Departament d’Economia Aplicada
Universitat Autònoma de Barcelona

WORKING PAPER 25-07

Causal Effects of Air Pollution on Child Health: Evidence from a Low-Pollution Setting

Toni Mora
Manuel Flores
David Roche

Departament d’Economia Aplicada
Edifici B
Universitat Autònoma de Barcelona
08193 Bellaterra, Barcelona, Spain.
d.econ.aplicada@uab.cat
www.uab.cat/en/applied-economics
@ecapuab.bsky.social

UAB
Universitat Autònoma
de Barcelona

Causal Effects of Air Pollution on Child Health: Evidence from a Low-Pollution Setting*

Toni Mora†, Manuel Flores‡ and David Roche§

July 2025

Abstract

To what extent does air pollution in low-pollution settings affect children's health? Which children benefit most from further reductions, and what factors moderate this relationship? We address these questions using the universe of administrative medical records from the universal public healthcare system in Catalonia (Spain) between 2013 and 2017. We combine these data with spatio-temporal kriging techniques to construct complete time-by-location data on several air pollutants and environmental confounders. We then instrument for local PM_{10} concentrations—the main reference pollutant for air quality policies at the time—using variation in local wind direction in a multiple fixed effects model. Our primary outcome is respiratory-related healthcare visits, a measure of child morbidity. We find that even at relatively low ambient levels, increases in PM_{10} concentrations raise the incidence of respiratory-related visits. Our preferred instrumental variables estimate indicates that a 1 $\mu g/m^3$ (or 4.5%) increase in PM_{10} leads to a 0.5% increase in overall respiratory visits, driven mainly by lower respiratory illnesses, which carry more serious health implications than other respiratory illnesses. We also find evidence of heterogeneous effects, with the youngest children (ages 0–5) and those exposed during hot or drier months being most affected. We estimate that the observed decline in PM_{10} concentrations during our sample period may have prevented approximately 16 million respiratory-related visits and saved around €800 million in direct healthcare costs. The results highlight the value of targeted public health interventions, particularly for young children and during periods of elevated environmental risk.

Keywords: air pollution; low-pollution setting; healthcare use; respiratory illnesses; children; spatio-temporal kriging; instrumental variables

JEL classification: I12, Q53

---

* We would like to thank María Bermúdez Ramos, Manuel F. Grela, Andrea Ichino, Almudena Sevilla, and participants at EuHEA 2024, 4th CESC, 3rd Applied Economics Conference, 46th SAEe, 41st AES, XI EvaluAES Workshop, and at seminars at the University of Manchester (HOPE), University of York (CHE), Universidad Pública de Navarra, and Universidade de Santiago de Compostela (DRIE) for their helpful comments and discussions. We also gratefully acknowledge financial support from the Ministry of Science and Innovation, grant PID2021-124067OB-C21.

† Research Institute for Evaluation and Public Policies (IRAPP), Universitat Internacional de Catalunya, Immaculada, 22, 08017 Barcelona (Spain). E-mail: tmora@uic.es (corresponding author).

‡ Serra Hunter Fellow, Department of Applied Economics, Universitat Autònoma de Barcelona. E-mail: manuel.flores@uab.cat.

§ Research Institute for Evaluation and Public Policies (IRAPP), Universitat Internacional de Catalunya. E-mail: droche@uic.es.

1. Introduction

Despite ongoing efforts to reduce air pollution, it remains a persistent concern for policymakers in both developed and developing countries. For instance, the European Union (EU) launched the European Green Deal in 2019, aiming to make Europe the first climate-neutral continent. Similarly, in 2020, India introduced the Bharat Stage VI emission standards—equivalent to Euro 6—to curb vehicular emissions.[1] A few recent studies in developed countries have shown that air pollution negatively affects health even at levels below current legal limits, particularly among children (Ferro et al., 2024; Jans et al., 2018; Schlenker & Walker, 2016; Simeonova et al., 2021). However, evidence on the health benefits of further reducing ambient pollution in already low-pollution settings remains limited. Moreover, it is unclear which children would benefit most from such reductions and what factors may moderate these health effects.

This paper presents new evidence on the contemporaneous and heterogeneous *causal* effects of air pollution concentrations on child morbidity in a low-pollution setting. Our analysis encompasses Catalonia, Spain's second-largest region, with a population of approximately eight million people.[2] Our main focus is on particulate matter (PM)—the mixture of solid and liquid particles suspended in the air, where the major components are sulfate, nitrates, ammonia, sodium chloride, carbon, mineral dust and water—and specifically on "coarse" particles with an aerodynamic diameter of 10 micrometers or less (PM_{10}). "Finer" $PM_{2.5}$ particles, generally considered more hazardous due to their ability to penetrate deeper into the lungs and bloodstream, and indoors, than "coarser" particles (EPA, 2024; WHO, 2021), were not systematically measured during our sample period (2013-2017).

[1] For the EU see https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/european-green-deal_en, and for India, see https://morth.nic.in/gsr-308e-regarding-emission-standards-bharat-stage-vi-bs-vi-quadricycle.

[2] Catalonia has a substantially greater population than most Nordic countries (e.g., Denmark, Finland and Norway).

Moreover, PM_{10} was, at the time, the primary reference pollutant for air quality policies in the EU. Nevertheless, we also collected data on other air pollutants (carbon monoxide (CO), nitrogen dioxide (NO_2), and sulfur dioxide (SO_2)).[3]

Children are a particularly sensitive group to air pollution due to biological and behavioral reasons (EPA, 2024; WHO, 2005). First, the immune system and lungs continue to develop during the early years of childhood. Second, children have relatively high breathing rates compared to adults, which may potentially result in a higher intake of air pollutants per unit of body weight. Third, kids typically spend more time outdoors than adults, particularly in the summer and late afternoon, when the concentrations (and impacts) of air pollution are generally higher. Moreover, as a recent literature shows, common "mild" adverse health shocks early in life, such as exposure to air pollution—and not just the less common "severe" shocks, such as the experience of famines—can have permanent adverse effects on later-life health (Almond et al., 2018).

Air pollution concentration is likely an endogenous regressor in health or healthcare use regressions due to issues such as non-random sorting, avoidance behavior, and measurement error.[4] To address these issues, earlier studies closely related to ours have used so-called multiple fixed effects (FE) models (e.g., Ferro et al. 2024; Janke, 2014; Neidell, 2004), which include local unit FE to control for residential sorting across these local units. Others, besides multiple FE, have relied on arguably exogenous weather regressors such as

[3] Spain (including Catalonia) follows the air quality standards set by the EU-directive 2008/50/EC, which sets limit values for various polluters, including the ones we study. For instance, for PM_{10} there are limit values for short-term (daily) exposure of 50 micrograms per cubic meter ($\mu g/m^3$), not to be exceeded more than 35 times in a calendar year, and for long-term (annual) exposure of 40 $\mu g/m^3$ (see https://environment.ec.europa.eu/topics/air/air-quality/eu-air-quality-standards_en). Limits recommended by the World Health Organization (WHO) are more stringent (e.g. 20 $\mu g/m^3$ for annual exposure to PM_{10}, see WHO (2006)). No legal limits for $PM_{2.5}$ were in place in the EU before 2015.

[4] Avoidance behavior and classical forms of measurement error have typically led to "significant downward bias in conventional dose-response estimates" (Schlenker & Walker, 2016, p.771). For reasons outlined below in this section, our "pollutions maps" are likely less prone to measurement error than those in typical Economics studies. Moreover, Aguilar-Gomez et al. (2022) argue that avoidance and mitigation behaviors will make pollution *exposure* endogenous, but "do not need to be included to estimate the concentration-response function properly", which is our goal in this paper.

thermal inversions (Jans et al., 2018) or on policy changes (e.g., Simeonova et al. (2021) exploit the introduction of a congestion tax in Stockholm). A final set of studies has implemented an instrumental variables (IV) approach using different types and combinations of weather variables to instrument for air pollution. For instance, Knittel et al. (2016) use traffic shocks interacted with weather conditions as IV for air pollution (PM_{10} and CO) to estimate its causal effect on infant mortality rates. We also employ an IV approach, constructing IV based on monthly changes in local wind direction (cf. Deryugina et al., 2019; Graff Zivin et al., 2023).

This paper contributes to the growing body of evidence on the detrimental health effects in children of air pollution within currently permitted legal levels, suggesting that direct exposure to even relatively low levels of air pollution has contemporaneous detrimental effects on child morbidity. Among its contributions, the first and most general one is to study the so-called "missing middle" (Almond et al., 2018; Currie, 2020). While a well-established literature exists on the long-lasting effects of early life conditions, most studies tend to focus on the fetal period and young children, leaving adolescence relatively understudied (cf. Flores & Wolfe, 2023). This applies as well to the studies more closely related to ours—those assessing the *causal* health effects of air pollution concentrations—which mainly focus on the prenatal and early childhood period (e.g. Arceo et al., 2016; Cesur et al., 2017; Chay & Greenstone, 2003; Currie & Neidell, 2005; Currie et al., 2009; Currie & Schwandt, 2016; Currie & Walker, 2011; Jayachandran, 2009; Knittel et al., 2016; Palma et al., 2022), with only a few covering the entire or most of the childhood period (Ferro et al., 2024; Janke, 2014; Jans et al., 2018; Neidell, 2004, 2009; Schlenker & Walker, 2016; Simeonova et al., 2021).[5]

[5] Schlenker & Walker (2016) and Simeonova et al. (2021) focus on young children but show results for older children separately. There is also an extensive epidemiological literature studying the *associations* between poor

A second and more specific contribution relates to our outcome variable. A large body of this literature has investigated the effects of pollution on either measures of health at birth, such as the incidence of low birth weight and prematurity (Currie et al., 2009; Currie & Schwandt, 2016; Currie & Walker, 2011; Palma et al., 2022) or child mortality (Arceo et al., 2016; Cesur et al., 2017; Chay & Greenstone, 2003; Currie & Neidell, 2005; Currie et al., 2009; Jayachandran, 2009; Knittel et al., 2016),[6] but less attention has been paid to child morbidity, which is arguably a "more common and sensitive indicator" (Simeonova et al., 2021). A few studies in this context (including ours) use healthcare visits driven by respiratory illnesses as a measure of child morbidity (Janke, 2014; Jans et al., 2018; Neidell, 2004, 2009; Simeonova et al., 2021). One issue here is what type of health providers' health records one uses. We include children's primary care visits, emergency room (ER) visits, and hospitalizations (inpatient visits) due to respiratory illnesses, using a universe of administrative public healthcare records. Other studies use hospitalizations *only* (e.g., Ferro et al., 2024; Janke, 2014; Neidell, 2004, 2009). Adding primary care visits is essential for measuring the short-term effects of air pollution levels in low-pollution settings, as most healthcare visits related to respiratory illnesses do not require a hospitalization (unless driven by, e.g., acute asthma attacks, as in Simeonova et al., 2021). Instead, we do not have access to specialist visits (outpatient visits). However, this is less problematic in settings like ours—the Spanish public healthcare system—which follows a gatekeeping model where General Practitioners (GPs) are the first point of contact for patients and are responsible for referring them to specialists when necessary.[7] Finally, and not done by these previous studies,

air quality indicators and health events in children (e.g. Nhung et al., 2017; Thurston et al., 2017). Here, we restrict ourselves to studies exploring the *causal* effects of air pollution concentrations on children's health.

