DISCUSSION PAPER SERIES

IZA DP No. 8916

Severe Air Pollution and Labor Productivity

Teng Li
Haoming Liu
Alberto Salvo

March 2015

Forschungsinstitut
zur Zukunft der Arbeit
Institute for the Study
of Labor

Severe Air Pollution and Labor Productivity

Teng Li
National University of Singapore

Haoming Liu
National University of Singapore and IZA

Alberto Salvo
National University of Singapore

Discussion Paper No. 8916
March 2015

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

IZA Discussion Paper No. 8916
March 2015

ABSTRACT

Severe Air Pollution and Labor Productivity

We examine day-to-day fluctuations in worker-level output over 15 months for a panel of 98 manufacturing workers at a plant located in an industrial city in Hebei province, north China. Long-term workers earn piece-rate wages, with no base pay or minimum pay, for homogeneous tasks performed over fixed 8-hour shifts. Over the sample period, ambient fine-particle (PM2.5) mass concentrations measured at an outdoor air monitor located 2 km from the plant ranged between 10 and 773 micrograms per cubic meter ($\mu g/m^3$, 8-hour means), variation that is an order of magnitude larger than what is observed in the rich world today. We document large reductions in productivity, of the order of 15%, over the first 200 $\mu g/m^3$ rise in PM2.5 concentrations, with the drop leveling off for further increases in fine-particle pollution. A back-of-the-envelope calculation suggests that labor productivity across 190 Chinese cities could rise by on average 4% per year were the distributions of hourly PM2.5 truncated at 25 $\mu g/m^3$. We also find reduced product quality as pollution rises. Our model allows for selection into work attendance, though we do not find particle pollution to be a meaningful determinant of non-attendance, which is very low in our labor setting. Subsequent research should verify the external validity of our findings.

JEL Classification: J24, Q51, Q52, Q53, O44, R11

Keywords: air pollution, labor productivity, labor supply, PM2.5, environmental damage

Corresponding author:

Haoming Liu
Department of Economics
National University of Singapore
1 Arts Link
Singapore 117570
E-mail: ecsliuhm@nus.edu.sg

1 Introduction

Among the world's capital cities, Beijing has become notorious for its air pollution. Between December 2013 and February 2014, the hourly mass concentration of particulate matter of diameter up to 2.5 microns (PM2.5) in Beijing's air, as measured at the US Embassy, averaged 129 micrograms per cubic meter (μg/m^3). Several hundred kilometers away, in an industrial city that is home to workers whom this paper examines, an air monitor located 2 km from their workplace recorded hourly PM2.5 concentrations that averaged 221 μg/m^3 over the same winter months, that is, over two-thirds higher than that in Beijing. The maximum reading during our study period, from April 2013 to June 2014, was reached on December 25, 2013 at 2 pm, namely 919 μg/m^3. This severe level of fine-particle pollution in ambient air is an order of magnitude higher than the World Health Organization's guidelines for human exposure sustained over a period of 24 hours—a mean concentration for PM2.5 of no more than 25 μg/m^3 is deemed "healthy."[1] While we cannot disclose the name of the city, we note that the city itself and its outskirts are home to several large coal-fired power generating units, steel mills, and cement kilns. The economic geography of the area, in part due to its abundance of coal deposits, as well as its physical geography, which includes a mountain range on one side, help explain the state of its air relative to that in more famous Beijing.

China's severe environmental degradation is the by-product of the nation's fast growth in economic activity, and fear of dampening this growth rate is widely regarded to be one key reason behind its government's revealed hesitation, to date, in taking swift action to abate air pollution (Ebenstein et al., 2015). While research on the social benefits of cleaner air has focused on public health, such as longer life expectancy (e.g., Chen et al., 2013), another potentially large source of gains, labor productivity, has received considerably less attention in the policy debate. We conjecture that the relative absence of labor productivity in policy discussions concerning the damage caused by China's

[1] See World Health Organization (2006). The National Ambient Air Quality Standards (NAAQS) set by the United States Environmental Protection Agency establish a PM2.5 "primary standard"—intended to provide public health protection including the health of vulnerable groups—of 12 μg/m^3 over one year of exposure, and a 24-hour standard of 35 μg/m^3. See http://www.epa.gov/air/criteria.html.

air owes at least in part to the paucity of rigorous empirical work on this potentially pervasive manifestation of morbidity.

To the best of our knowledge, only four previous studies have directly addressed the impact of air pollution on labor supply or productivity while attemping to overcome some major hurdles that the literature faces. First, people living in areas with higher ambient air pollution may differ from residents in areas with better air quality, making it difficult to use cross-sectional data to establish a causal relationship between air quality and labor productivity. The second challenge is that because most datasets do not record worker-level output or earnings on a daily basis, one cannot use daily variation in air quality to examine its effect on worker productivity while flexibly controlling for individual and seasonal heterogeneity. Nevertheless, the limited existing empirical evidence indeed suggests that air quality has a significant adverse effect on labor productivity and labor supply.[2]

Using a dataset that contains the daily work performance and operating environment of citrus pickers in southern California in 1973/74, Crocker and Horst (1981) find that in-sample ozone pollution reduces the productivity of the outdoor workers by up to 2%. This finding is confirmed by Graff Zivin and Neidell (2012) who use payroll data from a farm in California's Central Valley. Graff Zivin and Neidell (GZN hereafter) show that a 10 ppb (parts per billion) increase in average ozone exposure results in a 4% reduction in the productivity of workers picking berries and grapes. In their setting, workers are paid in proportion to their individual output, i.e., on a piece-rate basis, which acts as an incentive for workers to perform. In a recent working paper, Chang et al. (2014) find that outdoor air pollution also in California—but now examining fine-particle (PM2.5) rather than ozone pollution—also affects the indoor work environment. Chang et al. find that a 10 μg/m^3 increase in outdoor PM2.5 concentrations decreases worker productivity by roughly 6%, while having no significant effect on working hours. One challenge Chang

[2] Ostro (1983), for example, investigates the relationship between particle pollution (total suspended particles, TSP) and labor supply using cross-sectional survey-based data from the 1976 US Health Interview Survey. Workers responded to a survey question asking them how many days in the past two weeks had illness or injury prevented them from working. Ostro reports that "a 10% decrease in ambient levels of TSP is related to a 4.4% decrease in WLD (work loss days)."

et al. face is that the indoor work environment they study is naturally ventilated and, as the authors point out, variation in temperature, if not adequately controlled for, may confound variation in fine-particle concentrations, which is their main variable of interest. Finally, Hanna and Oliva (2014) exploit the closure of a large refinery in Mexico City, a policy that induced exogenous variation in air quality. Examining aggregate data from an administrative source, and using a difference-in-difference design, Hanna and Oliva find that the policy-induced reduction in air pollution led to a 4% increase in weekly working hours for workers living within a 5 km radius of the refinery.

