Journal of Environmental Economics and Management 114 (2022) 102680

Contents lists available at ScienceDirect

Journal of Environmental Economics and Management

journal homepage: www.elsevier.com/locate/jeem

The short-term impact of air pollution on medical expenditures: Evidence from Beijing[☆]

Fan Xia [a], Jianwei Xing [b,*], Jintao Xu [b], Xiaochuan Pan [c]

[a] State Key Laboratory of Pollution Control & Resource Reuse, School of Environment, Nanjing University, Nanjing, 210023, China
[b] National School of Development, Peking University, Beijing, 100871, China
[c] School of Public Health, Peking University, Beijing, 100871, China

ARTICLE INFO

JEL classification:
I15
Q51
Q53

Keywords:
Air pollution
Medical costs
Heterogeneous impacts
Pollution episodes

ABSTRACT

We identify the short-term effects of $PM_{2.5}$ concentrations on medical costs in Beijing by analyzing two datasets: one detailing daily air quality indexes over a four-year period and the other containing individual-level records of all health care visits and medical transactions that occurred under a government insurance program that covers most city residents. We find that both higher levels of air pollution and longer-lasting pollution episodes significantly increase health care visits and medical expenditures. An analysis of multiple-day pollution episodes shows that marginal health care visits and marginal health costs start to increase as the pollution event lasts for consecutive days. Omitting the variation in the magnitude of the marginal effects of pollution exposure over the course of a pollution episode would lead to the underestimation of the total health costs of air pollution. Our findings provide empirical evidence that both the intensity and the duration of pollution episodes are critical considerations when designing policies to reduce the health costs of air pollution.

1. Introduction

Air pollution is a major environmental risk to human health across the globe. Estimates from the World Health Organization (WHO) show that ambient air pollution caused 4.2 million premature deaths worldwide in 2016.[1] Air pollution affects people everywhere, but the burden is greatest in low- and middle-income countries where the air pollution levels are some of the highest in the world, the resources to address pollution are limited, and both the establishment and enforcement of related regulations are often lacking. Poor air quality imposes increased health costs related to cardiovascular and respiratory diseases, shortened life expectancy, and lower productivity. However, the exact extent of the medical costs imposed by severe air pollution is still not well understood because of both the lack of large-scale, high-quality medical expenditure data and the lack of variation in air pollution levels needed to detect and analyze the dose–response relationship.

A growing literature in economics has aimed to quantify the causal impact of air pollution on health. This literature has largely focused on mortality, especially among infants and elderly individuals (Chay and Greenstone, 2003; Currie and Neidell, 2005; Currie et al., 2009; Currie and Walker, 2011; Knittel et al., 2015; Deryugina et al., 2019). These studies either exploit geographic variation in pollution levels and/or rely on variation in air quality due to traffic regulations. Some studies have used health insurance records to investigate the impacts of air pollution on human health (Deschênes et al., 2017; Williams et al., 2019; Deryugina et al., 2019).

☆ We acknowledge the financial support of the Energy Foundation China (No. G-1905-29506).

* Corresponding author.
E-mail addresses: xiafan@nju.edu.cn (F. Xia), jerryxing@nsd.pku.edu.cn (J. Xing), xujt@nsd.pku.edu.cn (J. Xu), xcpan@bjmu.edu.cn (X. Pan).

[1] https://www.who.int/news-room/fact-sheets/detail/ambient-(outdoor)-air-quality-and-health.

https://doi.org/10.1016/j.jeem.2022.102680
Received 1 September 2020
Available online 25 May 2022
0095-0696/

Other research has investigated the impact of air pollution on contemporaneous health by estimating the impact of pollutants on hospital visits for respiratory or cardiovascular diseases (Moretti and Neidell, 2011; Schlenker and Walker, 2016). These studies have all focused on developed countries, and they often estimate the health costs of air pollution by assuming a constant marginal effect of pollution exposure. While some studies have explored defensive expenditures related to air pollution in the context of low- and middle-income countries (Zhang and Mu, 2018; Ito and Zhang, 2019), few studies have estimated the exact morbidity costs caused by acute air pollution — likely because of the absence of ideal, high-frequency health data from areas experiencing large-magnitude variations in air quality. One exception is the work of Barwick et al. (2018), who make a first attempt to estimate the morbidity cost of air pollution in China. Using the universe of credit- and debit-card transactions in Chinese health facilities between 2013 and 2015, they find that the morbidity cost of air pollution is approximately two-thirds of the mortality cost estimates. However, their data are restricted to bank card transactions in healthcare facilities, which does not include insurance payments or other payment methods, and they do not have information on diagnoses or detailed medical expenses.

By using transaction-level data on medical expenses with detailed treatment information from the city's largest medical insurance program, this paper estimates the short-term effects of exposure to air pollution on medical expenditures in Beijing, which suffers from notoriously high levels of air pollution. In 2013, annual fine particulate matter ($PM_{2.5}$) concentration readings in Beijing reached 89.5 micrograms per cubic meter ($\mu g/m^3$), 18 times the World Health Organization's guideline level in 2021. In mid-January 2013, Beijing's $PM_{2.5}$ readings sometimes exceeded 900 $\mu g/m^3$, which is far above the "beyond index"[2] level and is considered to be extremely hazardous. This striking degree of day-to-day variation in and the frequent occurrence of high levels of $PM_{2.5}$ concentrations in Beijing provide an opportunity to examine the short-term causal effects of exposure to high levels of ambient air pollution as evidenced by health care visits, overall medical spending, and spending on respiratory drugs, cardiovascular drugs, and all other western medications.

In addition, long-lasting pollution episodes are commonplace in Beijing and other metropolitan areas in China. Although environmental scientists have recognized the need for episode-based research (Guo et al., 2014; Tan et al., 2018), empirical economic studies of air pollution have yet to address the issues that may arise not only from the intensity of pollution but also from the duration of pollution episodes. The frequent occurrence of such pollution events in Beijing allows us to investigate the variation in the magnitude of the marginal impacts of air pollution on health outcomes over the course of a pollution episode.

We access detailed, transaction-level data on medical expenses incurred between 2013 and 2016 through the Urban Employee Basic Medical Insurance (UEBMI) program in Beijing, which covers the majority of the city's formal residents. The detailed information on various medication expenditures provided by the UEBMI dataset allows us to explore the impact of pollution exposure on different disease-specific medication categories. Having access to the universe of transactions for all beneficiaries covered by the program also enables us to estimate the health costs of pollution for a wide range of different population groups. For each date, we examine individual transactions for all beneficiaries enrolled in the UEBMI program in Beijing at the subpopulation-group level defined by age and gender. We next combine group-by-day-level medical data with daily air pollution readings averaged from across all monitoring stations in Beijing from 2013 to 2016. To capture individuals' short-term responses to pollution exposure, we focus on medical outcomes over a three-day period in our analysis. These combined data allow us to estimate the linear effects of air pollution on medical spending by controlling for weather conditions and for a rich set of time fixed effects. To address endogeneity concerns related to time-varying unobservables correlated with both air pollution and medical expenditures, we estimate causal effects by adopting an instrumental variables (IV) approach, with atmospheric thermal inversions as the instrument.

To examine the impact of pollution duration, we analyze and compare the effects on health care visits and medical expenditures that result from pollution events that last for a single day (*single polluted days*) and from those that continue over longer periods (*pollution episodes*). After identifying all air pollution episodes in Beijing during the 2013–2016 period, we adopt a flexible approach that allows the marginal effects of $PM_{2.5}$ concentrations on health care visits and costs to evolve over the course of an episode. With this episode-based specification, we further investigate (1) whether the health impacts of pollution episodes are heterogeneous across the population and (2) how pollution levels interact with the dynamics of health care utilization over pollution episodes.

The results of our linear effects specification reveal that a 10 $\mu g/m^3$ increase in $PM_{2.5}$ leads to a 0.387 percent increase in three-day health care visits and a 0.376 percent increase in three-day medical expenses, suggesting a marginal increase of 1.13 percent in total annual medical expenses. This magnitude of this estimate is close to that of the relative effects found in the United States for an elderly population facing much lower $PM_{2.5}$ concentrations than Beijing residents. As shown by Deryugina et al. (2019), inpatient admissions rates for the US elderly population rose by 0.66 percent, and inpatient spending for Medicare beneficiaries rose by 0.51 percent in response to a 10 $\mu g/m^3$ increase in $PM_{2.5}$. Our linear effects results are also comparable to findings from China by Barwick et al. (2018), who show that an increase of 10 $\mu g/m^3$ in $PM_{2.5}$ causes a 0.65 percent increase in the number of healthcare transactions in the short run and a 1.5 percent increase in the value of those transactions. The authors estimate a total healthcare cost that is approximately two-thirds the size of the mortality cost estimates in the literature.

The results from our episode-based analysis, which examines the varying marginal impacts of air pollution over the course of a pollution episode, reveal that larger responses emerge as the pollution event lasts for more consecutive days; this finding implies that efficiency in terms of medical costs – and benefits for human health – could be increased even further by taking steps to reduce air pollution events with a longer duration. A back-of-the-envelope calculation based on the episode-based analysis implies that the total medical cost of air pollution is equal to 3.15 billion USD per year,[3] relative to no pollution, for Beijing UEBMI beneficiaries.

[2] See the report by the New York Times. https://www.nytimes.com/2013/01/13/science/earth/beijing-air-pollution-off-the-charts.html.

[3] We use an exchange rate of 3.664 CNY per USD, the purchasing power parity rate for 2013 (OECD, 2021), throughout the analysis.

This estimate is approximately 70 percent larger than the cost estimated by using the linear effects specification (1.83 billion USD), suggesting that the omission of variation in the magnitude of the health effects over the course of an episode leads the actual health costs of pollution to be underestimated, especially in settings with frequent occurrences of high-level and long-lasting pollution events.

Our analysis contributes to the existing economic literature on air pollution in the following ways. First, the use of detailed information on medical expenses and prescriptions enables us to estimate the impacts of air pollution on expenses in different medication categories. Second, access to detailed personal data enables us to estimate the distribution of the health impacts of air pollution across a wide spectrum of people rather than members of vulnerable groups (infants and the elderly) only, as is largely the case in the existing literature. Third, unlike other relevant studies, our analysis explores the specific features of pollution episodes in China and identifies the variation in health costs over the course of a pollution episode, which greatly matters for an accurate estimation of the total health costs of air pollution. Our episode-based study allows us to provide empirical evidence that both the intensity and duration of air pollution episodes are critical considerations when designing policies to enhance air quality and reduce health costs.

The rest of the paper is organized as follows: Section 2 provides a brief description of the data on medical insurance, air quality and weather conditions. Section 3 introduces empirical models and identification strategies. Section 4 presents the estimation results for the linear effects of $PM_{2.5}$ concentrations on health care visits and medical costs, the estimated impact of air pollution from the episode-based analysis, and the heterogeneous effects among different demographic groups as well as different pollution levels. Section 5 concludes.

2. Background and data

To estimate the medical costs of air pollution, our study takes advantage of a dataset from the UEBMI program in Beijing. This administrative dataset records every hospital visit and transaction for each beneficiary.[4] As such, it provides rich information about medical expenses and allows us to investigate various outcomes of interest.

