



Regular Article

The effect of air pollution on migration: Evidence from China

Shuai Chen^a, Paulina Oliva^b, Peng Zhang^{c,*}^a China Academy for Rural Development (CARD), Academy of Social Governance, School of Public Affairs, Zhejiang University, China^b University of Southern California, BREAD, and NBER, USA^c School of Economics and Management, The Chinese University of Hong Kong, Shenzhen, Shenzhen Finance Institute, Shenzhen Research Institute of Big Data, China

ARTICLE INFO

JEL classification:

O15

Q53

Q56

J24

Keywords:

Air pollution

Migration

Human capital

Avoidance behavior

ABSTRACT

This paper looks at the effects of air pollution on migration in China using changes in the average strength of thermal inversions over five-year periods as a source of exogenous variation for medium-run air pollution levels. Our findings suggest that air pollution is responsible for large changes in inflows and outflows of migration in China. Specifically, we find that a 10 percent increase in air pollution, holding everything else constant, is capable of reducing population through net outmigration by about 2.8 percent in a given county. We find that these inflows are primarily driven by well-educated people at the beginning of their professional careers. We also find a strong gender asymmetry in the response of mid-age adults that suggests families are splitting across counties to protect vulnerable members of the household. Our results are robust to different specifications, including a spatial lag model that accounts for localized migration spillovers and spatially correlated pollution shocks.

1. Introduction

Air pollution has been shown to have causal impacts along an array of health and economic dimensions: infant and adult mortality, hospitalization rates, health expenditures, mental health, hours worked, labor productivity, labor market decisions, test scores, and income (e.g., see [Chay and Greenstone, 2003](#); [Graff Zivin and Neidell, 2012](#); [Hanna and Oliva, 2015](#); [Ebenstein et al., 2016](#); [Borgschulte et al., 2018](#); [Deryugina et al., 2019](#); [Fu et al., 2021](#); [Molina, 2021](#)). Many of these studies take place in middle-income countries, where air pollution is now considered the biggest environmental risk to human health.

These results suggest that the total cost of air pollution is quite large as a share of income per-capita, although an aggregation exercise is difficult due to differences in context, methodologies, and pollutant measures across studies. One solution to the aggregation problem is offered by the hedonic method which should reflect all costs of air pollution that are known to individuals ([Chay and Greenstone, 2005](#)). [Bayer et al. \(2009\)](#) note that costs associated with re-location might cause hedonic estimates to deviate from the willingness to pay for air pollution. In addition, housing markets and location decisions in the developing world are often distorted by market failures and regulation, causing further departures from the assumptions underlying hedonic methods. This is especially salient in China, where migration decisions

have been heavily constrained by the household registration (*hukou*) system ([Kinnan et al., 2018](#)) and land markets are subject to the discretion of government officials and corruption ([Chen and Kai-sing Kung, 2019](#)). However, the perception of air pollution costs is still likely to be reflected in the key economic decisions behind hedonic methods: re-location and migration.

Studying how migration decisions are affected by pollution in the developing world offers us a window into the air-pollution costs that are internalized by the population through semi-permanent adaptation measures. Also, zooming into the demographic composition of these flows helps us understand how the willingness to pay for air quality differs across socio-economic groups and how pollution-related migration can change the composition of the labor force across cities ([Hanlon, 2020](#); [Heblich et al., 2021](#)). Our results also contribute to a sizable literature on the factors that determine migration decisions ([Borjas, 1999, 2015](#)). In this literature, the emphasis has been placed on traditional economic factors, such as income, wages, and networks ([Clark et al., 2007](#); [Pedersen et al., 2008](#); [Kinnan et al., 2018](#)). Although recent literature has paid more attention to environmental factors, most of these studies focus on weather ([Feng et al., 2010, 2015](#); [Cai et al., 2016a](#); [Jesso et al., 2018](#)). As our results show, migration flows related to air pollution are of similar magnitude to those projected in response to plausible climate change scenarios ([Feng et al., 2015](#)).

* Corresponding author.

E-mail addresses: shuaichen@zju.edu.cn (S. Chen), olivaval@usc.edu (P. Oliva), zhangpeng@cuhk.edu.cn (P. Zhang).<https://doi.org/10.1016/j.jdevec.2022.102833>

Received 31 July 2021; Received in revised form 26 January 2022; Accepted 28 January 2022

Available online 2 February 2022

0304-3878/© 2022 Elsevier B.V. All rights reserved.

To the best of our knowledge, we are the first to estimate the causal effect of air pollution changes on migration flows.¹ The empirical challenges associated with studying migration responses to air pollution are two. First, as migration involves a large fixed cost and some of it is irretrievable, it is likely to respond slowly to changes in air pollution exposure. Thus, the empirical challenges of estimating the causal effects of air pollution on migration are similar to the challenges of estimating any medium to long-run impacts of air pollution: exogenous cross-sectional or mid-run variation in air pollution is hard to come by. In its absence, estimates may be confounded by unmeasured joint determinants of air pollution and migration. For example, economic activity, which has been shown to attract immigrants (Borjas, 1999; Clark et al., 2007), is also highly correlated with air pollution. Thus, as we demonstrate in this paper, an OLS regression of migration on air pollution yields a coefficient that could be (wrongly) interpreted as pollution attracting immigrants. The second challenge is data constraints when studying migration decisions. Data that can track the residence of an individual over time is hard to come by at the scale that would be required to pick up responses of migration to air pollution.

Our approach to overcoming the first empirical challenge is to use five-year variation in the average strength of thermal inversions within counties. A thermal inversion refers to an abnormal temperature-altitude gradient, where air gets hotter instead of cooler with altitude and traps pollutants near the ground. Thermal inversions have been used to study short-run effects of air pollution on infant and adult mortality (Arceo-Gómez et al., 2016; Jans et al., 2018), labor productivity (Fu et al., 2021), and mental health (Chen et al., 2018). Thermal inversions are a useful source of exogenous variation in air pollution as they emerge independently of pollution sources and human activity.²

We overcome the second challenge, the data constraints on migration decisions, by integrating aggregated and individual-level information from the Population Census in China to construct five-year flows of migration at the county level between 1996 and 2010. Using census questions that are common across all census rounds, we construct two separate measures of migration flows at the county level: net-outmigration and un-registered (floating) immigration.

A recent paper, Khanna et al. (2021), has also documented the relationship between migration and pollution in the context of China, and estimates the welfare consequences of migration decisions that result in misallocation of human capital. Like our paper, Khanna et al. (2021) also uses thermal inversions as one of their identification strategies. Our paper differs from Khanna et al. (2021) in two important ways (1) we study migration decisions since 1995 while they focus on recent migration, and (2) we restrict the variation in thermal inversions to the changes in frequency within a county. Our view is that cross-sectional variation in the frequency of thermal inversions can generate other differences across counties, e.g. in stringency of local environmental policy, which in turn may have an independent impact on migration flows.³

Importantly, the use of within-county changes in thermal inversion patterns as a source of variation in air pollution means that the response we observe is one that corresponds to changes in pollution that are out of synchrony with respect to changes neighboring counties. This is an important consideration when interpreting our results, as pollution was on an upward trajectory over most counties simultaneously during the period of our study. Our estimates can thus be interpreted as the partial

equilibrium effect of air pollution on migration.

Our findings suggest that air pollution is responsible for significant inflows and outflows of migration in China's counties. Specifically, we find that a 10 percent county-level increase in air pollution leads to a 2.8 percent reduction in population. Of this change, about half corresponds to reduced immigration by floating migrants (about 70% of migrants in our data). When applying our point estimates to all county-level changes in air pollution that were uncorrelated across counties, we find that pollution of this sort can produce a standard deviation (SD) in net-outmigration rates of 3 percentage points, while the SD in observed net-outmigration rates is 16.

We find that these migration responses are primarily driven by well-educated people at the beginning of their professional careers. We also find that females between 30 and 45 years of age, but not men, migrate in response to air pollution twice as much as the average adult. Our results are robust to different specifications, including a spatial lag model that allows for spillovers and spatial correlation, different weather controls, and different forms of error variance.

2. Empirical background

2.1. Migration and Household Registration System in China

Migration typically refers to the permanent or long-term changes of the place of residence. Unlike other countries in which people can usually migrate freely, China implements the Household Registration System, or *hukou* system. The *hukou* system keeps a record of legal address and family relations for every citizen from birth to death. Furthermore, it divides people into rural and urban citizens according to their parents or the place of birth, and those in the cities usually enjoy privileges of local employment, education, health care, and social welfare. There are certain requirements for changing registered residence, such as owning a permanent house in the area where a person has migrated to, having a stable occupation and stable income, and having good education and talents.⁴ Therefore, there are two types of migrants in China. The floating population, or people who move while leaving their *hukou* at the origin and the registered migrants, who change their *hukou* to match their destination.

In this paper, we have two measures of migration. The first is an approximate net-outmigration ratio over five years. Typically, the net-outmigration ratio is defined as the percent of population leaving the county net of new arrivals and deaths within a given period (Passel et al., 2004; Feng et al. 2010, 2015). The second measurement of migration is destination-based floating immigration, or immigration of those who are surveyed away from their *hukou*. Section 3.3 details how we construct these measures.