By examining a panel of individual workers at a textile mill located in the province of Hebei, in northern China, our paper adds to this sparse empirical literature. In particular, we are able to examine the effect of severe air pollution on labor productivity. As alluded to in the opening paragraph, the variation in air pollution observed in China today is much larger than what the extant literature has studied. For instance, daily mean PM2.5 concentrations range between 2 and 60 $\mu g/m^3$ in Chang et al. (2014); in our data, daily means vary between 25 and 687 $\mu g/m^3$. This large observed variation enables us to examine the magnitude and the shape of the relationship between particulate pollution and labor productivity, which can potentially be nonlinear. Knowing whether there is a nonlinear dose-response will allow policymakers to more accurately estimate the economic cost of air pollution to which workers in developing countries today, from China to India to Indonesia, are exposed—levels that are an order of magnitude higher than in developed countries in North America and Europe. We note that even at ambient air levels observed in the United States, PM2.5 is understood to be a major source of health damage, including minor restricted activity (Fann et al., 2012).

We gained access to daily payroll data for 98 machine operators working on parallel tasks at a common workplace—a department of an industrial plant—during 15 months. We argue that the labor institution we study coupled with the structure of the data are ideally suited to the the research question at hand, sharing several features of recent impactful studies, at lower levels of pollution, in particular, GZN. Like GZN, we are able to observe labor supply choices and outcomes for the same worker over time, as ambient air pollution fluctuates day in day out, allowing us to control for worker heterogeneity. Workers are paid on a piece-rate basis for their individual output, which rewards them for the level of effort that they individually supply (and has the added feature, as GZN argue, that this plausibly reduces measurement error in this relatively high-frequency output variable).[3] The piece rate does not vary across workers, consistent with the homogeneity across tasks and workstations.[4] Unlike GZN's outdoor setting, however, the indoor work environment we study is temperature controlled and sheltered from rain and wind,[5] which we argue enables us to directly control for an important possible confounder of the effect of pollution on labor productivity. An added benefit of temperature control at the workplace is that it provides us with an exclusion restriction to identify possible selection into work attendance, and thus control for selection bias, to the extent that outdoor temperature—and weather more generally—shifts the probability of showing up to work, e.g., by changing the value of the outside option such as leisure (Graff Zivin and Neidell, 2014).

Importantly, while temperature is controlled, indoor air that the workers are exposed to is not filtered and exchanges with the outdoor environment through one large main door (leading directly outside and through which yarn and fabric are wheeled in and out, respectively), as well as through a long array of windows that line one of the sidewalls (installed over thirty years ago). More generally, an environmental engineering literature reports tight correlation between outdoor and indoor concentrations for pollutants such as PM2.5, for typical indoor "microenvironments," as well as high indoor-outdoor ratios (see references in Section 2 and in Chang et al. (2014)). Finally, our labor output data consists of not only quantity produced (meters of fabric) but also the quality of

[3]To be precise, some of the workers in GZN are paid in proportion to their joint output, which is not the case in our labor market. Further, in our setting there is no minimum wage or base pay, and workers are on long-term contracts. Moreover, in contrast to Chang et al. (2014), the work shift in our setting is of fixed duration, so workers do not choose how many labor hours to supply, conditional on work attendance, which we model. Otherwise, it is conceivable that a worker might choose the number of hours worked in part as a function of the level of airborne contaminants, and this choice might be made jointly with the effort level. We also need not worry that workers might slacken to prolong their normal work hours to earn an overtime rate.

[4]Again for comparison, GZN observe workers picking different crops, with the piece rate varying across the crops. Since the authors pool observations across crops, they need to standardize output units.

[5]To compare, the pear-packing plant that Chang et al. (2014) study has no temperature control.

production (meters of defective fabric), allowing us to examine the effect of air pollution on the product of labor along this additional margin.

Our main finding is that higher mass concentrations of PM2.5 in outdoor air, as recorded at the air monitoring site situated 2 km from the plant, have a significant adverse impact on contemporaneously observed worker productivity. While the sign of the estimated impacts is consistent with what has been found in the small extant literature, the range of variation is larger than that, to the best of our knowledge, ever examined. Our main result is best described in Figure 4 below, where the scatter depicts data—mean output per attendant worker for every date by (8-hour) work shift in the sample—and the lines indicate alternative non-linear fits.[6] Starting at the sample minimum of 10 $\mu g/m^3$ (8-hour means), every additional 10 $\mu g/m^3$ of exposure to PM2.5 over a worker's shift reduces her total output by 4.3 meters of fabric, highly significant both statistically and economically, equivalent to about 0.9% of mean output in the sample (509 m/worker-shift). Estimated marginal effects are similar up to the 75th percentile of the PM2.5 distribution, at 149 $\mu g/m^3$, and halve thereafter (-2.0m of fabric per 10 $\mu g/m^3$ increase). Beyond the 90th percentile, at 230 $\mu g/m^3$, the estimated marginal effect is a low and only marginally significant -0.5m per 10 $\mu g/m^3$ increase.

In sum, integrating over the first 200 $\mu g/m^3$ increase in ambient PM2.5 concentrations, from 10 to 210 $\mu g/m^3$, output falls by 71m, equivalent to 14% of mean output (and with a standard error of 7m). These are large effects. The non-linearity of the relationship over such wide a range is another result that is new to the literature. We also find plausible effects of particle pollution on worker attendance which, while statistically significant, are economically small; irrespective of its determinants, worker non-attendance is low in our labor setting. For the subset of our sample for which we observe the quality of labor product, namely meters of defective fabric produced by worker by shift, we find that it rises with particle pollution.

To better interpret the implications of our findings, including the estimated non-linear

[6] As we subsequently explain, regression models include fixed effects for year-month, day-of-week, public holidays (that the workplace did not observe), time-of-day and individual worker, among other controls, including non-linear functions of co-pollutant concentrations, namely SO_2 and CO which, unlike ozone, are emitted from surrounding industry and may also penetrate indoors.

relationship over a wide range of particle pollution, we perform a back-of-the-envelope calculation. We take our estimated worker-day level model and predict aggregate output produced by the studied labor institution (i.e., the department of 98 individual workers and their characteristics) were it to be hypothetically transplanted to each of 190 major cities in mainland China, under currently observed ambient air levels of PM2.5. This provides a baseline for each Chinese city over the course of 12 (in-sample) months. We then repeat the exercise, again for each of the 190 cities, predicting output by the modeled labor institution hypothetically transplanted to the given city, under the counterfactual scenario that PM2.5 levels in the city were not to exceed 25 μg/m^3 at any hour during the course of the 12 months. Aggregating across the 190 cities and over the different seasons of the year, a comparison between counterfactual and actual scenarios suggests that aggregate output would rise by 3.8%. Of course this exercise has several limitations, not least that it ignores the costs of pollution control given the state of China today, but it serves to highlight one large and often overlooked benefit of abating air pollution.[7]

The rest of the paper is structured as follows. Section 2 discusses the labor institution and the environment workers are exposed to. Section 3 lays out the conceptual framework and the empirical model of worker choices. Sections 4 and 5 estimate the impact of PM2.5 pollution on worker non-attendance and worker output, respectively. Section 6 presents implications of our findings and concludes.