2.1. Beijing Urban Employee Basic Medical Insurance dataset

The universal basic medical insurance system developed by the Chinese central government maintains a 95 percent total enrollment rate (Yu, 2015). As one of the pillars of the basic medical insurance system, the UEBMI program is jointly funded by employees and employers, and it is the primary program that helps cover healthcare expenses for urban residents. As a government-run, mandatory social insurance program, UEBMI is administered at the municipal level and covers all urban residents employed in the formal sector, with coverage continuing after retirement. The family members of beneficiaries are not automatically covered by the program unless they themselves are employed in the formal sector. Therefore, the unemployed, those who work in the informal sector, and rural residents are not covered under this program.[5] Since UEBMI coverage is based on employment, individuals are eligible to join UEBMI based on the city where they are employed rather than the city where their registered residence is located under the *hukou* system.

In the Beijing UEBMI system, each beneficiary maintains an account, and each medical service payment that the beneficiary makes in a hospital, such as for registration fees, examination fees, or medication fees, is recorded. Nonhospital visits to medical personnel and the transactions that take place in nonhospital-affiliated pharmacies or drug stores are not included in the UEBMI data, nor are visits to hospitals or medications purchased outside Beijing. However, we believe that the expenditures recorded in the UEBMI data still capture the majority of medical-related spending among UEBMI beneficiaries for the following reasons. First, in China, health care services are provided mostly by hospitals, and most physicians are full-time employees at hospitals. Outpatient care in nonhospital clinics is rarely provided (Burns and Liu, 2017). Therefore, patients usually visit hospitals directly to seek treatment, including for nonemergency visits, and they usually prefer large hospitals with high-quality medical resources. Second, all hospitals in China operate their own pharmacies for outpatient drug sales. As a result, most patients purchase over-the-counter (OTC) or prescription medication directly from the hospital after seeing a doctor.[6] Medication purchases made in nonhospital-affiliated drug stores accounted for only 9.3 percent of overall drug consumption in 2014.[7]

We use confidential data from the Beijing UEBMI maintained by the Beijing Municipal Medical Insurance Bureau. We use data covering the 2013–2016 period.[8] According to our dataset, 19.7 million individuals were enrolled in the Beijing UEBMI at the end

[4] In China, health care services are mostly provided by hospitals. Visits to hospitals could thus include nonemergency visits to primary care physicians or other health care providers that would take place in a clinic or other nonhospital setting in countries such as the United States. See Section 2.1 for a more detailed introduction to the Chinese health care system and health insurance programs.

[5] The other two government-run complementary basic medical insurances programs are the Urban Resident Basic Medical Insurance (URBMI) program, which covers urban residents who are either unemployed or who work in the informal sector, including disabled individuals and children, and the New Rural Cooperative Medical Insurance (NRCMI) program, which covers rural residents. These two programs, together with UEBMI, cover the majority of residents in China.

[6] In China, patients have to register with their ID before receiving any medical services when they visit a hospital, and patients cannot purchase any drugs, including OTC medications, from hospital pharmacies without a prescription from a doctor. Therefore, all medical expenditures that take place in hospitals can be linked to a health care visit record.

[7] According to the Beijing Municipal Bureau of Statistics, total drug sales were 78.3 billion CNY in Beijing in 2014 (Source: http://www.beijing.gov.cn/gongkai/shuju/tjgb/201706/t20170608_1838187.html), while the sales in drug stores were 7.3 billion CNY that same year (Source: http://www.phirda.com/artilce_16010.html).

[8] Data for January and February 2016 are missing, and thus, those months are excluded from the study.

of 2016, which indicates a considerably high rate of coverage given the official estimate of 21.7 million residents in Beijing (Beijing Municipal Bureau of Statistics, 2017).[9]

There are several advantages of using the Beijing UEBMI data to investigate the impact of air pollution on medical spending. First, Beijing is one of the most developed cities in China, and its health care system is advanced, providing greater access to high-quality health care than is available in other regions. Medical records from facilities in less-developed regions may not accurately reflect the health outcomes of the local population because health facilities in such regions are comparatively scarce, and as a result, only a small share of the population is covered by the health care system. Thus, medical insurance records in other regions may not accurately reflect a population's health outcomes that result from its exposure to air pollution. In contrast, health insurance coverage is widespread in Beijing, and the UEBMI beneficiaries in our sample account for the majority of the adult residents in Beijing. Therefore, these data are likely to provide a more accurate estimate of the impact of air pollution on local residents' health.

Second, although the UEBMI data only cover those residents in Beijing who are officially employed, the medical expenditure records in the UEBMI data are likely to provide a more accurate estimate of the impacts of air pollution on local residents than is possible from using other options, such as transactions data collected from hospitals, which do not distinguish between local and nonlocal patients. As Beijing has some of the best medical resources in China and attracts patients from across the entire country, it would be difficult to separate out the hospital visits and medical spending of Beijing residents only if transactions data collected from hospitals with no detailed patient information were used.[10] Other studies similarly use health insurance records to examine the medical costs of air pollution (Deschênes et al., 2017; Williams et al., 2019; Deryugina et al., 2019), however, their analyses focus on linear effects of pollution and on effects in the United States, which has much lower levels of $PM_{2.5}$.

The information on enrolled beneficiaries included in the Beijing UEBMI data is updated monthly. The UEBMI dataset provides detailed information on variables in four key categories: **(1) hospital information**, including name and location (district); **(2) personal information**, including an anonymized individual identification number, gender, birth date, retirement status, and the date when the individual's insurance coverage began; **(3) medical expense settlement data**, which cover each transaction and report the individual identification number of the beneficiary, the identification number of the hospital visited, the date of the visit, the date of the settlement, the treatment category (inpatient, outpatient, or emergency room admission[11]), and a set of detailed expenditure variables (total medical spending, reimbursement amount and copay, and spending on different categories, including examinations, treatment, medications, and materials); and **(4) medication data**, which cover more than 20,000 different medications and include transaction identification numbers, medicine identification numbers, and the name, amount and cost of each medicine used during the corresponding hospital visit. Following the medicine classifications used by the Beijing Reimbursement Drug List for Basic Medical Insurance,[12] we assign every western drug to one of 14 major categories, focusing primarily on those medications used in respiratory and cardiovascular treatment.

We match all hospital, personal, expenditures, and medication information for each transaction using the transaction, individual, and hospital identification numbers. We are then able to aggregate daily costs and assign them to demographic subgroups. To analyze the heterogeneity in the relationships between medical costs and air pollution, we categorize individuals into different demographic groups based on age and gender. We use the following six age groups: younger than 30, 30–39, 40–49, 50–59, 60–69, 70 or older. We aggregate daily hospital-related expenditures for each of the 12 demographic groups and then divide the aggregated values by the number of beneficiaries in the corresponding demographic group for the corresponding month and thereby construct a balanced panel at the group-day level. Out of 19.7 million insurance beneficiaries in Beijing at the end of 2016, 60 percent were younger than 40, and 13.2 percent were over age 60. Appendix Table S2 summarizes the population distribution across each group defined by age and gender.[13]

To capture potential lagged effects or short-run "harvesting" effects, we focus on three-day (the current day and the following two days) totals for the health outcomes: health care visits; total medical expenditure; out-of-pocket expenditure; and expenditures for all medications and separately for cardiovascular medications, respiratory medications, and other western medications.[14]

Table 1 reports the three-day total number of health care visits and medical costs aggregated to the city level (amounting to a population-weighted average of the group-level variables). Over the period from 2013 to 2016, on average, there were 52.04 health care visits over a three-day period (17.35 visits per day) per 1,000 beneficiaries in Beijing; the three-day total medical expenditures were 30.56 CNY per beneficiary, including 18.83 CNY per beneficiary for medication expenses. Out-of-pocket costs accounted for 28.8 percent of total expenditure. The frequency of healthcare visits and the magnitudes of medication expenditures in our sample are close to the official statistics reported by National Health and Family Planning Commission in China (NHFPC) for the same data period.[15]

---

9 This is the officially reported population of formal and long-term residents living in Beijing in 2016, which has been imputed from survey samples and may deviate from a more accurate estimate based on census data.

10 According to the official reports, non-local patients (medical tourists from other cities) account for about one third of total patients in Beijing's hospitals. http://bj.people.com.cn/n2/2020/0518/c14540-34025741.html.

11 Outpatient, inpatient and emergency room visits account for 95.4 percent, 1.1 percent and 3.4 percent, respectively, of the daily average number of health care visits. Detailed summary statistics for the outcome variables of interest by treatment category are reported in Appendix Table S1.

12 Released by the Beijing Municipal Human Resources and Social Security Bureau, 2017.

13 The age distribution of the Beijing UEBMI program beneficiaries and that of Beijing permanent residents as well as of the total Chinese population are provided and compared in Appendix Table S2. The percentage of group members who are below the age of 30 is smaller for the program beneficiaries than for the resident population overall. This is reasonable, as the UEBMI program covers only residents who are employed, excluding children.

14 Since we are able to classify western medicines into disease categories based on the medicine classifications only, here, *other western medications* include expenditures on western medicines other than cardiovascular and respiratory medications. The difference between total medication expenditures and the sum of cardiovascular, respiratory, and other Western medications is due to expenditure on Chinese medicines.

15 Source: http://www.nhc.gov.cn/mohwsbwstjxxzx/tjtjnj/new_list.shtml.

Table 1
Summary statistics of health care visits and medical costs, daily $PM_{2.5}$ levels and weather variables in Beijing, 2013–2016.

	mean	sd	min	max	Count
Three-day totals of medical variables					
Health Care Visits	52.04	13.93	6.19	111.63	1380
Total Expenses	30.56	10.00	2.86	65.29	1380
Medication Expenses	18.83	5.80	1.72	41.51	1380
Out-of-pocket Expenses	8.81	3.07	1.20	20.88	1380
Cardiovascular Medication	2.85	0.86	0.26	6.01	1380
Respiratory Medication	0.34	0.11	0.05	0.93	1380
Other Western Medication	7.60	2.42	0.85	15.40	1380
Pollution and weather variables					
$PM_{2.5}$ (10 μg/m^3)	8.1	6.7	0.6	47.1	1384
Daily minimum temperature (°C)	10.7	9.3	−9.2	27.9	1384
Daily maximum temperature (°C)	20.1	10.4	−1.8	41.1	1384
Relative humidity (%)	54.1	19.8	8.0	99.0	1384
Precipitation (mm)	1.5	9.0	0.0	253.5	1384
Wind speed (m/s)	2.1	0.8	0.6	6.6	1384

Note: Statistics of three-day total medical outcomes are reported. The unit of health care visits is expressed as the number of visits per 1,000 beneficiaries. The units of cost variables are expressed as CNY per beneficiary. Note that observations of the above variables are not equal since a few more leads and lags of the weather and pollution variables are used in our analysis.

2.2. Air quality and meteorological data

The China National Environment Monitoring Centre (CNEMC) began systematic monitoring of ambient $PM_{2.5}$ in 2013. We aggregate hourly air quality data from all the national monitoring stations managed by CNEMC in Beijing to the daily level to obtain average daily pollutant concentrations from 2013 to 2016. Beijing is one of the most highly polluted cities in China. Compared with other regions in China, Beijing experiences worse pollution and has more days of high pollution levels, as indicated by the density distribution curves for daily $PM_{2.5}$ from 2013 to 2016 (Appendix Figure S1).