[6] See Almond et al. (2018) for a review.

[7] In this system, patients typically must see a GP before accessing specialized care. By using healthcare visits without specialist visits as our healthcare use measure, we avoid potential double counting of consultations for the same illness—since a specialist visit usually follows an initial GP consultation rather than being an independent event. This approach, however, does not account for total medical costs, which would include both GP and specialist visits (besides ER visits and hospitalizations). Our analysis thus focuses on healthcare utilization rather than on total economic burden.

we consider the severity of the respiratory disease by studying the effects of air pollution concentrations on both upper (less severe) and lower (more severe) respiratory illnesses, which, as our results show, is important.[8]

Our third contribution is more methodological and relates to our main variable of interest, PM_{10}, the other air pollutants (CO, NO_2, SO_2), and the controls that account for environmental confounding (weather variables). A unique feature of these data is the significant number of missing observations in the database, resulting from a misalignment between primary care units (our unit of analysis) and the available pollution (or weather) stations. To address this issue, we employ spatio-temporal kriging techniques to generate pollution (or weather) maps over time. This is common in Environmental studies (for instance, Ferro et al. (2024) and Stafoggia et al. (2019) use a random forest spatiotemporal model to predict ground level PM_{10} concentration levels), but less so in Economics, as most previous studies typically use some weighting procedures by distance, that do not explicitly take time into account for the interpolation (this includes all the previous studies that we discuss in Section 2; see Aguilar-Gomez et al. (2022) for more details). Thus, we expect measurement error to be less of an issue in our pollution and weather maps.

Our last contribution is to investigate the heterogeneity in the *causal* effect of air pollution on children's (respiratory) health. First, we examine heterogeneity by observable characteristics, such as the child's sex and age, as well as the underlying child health and socioeconomic status of the areas in which they live. Second, we assess heterogeneity by specific weather conditions. Our heterogeneity analyses are particularly relevant for understanding how weather conditions moderate the effects of air pollution on child health

[8] Some previous studies have examined (acute) asthma attacks (Neidell, 2004, 2009; Schlenker & Walker, 2016; Simeonova et al., 2021) or acute respiratory illnesses (Schlenker & Walker, 2016). A few have considered *all* healthcare visits, including both inpatient and outpatient visits (Jans et al., 2018; Simeonova et al., 2021). In our data, we cannot distinguish between planned and unplanned visits, though this is likely more relevant when analyzing acute outpatient visits (as in Simeonova et al. (2021)).

in low-pollution settings, as well as for informing the design of targeted health prevention and information policies for vulnerable groups.

In summary, this paper examines the impact of air pollution on children's health in a low-pollution setting and how this effect differ across various groups of children, geographic areas, and environmental conditions. It therefore uses the universe of administrative medical records from the universal public healthcare system in a large Spanish region (Catalonia) between 2013 and 2017, along with spatio-temporal kriging techniques to construct complete time-by-location data on several air pollution measures and other potential environmental confounders. It then instruments for local PM_{10} concentrations—the main reference pollutant for EU air quality policies at the time—using variation in local wind direction in a multiple FE model. The results show that even at relatively low ambient levels, increases in PM_{10} concentrations lead to a rise in the incidence of respiratory-related healthcare visits. The preferred IV estimate indicates that a 1 $\mu g/m^3$ (or 4.5%) increase in local PM_{10} concentrations results in a 0.5% increase in overall respiratory visits among children. This effect is primarily driven by lower respiratory illnesses, which have more serious health implications than other (upper) respiratory illnesses. The uninstrumented FE estimates are only slightly smaller (0.3%) than their IV counterparts (0.5%), in line with the expectation of less measurement error bias in pollution maps obtained from spatio-temporal kriging. There is also evidence of heterogeneous effects, with the youngest children (ages 0–5) and those exposed during hot or drier months being most affected.

Our findings are informative for policymakers. We estimate that the *observed* reduction in average annual PM_{10} concentrations during our sample period, from 24.8 in 2013 to 20.4 $\mu g/m^3$ in 2017, may have prevented approximately 16 million respiratory-related healthcare visits, with an associated direct cost savings of around €800 million. The results from our heterogeneity analyses highlight the value of targeted public health interventions, particularly for young children and during periods of elevated environmental risk. Moreover, the health benefits from these interventions for these children may extend into the long term (Almond et al., 2018). Finally, because the analysis focuses on PM_{10}, the estimated health impact may be conservative. For instance, finer particles like $PM_{2.5}$ can penetrate deeper into the body and may pose even greater health risks.

The rest of the paper is organized as follows. The second section discusses previous studies, with a focus on those most closely related to our work. The third section describes the various datasets used in this research. The fourth and fifth sections correspond to the empirical strategy and results sections, and the final section provides an overall discussion and concludes the paper.

2. Previous studies

Concerns on the harmful effects of air pollution on individuals' health date back at least to King Edward I of England, who in 1307 "banned the burning of coal to protect his citizens' bodily health" (Chay & Greenstone 2003, p. 1121). The literature on this topic is thus very broad, extending to disciplines beyond economics, such as environmental sciences and medicine. Here, we restrict our discussion of previous studies to those that try to isolate the causal effects of air pollution on children's health. Most have done so by employing a so-called multiple FE approach, with a few also adopting quasi-experimental research designs, such as the use of a policy change or IV. When explaining our empirical approach in section 4, we also discuss a few seminal papers that did not focus on children, but whose identification strategy we borrow.

More broadly, our findings contribute to the growing body of evidence on the harmful effects of air pollution on children's health—even at levels currently permitted by law. The results suggest that exposure to relatively low concentrations of air pollution has contemporaneous detrimental effects on child morbidity. These studies have proxied child morbidity using measures of respiratory illnesses (Ferro et al. 2024; Janke, 2014; Jans et al., 2018; Neidell, 2004, 2009; Schlenker & Walker, 2016; Simeonova et al., 2021).[9]

An early study related to ours is that by Neidell (2004), who investigated the impact of air pollution concentrations on child hospitalizations related to asthma episodes using zip-code-level panel data for California (USA) from the 1990s. Of the pollutants considered (CO, O_3, NO_2, and PM_{10}), only CO had a significant effect on asthma in children aged 1–18. Additionally, avoidance behavior—measured by smog alerts related to excess O_3 levels—reduced asthma-related hospitalizations. However, it left the coefficients measuring the effect of the pollution on childhood asthma relatively unchanged. Instead, these coefficients were greater for children from lower socio-economic status (SES) backgrounds, who also faced significantly higher exposure to pollution. This suggests that pollution may be one mechanism through which SES influences health.

In a follow-up paper, Neidell (2009) used a regression discontinuity design to compare outdoor activities on days just above and below the smog alert threshold for O_3 levels and found strong evidence of avoidance behavior in children aged 5–19. Moreover, estimates of the health effects of O_3 on hospitalizations due to asthma episodes were significantly larger when accounting for avoidance behavior than when this factor was omitted. Similar results were found among the elderly (65+), suggesting that both groups (children and the elderly) are particularly responsive to alerts and that studies omitting avoidance behavior likely underestimate the true health effects of pollution.

Janke (2014), who built on the work by Neidell (2009), used panel data at the English local authority level to estimate the effects of air pollution exposure on hospital emergency

[9] Schlenker & Walker (2016) also study heart-related issues.

admissions for respiratory diseases in children ages 5–19 during the mid-2000s. She found that a 1% increase in NO_2 or O_3 concentrations led to a slight increase in hospital admissions (0.1%). Moreover, she found that ignoring avoidance behavior did not lead to a statistically significant underestimation of air pollution's health effects.[10]

Another study by Jans et al. (2018) examined how air quality affects children from different SES groups in Sweden, a country with pollution levels and universal healthcare similar to England's. Inversion episodes—assumed to be exogenous—increased PM_{10} and NO_2 levels by 25% and 16%, respectively, and children's healthcare visits due to respiratory illnesses by 5.5%, with low-income children being particularly affected compared to their high-income peers. The study included *all* healthcare visits for respiratory illnesses from 2002 to 2007. Differences in baseline health appeared to be a key factor mediating the effect of pollution on the SES health gap. The effect of air pollution varied by illness type, with the largest impact on asthma—consistent with expectations, as asthma is highly sensitive to current exposure.

More recently, Simeonova et al. (2021) examined the effects of a congestion tax policy in central Stockholm (Sweden) on ambient air pollution (PM_{10} and NO_2) and children's health (acute asthma attacks) from 2004 to 2010. The tax was initially introduced as a trial in 2006 and made permanent in 2007. They focused on young children (ages 0–4), who have much higher rates of acute asthma attacks than older children, but reported results for the latter as well. Their health records included both hospital inpatient visits and acute (unplanned) outpatient visits, such as primary care, specialist, and ER visits that did not result in hospitalization. Ambient air pollution followed a stepwise decline, with a slight rebound during the temporary suspension of the congestion tax. In contrast, the rate of acute asthma

[10] One possible explanation is that, compared to California, O_3 levels in England rarely exceeded limit values, and public awareness of the health implications of poor air quality was limited. She also found no evidence that PM_{10} concentrations affected children's respiratory health, possibly due to measurement error problems.

attacks continuously declined from the trial period onward, possibly because some affected children never developed asthma in the first place. By 2010, ambient air pollution had declined by approximately 10–20%, while acute asthma attack rates had been cut in half. There were no comparable changes in accident rates or hospitalizations for non-respiratory conditions. Somewhat surprisingly, they also found a similar relative decline in acute asthma attacks among older children (see their Appendix Table 5).

Finally, Schlenker & Walker (2016) used an IV approach to estimate the contemporaneous (daily) local effects of air pollution caused by local airport congestion on ER visits and hospitalizations for respiratory illnesses (distinguishing asthma, acute, and overall conditions) and heart-related problems. Their study focused on areas within 10 km of airports in California (USA) between 2005 and 2007. Although they studied the overall population, they also presented results for children under the age of 5, those aged 5–19, and adults 65 and older.[11] Children aged 5–19 showed no sensitivity to pollution shocks; however, for those under 5, an increase in acute respiratory illnesses was linked to higher CO levels. Moreover, for these children, they found evidence that respiratory problems worsened disproportionately as pollution levels rose, indicating threshold effects and non-linearities in the pollution-health relationship.[12]

[11] At the local level, airport congestion (measured by airplane taxi time) and ambient pollution are influenced (i.e. confounded) by weather conditions. To address this issue, the authors instrumented for local airport congestion in California using network delays from the Eastern US coast. They also exploited the fact that wind speed and direction affect the transport of different air pollutants differently and included interactions between instrumented airplane taxi time, wind speed, and wind direction relative to California airports to isolate the effects of both CO and NO_2.