Fig. 1 depicts the migration patterns for each county in China over the period 1996–2010 measured by net-outmigration ratio. A positive net-outmigration ratio, (yellow) means that outflows are larger than inflows. A negative net-outmigration ratio (blue) means the opposite. In general, the metropolitan areas especially three economic regions in China—the Yangtze River Delta (Shanghai, Jiangsu, and Zhejiang), the Pearl River Delta (Guangdong), and the Jing-Jin-Ji Area (Beijing, Tianjin, and Hebei) and other coastal areas—attract a large share of migrants. There are a few exceptions to this pattern in the northwest (the Xinjiang Uyghur Autonomous Region, Qinghai, Gansu, and the Inner Mongolia Autonomous Region) and the Tibet Autonomous Region, where income is lower but inflows are plausibly driven by abundant natural resources and the China Western Development policy.

¹ We provide the formal mapping between our approach and a model of county-to-county flows in Online Appendix III.

² We confirm this in two ways: by showing that long-run thermal inversion and pollution trends are uncorrelated at the local level; and by comparing the frequencies of thermal inversions over the days of the week, which have different emission profiles (Jans et al., 2018).

³ See Section 2.2 for a summary of environmental policies in China in response to air pollution.

⁴ See http://www.gov.cn/xinwen/2014-07/30/content_2727331.htm (in Chinese).




Fig. 1. Migration in China measured by Net Outmigration Ratio (1996–2010). Notes: This figure depicts the average migration for each county in China over the period 1996–2010 measured by net outmigration ratio, which is the percent of population leaving the county net of new arrivals and approximate deaths.

2.2. Air pollution, awareness and avoidance behavior

Over the past decades, air quality has increasingly deteriorated in China, causing increasing concern on China's public health and economic development (Ebenstein et al., 2015). Fig. 2 plots the county-average concentrations measured in microgram per cubic meter ($\mu\text{g}/\text{m}^3$) of $\text{PM}_{2.5}$ in Panel A in China in each year over the period 1980–2015. Two red vertical lines highlight our study period: 1996–2010. The blue vertical line indicates the year of 2001, when China joined the World Trade Organization (WTO). The concentrations of $\text{PM}_{2.5}$ have significantly increased over the period, in particular after 2001, when China became “the world's factory”. In 2015, the average concentration is $66.90 \mu\text{g}/\text{m}^3$, which is nearly 7 times higher than the standard of $10 \mu\text{g}/\text{m}^3$ of annual mean recommended by the WHO (WHO, 2005).

Even though most regions in China experienced increases in pollution between 1996 and 2010, regional policy differences as well as differences in meteorological conditions led to substantial heterogeneity in pollution changes over time. Fig. 3 shows a map of local changes in pollution. As in our estimation we control for nation-wide changes as well as county fixed effects, this map is helpful to illustrate that there is a considerable amount of remaining variation in pollution. Out of this remaining variation, our IV strategy will ensure that we only use the one due to local variation in thermal inversion strength over time.

Concerns about air pollution spurred environmental policy as early as the mid 80's. In 1987, the Air Pollution Prevention and Control Law (APPCL) was enacted. In 1998, the Two Control Zone (TCZ) policy, which aimed to reduce SO_2 and acid rain was implemented. 175 out of 380 prefectures were designated as TCZ cities and faced tighter environmental regulations, such as using clean coal for power plants or installing sulfur-scrubbers.⁵ In 2000, China revised the APPCL, to control air pollution in 47 key cities, which are mainly the municipal cities, provincial capitals, coastal cities, and key tourist cities. 66 additional cities were included in 2003. In 2005, the Chinese Government included

⁵ See http://www.gov.cn/xgk/pub/govpublic/mrlm/201011/t20101122_62944.html (in Chinese). Several studies have estimated the effect of this policy on infant mortality (Tanaka, 2015) and foreign direct investment (Cai et al., 2016b).

“environmental protection” in the evaluation of government officials.⁶

The Government also mandated public information about air pollution starting in 1998. Importantly, the accuracy of early disclosed data has been put into question and recent papers suggest political motives behind the attempts to conceal high concentrations from official records (Ghanem and Zhang, 2014). In 1998, 113 key cities were required to disclose weekly pollution data, and in 2000, daily pollution data. During this period cities were only required to report the air pollution index (API) –a piece-wise linear transformation of PM_{10} , SO_2 , and NO_2 . In 2008, the U.S. Embassy started to report and publish $\text{PM}_{2.5}$ concentrations. In 2013, the Ministry of Environmental Protection (now called Ministry of Ecology and Environment) started to report daily concentration for the six air pollutants including $\text{PM}_{2.5}$, PM_{10} , O_3 , SO_2 , NO_2 , and CO across more than 1000 pollution stations.

The early adoption of some air pollution control policies as well as the data disclosure mandates, along with the efforts to conceal information, suggest that air pollution was a source of concern for the government and the public in China at least as early as the mid 90's. And, although in this study we will be measuring air pollution as $\text{PM}_{2.5}$, which was not widely monitored until 2008, we show in the Online Appendix I (Tables A1–A3) that our results are qualitatively similar if we use other measures of pollution, such as SO_2 and API, which were available earlier.

Concerns about air pollution in China and elsewhere have been shown to motivate changes in behavior. Several studies have demonstrated that people engage in short-run avoidance behaviors such as staying indoors (Neidell, 2009) or purchasing particulate-filtering facemasks (Zhang and Mu, 2018) in a highly polluted day. Recent research has also shown that pollution concentrations can motivate medium-run investments such as home air purifiers (Ito and Zhang, 2020). Our paper sheds light on the importance of migration, a long-run investment, in response to air pollution.

⁶ See http://www.gov.cn/zwgg/2005-12/13/content_125680.htm (in Chinese).




Fig. 2. Time Trend of PM_{2.5}, Thermal Inversions, and GDP in China (1980–2015). Notes: This figure depicts the national average of PM_{2.5} and thermal inversions in Panel A and GDP and thermal inversions in Panel B in each year during 1980–2015.

3. Empirical strategy

The goal of our empirical estimation is to capture the causal effect of 5-year pollution changes on migration, which poses two important challenges. First, air pollution and economic activity are highly correlated. Thus, it is likely that those cities with high economic activity that attract immigrants by offering highly paid jobs are also those experiencing high levels of air pollution. In fact, as we discuss in the results section, the correlation between air pollution and immigration is positive even when controlling for county fixed effects. Second, overcoming the first challenge requires finding a random source of variation for air pollution. However, most reliably exogenous determinants of air pollution in the literature provide short-run variation in air pollution (over the course of days, weeks, or months); and we observe changes in migration over five-year periods. Most sources of long-run variation in air pollution, such as changes in local policy or economic fluctuations in neighboring regions, are likely to have independent effects on migration as they may shift labor market conditions. The combination of these two issues poses an important challenge for identification.

3.1. Econometric model and identification

Our approach is to use medium-run random variation in air pollution

stemming from five-year fluctuations in the strength of thermal inversions in a given county. Thermal or temperature inversions are a common meteorological phenomenon that leads to higher concentrations of pollutants near the ground. The mechanism through which this occurs is the following: under a stable (normal) temperature gradient, temperature decreases as altitude increases. Since air moves from hot to cool regions, air pollutants can circulate vertically decreasing air pollution concentrations near the ground. A thermal inversion, however, occurs when temperature increases with height (an “inverted” temperature gradient). This temperature profile traps pollutants near the ground as it prevents vertical circulation. Thermal inversions that trap pollutants may emerge from different meteorological sources: earth’s thermal infrared radiation during the night (radiation inversions), air descending under a surface high-pressure system (large scale subsidence inversion), cool ocean breezes moving inland (marine inversion), and air flowing down a mountain slope (small scale subsidence inversion) (Jacobson 2002).

The idea to use thermal inversion as an instrumental variable for air pollution was first proposed by Arceo-Gómez et al. (2016), to estimate the effect of air pollution on infant mortality in Mexico City. This identification strategy has been subsequently used to explore the short-run effects of air pollution on children’s health in Sweden (Jans et al., 2018), on manufacturing labor productivity (Fu et al., 2021), and on mental health in China (Chen et al., 2018). This is the first study that uses thermal inversions to produce variation in air pollution over five-year periods.

We estimate the following 2SLS model

$$M_{ct} = \beta_0 + \beta_1 P_{ct} + f(W_{ct}) + \gamma_c + \sigma_t + \varepsilon_{ct} \tag{1}$$

$$P_{ct} = \alpha_0 + \alpha_1 TI_{ct} + f(W_{ct}) + \gamma_c + \sigma_t + \mu_{ct} \tag{2}$$

where M_{ct} denotes two measures of migration in county c and period t : the net-outmigration ratio, which is the fraction of people leaving a county minus new arrivals and deaths, and destination-based immigration ratio, which is the fraction of people entering a county but with their *hukou* in the origin.⁷ We define each period as a five-year interval. Thus, we have three periods in our study: 1996–2000, 2001–2005, and 2006–2010.

P_{ct} , measures the 5-year average concentration of PM_{2.5}, and we treat it as endogenous.⁸ Equation (2) shows the first stage of our empirical strategy. We instrument air pollution with the average strength of thermal inversions over each five-year period, TI_{ct} , conditional on flexible functions of weather variables (W_{ct}), county fixed effects (γ_c), and period fixed effects σ_t . Thermal inversion strength is defined using above-ground temperature minus ground temperature. A positive difference indicates the existence of a thermal inversion and the magnitude measures the inversion strength. A negative difference indicates the non-existence of a thermal inversion. We keep the positive difference and truncate the negative difference to zero within each 6-h period. The strength measures of individual inversions are then averaged from 6-h to five-year period. In Section 3.3, we provide a detailed description of the source of information for thermal inversions as well as migration and pollution measures.