We obtain daily weather data for Beijing from the National Daily Climatic Data maintained by the China Meteorological Administration.[16] We include variables for minimum temperature, maximum temperature, relative humidity, precipitation, and wind speed as controls for the direct impacts of weather on health.

Table 1 presents summary statistics of the daily $PM_{2.5}$ and weather variables. The mean $PM_{2.5}$ reading in Beijing from 2013 to 2016 was 81 μg/m^3, with a standard deviation of 67 μg/m^3. On the most polluted day during the period, the $PM_{2.5}$ concentration was 471 μg/m^3.

2.3. Air pollution episodes in Beijing

We examine the health impacts of air pollution episodes that extend beyond a single day by defining such episodes following Tan et al. (2018). Under this definition, an air pollution episode spans the period in which the daily average $PM_{2.5}$ concentration grows from less than 50 μg/m^3, to a peak value of over 100 μg/m^3, and then drops back below 50 μg/m^3. To be considered an air pollution episode, the duration, measured from the first day to the last day when $PM_{2.5}$ is above 50 μg/m^3, must be at least three days. In other words, if we refer to a day with a daily average $PM_{2.5}$ concentration over 50 μg/m^3 as a *polluted day*, three or more consecutive *polluted days* with at least one day having a $PM_{2.5}$ concentration over 100 μg/m^3 constitute a pollution *episode*. We define *single polluted days* as those *polluted days* that do not belong to any *episodes* as defined above. Days with $PM_{2.5}$ concentrations less than 50 μg/m^3 are referred to as *clear days*. To rule out potential lagged pollution effects captured by the marginal responses to $PM_{2.5}$ on clear days and single polluted days, the *clear days* and *single polluted days* in our sample do not include those that are preceded by a polluted day; such days are dropped in our episode-based analysis.

Fig. 1 plots the time series for daily $PM_{2.5}$ levels in each year in Beijing. The gray bars indicate every single air pollution episode, with the width of the bar indicating the duration of the episode. From 2013 to 2016, the city experienced 114 air pollution episodes in total, with a mean duration of 5.6 days and a mean concentration of 127 μg/m^3 (Appendix Table S3). The minimum duration was three days (by definition), and the longest episode lasted 16 consecutive days. The average $PM_{2.5}$ level was 75 μg/m^3 during the least-polluted episode and 297 μg/m^3 during the most heavily polluted episode. Most episodes lasted for three to five days, with an average $PM_{2.5}$ concentration of approximately 100 μg/m^3, as revealed by the frequency distribution in Appendix Figure S2.

Table 2 summarizes the daily $PM_{2.5}$ levels by type of day according to our definitions. Out of 1072 days included in the episode-based analysis, there were 332 *clear days* with an average $PM_{2.5}$ concentration of 27 μg/m^3, 101 *single polluted days* with a mean $PM_{2.5}$ of 77 μg/m^3, and 639 *episode days* with mean $PM_{2.5}$ levels by day of episode ranging from 89 μg/m^3 to 151 μg/m^3.

[16] Data source: https://data.cma.cn/en.

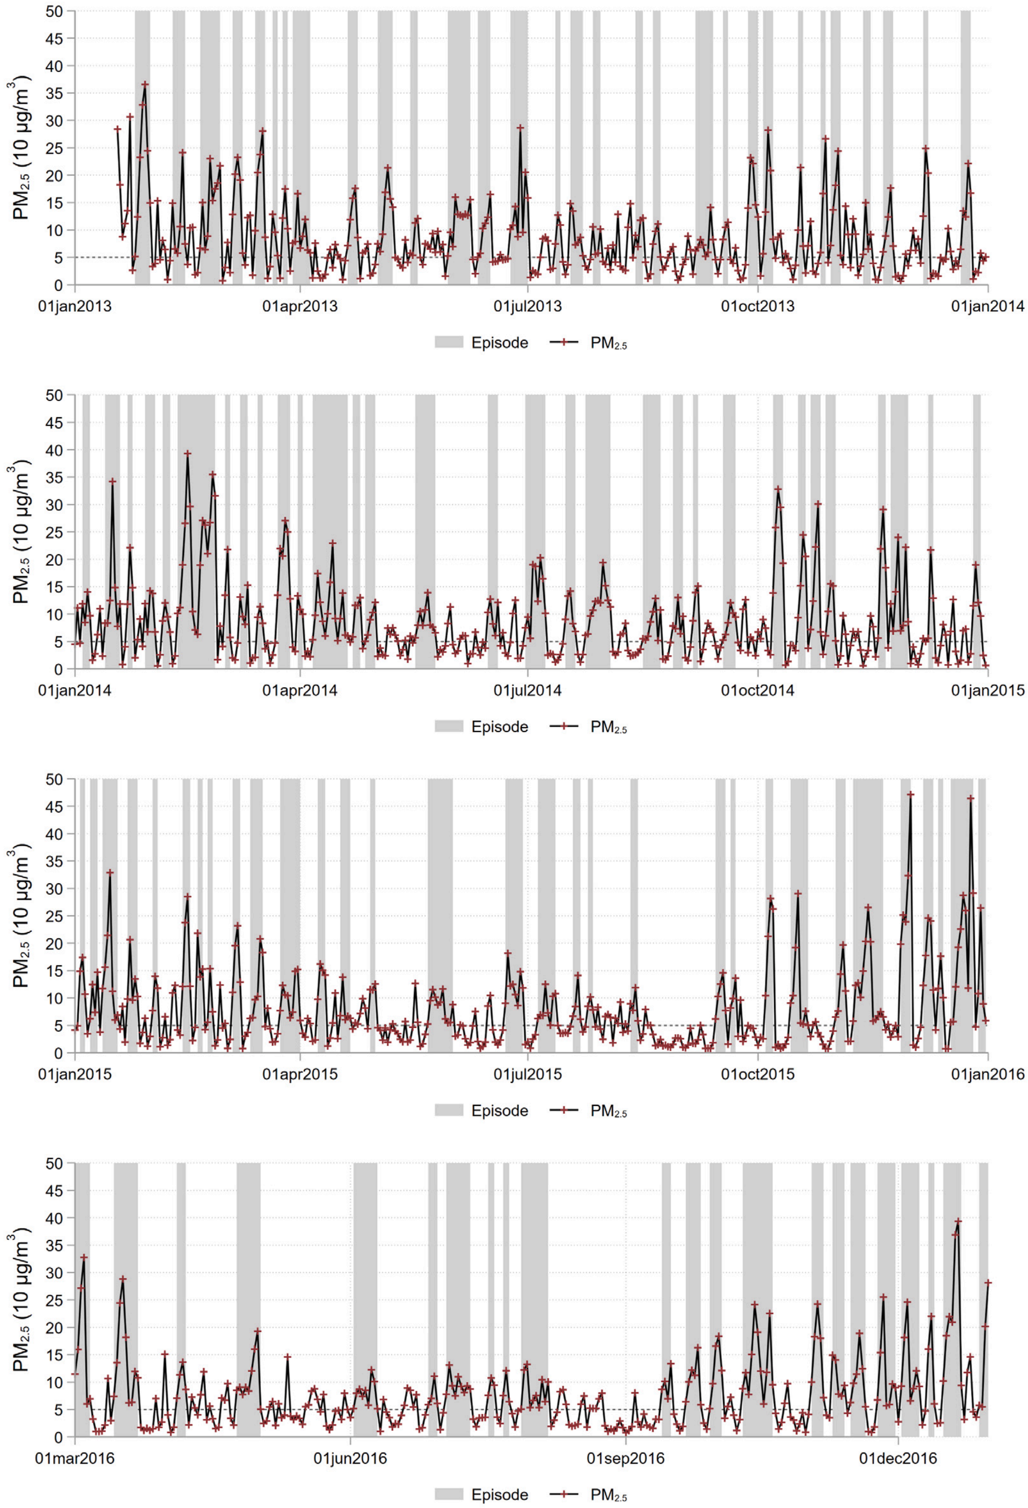

Fig. 1. Pollution episodes in Beijing, 2013–2016.

Note: Time series of daily $PM_{2.5}$ levels in Beijing, are marked by red pluses and connected by black lines, and pollution episodes are shaded by the gray rectangles covering days in each episode, from 2013 to 2016. A day with daily average $PM_{2.5}$ level above 50 μg/m^3 is referred to as a polluted day. An episode is defined as a series of at least three consecutive polluted days, and the $PM_{2.5}$ level on the most polluted day must exceed 100 μg/m^3. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 2
Summary statistics of daily $PM_{2.5}$ levels over pollution episodes in Beijing, 2013–2016.

	mean	sd	min	max	Count
Clear day	2.7	1.2	0.6	5.0	332
Single polluted day	7.7	2.5	5.0	15.3	101
Day 1	8.9	3.3	5.0	21.8	114
Day 2	13.0	5.4	5.3	26.4	114
Day 3	13.9	6.6	5.3	32.8	114
Day 4	15.1	8.0	5.1	34.2	90
Day 5	13.4	8.9	5.1	47.1	66
Day 6	12.2	7.5	5.0	39.4	41
Day 7+	13.0	7.5	5.0	46.4	100
Total	9.1	7.1	0.6	47.1	1072

Note: $PM_{2.5}$ is measured in the unit of 10 $\mu g/m^3$. A day with daily average $PM_{2.5}$ levels above 50 $\mu g/m^3$ is referred to as a polluted day. An episode is defined as a series of consecutive polluted days lasting at least three days, during which the $PM_{2.5}$ level on the most polluted day exceeds 100 $\mu g/m^3$. The indicator *clear day* refers to days with average $PM_{2.5}$ levels below 50 $\mu g/m^3$ that are not following a polluted day. The indicator *single polluted day* refers to polluted days that do not belong to any pollution episode, and do not follow a polluted day. Observations otherwise are dropped in the episode-based analysis.

3. Identification strategy

3.1. Linear short-term effects of daily $PM_{2.5}$ exposure

We first evaluate the linear health impacts of air pollution over the short run, as in other existing studies (Schlenker and Walker, 2016; Deryugina et al., 2019), and we specify our model as follows:

$$\ln Y_{gt} = \beta Pollution_t + W_t'\gamma + \delta_g + FE_{time} + \epsilon_{gt} \tag{1}$$

where $Pollution_t$ represents the daily average $PM_{2.5}$ concentration in the city on date t.[17] We examine the linear effects of $PM_{2.5}$ on the number of health care visits and spending at the group level for subpopulations defined by gender and age. Y_{gt} is the outcome variable for beneficiary group g on date t.

The impacts of air pollution can be dynamic in the short run. For example, the effect of pollution exposure on human health may be lagged, and individuals may delay seeking medical treatment. Some health effects could also be the result of forward displacement: among those with certain health conditions, acute air pollution exposure may trigger symptoms that would likely have occurred in the near future regardless of the presence of air pollution.