[12] Ferro et al. (2024) did not focus on the contemporaneous health effects of air pollution but on their medium-term impacts, up to age 10, in terms of hospitalizations and filled prescriptions. They specifically examined the effects of prenatal exposure to moderate air pollution. Their study included the universe of live births in Tuscany (Italy) between 2006 and 2018. Using multiple FE models, they showed that prenatal "exposure" to PM_{10} led to worse birth outcomes, more hospitalizations, and increased prescription use in the first 10 years of life, especially for children in the lowest health quantiles. For example, a one standard deviation (SD)-increase in PM_{10} concentrations during the prenatal period was associated with a 12% rise in hospitalizations and a 2% increase in prescribed medications during the first 10 years of life. Like our setting—as well as those in Janke (2014) and Jans et al. (2018)—theirs is characterized by pollution levels that generally comply with strict air quality regulations and a well-functioning universal healthcare system.

In this paper, we also employ an IV approach to estimate the population dose-response relationship between ambient pollution and child health outcomes, focusing on primary care, ER visits, and hospitalizations for respiratory illnesses. We follow the small IV literature that uses changes in local weather conditions as instruments for air pollution. Most of these studies (e.g., Deryugina et al., 2019; Graff Zivin et al., 2023; Knittel et al., 2016) exploit the fact that "wind speed and direction transport individual pollutants in different ways" (Schlenker and Walker, 2016, p.770).[13] We provide more details in Section 4.

3. Data and descriptive evidence

This study draws on data from multiple sources. Lacking individual residential addresses, we proxy each child's location by their assigned primary care unit (PCU). We thus define our unit of analysis as the PCU–month–year and construct a balanced panel accordingly. All air pollutants, weather variables, and other controls are measured at that same level.

Health data

To construct our health database, we start from the universe of administrative medical records from the universal public healthcare system for children in Catalonia, a large Spanish region, covering the period from January 2013 to December 2017. These records are maintained by the statutory National Health Service (NHS) and are provided by the Agency for Health Quality and Assessment of Catalonia (*AQuAS*). We have access to all doctor visits to primary care and ER, as well as all hospitalizations where the medically certified primary

[13] Graff Zivin et al. (2023) examined the effects of contemporaneous air pollution concentrations, measured by the Air Quality Index (AQI), on influenza hospitalizations across U.S. counties in 21 states from 2007 to 2017. While their focus was on the joint effect of air pollution and influenza on the overall population, they also reported that children under age 9 were disproportionately affected by air pollution compared to other age groups.

or secondary diagnosis of the visit (or the hospital stay) is a respiratory illness. Diagnostics for respiratory illnesses were established using records based on medically certified diagnoses coded according to the International Classification of Diseases, Ninth Revision, Clinical Modification (ICD-9-CM). Our focus is on visits whose *primary* diagnosis code is a lower respiratory illness (such as asthma, bronchitis, COPD, and pneumonia) or an upper respiratory illness (including rhinitis, nasal polyps, sinusitis, laryngotracheitis, and certain types of apneas). Since NHS coverage is universal, all residents of Catalonia—including all children—are potentially included in this database. According to our calculations, 99.1% of Catalonia's 7,555,830 residents in 2017 (IDESCAT, the Statistical Institute of Catalonia) were covered by the NHS.

While virtually all children are potentially covered by public health insurance, our analysis comprises only those who use public healthcare services. We include children with a primary diagnosis for any of the abovementioned respiratory conditions in any of the three earlier mentioned care levels (primary care, hospitalization, and ER) at any point in time during our five-year sample period (2013–2017). Visits to a specialist are not available from *AQuAS* for this period. But including these visits would probably introduce a double-counting of consultations in systems like ours, where GPs serve as "gatekeepers". Instead, adding primary care visits (which, as discussed in Section 2, many previous studies do not do) is important for measuring the overall short-term effects of air pollution levels in a low-pollution setting like ours. In total, we have comprehensive healthcare use data for nearly 88% of all children aged 0–17 years in Catalonia in 2017 (1,225,406 out of 1,399,850 children). For each child, we have information on all their healthcare visits to primary care, all hospitalizations, and all ER attendances related to a respiratory illness during our five-year sample period in the NHS in Catalonia. For each visit, we know the patient identifier (which is unique), the date of registry in the healthcare provider unit (primary, hospital, or ER care), and the medical diagnosis (ICD code) of the respiratory illness. The personal identifiers are used to link the health data files across healthcare provider units, as well as to an external source from *AQuAS* which includes additional socio-demographic information such as the child's gender, age, drug co-payment level (which is related to the SES of the household), nationality, date of death if not alive and health sanitary region (*Àrees Bàsiques de Salut*, ABS). The data also include an index of child comorbidity referred to as GMA (Adjusted Morbidity Groups, see Dueñas-Espín et al. (2016)). Since the GMA was provided for only the last year (2017), we exclude it from our regressions to avoid issues with "bad" controls. Instead, we use this variable for heterogeneity analyses, acknowledging that it may be affected by air pollution (as explained below, we divided the PCUs into quartiles based on their average GMA index in 2017 to perform heterogeneity analyses on the effects of air pollution on child health across areas, i.e., PCUs, with different grades of child comorbidity, see section 5).

Lacking individual residential addresses, we proxy each child's location by their assigned PCU and define our (spatial) unit of analysis at this level. Our database has 503 PCUs (the number of ABS is around 370 during our sample period). For all children assigned to a given PCU, we sum the number of visits to their GP, hospitalizations, and ER visits related to respiratory illnesses in each month. To avoid potential double-counting of visits for the same illness episode, we only record hospital admissions and ER attendances that occur at least 15 days after the initial GP consultation. For each PCU, we then compute a monthly rate of visits for all lower, all upper, and the total number of respiratory illnesses. These incidence rates are relativized to the average number of children who used the corresponding PCU during the sample period at least once and multiplied by 10,000 to obtain the number of visits per 10,000 children. In addition, to avoid potential double-counting of children across PCUs, we assign each child to its most frequent PCU, which accounts for about 92% of all individual visits.[14] Our final database is a balanced panel with 29640

[14] Less than half (44%) of all children use only one PCU, and about 79% and 93% use two and three PCUs, respectively. This is partly because some families spend time in their second residences on weekends, and

monthly-year observations for 494 PCUs (we drop nine PCUs as they are not observed during the entire five-year sample period). Figure 1 shows the geographical distribution of these PCUs in Catalonia, along with the pollution and weather stations from which we obtain the initial measurements of the air pollutants and weather variables.

[Insert Figure 1 about here]

Air pollution, weather, and pollen data

The corresponding air pollution, weather, and pollen data are obtained from various networks of stations spread throughout Catalonia, a territory of 32,108 km^2. The air pollution and weather networks are illustrated in Figure 1. A common difficulty with this type of data is the high number of missing observations due to a misalignment between the unit of analysis (the PCU in our study) and the available measuring stations. To address this issue—as explained further below—we apply spatio-temporal kriging techniques to generate maps of air pollution, weather, and pollen over time.

The air pollution dataset includes all official station records available in Catalonia during our sample period, obtained from the Department of Territory and Sustainability of the Catalan Government.[15] In total, 81 such stations collect multiple pollutants across the region (see Figure 1). We focus on four key pollutants: PM_{10}, CO, SO_2, and NO_2. The original hourly data were first collapsed into daily and then into monthly averages. Because individuals typically live near their assigned PCU, we take the PCU as the relevant unit for measuring children's exposure to air pollution. To address the misalignment between PCUs

especially during vacation periods. For instance, according to INE (the Spanish National Statistics Institute), 13.5% of all Catalan families had a second residence in 2021 (https://www.ine.es/dynt3/inebase/index.htm?padre=8981&capsel=9032).

[15] See https://analisi.transparenciacatalunya.cat/Medi-Ambient/Qualitat-de-l-aire-als-punts-de-mesurament-autom-t/tasf-thgu.

and monitoring stations, we apply spatio-temporal kriging—a geostatistical interpolation method that predicts pollutant levels at unmonitored locations and times by exploiting the spatial and temporal autocorrelation observed in the data. We use the same approach to generate time-by-location maps for our weather and pollen variables, as discussed next.[16]

We considered an extensive set of weather variables to control both for the direct effects of weather on health (Deschênes et al., 2009) and to leverage the quasi-experimental features of changes in local wind direction in distributing air pollution within locations (cf. Graff Zivin et al, 2023) (for more details, see our empirical model in Section 4). We collected daily data on rain, temperature, wind, humidity, and atmospheric pressure from all available registers of the Meteorological Service of Catalonia, using their 207 meteorological stations that were available for our entire sample period (see Figure 1). We then computed monthly averages for minimum and maximum temperature, atmospheric pressure, humidity, and accumulated precipitation. As for the wind, we calculated its average speed and direction using the 8-point compass rose. Additionally, for each month in our sample period, we calculated the number of days with mist and fog, which have been shown to impact PM_{10} concentration levels (Querol et al., 2019). Mist and fog data were retrieved from a subset of 172 stations of this same network.

[16] Specifically, we implement ordinary kriging with a spatio-temporal variogram model to interpolate monthly pollutant concentrations for each PCU based on data from nearby monitoring stations. Kriging generates smooth prediction surfaces by accounting for spatial distance, correlation structures, and observation density. This method enables us to assign pollution, weather, and pollen values to all PCU–month–year combinations in our sample period, minimizing measurement error due to missing observations.
Spatio-temporal kriging relies on a fitted variogram that captures spatial, temporal, and interaction-based (i.e. spatio-temporal covariance) variation in the data (Stein, 1999). We use the sum-metric model, the most flexible specification, which includes components for spatial and temporal variation, as well as an interaction term capturing residual covariance not accounted for by the two. Each of these three components is characterized by three parameters: the partial sill (the level at which the variogram flattens), the nugget (which relates to the amount of short-range variability in the data), and the range (the distance beyond which observations are uncorrelated). The interaction term also includes a fourth parameter: the anisotropy ratio of the residuals, which reflects directional differences in correlation length.
To estimate these parameters, we use evolutionary computation techniques, specifically differential evolution with standard settings (Storn and Price, 1997). Further technical details are available upon request.