As argued above, thermal inversions generate county-level variation in air pollution concentrations that is independent of structural sources of air pollution, including economic development. To illustrate the lack of correlation between thermal inversions and country-wide changes in air pollution, Panel A in Fig. 2 plots the county-average strength of thermal inversions in Celsius degrees (°C) in China over the period

⁷ Online Appendix III formally discusses how our first outcome variable, net-outmigration, relates to other measures of migration, such as county-to-county flows.

⁸ Tables A1–A3 in the Online Appendix I show the results with other measures of air pollution such as SO₂ and API.




Fig. 3. Pollution Changes in China (1996–2010). Notes: This figure depicts the changes in PM_{2.5} concentrations between the period 1996–2000 and 2005–2010 for each county in China.

1980–2015 along with PM_{2.5} levels. In contrast to air pollutants, there is no steep change in thermal inversion strength. This is especially important after 2001, when the steep increase in PM_{2.5} was tied to rapid economic growth, as shown in Panel B of Fig. 2. Although there appears to be no discernible trend in thermal inversions over time, we include period fixed effects, σ_t , in our main specification to be overly cautious.^{9,10}

There are a couple of additional considerations about thermal inversions that are relevant for identification. First, although there is no plausible direct mechanism through which temperature above ground level could affect human health or human behavior, thermal inversions often coincide with weather patterns on ground level.¹¹ Weather, in turn, may have direct impacts on our outcome of interest (Feng et al., 2010). Therefore, we control for very flexible functions of weather at the ground level, including 1 °C daily temperature bins, and second-degree polynomials in precipitation, sunshine duration, relative humidity, and wind speed.¹² Our identification strategy relies on the variation in the five-year average strength of thermal inversions net of weather variation at ground level.

Second, there are some regions that are more prone to thermal inversions than others, which causes permanent differences in air pollution concentrations across regions. Permanent differences in pollution may induce differences in environmental regulation stringency, industrial composition, and self-selection that could impact migration rates. For example, heavily-regulated areas could result in more layoffs and job displacement. Constraining the thermal inversion variation we use

⁹ Our results are robust to the exclusion of period fixed effects and to the inclusion of region-specific period fixed effects.

¹⁰ In addition, Figure A1 in Online Appendix I shows that 15-year-long changes in thermal inversion frequency or in thermal inversion changes are uncorrelated to 15-year-long changes in pollution and Figure A2 shows that pollution in China follows a very distinct pattern over the days of the week, whereas the thermal inversion frequency and strength looks identical over the days of the week. This test of exogeneity in thermal inversions is borrowed from Jans et al. (2018).

¹¹ See Figure A3 and its notes in Online Appendix I.

¹² We also explore the sensitivity of our results to variations in the functional forms of weather variables such as region-specific temperature effects.

to its deviations from county averages allows us to identify behavioral responses to pollution.¹³

When aggregating thermal inversion strength into five-year averages, we lose a substantial amount of temporal variation within counties. However, we show in Section 4 that there is still enough temporal variation to identify a strong first stage. This feature of thermal inversions is similar to other weather phenomena. For example, Burke and Emerick (2016) exploit time variation in 10-year temperature averages to study adaptation to climate change in agriculture.

Finally, we discuss two spatial considerations when using thermal inversions as a county-level instrument: spillovers and spatial correlation. First, a “treated county” (a county that experiences an abnormally strong spell of thermal inversions in a five-year period) could have a spillover effects over neighboring counties if the bulk of the migration in response to the pollution shock in question goes to a small number of nearby counties. If this were the case, our estimates of the response to an independent pollution shock would be biased, as some of the neighboring counties would in fact have some form of treatment. Second, assuming that the thermal inversion shocks are independent across space might be problematic as neighboring counties might share geographies and weather realizations that could make them similarly susceptible to a thermal inversion shock at a given time. To address these concerns, we estimate a spatial lag model, where we explicitly account for shocks to nearby counties in the estimation and explore several standard error structures that can account for spatial correlation in Section 5.

3.2. Thermal inversion-driven air pollution and migration decisions

Although we show that thermal inversions are capable of producing air pollution fluctuations that can last as long as 5 years, an important feature of the variation in air pollution stemming from thermal inversions is that it eventually reverts to the mean. Thus, not all rational models of avoidance behavior would predict a migration response to this specific source of air pollution variation. For example, consider a model

¹³ Figure A4 in Online Appendix I contrasts the cross-sectional variation with the within-county variation we are using in a map.

that assumes that (a) flow (per-period) utility is a function of contemporaneous pollution exposure (e.g. the health effects of air pollution are temporary); (b) that individuals have perfect information on the source of observed air pollution variation (i.e., they know whether any deviation from the local mean is temporary or long-lasting); and (c) that individuals know the dose-response function that relates air pollution to health, even for air pollution levels they have not yet experienced. In this model, individuals would move whenever the net present discounted value (NPDV) from moving to a new location exceeds the NPDV from staying in the current location. Temporary random shocks to pollution at the current location would not tip the balance in favor of moving away. However, models that allow for at least one of the following alternative assumptions would predict migration in response to the 5-year long fluctuations in air pollution generated by thermal inversions: (1) disutility from cumulative exposure, (2) imperfect information on the sources of stochastic changes in air pollution, and (3) imperfect information on the health effects of air pollution. We briefly explain these three alternative models below.

Model 1 offers an alternative to assumption (a). In this model, an individual who believes that her current health or her family's health depends on cumulative pollution exposure could be motivated to move earlier than she would have otherwise by a spike in pollution caused by thermal inversions. Similarly, a family who had planned to move from a polluted location could decide to stay longer after experiencing a five-year period with abnormally low pollution.¹⁴

Model 2 relaxes assumption (b) by assuming that individuals cannot decompose changes in air pollution into permanent and transitory, which is a reasonable assumption in the absence of knowledge from an atmospheric and pollution dispersion model. In the absence of source of pollution information, individuals could use past observations to forecast future pollution using Bayesian updating. However, it is important to keep in mind that pollution is a dynamic process over our period of study. In other words, in addition to the transitory changes (e.g. stemming from thermal inversions), the mean and trend of air pollution are also changing over time due to changes in permanent sources of air pollution (see Fig. 2). A Bayesian forecasting process (e.g. Harrison and Stevens 1976) shows that a rogue observation, such as one coming from a year with abnormally high thermal inversions, can result in very different expectations for the air pollution trajectory depending on the agent's beliefs about the source of variance of the underlying process. In a context where recent past deviations from the historical mean have signaled a shift in pollution trajectory, a new rogue observation stemming from a transitory change could be easily mistaken for a change in the mean or drift of the dynamic process.^{15,16}

Finally, Model 3 relaxes assumption (c) above. In this model, inversion-induced changes in air pollution could affect migration behavior if they inform individuals about the health effects of high pollution exposure. In a context where air pollution is increasing almost everywhere in China, a thermal inversion could provide a "window into the future" of pollution-related health damages that can motivate individuals to move to less polluted regions. Importantly, this model would only predict migration responses to "positive" shocks to air pollution in the presence of an upward sloping trend in air pollution (like the one observed during this period).¹⁷

¹⁴ We thank Solomon Hsiang for highlighting the response of individuals at the margin.

¹⁵ See Fig. 1 in Harrison and Stevens (1976).

¹⁶ Yet another reason for why individuals may move in response to a 5-year period of abnormally high/low air pollution is projection bias, which has been well documented in the literature (see Conlin et al., 2007; Simonsohn 2009; Busse et al., 2015; Chang et al., 2018). This non-rational explanation for the behavior we observe would be applicable even in the context of a static random process for air pollution.

¹⁷ We thank Catherine Wolfram for highlighting this potential mechanism.

A positive response of migration to thermal-inversion induced pollution shocks would suggest that one of the three models above is at play, but without information on expectations and beliefs, it is difficult to test across them. Nevertheless, in Section 7 we discuss suggestive evidence that rejects Model 3 by comparing responses to "positive" and "negative" pollution shocks generated by thermal inversions.

The source of variation in air pollution that we use is also relevant for the interpretation of the magnitude of the results. Note that the thermal-inversion-related air pollution shocks that each county experiences are independent across counties. In fact, we test for independence of these shocks using a spatial lag model (Section 5). Thus, the effect we find can be interpreted as the migration response to a pollution shock in one county, everything else equal (i.e., keeping pollution constant everywhere else). In reality, the bulk of pollution changes were in the form of long-run permanent trends that were highly correlated across counties. But, our estimates are only applicable to the variation in pollution that is uncorrelated across counties. We apply our estimates to the relevant variation in air pollution in Section 7 and we discuss the magnitude of the implied effects.

3.3. Data sources and summary statistics

3.3.1. Migration

As discussed in Section 2, there are two types of migrants in China: those who migrate to a new county but do not possess the local household registration (floating migrants), and those who migrate and possess the local household registration (registered migrants).

We use population and death counts from the population census in China to calculate two measures of migration: net outmigration flows of all types of migration and immigration flows of floating migrants. For our study, we use 1% and 20% individual-level data randomly drawn from the 2000 and 2005 censuses respectively, and county-aggregated data in 1995 and 2010 from National Bureau of Statistics (NBS) of China.¹⁸

The first migration measure, the net-outmigration ratio, is the percent of population leaving the county net of new arrivals and deaths. Since the population herein is based on individual's physical presence in that county, the net-outmigration ratio essentially measures the migration of both floating and registered migrants. We use the residual approach to calculate net-outmigration (see Passel et al., 2004; Feng et al., 2010; Feng et al., 2015). Specifically, we calculate:

$$NetOutmig[15,60]_{c,t} = \frac{Pop[15,60]_{c,t} - Pop[20,65]_{c,t+5} - D[\widehat{15,60}]}{Pop[15,60]_{c,t}} \times 100\% \quad (3)$$

where $NetOutmig_{c,t}$ is the net-outmigration ratio for those aged [15, 60] during the five-year interval starting from year t in county c ; $Pop[15,60]_{c,t}$ indicates the total population aged [15, 60] in county c at the beginning of the five-year interval that started in year t , while $Pop[20,65]_{c,t+5}$ denotes the population of the same cohort five years later, and $D[\widehat{15,60}]$ represents an approximate measure of deaths for the same population during the five-year interval. Below we explain the data constraints on deaths and our approach to ensure that these constraints do not affect our results.