To address these possible short-run lagged or forward displacement effects, following Deryugina et al. (2019), we use the total number of health care visits and medical expenditure over a three-day time window as our outcome variables, i.e., the health outcome variables Y_{gt} are three-day totals, beginning on day t and including the following two days, $t+1$ and $t+2$. More specifically, the outcome variables include three-day totals for (1) health care visits, (2) total medical expenditures, (3) drug expenditures, and (4) out-of-pocket expenditures, as well as for three categories of drug expenditure: (5) cardiovascular medications, (6) respiratory medications, and (7) other western medication expenditures.[18] To investigate whether this three-day specification adequately captures the lagged effects, we estimate a distributed lag regression model with an increasing number of lags, as reported in Table 3. The results show that the largest response occurs on the concurrent day and that the cumulative effects of pollution exposure stabilize beyond a span of three days, which boosts our confidence in our choice of a three-day specification to estimate the short-term impacts of air pollution in our setting.[19]

β is the coefficient of interest, which captures the short-term linear effects of $PM_{2.5}$ fluctuations on the health outcome Y. All the dependent variables are in log form. Thus, β should be interpreted as a percentage change in health care visits or expenditures (i.e., the relative impact of increased air pollution), holding other factors constant.

The matrix of weather controls, W_t', consists of flexible functions of the daily maximum temperature, minimum temperature, relative humidity, precipitation, and wind speed.[20] Following Deryugina et al. (2019), we also control for two leads of these weather variables to rule out the effects on health outcomes from weather conditions over the next two days. In addition, two leads and two lags of $PM_{2.5}$ are included for the OLS estimates, and two leads and two lags of the instruments are included for the IV estimates.[21]

[17] Air pollution varies only along a temporal dimension because different groups living within the same city are assumed to be exposed to the same concentration of ambient air pollution on the same day.

[18] The unit of measure for health care visits is visits per 1,000 beneficiaries. The unit of measure for the cost variables is CNY per beneficiary.

[19] As a robustness check, we also estimate our model with a seven-day specification (Panel A in Appendix Table S6), and we find a similar magnitude for the marginal effect of pollution.

[20] More specifically, we include a set of daily weather indicators that includes maximum and minimum temperature bins (in 5 °C bins, starting from −15 °C), and second order polynomials of relative humidity, precipitation, and wind speed.

[21] Since we use three-day health outcomes as our dependent variables, as stated in Deryugina et al. (2019), two lags of $PM_{2.5}$ (OLS) or two lags of the IVs are included to minimize concerns about autocorrelation.

Table 3
Dynamic effects of daily $PM_{2.5}$ levels on medical expenses.

	(1) No lags	(2) 1 lag	(3) 2 lags	(4) 3 lags	(5) 4 lags	(6) 5 lags	(7) 6 lags
OLS estimates							
$PM_{2.5}$(t)	0.0247*** (0.00866)	0.0223** (0.00955)	0.0258*** (0.00841)	0.0249** (0.01011)	0.0260** (0.01274)	0.0245* (0.01324)	0.0247* (0.01365)
$PM_{2.5}$(t-1)		0.0109 (0.01602)	−0.00185 (0.00875)	0.000606 (0.01032)	−0.000896 (0.00630)	0.00367 (0.00508)	0.000904 (0.00698)
$PM_{2.5}$(t-2)			0.0225 (0.01796)	0.0114 (0.00843)	0.0155 (0.00990)	0.0119 (0.01053)	0.0108 (0.01228)
$PM_{2.5}$(t-3)				0.0131 (0.01406)	0.00557 (0.01566)	0.0120 (0.01352)	0.0102 (0.01343)
$PM_{2.5}$(t-4)					0.00359 (0.01394)	−0.00410 (0.00918)	0.000430 (0.00772)
$PM_{2.5}$(t-5)						0.00292 (0.01257)	−0.00590 (0.01746)
$PM_{2.5}$(t-6)							0.00522 (0.01233)
cum. eff.	0.0247*** (0.00866)	0.0332** (0.01556)	0.0465* (0.02510)	0.0500* (0.02712)	0.0498** (0.02101)	0.0509** (0.02358)	0.0464* (0.02404)
Observations	16608	16584	16560	16536	16512	16488	16464

Note: Daily medical expenses (in CNY per beneficiary) are regressed on same-day and lagged $PM_{2.5}$ (in 10 $\mu g/m^3$) simultaneously using specifications of distributed lag models. Number of lags involved varies from zero to six. Cumulative effects are the sum of coefficients of same-day and all lagged $PM_{2.5}$ included. All the regressions control for group fixed effects, time fixed effects of holiday, year–month, and day-of-week, and daily weather variables. The weather variables include maximum and minimum temperature bins (in 5°C bins starting from -15°C), second-order polynomials of relative humidity, precipitation, and wind speed. Standard errors are in parentheses and clustered by group-year. Significance levels are indicated by *** 1%, ** 5%, * 10%.

δ_g refers to demographic group fixed effects. Regressions are weighted by the sample size of each group — that is, the number of UEBMI participants by the end of the corresponding month. We also control for several sets of time fixed effects, FE_{time}, to address endogeneity arising from the temporal patterns in human activities that simultaneously affect air quality and health care visits. We include holiday indicators, day-of-week fixed effects, and year-by-month fixed effects.[22]

Appendix Figure S3 presents scatter plots of the three-day total number of visits and total medical expenses versus pollution levels. The figure reveals that after controlling for weather factors and time fixed effects, the health care utilization and expenditure variables exhibit significant positive correlations with the daily $PM_{2.5}$ concentrations, providing strong preliminary evidence of the short-term health effects of air pollution.

3.2. Instrumental variables

Although air pollution is relatively random in the short run compared to the long run, daily $PM_{2.5}$ exposure can still be endogenous. Endogeneity can arise through unobservables that simultaneously affect daily air pollution levels and health care behaviors. For example, traffic congestion worsens air quality and increases travel costs, which might prevent people from visiting hospitals.

To identify the causal relationship between ambient fine particulate matter pollution and short-term health care visits and expenditure, we exploit the day-to-day variation in $PM_{2.5}$ and instrument for $PM_{2.5}$ concentrations with thermal inversions. A thermal inversion is a meteorological phenomenon that occurs when warmer air is held above cooler air. Normally, the temperature of the atmosphere decreases as altitude increases, but when a thermal inversion occurs, the polluted air below is cooler than the clean air above and cannot rise and mix with that warmer, clean air. Pollutants are therefore trapped close to the ground. A number of economic studies have used thermal inversions as an instrument to study the effects of air pollution (Arceo et al., 2016; Chen et al., 2018; He et al., 2019; Fu et al., 2021). In our study, we exploit the day-to-day variation and short-term changes in thermal inversions, which we believe to be exogeneous and uncorrelated with long-term city development and other economic conditions that might also affect human health. Thus, we estimate model (1) using two-stage least squares (2SLS) regression analysis. However, thermal inversions cannot be used to account for the avoidance behaviors that people might practice to reduce their exposure to air pollution. Therefore, our estimates of the health costs of pollution do not include the cost of defensive behaviors and are conditional on the avoidance behaviors practiced by individuals.

To construct the thermal inversion variables, we obtain atmospheric temperature gradient data from the MERRA-2 product M2I6NPANA (Version 5.12.4), developed by NASA Goddard Earth Sciences (GES) Data and Information Services Center (DISC),

[22] The year-by-month fixed effects throughout this paper are dummies indicating each year–month combination during our sample period, e.g., 2013-January, 2013-February, . . . , 2016-November, 2016-December.

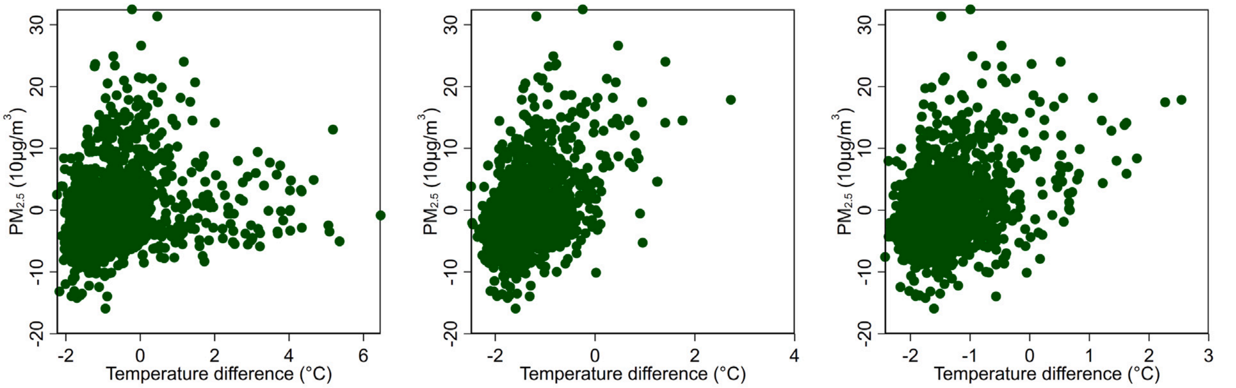

(a) 2nd-1st layer (975hPa - 1000hPa) (b) 3rd-2nd layer (950hPa - 975hPa) (c) 4th-3rd (925hPa - 950hPa)

Fig. 2. Correlation between thermal inversion and $PM_{2.5}$.
Note: Figures show scatterplots of the relationships between residuals for daily average $PM_{2.5}$ (10 μg/m^3) and temperature differences (°C) between atmosphere layers at pressures (or heights) of (a) 975 hPa (320 m) and 1000 hPa (110 m), (b) 950 hPa (540 m) and 975 hPa (320 m), (c) 925 hPa (990 m) and 950 hPa (540 m), separately, after time fixed effects of holiday, year–month, day-of-week, are accounted for.

which supplies temperatures every 6 h at 42 pressure levels with a spatial resolution of 0.5° × 0.625°.[23] We extract data from the grids covering Beijing, and we aggregate the temperatures in each layer to daily averages.

Although an inversion blocks atmospheric flow and pollution diffusion, the extent of its effects might also depend on the altitude at which it occurs. Thus, we focus on the first four atmospheric layers closest to the ground, i.e., layers with pressure 1000 hPa (110 m, first layer), 975 hPa (320 m, second layer), 950 hPa (540 m, third layer), and 925 hPa (990 m, fourth layer). Next, we calculate the temperature differences between adjacent layers, i.e., the second layer minus the first layer, the third layer minus the second layer, and the fourth layer minus the third layer. Positive values translate to an inversion, and negative values suggest no inversion. As depicted in Fig. 2, the temperature gradients for all four layers are strongly correlated with the daily air pollution levels. In general, $PM_{2.5}$ concentrations are higher when temperature gradients are larger, that is, when the upper air is warmer than the lower air.

Finally, we construct three dummies that indicate the existence of an inversion in the first, second, or third layer based on the temperature gradients. For example, the indicator for a first-layer inversion equals one if the temperature in the second layer minus the temperature in the first layer is positive and equals zero otherwise. By including these three dummies as IVs, we are able to distinguish between the influences of thermal inversions at varying atmospheric heights on air pollution. We also check the robustness of our results by using alternative IV specifications, such as temperature differences, in Section 4.1.2; the results are robust.