We also collected pollen information, as this may mediate the effect of air pollution on respiratory health. For instance, Schlenker & Walker (2016, p. 783) report that "A body that is weakened by the elevated pollen levels might be more (or less) susceptible to pollution shocks.". We included five types of pollen related to allergies in our setting: *Platanus*, *Poaceae*, *Cupressaceae*, *Cupressa*, and *Cupressace* (Pereira et al., 2006). These data were obtained for our entire sample period from the nine stations of the Catalan Network of Aerobiology (*Xarxa Aerobiològica de Catalunya*).[17]

Additional data on SES-related factors

To capture a SES-related factor, we created an income area-level measure at the postal code level. This variable is available for only one year (2013), which is why we use it to detect heterogeneous effects rather than as a control variable. To obtain this variable, we matched our PCUs with microdata corresponding to a random sample of declarants of the 2013 income tax provided by the Fiscal Studies Institute (https://sede.hacienda.gob.es/es-es/sedes/ief/paginas/ief). This tax file contains declarants from 1,112 postcodes for the region of Catalonia. We calculated the average tax base in each postcode and allocated that value as a proxy of the income of the PCU that was in that same postcode. We then divided the PCUs into quartiles based on their average tax base in 2013 to perform heterogeneity

[17] The airborne pollen and fungal spore data we use were provided by the Catalan Network of Aerobiology (*Xarxa Aerobiològica de Catalunya*, XAC, https://aerobiologia.cat) at *Universitat Autònoma de Barcelona* (UAB). These data have been obtained following standard procedures. Daily samples are obtained at the sampling sites using Hirst traps (Hirst, 1952). Samples are then transported to the XAC laboratory at the UAB, processed and analysed under a light microscope to recognise and count each of the pollen and spore particles. Daily results expressed as pollen and fungal spores concentrations (Pollen/m^3 and spores/m^3) are then converted into Monthly Integrals (MPIn expressed as Pollen*day/m^3 and MSIn as spores*day/m^3) by summing up the daily concentrations corresponding to each month and Annual Integrals (APIn expressed as Pollen*day/m^3 and ASIn as spores*day/m^3) by summing up the daily concentrations corresponding to each year (Galán et al., 2014).

analyses on the effects of air pollution on children's health across areas (PCUs) with different income levels.

Descriptive evidence

We start the descriptive analysis with our outcome variables. Figure 2 shows the evolution of the incidence rates of visits per 10,000 children for lower, upper, and total respiratory illnesses in Catalonia during our sample period. A few patterns are evident. First, there is a strong seasonality in the respiratory visits (they are lower in the summer and higher during the winter). Second, lower respiratory illnesses (which have more severe health implications) show a much higher incidence than upper respiratory illnesses. Finally, there is no clear trend in these incidence rates during our sample period.

[Insert Figure 2 about here]

Table 1 presents the overall average and median rates of visits for lower, upper, and total respiratory illnesses during our sample period. Descriptive statistics for air pollutants, weather, and other potential confounders, controls, or moderators are also included.

[Insert Table 1 about here]

Figure 3 focuses on our main variable of interest (PM_{10}) and the other air pollutants (NO_2, SO_2, and CO). The concentration of these air pollutants is expressed in micrograms (one-millionth of a gram) per cubic meter of air, or $\mu g/m^3$, except for CO, which is measured in milligrams per cubic meter (mg/m^3). We highlight the following patterns. First, the figure illustrates that our setting, Catalonia (Spain), is representative of an area with low ambient pollution levels, which comply with relatively strict pollution regulations. For instance, PM_{10} does, on average, never exceed the EU's legal limit of 40 $\mu g/m^3$ for annual exposure, but it typically exceeds the WHO's safe concentration level of 20 μg/m³ for annual exposure. Second, during our sample period, we see a slightly negative trend in three out of four air pollutants (PM_{10}, NO_2, and SO_2), with the other (CO) being relatively stable. For instance, average annual PM_{10} concentrations decreased from 24.8 μg/m³ in 2013 to 20.4 μg/m³ in 2017. Finally, some air pollutants also exhibit a seasonal pattern (like the incidence rates of visits). For instance, PM_{10} typically peaks during the winter months.

[Insert Figure 3 about here]

The last figure in this section helps illustrate the type of variation that we use to construct our IV. Figure 4 presents polar graphs for four different locations, based on historical data on PM_{10} concentrations, wind direction, and wind speed. In each location, PM_{10} concentration varies according to wind patterns, but in different ways. This variation depends on factors such as the presence of internal sources of PM_{10} emissions, the location's proximity and orientation relative to external sources, and the local orography. For example, in the city of Barcelona (top-left panel), which generally experiences high pollution levels, PM_{10} concentration is higher when winds blow from the east at relatively high speeds. This is intuitive, as the city has a large port on its eastern side, substantial sea traffic, and a mountain range to the west. In contrast, the PM_{10} level is much lower when winds blow from the northwest at similarly high speeds.

Since individuals may sort themselves based on prevailing wind patterns, we do not use these levels directly as IV for PM_{10}. Instead, we rely on monthly changes in wind direction as our source of variation that shifts pollution, net of average pollution in that location (Graff Zivin et al. 2023). Further details are provided in the next section.

[Insert Figure 4 about here]

4. Empirical strategy

We assessed the impact of air pollution on the incidence of visits using the following specification estimated on PCU (*p*)-by-month (*m*)-by-year (*y*) level data:[18]

$$visits_{pmy} = \alpha + \beta PM_{10,pmy} + \boldsymbol{Pollut'_{pmy}}\,\boldsymbol{\gamma} + \boldsymbol{W'_{pmy}}\,\boldsymbol{\delta} + \boldsymbol{X'_{pmy}}\,\boldsymbol{\lambda} + \theta_{pqy} + \mu_m + \varepsilon_{pmy} \tag{1}$$

where $visits_{pmy}$ is our outcome variable representing the rate of visits per 10,000 children related to the conditions we study in the local unit *p* in month *m* of year *y*. Our main parameter of interest is β, which measures the effect of PM_{10} concentrations on children's healthcare visits for respiratory illnesses. In our empirical analyses, we controlled for both correlated observable and unobservable effects defined at the PCU-month-year level. $\boldsymbol{Pollut}$ is a vector with other air pollutants (CO, SO_2 and NO_2) that are potentially correlated with PM_{10} and child morbidity (Aguilar-Gomez et al., 2022; Janke, 2014; Neidell, 2009; Schlenker & Walker, 2016). $\boldsymbol{W}$ is a vector of monthly controls that includes weather variables—namely, accumulated precipitation, maximum and minimum temperature, atmospheric pressure, humidity, wind speed, and the number of days affected by correspondingly mist and fog—as well as the presence of common types of pollen in our setting (*Platanus*, *Poaceae*, *Oleaceae*, *Urticaceae*, and *Cupressaceae*). Controlling for weather conditions is important. Extreme weather conditions directly affect health (Deschênes & Moretti, 2009) and moderate the impact of air pollutants on children's respiratory health (Knittel et al., 2016). Weather conditions are also correlated with time spent outdoors by children (Janke, 2014), influencing their actual exposure to air pollution concentration levels. $\boldsymbol{X}$ is a vector of individual-level controls (aggregated at the local-month-year level) referring to the child (age, gender, nationality, and mortality) or household (copayment level for prescription drugs). We also

[18] For instance, Simeonova et al. (2021) and Graff Zivin et al. (2023) also estimate their models on local level-monthly data. We use here primary care unit (PCU) and local unit as synonyms.

included PCU-by-quarter-by-year fixed effects (θ_{pqy}) to control for any remaining unobserved differences between our local units p. These differences were allowed to change every three months capturing thereby unobserved local seasonal patterns. For instance, the monitoring of some public healthcare services across the health sanitary regions in Catalonia is done quarterly.[19] Quarters might also represent the seasons that coincide with changes in health behaviors (e.g., time spent outdoors is lower during winter and increases in springtime as days become longer) or in economic activity. These changes were allowed to differ between local units (i.e., PCUs). Finally, we also included monthly dummies (μ_m) to control for any remaining seasonality and general monthly trends within each year in both respiratory health and air pollution concentrations. For instance, particulate matter and respiratory illnesses peak in winter months (see Figures 1 and 2); month fixed effects capture such seasonality. In robustness checks, we examined models using specifications with alternative sets of control variables.

We estimated Equation (1) by weighted least squares (WLS) using the number of children that are users of the PCU as weights and clustered standard errors at the PCU to account for arbitrary correlated errors within the PCU across time (in robustness checks, we adjusted our standard errors for two-way clustering across both PCUs and time). The coefficient β measures the short-run (i.e., monthly) health effect of a one-unit increase in PM_{10}. The identification assumption is that within-local unit monthly changes in PM_{10} are exogenous, after controlling for other air pollutants, weather and pollen variables in the same cell-level, quarter-by-year unobserved differences between local units, and (general) seasonality effects (identification, therefore, comes from changes in PM_{10} concentrations within each month of the year and within each quarter of the year within a local unit).

[19] See https://catsalut.gencat.cat/web/.content/minisite/catsalut/publicacions/docs/sistema-atencio-salut-catalunya-cast.pdf.

The least squares FE estimates, however, may still be prone to bias from time-varying unobserved confounding within each PCU-by-quarter-by-year cell, measurement error in pollution assignment, and defensive or avoidance behavior (e.g., Deryugina et al., 2019; Graff Zivin et al., 2023; Neidell, 2009; Schlenker & Walker, 2016). We, therefore, implemented an IV approach. Other studies closely related to ours (besides Schlenker and Walker (2016) and Graff Zivin et al. (2023), which we discussed in Section 2), have used weather conditions to construct instruments for air pollution. For instance, Knittel et al. (2016) use traffic shocks interacted with weather conditions as IV for air pollution (PM_{10} and CO) to estimate its causal effect on infant mortality rates. Our approach draws on Deryugina et al. (2019) and Graff Zivin et al. (2023) in that it exploits changes in wind direction to construct instruments for PM_{10}.[20] Here, we used the approach of Graff Zivin et al. (2023) because, compared to Deryugina et al. (2019), it preserves the number of instruments, while allowing pollution shocks to vary in magnitude and direction at the local level, depending on proximity to pollution sources and prevailing wind patterns.[21] To build our instrument, we first created a new variable, $\widetilde{PM10}^{q_p}$, which is PM_{10} concentrations in PCU p averaged over the entire sample period when the wind is blowing from direction q in unit p, demeaned by the average PM_{10} level in unit p (see Equation (2)):

$$\widetilde{PM10}^{q_p} = \frac{\sum_{i\in q_p} PM10_i^{q_p}}{\sum_{i\in q_p}} - \frac{\sum_{i\in p} PM10_i}{\sum_{i\in p}} \tag{2}$$

We then used $\widetilde{PM10}^{q_p}$ to generate a set of instruments $\boldsymbol{Z}_i^q$, where each instrument corresponds to a particular wind direction q if a particular observation i belongs to local unit p and the wind in that particular year-month in this local unit is blowing from q (see Equation (3)). Differently from

[20] Our IV approach will alleviate the bias from avoidance behavior if individuals do not correctly predict the impact of local *changes* in wind direction on local pollution levels. It is also well-known that IV only solves the bias from measurement error in the independent variable when the measurement error is classical, namely mean zero and i.i.d. (Griliches and Hausman, 1986).