Because NBS only surveys deaths during the survey year, we are not able to obtain the death counts in the whole five-year period. Thus, we instead subtract deaths in the survey year times five. This approximation creates measurement error and will bias our estimates upwards if pollution is positively correlated with deaths. To evaluate the potential bias, we estimate the effect of air pollution on deaths for different age

¹⁸ To the best of our knowledge, no individual-level census data in 1995 and 2010 are publicly available.

groups using the years for which deaths data are available (2000, 2005, and 2010). Specifically, we estimate model in Equations (1) and (2) with deaths in each of these years as the dependent variable and pollution in the past 1–5 years as explanatory variables. Results of these specifications are shown in Tables A4. We find that air pollution exposure within the last four and five years has a positive and significant effect on current year deaths of total population (all ages), population under 15 years of age, and population above 60 years of age. However, we find a small and non-statistically significant response for those between 15 and 60 years of age. These findings across age groups are consistent with prior literature on the effects of air pollution by age (Chen et al., 2013; Deryugina et al., 2019) and suggest that the bias caused by the measurement error in our net-outmigration measures should be minimal and statistically undetectable.

Descriptive statistics are shown in Table 1. The mean five-year death rate in our period is 1.28 per thousand. On average the net-outmigration ratio is negative (−9.17 per 100 inhabitants). Because this average is unweighted by population, the sign likely means that less populated counties, which are also more numerous, are predominantly experiencing net inflows. The standard deviation of the net-outmigration is 16.15, and shows there is substantial heterogeneity across counties. This is also clear from Panel A in Figure A5, which depicts the histogram of the net-outmigration. Although average net changes in population are modest, the top five percent of counties experience increases in population of 40 percent or more due to migration. Two advantages of this measure are that it captures both floating and registered migrants, and is not subject to misreporting *hukou*.

Our second measurement on migration captures destination-based immigrants whose *hukou* are in their origin, or floating immigration. This measure is calculated from individual-level census in 2000 and 2005, and county-level aggregated census in 2010.¹⁹ Importantly, this measure of migration is estimated directly from individual responses to a migration question, rather than from population counts at different points in time. Specifically, individuals who say their *hukou* does not correspond to their present location are subsequently asked when did they move to their present location. From previous work on Chinese migration (Ebenstein and Zhao, 2015) and from our calculations, we know that about 70 percent of migrants constitute floating migrants.²⁰ Since the majority of migrants do not transfer their *hukou*, our destination-based immigration captures the bulk of the response to air pollution. Table 1 shows an increasing trend in destination-based floating immigration during the period of our study, with an average of 6 floating immigrants per 100 people in our whole period. Studying this measure of migration has multiple purposes: it allows us to check for the pull effect of air quality, i.e., whether individuals pay attention to recent pollution levels at their destination. Second, it relies on individual's direct answer to a migration question rather than on an accounting exercise using population counts. Third, it helps us understand whether migration flows in response to air pollution are driven by registered migrants or are also driven by floating migrants.

3.3.2. Air pollution

The data on air pollution are derived from remote sensing AOD retrievals, as these data are available for the full period of our study. AOD essentially measures the amount of sunshine duration that are absorbed, reflected, and scattered by the particulates suspended in the air, and can be used to estimate particular matter concentrations. The AOD-based pollution data closely match the ground-based monitoring station measures (Gupta et al., 2006; Kumar et al., 2011).

¹⁹ Details are available in Online Appendix II. Note that we cannot calculate origin-based outmigrants because aggregated data in 2010 only report the destination-based immigrants.

²⁰ Our calculation comes from the 2000 Census, which has information on both floating and registered migration.

We obtain the AOD data from the product M2TMNXAER version 5.12.4 from the Modern-Era Retrospective analysis for Research and Applications version 2 (MERRA-2) released by the National Aeronautics and Space Administration (NASA) of the U.S.²¹ The data are reported at each $0.5^\circ \times 0.625^\circ$ (around 50 km \times 60 km) latitude by longitude grid in each month since 1980. The concentration of PM_{2.5} is calculated following Buchard et al. (2016).²² The monthly pollution data are then aggregated from grid to county. We then average to annual level across all months and further average to each five-year period for each county.

We compare our AOD-based data with ground-based data during the period 2013–2015, when CNEMC and US Embassy started to report hourly concentration specific air pollutants and manipulation is not a major concern.²³ We find no statistical difference between them conditional on county fixed effects. The details are discussed in Online Appendix II. The average concentration of PM_{2.5} during 1996–2010 is 53.08 $\mu\text{g}/\text{m}^3$, which is five times larger than the WHO's standard.

3.3.3. Thermal inversions and weather

The data on thermal inversions are also from the MERRA-2. In particular, we utilize the product M2I6NPANA version 5.12.4, which divides the earth by $0.5^\circ \times 0.625^\circ$ (around 50 km \times 60 km) grid, and records the 6-h air temperature at 42 layers, ranging from 110 m to 36,000 m.²⁴ We aggregate all data from grid to county. Within each 6-h period, we calculate the temperature difference between the second layer (320 m) and the first layer (110 m). If the difference is positive, there exists a thermal inversion and the difference measures the inversion strength. If the difference is negative, we code it as zero. We then average the inversion strength across all 6-h lapses within each five-year period. The average strength during our study period is 0.22 °C. During 1996–2010, average thermal inversion strength appears to be increasing at a very slow pace.

The weather data are obtained from the China Meteorological Data Sharing Service System (CMDSSS), which records daily minimum, maximum, and average temperatures, precipitation, sunshine duration, relative humidity, and wind speed for 820 weather stations in China.²⁵

4. Results

Table 2 presents the first-stage estimates of the effect of thermal inversions on PM_{2.5} concentrations (Equation (2) in Section 3). Column (1) shows the results without population weights, while column (2) uses population aged 15 to 60 in 1995 to weight the regression. Table 2 also reports the Kleibergen-Paap (KP) *F*-statistics, and all of them are well above Stock and Yogo's 10% maximal bias threshold of 16.38. All regressions control for county and period fixed effects as well as weather controls, so it is most helpful to think about residualized changes in pollution in response to residualized changes in thermal inversions when interpreting the magnitude of these coefficients. One can multiply the point estimates by 0.004 (0.22/53.08) in order to obtain an elasticity. The point estimates in column (1) suggest that a 1 percent increase in residualized average thermal inversion strength leads to a 0.3

²¹ The data can be downloaded at https://disc.gsfc.nasa.gov/datasets/M2TMNXAER_5.12.4/summary. Although satellite-based AOD measures are only available from 2000 on, the MERRA-2 product that we use for this paper also incorporates AOD measures from Advanced Very High Resolution Radiometer (AVHRR) as inputs, which are available for the previous decade.

²² MERRA-2 also reports SO₂ and we can compute API from its components.

²³ For real-time air pollution and the geographic locations of the eight monitoring stations, please see <http://www.cnemc.cn/from/CNEMC> and <http://www.stateair.net/web/historical/1/1.html> from US Embassy.

²⁴ The data can be downloaded at https://disc.gsfc.nasa.gov/datasets/M2I6NPANA_V5.12.4/summary.

²⁵ CMDSSS was developed and is currently managed by the Climatic Data Center, National Meteorological Information Center, China Meteorological Administration. See <http://data.cma.cn/> for details.

Table 1
Summary statistics.

Variable	Unit	1996–2010		1996–2000		2001–2005		2006–2010	
		Mean	SD	Mean	SD	Mean	SD	Mean	SD
Migration									
Net-outmigration ratio (death adjustment)	%	−9.17	16.15	−6.61	9.89	−5.82	16.39	−15.08	19.07
Immigration ratio	%	6.01	10.73	3.41	6.16	4.20	7.49	10.40	14.90
Immigration ratio by origin									
Within county	%	4.45	4.53	2.50	2.23	3.02	3.48	7.84	5.20
Across county within province	%	3.66	6.72	2.37	4.10	2.36	4.16	6.25	9.56
Across county outside province	%	2.58	5.93	1.72	4.06	1.88	4.71	4.15	7.94
Death rates	‰	1.28	0.54	1.31	0.39	1.35	0.61	1.19	0.57
Air pollution									
PM _{2.5}	µg/m ³	53.08	27.93	42.68	19.85	50.89	24.53	65.67	32.76
Thermal inversion									
Strength	°C	0.22	0.19	0.21	0.19	0.22	0.20	0.23	0.20
Number of inversions	Days	107.65	59.00	107.30	56.87	107.03	59.63	108.62	60.47

Notes: The unit of observation is county-period (five years). Number of observations is 7911. Net outmigration ratio is defined as the percent of population aged 15 to 60 leaving the county net of new arrivals and approximate deaths. Immigration ratio is defined as the percent of population aged 15 to 60 entering the county with their *hukou* in the origin. Death rates are for population aged 15 to 60. Pollution data are reported at monthly level, and then are averaged to each year and further to each period. Thermal inversion strength is calculated using the temperature difference in altitudes of 110 and 330 m within each 6-h period, and then is averaged for each period. Positive difference indicates an existence of a thermal inversion with magnitude representing the strength, while negative difference indicates a non-existence of a thermal inversion and is truncated to zero. Number of inversions is calculated using annual days with thermal inversions, and then averaged to the five-year period.