2 Data and descriptive analysis

Our data on worker output is made available thanks to a special agreement with management at a textile operation in a city in the northern province of Hebei, China. We do not name the city to protect the operation's confidentiality, as agreed with management.[8] The operation we examine—namely, the textile department—is part of a larger industrial complex. The textile department obtains yarn from an upstream operation

[7]A second back-of-the-envelope calculation in the Appendix considers whether the principal at the workplace would privately benefit from installing indoor pollution control, as in a semiconductor plant.

[8]For reference, Hebei province surrounds Beijing. By management, we refer to several employees at the firm in managerial, administrative or technical positions with whom we have developed a relationship.

within the firm and processes the input into fabric rolls, which are then sold to buyers across China and abroad.[9] The department runs around the clock, seven days per week, but shuts down over short multiday periods which typically include or overlap with public holidays such as the weeks of Chinese New Year and National Day. Our sample comprises the period between April 1, 2013 and June 30, 2014. During this period there were 456 days, 54 days of which (12%) the department produced no output. Excluding these plant holidays, our sample thus consists of 402 dates.

Labor organization. The department operates three shifts of fixed, 8-hour duration, starting at: (i) 0 am (to 8 am), (ii) 8 am (to 4 pm), and (iii) 4 pm (to 0 am the next day). Workers are divided into three teams which rotate, as a team, among these shifts every four workdays. For example, in our sample, in 2013, Team 1 worked the 0 am shift from April 6 to 9, the 4 pm shift from April 10 to 13, the 8 am shift from April 14 to 17, and so on. Thus, we observe each worker and her team repeatedly working in each of the three shifts. This allows us to control for effects that may be specific to the shift, such as the likelihood that a worker is absent on a given date, or any shift-specific variation in productivity. As we describe below, we observe each worker's quantity of output, total and defective, expressed in meters of fabric, by date (and we know the shift). According to management, workers typically work the fixed 8-hour shift along with their team, and do not select the number of hours worked; for this reason, this measure of hours is not recorded. We return to this institutional feature below.

A shift is staffed by about 30 workers and one supervisor. Each workstation is operated by one worker, and workstations operate in parallel. Each worker's workstation consists of 10 machines, or looms, valued at about US$ 30,000 each. The task being performed is quite homogeneous across workers and workstations. The job description is to walk up and down the workstation and attentively observe the looms as they weave threads into fabric. Typically, 1 to 2 threads will break apart per machine per hour,

[9] There are two upstream operations within the firm. First, a preparatory department purchases cotton from the nearby provinces of Shandong and Jiangsu and the northwestern province of Xinjiang, which it then cleans and draws. Next, a yarn department spins and passes cotton into yarn. The textile department that we study takes the yarn and weaves threads into fabric. We are unaware of any input shortages, including electricity, during the sample period. Also, in view of its widespread geographic nature, demand for the operation's produce is unlikely to depend on local economic activity.

requiring skill and effort from the worker to reconnect the thread, a task that might take several minutes. Therefore, a worker will typically reconnect 10-20 threads across the 10 machines every hour, and in extreme cases this number can rise to 50-60 threads per hour. Day in, day out, a worker tends to return to the same workstation. There is some, albeit limited, variation in composition across workstations, in that some workstation's machines are programmed to produce fabric of type 133, others produce fabric type 134. In practice, "standard" output rates vary slightly according to whether machines are set up to produce fabric type 133 or 134: 510 and 495 meters of fabric per 10-machine workstation per 8-hour shift, respectively.[10] Due to setup costs, machine composition changes only rarely within workstation.

The machines are fairly new, having been purchased in block when the firm was privatized in the early 2000s. We do not observe machine breakdowns. Management informed us that machines, likely because they are only a decade old, break down rarely—preserving machinery is another reason justifying indoor temperature control in this firm and the wider industry. In the event that a machine breaks down, the worker's variable pay is pro-rated based on her output using the functioning machines. (We subsequently detail both variable pay and temperature control.) Every batch of 10 machines, comprising a workstation, undergoes planned maintenance every fortnight, which may include a simple inspection by a maintenance crew member.

While work in the textile department is capital intensive, productivity depends critically on the quality of its workers. Our understanding is that the quality of labor is a function of the skill (experience) and attentiveness (effort level) of the worker. The appropriate unit of analysis is the individual worker, as there is minimal complementarity across workers.

There are three types of employees operating the machines. What we label "department" workers form the bulk of our sample. These workers are assigned continuously to work at the textile department. When we observe a department worker produce zero output on a given date, other than a plant holiday, our interpretation is that the worker

[10] The numbers 133 and 134 figures correspond to the fabric's warp. The weft in both cases is 72.

did not attend work. (We also observe some department workers who, before the sample period ends, are reassigned to other departments or leave the firm, and others who, after the sample period begins, join as department workers, including those coming from other departments within the firm.) In total, our sample includes 98 department workers. The second type of employee who operates the machines is a "cross-department" worker. These workers work across departments (e.g., preparatory, yarn, textile), to smoothen fluctuations in aggregate output brought about by planned and unplanned absences in the workforce. Our sample includes 12 cross-department workers, and for each of these workers we observe intermittent (though often adjacent) dates with non-zero output. Finally, the third type of machine operator is labeled a "substitute" worker. These workers are also brought in to meet demand and to keep the machines from staying idle. Our data does not include information on these workers.

We base our analysis on the panel of 98 regular department workers.[11] For perspective, the median worker was born in 1970, is female and Han (race), has nine years of schooling, was hired by the firm in 1991, and is local to the same city as the firm (according to the department of human resources' records). This median worker is married, with 1994 being the median year of marriage, and has one child, with 1996 being the median birth year among first children. Thus, workers tend to be middle aged, long-term employees, and have older children. The median worker typically lives in the vicinity of the plant, so commuting costs are low. She is also currently on a five-year contract with the firm. Throughout the entire sample, we observe 33, 40 and 25 department workers attached to Teams 1, 2 and 3, respectively. As we show in the Appendix, Team 2 experienced more attrition (and entry) over the sample period: counts of department workers who are attached to Teams 1, 2 and 3 through the end of the sample, on June 30, 2014, are 27, 28 and 21 respectively.

With worker attention being critical and complementarity across workers being minimal, workers are paid by the amount of good-quality fabric they individually produce in each shift. During the sample period, the piece rate was increased once, from CNY

[11] Our results are robust to including the intermittent output by cross-department workers.

0.07 per meter of (133-equivalent) fabric prior to September 29, 2013 to CNY 0.1 subsequently. (By 133-equivalent meters we mean that any output of 134-type fabric is multiplied by 510/495 to account for the 3% higher throughput of machines when producing 133 relative to 134.) The large increase in the piece rate is one reason why we control for calendar month fixed effects in the empirical analysis.[12] This piece rate does not vary across workers, consistent with the homogeneity across tasks and workstations (with the exception of the 133 versus 134-type fabric adjustment explained above). Other than through benefits such as subsidized housing and health insurance, there is no base pay and all compensation is variable, again, reflecting the importance of skill and effort as inputs to the production function, with the worker's "skin in the game." Similarly, there is no minimum wage or threshold level above which variable compensation applies (as there is, for example, in GZN's study of pollution in California). At the end of her shift, a worker marks the point at which her production ended and at which the production of her colleague, working the subsequent shift at the same workstation, begins. The rolls of fabric, once complete, are subsequently inspected for defects and payroll records are updated. The most common defect is fabric that is short of threads (e.g., warps) that the worker—or machine—failed to detect. While workers are paid based on the amount of defect-free production, they do not receive pecuniary punishment for producing defects. To the best of our knowledge, there are no disputes over what constitutes defect-free fabric, or how individual-level output is recorded. It is these payroll records that we gained access to. To provide perspective on how much workers earn, a worker who produces at the sample median of 506 m in a shift earns (after the pay rise) CNY 51, equivalent to about US$ 9, in that shift.[13]

The key aspect of the individual level productivity data is its longitudinal structure. Similar to GZN, we are able to follow the same worker, date-shift by date-shift, which allows us to control for worker heterogeneity. In addition to the quantity of individual

[12]Our data indicates that the month in our sample with the highest attrition in the department was September 2013, with 13 departures alone.