[21] Allowing for location-and-wind direction specific impacts on air pollution implies that the number of instruments grows larger than the number of clusters (PCUs), leading to inefficient standard errors (see Graff Zivin et al. (2023) for more details).

Graff-Zivin et al. (2023) we used the eight quadrants as wind direction bins instead of the four quadrants.

$$\boldsymbol{Z}_i^q = \begin{Bmatrix} \widetilde{PM10}^{q_p} \; if \; wind_{dir}{}_i^q = q \;\&\; i \in p \\ 0, \quad\quad otherwise \end{Bmatrix} \tag{3}$$

Going back to our example in Figure 4: if the prevailing wind in a given month in Barcelona changes from the east to the northwest, our instrument would shift PM_{10} concentrations down for that month, net of average PM_{10} levels in Barcelona for that year-quarter, and conditional on our multiple FE and controls. Wind blowing from a particular direction can redistribute pollution locally or bring in external pollution—or cleaner air. The identifying assumption of our IV approach is that, after flexibly controlling for multiple FE, other air pollutants, weather and pollen variables, local changes in wind direction are unrelated to local incidence rates of respiratory illnesses, except through their effects on PM_{10} concentrations.[22] The specification of our first-stage equation is identical to that of Equation (1), including the vectors $\boldsymbol{Pollut}$, $\boldsymbol{W}$, $\boldsymbol{X}$, the fixed effects θ_{pqy} and μ_m, and the instruments $\boldsymbol{Z}$. In robustness checks, we again examined models using alternative specifications.

5. Results

Main Results

Table 2 presents our OLS results. We find that higher PM_{10} concentrations are associated with an increased incidence of respiratory illnesses among children. These associations are driven by lower respiratory illnesses—which are more severe than upper respiratory illnesses—and remain robust to the inclusion of multiple FE, other air pollutants, weather

[22] As noted in Deryugina et al. (2019) and Graff Zivin et al. (2023), identification does not use the prevailing wind direction, which would not change across time and could lead to sorting or strategic placement of air pollution monitors. Instead, our instrument uses month-to-month changes in wind patterns in a given PCU-by-year-by-quarter cell, which should not affect sorting or monitor placement.

and pollen variables, and sociodemographic controls. In our preferred specification (columns 4), a 1 $\mu g/m^3$ (or 4.5%) increase in PM_{10} is associated with approximately 0.7 additional visits for lower and total respiratory illnesses per 10,000 children, corresponding to a 0.3% and 0.4% increase, respectively.

[Insert Table 2 about here]

Table 3 shows the IV estimates of PM_{10} on children's respiratory illnesses. The first-stage regressions yield high *F*-statistics, suggesting weak instrument bias is not a concern in our setting. Wind direction is a strong predictor of PM_{10} concentrations, as indicated by the large first-stage *F*-statistics reported in the bottom panel of the table.[23] Importantly, our instruments also appear valid—that is, uncorrelated with the error term in Equation (1)—especially after accounting for multiple FE, other air pollutants, and environmental confounders. As shown in the bottom panel, the *p*-value of Hansen's *J*-statistic of overidentifying restrictions in columns (3)-(4) is around 0.23-0.28 in the models of low and total respiratory illnesses, so we cannot reject the validity of the model.

The IV estimates of PM_{10} remain robust across different specifications and, if anything, increase in magnitude when controlling for other air pollutants and environmental confounders. As with the OLS results, the overall effects are driven by the more severe conditions—the lower respiratory illnesses. In our preferred specification (column 4), a 1 $\mu g/m^3$ (or 4.5%) increase in PM_{10} leads to approximately 0.8 and 1 additional visits per 10,000 children for lower and total respiratory illnesses, respectively—equivalent to a 0.5% increase.

[23] Table 3 reports first-stage *F*-statistics assuming homoskedastic errors, namely the Cragg-Donald (CD) *F*-statistic, which allows for direct comparison with the Stock & Yogo (2005) critical values, valid under the assumption of homoskedasticity. However, since our empirical models allow for serial correlation and heteroskedasticity in the error structure, we also compute robust first-stage *F*-statistics—the Kleibergen-Paap rk Wald *F*-statistic (Kleibergen & Paap, 2006). Across all specifications, these robust statistics are consistently larger than the CD values, providing stronger evidence of instrument relevance under more general error structures.

Previous studies (e.g., Deryugina et al., 2019; Graff-Zivin et al., 2023; Schlenker & Walker, 2016) have reported substantially larger IV estimates compared to OLS, attributing this to attenuation bias from measurement error in FE models. In contrast, our IV estimates are only somewhat larger than the OLS ones—possibly because our use of spatio-temporal kriging techniques reduces measurement error in pollution assignment.

[Insert Table 3 about here]

Heterogeneity

Our analyses so far have relied on the assumption that the relationship between air pollution and children's health is homogeneous across the population. However, previous research suggests that children may differ in their susceptibility to air pollution (e.g., Jans et al., 2018; Knittel et al., 2016; Neidell, 2004; Schlenker & Walker, 2016), as well as in the preventive or avoidance behaviors of their parents (e.g., Janke, 2014; Neidell, 2009). For instance, Schlenker & Walker (2016) found that only the youngest children (under age 5) were sensitive to pollution shocks, while Jans et al. (2018) reported that low-income children were particularly affected by inversion episodes—likely due in part to differences in baseline health, which may mediate the relationship between pollution and the SES-health gradient. Additionally, the health effects of PM_{10} concentrations can vary with weather conditions, being more pronounced during hot periods, when environmental conditions exacerbate respiratory vulnerability (Knittel et al., 2016).

To explore the potential heterogeneity in the health effects of PM_{10}, we conducted two complementary analyses, the results of which are presented in Figure 5. First, in panels A to C, we examine heterogeneity by observable characteristics, such as the child's sex and age, as well as the underlying child health and socioeconomic status of the areas in which they live. Second, in panels D to F, we assess heterogeneity by specific weather conditions. In both cases, we split the sample according to the relevant characteristic or condition and re-estimate the health effects of PM_{10} using our IV approach.

Regarding children's characteristics, the effects on total respiratory illnesses in Panel C are primarily driven by younger children—particularly those aged 0–5, consistent with Schlenker & Walker (2016)—and are present for both boys and girls. When sorting PCUs by income quartiles, we do not find evidence of a clear socioeconomic gradient. The pattern of effects is U-shaped, with the largest point estimates observed in the lowest and highest income quartiles; however, the confidence intervals overlap, and average PM_{10} levels are similar across quartiles (approximately 22.3 $\mu g/m^3$). Sorting areas by quartiles of child comorbidity yields a similar result: overlapping confidence intervals. Still, the point estimates suggest weaker effects in areas with higher comorbidity, possibly reflecting greater avoidance or protective behavior among parents of children with pre-existing health conditions (average PM_{10} levels are slightly higher in these areas—around 22.6 $\mu g/m^3$ in GMA-Q3 and Q4, compared to 22.3 $\mu g/m^3$ in GMA-Q2 and 21.9 $\mu g/m^3$ in GMA-Q1). The patterns for lower respiratory illnesses, which are more harmful, closely mirror those for total respiratory illnesses (see Panel A). By contrast, we find no evidence of heterogeneity in the effects on upper respiratory illnesses across child or area characteristics; the estimates are consistently close to zero and not statistically significant at the 5% level (see Panel B).

Turning to weather conditions, our results suggest that temperature and humidity moderate the effects of PM_{10} concentrations on children's health. The point estimates on total respiratory illnesses in Panel F are largest during months with higher average maximum (≥20ºC) and minimum temperatures (≥10ºC), as well as during months with lower humidity levels (<70%).[24] Notably, when comparing months with average maximum temperatures above and below the 20ºC threshold, the confidence intervals barely overlap, indicating a significant difference in effect size. A similar pattern is observed for upper respiratory illnesses (see Panel E). For lower respiratory illnesses, the point estimates follow a similar trend, but the confidence intervals always overlap, so we cannot reject the null hypothesis of homogeneous effects across weather conditions (see Panel D). Finally, we find no evidence that fog or mist moderates the effects of PM_{10} on children's health.

[Insert Figure 5 about here]

Robustness

We conducted various robustness checks to corroborate the validity of our main findings. First, as shown in Tables 2 and 3, our results are robust to the inclusion of high-dimensional FE (PCU-by-quarter-by-year FE) and month FE, and to the inclusion and exclusion of weather and pollen variables, sociodemographic controls, and importantly also other air pollutants that affect health (CO, SO_2, NO_2) and co-vary with PM_{10} concentrations. Second, we tested the robustness of our main estimates to alternative clustering choices. In particular, we adjusted our standard errors for two-way clustering across both location (PCUs) and time (month-by-year), which is the level of variation in our IV (cf. Schlenker & Walker, 2016; Knittel et al., 2016). This decreases our first-stage Kleibergen-Paap rk Wald *F*-statistics to around 21-23 for PM_{10} (the Cragg-Donald *F*-statistics are unchanged by definition). However, the second-stage coefficient remains statistically significant at the 1% level in all

[24] There are several reasons why the health effects of PM_{10} may be more adverse when humidity levels are low. Dry air allows PM_{10} to remain suspended in the atmosphere for longer, increasing the risk of inhalation. In addition, mucosal surfaces—such as nasal and bronchial linings—tend to dry out, reducing their ability to trap and clear inhaled particles. Finally, particle dispersion is more efficient in dry air, which may facilitate deeper penetration into the lungs (see, e.g., Kudo et al., 2019).

models for lower and total respiratory illnesses (results available upon request). Finally, we conducted a placebo test. Janke (2014), Simeonova et al. (2021), and Schlenker & Walker (2016) use visits for non-respiratory diseases as a placebo. We have access to overall and specific mental health diseases. As a placebo, we used Attention-Deficit/Hyperactivity Disorder (ADHD)-related visits, as to the best of our knowledge, there is no evidence of a causal link with air pollution. We also used a subset of mental health-related visits where we excluded depression and anxiety, as there is causal evidence relating these two to air pollution (Chang et al., 2019). The incidence rates of these illnesses are high (269 for ADHD and 560 for this subset of mental health illnesses per 10000 children). Across all our specifications, we do not find any evidence of an effect of PM_{10} on these outcomes (results available upon request).