Table 2
The effect of inversions on PM_{2.5} (first stage).

	(1)	(2)
Thermal inversions	78.6938*** (5.4654)	82.0176*** (5.1621)
KP <i>F</i> -statistics	865.2	984.9
Observations	7911	7911
Weighting	No	Yes

Notes: The dependent variable is PM_{2.5}. Regression models are estimated using Equation (2) and include county FE and period FE. Weather controls include temperature bins within 1 °C, second-order polynomial in precipitation, humidity, wind speed, and sunshine durations. Regression models are weighted using population aged 15 to 60 in 1995 in column (2). Standard errors are listed in parentheses and clustered at county level. **p* < 0.10, ***p* < 0.05, ****p* < 0.01.

percent increase in PM_{2.5}.

To put these estimates in perspective, average PM_{2.5} pollution in China has increased by about 7.6 µg/m³ every five years from 1995 to 2010, with the bulk of this change happening between 2000 and 2010. A within-county change in (residualized) 5-year average thermal inversion strength from the 25th (5th) percentile to the 75th (95th) percentile would have produced an increase in average air pollution over a five-year period of 1.23 (3.38) µg/m³ of PM_{2.5}; that is 16 (44) percent of the observed average change in air pollution occurring over the same period. This shows that, while thermal inversions are obviously not the predominant driver of air pollution changes in this period, their effects are not negligible compared to overall changes.

Table 3 reports the estimated effects of air pollution on migration flows. Panel A shows estimates of the effect of air pollution on net-outmigration adjusting for deaths.²⁶ Column (1) reports the fixed effects (FE) estimates, while columns (2)–(4) report the IV estimates. Columns (3) and (4) include population weights. Column (4) includes period-by-province FE while the rest include only period FE.

The FE estimates suggest a positive and significant correlation between air pollution and net-outmigration after controlling for weather variables as well as county and period fixed effects. Note that pollution is endogenous in this specification and may be correlated with other

²⁶ Results look similar without the death adjustment, which is consistent with the lack of significant effects of air pollution on deaths (Section 3). Results without dead adjustment are shown in Table A6.

Table 3
The effect of PM_{2.5} on migration.

	FE	IV		
	(1)	(2)	(3)	(4)
Panel A: Net-Outmigration ratio				
PM _{2.5}	0.2967*** (0.0589)	0.5944*** (0.1657)	0.5314*** (0.1761)	0.9348** (0.4119)
Panel B: Immigration Ratio				
PM _{2.5}	0.1258*** (0.0299)	−0.2516*** (0.0679)	−0.3246*** (0.0820)	−0.4275** (0.1675)
Observations	7911	7911	7911	7911
KP <i>F</i> -statistics	–	205.9	250.7	120.4
Period FE	Yes	Yes	Yes	No
Per. by Province FE	No	No	No	Yes
Weighting	Yes	No	Yes	Yes

Notes: The dependent variables are net outmigration ratio in Panel A and destination-based immigration ratio in Panel B. Net-outmigration ratio is defined as the percent of population aged 15 to 60 leaving the county net of new arrivals and approximate deaths. The destination-based immigration ratio is defined as the percent of population aged 15 to 60 entering the county with their *hukou* in the origin. Column (1) presents fixed effects estimates, and columns (2)–(4) present IV estimates in which we instrument PM_{2.5} using thermal inversions strength. Weather controls include temperature bins within 1 °C, second-order polynomial in precipitation, humidity, wind speed, and sunshine durations. All regressions include county FE. Regression models are weighted using population aged 15 to 60 in 1995 in columns (1), (3), (4), and (6). Standard errors are listed in parentheses and clustered at county level. **p* < 0.10, ***p* < 0.05, ****p* < 0.01.

determinants of migration that vary over time within counties—such as wage, GDP, job opportunities, and infrastructure—which would result in omitted-variable bias. As many of these omitted factors are likely to attract migrants, the bias is likely negative. The FE estimates may also be biased downwards due to reverse causality, as positive net-outmigration flows may bring down pollution.

Consistent with the expected bias just discussed, the IV estimates of the effect of air pollution on net-outmigration are larger in magnitude. Our preferred specification is in column (3), which weights the regressions using population in 1995. Weighting is important as it results in a more efficient variance estimator and in point estimates that reflect the migration flows faced by a representative individual. The point estimate for the weighted IV effect is 0.53, nearly twice the size of the FE effect. Column (2) of Table 3 shows the same specification without

population weighting. Results vary little with respect to column (3), suggesting heterogeneity by population density is not very important. Finally, column (4) shows the IV results when we include period-by-province FE. This specification limits the time variation in pollution to changes net of province-level changes. We find that the results are 75% larger than when we control for period FE. If costs of migration are smaller within the same province, it is reasonable to see a stronger migration response to pollution changes relative to province-level changes.²⁷

We now turn to our results on floating immigration presented in Panel B of Table 3. There are two important differences in the interpretation of these results with respect to Panel A. First, destination-based immigration corresponds to floating immigration only. This means that registered immigration, which is costlier (Kinnan et al., 2018), is excluded from this measure (see Section 3). Second, the fact that we are now measuring only the inflow as opposed to the net-outflow has a couple of important implications: A) if individuals value air quality, then we expect the opposite sign to the effect of pollution in Table 3, and B) finding a causal response of immigration flows to destination air pollution requires individuals to be aware of pollution changes in the place where they are planning to move to, as opposed to pollution changes in the county where they live. As we show next, we find results that are consistent with people moving to counties where pollution has improved.

Panel B has a similar structure to Panel A. The FE estimates in column (1) suggest a significantly positive relationship between air pollution and immigration, the opposite sign to what one would expect from the causal relationship. It appears that in the case of immigration, the bias stemming from confounding factors is large enough to flip the sign of the expected causal relationship. When we instrument air pollution using strength of thermal inversions, we find significantly negative effects of air pollution on immigration, which is consistent with individuals valuing clean air. Our preferred estimates in column (3) imply that a 10 percent reduction in PM_{2.5} (5.31 µg/m³) brings in 1.7 people per 100 inhabitants. The smaller magnitude of the effects compared to Panel A is expected as net-outmigration captures the effect of air pollution on both inflows and outflows, while immigration only captures the effect on inflows. In addition, registered migration (which is not captured by this measure) is less than one third of overall migration (see Section 3.3.1). Column (2) shows the results without population weights. In the case of destination-based immigration, the immigration effects faced by the average person in China seem to be slightly larger than those faced by the average county. The results in column (4), which further restrict pollution and thermal inversion variation over time, show a similar pattern to net-outmigration flows: the inflow of immigrants is larger when a county presents an air pollution reduction that is out-of-the-norm with respect to the same province, as opposed to the country as a whole.

Taken together, the results of Table 3 imply that individuals respond both to changes in air pollution at their home county as well as changes in air pollution in destination counties. In addition, responses appear to be stronger when the change in air pollution is out-of-the-norm with respect to counties in the same province. Section 7 interprets the magnitudes of the migration response effects we find in the context of air pollution variation over time and observed migration rates. It also discusses additional results that help us understand the mechanisms driving this causal relationship.

5. Robustness

Here we discuss the results of several robustness checks. Tables A7 in the Online Appendix I shows several of these additional results and

²⁷ The fact that migration responses within a province are stronger than across provinces is confirmed by our immigration results by origin shown in Table 5.

compares them to our baseline. We first explore the robustness of our results to alternative forms of clustering (column (2)). Clustering errors at the prefecture level (which has 10–20 counties) instead of the county level results in standard errors that are about twice as large, but the effects are still statistically significant at the 5% level. The KP *F*-statistic is also still above the Stock-Yogo critical value for 10% relative bias. Our results are also robust to replacing the baseline weights (1995 population) with average population during the 1996–2010 period (column (3)).

We also show that our results are robust to using different layers of temperature to calculate thermal inversions and to using number of days with thermal inversion as our instrumental variable instead of inversion strength (columns (4) and (5)). Our results are robust to these two alternative definitions of inversions.

We also test the robustness of our results to region specific weather functions. As Figure A3 in Online Appendix I demonstrates, the relationship between weather and inversions is different for different regions. If there was a spurious correlation pattern of temperature and migration that coincided with this heterogeneity, temperature could still bias our results. When we include interactions between all of our standard weather controls and six region dummies, our results are very similar to our baseline results (column (6)).

Finally, Tables A8–A10 show the results of a spatial lag model that addresses several potential issues with our estimates stemming from the spatial proximity of some of these counties. First, if the bulk of migration in response to an air pollution shock goes to (or comes from) a small set of counties in close proximity to the shock-receiving county, our identification strategy would violate the stable unit treatment value assumption (SUTVA). If on the other hand, the migration response is dispersed over numerous counties, then SUTVA would not be violated. Controlling for a spatial lag of migration is a useful to test this assumption. Because migration shocks to nearby counties are also endogenous, we generate an instrument for the spatial lag of migration using the analogous spatial lag of thermal inversions. The results of this model are presented in Tables A8, where the spatial lag is defined as the inverse distance weighted average of migration flows to neighboring counties. Columns (1)–(4) show the results of this model with net outmigration as the dependent variable and columns (5)–(8) show the results for immigration of floating migrants.²⁸ The effect of local pollution on migration in all models remains of roughly the same magnitude as in our baseline and has similar levels of significance. In addition, the effect of the spatial lag of migration flows is insignificant in all models. Together, these results suggest that thermal inversion shocks are well spread out across counties as opposed to concentrated in a few neighboring counties.