However, a concern arises. Thermal inversions are usually correlated with weather conditions that may directly affect the health outcomes of interest (Chen et al., 2017, 2018). Appendix Figure S4(a) reveals the seasonal pattern in the thermal inversions in Beijing. Such inversions occur more frequently in winter than in summer. To address this potentially confounding issue, we control for daily weather conditions in a flexible way following Chen et al. (2017) by including maximum and minimum temperature bins specified over 5 °C intervals (see Appendix Table S4 for summary statistics on the temperature bins) and second-order polynomial functions for relative humidity, precipitation, and wind speed.

As shown in Appendix Figure S4(b), the seasonality in the thermal inversions fades when conditioned on the set of weather measures. The underlying assumption for our IV approach is that once weather variables as well as the time fixed effects are controlled for, the occurrence of thermal inversions within a short time window is uncorrelated with human health or health care behaviors in any way other than through the effects of air pollution. We also check the robustness of our results using an alternative method for parametrizing our weather controls (see Section 4.1.2); we create bins for all the weather variables, including the maximum and minimum temperature bins specified at every 5 °C interval, relative humidity bins specified at every 10 percent interval, indicators for no precipitation and deciles for positive precipitation, and wind speed deciles. The results are robust to these alternative specifications.

The correlation between air pollution and our thermal inversion instruments is reported in Appendix Table S5. The F test for the joint significance of the instruments shows that the instruments are highly significant.

[23] Data are available at https://disc.sci.gsfc.nasa.gov/datasets/M2I6NPANA_5.12.4/summary?keywords=M2I6NPANA.

3.3. Episode-based analysis

In addition to exceedingly high ambient pollutant concentrations on single days, metropolitan areas suffer from frequent outbreaks of pollution episodes that can last for days or even over a week at a time. To investigate the health impacts of pollution that lasts longer than a single day, we conduct an episode-based analysis. This approach is important for understanding the ramifications of air pollution on health care spending for at least three reasons. First, periodic cycles of high levels of pollution that last well beyond a single day are a defining characteristic of air pollution in developing countries. Second, the duration of these pollution episodes matters greatly for human health. For example, a person might be able to withstand a one-day pollution event but not a three-day pollution episode, with damages accumulating over the course of the event. Third, a clear understanding of how pollution events affect public health has significant implications for relevant public policies meant to combat air pollution. If long-lasting periods of pollution are most costly, then measures to prevent such episodes should receive higher priority than measures meant to reduce the average concentration of $PM_{2.5}$. Beijing, among other cities, has launched a system to alert the public about high levels of air pollution, using three warning levels based on the expected concentration and duration of pollution episodes.[24] Under this warning system, corresponding actions, such as shutting down the sources of pollution and closing schools, are compulsory when conditions are severe. Episode-based estimations of medical costs are essential for evaluating and justifying the costs and benefits of such policies since the health costs of air pollution are underestimated if they do not address the characteristics of frequent and extended pollution events when health costs increase as the pollution event continues over consecutive days.

In this section, we explore how the marginal effects of daily $PM_{2.5}$ concentrations on health care visits and medical costs evolve over the course of an episode. Instead of treating each polluted day identically as in the baseline model, model (1), we allow the marginal effects of $PM_{2.5}$ on health care visits and costs to vary over the evolution of each episode in model (2), with all controls remaining the same as those in model (1).

$$\ln Y_{gt} = (\beta^{c} Clear\ Day_t + \beta^{p} Single\ Polluted\ Day_t + \beta^{1} Day_t^1 + \beta^2 Day_t^2 + \cdots + \beta^N Day_t^N) \times Pollution_t + W'_t\gamma + \delta_g + FE_{time} + \epsilon_{gt} \quad (2)$$

where $Day_t^n, n = 1, 2, \ldots, N$ are a set of dummies indicating that date t is the nth day of an episode if it falls during an episode. In this study, we set $N = 7$, and Day_t^7 indicates the seventh day and all subsequent days. To compare the marginal response to air pollution on a single polluted day and during multiday pollution episodes, the regression also includes two types of "non-episode" days, *clear days* and *single polluted days*, as defined in Section 2.3. To prevent the potential lagged impacts of air pollution from confounding the estimates of *clear days* and *single polluted days*, we drop those days for which the previous day was polluted (daily $PM_{2.5}$ over 50 μg/m^3). As in model (1), to capture potential lagged effects, Y_{gt} denotes the three-day totals for various health outcomes. As a result, two leads of whether variables are controlled, and two leads and two lags of instruments are included in the IV estimation.

Due to the interaction with $Pollution_t$, the coefficient β^c represents the short-term marginal impact of $PM_{2.5}$ levels on *clear days*, β^p captures those during polluted days that are not part of an episode, and $\beta^n (n = 1, 2, \ldots, N)$ measures the varying short-term marginal effects of $PM_{2.5}$ levels conditional on exposure to pollution lasting n consecutive days.[25]

4. Results

4.1. Linear short-term effects of $PM_{2.5}$ concentrations on health care visits and medical costs

4.1.1. Baseline results from the linear effects model

We first examine the linear short-term impacts of $PM_{2.5}$ concentrations on health care utilization for the population enrolled in the UEBMI system. Table 4 presents the OLS and IV estimates of the short-run effects of daily $PM_{2.5}$ exposure on the three-day totals for health care visits, medical expenditures, medication expenditures, out-of-pocket expenditures, and category-specific medication expenses (cardiovascular, respiratory, and other western medications).

As reported in column (1) of Table 4, the OLS estimations show that a 10 μg/m^3 increase in daily average $PM_{2.5}$ exposure is associated with a 0.116 percent increase in the number of health care visits. Exploiting the variation in air pollution induced by the random occurrence of thermal inversions, a 10 μg/m^3 increase in daily $PM_{2.5}$ results in 0.387 percent more health care visits. This

[24] According to the "Heavy Air Pollution Contingency Plan for the Municipality of Beijing", alert levels depend not only on the pollution level but also on the duration of the pollution episode. For example, a Red Alert is issued for any forecasted citywide mean daily air quality index greater than 200 and lasting for four days (96 h) or more for any such index over 300 for two days (48 h) or more, or for any such index reaching 500. Source: http://english.beijing.gov.cn/latest/lawsandpolicies/202007/t20200723_1957677.html.

[25] Adding these interactions increases the number of endogenous variables in our regression, but the endogeneity arises only through the variable $Pollution_t$. We address this endogeneity issue as follows: in the first stage, as in the first stage for the 2SLS estimation of model (1), $PM_{2.5}$ is regressed on the instruments and the controls specified in the baseline regressions in Table 4 (i.e., time fixed effects for holidays, the day of the week, and the year and month, weather variables including the maximum and minimum temperature bins and second-order polynomials for relative humidity, precipitation and wind speed, and two lags of the instruments, as well as two leads of the instruments and weather variables), and the residuals are obtained. The first stage F-Statistic is 19.75, which is reported in Appendix Table S5. In the second stage, equation (2) is estimated by including the residuals obtained in the first stage as a control. Bootstrapped standard errors are clustered at the group-year level and are obtained via 200 replications of the two-stage procedure. Our results are robust to different clustering choices, including two-way clustering at the group-year and the date levels and two-way clustering at the group-year and year–month levels. The estimates using alternatively clustered standard errors are reported in Appendix Figure S5. We find little evidence that the statistical precision of our estimates is greatly affected by serial correlation across groups.

Table 4
Overall short-term effects of daily $PM_{2.5}$ levels on health care visits and costs.

	(1) Health care visits	(2) Total expenses	(3) Medication expenses	(4) Out-of-pocket expenses	(5) Cardiovascular medication	(6) Respiratory medication	(7) Other western medication
OLS estimates							
$PM_{2.5}$	0.00116	0.00163*	0.00141	0.00163*	0.00155*	0.00123	0.00140
	(0.00076)	(0.00092)	(0.00085)	(0.00091)	(0.00086)	(0.00076)	(0.00088)
level dep. mean	52.04	30.56	18.83	8.81	2.85	0.34	7.60
level eff.	0.06	0.05	0.03	0.01	0.00	0.00	0.01
N	16512	16512	16512	16512	16512	16512	16512
IV estimates							
$PM_{2.5}$	0.00387*	0.00376	0.00399*	0.00407*	0.00643**	0.00474**	0.00432*
	(0.00222)	(0.00242)	(0.00237)	(0.00235)	(0.00276)	(0.00193)	(0.00249)
level dep. mean	52.04	30.56	18.83	8.81	2.85	0.34	7.60
level eff.	0.20	0.11	0.08	0.04	0.02	0.00	0.03
N	14040	14040	14040	14040	14040	14040	14040
F-stat.	253	253	253	253	253	253	253

Note: Three-day total health care utilization outcomes (sums of current day and forward two days), based on the population enrolled in Beijing Basic Medical Insurance for Urban Employees from 2013 to 2016, are regressed on daily $PM_{2.5}$ levels (in 10 μg/m^3). The unit of observation is daily aggregates for each demographic group. All the dependent variables are in logarithm forms for health care visits (counts per thousand beneficiaries) and different types of expenditure (CNY per beneficiary) indicated by column titles. Means of dependant variables in level forms and the imputed level effects are also reported. IV estimates are obtained using thermal inversions as the instrument variables. All regressions control for group fixed effects, time fixed effects of holiday, year–month, and day-of-week, and daily weather variables. The weather variables include maximum and minimum temperature bins (in 5°C bins starting from -15°C), second-order polynomials of relative humidity, precipitation, and wind speed. Two leads of weather variables are also included as controls. For OLS estimates, two lags and two leads of $PM_{2.5}$ are included, while for IV estimates, two lags and two leads of instrument variables are included. Instrument variables are thermal inversions measured as three dummies indicating existence of inversion at three atmosphere levels from the ground. The first-stage Cragg–Donald Wald F-statistic is reported for IV estimates. All regressions are weighted by the number of beneficiaries in each group that month. Standard errors are in parentheses and clustered by group-year. Significance levels are indicated by *** 1%, ** 5%, * 10%.

IV estimate is more than three times the magnitude of the OLS estimate. Of the two, we believe that the IV estimate is more realistic because the OLS estimate is biased downward due to the endogeneity caused by omitted variables (described in Section 3.2).

Columns (2)–(4) of Table 4 report the marginal percentage growth in per beneficiary total health care, medication, and out-of-pocket expenses in response to 10 μg/m^3 increments of ambient $PM_{2.5}$. The results for these three expense categories indicate increases of similar magnitudes: 0.376 percent, 0.399 percent, and 0.407 percent. Based on the linear effects estimate of a 0.376 percent increase in three-day total medical expenditures, the short-term medical cost of a 10 μg/m^3 increase in $PM_{2.5}$ exposure amounts to 791.0 million CNY (215.9 million USD) per year for the UEBMI beneficiaries in Beijing. If compared with expenditures under zero pollution, the total medical cost of air pollution for the 19.7 million Beijing UEBMI beneficiaries is 1.83 billion USD per year.

