month in Figures A.7 and A.8 for the two largest French cities, Paris in the North and Marseille in the South-East.

For multi-establishment firms, we generate a plausibly exogenous predicted pollution exposure using the results from a first stage equation specified at the municipality level. We regress PM_{gyt} on the same variables as in equation (10), except that we replace firm-by-year fixed effects with municipality-by-year fixed effects, and month-by-year-by-industry fixed effects with month-by-year fixed effects. After saving the vector of estimated $\widehat{\beta^j}$, we compute the predicted pollution exposure in each municipality as $\widehat{PM_{gyt}} = \sum_{j=1}^{4} \widehat{\beta^j} \mathrm{Z}_{gyt}^{j}$. We then compute the firm-level predicted pollution exposure, $\widehat{PM_{fyt}}$, as the weighted average of $\widehat{PM_{gyt}}$ across municipalities g where firm f owns establishments in year y using labor shares as weights. We use $\widehat{PM_{fyt}}$ as an instrument for PM_{fyt} in equation (7).[31]

In our main analysis, we include both single- and multi-establishment firms and we instrument pollution exposure with the predicted pollution measure. For robustness, we restrict the sample to single-establishment firms and directly use the four wind direction instruments. We cluster the standard errors at the Copernicus grid cell level based on the location of the single establishment or of the headquarter of larger firms. Our final dataset includes 1,090 such grid cells. Figure A.12

[31] In OLS models, inference using predicted regressors should be corrected for first-stage sampling variance. When the predicted regressor is used as an instrumental variable, like we do here, the standard errors of the 2SLS regression are unbiased under a set of weak assumptions (Wooldridge, 2010). Predicted regressors have similarly been used as instruments in Schlenker and Walker (2016) and Dahl and Lochner (2012).

shows that the point estimates remain quite precise when we cluster the standard errors at the firm level, at the county level and two-way at the Copernicus grid cell and time level.[32]

4.4 Identification assumptions

Our research design requires three conditions. First, our wind instruments must be correlated with $PM_{2.5}$ (instrument relevance). Second, they must be uncorrelated with the error term from the second stage, ϵ_{fiyt} (instrument validity). Third, with plausibly heterogeneous treatment effects, the two-stage least squares estimates can be interpreted as local average treatment effects (LATE) only if the monotonicity assumption holds.

Instrument relevance. Table 2 reports the first-stage results and shows that the effects of the wind instruments on pollution exposure are similar whether the regression is estimated at the municipality level (column 1) or at the firm level for single-establishment firms (column 2). The estimated coefficients $\widehat{\beta^j}$ are all positive because Z^j_{gyt} is constructed to take negative values when wind from direction j reduces pollution in municipality g. All coefficients are positive and highly significant.

To assess instrument strength, we use two complementary weak-IV diagnostics applied to the sample of single-establishment firms. First, we compute the Montiel Olea and Pflueger (2013) effective F-statistic, appealing in a setting with many instruments and heteroskedasticity (Andrews et al., 2019). Given our large sample, high-dimensional fixed effects, and the constraint of running the analysis on a secure server, we compute it based on a random 2% sample of single-establishment firms.[33] The effective F-statistic (reported in Table 3) equals 365, which exceeds the 5% and 10% worst-case bias critical values (26 and 16, respectively), allowing us to rule out weak instruments. Because the effective F-statistic does not accommodate multiple endogenous regressors and is computationally costly in large samples, we also report the Kleibergen–Paap Wald rk (K-P) F-statistic for the full sample of single-establishment firms, although it might be less reliable when the instrument set is large or heterogeneous. The K-P F-stat is 167. Overall, both diagnostics confirm the strength of our wind-based instruments.

Instrument validity. The validity of the instruments relies on two assumptions. First, the wind direction instruments must be as-good-as-randomly assigned, meaning no weather or seasonal patterns influencing sales should co-vary with the instruments. To address this, we control for wind speed, temperature, and precipitation, which may correlate with wind direction and affect sales, and include quarter-by-county fixed effects to account for location-specific seasonality and quarter-specific wind and sales patterns. The remaining variation in the instruments is assumed to be random, as no other weather variables are known to influence both sales and the instruments.

[32] Implementing Conley standard errors to address concerns of spatial autocorrelation in wind patterns is too computationally demanding due to the combination of high-dimensional fixed effects, instrumental variables, and a large sample size, given the constraint of accessing the data on a secure server.

[33] Since the effective F-statistic is defined for one endogenous regressor only, we instrument for pollution at $t-1$—our main parameter of interest—and control for the wind instruments at t and $t+1$.

Table 2: First stage results

	$PM_{2.5}$ exposure	
	Municipality aggregation (1)	Firm aggregation (2)
$Z_{\text{South}\,gyt}$	1.432***	1.468***
	(0.097)	(0.152)
$Z_{\text{West}\,gyt}$	0.529***	0.575***
	(0.0635)	(0.148)
$Z_{\text{North}\,gyt}$	1.112***	1.231***
	(0.0484)	(0.055)
$Z_{\text{East}\,gyt}$	1.645***	1.610***
	(0.0481)	(0.0748)
Municipality-by-year FE	Yes	No
Firm-by-year FE	No	Yes
Month-by-year FE	Yes	No
Month-by-year-by-industry FE	No	Yes
N	391,234	6,322,128
R-squared	0.93	0.93

Notes: Table reports the first stage results based on equation (10), using municipality-month data (column (1)), and single establishment firm-month data (column (2)). Holiday and weather controls and quarter by-county-fixed effects are included. We report standard errors in parentheses, clustered at the Copernicus grid cell. We denote * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Second, the exclusion restriction must hold: the wind instruments should affect firms' sales only through their impact on $PM_{2.5}$. This assumption could be violated if other pollutants affecting health and productivity co-vary with wind direction. Among the four other regulated air pollutants (SO_2, NO_2, PM_{10}, and ozone), SO_2 and NO_2 are primary pollutants that convert to particulate matter within two to three days. With monthly data, we capture their effects as part of $PM_{2.5}$. PM_{10} is highly correlated with $PM_{2.5}$ ($\rho = 0.93$) and includes $PM_{2.5}$, so our estimates also reflect PM_{10}'s impact. Ozone, however, is typically anti-correlated with these pollutants due to its atmospheric formation process.[34] In our data, $PM_{2.5}$ and ozone have a Pearson correlation coefficient of -0.3. To address concerns about ozone effects, in section 5.3 we conduct a robustness check where we replace $PM_{2.5}$ with a multi-pollutant air quality index that includes ozone, which does not alter the results.

Instrument monotonicity. We test for instrument monotonicity by plotting the relationship between residualized instruments and residualized $PM_{2.5}$ exposure. Figure A.9 presents binned scatter plots of these variables using the panel of single-establishment firms, showing a predominantly monotonically increasing and approximately linear relationship, except at the distribution tails. Figure A.10 displays the distribution of residualized predicted $PM_{2.5}$ and its relationship with residualized firm-level $PM_{2.5}$ exposure, confirming that the monotonicity assumption holds for this instrument.

Spillovers as potential threat to identification. Our identification relies on comparing firm-

[34]Ozone forms through reactions involving solar radiation, nitrogen oxide, and volatile organic compounds (Nasa Earth Observatory, 2003). Figures A.2 and A.4 illustrate this anti-correlation, showing reverse seasonality between ozone and $PM_{2.5}$ or NO_2.

months exposed to plausibly exogenous air pollution shocks driven by wind direction changes with those less exposed, under the assumption of stable unit treatment values (SUTVA), meaning no spillovers between exposed and non-exposed firms. While spillovers cannot be ruled out *a priori*, such as competitors gaining market shares from firms experiencing sales declines due to pollution, the low saliency and temporary nature of monthly air pollution exposure make it unlikely that firms adjust to competitors' shocks on a month-to-month basis. The high frequency of shocks reduces the likelihood of spillovers, and any effects occurring over a longer horizon are absorbed by firm-year fixed effects. Moreover, firms serving the same local demand—e.g., in business-to-consumer sectors—typically face the same pollution shocks owing to geographic proximity, which limits competitive advantages. Firms experiencing lower pollution exposure within the same industries are likely geographically distant, reducing direct competition and further minimizing spillover risks.

5 Main Results

5.1 Impact of $PM_{2.5}$ on Firms' Sales

All Sectors. Table 3, panel A, shows that firms' sales decline in response to contemporaneous and lagged monthly $PM_{2.5}$ exposure when pollution exposure is instrumented. By contrast, column (1) shows a positive (non-significant) association between $PM_{2.5}$ and sales in the OLS specification, likely driven by reverse causality and omitted variables. Within a firm-year, even after accounting for industry-specific time-varying shocks and local seasonality in sales and pollution, months with positive local economic shocks tend to be more polluted.[35] When instrumenting pollution with changes in wind direction (column 2), the effect of lagged $PM_{2.5}$ exposure on sales becomes negative and statistically significant at the 1% level. A one-unit (1 µg/m^3) increase in lagged firm-level $PM_{2.5}$ reduces sales over the following two months by 0.45 percent ($p < 0.001$). Given a baseline average pollution of 15.3 µg/m^3, the corresponding elasticity is -0.069, implying that a 10 percent increase in pollution exposure lowers sales by 0.69 percent on average. The elasticity for contemporaneous exposure is smaller, at -0.023 ($p = 0.025$), indicating that the sales response to pollution is more strongly driven by lagged rather than current exposure. Columns (3) and (4) confirm similar results in the subsample of single-establishment firms.

Table A.2 shows how magnitudes vary with fixed effects. The IV point estimates remain consistently significant and negative across specifications. The elasticity of sales to lagged exposure ranges from -0.089 to -0.069 (in our preferred specification), with smaller magnitudes when controlling for quarter-by-county fixed effects in addition to firm-by-year and month-by-year-by-industry fixed effects. Quarter-by-county fixed effects account for any systematic correlation between wind seasonality and local economic activity, which, if not controlled for, could violate the exclusion restriction.

[35] In the Chinese context, Fu et al. (2021) similarly finds a positive association between $PM_{2.5}$ and manufacturing value added per worker in their OLS regression, and a negative effect in the IV regression.

Our results can be compared to the ones obtained for Europe by Dechezleprêtre et al. (2020), which reports that a 1 $\mu g/m^3$ increase in $PM_{2.5}$ lowers regional GDP (NUTS3) by 0.80 percent, using thermal inversions as an instrument and only controlling for country-year fixed effects. While the latter effect is larger in magnitude, the difference could be due to the level of aggregation, outcome definition (including agriculture production), and empirical strategy, with fewer sources of time-varying unobserved heterogeneity being absorbed in their study. Taken together, the evidence consistently points to sizable economic costs of $PM_{2.5}$.

Heterogeneous Response by Sector. Examining how sales respond to air pollution by sector provides insights into which sectors might benefit most from air quality improvements. Panels B to E of table 3 show that, after instrumenting for $PM_{2.5}$ exposure, firm sales decrease with lagged exposure in all sectors, with sales in construction and business-to-consumer retail and services also responding to contemporaneous exposure. Results are broadly consistent for the subsample of single-establishment firms (column 4).

Expressed as elasticities, these results indicate that a 10 percent increase in pollution exposure at $t-1$ decreases sales in the following two months by 0.28 percent in manufacturing, 0.36 percent in construction (0.24 for exposure at t), 0.44 percent in business-to-business trade and services, and 1.17 percent in business-to-consumer retail and services (0.35 for exposure at t). Fu et al. (2021) found a higher elasticity (-0.44) for manufacturing firms that could be due to higher levels of pollution in China, to their less granular data (annual level, county pollution exposure) or due to a different outcome variable (value added per worker) and a different instrument (thermal inversions).

Why do magnitudes differ across sectors? As discussed in section 2.2, both supply-side and demand-side mechanisms likely influence sectoral responses differently. On the supply side, workers in different sectors may have varying vulnerabilities to pollution shocks. Although workers in sectors like construction and manufacturing face more direct and cumulative exposures, they also self-select into these occupations and could have adapted to higher levels of pollution. We explore the heterogeneity of workers' absenteeism response to air pollution shocks by sector in the next section.

On the demand side, business customers and final consumers are exposed to the same pollution shocks as workers only in industries serving local demand. This is the case predominantly for business-to-consumer retail and services—the sector incurring the largest sales losses. In the next section, we shed light on a potential demand channel that affects purchases of different types of consumer goods differently. By contrast, manufacturing firms typically serve distant markets, which may counteract the negative impact of local demand shocks (Almunia et al., 2021).

The larger sales elasticities to pollution observed in business-to-consumer retail and services can also be understood through the lens of the theoretical framework in Section 2.2. Sectors characterized by higher elasticities of substitution across varieties or lower profit margins are more sensitive to both supply- and demand-side disruptions. Given that retail typically operates on thin profit margins, this amplifying mechanism likely contributes to the heightened responsiveness in the consumer-oriented sector.

Table 3: The effect of $PM_{2.5}$ on firms' sales, overall and by sector

	All firms		Single-establishment firms	
	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)
Panel A: All Sectors				
PM_{t-1}	0.00766	-0.448***	0.0250	-0.490***
	(0.0207)	(0.108)	(0.0231)	(0.110)
PM_t	0.0813***	-0.151**	0.0949***	-0.116*
	(0.0254)	(0.0675)	(0.0279)	(0.0694)
Effective F-stat of first stage				365
K-P F-stat of first stage				167
N	9,411,781	9,411,781	6,072,012	6,072,012
R-squared	0.9469	0.9469	0.9338	0.9338
Panel B: Manufacturing				
PM_{t-1}	0.00665	-0.180**	-0.0101	-0.149*
	(0.0210)	(0.0717)	(0.0247)	(0.0779)
PM_t	0.0588***	-0.0527	0.0475**	0.00173
	(0.0201)	(0.0528)	(0.0221)	(0.0556)
K-P F-stat of first stage				150
N	1,880,380	1,880,380	1,233,990	1,233,990
R-squared	0.9641	0.9641	0.9535	0.9535
Panel C: Construction				
PM_{t-1}	-0.0201	-0.232***	-0.0226	-0.283***
	(0.0219)	(0.0633)	(0.0248)	(0.0685)
PM_t	-0.0174	-0.158***	-0.0231	-0.132*
	(0.0216)	(0.0584)	(0.0240)	(0.0679)
K-P F-stat of first stage				131
N	1,531,685	1,531,685	1,074,583	1,074,583
R-squared	0.9353	0.9353	0.9165	0.91625
Panel D: Business-to-Business Trade and Services				
PM_{t-1}	-0.0443**	-0.288***	-0.0146	-0.258***
	(0.0217)	(0.0726)	(0.0273)	(0.0858)
PM_t	0.0491**	-0.0195	0.0505*	-0.00856
	(0.0214)	(0.0573)	(0.0266)	(0.0651)
K-P F-stat of first stage				104
N	2,875,213	2,875,213	1,498,367	1,498,367
R-squared	0.9338	0.9338	0.9155	0.9155
Panel E: Business-to-Consumer Retail and Services				
PM_{t-1}	0.0840**	-0.757***	0.113***	-0.788***
	(0.0404)	(0.196)	(0.0419)	(0.186)
PM_t	0.190***	-0.225*	0.226***	-0.113
	(0.0462)	(0.128)	(0.0488)	(0.136)
K-P F-stat of first stage				108
N	3,124,500	3,124,500	2,265,070	2,265,070
R-squared	0.9457	0.9457	0.9342	0.9342

Notes: Table reports the OLS and IV estimates of the effect of a one unit increase in $PM_{2.5}$ at $t-1$ and t on the sales outcome at t based on equation (7), for all firms in columns (1) and (2) and for single-establishment firms in columns (3) and (4), overall and by sector. All regressions include firm-by-year fixed effects, month-by-year-by-industry fixed effects and quarter-by-county fixed effects, as well as weather and holidays controls at $t-1$, t and $t+1$ and (instrumented) pollution at $t+1$. The instruments are either the predicted firm-level pollution measure (column 2) or the 4 wind directions (column 4). The effective F-stat reported in column (4) of Panel A was computed based on a 2% random sample of single establishment firms, due to computational constraints. Standard errors are clustered at the Copernicus grid cell level of the firm's headquarter. Coefficients (and standard errors) have been multiplied by 100 for readability. We denote * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Finally, sectoral differences in the timing of production and sales recording may explain some of the heterogeneity. For instance, business-to-business services often involve payment delays, meaning that the timing of sales responses may not align with the two-month period used in our analysis. We investigate dynamic effects up to five month following the pollution shock in the next section.

Heterogeneity by firm size. Smaller firms are generally found to be less productive and more vulnerable to financial shocks than larger firms (Gertler and Gilchrist, 1994). We investigate this heterogeneity by comparing the impact of air pollution on sales for firms below the median size (fewer than 15 employees) versus for those above the median size (with 15 employees or more). As shown in Table A.4, smaller firms are more adversely affected overall and across all sectors, except in business-to-consumer retail and services. In manufacturing and construction, larger firms exhibit no significant sales losses from lagged pollution exposure, but losses are statistically significant for contemporaneous pollution in construction. In these sectors where supply-side mechanisms are more likely—because demand is primarily non-local— large firms may be able to offset productivity losses or absenteeism through adaptation strategies such as task reallocation or flexible scheduling (Adhvaryu et al., 2022). In contrast, the fact that even large firms appear vulnerable in the business-to-consumer sector suggests that demand-side mechanisms, less amenable to firm-level mitigation, may be the dominant channel in this context.

5.2 Dynamic Effects on Sales

Figure 2 plots the dynamic effects of pollution on sales from the month of exposure (t) through five months afterward ($t+5$), estimated using the polynomial distributed lag specification for the full sample (panel a), single-establishment firms (panel b), and all firms in each sector (panels c-f). In the full sample, sales decline in the month of exposure, reach their lowest point one month later, and gradually return to zero effect after about 4 months; the pattern for single-establishment firms is nearly identical.