A second issue generated by spatial proximity of the counties is that thermal inversions could be spatially correlated. Although clustering at the prefecture level (like we do in Table A7) could help approximate the correct error structure in this case, a spatial lag model offers an alternative way of accounting for this spatial correlation. Tables A9 and A10 estimate the first stage and reduced forms of our IV model in Table A8 using the correct variance structure. In all cases the main results do not differ from our baseline, suggesting that localized spillovers and spatial correlation in the error term are not affecting our main estimates.

²⁸ Standard errors are clustered at either the county level (columns (1), (2), (5), and (6)) or at the prefecture level (columns (3), (4), (7), and (8)). We do not use a spatially decaying correlation structure for the variance as the code that incorporates both IV and this variance structure is currently unavailable. However, a variance structure with a spatially decaying correlation is studied in Tables A9 and A10, which correspond to the first stage and the reduced form of the IV model in Table A8.

6. Heterogeneity by demographic groups and origin

Different individuals may have different tradeoffs between perceived harm from air pollution and economic opportunities. For example, we expect that highly educated individuals will be better informed about potential harm from air pollution exposure and will also have lower costs of migration, as registered migration is within reach for this demographic group. Heterogeneity in the response could also stem from differences in vulnerability to air pollution: children and elderly face higher health impacts from poor air quality.

In this section we exploit information on demographic characteristics to explore whether our main result masks any heterogeneity that is consistent with these relative tradeoffs. Table 4 shows results by gender, education and age (the categories for which net-outmigration flows could be constructed from the census). As usual, the units of the coefficients are in net-outmigrants per 100 people associated with one $\mu\text{g}/\text{m}^3$ of $\text{PM}_{2.5}$. We focus on three important observations from these results. The first observation is that education significantly increases the migration response to air pollution. Having a college degree makes it nearly twice as likely to migrate in response to air pollution compared to the average person (0.93 vs. 0.53). And although the response to air pollution is significant across all education levels for female migrants, male migrants with primary education or less show no statistically significant response. The education gradient of our response is consistent with either the perceived benefits of air pollution rising with education, the cost of migration falling with education, or both. Importantly, these findings suggest that air pollution may have important effects on the composition of the labor force, causing a “brain-drain effect” (Fischer, 2003). In addition, our findings support other recent literature that finds air pollution changes the socioeconomic composition of neighborhoods (Hanlon, 2020; Heblich et al., 2021).

The second observation is that male and female responses are not consistent across age groups. The starkest difference is between male and female responses between the ages of 30 and 45 years of age. For this age group, we find the smallest (and non-significant) response from men (0.21) and the highest response from women (1.12). This pattern could emerge from a lower labor force participation of women combined with the ability of families to live in separate counties (Chang et al., 2011). Since young children are particularly vulnerable to air pollution, splitting across counties would allow families to maximize both health benefits and economic gains, conditional on low female labor participation.

The third and last observation is that the age gradient has the opposite sign to what we would expect based on vulnerability to air pollution for men: the youngest working-age individuals are more likely to migrate compared to the oldest (0.79 vs. 0.36). This suggests that the cost of migration, which is likely smaller for those who are just entering the labor force, plays an important role in migration decisions.

Our second measure of migration flows, floating immigration, allows us to further explore the influence of migration costs in responses to air pollution. The data on destination-based immigration classifies the origins of immigrants by whether they come from the same province or from other provinces. Table 5 reports our results by these two categories. We find significantly negative effects of pollution on destination-based immigration for both movements across counties within a province and movements across provinces (cols. (2) and (3)). However, migration within a province appears to be twice as large as migration across provinces.²⁹ This is also consistent with our observation on the gender imbalance, which suggests that heads of household may be staying

²⁹ Floating immigration within county is also available. Since our air pollution level is at the county level, we would not expect an effect on within county migration, and we indeed find no empirical response. This result serves as a robustness test on the validity of the exclusion restriction of our instrument. Results are not reported, but available upon request.

behind. If this were the case, remaining within the same province might be less costly for the family.

The fact that internal migration in response to air pollution tends to happen disproportionately within a province also informs the interpretation of the magnitude of our main point estimates. These results suggest that individuals pay more attention to deviations from province-level changes in air pollution compared to deviations from country-level changes. In the next Section we apply our main point estimates to changes in air pollution that deviate from province level changes in order to assess the amount of migration that can be attributed to air pollution.

7. Magnitude and interpretation

Our preferred estimate from Panel A in Table 3 indicates that a ten percent increase in $\text{PM}_{2.5}$ ($5.31 \mu\text{g}/\text{m}^3$) reduces population by 2.8 per 100 inhabitants. This is a large effect on migration when compared to the standard deviation of net-outflows: 16.15%.³⁰ Although it is tempting to multiply this point estimate by the increase in air pollution that the average county in China experienced over the 15 years of our study: $27 \mu\text{g}/\text{m}^3$ (60%), we refrain from this extrapolation exercise as our estimates are not predictive of migration in response to country-wide changes in air pollution. Because the time variation in the frequency of thermal inversions is as if random, our estimate corresponds to the response to pollution changes that are out-of-the-norm with respect to country-wide changes.

To assess the amount of migration that could be attributed to pollution shocks of a similar nature to our identifying variation, we instead apply our point estimates to the changes in air pollution that remain after removing the country-wide trend in pollution and the province-level trends in pollution. Panel A of Fig. 4 shows the distributions of predicted migration responses when using only the changes in air pollution that exclude country-wide level changes (black dashed line) and province-level changes (gray dotted line). When compared to the distribution of naïve predicted responses (solid gray line), which uses observed changes in air pollution, these two distributions have a higher mass of small (close to zero) migration responses.

Panel B of Fig. 4 plots the average prediction of the migration response to pollution for each corresponding change in air pollution between periods. The plus-sign markers show a naïve prediction obtained by multiplying our point estimate of the effect of air pollution on net-outmigration (0.53 from Table 3) times the change in air pollution concentration (the number in the horizontal axis). This prediction is naïve because the bulk of these changes in air pollution happened simultaneously across the country and therefore would not generate different changes in pollution across counties. In contrast, the triangles and circles predict average net-outmigration effects using only the left-over variation in air pollution after subtracting national and province level changes respectively. Since, the bulk of the migration response to air pollution happens within a province (Table 5), we think that using only changes in air pollution relative to other counties within the same province (circles) is a reasonable way to illustrate the magnitude of our estimates. Each point marked by a triangle and circle is constructed by multiplying residual changes in air pollution times the point estimate that controls for period by province FE (0.9348 in col. (4) of Table 3), and then averaging these county-period predictions within each bin of observed air pollution changes. To facilitate the interpretation, the graph also shows the share of county-period observations that experienced those changes (gray bars) and the observed (raw) net-outmigration rates averaged over the same bins (solid line with no markers).

First, let us interpret raw data depicted by the solid line. This line

³⁰ The average net-outflow is regularly around zero, as most migration happens within the country's boundaries.

Table 4
The effect of PM_{2.5} on net outmigration ratio: By education, gender, and age.

Panel A: By Education and Gender												
	All gender				Males				Females			
	Full sample	Primary and below	Middle school	College and above	Full sample	Primary and below	Middle school	College and above	Full sample	Primary and below	Middle school	College and above
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
PM _{2.5}	0.5314*** (0.1761)	0.4723** (0.2350)	0.5736** (0.2427)	0.9314** (0.4433)	0.3600* (0.1935)	-0.2229 (0.3219)	0.6480*** (0.2429)	0.9192** (0.4315)	0.6977*** (0.1846)	0.6975*** (0.2169)	0.7910*** (0.2742)	1.2369** (0.5347)
KP F-stat.	250.7	300.9	224.2	169.7	250.7	305.3	226.6	171.6	250.7	304.1	224.4	172.1
Obs.	7911	7662	7673	7617	7911	7699	7692	7621	7911	7712	7655	7595
Mean	-9.17	-1.59	-11.20	-17.03	-10.07	-7.18	-7.05	-10.01	-8.68	-0.04	-12.40	-18.97
[SD] of D.V.	[16.15]	[10.06]	[24.64]	[40.44]	[18.84]	[40.02]	[22.75]	[32.91]	[16.64]	[25.49]	[27.52]	[43.68]

Panel B: By Age and Gender												
	All gender				Males				Females			
	Age 15-60	Age 15-30	Age 30-45	Age 45-60	Age 15-60	Age 15-30	Age 30-45	Age 45-60	Age 15-60	Age 15-30	Age 30-45	Age 45-60
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
PM _{2.5}	0.5314*** (0.1761)	0.7446** (0.3173)	0.6905*** (0.2089)	0.0470 (0.2555)	0.3600* (0.1935)	0.7865** (0.3774)	0.2198 (0.2282)	0.3536 (0.2737)	0.6977*** (0.1846)	0.8583*** (0.3274)	1.1164*** (0.2295)	-0.2131 (0.3034)
KP F-stat.	250.7	269.3	255.5	281	250.7	269.3	255.5	281.1	250.7	269.2	255.5	281
Obs.	7911	7911	7911	7911	7911	7911	7911	7911	7911	7911	7911	7911
Mean	-9.17	-1.07	-7.62	-13.24	-10.07	-4.25	-8.61	-16.21	-8.68	-0.54	-9.37	-16.65 [13.02]
[SD] of D.V.	[16.15]	[31.16]	[23.15]	[13.21]	[18.84]	[40.15]	[25.31]	[14.65]	[16.64]	[30.65]	[26.98]	

Notes: The dependent variable is net outmigration ratio by each group. Net outmigration ratio is defined as the percent of population aged 15 to 60 leaving the county net of new arrivals and deaths. Regression models are estimated using Equation (1) and include county fixed effects, period fixed effects, and weather controls. Regression models are weighted using population for each group in 1995. Standard errors are listed in parentheses and clustered at county level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5
The effect of PM_{2.5} on immigration ratio: By origins.