6. Discussion and Conclusions

In this paper, we sought to determine the extent to which air pollution in low-pollution settings affects children's health, identify which children would benefit from further reductions in ambient pollution, and investigate whether environmental factors in such settings moderate these effects. To address these questions, we used the universe of administrative medical records from the universal public healthcare system between 2013 and 2017 for children in a large Spanish region (Catalonia). We combined these data with spatio-temporal kriging techniques to construct complete time-by-location data on several air pollution measures and other potential environmental confounders. We then instrumented for local PM_{10} concentrations—the main reference pollutant for air quality policies at the time—using variation in local wind direction in a multiple FE model.

We found that even at relatively low ambient levels, increases in PM_{10} concentrations raised the incidence of respiratory-related healthcare visits among children. Our preferred IV estimate indicated that a 1 $\mu g/m^3$ (or 4.5%) increase in local PM_{10} concentrations resulted in a 0.5% increase in overall respiratory visits among children. This effect was driven by lower respiratory illnesses, which have more serious health implications. The IV estimate was also only slightly larger than the uninstrumented FE estimate (0.3%), consistent with reduced measurement error from the use of spatio-temporal kriging. We also found evidence of heterogeneous effects: the youngest children (ages 0–5), and those exposed during hot or drier months were most affected.

Our setting (Catalonia, Spain) is subject to air quality standards set by the EU (Ambient Air Quality Directive 2008/50/EC). While there is a consensus in favor of more stringent regulations in heavily polluted environments, in more regulated settings the trade-off between the costs and benefits of further air quality improvements is less clear cut, as the marginal cost of additional restrictions is higher and the resulting health benefits are less obvious (Ferro et al., 2024; Gehrsitz, 2017). Understanding the impact of air pollution at these levels is crucial to inform policymakers in advanced economies of this trade-off. Our reduced form results suggest that tightening air quality standards, even in a low-pollution setting like ours, would yield additional benefits for society in terms of improved health for children. Back-of-the-envelope calculations indicate that the reduction in average annual PM_{10} concentrations during our sample period, from 24.8 in 2013 to 20.4 $\mu g/m^3$ in 2017, may have prevented approximately 16 million respiratory-related healthcare visits, with an associated direct cost savings of around €800 million over this period.[25] Moreover, the health benefits for these children may extend into the long term (Almond et al., 2018).

[25] Our IV estimates in Table 3, column 4, shows that a 1 $\mu g/m^3$ increase in PM_{10} causes 1 additional healthcare visit per month per 10,000 children who use public healthcare services across the PCUs in our sample. For the total population of children who use public healthcare services in Catalonia (1,225,406) this implies 122.54 additional visits per month and PCU. Aggregating the estimated visits increase across all PCUs (494) and

Our results from heterogeneity analyses are particularly relevant for understanding how weather conditions moderate the effects of air pollution on child health in low-pollution settings, and for informing the design of targeted policies for vulnerable groups. For example, our findings suggest that preventive public health interventions—such as targeted alerts for vulnerable populations during periods of elevated environmental risk—may be an effective tool for improving population health.

Our estimates based on PM_{10} are likely conservative for several reasons. First, because finer PM such as $PM_{2.5}$ would likely yield more adverse health effects on (lower) respiratory illnesses as a more refined PM can reach deeper into the lungs and into other organs of our body, including the cardiovascular system (Schlenker & Walker, 2016; Hayes et al., 2020). Second, individuals may engage in avoidance behavior or make defensive health investments (e.g., keeping children indoors), which can attenuate observed impacts (e.g., Neidell, 2009). Third, because we only count children who are sick enough to seek care, potentially underestimating the full burden of illness. Finally, air pollution affects more than just health: outcomes critical for human capital accumulation—such as fatigue, lack of focus, and memory impairment—are also likely to be affected (Aguilar-Gomez et al., 2022). Consistent with these findings, Chang et al. (2019) report an effect on worker productivity. That said, caution is warranted when generalizing these findings to other low-pollution settings with healthcare systems that differ substantially from ours—particularly those that are not universal, free at the point of use, or of comparable quality. Still, our findings highlight

months (60) months yields a cumulative burden of 3,632,326 additional healthcare visits attributable to a 1 $\mu g/m^3$ increase in PM_{10}. Applying observed proportions of care types—94.17% primary care, 5.05% emergency care, and 0.78% hospital admissions—and using average public tariffs from Mora et al. (2023) as costs (€41 for primary care, €85.9 for emergency care, and €880.2 for hospital admissions), the estimated cost burden is approximately €182 million. Multiplying this by the observed 4.4 $\mu g/m^3$ decline in PM_{10} between 2013 and 2017 leads to a saving of around 16 million respiratory-related healthcare visits, with an associated direct cost saving of around €800 million. These estimates do not include specialist visits and children who rely exclusively on private healthcare services.

the health and policy relevance of air pollution even in settings that meet current regulatory standards.

References

Aguilar-Gomez, S., Dwyer, H., Graff Zivin, J., & Neidell, M. (2022). This is air: The "nonhealth" effects of air pollution. *Annual Review of Resource Economics*, 14(1), 403-425.

Almond, D., Currie, J. & Duque, V. (2018). 'Childhood circumstances and adult outcomes: Act II', *Journal of Economic Literature*, Vol. 56, pp. 1360–1446.

Arceo, E., Hanna, R., & Oliva, P. (2016). Does the effect of pollution on infant mortality differ between developing and developed countries? Evidence from Mexico City. *The Economic Journal*, *126*(591), 257-280.

Chang, T. Y., Graff Zivin, J., Gross, T., & Neidell, M. (2019). The effect of pollution on worker productivity: evidence from call center workers in China. *American Economic Journal: Applied Economics*, 11(1), 151-172.

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

Cesur, R., Tekin, E., & Ulker, A. (2017). Air pollution and infant mortality: evidence from the expansion of natural gas infrastructure. *The Economic Journal*, 127(600), 330-362.

Currie, J. (2020). Child health as human capital. *Health Economics*, 29(4), 452–463.

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

Currie, J., Neidell, M., & Schmieder, J. F. (2009). Air pollution and infant health: Lessons from New Jersey. Journal of health economics, 28(3), 688-703.

Currie, J., & Schwandt, H. (2016). The 9/11 dust cloud and pregnancy outcomes: A reconsideration. *Journal of Human Resources*, 51(4), 805-831.

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

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

Deschênes, O., & Moretti, E. (2009). Extreme weather events, mortality, and migration. *The Review of Economics and Statistics*, 91(4), 659-681.

Deschênes, O., Greenstone, M., & Guryan, J. (2009). Climate change and birth weight. *American Economic Review*, 99(2), 211-217.

Dueñas-Espín, I., Vela, E., Pauws, S., Bescos, C., Cano, I., Cleries, M., ... & Roca, J. (2016). Proposals for enhanced health risk assessment and stratification in an integrated care scenario. *BMJ Open*, *6*(4), e010301.

Ferro, S., Palma, A., Serra, C., & Stafoggia, M. (2024). Beyond birth: The medium-term health impact of prenatal exposure to air pollution. *Journal of Environmental Economics and Management*, 103009.

Flores, M., & Wolfe, B. L. (2023). The influence of early-life health conditions on life course health. *Demography*, 60(2), 431-459.

Galán, C.; Smith, M.; Thibaudon, M.; Frenguelli, G.; Oteros, J.; Gehrig, R.; Berger, U.; Clot, B.; Brandao, R.; the EAS QC Working Group. (2014). Pollen monitoring: minimum requirements and reproducibility of analysis. *Aerobiologia*, 30(4), 385-395.

Gehrsitz, M. (2017). The effect of low emission zones on air pollution and infant health. *Journal of Environmental Economics and Management*, 83, 121-144.

Graff Zivin, J., Neidell, M., Sanders, N. J., & Singer, G. (2023). When externalities collide: Influenza and pollution. *American Economic Journal: Applied Economics*, 15(2), 320-351.

Hayes, R. B., Lim, C., Zhang, Y., Cromar, K., Shao, Y., Reynolds, H. R., ... & Thurston, G. D. (2020). PM2.5 air pollution and cause-specific cardiovascular disease mortality. *International Journal of Epidemiology*, 49(1), 25-35.

Hirst, J. M. (1952). An automatic volumetric spore trap. *Annals of Applied Biology*, 39(2), 257–265.

Jans, J., Johansson, P., & Nilsson, J. P. (2018). Economic status, air quality, and child health: Evidence from inversion episodes. *Journal of Health Economics*, *61*, 220-232.

Jayachandran, S. (2009). Air quality and early-life mortality: Evidence from Indonesia's wildfires. *Journal of Human Resources*, 44(4), 916-954.

Kleibergen, F., & Paap, R. (2006). Generalized reduced rank tests using the singular value decomposition. *Journal of Econometrics*, 133(1), 97-126.

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

Kudo, E., Song, E., Yockey, L. J., Rakib, T., Wong, P. W., Homer, R. J., & Iwasaki, A. (2019). Low ambient humidity impairs barrier function and innate resistance against influenza infection. *Proceedings of the National Academy of Sciences*, *116*(22), 10905-10910.

Mora, T., Puig-Junoy, J., Jacobs, R., & Cid, J. (2023). Differential costs for the non-adult ADHD population in Catalonia. *Health Economics Review*, 13(1), 24.

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

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

Nhung, N. T. T., Amini, H., Schindler, C., Joss, M. K., Dien, T. M., Probst-Hensch, N., ... & Künzli, N. (2017). Short-term association between ambient air pollution and pneumonia in children: A systematic review and meta-analysis of time-series and case-crossover studies. *Environmental Pollution*, 230, 1000-1008.

Palma, A., Petrunyk, I., & Vuri, D. (2022). Prenatal air pollution exposure and neonatal health. *Health Economics*, 31(5), 729-759.

Pereira, C., Valero, A., Loureiro, C., Davila, I., Martinez-Cocera, C., Murio, C., ... & Palomino, R. (2006). Iberian study of aeroallergens sensitisation in allergic rhinitis. European annals of allergy and clinical immunology, 38(6), 186.

Querol, X., Perez, N., Reche, C., Ealo, M., Ripoll, A., Tur, J., ... & Alastuey, A. (2019). African dust and air quality over Spain: is it only dust that matters? *Science of the Total Environment*, 686, 737-752.

Schlenker, W., & Walker, W. R. (2016). Airports, air pollution, and contemporaneous health. *The Review of Economic Studies*, 83(2), 768-809.

Simeonova, E., Currie, J., Nilsson, P., & Walker, R. (2021). Congestion pricing, air pollution, and children's health. *Journal of Human Resources*, 56(4), 971-996..

Stein, M. L. (1999). *Interpolation of Spatial Data*. New York: Springer.

Stafoggia, M., Bellander, T., Bucci, S., Davoli, M., De Hoogh, K., De'Donato, F., ... & Schwartz, J. (2019). Estimation of daily PM10 and PM2.5 concentrations in Italy, 2013–2015, using a spatiotemporal land-use random-forest model. *Environment International*, 124, 170-179.