In terms of magnitude, the first two estimates in panel (a) imply elasticities of -0.035 for contemporaneous exposure and -0.046 for one-month lagged exposure, which are close to the corresponding elasticities of -0.023 and -0.069 in Table 3. By imposing a smooth polynomial form on the lag structure, the PDL specification produces a slightly larger contemporaneous effect and a somewhat smaller lagged effect. The largest effects occur at $t+1$ in nearly all panels (in manufacturing, the estimates at $t+1$ and $t+2$ have almost the same magnitude), justifying our emphasis on pollution exposure at t and $t-1$ in our main specification.

The close agreement between the one-lag distributed lag model and the PDL specification is reassuring. The PDL helps mitigate serial correlation in pollution (and wind) exposure and imposes structure on the lag profile, but it may also allocate more weight to peak lags, inflating magnitudes or introducing misspecification. For these reasons, we rely on the simpler specification with just one lag for our main results. Because this approach deliberately restricts attention to effects within three months of exposure—thereby omitting potentially non-negligible longer-run responses—it yields conservative estimates of the dynamic impact of pollution on sales. As shown in Table A.3,

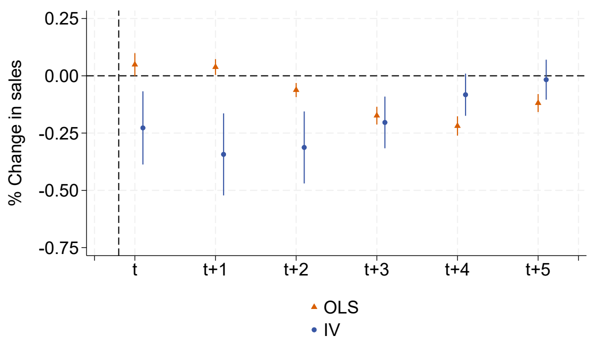

(a) All firms

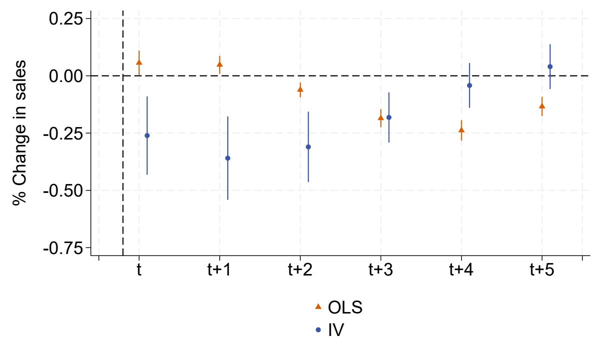

(b) All single-establishment firms

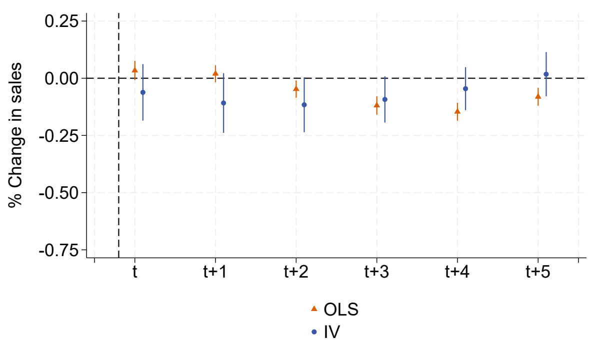

(c) Manufacturing

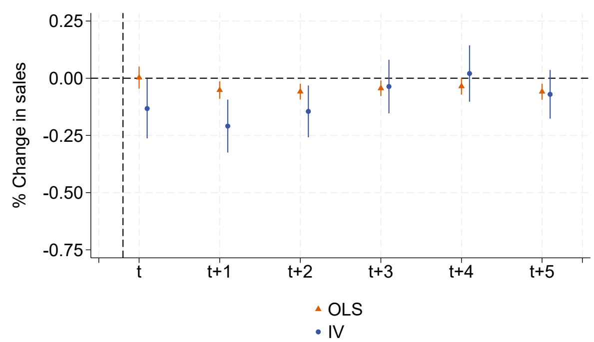

(d) Construction

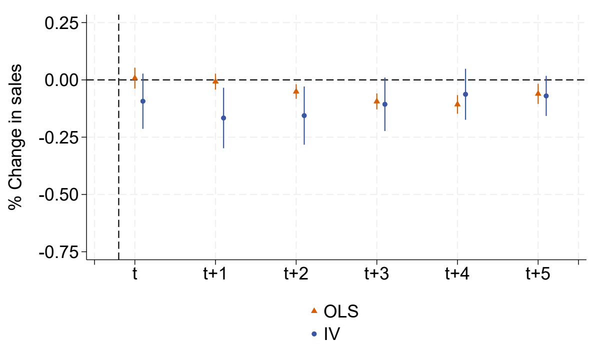

(e) Business-to-Business Trade & Services

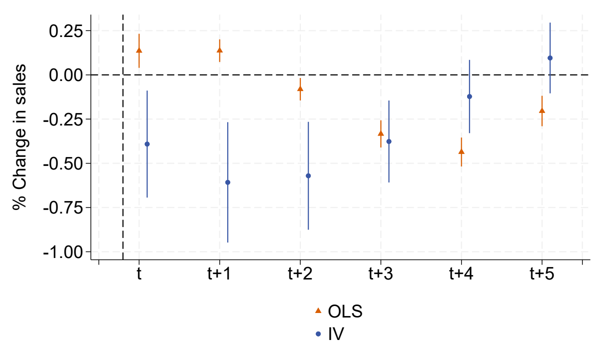

(f) Business-to-Consumer Retail & Services

Figure 2: Dynamic effects of $PM_{2.5}$ on sales for all firms, overall and by sector

Notes: Figure shows the OLS and IV point estimates and 95% confidence intervals for the effect of contemporaneous and lagged $PM_{2.5}$ (up to $t-5$) on firms' sales at t using a polynomial distributed lag specification based on equation (7) for all firms (a) and for single-establishment firms (b), as well as by sector: manufacturing (c), construction (d), B-to-B (e) and B-to-C (f). All regressions include month-by-year-by-industry fixed effects, firm-by-year fixed effects, quarter-by-county fixed effects, and weather and holidays controls from $t-1$ to $t+1$. The confidence intervals are based on standard errors clustered at the Copernicus grid cell level of the firm's headquarter. Coefficients (and standard errors) have been multiplied by 100 for readability.

the estimates at t and $t-1$ are not sensitive to the number of lags introduced in equation (7).

In all panels from Figure 2, the OLS estimates are either at zero or slightly positive in the first few months after exposure and become negative afterward (after $t+3$). By contrast, the IV estimates follow similar patterns for the full sample and in each sector, with the same differences in magnitude across sectors as the ones reported in Table 3. While manufacturing and business-oriented sectors respond with a lag of one or two months, construction and consumer-oriented sectors experience a sales decline in the month of exposure and recover within three months for construction and only after four months for the consumer-oriented sector.

The delayed onset of sales declines following pollution exposure can be explained by three mechanisms. First, in manufacturing and, to some extent, in business-to-business trade, there is a natural lag between production—when workers are affected—and the recording of sales upon delivery. Second, in construction and services, where sales are recorded at the time of payment, delays between service delivery and payment can lengthen the lag. Third, in consumer-oriented sectors, demand responses may materialize gradually: income losses from uncompensated sick leave typically become salient only at the end of the month when wages are paid and some consumers may temporarily smooth consumption by drawing on savings despite healthcare spending and reduced income. The absence of any rebound even six months after a pollution shock indicates that sales are not simply postponed. Instead, the decline appears persistent, consistent with a reduction in disposable income driven by both increased uncompensated absenteeism and higher healthcare expenditures.

5.3 Robustness checks

We assess the validity of the identification assumptions and the robustness of our main results in various ways. First, we run a falsification test using future pollution exposure to rule out that our effect is driven by spurious correlation. Table A.5 shows that future pollution exposure, at time $t+2$, has small and insignificant effect on sales at time t, for all sectors taken together and for each sector.

Second, we consider the risk of violation of the exclusion restriction due to ozone pollution and adopt a broader measure of air pollution using the Air Quality Index (AQI). Column (1) of Table 4 reports the main result, replicating the primary specification reported in column (2) of Table 3, panel A. Column (2) presents the effect of the AQI, instrumented by the same four wind instruments. The magnitude of the coefficient cannot be directly compared to our main $PM_{2.5}$ estimate due to differences in scale. However, expressed in standard deviations, the results are similar: a 1-SD increase in lagged $PM_{2.5}$ (SD=6.2 µg/m^3) reduces sales by 2.8%, while a 1-SD increase in AQI (SD=0.41) reduces sales by 2.5%. The slightly lower AQI estimate may reflect its dependence on $PM_{2.5}$ in fall/winter and on ozone in spring/summer. If ozone has no significant effect on sales, using AQI dilutes the impact.

Third, we check that our results are not driven by air quality alerts and the avoidance behaviors that they may induce. To do so, we replicate the analysis on a sample excluding months with PM_{10}

alerts. While no $PM_{2.5}$ alerts exist in France, PM_{10} alerts—triggered by daily PM_{10} concentrations exceeding regulatory thresholds—are highly correlated with high $PM_{2.5}$ levels.[36] Column (3) of Table 4 shows that the estimated coefficient remains consistent with the main result.

Table 4: Robustness checks for the effect of $PM_{2.5}$ on firm-level sales

	Baseline	AQI	No AQ alerts	Winsorized outcome	Quadratic weather controls	Weather incl. humidity	Flu incidence control
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
PM_{t-1}	-0.448***		-0.459***	-0.431***	-0.677***	-0.396***	-0.453***
	(0.108)		(0.111)	(0.114)	(0.148)	(0.111)	(0.108)
PM_t	-0.151**		-0.166**	-0.155**	-0.344***	-0.170**	-0.150**
	(0.0675)		(0.0796)	(0.0721)	(0.0920)	(0.0830)	(0.0685)
AQI_{t-1}		-6.14***					
		(1.81)					
AQI_t		-2.08*					
		(1.15)					
N	9,411,781	9,411,781	9,048,811	9,460,260	9,411,787	9,411,764	9,411,781

Notes: Table reports the IV estimates of the effect of a one-unit increase in $PM_{2.5}$ (or AQI) at $t-1$ and t on the sales outcome at t based on equation (7) for all firms. All regressions include weather and holidays controls at $t-1$, t and $t+1$, instrumented pollution at $t+1$, firm-by-year fixed effects, quarter-by-county fixed effects and industry-by-month-by year fixed effects. Standard errors are clustered at the Copernicus grid cell level of the firm's headquarter. Coefficients (and standard errors) have been multiplied by 100 for readability. We denote * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Fourth, we test the sensitivity of our results to outliers and to the specification of time-varying controls. Column (4) of Table 4 shows that winsorizing sales at the 2nd and 98th percentiles of the monthly sales distribution does not affect our main estimate. In column (5), we control for weather using a quadratic function for average temperature, wind speed, and rainfall, instead of our flexible set of controls. The larger estimated coefficient on pollution suggests that our main estimate is conservative. In column (6), we add monthly averages of daily maximum humidity to the set of flexible weather controls —splitting the variable into quintiles and generating indicators for all the possible interactions with other weather categories.[37] In column (7), we control for county-level flu cases per 100,000 inhabitants to account for seasonal illnesses that affect absenteeism and economic activity and may correlate with wind patterns. Because air pollution exacerbates influenza (Graff Zivin et al., 2023), flu incidence is likely a bad control, which is why it is omitted from the main specification. Our results remain similar when including either humidity or flu controls.

Fifth, we check that our results are not driven by the design of our wind instruments. Table A.6, Panel A shows the results for all firms. Panel B shows the results for single-establishment firms

[36] The threshold is 80 µg/m^3 before November 2014 and 50 µg/m^3 after. Alerts give rise to awareness messaging advising against physical activity for vulnerable populations. Even in Paris, the most polluted city, alerts occurred on only 4% of the days in 2009. The more severe alert level, giving rise to driving restrictions, occurred only 0.7% of the days.

[37] We define humidity as the dewpoint temperature in °C, which corresponds to the temperature that the air needs to be cooled to, at constant pressure, in order to achieve a relative humidity (RH) of 100%.

only—the sample on which the Kleibergen-Paap first stage F-statistic is computed. The first alternative instrument computes component B using data for 2009 only to alleviate concerns of potential endogeneity of this component with respect to economic activity in later periods. The results show comparable magnitudes to the baseline for a sample period starting in 2009 (column 2) or only in 2010 (column 5). Second, we only exploit component A for the wind instruments, while constraining the first stage coefficients to be the same for all municipalities within a 100km-by-100km grid cell, in an approach similar to Deryugina et al. (2019). Results in column (3) show that for all firms (Panel A), the negative sign and statistical significance are preserved, with a substantially smaller magnitude reflecting both the loss of identifying variation when shares are removed and a different LATE. For single-establishment firms (Panel B), the point estimate shrinks and falls below conventional thresholds; the instrument remains strong (KP F = 231), pointing to a loss of remaining identifying variation under this smaller sample, rather than a weak-instrument problem.

Sixth, we assess the robustness of our results to alternative pollution measures. Column (2) of Table A.7 replicates our main analysis using satellite-based monthly $PM_{2.5}$ data from van Donkelaar et al. (2021); Shen et al. (2024). These data rely on satellite-based Aerosol Optical Depth (AOD), combined with a chemical-transport model to map AOD to $PM_{2.5}$, and are cross-validated using ground monitoring stations. Satellite-based exposure is highly correlated with our main exposure measure ($\rho = 0.90$), and the estimated coefficient remains of similar magnitude to the baseline in column (1).

Columns (3)-(5) of Table A.7 compare the reanalysis-based estimates with $PM_{2.5}$ data from monitoring stations for 2011-2015.[38] Using monitor data helps rule out the concern that the first stage linking wind directions to $PM_{2.5}$ is driven by weather inputs embedded in the reanalysis model. In column (4), municipal exposure is constructed as an inverse-distance–weighted average of nearby monitors within 150 km, following standard practice. Column (5) instead uses only the nearest monitor.[39] Both monitor-based measures are highly correlated with the reanalysis data ($\rho = 0.95$) and yield point estimates closely aligned with the reanalysis-based estimate for the 2011-2015 period, reported in column (3).

Finally, we examine how our results' precision varies with the clustering level in Figure A.12. The top estimate shows our baseline with one-way clustering at the Copernicus grid cell of the firm's headquarter—the spatial scale at which the instrument varies for single-establishment firms. Clustering at the firm level (second line), the scale relevant for multi-establishment firms, yields smaller standard errors because it ignores spatial correlation in wind exposure across nearby firms. Clustering at the county level (third line), and thus indirectly capturing broader spatial correlation—the average county includes 10 wind grid cells—produces slightly smaller standard errors than the baseline, which remains the most conservative. The final three estimates apply two-way clustering by space and time, using the Copernicus grid cell (fourth line), firm (fifth) or county (sixth) as

[38] Monitoring stations data cover the period 2011-2015 and are available at: https://eeadmz1-downloads-webapp.azurewebsites.net/.

[39] In our sample, the average distance to the nearest monitor is 25 km and the median is 16 km.

the spatial dimension and month-by-year as the temporal dimension. This accounts for correlation across observations within the same month. While these specifications reduce precision slightly, the effect remains significant at the 5% level.

6 Mechanisms

The temporary decline in sales following a month of high $PM_{2.5}$ may stem from several mechanisms identified in our analytical framework, namely reduced labor supply, reduced worker productivity, and lower demand. In this section, we explore each of these potential channels.

6.1 Sickness-induced absenteeism

Table B.1, panel A, reports the main OLS and IV estimates of the contemporaneous effect of $PM_{2.5}$ on sick leave using equation (8), for the sample of workers whose firm is included in our sales data. The OLS estimate in column (1) shows that a one-unit increase in average $PM_{2.5}$ exposure is associated with a 0.07 increase in sick leave per 1,000 workers. The IV estimate (reproduced in Figure 3) in column (2) is twice as large, at 0.15, suggesting that the OLS estimate is downward biased due to omitted variables and classical measurement error. Both estimates are statistically significant at least at the 5% level. With a baseline average of 23 per 1,000 workers, our IV results imply that a 10 percent increase in monthly $PM_{2.5}$ raises absenteeism by 1 percent, corresponding to a 0.1 elasticity of sick leave to pollution. This elasticity is similar to the literature, despite differences in the type of pollutant and time horizon: an analysis based on weekly data from Spain, focusing on PM_{10} pollution in urban areas, finds an increase in sick leave corresponding to an elasticity of 0.08 (Holub et al., 2021). Figure A.14 shows the dynamic effects of pollution using the same PDL specification as for sales: the impact of air pollution on sick leave is concentrated in the month of exposure, quickly dissipating to zero within two months.

Can the pollution-induced reduction in labor supply due to sick leave explain most of the observed decline in sales? Applied to the sample-mean monthly sales of €1,316,000, our estimated sales elasticity to lagged pollution implies that a 10% increase in $PM_{2.5}$ lowers the average firm's monthly sales by approximately €9,080. We benchmark this figure against the sales loss implied by increased absenteeism alone. The baseline absenteeism rate due to sick leave is 0.023 in our sample, and the average firm employs 60 workers, of whom 58.6 are on the job in a typical month ($(60 \times (1 - 0.023) = 58.6$). A 10% increase in $PM_{2.5}$ generates 0.0138 additional absent workers in such a firm; valued at monthly sales of €22,450 per on-the-job worker (€1,316,000/58.6), the implied loss is $0.0138 \times$€22,450$\approx$€310, or about 4% of the €9,080 total. The remaining 96% must therefore operate through channels other than our measured channel of sick leave-driven absenteeism.