	Total immigration	Across county within province	Across county outside province
	(1)	(2)	(3)
PM _{2.5}	-0.3246*** (0.0820)	-0.1780*** (0.0595)	-0.0851** (0.0409)
Observations	7911	7911	7911
Mean [SD] of D.V.	6.01 [10.73]	3.66 [6.72]	2.58 [5.93]

Notes: The dependent variable is destination-based immigration ratio, which is defined as the percent of population aged 15 to 60 entering the county with their hukou in the origin. Column (1) includes all migrants regardless of origins. Column (2) includes migrants whose origins and destinations are in the same province. Column (3) includes migrants whose origins are outside the province of the destination. Regression models are estimated using Equation (1) and include county fixed effects, period fixed effects, and weather controls. Regression models are weighted using population aged 15–60 in 1995. Standard errors are listed in parentheses and clustered at county level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.


illustrates the observed correlation between net-outmigration and changes in air pollution. Note that the median county experienced a net inflow of population (negative net outmigration). Also note that a positive correlation between changes in air pollution and net outmigration is only evident for very large changes in air pollution.

Second, we can compare the predictions based on deviations from the national and province-level changes in air pollution. The predicted migration responses to pollution of county-periods that had changes between 0 and 20 $\mu\text{g}/\text{m}^3$ are balanced around zero, especially when we


remove the province-level variation (circles). In contrast, the predicted net-outmigration responses in counties that experienced overall changes in air pollution above 25 $\mu\text{g}/\text{m}^3$ are positive and increasing in the magnitude of the change. This is because large (and very small or negative) changes are more likely to be out-of-the-norm with respect to other counties in China and counties within the same province. Again, the average predicted responses to pollution mirror what happened with observed net-outmigration rates for these range of changes in air pollution.

Having discussed the magnitude of migration flows implied by our estimates, we now turn to some additional results that help us understand the motivation behind the observed migration response. First, we explore the role of available air pollution information. Second, we explore whether people react symmetrically to positive and negative changes in air pollution, which sheds light on whether individuals have knowledge on the dose-response function that maps air pollution to health (see Section 3). Finally, we explore whether migration responses are driven by individually motivated behavior or by employers changing location.

To study the role of available information on air pollution, we explore whether counties that introduced air pollution monitors over time experienced sharper responses to air pollution. Note that these results are only suggestive, as the presence of monitors could also capture heterogeneity along other dimensions, such as stringency of air pollution regulation. In addition, recall that there is ample evidence that some provinces manipulated air pollution information (Ghanem and Zhang 2014). The number of cities with API data available increased



Panel A: Distribution of Predicted Effect of Air Pollution on Net-Outmigration



Panel B: Average Pollution Effects Across Different 5-Year County Level Changes in PM_{2.5}

Fig. 4. Effect Magnitude. Notes: Panel A shows the distribution of the predicted migration responses to PM_{2.5}. The dotted (dashed) line shows the distribution of predicted effects on net-outmigration to changes in PM_{2.5} net of national (province)-level changes. The gray line shows the corresponding distribution if we apply our point estimates to raw changes in PM_{2.5}. Panel B shows that predicted net-outmigration responses to air pollution are only large at the right tail of the distribution of PM_{2.5} changes. The triangle (circle) markers show averages of predicted net-outmigration responses to PM_{2.5} net of national (province)-level changes. The plus-sign markers show the corresponding averages if we use raw 5-year changes in air PM_{2.5}. The bars show the share of county-period observations in each raw 5-year PM_{2.5} change bin.

from 47 in 2000 to 86 in 2010. Consistent with information on air pollution influencing migration decisions, we find a much larger effect for cities with API data.³¹ For example, we find that the effect on net-outmigration ratio is twice as large for the 86 cities with API data in 2010 than in the remaining 250 cities. The same conclusion holds for the immigration ratio.

Even when individuals have information on the air pollution they are exposed to, they could lack information on the dose-response function that maps air pollution to health. Being exposed to air pollution shocks that are out-of-the-norm could provide a window into the health they will experience as pollution levels continue to increase. As discussed in Section 3, this is one plausible reason why transitory shocks to air pollution are observed to cause migration responses (Model 3). We can test for the presence of this information channel by comparing migration

responses to positive (informative about health response to future exposure) and negative (non-informative about health response to future exposure) shocks to air pollution generated by thermal inversions. Table A12 in Online Appendix I tests for this asymmetry using the reduced form specification. Here, we distinguish positive from negative shocks by comparing the 5-year average strength or frequency of thermal inversions with the county-specific mean for the whole 15-year period. Our results show no differential response between positive and negative thermal inversion shocks. There is also no discrete jump in migration for thermal inversion shocks above the mean with respect to shocks below the mean. Thus, by ruling out Model 3, these results suggest that individuals are either not able to distinguish between temporary and permanent changes to pollution (Model 2) or that migration decisions depend on cumulative exposure (Model 1).

Finally, we look for evidence on whether relocation decisions are driven by individual motivation or are responding to firms (employers) or government policy. We utilize the Chinese Industrial Enterprises

³¹ Results are available in Table A11 of Online Appendix I.

Database, also used in Fu et al. (2021) to study the effects of pollution on firm-level productivity. This dataset provides firm location and covers all state-owned enterprises and non-state firms with sales above CNY 5 million. Similar to our migration measure, we define five years as a period, namely, 1998–2002 as the first period, and 2003–2007 as the second period. We construct the outmigration (immigration) ratio for each county as the ratio between number of firms moving out (in) and total number of firms at the beginning of each period. We calculate net-outmigration ratio as the difference between outmigration and immigration ratio. We then estimate the effect of air pollution on three measures of firm migration using thermal inversions as the instrument and include county fixed effects, period fixed effects, and weather controls. Table A13 in Online Appendix I reports the estimates. Similar to Fu et al. (2021), we do not find any significant effects of air pollution on any measure of migration for any type of ownership. Our results suggest that the migration in response to air pollution that we are capturing is predominantly driven by decisions at the individual or household level.

8. Conclusions

Our findings suggest that pollution changes are an important determinant of internal migration in China. A county-level independent shock to air pollution of 10 percent of the average concentration will reduce the population in that county by 2.8 percent through a combination of less immigration and more outmigration. A significant share (close to half) of that response seems to be produced by reduced immigration of floating immigrants; i.e. immigrants that do not change their *hukou* or official residence when they move. This suggests that individuals keep track of air pollution levels not only in their county of origin but also in potential destination counties.

When interpreting the magnitude of our results, it is important to account for the independence of the shocks that we use to identify our effects. This is relevant because pollution changes in China in the period of our study were highly correlated across counties. Specifically, out of the average time variation in our pollution data (that is, the average variation left after subtracting the cross-sectional variation), only 26 percent is uncorrelated across counties. Therefore, extrapolating our estimates to the total changes in air pollution that the average county experienced would likely overestimate the movement in population that air pollution was responsible for in this period. However, when applying our estimates to the within county variation net of province wide period-to-period changes, we find that reductions in population of the order of 1.7–3.5 percent in response to air pollution are not rare: these correspond to the 75th and 90th percentiles of predicted changes.

The magnitude of these flows is especially important when considering their demographic composition. Our results show that responses to air pollution are predominantly driven by women in childbearing and child-rising age and that their male counterparts migrate at lower rates (and only when they are very young). This suggests that families are choosing to split between different locations in response to air pollution; a result that had not been documented in the literature. In addition, the migration response to air pollution has a steep education gradient. This has the potential to reshape the labor force composition across counties.

Data availability

The data that has been used is confidential.

Acknowledgement

All authors contributed equally and are ordered alphabetically. Shuai Chen thanks the Natural Science Foundation of Zhejiang Province (LR22G030003), the Cyrus Tang Foundation, the National Natural Science Foundation of China (72134006), the Ministry of Science and Technology of China (2020YFA0608600), the Ecological Civilization Project of Zhejiang University, the Fundamental Research Funds for the

Central Universities, and the *Efd* Initiative of the University of Gothenburg through Sida. Peng Zhang thanks the National Social Science Foundation of China (21ZDA065) and National Natural Science Foundation of China (72033005). We thank Hunt Allcott, Tamma Carleton, Olivier Deschênes, Andrew Foster, Joshua Graff Zivin, Solomon Hsiang, Peter Kuhn, Nick Kuminoff, Juan Pablo Rud, Christopher Timmins, Catherine Wolfram, and Yanos Zylberberg; seminar participants at Beijing Institute of Technology, Guangzhou University, Lingnan University, Jinan University, Peking University, The Hong Kong Baptist University, The Hong Kong Polytechnic University, University of Pittsburgh, University of Texas at Austin, University of Washington, Xi'an Jiaotong University, and Zhejiang University; and conference participants at the 2018 China Meeting of the Econometric Society, the 11th China Economics Summer Institute, the 6th World Congress of Environmental and Resource Economists, the 2017 International Symposium on Contemporary Labor Economics, and the 16th Occasional Workshop in Environmental and Resource Economics for helpful comments. Jianghao Wang provided excellent research assistance.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.jdeveco.2022.102833>.