Stock, J.H., & Yogo, M. 2005. "Testing for Weak Instruments in Linear IV Regression." In Identification and Inference for Econometric Models, edited by Donald W. K. Andrews and James H. Stock, 80–108. Cambridge: Cambridge University Press.

Storn, R., & Price, K. (1997). Differential evolution–a simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, 11(4), 341-359.

Thurston, G. D., Kipen, H., Annesi-Maesano, I., Balmes, J., Brook, R. D., Cromar, K., ... & Brunekreef, B. (2017). A joint ERS/ATS policy statement: what constitutes an adverse health effect of air pollution? An analytical framework. *European Respiratory Journal*, 49(1).

US Environmental Protection Agency (EPA). (2024). Particle Pollution and Your Health. [Document]. https://nepis.epa.gov/Exe/ZyPDF.cgi?Dockey=P1001EX6.txt (accessed 13 February 2025)

World Health Organization (WHO). (2005). *Effects of Air Pollution on Children's Health and Development – A Review of the Evidence*. WHO, European Centre for Environment and Health, Bonn. Available at: https://iris.who.int/bitstream/handle/10665/107652/E86575.pdf

World Health Organization (WHO). (2006). *Air quality guidelines: global update 2005: particulate matter, ozone, nitrogen dioxide, and sulfur dioxide*. WHO.

World Health Organization (WHO). (2021). WHO global air quality guidelines: particulate matter (PM2. 5 and PM10), ozone, nitrogen dioxide, sulfur dioxide and carbon monoxide. WHO.

Figure 1. Geographical distribution of primary care units, pollution and weather stations

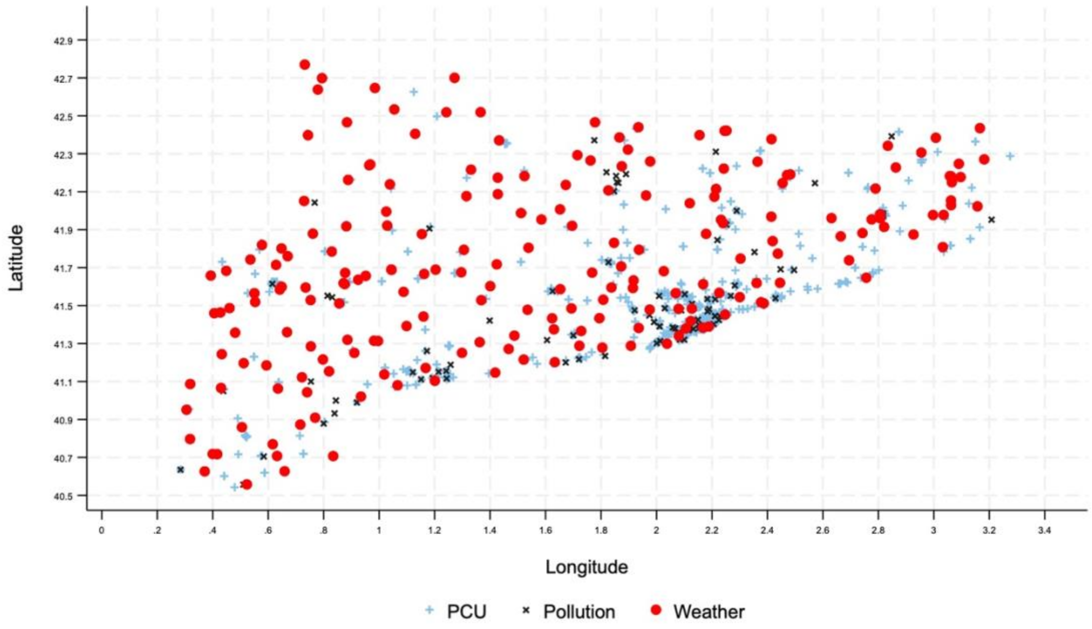

Source: Authors' elaboration.

Notes: The figure shows the geographical distribution of the included primary care units (*PCU*), air pollution stations (*Pollution*), and weather stations (*Weather*). The shape of the figure shows the Catalan territory.

Figure 2. Monthly incidence rates of lower, upper and total respiratory illnesses

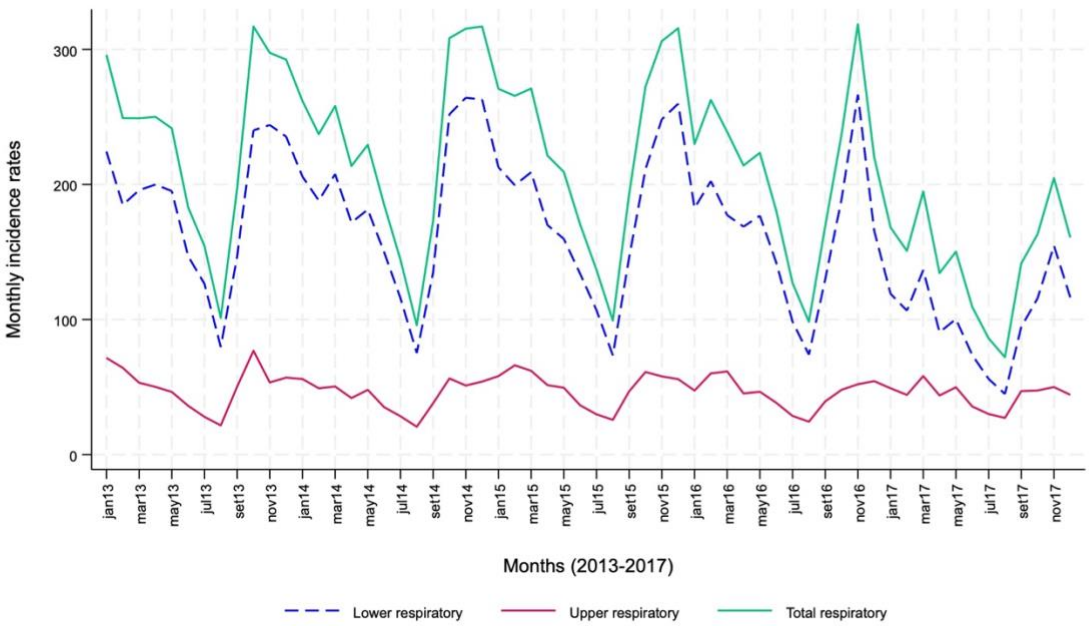

Source: Authors' calculations using the final estimation sample.

Notes: The figure displays average incidence rates of lower, upper, and total respiratory illnesses across primary care units in the final estimation sample, by year and month. See Section 3 for definitions.

Figure 3. Monthly evolution of air polluters

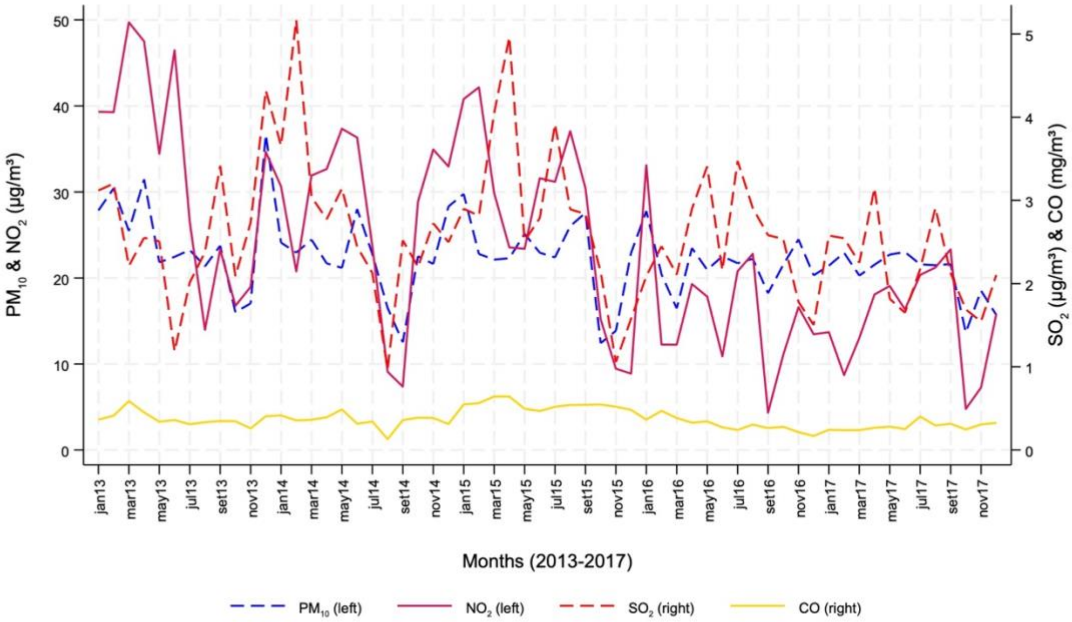

Source: Authors' calculations using the final estimation sample.

Notes: The figure displays average concentrations of particles with an aerodynamic diameter of 10 micrometers or less (PM_{10}), carbon monoxide (CO), nitrogen dioxide (NO_2), and sulfur dioxide (SO_2) across primary care units in the final estimation sample, by year and month. All pollutants are measured in micrograms per cubic meter (µg/m³), except CO, which is measured in milligrams per cubic meter (mg/m³).

Figure 4. PM_{10} concentration by wind direction and speed: Polar graphs for selected locations

Barcelona (East)	Berga (North)
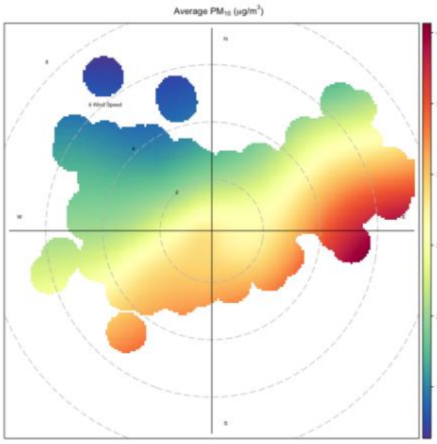	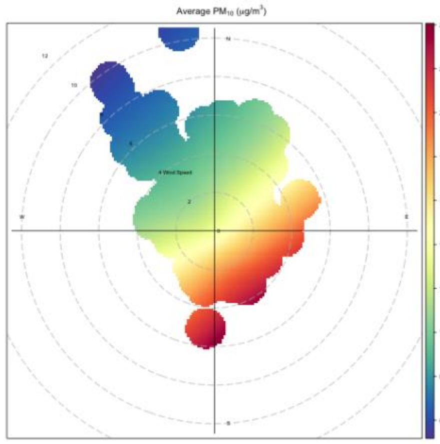
Lleida (West)	**Tarragona (South)**
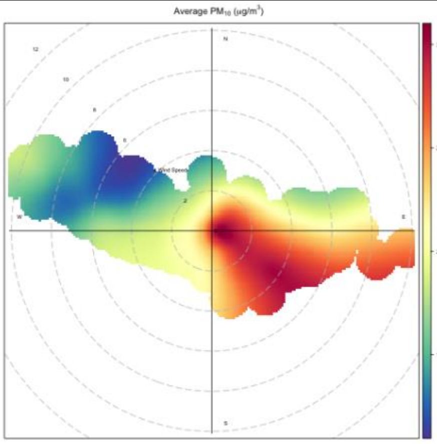	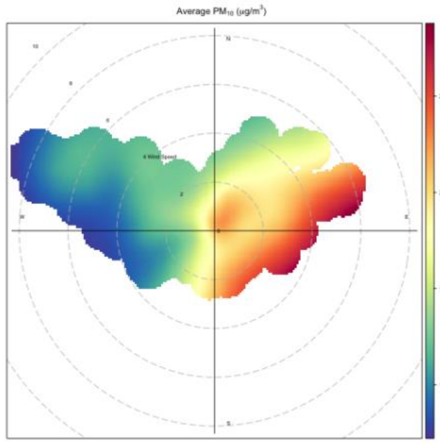

Source: Authors' elaboration.