Several assumptions underlie this calculation, which could make the 4% figure either an upper or a lower bound for the absenteeism channel. First, the calculation attributes a full month's output loss to every worker beginning a sick-leave episode during the month, while the average episode shorter than three months lasts 16 days in our sample. Second, it rests on a linear approximation that values each additional absence at the firm's average sales per on-the-job worker.[40] Since observed sales already absorb the effect of the sample-average rate of absenteeism, our exercise captures the marginal impact of pollution-induced absences around this baseline. A final assumption is that the productive disruption caused by a pollution-induced absence is comparable to that of an absence from other causes.

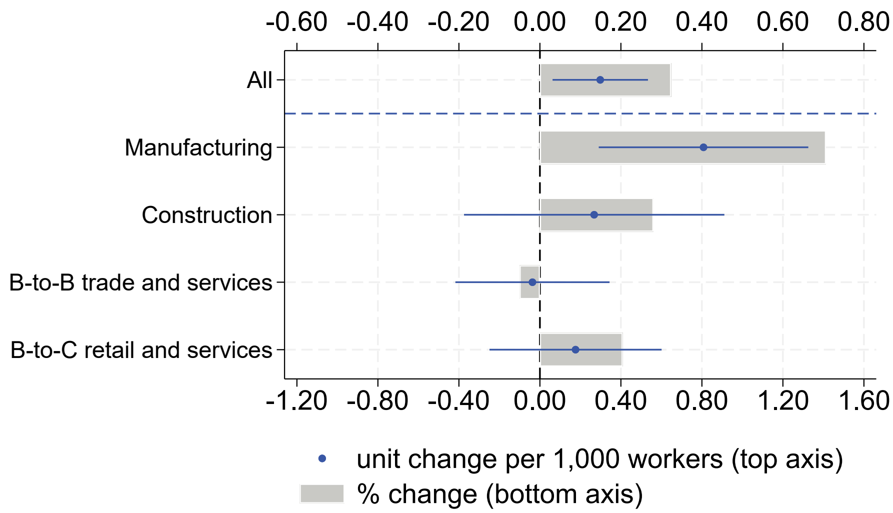

Figure 3: Contemporaneous effect of a one-unit increase in $PM_{2.5}$ on sick leave episodes

Notes: Figure shows IV point estimates and 95% confidence intervals for the effect of $PM_{2.5t}$ on the number of workers starting a sick leave per 1,000 workers at the establishment level, overall and by sector, based on equation (8). All regressions include industry-by-month-of-sample, establishment, and quarter-by-county fixed effects, as well as weather and holidays controls. Observations are weighted by the number of workers in each establishment. Standard errors are clustered at the Copernicus grid cell level of the establishment.

Figure 3 reports heterogeneous effects by sector and reveals that pollution-induced sick leave is mainly driven by manufacturing (the only sector with a statistically significant effect) and, to a lesser extent, construction. In contrast, the sick leave response is minimal in the consumer- and business-oriented trade and service sectors. Higher absenteeism rates are thus observed in sectors with greater pollution exposure at work and/or strong collective agreements (ensuring higher wage replacement rates), such as manufacturing and construction, whereas services sectors experience lower absenteeism.[41] Because our measure of absenteeism only captures recorded sick leave episodes, these lower responses in service sectors are also possibly due to the ability to work remotely or take leave without a medical certificate, or reflect higher incentives to work while sick due to lower replacement rates.

[40]Formally, let $y(a)$ denote a firm's monthly sales as a function of its absenteeism rate a. The exercise is a linear approximation of y around the sample mean $\bar{a} = 0.023$: $\Delta y \approx y'(\bar{a})\,\Delta a$, with $y'(\bar{a})$ approximated by $-\bar{y}/(1-\bar{a})$. The neglected second-order term captures the curvature in $y(a)$: concavity if congestion makes additional absences increasingly costly, convexity if substitution by the remaining workforce makes them less so. Quantitatively, however, this correction is negligible: since the perturbation $\Delta a/\bar{a} \approx 0.01$ is small relative to baseline, the second-order term amounts to less than one percent of the first-order term for any plausibly bounded curvature-to-slope ratio $|y''(a)/y'(a)|$.

[41]Pollak (2015) shows that employees in retail trade and construction are less likely to receive full wage replacement during sick leave than those in manufacturing and other service sectors.

There is little correspondence between sector-level effects on sick leave and on sales: comparing Table 3 and Figure 3, manufacturing and business-to-business trade and services exhibit sales elasticities of –0.028 and –0.044, respectively, yet their absenteeism elasticities move in the opposite direction (0.22 for manufacturing versus -0.015 for business-to-business trade and service sectors). Additionally, the business-to-consumer trade and services sector experiences large pollution-induced sales losses but small pollution-induced absenteeism, indicating that absenteeism is not an important channel in that sector.

With the same calculation and assumptions as above, we quantify the share of pollution-induced sales losses that can be attributed to increased absenteeism in the most affected sector, manufacturing. For this sector, a 10% increase in $PM_{2.5}$ reduces the average firm's sales by approximately €9,336 per month. Of this total, about €2,096 can be explained by increased absenteeism due to pollution. This represents roughly 22% of the pollution-induced sales losses. This back-of-the envelope calculation highlights sickness-driven absenteeism as a significant transmission channel through which air pollution affects firm sales, but not the main one. Other channels such as lower productivity and demand-side effects likely play an important role in driving the economic costs of pollution shocks.

Robustness checks Table B.1, panel B, shows that the effect of pollution on sick leave is robust to aggregating at the municipality rather than the establishment level, using municipality fixed effects and month-by-year fixed effects. Table B.2 shows that it is not driven by a confounding effect of ozone, by air quality alerts, or by the specification of weather controls. Table B.3 shows that the results are largely insensitive to the source of pollution data. Figure A.15 shows that our estimates (left) are comparable to the OLS and IV estimates for the representative sample of workers (right).

6.2 The role of productivity and demand

Productivity. We first consider sales per worker as an economy-wide proxy for productivity. Figure 4 reports the estimated effects of a one-unit increase in quarterly $PM_{2.5}$ on contemporaneous and lagged sales and sales per worker, overall and by sector. The sample is restricted to single-establishment firms for which we observe quarterly employment data. Given the sampling of the underlying quarterly employment data — including all firms with more than 50 workers and only a random sample of those with fewer employees —the average firm size is roughly twice the size of the average single-establishment firm in our main analysis sample. We find that within every sector and in the pooled specification, the responses of sales and sales per worker are quantitatively close and statistically indistinguishable: a $1\,\mu g/m^3$ increase in $PM_{2.5}$ lowers contemporaneous sales by about 0.20% and contemporaneous sales per worker by about 0.25%, with lagged effects exhibiting the same pattern at a slightly smaller magnitude. The near co-movement of sales and sales per worker across all specifications confirms that the per-worker normalization absorbs essentially the same variation as sales itself and therefore cannot independently identify the productivity channel.

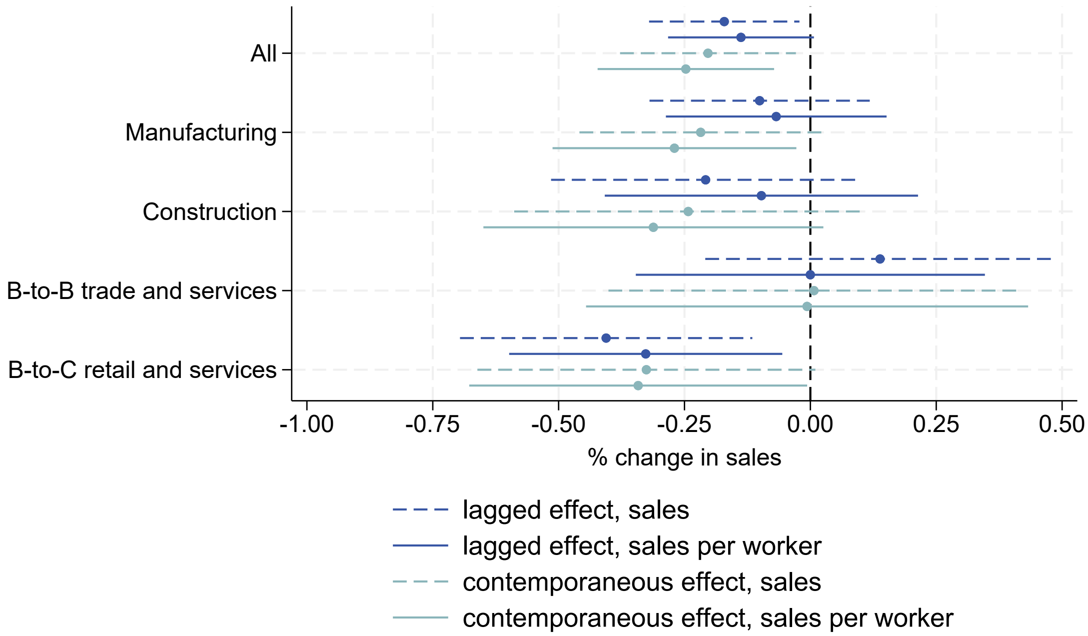

Figure 4: Effect of a one-unit increase in quarterly $PM_{2.5}$ on firms' lagged and contemporaneous quarterly sales and sales per worker

Notes: Figure shows IV point estimates and 95% confidence intervals for the effect of $PM_{2.5}$ at quarter $t-1$ (lagged effect) and t (contemporaneous effect) on firms' quarterly sales and sales per worker at t by sector. All regressions include quarter-by-year-by-industry fixed effects, firm-by-year fixed effects, weather and holidays controls at $t-1$ and t. The confidence intervals are based on standard errors clustered at the Copernicus grid cell level of the firm's headquarter.

Next, we provide suggestive evidence of the importance of the productivity channel by focusing on manufacturing. Our previous decomposition attributes 22% of the manufacturing sales decline to absenteeism, leaving 78% to productivity or demand channels. Since manufacturing firms typically serve national or international markets, the local demand channel should play a limited role, and the residual is therefore likely dominated by productivity losses. We test this conjecture by exploiting cross-industry heterogeneity in inventory.

Using a 2004 survey of 2,058 manufacturing establishments that records inventory levels in days of production, we classify industries as high- or low-stock relative to the median.[42] The logic of the test is as follows: high-stock industries can buffer transitory supply-side disruptions by drawing down inventories, whereas inventories provide no comparable cushion against demand shocks. A supply-driven mechanism therefore predicts a sharper sales response among low-stock industries, while a demand-driven mechanism would generate comparable responses across the two groups.

Columns (1)–(3) of Table 5 support the supply-side interpretation: based on our monthly data,

[42] Industries with high stock are: production of textile, clothing, shoes and leather; chemicals; pharmaceuticals; other non metallic mineral products; machine and equipment; transport equipment outside car industry; furniture; other manufacturing industry; repair and installation of machines. Industries with low stock are: food industry; production of beverages; tobacco products; wood products; paper; printing and recording industry; refineries; plastic and rubber; metal industry; other metal products; electronic, optic and IT equipment; electric equipment; car industry.

pollution sharply reduces sales among firms in low-stock industries, while the effect is statistically and economically negligible in high-stock industries. Columns (4)–(6) show that absenteeism responses are nearly identical across the two groups—if anything, slightly larger in high-stock industries—so the divergence in sales cannot reflect differences in worker availability. The two groups also display comparable average sales and employment, ruling out firm-size explanations. Together with the inability of inventories to buffer demand shocks, this pattern suggests that the residual sales decline in manufacturing operates mostly through on-the-job productivity losses, which firms partially attenuate by drawing on existing stocks.

Table 5: Productivity Channel: Heterogeneous effects of air pollution on sales and worker absenteeism in the manufacturing sector, by stock level

	Sales effect			Absenteeism effect		
	(1)	(2)	(3)	(4)	(5)	(6)
	All firms	Low stock	High stock	All firms	Low stock	High stock
PM_{t-1}	-0.180**	-0.314***	-0.00625			
	(0.0717)	(0.0971)	(0.103)			
PM_t	-0.0527	-0.132*	0.0673	0.313*	0.316	0.378*
	(0.0528)	(0.0691)	(0.0895)	(0.165)	(0.194)	(0.226)
Avg Nb. employees	90	83	96	90	83	96
Median Nb. employees	27	25	29	27	25	29
Avg. sales	2,315,972	2,160,235	2,368,296	2,315,972	2,160,235	2,368,296
N	1,880,380	1,151,685	629,076	1,351,931	865,271	486,658
R-squared	0.9640	0.9708	0.9530	0.1273	0.1279	0.1271

Notes: Columns 1-3 report the IV estimates of the effect of a one unit increase in $PM_{2.5}$ at $t-1$ and t on the sales outcome at t based on equation (7) for manufacturing firms. All regressions include weather and holidays controls at $t-1$, t and $t+1$, as well as instrumented pollution at $t+1$, and firm-by-year, quarter-by-county and industry-by-month-by-year fixed effects. Columns 4-6 report the IV estimates of the effect of a one unit increase in $PM_{2.5}$ at t on absenteeism at t, controlling for weather and holidays controls at t and industry-by-month-of-sample, establishment, and quarter-by-county fixed effects. Standard errors are clustered at the Copernicus grid cell level of the firm's headquarter (1-3) or of the establishment (4-6). Coefficients (and standard errors) have been multiplied by 100 for readability for the sales outcome. We denote * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Demand. The impact of $PM_{2.5}$ on sales is particularly pronounced in business-to-consumer trade and services (Table 3), and unlike in other sectors the effect is not attenuated among large firms (Table A.4). Because firm size buffers productivity shocks—through, for instance, the ability to hire temporary workers or reallocate tasks—but offers no protection against demand shocks, this asymmetry points to a demand-side channel: retail and consumer services serve local customers who are themselves exposed to the same pollution shocks as workers.[43] Consumers affected by pollution may face reduced disposable income—because sick leave is only partially compensated and because health shocks raise medical expenses—leading them to cut back more on discretionary

[43]Consumers in these sectors were predominantly local during our sample period: in 2012, online sales accounted for only 3.6% of French retail (versus 8% in the U.S.) and food delivery for 0.6% of restaurant sales (Source: INSEE). This share has likely grown as e-commerce expanded in later years.

purchases, such as clothing, than on essential goods like groceries or necessary repairs.

Figure 5 disaggregates the consumer-oriented sector by industry and reveals an ordering consistent with this mechanism. Sales declines are smallest in essential categories: vehicle and goods repairs show no significant response, and supermarkets and the health-and-beauty segment (pharmacies and veterinary services) display only modest reductions (our estimates imply a lagged supermarket elasticity of about -0.08). They are largest in discretionary categories, with clothing and restaurants exhibiting the steepest drops (the associated lagged restaurant elasticity reaches -0.19, while the contemporaneous effect is more muted). Intermediate categories—specialized food retail, furniture, and car dealers—fall between these poles. Although confidence intervals partially overlap, the consistent ordering from essential to discretionary categories supports the interpretation that consumer behavior, shaped by both health concerns and temporary financial constraints, materially mediates the economic impact of pollution.

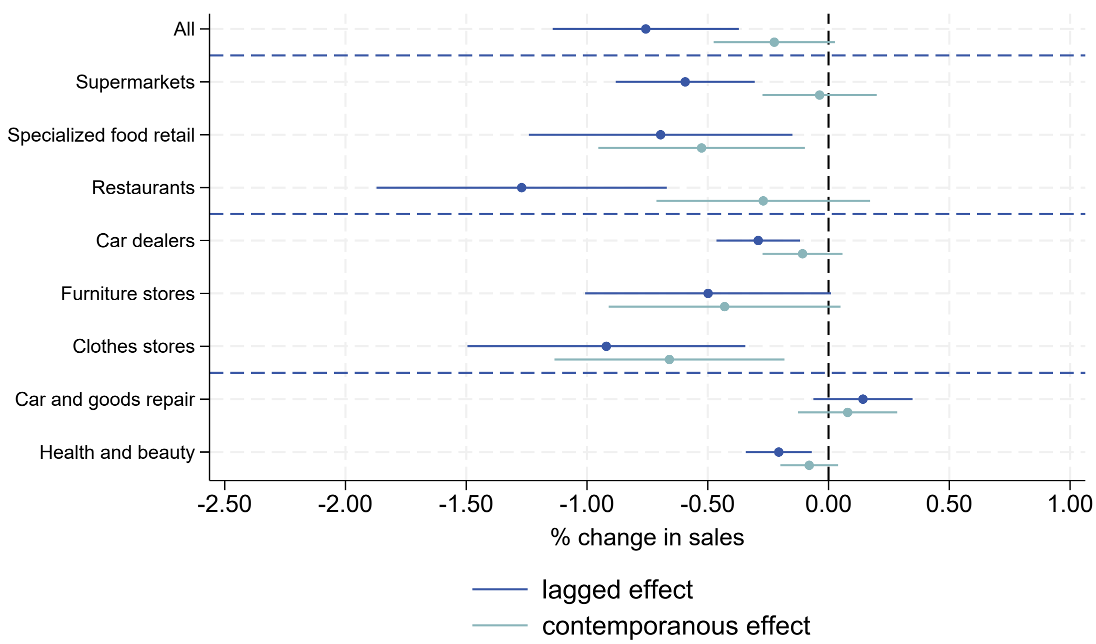

Figure 5: Effect of a one-unit increase in $PM_{2.5}$ on firms' sales by industry within the business-to-consumer retail and services sector

Notes: Figure shows IV point estimates and 95% confidence intervals for the effect of $PM_{2.5}$ at $t-1$ (lagged effect) and t (contemporaneous effect) on firms' sales at t by industry, based on equation (7). All regressions include month-by-year-by-industry fixed effects, firm-by-year fixed effects, quarter-by-county fixed effects, weather and holidays controls at $t-1$, t, and $t+1$, as well as instrumented pollution at $t+1$. The confidence intervals are based on standard errors clustered at the Copernicus grid cell level of the firm's headquarter.