[13]For a rough conversion into US$, divide CNY by 6. Also for perspective, the Hebei Statistics Bureau reports average annual earnings in the city to be a little over 35,000 CNY (www.hetj.gov.cn/hetj/tjsj/ndsj/101400644755604.html).

output produced by each worker (ID) over each shift, we observe worker characteristics, which we can use to learn more about worker heterogeneity. As we illustrate below, our source of identification, as in GZN, is the day-to-day covariance between the state of outdoor air pollution and the individual worker's productivity, once we control for seasonality and other potential confounders such as temperature in the workplace.

As in other studies (e.g., GZN), a worker occasionally chooses to not attend work. On the one hand, planned leaves, like plant holidays, are predetermined, so they are unlikely to depend on day-to-day variation in environmental quality. On the other hand, a worker might not attend work due to herself or a family member falling sick, or meteorological and other shocks raising commuting costs or shifting the value of the outside option, e.g., a leisure day spent outdoors. Such unplanned non-attendance may in part depend on the state of air (e.g., pollution might raise the likelihood of an asthma attack), or unplanned non-attendance and pollution might be correlated through weather shocks (i.e., these might directly affect health and air). We therefore model the worker's selection into work. We note that while health may drive work attendance, shifts in commuting costs and the value of leisure driven by weather and/or pollution shocks are unlikely to be sizable, given the setting: about 90% of workers live in nearby housing subsidized by the employer, and the employer has much information about the employee. As we show below, worker non-attendance rates are low, of the order of one day per month (on top of plant holidays, common to all workers). This includes both planned and unplanned leaves (we are unable to distinguish between the two).

Worker productivity. To illustrate the importance of worker heterogeneity, for every one of the 98 department workers in our sample we calculate her mean individual output per shift worked over the entire sample period. The distribution of mean worker performance is plotted in panel (a) of Figure 1. The figure shows workers' total output, i.e., including any defective fabric. The mode is just over 510 m per worker-shift. The panel indicates that the most productive workers can sustain a production rate that is up to 40% higher than that of the least productive workers.

Besides comparing mean performance across workers, panels (b) to (d) of Figure 1 compare mean performance across dates, separately for the 0 am, 8 am and 4 pm shifts. (One can think of shift here as the time of day.) To prepare these plots, we compute the mean output per worker for each of the 402×3 date-shift combinations in our sample. This day-to-day variation in the average productivity of the workforce is of key importance to our empirical strategy—our task is to uncover the extent to which this temporal variation in output is driven by variation in ambient PM2.5 concentrations.

Figure 2 provides the same descriptive statistics as Figure 1 but for defective, rather than total, output. We note that for one of the three teams of workers, Team 3, daily records on defective output are missing, so the figure considers only the 33+40=73 workers attached to Teams 1 and 2.[14] Somewhat surprisingly, some workers are reported to produce 0 m of defective fabric during the sample period, whereas others produce as much as 7 m of defects per shift on average. The variation of defective output is similar across shifts, with a mode at about 2 m per worker, less than 0.5% of the modal output. Defective quantity is low presumably because workers reduce the rate of output to prevent defects from being produced in the first place.

Environment workers are exposed to. The department is located inside a single-storey factory building. (See the Appendix for further details, including pictures.) Air exchanges between the outdoor and indoor environments through windows that are installed at the top of sidewalls, and one main large door which leads directly outside. The building was built in 1982 and has not gone through any major remodeling. This being northern China, winter conditions require central heating (Chen et al., 2013). Outdoor summer temperatures can rise above 35 degrees Celsius, and the indoor microenvironment is also air-conditioned. Indoor ambient air temperatures are recorded but are not available to us over the sample period. To gain perspective, however, we obtained copies of records on specific more-recent dates, such as July 14 and 20, 2014. While outdoor temperatures, which we observe, were in the range of 36 to 41 degrees Celsius, indoor temperatures were recorded between 25 and 31 degrees Celsius. We thus assume that workers are exposed to a work environment with reasonably well-functioning tempera-

[14]To complete the description of missing output records, we do not observe Team 3 worker output in June 2014. Defective output records are also missing on five specific dates surrounding plant holidays.

ture controls (we qualify the statement given the age of the system). This allows us to directly control for a typical confounder—ambient temperature—encountered in studies of air pollution on health and labor outcomes (e.g., Crocker and Horst, 1981). According to management, heating and air-conditioning systems are the norm in this industry and this part of China, in part because extreme temperatures might damage inputs (e.g., machines, yarn, labor) and outputs (fabric).

We have access to outdoor meteorological conditions recorded in the same city as the plant for every three-hour interval in the sample. Temperatures (three-hourly means) fell below -10 degrees Celsius on two occasions in February 2014, in the early hours of the morning. Temperatures exceeded 40 degrees Celsius in summer 2013 and May/June 2014, on 13 dates, typically in the early afternoon hours. The northern province of Hebei is dry year-round. Between April 2013 and June 2014, 0 mm of precipitation was recorded on more than nine-tenths of all three-hourly intervals; among the rare non-dry intervals, the median rate of rainfall is a low 0.15 mm/hour, with rain being less rare in warmer months than in colder ones. Humidity is significantly lower in January than in July, namely means of 37% versus 62%, respectively. Snow was recorded only on four dates, and it was labeled light to moderate. Wind speed, a determinant of ambient air pollutant concentrations, tends to be higher from April to September (an average of 6.6 miles per hour) compared with the remaining colder months (6.1 mph). Wind tends to blow more strongly in the afternoon hours. In terms of wind direction, winter temperatures tend to fall when wind blows from the north rather than the south. Wind direction is quite variable within season but the pattern is quite stable over the four seasons of the year. Wind blowing along the north-south axis is more common than along the east-west one, likely due in part to the mountain range to the west of the city. Daily means obtained from a second data source are highly consistent with the higher-frequency dataset. In sum, patterns in the meteorological data are highly plausible.