It has been widely documented that fine particulate matter is especially harmful to the respiratory and cardiovascular systems (Dockery et al., 1993; Dominici et al., 2006; Brook et al., 2009). To test whether the impacts of daily $PM_{2.5}$ exposure on cardiovascular and respiratory diseases are evident to a greater degree than those on other diseases, columns (5)–(7) in panel B of Table 4 display estimates for expenditures on cardiovascular medications, respiratory medications, and other western medications. When daily $PM_{2.5}$ worsens by 10 μg/m^3, expenditures on cardiovascular medications and respiratory medications grow significantly — by 0.643 percent and 0.474 percent, respectively. Moreover, expenditures on all other western drugs increase by 0.432 percent, which is less than for cardiovascular and respiratory medications but still significant. This finding is similar to those of Deschênes et al. (2017), who report that air pollution increases expenditures on non-respiratory and non-cardiovascular medications.

4.1.2. Robustness checks

We also conduct a variety of other robustness checks. In the first of these, we examine medical outcomes during an extended time window of seven days, reported in Panel A of Appendix Table S6, and find that the magnitudes of the impacts are close to the three-day benchmark estimates, which suggests that the short-term impacts estimated in the three-day specification are neither obscured by harvesting effects nor seriously underestimated because of lagged effects.

Second, to address the potential concern that some hospital admissions could be prearranged and less influenced by short-term air pollution changes, we exclude inpatient visits from the sample and re-estimate the model using outpatient and ER visits; the results are summarized in Panel B of Appendix Table S6. However, outpatient visits could also be scheduled beforehand, especially in tertiary hospitals with heavy patient loads. Therefore, we report the results from regressions that use only emergency room visits, which are certainly not prearranged, and the coefficients become slightly larger and more statistically significant (Appendix Table S7).

Another issue that might undermine our goal to estimate the short-term impacts of air pollution on medical expenditures is that medications are not necessarily purchased for treatment of acute disorders or immediate use and thus may not respond to short-term pollution. Although we are not able to separate medication expenditures into treatments for acute or chronic disorders, as many medications treat both acute and chronic symptoms, as a robustness check, we instead focus on the medication expenditures for acute upper respiratory infections (AURIs) and use these expenditures as a proxy for spending on immediate-use medications. By following the medicine classifications used by the Beijing Reimbursement Drug List for UEBMI, we are able to identify all the drugs that specifically treat AURIs and aggregate the corresponding daily expenditures in our sample.[26] As reported in Appendix Table S7, AURI medication costs are estimated to increase by 0.695 percent on average in response to a 10 $\mu g/m^3$ higher level of $PM_{2.5}$, which is slightly larger than the estimate using total medication expenses.

Next, we estimate the results using alternative functions for weather controls. Instead of second-order polynomials, we construct bins for relative humidity, precipitation, and wind speed (Panel C of Appendix Table S6).[27]

We then employ alternative approaches to the construction of the thermal inversion variables used in the baseline IV estimation. Instead of using dummy variables indicating whether an inversion has occurred, we check the robustness of our results by directly using the temperature differences between adjacent layers to measure inversions, with larger positive temperature gaps between upper and lower layers indicating stronger thermal inversions that constrain the diffusion of pollution (Panel D of Appendix Table S6).

Finally, we make other adjustments that take into account the potential ramifications of influenza. These controls account for the fact that both influenza and air pollution cause similar respiratory symptoms and that both influenza and air pollution peak in autumn and winter. To account for these patterns, we control for the percentage of influenza-like illnesses (ILI), the major indicator of influenza surveillance provided by the Chinese National Influenza Center (Panel E of Appendix Table S6).

The results of the above checks are robust and consistent with our previous findings. The exact estimates are summarized in Appendix Tables S6–S7.

4.2. Impact of pollution episodes on health care utilization and medical costs

As the above results show, daily air pollution has significant average impacts on health care utilization and corresponding medical expenditures. However, air pollution in Beijing and in many other regions in China is often present for multiple days until strong winds or rainfall clears the air. Negative health effects might accrue over consecutive days of exposure to pollutants. The episode-based analysis, specified in Eq. (2), can inform us about the health impacts of multi-day pollution episodes and can offer some insights into the implied benefits of pollution reductions, not just in terms of daily pollution concentrations but also in terms of the duration of such events.

To present the dynamics of the marginal damages that accrue over a pollution episode more intuitively, Fig. 3 plots the estimates of Eq. (2) and shows results for health care visits, medical expenditures, cardiovascular medication expenditures, and respiratory medication expenditures.[28] In sum, the figure shows the evolution of the marginal impacts of pollution on health over the course of an *episode* as well as the marginal impacts during clear days and single polluted days. There are three main findings. First, $PM_{2.5}$ exposure is found to have a statistically significant impact in most instances. Second, within a pollution episode, starting on the second polluted day, the marginal effects of $PM_{2.5}$ concentrations are larger than those on *Day 1* of the same episode. Third, the increases in health care visits and medical costs caused by a marginal increase in ambient $PM_{2.5}$ levels are the greatest on *Day 7 and later* of a pollution episode and are also larger than those on a *single polluted day* that is not part of a longer-lasting episode of pollution. As shown in Appendix Figure S7, most of the later days during a pollution episode have statistically larger effects than *Day 1*, and *Day 7 and later* has the largest pollution effects out of all other days within an episode, *clear days*, and *single polluted days*. In addition, we provide a statistical test in Appendix Table S8 in which we include an interaction term between pollution and the day of the episode, along with the main pollution term, and we find that the effects of pollution increase significantly as the number of days for the pollution event increases.

These patterns are consistent across our different health outcomes of interest. The results of the episode-based regressions for all medical outcomes are reported in Table 5. Taking cardiovascular medication expenditures as an example, every 10 $\mu g/m^3$ increase in daily $PM_{2.5}$ concentrations causes spending to increase on average by 0.309 percent on *Day 1* of an *episode*, although insignificantly. However, this marginal impact grows to a significant 0.793 percent on *Day 2* of the *episode*, and for the subsequent episode days from *Day 3* to *Day 5*, the marginal impacts remain high, with magnitudes of approximately 0.6–0.8 percent. If a pollution episode continues to *Day 7 and later*, expenditures on cardiovascular medication increase by 1.31 percent in response to an increase of 10 $\mu g/m^3$ in the level of $PM_{2.5}$ concentrations, which is a much larger increase than those on previous episode days.

These results suggest that the health costs induced by air pollution increase as the duration of the pollution event increases over consecutive days. This underscores the rationale behind issuing public health pollution alerts based on both the intensity and the duration of air pollution episodes. Incorporating all *clear days*, *single polluted days* and *pollution episodes* in our sample, a back-of-the-envelope calculation based on the estimates from the episode-based analysis implies that the total medical cost of air pollution, relative to no pollution, for the 19.7 million Beijing UEBMI beneficiaries amounts to 3.15 billion USD per year during our study period.[29] This estimate is 70 percent larger than the linear effects estimates obtained from model (1) (1.83 billion USD),

[26] AURI medicines are recorded as a subcategory under the category of medicines for respiratory diseases.

[27] Available at http://www.chinaivdc.cn/cnic/.

[28] Plots for the other outcome variables (i.e., medication expenditures, out-of-pocket expenditures, and other western medication expenditures) are presented in Appendix Figure S6.

[29] This calculation is based on the estimates in Table 5, column (2). We use the average $PM_{2.5}$ levels for each type of day (*clear days*, *single polluted days* and *pollution episodes*), which are reported in Table 2.

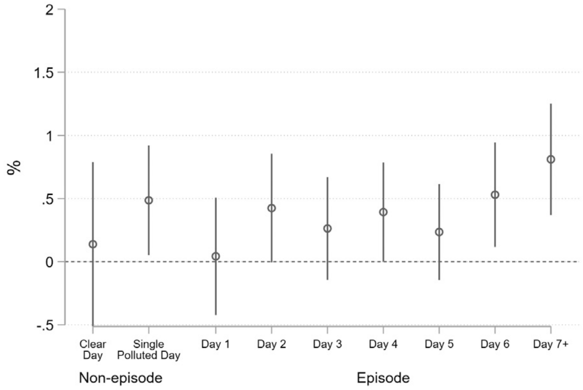

(a) Health care visits

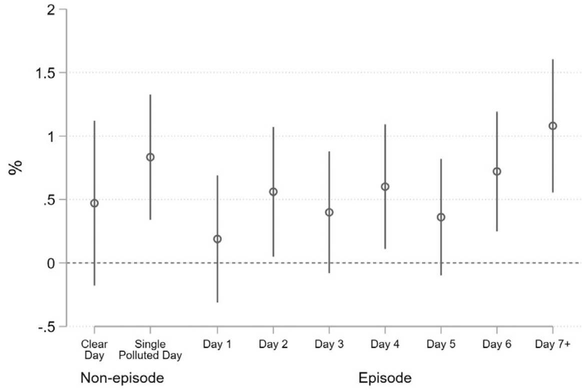

(b) Total medical expenses

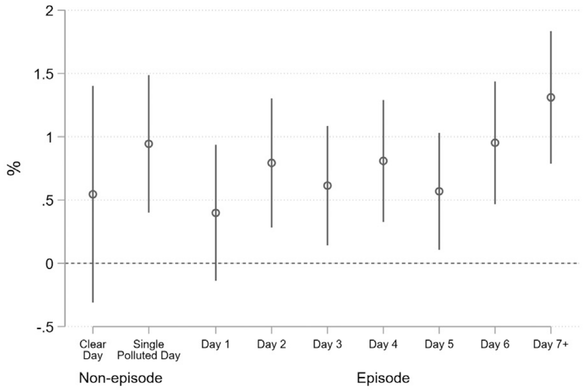

(c) Cardiovascular medication expenses

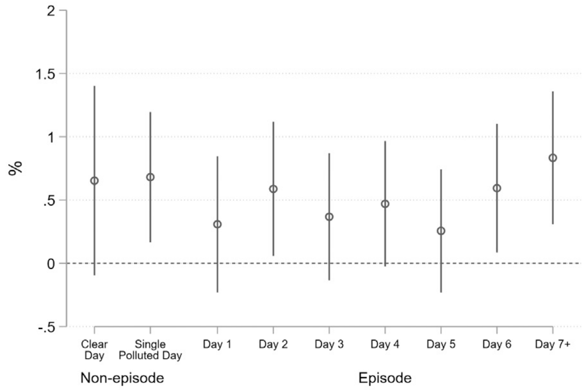

(d) Respiratory medication expenses

Fig. 3. Marginal effects of $PM_{2.5}$ levels on health care visits and costs over pollution episodes.
Note: Figures plot the coefficients of $PM_{2.5}$ levels (measured in 10 $\mu g/m^3$) interacting with indicators of the *clear day*, the *single polluted day*, the nth *day* of an air pollution episode (see definitions of these indicators in Section 2.3), for three-day totals of (a) health care visits, (b) total medical expenses, (c) cardiovascular medication expenses, and (d) respiratory medication expenses, in the logarithm forms. Estimates are obtained by a two-stage approach. In the first stage, $PM_{2.5}$ is regressed on inversion instruments, controlling for group fixed effects, time fixed effects, weather variables, two lags of instruments, and two leads of instruments and weather variables. Time fixed effects and weather controls are the same as those specified in regressions in Table 4. In the second stage, three-day health outcomes are regressed on interactive terms of concurrent $PM_{2.5}$ and all episode-related indicators, controlling for residuals of $PM_{2.5}$ predicted from the first stage, as well as identical controls in the first stage. Standard errors are obtained using bootstrap with 200 replications of the two-stage procedure, clustered by group-year. Whiskers stand for 95% confidence intervals.

which suggests that omitting the variation in the magnitude of the marginal effects over the course of pollution episodes results in a substantial underestimation of the actual health costs of air pollution in settings that frequently experience long-lasting and high-level pollution events.