7 Discussion on the overall effects of $PM_{2.5}$ reductions

This section quantifies the economic benefits to the French private sector of a policy capping daily $PM_{2.5}$ concentrations at the WHO target of 15 µg/m^3. The estimated benefits are substantial—

comparable in magnitude to upper-bound estimates of abatement costs and to the mortality benefits implied by the same policy.

Our counterfactual exercise combines the sector-specific sales elasticities from Table 3 with firm-specific measures of the implied reduction in pollution exposure, and proceeds in three steps. First, for each municipality, we compute the daily $PM_{2.5}$ reductions that would have occurred over our seven-year sample period if concentrations had been capped at 15 $\mu g/m^3$, and aggregate them to the monthly level; enforcing the cap would have lowered workers' average monthly exposure from 15.4 to 11.5 $\mu g/m^3$, a 25% reduction. Second, we translate municipality-level reductions into firm-month exposure changes, accounting for intrafirm network structure. Third, we multiply each firm-month exposure change by the relevant sectoral elasticity from Table 3—restricted to contemporaneous and one-month-lagged estimates significant at the 5% level—and by firm-month sales.[44]

Summing over all firm-months, we find that meeting the WHO target would have avoided €25 billion in annual lost sales—1.4% of average total sales in our sample.[45] Applying the average value-added-to-sales ratio of 27% (INSEE, 2015), this corresponds to €6.9 billion in annual foregone value added, abstracting from longer-term and general-equilibrium effects. The losses are unevenly distributed: €1.2 billion in manufacturing (0.7% of sectoral value added), €0.43 billion in construction (1.4%), €2.1 billion in business-to-business trade and services (1.2%), and €3.1 billion in business-to-consumer retail and services (2.8%). The largest relative gains accrue to retail, consumer services, and construction; manufacturing experiences the smallest relative loss, despite—along with construction—contributing the largest share of direct emissions.

Incorporating dynamic effects up to five months after exposure, using the sector-specific PDL estimates in Figure 2 significant at the 5% level, raises the estimated annual foregone sales to €38 billion and value-added losses to €10.3 billion. Depending on whether dynamic effects are included, failing to meet the WHO standard cost the French private sector between €6.9 and €10.3 billion in annual value added. For comparison, the European Commission's 2030 air-quality revision sets a 24-hour $PM_{2.5}$ limit of 25 $\mu g/m^3$; complying with this weaker standard would have delivered only about 40% of the benefits associated with the WHO guideline.

To weigh these benefits against costs, we follow Dechezleprêtre et al. (2020) and use estimates of the cost of reducing $PM_{2.5}$ emissions (rather than concentrations) from the European Commission (2013). Achieving a 33% reduction in French $PM_{2.5}$ emissions—more than enough to meet the WHO threshold—is estimated to cost around €7.7 billion per year in investment and maintenance for abatement equipment (option 6D, Table A7.3, p. 187). This represents 0.31% of French GDP, close to the average across European countries (0.3% of EU-28 GDP). The economic gains from cleaner air are therefore of the same magnitude as, and may exceed, this upper-bound estimate of

[44] This relies on the assumption that the linear specification provides a good approximation to the relationship between residualized log sales and the residualized $PM_{2.5}$ instrument. Figure A.13 supports this assumption.

[45] These benefits represent a lower bound because our sample covers only 56% of annual private-sector sales. However, our quantification may overstate aggregate welfare loss to the extent consumers substitute to establishments not covered by our sample. Additionally, ignoring input–output linkages may introduce some double counting at the aggregate level, since lower sales upstream mechanically reduce sales downstream—though such spillovers likely operate over horizons longer than the contemporaneous and one-month lags we use.

abatement costs. We expect the comparison to extend to other high-income European countries with similar sectoral composition and pollution levels.[46]

We finally contextualize these gains against the mortality benefits of the same policy. Combining Deryugina et al. (2019)'s estimates of the short-term mortality effects of $PM_{2.5}$ on the U.S. elderly with the French Value of a Statistical Life Year (VSLY) of €115,000 in 2010, we estimate that each one-unit decrease in $PM_{2.5}$ yields about €1.6 billion in annual mortality-reduction benefits in France.[47] Applied to our scenario, the implied mortality benefits amount to about €6.1 billion per year—comparable in magnitude to the avoided sales losses we document.

8 Conclusion

This paper examines how fine particulate matter ($PM_{2.5}$) affects economic activity in the French private sector. We show that higher pollution concentrations significantly reduce firm sales over a two-month horizon, with a one-month-lagged elasticity of -0.069 and a contemporaneous elasticity of -0.023. Three mechanisms can theoretically underpin these effects: a decrease in workers' labor supply, in their productivity, and a decrease in consumer demand. We find evidence of a labor supply channel, as air pollution raises sickness-related absenteeism (elasticity of 0.10). But even in the most affected sector—manufacturing—this channel accounts for at most a quarter of sales loss. Second, we find suggestive evidence for on-the-job productivity losses, most evident among manufacturing firms with low inventories, which cannot buffer transitory supply shocks by drawing down stocks. Third, heterogeneity across sectors and within the retail sector points to a demand-side channel: firms serving local consumers experience the sharpest sales decline, and among them, especially those selling discretionary goods. The effects are persistent: sales remain depressed for four to five months with no subsequent rebound, a pattern consistent with reductions in consumers' disposable income and related demand-side channels.

Our findings have direct implications for the design of air-quality regulation. Standard cost-benefit analyses that omit the impact of pollution on firm sales likely understate the net benefits of stricter standards. A notable feature of our results is that these gains accrue across all private sectors—not only to manufacturing and construction. Tightening $PM_{2.5}$ standards to meet WHO guidelines would generate economic benefits of the same magnitude as, and potentially exceeding, available estimates of abatement costs. Once the well-documented population-wide health benefits are added, the case for stricter air-quality standards becomes considerably stronger.

[46] While our estimated benefits may partially capture the effect of co-pollutants, the cost estimates correspond to a scenario that simultaneously reduces VOC and NH_3 emissions by 33%, NOx emissions by 22%, and SO_2 emissions by 19% (Appendix 7.1 of European Commission (2013), pp. 203–212).

[47] The calculation uses Deryugina et al. (2019)'s point estimate of 2.991 life-years gained per million elderly (65+) per unit decrease in daily $PM_{2.5}$, assumes annual effects scale linearly, converts the VSLY to 2013 euros, and applies France's 11.7 million elderly population in 2013. The French VSLY of €115,000 per life-year is comparable to the \$100,000 value used for the U.S. in Deryugina et al. (2019).

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. Graff Zivin, and M. Neidell (2022). This Is Air: The "Nonhealth" Effects of Air Pollution. *Annual Review of Resource Economics 14*(1), 403–425.

Almunia, M., P. Antràs, D. Lopez-Rodriguez, and E. Morales (2021). Venting out: Exports during a domestic slump. *American Economic Review 111*(11), 3611–62.

Andrews, I., J. H. Stock, and L. Sun (2019). Weak Instruments in Instrumental Variables Regression: Theory and Practice. *Annual Review of Economics* (Volume 11, 2019), 727–753.

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.

Barwick, P. J., S. Li, D. Rao, and N. B. Zahur (2024). The Healthcare Cost of Air Pollution: Evidence from the World's Largest Payment Network. *The Review of Economics and Statistics*, 1–52.

Borgschulte, M., D. Molitor, and E. Y. Zou (2024). Air Pollution and the Labor Market: Evidence from Wildfire Smoke. *Review of Economics and Statistics 106*(6), 1558–1575.

Borusyak, K., P. Hull, and X. Jaravel (2025). A Practical Guide to Shift-Share Instruments. *Journal of Economic Perspectives 39*(1), 181–204.

Bruyneel, L., W. Kestens, M. Alberty, G. Karakaya, R. Van Woensel, C. Horemans, E. Trimpeneers, C. Vanpoucke, F. Fierens, T. S. Nawrot, and B. Cox (2022). Short-Term exposure to ambient air pollution and onset of work incapacity related to mental health conditions. *Environment International 164*, 107245.

Calderón-Garcidueñas, L., A. Mora-Tiscareño, E. Ontiveros, G. Gómez-Garza, G. Barragán-Mejía, J. Broadway, S. Chapman, G. Valencia-Salazar, V. Jewells, R. R. Maronpot, C. Henríquez-Roldán, B. Pérez-Guillé, R. Torres-Jardón, L. Herrit, D. Brooks, N. Osnaya-Brizuela, M. E. Monroy, A. González-Maciel, R. Reynoso-Robles, R. Villarreal-Calderon, A. C. Solt, and R. W. Engle (2008). Air pollution, cognitive deficits and brain abnormalities: a pilot study with children and dogs. *Brain and Cognition 68*(2), 117–127.

Champalaune, P. (2020). Inequality in Exposure to Air Pollution in France: Measurement and Impact of a City-Level Public Policy. pp. 67.

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.

Currie, J., J. Voorheis, and R. Walker (2023). What Caused Racial Disparities in Particulate

Exposure to Fall? New Evidence from the Clean Air Act and Satellite-Based Measures of Air Quality. *American Economic Review 113*(1), 71–97.

Dahl, G. B. and L. Lochner (2012). The Impact of Family Income on Child Achievement: Evidence from the Earned Income Tax Credit. *American Economic Review 102*(5), 1927–1956.

Dechezleprêtre, A., N. Rivers, and B. Stadler (2020). The economic cost of air pollution: Evidence from Europe. *OECD Economics Department Working Papers 1584*.

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.

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

European Commission (2013). Comission staff working document: Impact assessment on the proposal for a directive of the european parliament and of the council on the reduction of national emissions of certain atmospheric pollutants and amending directive 2003/35/ec.

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

France Stratégie and Inspection générale des Finances (2021). Comité de suivi et d'évaluation des mesures de soutien financier aux entreprises confrontées à l'épidémie de covid-19 - Rapport final. Technical report.

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.

Gertler, M. and S. Gilchrist (1994). Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms∗. *Quarterly journal of Economics 109*(2), 309–340.

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

Graff Zivin, J., M. Neidell, N. J. Sanders, and G. Singer (2023). When Externalities Collide: Influenza and Pollution. *American Economic Journal: Applied Economics 15*(2), 320–51.

Griliches, Z. and J. A. Hausman (1986). Errors in variables in panel data. *Journal of Econometrics 31*(1), 93–118.

Hanna, R. and P. Oliva (2015). The effect of pollution on labor supply: Evidence from a natural experiment in Mexico City. *Journal of Public Economics 122*, 68–79.

Harrigan, J., A. Reshef, and F. Toubal (2024). Techies and Firm Level Productivity. *mimeo*.

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.

Hersbach, H., B. Bell, P. Berrisford, G. Biavati, A. Horányi, J. Muñoz Sabater, J. Nicolas, C. Peubey, R. Radu, I. Rozum, D. Schepers, A. Simmons, C. Soci, D. Dee, and J.-N. Thépaut (2018). Era5 hourly data on single levels from 1959 to present. *Copernicus Climate Change Service (C3S) Climate Data Store (CDS) (Accessed on 20-JUL-2022).*

Hoffmann, B. and J. P. Rud (2024). The Unequal Effects of Pollution on Labor Supply. *Econometrica 92*(4), 1063–1096.

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

INSEE (2021). Déplacements domicile-travail – Même sur de très courts trajets, l'usage de la voiture reste majoritaire - Insee Flash Provence-Alpes-Côte d'Azur - 70.

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

Lee, S. and S. Zheng (2025). Extreme Temperatures, Adaptation Capacity, and Household Retail Consumption. *Journal of the Association of Environmental and Resource Economists*.

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 (2024). Fresh air eases work—the effect of air quality on individual investor activity. *Review of Finance 28*(3), 1105–1149.

Montiel Olea, J. L. and C. E. Pflueger (2013). A Robust Test for Weak Instruments. *Journal of Business and Economic Statistics*.

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

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

Prost, J.-M. and J.-P. Villetelle (2018). Rapport annuel de l'Observatoire des delais de paiement 2017. Contre les retards de paiement: Les entreprises inegalement mobilisees. *Banque de France publication*.

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

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

Schwartz, J. (2000). The distributed lag between air pollution and daily deaths. *Epidemiology (Cambridge, Mass.) 11*(3), 320–326.

Shen, S., C. Li, A. van Donkelaar, N. Jacobs, C. Wang, and R. V. Martin (2024). Enhancing Global Estimation of Fine Particulate Matter Concentrations by Including Geophysical a Priori Information in Deep Learning. *ACS EST Air 1*(5), 332–345.

Sicard, P., E. Agathokleous, A. De Marco, E. Paoletti, and V. Calatayud (2021). Urban population exposure to air pollution in Europe over the last decades. *Environmental Sciences Europe 33*(1), 28.

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

van Donkelaar, A., M. S. Hammer, L. Bindle, M. Brauer, J. R. Brook, M. J. Garay, N. C. Hsu, O. V. Kalashnikova, R. A. Kahn, C. Lee, R. C. Levy, A. Lyapustin, A. M. Sayer, and R. V. Martin (2021). Monthly Global Estimates of Fine Particulate Matter and Their Uncertainty. *Environmental Science and Technology 55*(22), 15287–15300.

Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data*. The MIT Press.

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

Zhang, W., H. Sun, S. Woodcock, and A. H. Anis (2017). Valuing productivity loss due to absenteeism: firm-level evidence from a Canadian linked employer-employee survey. *Health Economics Review 7*, 1–14.

Additional Materials – For online publication only

A Additional Figures and Tables

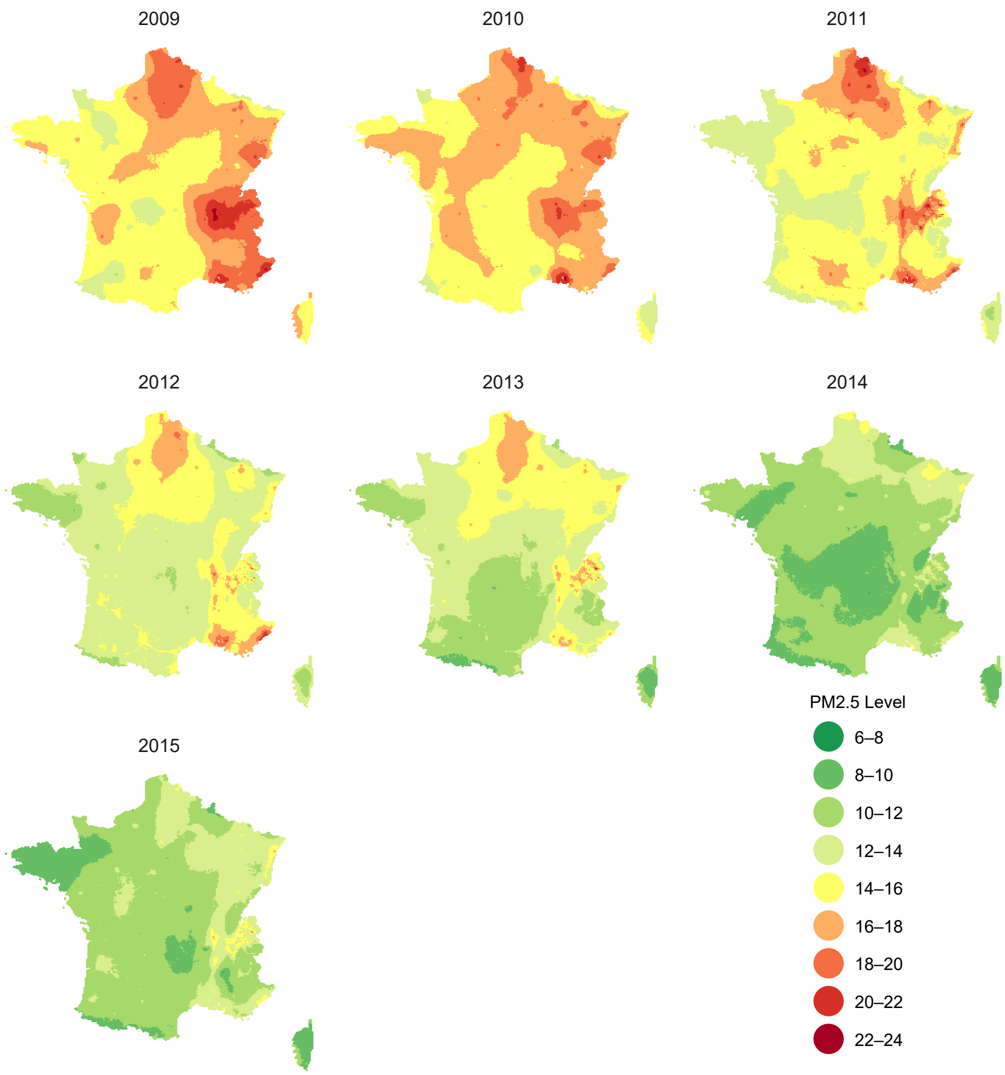

Figure A.1: Average annual concentrations of $PM_{2.5}$ (µg/m^3)

Notes: Figure shows the average annual concentration of $PM_{2.5}$ measured at the 4km x 4 km grid cell level using the reanalysis pollution data. There are 33,252 such grid cells in metropolitan France.