We obtain ambient air pollutant concentrations recorded by the Chinese Ministry of Environmental Protection. Mass concentrations, in $\mu g/m^3$, for PM2.5, PM10, SO_2, CO, NO_2, and O_3, were recorded every hour at an outdoor monitoring site located only 1.7

km from the plant we analyze. The times series are quite complete, e.g., PM2.5 and CO measurements are missing or invalid for only 8% and 9%, respectively, of the possible $456 \times 24 = 10944$ hourly observations between April 2013 and June 2014. These missing values are fairly evenly distributed throughout the sample and do not cluster on specific dates, i.e., missing values for PM2.5 on eight or more consecutive hours occur only on 24 occasions in the sample (e.g., April 22, 2013, May 18, 2013, etc), suggesting the site is well maintained. Where hourly observations are missing, we use the mean concentration for the pollutant at the given date-hour recorded at three other outdoor monitoring sites located between 3 and 5 km of the plant, also maintained by the Ministry. Air measurements at the city's four sites are highly correlated. The correlation coefficients for hourly PM2.5 concentrations at the closest site relative to that at each of the other three sites are 0.88, 0.89 and 0.91.

We have inspected the annual, weekly and diurnal cycles for PM2.5 and other pollutant concentrations. Patterns are consistent with measurements elsewhere, even if for other atmospheric systems, e.g., Davis (2008), Salvo and Geiger (2014). We briefly describe these patterns here, and include descriptive regressions below, as this informs one of our subsequent identification strategies, based on instrumental variables. PM2.5 levels tend to be considerably higher in the colder months over the warmer months—as much as several hundred μg/m^3—and slightly higher in the morning hours than in the afternoon—a few dozen μg/m^3. PM2.5 as a proportion of PM10 mass concentrations range between 0.4 and 0.8, with higher ratios being observed in the winter. SO_2 concentrations also peak in the winter, rising all the way from late afternoon and through the night, likely due in part to high-sulfur coal-fired power generation responding to demand as temperatures drop. CO concentrations show a similar pattern to SO_2. Ozone concentrations are higher in the warmer months and at 3 pm, when radiation and temperatures increase, and are inversely correlated with NO_2 concentrations, consistent with ozone chemistry (Madronich, 2014). In terms of weekly cycles, NO_2, SO_2 and CO concentrations are somewhat lower on Sundays compared with other days of the week, consistent with, e.g., nitrogen dioxide's role as a signature of anthropogenic sources (Beirle et al., 2003). In contrast, PM2.5 concentrations exhibit a less pronounced weekly cycle, given their longer life-time in the atmosphere, consistent with observations elsewhere (Salvo et al., 2015). In short, we deem the Ministry's pollution data to be reliable.

We do not observe PM2.5 mass concentrations inside, right by the workstations. We rely on a literature in environmental sciences, engineering and epidemiology that finds that fine (including ultrafine) particles penetrate indoors, e.g., Morawska et al. (2001), Cyrys et al. (2004), Gupta and Cheong (2007). For example, Cyrys et al. (2004) report, for a given microenvironment they study, that with "closed windows, the I/O (indoor-outdoor) ratios for PM2.5 are ... 0.63 ... (and) that more than 75% of the daily indoor variation could be explained by the daily outdoor variation for those pollutants." The workplace we study is set in a building that is over 30 years old and is directly linked to a ventilated outdoor environment by way of a long row of closed (though likely imperfectly sealed) windows and a large open door, through which yarn (input) and fabric (output) is regularly wheeled in and out (again, see pictures). For comparison, Chang et al. (2014) also proxy for the quality of indoor air using measurements from official outdoor monitors. With observational data, high-frequency indoor measurements are unlikely to be available. In studies where indoor air is measured, this is likely on an experimental basis and may affect behavior; further, measurements (e.g., on handhelds, measuring particle counts) may be of inferior quality compared to an official monitoring site that is regularly subject to standard QA/QC procedures (quality assurance/quality control). It is also conceivable that the worker may arrive to work already feeling unwell due to her exposure to pollution in the preceding hours, while at home or elsewhere (PM2.5 concentrations tend to be persistent over adjacent hours compared to variation across multiple days, when, e.g., wind conditions change).[15]

Table 1 reports sample statistics, consistent with the discussion above, at the worker-date-shift level (e.g., output), the worker level (demographics), and the date-shift level (environment). Given the frequency of the worker output data, we compute means of

[15]A possibility we are considering is to conduct simultaneous PM2.5 measurements, based on 8-hour filters, or perhaps less conspicuous high-frequency handhelds, at different locations of the indoor and adjacent outdoor environments.

hourly pollutant concentrations within each 8-hour (i.e., 402×3) date-shift combination in the sample. Again, PM2.5 concentrations are those measured at a monitor located 2 km away (and we impute missing hourly observations using available observations from the three other nearby monitors that same hour), and these tend to exceed those measured by the US Embassy in Beijing.

The mean PM2.5 mass concentration is 124 μg/m^3. Figure 3 plots the kernel density function of 8-hour mean PM2.5 concentrations across all date-shifts in the sample period, in panel (a), as well as residuals of mean PM2.5 concentrations when these are regressed on calendar-month, day-of-week and time-of-day (shift) fixed effects. The wide variation in ambient fine-particle pollution, even when seasonality is accounted for, as is evident in panel (b), will enable us to compare the productivity benefits from abating pollution over the wide range of pollution levels we observe in our sample.[16] For comparison, panel (c) plots the kernel density function of 8-hour mean PM2.5 concentrations measured by the US Embassy in Beijing, over the same sample period. Relative to levels recorded by the Ministry 2 km from the plant, concentrations in Beijing several hundred km away tend to be lower (a mean of 95 against 124 μg/m^3 over the study period)—this happens particularly in the mornings during the colder months.

Observed determinants of fine-particle pollution. Given the central importance of environmental variables to our study, we now examine the observed determinants of PM2.5 in ambient air. The purpose is twofold. First, an analysis of the covariance of PM2.5, season and meteorology serves as a check on the quality of the data. Second, the analysis provides the first-stage to a 2SLS estimator we subsequently report on.

Table 2 regresses mean PM2.5 concentrations recorded at the monitor 2 km away on time-varying fixed effects and meteorological conditions recorded in the same city. An observation is a date by shift combination over the sample period April 1 2013 to June 30 2014. PM2.5 concentrations are missing for 5 combinations, thus there are 456×3-5=1363 observations. Column (1) shows the importance of seasonal, weekly and within-day

[16]For perspective, (24-hour) PM2.5 mass concentrations vary between 2 to 60 μg/m^3 in Chang et al. (2014)'s sample, in California. In contrast, (8-hour) PM2.5 levels exceed 60 μg/m^3 in about three-quarters of our sample. See the final section on implications.

cycles, both anthropogenic and natural. The predictive power of year-month, day-of-week and time-of-day (i.e., for 8-hour shifts starting at 0 am, 8 am and 4 pm) indicators is such that the R^2 is 38%. (We allow the shift fixed effect to vary by the quarter of the year, and could do likewise with day of week.) Adding meteorological conditions in column (2) raises R^2 to 45%. PM2.5 concentrations are increasing in temperature and humidity (maximum and minimum recorded for the date-shift observation). PM2.5 concentrations are increasing in atmospheric pressure. Relative to zero rain, precipitation averaging as low as between 0 and 0.5 mm/hour is associated with a reduction of 19 μg/m^3 in PM2.5 levels. A similar downward effect is observed for precipitation in excess of 0.5 mm/hour, which is the case for only 2% of observations. Relative to weak winds, wind speeds between 5 and 10 miles/hour reduce PM2.5 concentrations by 7 μg/m^3; even stronger winds lower PM2.5 concentrations by 16 μg/m^3. The direction of wind also has systematic effects. Not only are the point estimates intuitively signed—they tend to be highly statistically significant.