4.3. Heterogeneous health care responses to pollution episodes

In the episode-based analysis, the average short-term health impacts of $PM_{2.5}$ exposure over the course of a pollution episode are estimated for the whole population enrolled in the Beijing UEBMI program. Using the detailed information on personal characteristics provided, we further investigate the short-term effects of air pollution episodes on medical costs for different demographic groups to understand the heterogeneity in responses to air pollution, which is crucial for estimating the distribution of health costs among different populations and accurately evaluating the social costs of air pollution.

There are at least two potential sources of heterogeneity. First, actual exposure to air pollution might differ significantly across individuals because of differences in working environments and defensive behaviors practiced. Second, people's inherent susceptibility to the same exposure levels might also diverge due to biological differences resulting from gender, age, and unobserved genetic factors.

To investigate the heterogeneity in the health impacts arising from $PM_{2.5}$ exposure over the course of a pollution episode, we divide the studied population into two groups (1) by age and (2) by gender and allow the βs in model (2) to differ between groups by interacting either the age group indicator (below/above age of 50) or the gender indicator with the pollution terms. Fig. 4

Table 5
Marginal effects of $PM_{2.5}$ levels on health care visits and costs over pollution episodes.

	(1) Health care visits	(2) Total expenses	(3) Medication expenses	(4) Out-of-pocket expenses	(5) Cardiovascular medication	(6) Respiratory medication	(7) Other western medication
Clear day	0.00139	0.00471	0.00359	0.00256	0.00546	0.00653*	0.00696**
	(0.00332)	(0.00331)	(0.00374)	(0.00340)	(0.00437)	(0.00382)	(0.00313)
Single polluted day	0.00487**	0.00834***	0.00731***	0.00747***	0.00944***	0.00681***	0.00897***
	(0.00222)	(0.00252)	(0.00271)	(0.00261)	(0.00277)	(0.00263)	(0.00242)
Day 1	0.000429	0.00189	0.000922	0.00121	0.00399	0.00309	0.00322
	(0.00237)	(0.00255)	(0.00261)	(0.00253)	(0.00274)	(0.00274)	(0.00234)
Day 2	0.00425*	0.00561**	0.00502*	0.00503**	0.00793***	0.00588**	0.00632***
	(0.00219)	(0.00261)	(0.00259)	(0.00252)	(0.00260)	(0.00271)	(0.00242)
Day 3	0.00263	0.00399	0.00337	0.00360	0.00614**	0.00368	0.00482**
	(0.00207)	(0.00245)	(0.00237)	(0.00238)	(0.00240)	(0.00256)	(0.00225)
Day 4	0.00393**	0.00601**	0.00518**	0.00575**	0.00809***	0.00470*	0.00651***
	(0.00200)	(0.00250)	(0.00239)	(0.00241)	(0.00246)	(0.00253)	(0.00234)
Day 5	0.00235	0.00361	0.00300	0.00357	0.00569**	0.00257	0.00389*
	(0.00193)	(0.00234)	(0.00237)	(0.00234)	(0.00236)	(0.00248)	(0.00221)
Day 6	0.00531**	0.00721***	0.00642**	0.00674***	0.00952***	0.00594**	0.00709***
	(0.00211)	(0.00241)	(0.00258)	(0.00251)	(0.00248)	(0.00259)	(0.00232)
Day 7+	0.00811***	0.0108***	0.00969***	0.00905***	0.0131***	0.00834***	0.00983***
	(0.00225)	(0.00268)	(0.00276)	(0.00271)	(0.00267)	(0.00268)	(0.00255)
Observations	10932	10932	10932	10932	10932	10932	10932

Note: Table presents the coefficients of daily $PM_{2.5}$ levels (in 10 $\mu g/m^3$) interacting with indicators of the *clear day*, the *single polluted day*, the *n*th day of an air pollution episode (see definitions of these indicators in Section 2.3), for dependent variables of health care visits and different types of expenditure in the logarithm forms. Estimates are obtained by a two-stage approach. In the first stage, $PM_{2.5}$ is regressed on inversion instruments, controlling for group fixed effects, time fixed effects, weather variables, two lags of instruments, and two leads of instruments and weather variables. Time fixed effects and weather controls are the same as those specified in regressions in Table 4. In the second stage, three-day health outcomes are regressed on interactive terms of current-day $PM_{2.5}$ and all episode-related indicators, controlling for residuals of $PM_{2.5}$ predicted from the first stage, as well as identical controls in the first-stage. Standard errors are in parentheses and are obtained using bootstrap with 200 replications of the two-stage procedure, clustered by group-year. Significance levels are indicated by *** 1%, ** 5%, * 10%.

depicts the heterogeneity in the dynamic health effects of $PM_{2.5}$ exposure over a pollution episode on medical costs for different population groups. In terms of the percentage changes in medical cost resulting from marginal increases in $PM_{2.5}$ levels, the older group (individuals aged 50 and above) experienced slightly and insignificantly larger impacts than the younger group (individuals under 50), while no heterogeneity was found between the male and female groups.

4.4. The role of $PM_{2.5}$ levels in health impacts over a pollution episode

The results of the episode-based estimation reveal that the marginal impacts of $PM_{2.5}$ exposure are not identical across different types of polluted days, which suggests that the marginal effects of air pollution vary with the episode duration. This variation in the magnitude of marginal impacts matters greatly for accurately estimating the actual cost of air pollution in settings with frequent high-level and long-lasting pollution events. As the main purpose of the episode analysis is to investigate the variation in the effects of pollution over the course of a pollution episode and to provide a more accurate policy-relevant estimate of the health costs, we do not intend to draw any strong conclusions on the medical reasons why such variation in the effect of pollution exists. The exact reasons could be related to either the accumulated effects of pollution exposure in the short term or to the effects of higher pollution levels during the later days of an episode due to the pollution accumulation itself. In Appendix Figure S8, we plot the distribution of pollution levels for each day of the pollution episodes, which reveals that pollution often accumulates and worsens during the later days of an episode and then starts to dissipate. Therefore, the higher marginal effects during the later days of an episode could reflect the mixed effects of higher concentration levels and accumulated exposure due to the longer duration of the pollution event.

To investigate whether pollution concentration has a nonlinear effect on health outcomes, we first estimate a dose–response curve for health care visits and total expenditures separately based on pollution level bins and report the results in Appendix Figure S9. The first bin (daily $PM_{2.5} \leq 15\ \mu g/m^3$) is the omitted category, and the other coefficients are estimates of the differences in the effect of pollution between higher-pollution days and the days that fall in the first bin. The pattern in the figure reveals that the effect of pollution first increases, then flattens when it passes through the third bin, and then increases sharply again when pollution rises to the level in the largest bin.[30]

[30] A similar pattern between mortality and pollution exposure has been documented in other epidemiological studies (Pope III et al., 2011, 2015).

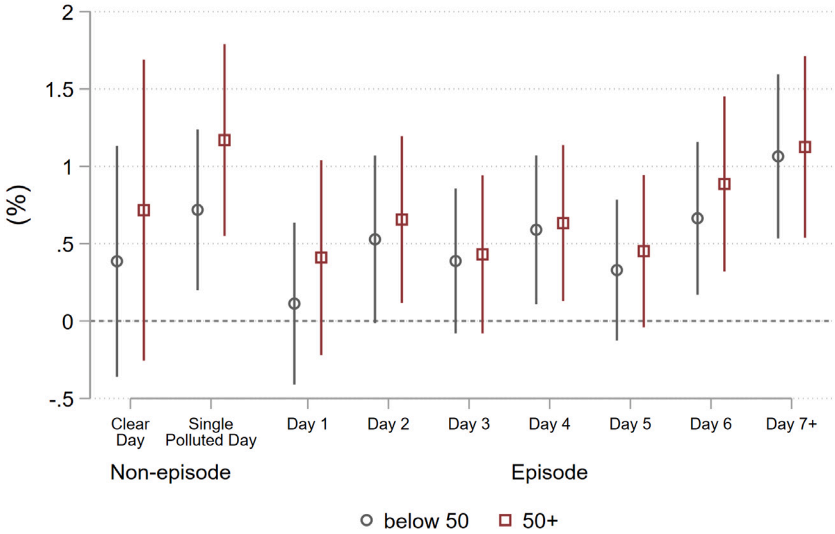

(a) Heterogeneity by age

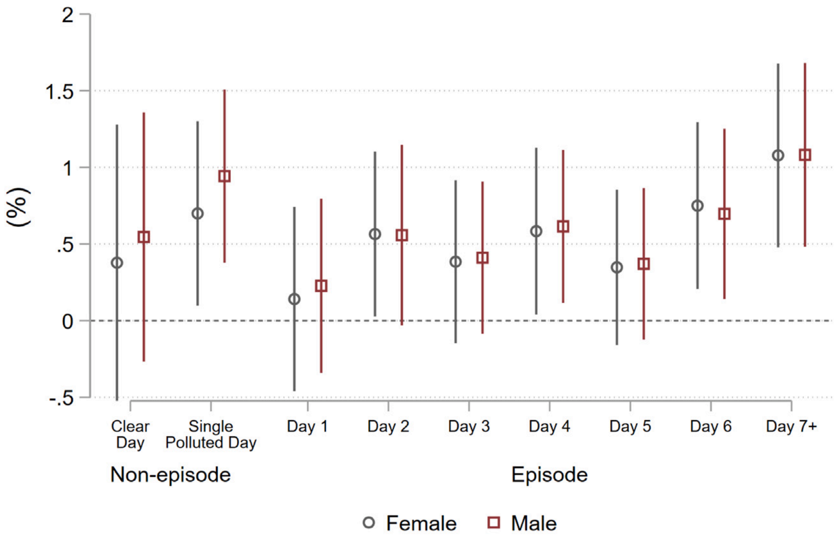

(b) Heterogeneity by gender

Fig. 4. Heterogeneity analysis of short-term health effects of pollution episodes by age and by gender.
Note: Estimates of marginal short-term effects of $PM_{2.5}$ levels (measured in 10 μg/m^3) on three-day total medical costs over pollution episodes are plotted for subpopulations (a) by age and (b) by gender. We interact the group indicator based on age or gender with pollution episode terms specified in Eq. (2) to allow coefficients of interest to differ between groups. Whiskers stand for 95% confidence intervals.

Although our primary focus is to estimate the variation in the marginal effects of pollution across the different days of an episode, we provide an additional analysis by examining whether the nonlinear effects of pollution concentrations could affect the varying marginal effects of pollution exposure that we find over the course of an episode. More specifically, we first divide the episodes into more-polluted episodes and less-polluted episodes based on the average pollution level within each episode. We then investigate the differences in the health response over the course of the episode between the two types of episodes, which are summarized in Fig. 5. As the figure reveals, from *Day 3* to *Day 6* of a pollution episode, the marginal impacts of $PM_{2.5}$ are slightly larger for those days that belong to episodes with an average $PM_{2.5}$ concentration above 148 μg/m^3 (the 75th percentile of the episode-average $PM_{2.5}$ levels) compared with those belonging to less polluted episodes. This discrepancy is larger for *Day 2* and for *Day 7 and later*. However, even for those episodes with much higher pollution levels, there is still a variation in the health effects over the course of an episode, and thus the duration of pollution episodes is an important dimension of the non-constant effects of pollution exposure, the main finding in the episode analysis.