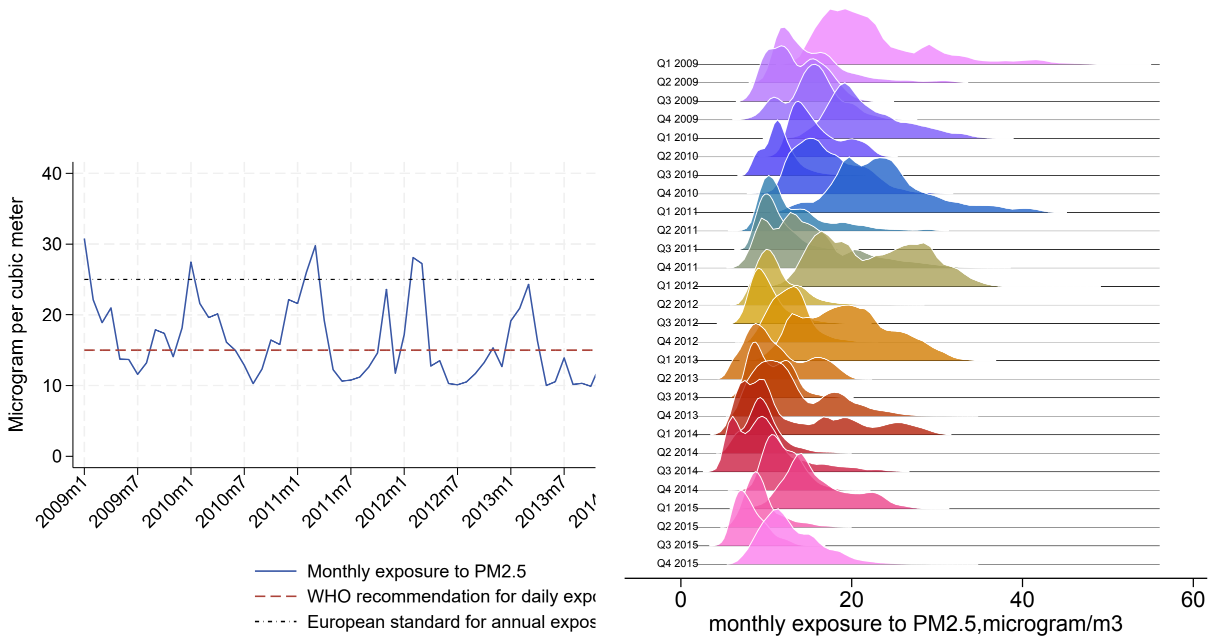

(a) Monthly average exposure to $PM_{2.5}$ (µg/m^3) (b) Distribution from Q1 2009 to Q4 2015

Figure A.2: Monthly exposure to $PM_{2.5}$ (µg/m^3)

Notes: Figure a) shows municipality-level $PM_{2.5}$ exposure in 2009-2015, weighted by the number of workers employed in each municipality in the absenteeism dataset. Figure b) shows the unweighted distribution of monthly exposure to $PM_{2.5}$.

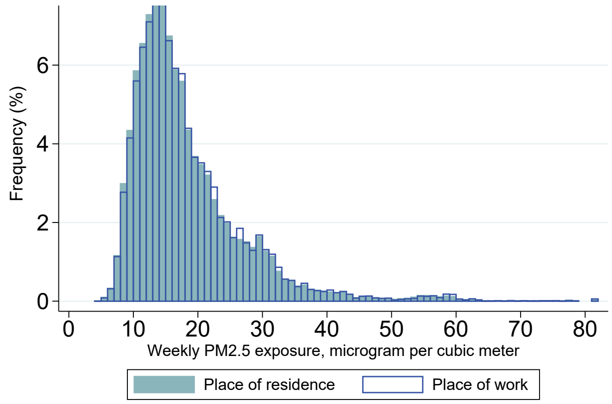

(a) All

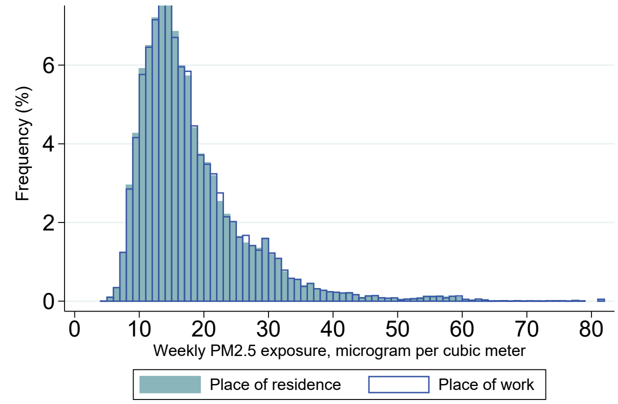

(b) Q1

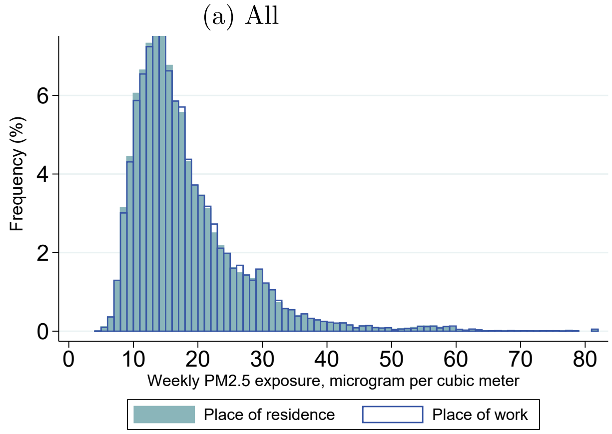

(c) Q2

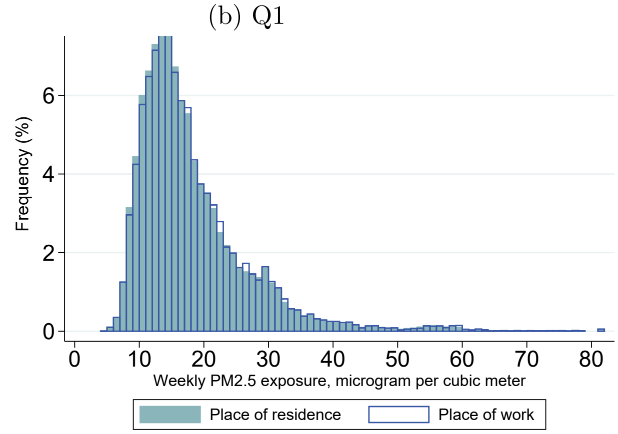

(d) Q3

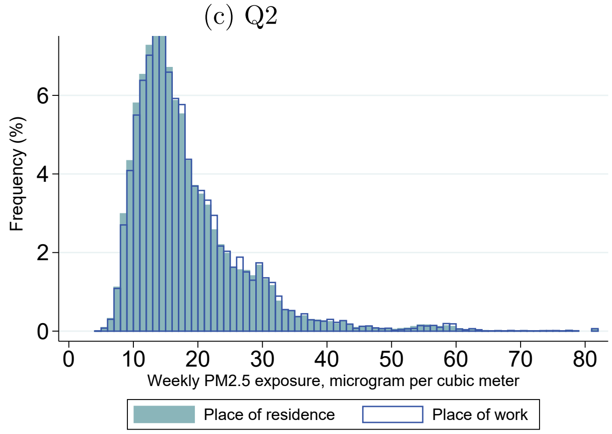

(e) Q4

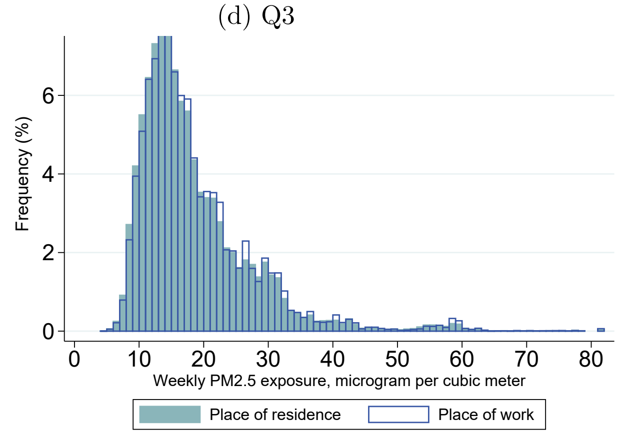

(f) Q5

Figure A.3: Distribution of $PM_{2.5}$ pollution exposure at the municipality of residence and at the municipality of workplace in 2009, for all private sector workers and by wage quintile

Notes: Data source: exhaustive matched-employer employee data for all private sector workers in 2009.

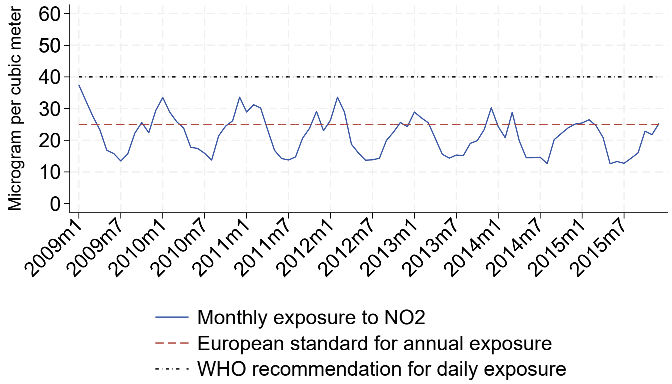

(a) NO_2

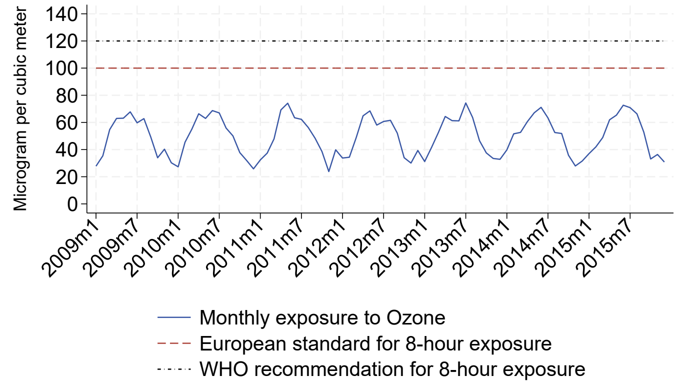

(b) Ozone

Figure A.4: Average monthly exposure to nitrogen dioxide (NO_2 and ozone

Notes: Figure shows the monthly average of workers' exposure to $PM_{2.5}$ at the municipality of their workplace. For NO_2, the European standard for annual exposure is 40µg/m^3 while the WHO's recommendation for daily exposure is 25µg/m^3. For ozone, the European standard for 8-hour exposure is 120µg/m^3 while the WHO's recommendation for 8-hour exposure is 100µg/m^3. Exposure in each municipality is weighted by the number of workers employed in each municipality in the absenteeism dataset.

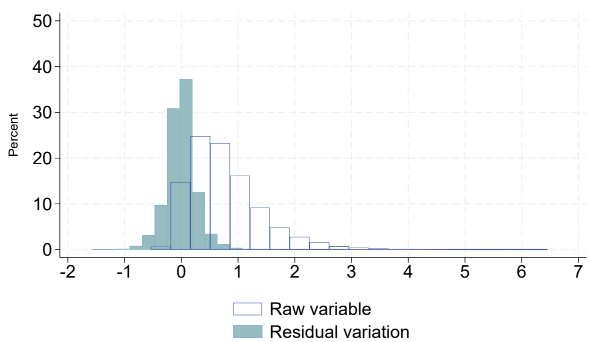

(a) East wind

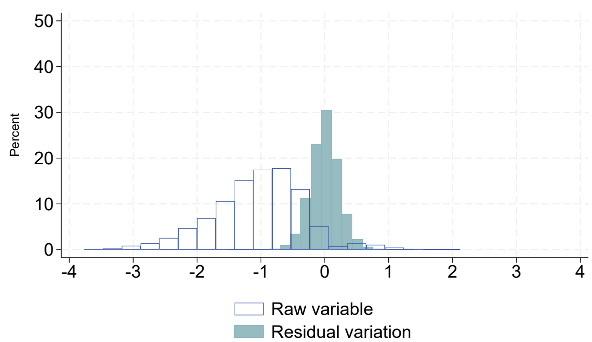

(b) West wind

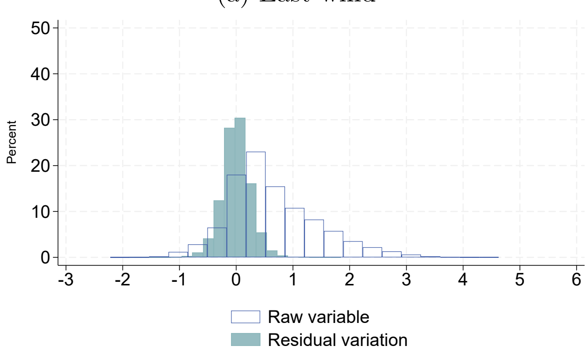

(c) North wind

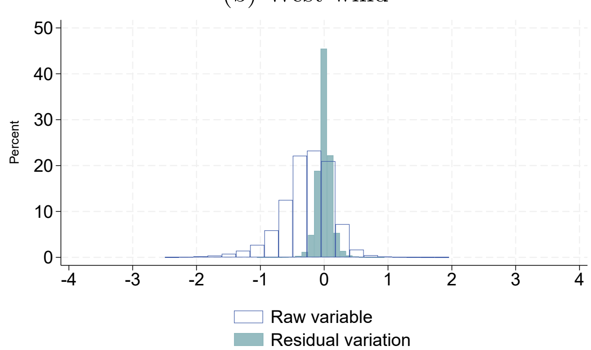

(d) South wind

Figure A.5: Identifying variation - Distribution of raw and residualized wind instrument

Notes: Residualized variables are obtained by regressing each lagged wind instrument value on the right-hand side variables of equation (10) for the sample of single-establishment firms: weather and holiday controls at $t-1$, t, and $t+1$, industry-by-month-by-year fixed effects, quarter-by-county fixed effects, firm-by-year fixed effects, and controls for the wind instrument at t and $t+1$.

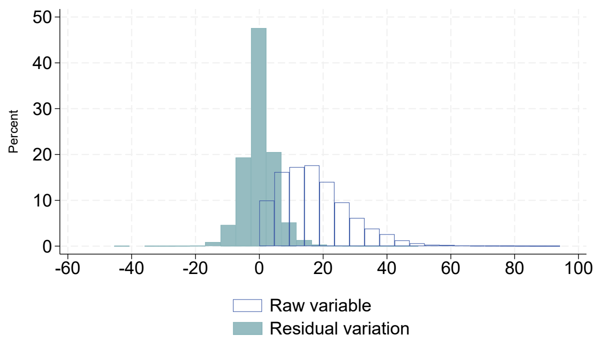

(a) East wind

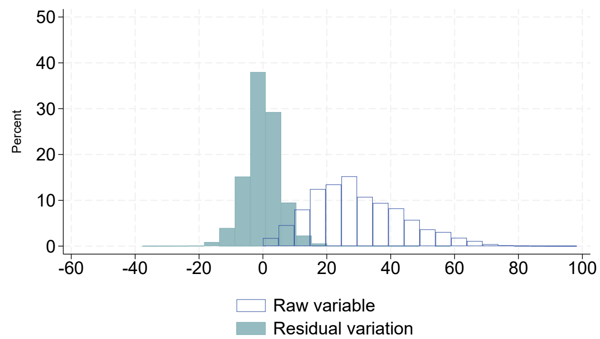

(b) West wind

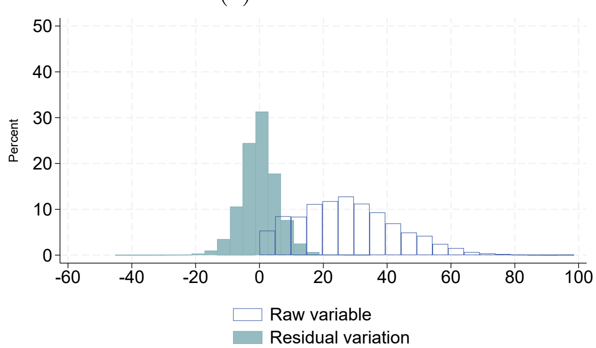

(c) North wind

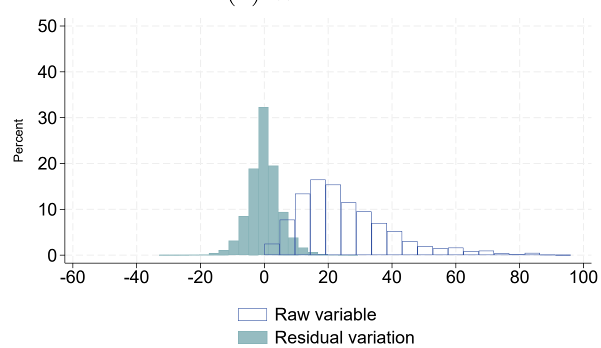

(d) South wind

Figure A.6: Identifying variation - Distribution of raw and residualized share of each wind's direction

Notes: Residualized variables are obtained by regressing monthly shares of each lagged wind direction (at $t-1$) on the right-hand side variables of equation (10) for the sample of single-establishment firms: weather and holiday controls at $t-1$, t, and $t+1$, industry-by-month-by-year fixed effects, quarter-by-county fixed effects, firm-by-year fixed effects, and controls for wind directions at t and $t+1$.