While column (2) restricts meteorological effects to be the same year-round, column (3) reports that allowing quarterly variation in meteorological effects adds 5 percentage points to the R^2 (estimates are omitted for brevity). Column (4) indicates that lagged meteorology also determines contemporaneous PM2.5 concentrations. For example, rain in the preceding 8 hours at mean rates of 0-5 mm/hour and 5+ mm/hour significantly lowers PM2.5 by 27 and 53 μg/m^3, respectively, whereas contemporaneous rain at these rates has a smaller effect (estimates not shown). Column (5) repeats the specification shown in column (2) taking PM2.5 concentrations in logs as the dependent variable. For example, precipitation up to 0.5 mm/hour and wind speeds above 10 mph significantly lower PM2.5 by 15% and 16%, respectively, relative to no rain and weak winds.

Worker output against fine-particle pollution in the raw data. Figure 4 plots mean worker output per shift against the contemporaneous 8-hour mean PM2.5 concentration measured 2 km from the plant. The top and bottom panels use linear and logarithmic ordinates, respectively, for PM2.5. Each observation in the scatterplot is a date-shift pair in the sample, and we average total output (i.e., including any defects) across workers in that date-shift (we return to the fitted lines in Section 5). The relationship shown in the figure is striking. When fine-particle pollution is very high, mean output per worker per shift lies at the lower end of the productivity distribution. Output rarely exceeds the sample mean of 509 m per worker per shift when PM2.5 levels rise above 250 μg/m^3; in particular, shift output per worker exceeds the sample mean only four times out of the 100 date-shifts with PM2.5 levels above 250 μg/m^3. Moreover, among date-shifts with PM2.5 concentrations below 250 μg/m^3 (but still mostly severe), there is clearly a steep negative relationship between labor productivity and PM2.5.

Figure A.8 in the Appendix shows that the production of defective fabric, while low, tends to rise with air pollution. For perspective, when PM2.5 levels exceed 250 μg/m^3, the mean quantity of defects per worker lies below the sample mean of 3.8 m only 12 times out of the 100 date-shifts. For date-shifts with PM2.5 levels higher than 500 μg/m^3, defects never fall below 5 m per worker, and for date-shifts with PM2.5 exceeding 600 μg/m^3, defects never fall below 10 m.

3 Conceptual framework and empirical model

3.1 Conceptual framework

To fix ideas, consider a worker of ability a who works individually over a fixed shift of eight hours. During this shift, the worker is exposed to ambient air pollution Ω^I (the superscript denotes the *indoor* work microenvironment), chooses effort level e, and produces gross output quantity $\tilde{q}$, a fraction $0 \leq \xi \leq 1$ of which comes out free of defects. The worker incurs an effort cost given by the function $c(e, \Omega^I)$, which exhibits the following properties:

$$\frac{\partial c(\cdot)}{\partial e} > 0, \frac{\partial c(\cdot)}{\partial \Omega^I} \geq 0, \frac{\partial^2 c(\cdot)}{\partial e^2} \geq 0, \frac{\partial^2 c(\cdot)}{\partial e \partial \Omega^I} \geq 0.$$

These conditions state that the cost of working increases in effort and pollution (strictly and weakly, respectively), and the positive marginal cost of effort weakly increases in both effort and pollution. We do not make any assumption on the sign of $\partial^2 c(\cdot)/\partial {\Omega^I}^2$.

We specify net, or defect-free, output quantity $q = \tilde{q}\xi = q(e,a)$ as an increasing and concave function of the effort level, with production and marginal product increasing strictly and weakly, respectively, in ability. Formally,[17]

$$\frac{\partial q(\cdot)}{\partial e} > 0, \frac{\partial^2 q(\cdot)}{\partial e^2} < 0, \frac{\partial q(\cdot)}{\partial a} > 0, \frac{\partial^2 q(\cdot)}{\partial e \partial a} \geq 0.$$

As discussed, the worker is paid a piece rate p per unit of defect-free output, does not receive pecuniary punishment for producing defects, and there is no daily minimum wage. The piece rate is invariant to air quality and does not vary across workstations, with workers performing the same parallel tasks on machines of similar vintage and quality.[18]

Conditional on coming to work, the worker solves:

$$\arg\max_{e} pq(e,a) - c(e,\Omega^I), \tag{1}$$

The optimal effort level, $e^* = e(\Omega^I, a)$, satisfies:

$$\left(\frac{p\partial q(e,a)}{\partial e} - \frac{\partial c(e,\Omega^I)}{\partial e}\right)|_{e=e^*} = 0 \tag{2}$$

The first term of first-order condition (2) captures the benefit from marginally exerting more effort while the second term depicts the marginal cost. The total derivative of (2), with respect to the cost- and output-shifters Ω^I and a, yields:

$$\left(\left(p\frac{\partial^2 q(e,a)}{\partial e^2} - \frac{\partial^2 c(e,\Omega^I)}{\partial e^2}\right)\begin{pmatrix} \frac{\partial e(\Omega^I,a)}{\partial \Omega^I} \\ \frac{\partial e(\Omega^I,a)}{\partial a} \end{pmatrix}' + \begin{pmatrix} -\frac{\partial^2 c(e,\Omega^I)}{\partial e \partial \Omega^I} \\ p\frac{\partial^2 q(e,a)}{\partial e \partial a} \end{pmatrix}'\right)\begin{pmatrix} d\Omega^I \\ da \end{pmatrix} = 0$$

Consider an increase in pollution $d\Omega^I > 0$ (and fix the worker, $da = 0$). The worker's optimal response to this shift in the environment is to reduce effort, and the magnitude

[17] Where ξ is not observed, specify the properties of total output, $\tilde{q}(e,a)$, similarly.

[18] Where relevant, we adopt the notation in GZN. For comparison, GZN specify ambient pollution as impacting the output function, rather than the cost of effort, and they do not model ability. One can extend our model to incorporate the probability of job retention as increasing in the level of output (as GZN do). In our setting most workers have worked for the firm over many years.

of the effort reduction depends on the magnitude of the shift in marginal cost, $\frac{\partial^2 c(\cdot)}{\partial e \partial \Omega^I} \geq 0$:

$$\frac{\partial e(\Omega^I, a)}{\partial \Omega^I} = \frac{\partial^2 c(e, \Omega^I)/\partial e \partial \Omega^I}{p\partial^2 q(e,a)/\partial e^2 - \partial^2 c(e, \Omega^I)/\partial e^2} \leq 0, \tag{3}$$

noting that the denominator (namely, the rate at which the difference between marginal product and marginal cost changes in effort) is negative. This is illustrated in Figure 5, panel (a). (As drawn, both production and cost functions can be reasonably approximated by quadratic functions in the neighborhood of e^*, such that the denominator of (3) is approximately constant.) Thus, a worker who is more sensitive to pollution, i.e., for whom the shift in marginal cost $\frac{\partial^2 c(\cdot)}{\partial e \partial \Omega^I}$ is larger, will reduce effort to a greater extent than a less sensitive worker.[19] This heterogeneous sensitivity to pollution is likely to depend on the level of pollution. For example, starting from low levels of pollution, a worker who suffers from asthma might incur a larger shift in marginal cost—and thus reduce effort more—than a non-asthmatic worker, who is hardly affected; on the other hand, at higher levels of pollution, the *change* in effort from additional pollution might be larger for the non-asthmatic worker, who is newly affected. Similarly, in a population of workers the effect of pollution on effort and output will be non-linear to the extent that $\frac{\partial^2 c(\cdot)}{\partial e \partial \Omega^I}$ varies over Ω^I (within worker).