References

- Arceo-Gómez, Eva, Hanna, Rema, Oliva, Paulina, 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.
- Bayer, Patrick, Keohane, Nathaniel, Timmins, Christopher, 2009. Migration and hedonic valuation: the case of air quality. *J. Environ. Econ. Manag.* 58 (1), 1–14.
- Borgschulte, Mark, Molitor, David, Eric, Y., Zou, 2018. Air Pollution and the Labor Market: Evidence from Wildfire Smoke (Working Paper).
- Borjas, George J., 1999. The economic analysis of immigration. *Handb. Labor Econ.* 3, 1697–1760.
- Borjas, George J., 2015. Immigration and globalization: a review essay. *J. Econ. Lit.* 53 (4), 961–974.
- Buchard, V., da Silva, A.M., Randles, C.A., Colarco, P., Ferrare, R., Hair, J., Hostetler, C., Tackett, J., Winker, D., 2016. Evaluation of the surface PM_{2.5} in version 1 of the NASA MERRA aerosol reanalysis over the United States. *Atmos. Environ.* 125, 100–111.
- Burke, Marshall, Emerick, Kyle, 2016. Adaptation to climate change: evidence from US agriculture. *Am. Econ. J. Econ. Pol.* 8 (3), 106–140.
- Busse, Meghan R., Devin, G., Pope, Jaren C. Pope, Silva-Risso, Jorge, 2015. The psychological effect of weather on car purchases. *Q. J. Econ.* 130 (1), 371–414.
- Cai, Ruohong, Feng, Shuaizhang, Oppenheimer, Michael, Pytlíkova, Mariola, 2016a. Climate variability and international migration: the importance of the agricultural linkage. *J. Environ. Econ. Manag.* 79, 135–151.
- Cai, Xiqian, Lu, Yi, Wu, Mingqin, Yu, Linhui, 2016b. Does environmental regulation drive away inbound foreign direct investment? Evidence from a quasi-natural experiment in China. *J. Dev. Econ.* 123, 73–85.
- Chang, Hongqin, Dong, Xiao-yuan, MacPhail, Fiona, 2011. Labor migration and time use patterns of the left-behind children and elderly in rural China. *World Dev.* 39 (12), 2199–2210.
- Chang, Tom Y., Huang, Wei, Wang, Yongxiang, 2018. Something in the air: pollution and the demand for health insurance. *Rev. Econ. Stud.* 85 (3), 1609–1634.
- Chay, Kenneth Y., Greenstone, Michael, 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.
- Chay, Kenneth Y., Greenstone, Michael, 2005. Does air quality matter? Evidence from the housing market. *J. Polit. Econ.* 113 (2), 376–424.
- Chen, Ting, Kai-sing Kung, James, 2019. Busting the ‘Pricelings’: the campaign against corruption in China’s primary land market. *Q. J. Econ.* 134 (1), 185–226.
- Chen, Yuyu, Ebenstein, Avraham, Greenstone, Michael, Li, Hongbin, 2013. Evidence on the impact of sustained exposure to air pollution on life expectancy from China’s Huai river policy. *Proc. Natl. Acad. Sci. Unit. States Am.* 110 (32), 12936–12941.
- Chen, Shuai, Oliva, Paulina, Zhang, Peng, 2018. Air Pollution and Mental Health: Evidence from China. NBER Working Paper w24686.
- Clark, Ximena, Hatton, Timothy J., Williamson, Jeffrey G., 2007. Explaining US immigration, 1971–1998. *Rev. Econ. Stat.* 89 (2), 359–373.
- Conlin, Michael, O’Donoghue, Ted, Timothy, J., Vogelsang, 2007. Projection bias in catalog orders. *Am. Econ. Rev.* 97 (4), 1217–1249.
- Deryugina, Tatyana, Heutel, Garth, Miller, Nolan H., Molitor, David, Reif, Julian, 2019. The mortality and medical costs of air pollution: evidence from changes in wind direction. *Am. Econ. Rev.* 109 (12), 4178–4219.
- Ebenstein, Avraham, Zhao, Yaohui, 2015. Tracking rural-to-urban migration in China: lessons from the 2005 inter-census population survey. *Popul. Stud.* 69 (3), 337–353.

- Ebenstein, Avraham, Fan, Maoyong, Greenstone, Michael, He, Guojun, Yin, Peng, Zhou, Maigeng, 2015. Growth, pollution, and life expectancy: China from 1991-2012. *Am. Econ. Rev.* 105 (5), 226–231.
- Ebenstein, Avraham, Lavy, Victor, Roth, Sefi, 2016. The long-run economic consequences of high-stakes examinations: evidence from transitory variation in pollution. *Am. Econ. J. Appl. Econ.* 8 (4), 36–65.
- Feng, Shuaizhang, Krueger, Alan B., Oppenheimer, Michael, 2010. Linkages among climate change, crop yields and Mexico-US cross-border migration. *Proc. Natl. Acad. Sci. Unit. States Am.* 107 (32), 14257–14262.
- Feng, Shuaizhang, Oppenheimer, Michael, Schlenker, Wolfram, 2015. Weather Anomalies, Crop Yields and Migration in the US Corn Belt. " NBER Working Paper w17734.
- Fischer, Stanley, 2003. Globalization and its challenges. *Am. Econ. Rev.* 93 (2), 1–30.
- Fu, Shihe, Brian Viard, V., Zhang, Peng, 2021. Air pollution and manufacturing firm productivity: nationwide estimates for China. *Econ. J.* 131 (640), 3241–3273.
- Ghanem, Dalia, Zhang, Junjie, 2014. Effortless Perfection": do Chinese cities manipulate air pollution data? *J. Environ. Econ. Manag.* 68 (2), 203–225.
- Gupta, Pawan, Christopher, Sundar A., Wang, Jun, Gehrig, Robert, Lee, Y.C., Kumar, Naresh, 2006. Satellite remote sensing of particulate matter and air quality assessment over global cities. *Atmos. Environ.* 40 (30), 5880–5892.
- Hanlon, W. Walker, 2020. Coal smoke, city growth, and the costs of the industrial revolution. *Econ. J.* 130 (626), 462–488.
- Hanna, Rema, Oliva, Paulina, 2015. The effect of pollution on labor supply: evidence from a natural experiment in Mexico City. *J. Publ. Econ.* 122, 68–79.
- Harrison, P.J., Stevens, C.F., 1976. Bayesian forecasting. *J. Roy. Stat. Soc. B* 38 (3), 205–247.
- Heblich, Stephan, Trew, Alex, Zylberberg, Yanos, 2021. East-side story: historical pollution and persistent neighborhood sorting. *J. Polit. Econ.* 129 (5), 1508–1552.
- Ito, Koichiro, Zhang, Shuang, 2020. Willingness to pay for clean air: evidence from air purifier markets in China. *J. Polit. Econ.* 128 (5), 1627–1672.
- Jacobson, Mark Z., 2002. Atmospheric Pollution. History, Science and Regulation. Cambridge University Press.
- Jans, Jenny, Johansson, Per, Peter Nilsson, J., 2018. Economic status, air quality, and child health: evidence from inversion episodes. *J. Health Econ.* 61, 220–232.
- Jessoe, Katrina, Manning, Dale T., Edward Taylor, J., 2018. Climate change and labor allocation in rural Mexico: evidence from annual fluctuations in weather. *Econ. J.* 128 (608), 230–261.
- Khanna, Gaurav, Liang, Wenquan, Mobarak, Ahmed Mushfiq, Song, Ran, 2021. The Productivity Consequences of Pollution-Induced Migration in China. " NBER Working Paper w28401.
- Kinnan, Cynthia, Wang, Shing-Yi, Wang, Yongxiang, 2018. Access to migration for rural households. *Am. Econ. J. Appl. Econ.* 10 (4), 79–119.
- Kumar, Naresh, Chu, Allen D., Foster, Andrew D., Peters, Thomas, Willis, Robert, 2011. Satellite remote sensing for developing time and space resolved estimates of ambient particulate in Cleveland, OH. *Aerosol. Sci. Technol.* 45 (9), 1090–1108.
- Molina, Teresa, 2021. Pollution, ability, and gender-specific investment responses to shocks. *J. Eur. Econ. Assoc.* 19 (1), 580–619.
- Neidell, Matthew, 2009. Information, avoidance behavior, and health: the effect of ozone on asthma hospitalizations. *J. Hum. Resour.* 44 (2), 450–478.
- Passey, Jeffrey S., Van Hook, Jennifer, Bean, Frank D., 2004. Estimates of the Legal and Unauthorized Foreign-Born Population for the United States and Selected States, Based on Census 2000. " *US Bureau of the Census*, Washington, DC.
- Pedersen, Peder J., Pytlikova, Mariola, Smith, Nina, 2008. Selection and network effects—migration flows into OECD countries 1990-2000. *Eur. Econ. Rev.* 52 (7), 1160–1186.
- Simonsohn, Uri, 2009. Weather to go to college. *Econ. J.* 120 (543), 270–280.
- Tanaka, Shinsuke, 2015. Environmental regulations on air pollution in China and their impact on infant mortality. *J. Health Econ.* 42, 90–103.
- Zhang, Junjie, Mu, Quan, 2018. Air pollution and defensive expenditures: evidence from particulate-filtering facemasks. *J. Environ. Econ. Manag.* 92, 517–536.
- Graff Zivin, Joshua, Neidell, Matthew, 2012. The impact of pollution on worker productivity. *Am. Econ. Rev.* 102 (7), 3652–3673.