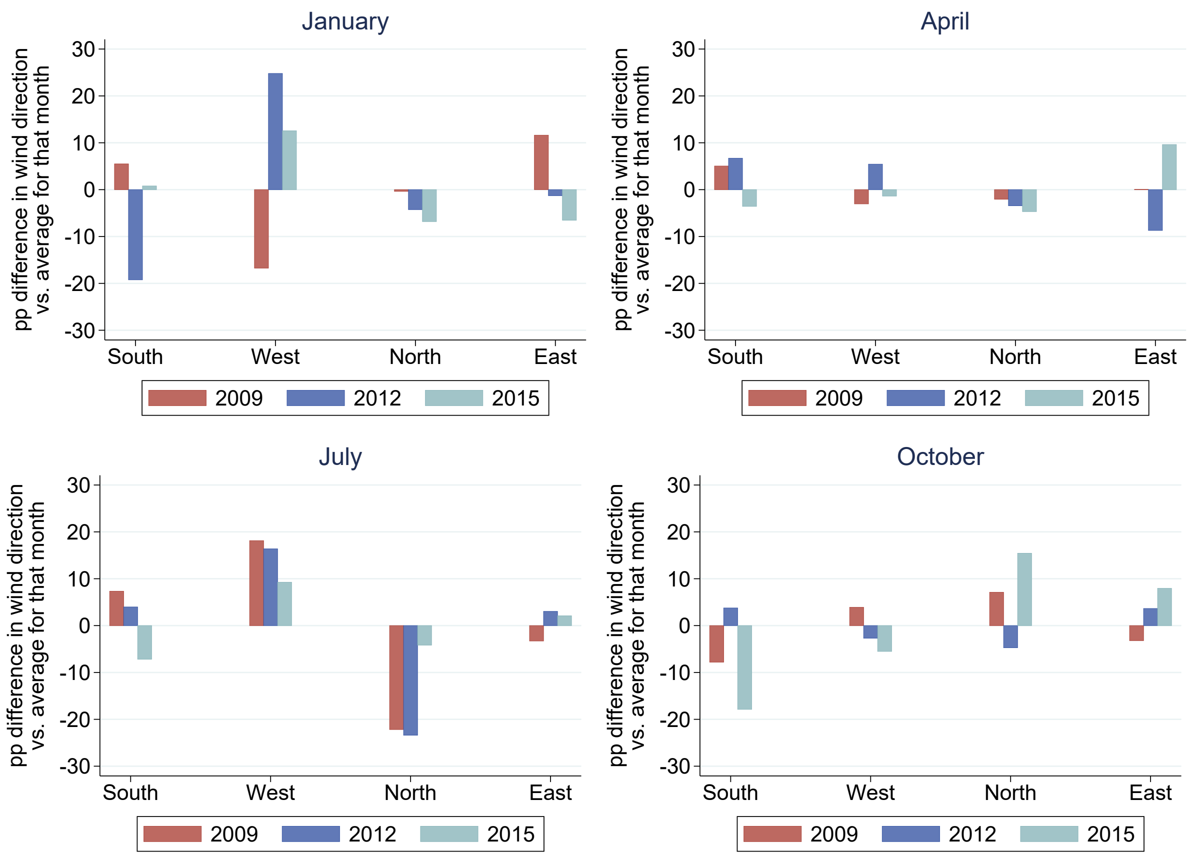

Figure A.7: Within-calendar month variation in wind direction, Paris

Notes: Figure shows the share of hours in a month in which the wind blows from a given direction, demeaned by the average for the month, for four different months (Month 3=March, Month 6=June, Month 9=September, Month 12=December) and three different years (2009, 2012, 2015).

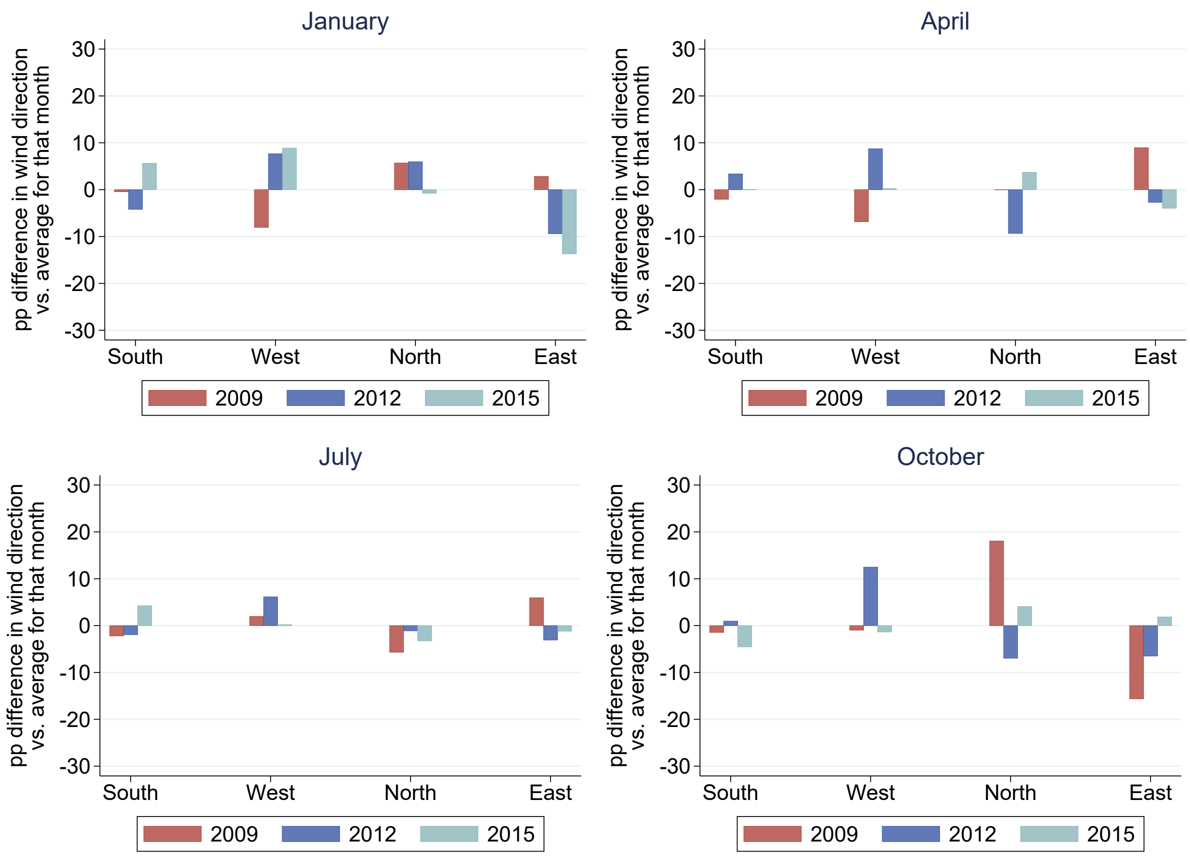

Figure A.8: Within-calendar month variation in wind direction, Marseille (South-East of France)

Notes: Figure shows the share of hours in a month in which the wind blows from a given direction, demeaned by the average for the month, for four different months (Month 3=March, Month 6=June, Month 9=September, Month 12=December) and three different years (2009, 2012, 2015).

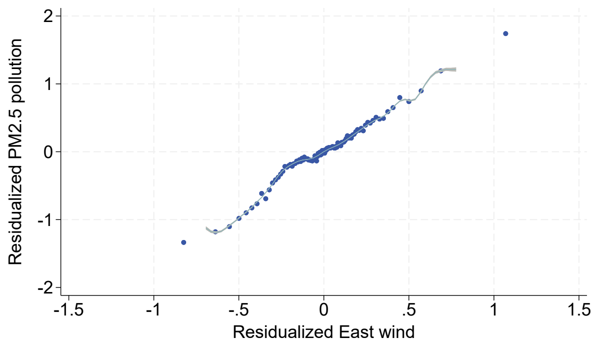

(a) East wind

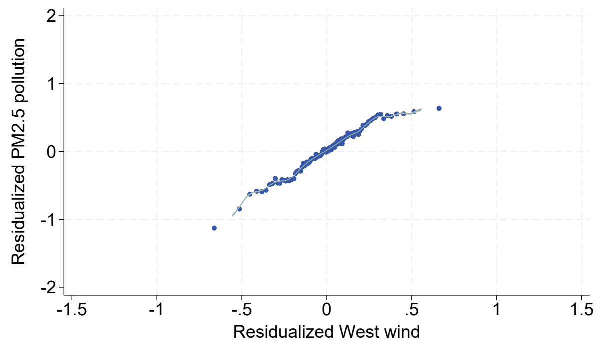

(b) West wind

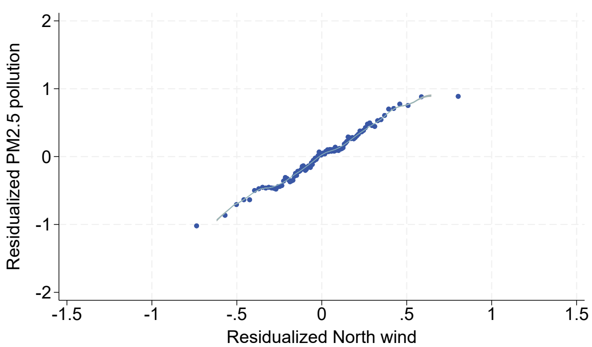

(c) North wind

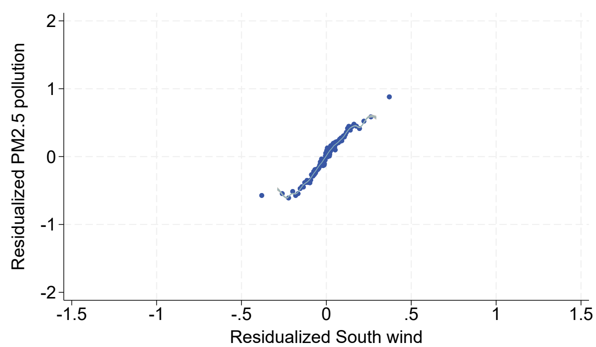

(d) South wind

Figure A.9: Residualized binned scatter plot between wind instruments and $PM_{2.5}$ concentrations and local polynomial fit

Notes: Figure is based on the sample of single-establishment firms. x-axis residuals (resp. y-axis residuals) are obtained by regressing each wind instrument value (resp. $PM_{2.5}$) on the right-hand side variables of equation 10: weather and holiday controls, industry-by-month-by-year fixed effects, quarter-by-county fixed effects and firm-by-year fixed effects. Observations are grouped in equal-sized bins (centiles) based on the value of the x-variable. Each dot shows the mean value of that bin for the x-axis and y-axis residuals. The solid blue line shows a local polynomial regression fit of residualized $PM_{2.5}$ on residualized wind direction, with the grey area around showing 95% confidence bands (leaving out the top and bottom 1% of the distribution).

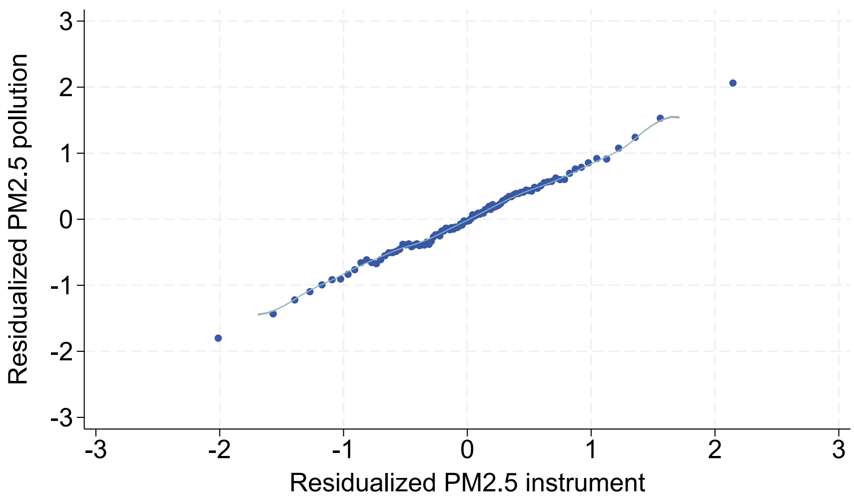

Figure A.10: Residualized binned scatter plot between wind instruments and $PM_{2.5}$ concentrations and local polynomial fit

Notes: Figure is based on the sample of all firms. x-axis residuals (resp. y-axis residuals) are obtained by regressing the predicted $PM_{2.5}$ variable $\widehat{PM_{fyt}}$ (resp. the endogenous $PM_{2.5}$ variable) on the right-hand side variables of equation 10: weather and holiday controls, industry-by-month-by-year fixed effects, and firm-by-year fixed effects. Observations are grouped in equal-sized bins (centiles) based on the value of the x-variable. Each dot shows the mean value of that bin for the x-axis and y-axis residuals. The solid blue line shows a local polynomial regression fit of residualized $PM_{2.5}$ on residualized wind direction, with the grey area around showing 95% confidence bands (leaving out the top and bottom 1% of the distribution).

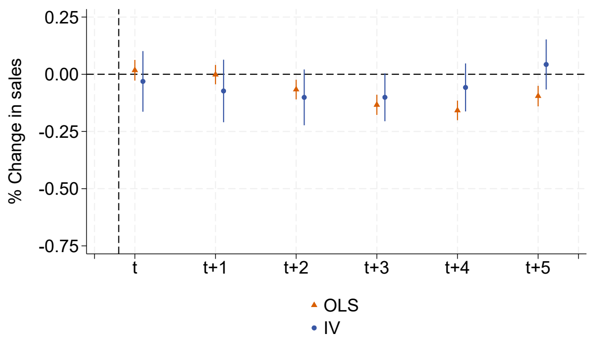

(a) Manufacturing

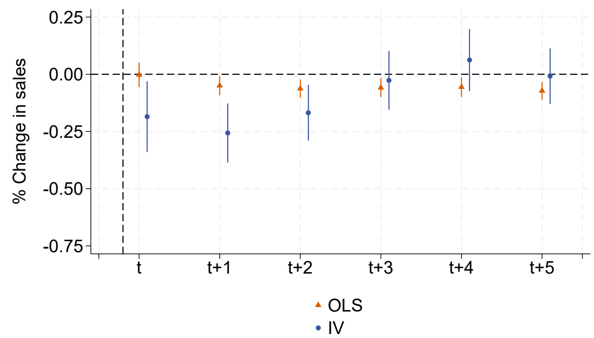

(b) Construction

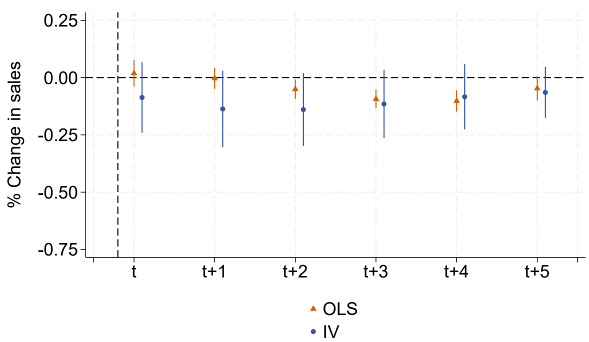

(c) Business-to-Business Trade & Services

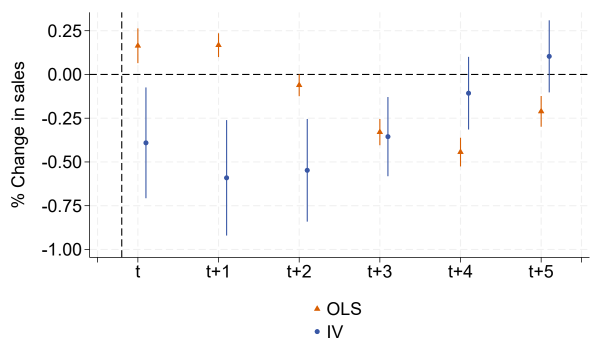

(d) Business-to-Consumer Retail & Services

Figure A.11: Dynamic effects of $PM_{2.5}$ on sales of single-establishment firms, by sector

Notes: Figure shows the OLS and IV point estimates and 95% confidence intervals from equation (7) for the effect of contemporaneous and lagged $PM_{2.5}$ (up to $t-5$) on firms' sales outcome at t by sector, using the polynomial distributed lag method. All regressions include month-by-year-by-industry, firm-by-year, and quarter-by-county fixed effects, as well as weather and holidays controls from $t-1$ to $t+1$. Standard errors are clustered at the Copernicus grid cell level. Coefficients (and standard errors) have been multiplied by 100 for readability.

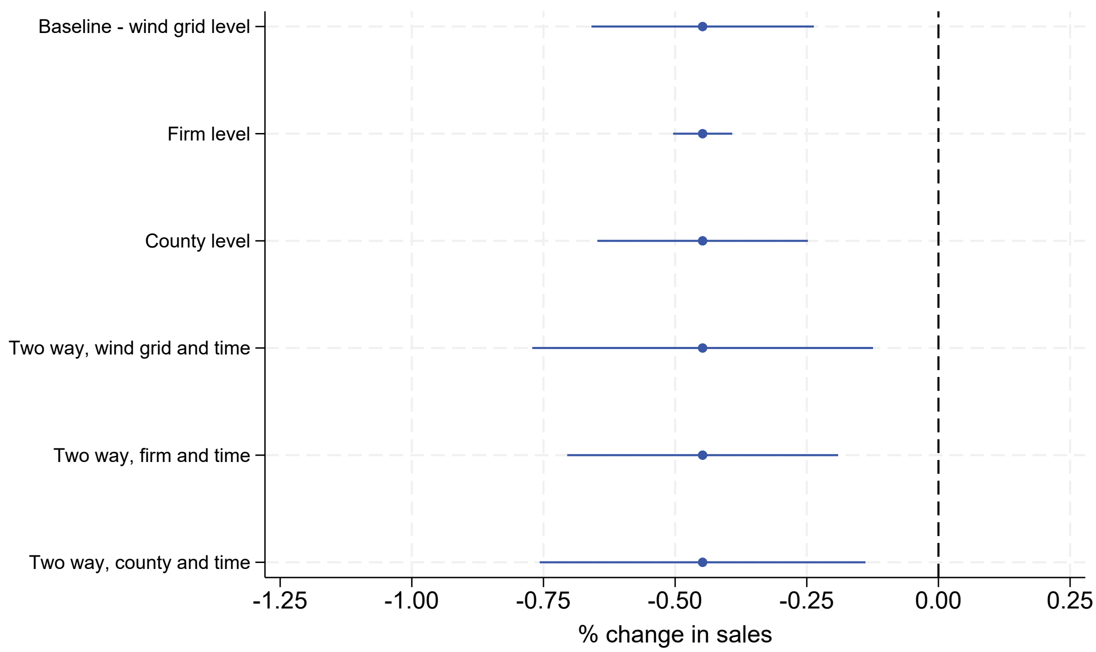

Figure A.12: Sensitivity of the results to the level of clustering for standard errors

Notes: Figure shows IV point estimates and 95% confidence intervals for the effect of $PM_{2.5}$ at $t-1$ and t on firms' sales at t, based on equation (7). All regressions include month-by-year-by-industry fixed effects, firm-by-year fixed effects, quarter-by-county fixed effects, weather and holidays controls at $t-1$, t, and $t+1$, as well as instrumented pollution at $t+1$. Coefficients (and standard errors) have been multiplied by 100 for readability.