Notes: The figure displays polar plots for four locations in Catalonia—Barcelona, Berga, Lleida, and Tarragona—based on historical data. The plots illustrate how PM_{10} concentration varies with wind direction and wind speed at each site.

Figure 5. Heterogenous effects on lower, upper, and total respiratory illnesses by socioeconomic and health characteristics (panels A to C), and by weather conditions (panels D to F)

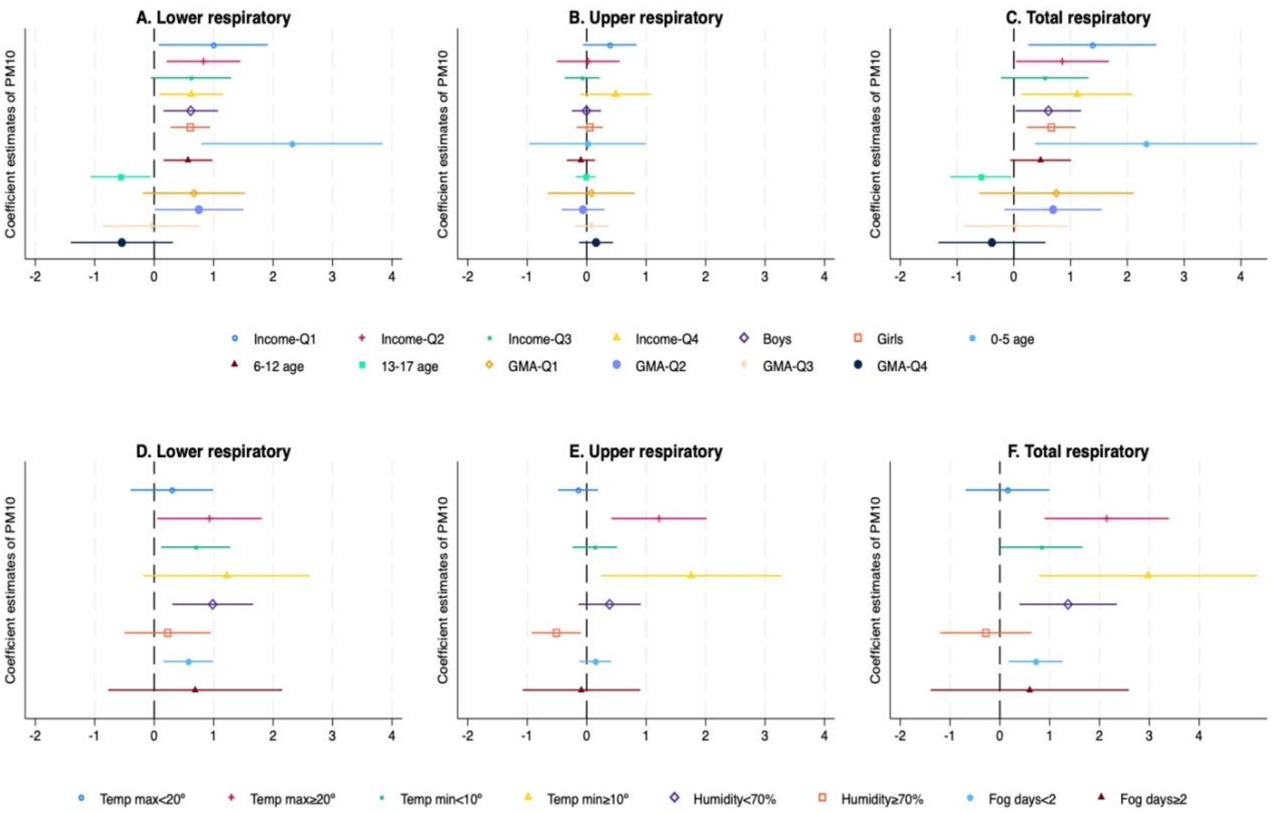

Notes: The figure shows coefficient estimates for PM_{10}, along with 95% confidence intervals, obtained from our preferred instrumental variables regressions, as in column (4) of Table 3. Each panel presents results for a different outcome. Each coefficient is from a separate regression, estimated on subsamples defined by the characteristic or condition indicated in the legend.

Table 1. Descriptive statistics for dependent and key explanatory variables

	Mean	Median	SD
Dependent variables			
Lower respiratory	162.38	112.749	207.682
Upper respiratory	46.883	29.621	69.599
Total respiratory	209.263	155.874	238.652
Air pollutants			
PM_{10} (μg/m^3)	22.32	22.213	6.189
CO (mg/m^3)	.372	.35	.118
SO_2 (μg/m^3)	2.598	2.505	1.165
NO_2 (μg/m^3)	23.626	22.523	12.595
Weather variables			
Accumulated precipitation	48.296	41.181	29.792
Maximum temperature (degrees)	19.605	20.097	4.543
Minimum temperature (degrees)	8.289	8.404	4.513
Atmospheric pressure (atm)	11.371	11.443	2.128
Mist (days)	.726	.226	1.409
Fog (days)	1.787	1.2	1.732
Humidity (%)	69.091	69.064	6.373
Wind speed at 10 m (m/s)	.997	.8	.868
Pollen variables (Pollen·day/m^3)			
Platanus	616.768	159.364	1326.703
Poaceae	184.532	93.639	314.872
Oleaceae	315.258	102.773	776.983
Urticaceae	177.478	124.559	172.4
Cupressaceae	990.65	423.226	1888.415
Sociodemographic variables			
Age (years)	9.323	9.341	1.961
Female (%)	.416	.419	.147
Individual passed away (%)	0	0	.004
Spanish nationality (%)	.895	.917	.111
Drug copayment exemption (%)	.065	.049	.082
Drug copayment 10% (%)	.06	.049	.073
Drug copayment 40% (%)	.546	.541	.164
Drug copayment 50% (%)	.329	.327	.165

Source: Authors' calculations using the final estimation sample.

Notes: SD stands for standard deviation. Dependent variables are measured as monthly healthcare visits for, correspondingly, lower, upper or any respiratory illness, per 10,000 children who use public healthcare services.

Table 2. OLS estimates of the effect of PM_{10} on respiratory illness visits

	Lower respiratory				Upper respiratory				Total respiratory			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
PM_{10} (μg/m3)	0.624***	0.721***	0.680***	0.669***	-0.050	0.003	0.022	0.022	0.574***	0.724***	0.702***	0.690***
	(0.060)	(0.062)	(0.062)	(0.062)	(0.039)	(0.038)	(0.040)	(0.040)	(0.081)	(0.082)	(0.083)	(0.082)
PCU-by-Quarter-by-Year FE	X	X	X	X	X	X	X	X	X	X	X	X
Month FE	X	X	X	X	X	X	X	X	X	X	X	X
Other pollutants (CO, SO_2, NO_2)		X	X	X		X	X	X		X	X	X
Weather and Pollen controls			X	X			X	X			X	X
Sociodemographic controls				X				X				X
N	29640	29640	29640	29640	29640	29640	29640	29640	29640	29640	29640	29640
Adjusted R2	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.91	0.91	0.91	0.91
R2	0.93	0.94	0.94	0.94	0.93	0.93	0.93	0.93	0.94	0.94	0.94	0.94
Mean-Y	151	151	151	151	52	52	52	52	203	203	203	203

Notes: The table reports ordinary least squares (OLS) estimates of the effect of PM_{10} concentrations on lower, upper, and total respiratory illnesses across different specifications. Standard errors, clustered at the primary care unit level, are reported in parentheses. Significance levels: * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

Table 3. IV estimates of the effect of PM_{10} on respiratory illness visits

	Lower respiratory				Upper respiratory				Total respiratory			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
PM_{10} (μg/m3)	0.620***	0.705***	0.802***	0.804***	0.087	0.158	0.201	0.203	0.707***	0.864***	1.003***	1.007***
	(0.144)	(0.171)	(0.171)	(0.171)	(0.106)	(0.124)	(0.125)	(0.125)	(0.202)	(0.239)	(0.238)	(0.238)
PCU-by-Quarter-by-Year FE	X	X	X	X	X	X	X	X	X	X	X	X
Month FE	X	X	X	X	X	X	X	X	X	X	X	X
Other pollutants (CO, SO_2, NO_2)		X	X	X		X	X	X		X	X	X
Weather and Pollen controls			X	X			X	X			X	X
Sociodemographic controls				X				X				X
N	29640	29640	29640	29640	29640	29640	29640	29640	29640	29640	29640	29640
CD *F*-statistic (8 df)	544.0	455.4	451.9	451.7	544.0	455.4	451.9	451.7	544.0	455.4	451.9	451.7
KP *F*-statistic (8 df)	667.5	524.1	474.3	470.2	667.5	524.1	474.3	470.2	667.5	524.1	474.3	470.2
Hansen *J*-statistic (*p-value*)	0.052	0.070	0.273	0.233	0.077	0.066	0.022	0.024	0.127	0.121	0.275	0.269
Mean-Y	151	151	151	151	52	52	52	52	203	203	203	203

Notes: The table reports instrumental variables (IV) estimates of the effect of PM_{10} concentrations on lower, upper, and total respiratory illnesses across different specifications. Eight instruments are included (see Equations (2) and (3)). The bottom panel reports two first-stage *F*-statistics used to assess instrument relevance: the Cragg–Donald (CD) *F*-statistic, computed under homoskedasticity and comparable to Stock–Yogo critical values, and the Kleibergen–Paap rk Wald (KP) *F*-statistic, which is robust to heteroskedasticity and serial correlation. The Hansen *J*-statistic tests the null hypothesis that the instruments are valid—that is, uncorrelated with the error term and correctly excluded from the main equation (Equation (1)). Standard errors, clustered at the primary care unit level, are reported in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.