Now compare two workers, H and L, with different levels of ability, $a_H > a_L$. It is trivial to show that the effort choice is weakly increasing in ability, $\frac{\partial e(\Omega^I, a)}{\partial a} \geq 0$, and this relationship is strict if marginal output strictly increases in ability, $\frac{\partial^2 q(\cdot)}{\partial e \partial a} > 0$. As illustrated in Figure 5, panel (b), the equilibrium marginal cost and revenue product are higher for the higher ability worker. As drawn, to the extent that variation in pollution shifts the marginal cost of either type of worker similarly (and marginal product schedules are approximately parallel across different ability levels), a change in pollution $d\Omega^I$ may lead to an optimal effort response of similar magnitude across the workers of different ability, i.e., $de^*_H \approx de^*_L$. The assumed properties of the production function then imply that the output response to pollution for the higher ability worker will exceed that of her lower ability counterpart. This can be seen in the panel by comparing the

[19]We do not index the cost function by i, to denote worker, to save on notation.

areas of the shaded trapezoids (of similar base). This discussion highlights that worker ability is another potential source of heterogeneity in the individual response of output to pollution, as is worker sensitivity in the preceding paragraph.

We can further model the worker's choice of attending work, with the reservation utility ϕ being a possible function of *outdoor* air pollution Ω^O and family composition F. For example, poor air quality may affect the health of the worker or of a family member who demands home care from the worker (e.g., Diette et al., 2000), and this demand for home care is likely to be pronounced for workers with young children. Alternatively, good air quality may raise the value of outdoor leisure relative to work. As discussed for the microenvironment we study, fine particles in outdoor air are likely to largely penetrate the workplace, in which case Ω^I and Ω^O are highly correlated.

In addition, we posit that meteorological conditions Λ, such as temperature or precipitation, may drive selection into work attendance, since meteorology may directly impact own or family health, the relative value of leisure, or the cost of commuting to work. Importantly, we posit that in our setting meteorology does not affect worker productivity directly (i.e., conditional on attendance), given that the workplace is temperature controlled and sheltered from rain and wind. This assumption provides an exclusion restriction to identify an overall selection-plus-productivity effect of pollution. The exclusion of meteorology from the production and cost functions above—and thus from the output equation we specify below—in principle allows us to better control for an otherwise potentially important confounder. Further, given that meteorology shifts fine particle pollution (per Section 2) but does not affect output directly, we can use meteorology to instrument for pollution in the output equation, to control for potential measurement error in PM2.5 or omitted determinants of output.

The worker's problem is then:

$$\max \left\{ \phi(\Omega^O, \Lambda, F), \max_e pq(e, a) - c(e, \Omega^I) \right\} \tag{4}$$

with the worker choosing between an unplanned absence or non-attendance (i.e., excluding predetermined leave days, such as consecutive days of vacation) or attending work as planned, and, conditional on attending work, optimizing over the effort level. We note that one feature of this labor market, that the work shift is fixed at eight hours and at a predetermined start time for the worker's team, implies that we need not model this additional margin of labor supply.

3.2 Empirical model

In our setting, the empirical counterpart to the framework above can be written:

$$q_{ijt} = \alpha_0 + \gamma\Omega_{jt}^O + X_{ijt}'\alpha_1 + \alpha_t + \alpha_j + \alpha_i + \epsilon_{ijt}, \tag{5}$$

$$d_{ijt}^* = \beta_0 + \delta\Omega_{jt}^O + \tilde{X}_{ijt}'\beta_1 + \beta_t + \beta_j + \beta_i + \zeta_{ijt}; \quad d_{ijt} = 1[d_{ijt}^* > 0] \tag{6}$$

where 1[.] is an indicator function, which equals one if its argument is true, and zero otherwise. In terms of subscripts, q_{ijt} is worker i's defect-free output quantity (alternatively, total quantity $\tilde{q}$) observed during the 8-hour shift j on date t, which is only observable if $d_{ijt} = 1$. In the output equation, Ω_{jt}^O is the mean PM2.5 concentration contemporaneously recorded outdoors during the 8-hour work shift, assumed to correlate tightly with unobserved indoor PM2.5 concentrations, $\Omega_{jt}^I = \iota\Omega_{jt}^O + \omega_{jt}$, where ι captures the "indoor-outdoor ratio" and ω is classical measurement error. Other controls include: (i) a vector of worker characteristics, X_{ijt}, to capture observed heterogeneity, such as the worker's schooling and tenure at the firm; (ii) year-month (i.e., 15 months in sample) and day-of-week fixed effects, α_t, to capture trends and seasonality over the year and within the week; (iii) shift fixed effects, α_j, to account for time of day; and (iv) individual worker fixed effects, α_i, to account for unobserved worker productivity. The random error term is denoted ϵ_{ijt}. Parameters α_1 cannot be identified if we control for worker fixed effects in the regression, but one can subsequently regress fitted worker fixed effects on observed individual characteristics to obtain mean effects. Other variations are possible such as allowing time-of-day fixed effects to vary by season (e.g., quarter), α_{jt}, or interacting them with worker fixed effects, α_{ij}. Another possibility is that the impact of fine particle pollution on labor productivity might be nonlinear. To address this possibility, we can include either a third-order polynomial function or a spline function of PM2.5 concentrations in the regression.

The selection equation contains all the control variables of the output equation (5) in addition to other variables that affect d^*_{ijt} but not q_{ijt}, such as meteorological conditions. As discussed, meteorology may affect a worker's reservation utility but, in view of the temperature controlled work environment, it should not affect workers' productivity. β_i is an unobserved time invariant worker specific effect that affects a worker's attendance decision, and ζ_{ijt} is a random error term.

In a regression of (5), but without the worker fixed effects α_i, a sufficient condition to obtaining a consistent estimate of γ, the effect of pollution on output, using pooled OLS is (e.g., Dustmann and Rochina-Barrachina, 2007):

$$E[\alpha_i + \epsilon_{ijt} | \Omega^O_{jt}, X_{ijt}, \alpha_t, \alpha_j, d_{ijt} = 1]) = 0. \tag{7}$$

The OLS estimator will be biased, for example, if workers with higher α_i (higher ability) are more likely to take leave when pollution is high, or if workers with lower α_i are more likely to take leave when pollution is low. Or say that pollution impacts the health of a particular group of workers, such as asthmatics, who are then more likely to call in sick during polluted days; and these workers happen to be less (resp., more) productive than other workers, then the estimate for γ will biased downward (resp., upward).