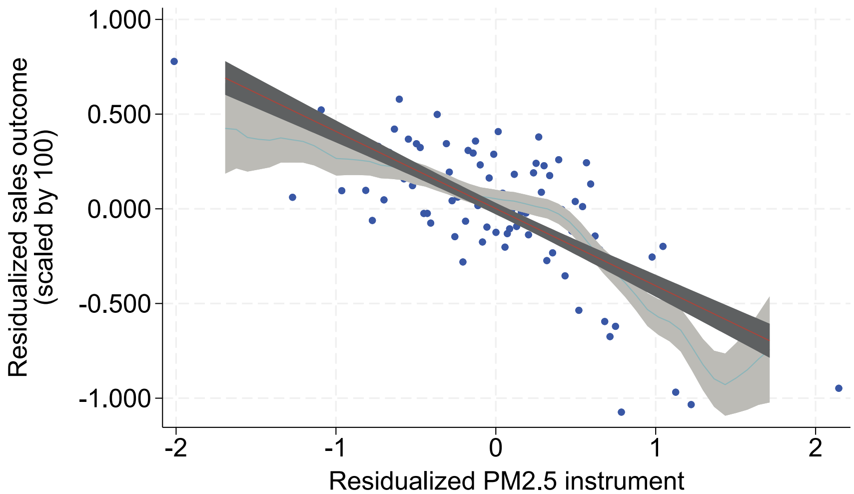

Figure A.13: Residualized binned scatter plot between sales and $PM_{2.5}$ instrument and local polynomial fit

Notes: Figure is based on the sample of single-establishment firms. x-axis residuals (resp. y-axis residuals) are obtained by regressing the sales outcome (the predicted $PM_{2.5}$ variable, $\widehat{PM_{fyt}}$) on the right-hand side variables of equation 10: weather and holiday controls, industry-by-month-by-year fixed effects, and firm-by-year fixed effects. Observations are grouped in equal-sized bins (centiles) based on the value of the x-variable. Each dot shows the mean value of that bin for the x-axis and y-axis residuals. The solid blue line shows a local polynomial regression fit of residualized $PM_{2.5}$ on residualized wind direction, with the grey area around showing 95% confidence bands (leaving out the top and bottom 1% of the distribution). The solid red line shows a linear regression fit, with the dark grey area around showing 95% confidence bands.

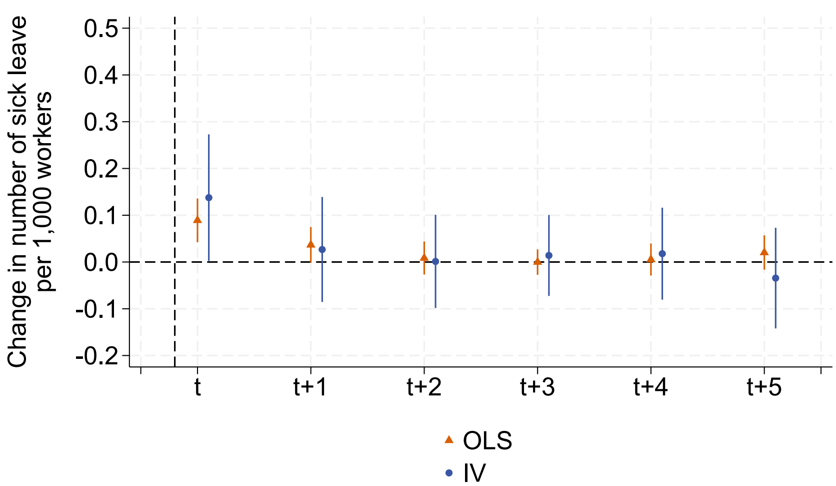

Figure A.14: Dynamic effects of $PM_{2.5}$ on absenteeism

Notes: Figure shows the OLS and IV point estimates and 95% confidence intervals based on equation (8) for the effect of contemporaneous and lagged $PM_{2.5}$ (up to $t-5$) on the number of workers entering sick leave at t per 1,000 workers, using the polynomial distributed lag method. All regressions include month-by-year-by-industry, establishment, and quarter-by-county fixed effects, as well as weather and holidays controls from $t-1$ to $t+1$. Standard errors are clustered at the Copernicus grid cell level.

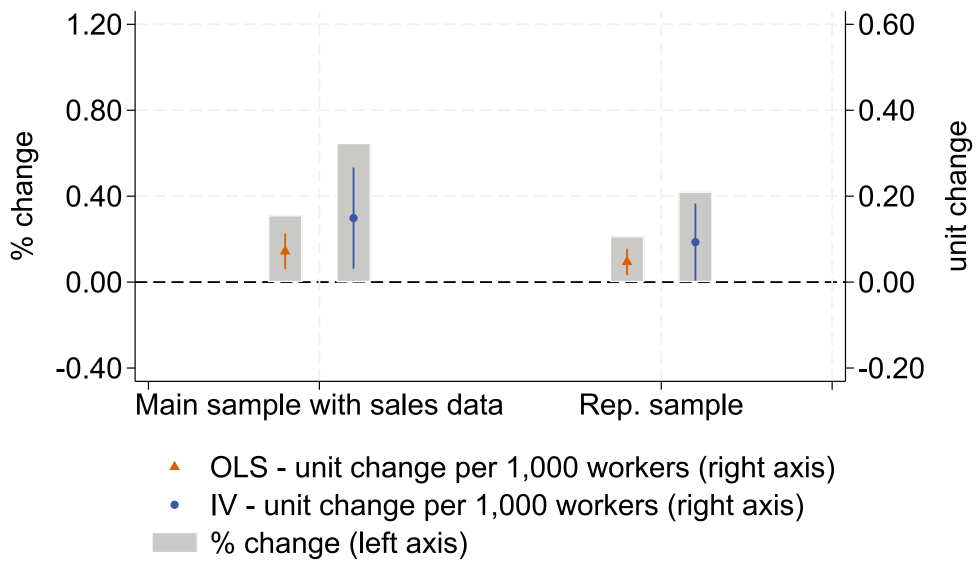

Figure A.15: Absenteeism results for our main sample, only including workers employed in the firms of our sales sample (left), vs. the representative sample including all workers (right)

Notes: Figure shows the OLS and IV point estimates and 95% confidence intervals for the effect of contemporaneous $PM_{2.5}$ on the number of workers entering sick leave at t per 1,000 workers, based on equation (8). All regressions include month-by-year-by-industry, establishment, and quarter-by-county fixed effects, as well as weather and holidays controls. Standard errors are clustered at the Copernicus grid cell level.

Table A.1: Workers' characteristics (aggregated at establishment level), 2009-2015

Sample	Main sample establishments with sales data		All establishments with absenteeism	
	Mean	Sd	Mean	Sd
Age	40.2	8.7	40.4	8.9
Annual wage	28,542.0	20,576.1	25,911.0	20,547.4
Annual total medical expenditures	442.0	809.8	462.5	819.8
Works in a single-establishment firm	41%	0.49	-	-
Works in: Manufacturing	28%	0.45	17%	0.37
Construction	12%	0.32	7%	0.26
Business-to-business services	33%	0.47	20%	0.40
Business-to-consumer services	27%	0.39	16%	0.32
Others	0%	-	40%	0.49
Exposure to $PM_{2.5}$ ($\mu g/m^3$)	**15.4**	6.3	**15.3**	6.3
Workers falling sick each month (per 1,000)	24.7	113.4	23.9	111.3
incl: for <93 days	**23.0**	109.2	**22.1**	107.0
N	8,233,440		16,409,124	

Notes: Table reports descriptive statistics on workers included in the absenteeism dataset, aggregated at the establishment level applying worker weights. The left panel is based on our main sample of analysis for whom we have sales data and the right panel on the entire representative sample of private sector employees.

Table A.2: The Effect of $PM_{2.5}$ on Firm-level Sales, Adding Fixed Effects progressively

	(1)	(2)	(3)	(4)	(5)
Panel A: OLS results					
PM_{t-1}	-0.0404*	0.0335	0.0325	0.00147	0.00766
	(0.0216)	(0.0258)	(0.0251)	(0.0213)	(0.0207)
PM_{t}	0.111***	0.145***	0.137***	0.0786***	0.0813***
	(0.0228)	(0.0284)	(0.0297)	(0.0235)	(0.0254)
N	9,412,076	9,411,967	9,403,047	9,411,967	9,403,047
R-squared	0.9208	0.9456	0.9468	0.9459	0.9470
Panel B: IV results					
PM_{t-1}	-0.582***	-0.576***	-0.539***	-0.461***	-0.448***
	(0.131)	(0.114)	(0.116)	(0.103)	(0.108)
PM_{t}	-0.266***	-0.275***	-0.257***	-0.154**	-0.151**
	(0.0832)	(0.0827)	(0.0791)	(0.0673)	(0.0675)
N	9,412,076	9,411,945	9,411,781	9,411,945	9,411,781
R-squared	0.9208	0.9456	0.9467	0.9458	0.9469
Firm FE	Yes	No	No	No	No
Firm-by-year FE	No	Yes	Yes	Yes	Yes
Month-by-year FE	Yes	Yes	No	Yes	No
Month-by-year-by-industry FE	No	No	Yes	No	Yes
Quarter-by-departement FE	No	No	No	Yes	Yes

Notes: Table reports the OLS and IV estimates of the effect of a one unit increase in $PM_{2.5}$ at $t-1$ and t on the sales outcome at t based on equation (7) for all firms in all sectors. All regressions include weather and holidays controls at $t-1$, t, and $t+1$, as well as instrumented pollution at $t+1$. Standard errors are clustered at the Copernicus grid cell level. Coefficients (and standard errors) have been multiplied by 100 for readability. We denote * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Table A.3: The Effect of adding lags of $PM_{2.5}$

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: OLS results						
PM_t	0.0527*	0.0819***	0.0949***	0.0851***	0.0791***	0.0680***
	(0.0286)	(0.0253)	(0.0259)	(0.0239)	(0.0247)	(0.0246)
PM_{t-1}		0.00794	0.0275	0.0220	0.00412	-0.00419
		(0.0207)	(0.0202)	(0.0214)	(0.0224)	(0.0238)
PM_{t-2}			-0.0198	0.00382	-0.0150	-0.0290
			(0.0179)	(0.0183)	(0.0175)	(0.0185)
PM_{t-3}				-0.165***	-0.154***	-0.162***
				(0.0235)	(0.0222)	(0.0223)
PM_{t-4}					-0.221***	-0.240***
					(0.0213)	(0.0240)
PM_{t-5}						-0.105***
						(0.0203)
N	9,585,132	9,403,025	9,220,160	9,036,731	8,853,170	8,669,685
R-squared	0.9461	0.9470	0.9476	0.9480	0.9485	0.9489
Panel B: IV results						
PM_t	-0.175***	-0.148**	-0.163**	-0.200**	-0.210***	-0.190**
	(0.0640)	(0.0672)	(0.0728)	(0.0798)	(0.0806)	(0.0810)
PM_{t-1}		-0.447***	-0.423***	-0.425***	-0.448***	-0.444***
		(0.107)	(0.101)	(0.100)	(0.105)	(0.110)
PM_{t-2}			-0.243***	-0.245***	-0.257***	-0.250***
			(0.0788)	(0.0774)	(0.0764)	(0.0763)
PM_{t-3}				-0.207***	-0.225***	-0.213***
				(0.0633)	(0.0629)	(0.0637)
PM_{t-4}					-0.129***	-0.139***
					(0.0464)	(0.0497)
PM_{t-5}						0.00373
						(0.0458)
N	9,585,132	9,403,025	9,220,160	9,036,731	8,853,170	8,669,685
R-squared	0.9461	0.9470	0.9476	0.9480	0.9485	0.9489
Weather and holiday controls at t and $t+1$	Yes	Yes	Yes	Yes	Yes	Yes
Weather and holiday controls at $t-1$	No	Yes	Yes	Yes	Yes	Yes

Notes: Table reports the OLS and IV estimates of the effect of a one unit increase in $\text{PM}_{2.5}$ at different lags on the sales outcome at t based on equation (7) for all firms in all sectors. All regressions include weather and holidays controls at $t-1$, t, and $t+1$, as well as instrumented pollution at $t+1$. Standard errors are clustered at the Copernicus grid cell level. Coefficients (and standard errors) have been multiplied by 100 for readability. We denote * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Table A.4: Heterogeneous sales responses to $PM_{2.5}$, by firm size

	(1) Below 15 workers	(2) Above 15 workers
Panel A: All firms		
PM_{t-1}	-0.499***	-0.367***
	(0.119)	(0.0957)
PM_t	-0.133	-0.151**
	(0.0846)	(0.0606)
N	4,518,389	4,884,519
R-squared	0.8527	0.9386
Panel B: Manufacturing		
PM_{t-1}	-0.367***	-0.0546
	(0.123)	(0.0644)
PM_t	-0.0229	-0.0496
	(0.0964)	(0.0587)
N	603,306	1,272,780
R-squared	0.8395	0.9559
Panel C: Construction		
PM_{t-1}	-0.363***	-0.0659
	(0.0906)	(0.0737)
PM_t	-0.135	-0.159**
	(0.0835)	(0.0679)
N	837,109	693,235
R-squared	0.8026	0.9253
Panel D: Business-to-Business Trade and Services		
PM_{t-1}	-0.304***	-0.260***
	(0.0972)	(0.0870)
PM_t	-0.0458	0.000512
	(0.0735)	(0.0701)
N	1,226,875	1,646,457
R-squared	0.8492	0.9281
Panel E: Business-to-Consumer Retail and Services		
PM_{t-1}	-0.665***	-0.872***
	(0.186)	(0.216)
PM_t	-0.149	-0.320**
	(0.140)	(0.132)
N	1,847,214	1,275,900
R-squared	0.8737	0.9377

Notes: Table reports the IV estimates of the effect of a one unit increase in $PM_{2.5}$ at $t-1$ and t on the sales outcome at t based on (7) for firms with fewer than 15 workers on average (column 1) and those with at least 15 workers (column 2). All regressions include firm-by-year, month-by-year-by-industry and quarter-by-county fixed effects, weather and holidays controls at $t-1$, t, and $t+1$, as well as instrumented pollution at $t+1$. The confidence intervals are based on standard errors clustered at the Copernicus grid cell level. Coefficients (and standard errors) have been multiplied by 100 for readability. We denote * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Table A.5: Falsification test: Effect of future air pollution shocks on contemporaneous sales

	All	Manufact.	Construction	Business-to-business trade and services	Business-to-consumer trade and services
PM_{t+2}	-0.0197	-0.0839	0.000437	-0.0323	-0.0196
	(0.0441)	(0.0527)	(0.0607)	(0.0498)	(0.101)
N	9,402,279	1,880,385	1,531,601	2,874,733	3,124,309
R-squared	0.9470	0.9643	0.9354	0.9339	0.9460

Notes: Table reports the IV estimates of the effect of a one unit increase in $PM_{2.5}$ at $t+2$ on the sales outcome at t based on equation (7) for all firms and by sector. All regressions include firm-by-year fixed effects, month-by-year-by-industry fixed effects and quarter-by-county fixed effects, as well as weather and holidays controls at t. Standard errors are clustered at the Copernicus grid cell level of the firm's headquarter. Coefficients (and standard errors) have been multiplied by 100 for readability. We denote * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Table A.6: Sensitivity of the results to the choice of instrument

	(1) Baseline 2009-2015	(2) Early-period component B 2009-2015	(3) Component A only 2009-2015	(4) Baseline 2010-2015	(5) Early-period component B 2010-2015
Panel A: All firms					
PM_{t-1}	-0.448***	-0.357***	-0.217***	-0.435***	-0.311***
	(0.108)	(0.110)	(0.0752)	(0.107)	(0.108)
N	9,411,781	9,381,735	9,411,781	8,172,701	8,146,573
Panel B: Single establishment firms					
PM_{t-1}	-0.482***	-0.400***	-0.136	-0.488***	-0.369***
	(0.111)	(0.118)	(0.086)	(0.112)	(0.116)
First stage KP F-stat	120	60	231	92	61
N	6,072,012	6,048,456	6,072,012	5,277,251	5,256,729

Notes: Table reports the IV estimates of the effect of a one-unit increase in $PM_{2.5}$ at $t-1$ on the sales outcome at t based on equation (7), for all firms in all sectors in Panel A and for single-establishment firms only in Panel B. All regressions include weather and holidays controls at $t-1$, t, and $t+1$, firm-by-year fixed effects, quarter-by-county fixed effects and industry-by-month-by year fixed effects. All regressions in panel A use predicted pollution as instrument to accommodate for multi-establishment and hence multi-location firms, while those in panel B use directly the wind direction instruments. Regressions in all columns in panel A, and in columns (1), (2), (4) and (5) in panel B include instrumented pollution at t and $t+1$, while the regression in column (3) of panel B simply controls for the instrument at t and $t+1$ given computational constraints. Column (1) shows the baseline result. Column (2) shows the result with a modified instrumental variable where component B of each instrument is calculated using data from the first period only, 2009. Column (3) shows the result with an alternative instrumental variable leveraging only variation in component A (wind direction changes). Column (4) shows the baseline result restricting the sample to 2010-2015, and column (5) shows the result using the same modified instrumental variable as in columns (2), but restricting the sample to the 2010-2015 period. Standard errors are clustered at the Copernicus grid cell level of the firm's headquarter. Coefficients (and standard errors) have been multiplied by 100 for readability. We denote * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Table A.7: Sensitivity of the results to the source of pollution data

	(1) Reanalysis Baseline	(2) Satellite -based	(3) Reanalysis 2011-2015	(4) Monitor -based inv. distance 2011-2015	(5) Monitor -based nearest 2011-2015
PM_{t-1}	-0.448***	-0.553***	-0.318***	-0.278***	-0.267***
	(0.108)	(0.162)	(0.0987)	(0.0923)	(0.0892)
N	9,411,781	9,411,781	6,797,845	6,693,030	6,489,037

Notes: Table reports the IV estimates of the effect of a one-unit increase in $PM_{2.5}$ at $t-1$ on the sales outcome at $t+1$ based on equation (7) for all firms in all sectors. All regressions include weather and holidays controls at $t-1$, t, and $t+1$, instrumented pollution at t and $t+1$, firm-by-year fixed effects, quarter-by-county fixed effects and industry-by-month-by year fixed effects. Column (3) shows the baseline result when the sample is restricted to the 2011-2015 period to enable comparison with columns (4) and (5), where the $PM_{2.5}$ data is based on observations from monitoring stations only, which are more readily available starting in 2011. The median distance to the nearest $PM_{2.5}$ monitoring station is 16 kilometers and the average is 25 kilometers. Standard errors are clustered at the Copernicus grid cell level of the firm's headquarter. Coefficients (and standard errors) have been multiplied by 100 for readability. We denote * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

B Robustness checks for the results on absenteeism

We first present the main results and then perform the same set of robustness checks as for the effect on sales to validate the evidence of a causal effect of $PM_{2.5}$ concentrations on sick leave episodes.