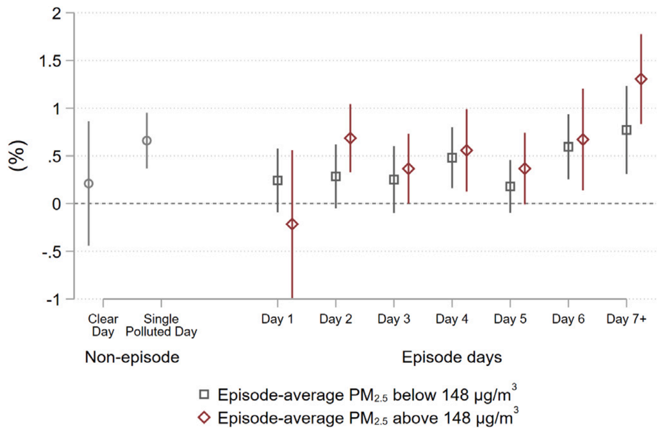

Fig. 5. Heterogeneous marginal health effects of pollution episodes based on episode-average $PM_{2.5}$ levels.
Note: Estimates of marginal effects of $PM_{2.5}$ levels (measured in 10 μg/m^3) on three-day total medical cost over episode days are plotted based on episode-average pollution level. We interact the indicator of whether episode-average $PM_{2.5}$ level exceeding 148 μg/m^3 (the 75th percentile value) with pollution episode terms specified in Eq. (2) to allow coefficients of interest to vary depending on the episode pollution level. Health effects of clear days and single polluted days are also plotted in the figure for better comparison of the magnitudes of pollution effects. Whiskers stand for 95% confidence intervals.

5. Conclusions

We estimate the medical costs of exposure to air pollution by using detailed data that document the medical expenses incurred by beneficiaries enrolled in the Urban Employee Basic Medical Insurance (UEBMI) program, which covers 19.7 million residents in Beijing. We find a strong, positive correlation between air pollution and health care use in the short term. Our linear effects estimation shows that a 10 μg/m^3 increase in $PM_{2.5}$ concentrations corresponds to a 0.376 percent per beneficiary increase in three-day medical expenditures.

Importantly, our episode-based analysis implies that the short-term health impacts of $PM_{2.5}$ escalate as pollution lingers. When relatively high air pollution levels persist for more days within a multi-day pollution episode, larger health responses emerge, as evidenced by health care visits and medical costs. A back-of-the-envelope calculation based on our episode analysis shows that air pollution generated approximately 3.15 billion USD in medical expenditures (for UEBMI beneficiaries and their covered expenses relative to zero pollution) annually in the city. This health-cost estimate is approximately 70 percent larger than the estimated cost based on a linear effects model. This discrepancy suggests that it is critical to account for the variation in the marginal effects of pollution exposure over the course of the episode in order to accurately estimate total health costs in a setting in which high-level and long-lasting pollution events occur frequently.

If any intervention could prevent a pollution episode from lasting more than six days or could reduce the pollution level on *Day 7 and later* during our defined pollution episodes to the average level on *clear days*, the short-term cost of air pollution on medical expenditures would drop to 2.43 billion USD a year, a 23 percent decrease compared with the status quo. As a comparison, if the same number of *single polluted days* were turned into an average *clear day*, the reduction in the medical cost of air pollution would be nine percent. Therefore, to reduce the health costs of air pollution, the prevention of long-lasting pollution events should be given greater priority in combination with methods for reducing average pollution levels. Reducing the number of long-lasting pollution events may require more prudent strategies and preventive methods (such as issuing warnings a few days earlier), as pollutants often accumulate and take time to dissipate. Governments should consider the effects of both the intensity and duration of air pollution when designing and adjusting warning systems to reduce the public's exposure to pollution.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jeem.2022.102680.

References

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. Econ. J. 126 (591), 257–280. http://dx.doi.org/10.1111/ecoj.12273.

Barwick, P.J., Li, S., Rao, D., Zahur, N.B., 2018. The Healthcare Cost of Air Pollution: Evidence from the World's Largest Payment Network. Technical Report w24688, National Bureau of Economic Research, Cambridge, MA, http://dx.doi.org/10.3386/w24688.

Brook, R.D., Urch, B., Dvonch, J.T., Bard, R.L., Speck, M., Keeler, G., Morishita, M., Kaciroti, N., Harkema, J., Corey, P., Silverman, F., Wellenius, G., Mittleman, M.A., Rajagopalan, S., Brook, J.R., 2009. Insights into the mechanisms and mediators of the effects of air pollution exposure on blood pressure and vascular function in healthy humans. Hypertension 54 (3), 659–667. http://dx.doi.org/10.1161/HYPERTENSIONAHA.109.130237.

Burns, L.R., Liu, G.G., 2017. China's healthcare industry: A system perspective. In: Liu, G.G., Burns, L.R. (Eds.), China's Healthcare System and Reform. Cambridge University Press, Cambridge, pp. 3–30. http://dx.doi.org/10.1017/9781316691113.003.

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. Q. J. Econ. 118 (3), 1121–1167. http://dx.doi.org/10.1162/00335530360698513.

Chen, S., Oliva, P., Zhang, P., 2017. The Effect of Air Pollution on Migration: Evidence from China. Technical Report w24036, National Bureau of Economic Research, http://dx.doi.org/10.3386/w24036.

Chen, S., Oliva, P., Zhang, P., 2018. Air Pollution and Mental Health: Evidence from China. Technical Report w24686, National Bureau of Economic Research, Cambridge, MA, http://dx.doi.org/10.3386/w24686.

Currie, J., Neidell, M., 2005. Air pollution and infant health: What Can We Learn from California's recent experience? Q. J. Econ. 120 (3), 1003–1030. http://dx.doi.org/10.1093/qje/120.3.1003.

Currie, J., Neidell, M., Schmieder, J.F., 2009. Air pollution and infant health: Lessons from New Jersey. J. Health Econ. 28 (3), 688–703. http://dx.doi.org/10.1016/j.jhealeco.2009.02.001.

Currie, J., Walker, R., 2011. Traffic congestion and infant health: Evidence from E-ZPass. Am. Econ. J.: Appl. Econ. 3 (1), 65–90. http://dx.doi.org/10.1257/app.3.1.65.

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. Amer. Econ. Rev. 109 (12), 4178–4219. http://dx.doi.org/10.1257/aer.20180279.

Deschênes, O., Greenstone, M., Shapiro, J.S., 2017. Defensive investments and the demand for air quality: Evidence from the NOx budget program. Amer. Econ. Rev. 107 (10), 2958–2989. http://dx.doi.org/10.1257/aer.20131002.

Dockery, D.W., Pope, C.A., Xu, X., Spengler, J.D., Ware, J.H., Fay, M.E., Ferris, B.G., Speizer, F.E., 1993. An association between air pollution and mortality in six U.S. cities. N. Engl. J. Med. 329 (24), 1753–1759. http://dx.doi.org/10.1056/NEJM199312093292401.

Dominici, F., Peng, R.D., Bell, M.L., Pham, L., McDermott, A., Zeger, S.L., Samet, J.M., 2006. Fine particulate air pollution and hospital admission for cardiovascular and respiratory diseases. JAMA : J. Am. Med. Assoc. 295 (10), 1127–1134. http://dx.doi.org/10.1001/jama.295.10.1127.

Fu, S., Viard, V.B., Zhang, P., 2021. Air pollution and manufacturing firm productivity: Nationwide Estimates for China. Econ. J. 131 (640), 3241–3273. http://dx.doi.org/10.1093/ej/ueab033.

Guo, S., Hu, M., Zamora, M.L., Peng, J., Shang, D., Zheng, J., Du, Z., Wu, Z., Shao, M., Zeng, L., Molina, M.J., Zhang, R., 2014. Elucidating severe urban haze formation in China. Proc. Natl. Acad. Sci. 111 (49), 17373–17378. http://dx.doi.org/10.1073/pnas.1419604111.

He, J., Liu, H., Salvo, A., 2019. Severe air pollution and labor productivity: Evidence from industrial towns in China. Am. Econ. J.: Appl. Econ. 11 (1), 173–201. http://dx.doi.org/10.1257/app.20170286.

Ito, K., Zhang, S., 2019. Willingness to pay for clean air: Evidence from air purifier markets in China. J. Polit. Econ. 128 (5), 1627–1672. http://dx.doi.org/10.1086/705554.

Knittel, C.R., Miller, D.L., Sanders, N.J., 2015. Caution, drivers! children present: Traffic, pollution, and infant health. Rev. Econ. Stat. 98 (2), 350–366. http://dx.doi.org/10.1162/REST_a_00548.

Moretti, E., Neidell, M., 2011. Pollution, health, and avoidance behavior: Evidence from the ports of Los Angeles. J. Hum. Resour. 46 (1), 154–175. http://dx.doi.org/10.1353/jhr.2011.0012.

Pope III, C.A., Burnett, R.T., Turner, M.C., Cohen, A., Krewski, D., Jerrett, M., Gapstur, S.M., Thun, M.J., 2011. Lung cancer and cardiovascular disease mortality associated with ambient air pollution and cigarette smoke: Shape of the exposure–response relationships. Environ. Health Perspect. 119 (11), 1616–1621. http://dx.doi.org/10.1289/ehp.1103639.

Pope III, C.A., Cropper, M., Coggins, J., Cohen, A., 2015. Health benefits of air pollution abatement policy: Role of the shape of the concentration–response function. p. 8.

Schlenker, W., Walker, W.R., 2016. Airports, air pollution, and contemporaneous health. Rev. Econom. Stud. 83 (2), 768–809. http://dx.doi.org/10.1093/restud/rdv043.

Tan, T., Hu, M., Li, M., Guo, Q., Wu, Y., Fang, X., Gu, F., Wang, Y., Wu, Z., 2018. New insight into PM2.5 pollution patterns in Beijing based on one-year measurement of chemical compositions. Sci. Total Environ. 621, 734–743. http://dx.doi.org/10.1016/j.scitotenv.2017.11.208.

Williams, A.M., Phaneuf, D.J., Barrett, M.A., Su, J.G., 2019. Short-term impact of PM2.5 on contemporaneous asthma medication use: Behavior and the value of pollution reductions. Proc. Natl. Acad. Sci. 116 (12), 5246–5253. http://dx.doi.org/10.1073/pnas.1805647115.

Yu, H., 2015. Universal health insurance coverage for 1.3 billion people: What accounts for China's success? Health Policy 119 (9), 1145–1152. http://dx.doi.org/10.1016/j.healthpol.2015.07.008.

Zhang, J., Mu, Q., 2018. Air pollution and defensive expenditures: Evidence from particulate-filtering facemasks. J. Environ. Econ. Manag. 92, 517–536. http://dx.doi.org/10.1016/j.jeem.2017.07.006.