Column (1) of table B.2 shows the baseline estimate for the specification at the establishment level (same as column (2) of table B.1). Column (2) shows that a one-unit increase in the AQI index increases the number of workers entering sick leave that month by 2.1 per 1,000 workers. The effect in terms of SD increase is 0.85, while the effect of a one-SD increase in $PM_{2.5}$ is 0.93, a similar order of magnitude. Columns (3) to (6) show that the estimated effect of $PM_{2.5}$ on the number of workers starting a sick leave is robust to discarding months with PM_{10} alerts, winsorizing the absenteeism outcome, changing the specification of weather controls and controlling for flu incidence.

Column (2) of table B.3 shows that the order of magnitude of the effect holds if we use satellite-derived $PM_{2.5}$ data instead of reanalysis $PM_{2.5}$ data. Columns (3) and (4) show that the results are the same if we use monitoring station data only (restricting the period of analysis to 2011-2015).

Table B.1: The contemporaneous effect of $PM_{2.5}$ on sick leave (per 1,000 workers), all sectors

	OLS (1)	IV (2)
Panel A: Main Establishment Panel		
PM_t	0.0703***	0.147**
	(0.0212)	(0.0603)
N	8,238,888	8,238,888
R-squared	0.0637	0.0637
Dep. var. mean	23	23
First-stage effective F-statistic		306
Panel B: Aggregating data at the municipality level		
PM_t	0.0644**	0.148**
	(0.0208)	(0.0613)
N	369,190	369,190
R-squared	0.1602	0.1602
Dep. var. mean	23	23
First-stage effective F-statistic		268

Notes: Table reports OLS and IV estimates based on equation (8) for the effect of $PM_{2.5}$ at t on the number of workers starting a sick leave per 1,000 workers at the establishment level, either using the main sample at establishment level or one aggregated at the municipality level. Regressions for panel A include industry-by-month-of-sample, establishment, and quarter-by-county fixed effects, as well as weather and holidays controls. Observations are weighted by the number of workers observed in the Hygie dataset in each establishment. The effective F-statistic is based on a 2% random sample of single-establishment firms. Regressions for panel B include month-of-sample, municipality, and quarter-by-county fixed effects, as well as weather and holidays controls. Observations are weighted by the number of workers in each municipality. Standard errors are clustered at the Copernicus grid cell level. The effective F-statistic is based on a subsample of single-establishment firms aggregated at the municipality level. *p<0.10, **p<0.05, ***p<0.01.

Table B.2: The Effect of $PM_{2.5}$ on worker absenteeism, all sectors, robustness checks

	(1) Baseline	(2) AQI	(3) No AQ alerts	(4) Winsorized outcome	(5) Linear and quadratic weather controls	(6) Weather incl. humidity	(7) Flu incidence control
PM_t	0.147**		0.156**	0.157***	0.155**	0.117*	0.131**
	(0.0603)		(0.0650)	(0.0496)	(0.0611)	(0.0637)	(0.0604)
AQI index_t		2.151**					
		(0.868)					
N	8,238,888	8,238,888	7,890,564	8,238,888	8,238,888	8,238,887	8,238,888

Table reports IV estimates based on equation (8) for the effect of $PM_{2.5}$ on the number of workers starting a sick leave, per 1,000 workers. All regressions include industry-by-month-by-year fixed effects, quarter by county fixed effects, establishment fixed effects, weather and holidays controls. Observations are weighted by the number of workers for which we observe sick leave status in each establishment. Standard errors in parentheses are clustered at the Copernicus grid cell level. We denote * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

Table B.3: Sensitivity of the results on sick leave to the source of pollution data

	(1) Reanalysis Baseline	(2) Satellite -based	(3) Reanalysis 2011-2015	(4) Monitor -based inv. distance 2011-2015
PM_t	0.147^{**}	0.191^{**}	0.189^{***}	0.189^{***}
	(0.0603)	(0.0818)	(0.0605)	(0.0612)
N	8,238,888	8,238,888	5,796,540	5,796,540

Notes: Table reports the IV estimates of the effect of a one-unit increase in $PM_{2.5}$ at t on the number of workers starting a sick leave at t per 1,000 workers, based on equation (8). All regressions include industry-by-month-by-year fixed effects, quarter by county fixed effects, establishment fixed effects, weather controls, and holidays controls. Observations are weighted by the number of workers for which we observe sick leave status in each establishment. Standard errors in parentheses are clustered at the Copernicus grid cell level. We denote * for $p < 0.10$, ** for $p < 0.05$, and *** for $p < 0.01$.

C Data Appendix

C.1 Sick Leave Episodes

We obtain data on sick leave episodes (SLE) from the Hygie dataset, which follows approximately 900,000 employees during the period 2009-2015. The Hygie dataset 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).

The main subsample of interest is based on an exact match between the firm identifier of the establishment where the worker is employed and observing that firm in the sales data. When we instead consider the representative sample of private sector employees (such as in figure A.15), we make three restrictions to the Hygie sample. First, we only keep individuals to whom we are able to assign a place of work based on the establishment's unique identifier. This makes us discard individuals with no employment history declared between 2009 and 2015, who represent 25% of the sample. Although we cannot check the exact reason for missing information, these individuals are probably retired, unemployed or out of the labor force over the whole period. Two-thirds of them should be retired in 2009 given their age. We also discard 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 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 they are employed by private individuals).

Second, we discard individuals whose establishment identifier corresponds to a public institution such as hospital or schools, because we want to focus the analysis on private sector employees. Some individuals working in these institutions have a private sector type of contract and are thus eligible to enter the Hygie sample. Third, we discard a few individuals who did not work enough to contribute to the public pension system for any of the years included in the period. 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 labor supply, they are unlikely to experience sick leave episodes.

We assign each worker to the municipality of her workplace (there are around 6,000 municipalities in France). Figure C.1 shows the geographic distribution of the employees' workplaces in 2009, which is consistent with the distribution of the French population across the territory.

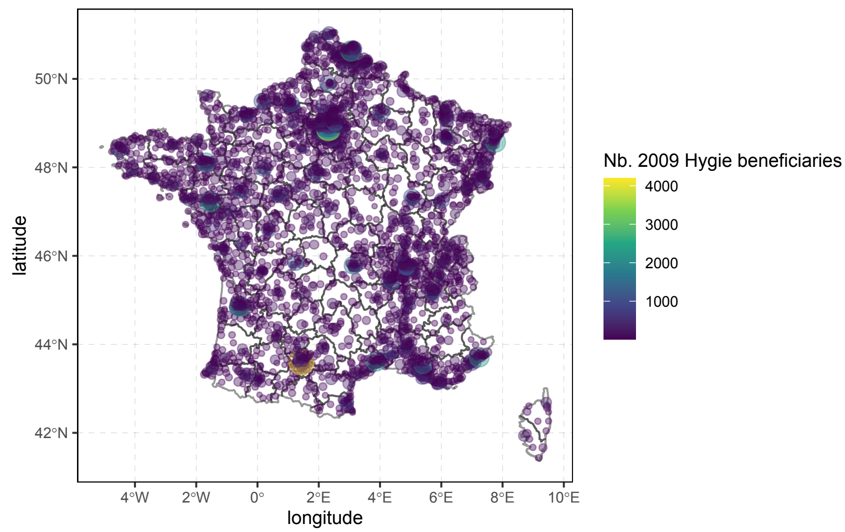

Figure C.1: Location of workers from the Hygie dataset based on the workplace municipality, in 2009

We use the exhaustive matched employer-employee data (DADS-Postes) to compare the characteristics of our representative sample of workers to the characteristics of the whole population of private sector employees. Applying the same restrictions as in the Hygie dataset,[48] 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. Thus, the average individual in our final worker sample – as shown in Table A.1 – is very close to the average private sector employee.

[48]Namely, we keep private sector employees born between 1935 and 1989, less 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.

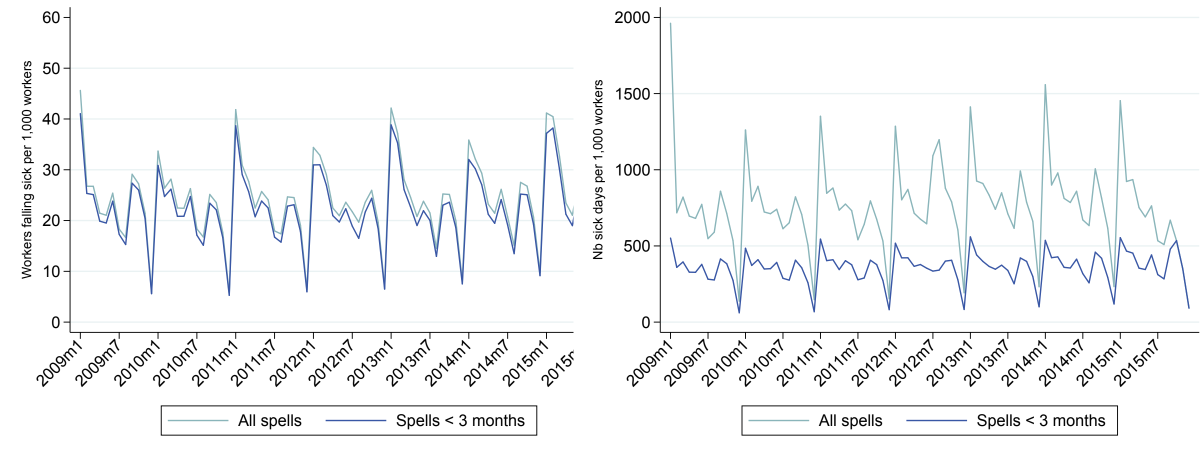

(a) Number of workers falling sick

(b) Number of sick days

Figure C.2: Number of workers falling sick and number of sick days per 1,000 workers

Notes: Figure presents the average number of workers falling sick and average number of sick days per 1,000 workers over time. While the spells larger than 3 months represent a small proportion of total spells, their tend to strongly increase the average number of sick days.

C.2 Firm-Level Sales

We compute firms' monthly sales by adding up different components included in the VAT records, following the methodology of France Stratégie and Inspection générale des Finances (2021). In the raw data, total sales are broken down into different components based on two main criteria that determine VAT liability: the location of the buyer (whether in France, in another EU country, or in a non EU-country) and whether the buyer is herself liable to VAT. In addition, the sales value of goods and services subject to specific tax rules is reported separately.[49] Our measure of sales includes both domestic sales and exports to EU and non-EU countries. The French tax administration imposes monthly declarations to firms with annual sales above €818,000 for the manufacturing sector and the hospitality industry and to those with annual sales above €247,000 for the other sectors. Firms below this threshold are allowed to fill declarations on a quarterly basis.

We discard the entire firm-year series for firms not reporting sales each month within a year. However, we make one exception for zero sales records in July since it is a relatively common pattern in the data. A large number of French firms close for vacation during some weeks in August, the month where the July VAT declaration is expected (the VAT declaration corresponding to the business month t is typically made on month $t + 1$). French tax authorities allow firms to report their July sales together with the August sales.[50] We indeed observe in the data that when the sales are 0 in July, the sales for August are frequently twice as high as those in June or September.

[49]For instance, the sales of natural gas and electricity is subject to a specific VAT rule in the French tax code, so they have their own subcomponent in the VAT records. See https://www.impots.gouv.fr/sites/default/files/formulaires/3310-ca3-sd/2022/3310-ca3-sd_3947.pdf

[50]See https://shorturl.at/TAZjH.

We re-allocate sales for July and August by splitting August sales in two.

We determine sectors of activity based on the sectoral classification available at the establishment level and we use the mode of sector categories across establishments for multi-establishment firms. We define 4 sectors of interest: manufacturing, construction, business-to-consumer retail and services, and business-to-business services. We discard firms belonging to the financial services sector, to the health, education and charitable sectors, which are often not-for-profit, as well as business-to-consumer services for which the timing and location of sales is often disconnected from the timing and location of consumption: hotels and transportation activities.

We check the quality of the reported data in two different ways. First, for a few large French companies for which annual financial reports are publicly available, we manually check that the sum of monthly sales of a given year is close to the official annual sales value. Second, we compare the time series of monthly sales value aggregated by economic sector to the data published at the industry level by the French statistical institute, using the same source. Figure C.3c shows the time series of monthly sales in construction (C.3a), manufacturing (C.3b) and all services (C.3c) as constructed from the VAT micro-data compared with the INSEE index. Differences may arise between our sales value and the statistical agency's because of different choices in data cleaning or the subcomponents entering the sales variable, but the correlation between the two series are above 0.9 for the three broad sectors.

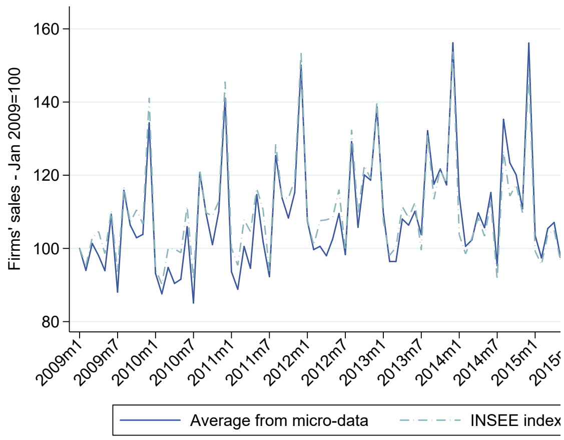

(a) Construction

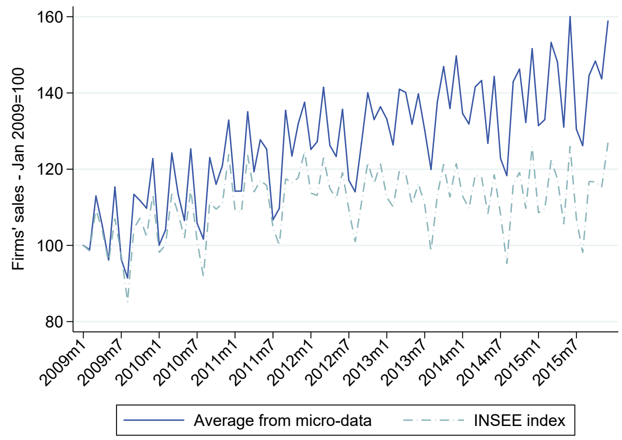

(b) Manufacturing

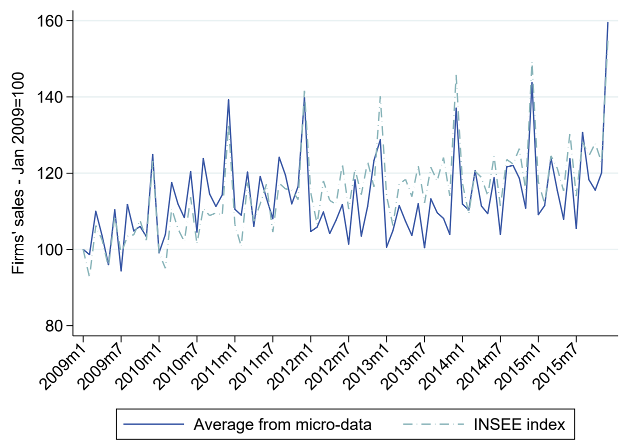

(c) Services

Figure C.3: Average firms' nominal sales in construction, manufacturing and service sector, 2009=100

Notes: Figure presents the average nominal sales from our VAT micro-data in blue for construction, manufacturing, and services and the INSEE sales index in dashed green, using January 2009 as the reference point. We exclude several service industries (trade - sector G in NACE classification, banking - sector K and health - sector Q) to compare with the INSEE index which also excludes these industries.