A model with worker fixed effects can remove the bias caused by selection on α_i. To obtain a consistent estimate of γ, a sufficient condition is:

$$E[\epsilon_{ijt} - \epsilon_{ijs} | \Omega^O_{jt}, X_{ijt}, \alpha_t, \Omega^O_{js}, X_{ijs}, \alpha_s, d_{ijt} = d_{ijs} = 1]) = 0, \quad s \neq t, \tag{8}$$

for time periods $s \neq t$. Now, this condition will be violated if there is selection on ϵ_{ijt}, for example, if a worker who experiences a positive productivity shock, i.e., higher ϵ_{ijt}, is more likely to take leave when pollution is high compared to when it is low (and, of course, observes her productivity shock prior to making the work versus non-work choice). To account for this possibility, Wooldridge (1995) recommends deriving an expression for (7), and adding it as an additional regressor to output equation (5).

4 The impact of PM2.5 pollution on non-attendance

Whereas plant holidays are predetermined, workers do occasionally choose to not attend work. Among workers who are in the sample throughout all 15 months, the mean number of non-attendance days is 15 per worker, i.e., an average of one day per month (on top of plant holidays, averaging just under 4 days per month). Defining each set of adjacent days of non-attendance by a worker as a single non-attendance event, or spell, the mean number of non-attendance spells is only 10 per worker during the 15 months. This labor market is thus characterized by limited non-attendance. In any case, we seek to understand its determinants, including the extent to which PM2.5 pollution may drive non-work decisions, and thus possibly change the composition of the workforce present each day.

There are a total of 866 worker non-attendance spells in our sample. Of these, 577 lasted for only one day, 161 for two days, and 65 for three days. We exclude the 63 non-attendance spells that are longer than three days from our analysis, as factors that lead a worker to take leave for a relatively long period likely differ from those that keep her from working for only a few days. In particular, longer non-attendance events are likely to be planned leaves, which are unlikely to be triggered by day-to-day shifts in pollution (or to correlated with pollution via weather).[20] More broadly, when we examine the relationship between particle pollution and worker output, we can control for the work choice probability on each date we observe a worker choosing to work.

Table 3 reports estimates for variations of our selection equation, either a linear probability model (all columns but (3)) or a probit model (column (3)). Across workers, we have 803 date-shifts of non-attendance (the first day of a non-attendance spell of at most three days) plus 27,776 date-shifts of work attendance, thus 28,579 observations. The dependent variable is 1 if the worker chose to start a non-attendance spell and 0 if she chose to attend work on the given date; the mean value is thus 803/28579=0.028. We consider the mean PM2.5 mass concentration over the 24 hours that immediately precede

[20] If we observed which non-attendance events had been agreed in advance, we would not seek to explain these.

the start of a worker's shift, since the decision to start an (unplanned) non-attendance spell is likely to occur at this time. We allow for a non-linear relationship by specifying a linear spline function of PM2.5 concentrations, with three knots set at the first, second and third quartiles of the PM2.5 distribution, respectively, 62, 94 and 149 μg/m^3.[21]

Column (1) suggests that the impact of (past 24-hour mean) PM2.5 on non-attendance is highly non-linear in our sample, associated with falling non-attendance as pollution rises from low levels, i.e., over the first quartile up to 62 μg/m^3, while associated with rising non-attendance when pollution rises above some already severe threshold, over the fourth quartile beyond 149 μg/m^3 (to a maximum of 687 μg/m^3 in the sample of 24-hour means). For example, estimates suggest that as PM2.5 increases from (close to) 0 to 62 μg/m^3, a worker is 4 percentage points (0.62×0.063) less likely to choose non-work over work. Of note, had we controlled for PM2.5 levels linearly, the point estimate would be small and statistically insignificant. Figure A.9 in the Appendix plots daily rates of non-attendance against mean PM2.5 concentrations.

In column (2), we control for year-month, day-of-week and time-of-day (0 am, 8 am or 4 pm shift), as both non-attendance and PM2.5 may follow seasonal/cyclical patterns. Day-of-week includes indicators for any days on the public holiday calendar that the department was working (see Appendix Table A.1).[22] We allow time-of-day effects to differ across the four quarters of the year. We control for days that immediately precede or come after plant holidays and find that non-attendance is significantly higher one day before, and one day after, plant holidays, by 6 percentage points in both cases (estimates are not reported). To save space, column (2) already includes contemporaneous meteorological controls, namely the same ones we included in the PM2.5 regressions of Table 2 (columns (3) or (4)). The modeling assumption, as per Section 3, is that meteorological conditions Λ, by changing the attractiveness of outdoor activities or of

[21] We base the knots, or kink points, on the quartiles of the distribution of 8-hour means, to maintain consistency throughout. Using quartiles of the distribution of means over the preceding 24 hours would yield similar knots, namely, 66, 96 and 146 μg/m^3. We sometimes loosely refer to variation from the sample minimum to the 25th percentile as the first (or bottom) quartile, the 25th percentile to the median as the second quartile, and so on.

[22] As one might expect, we obtain that non-attendance on such dates is significantly higher (most often the case) or insignificantly different from zero. We do not report for brevity.

running outside errands, shift reservation utility ϕ directly (in addition to indirectly, as Λ impacts pollution Ω^O). Further, meteorology, like pollution, may impact health, shifting the value of work versus non-work. Adding all of these controls has little impact on the estimated effect of PM2.5 on non-attendance. We also include a vector of observed worker characteristics. We find that mothers with young children are 2 percentage points more likely to not attend work, but this is not the case for fathers. Married workers (excluding divorcees and widows) display a higher probability of non-attendance, of 1 percentage point, consistent with a higher reservation utility, whereas local workers are 2 percentage points less likely to not attend work, perhaps because they enjoy the support of local family to perform household chores. These findings are consistent with family composition F shifting reservation utility ϕ, as we posited in Section 3.

The marginal effects for a probit model, in column (3), are almost identical to the estimates from the linear probability model of column (2), with common regressors. In column (4), we replace the worker characteristics of column (2) by a full set of worker fixed effects, to allow for unobserved heterogeneity. This significantly raises explanatory power but again has little impact on PM2.5 estimates. Finally, column (5) tests robustness in two directions. First, we specify additional flexible meteorological controls (Auffhammer and Kellogg, 2011), namely: (i) cubic polynomials in the maximum and minimum temperature, humidity and wind speed, and the maximum precipitation rate, in the contemporaneous 8-hour shift, allowing these cubic polynomials to vary by time-of-day; (ii) pairwise interactions for all linear terms in (i); and (iii) linear controls for the maximum and minimum temperature, humidity and wind speed, and the maximum precipitation rate, observed in the 24-hour period that precedes the start of the shift. The second direction in which we test robustness is to control for co-pollutants. Specifically, column (5) also includes linear splines of SO_2 and CO concentrations, each with three knots set at the first, second and third quartiles of the respective distribution.[23]

The estimated non-linear relationship between PM2.5 and non-attendance thus survives the inclusion of controls for seasonality, worker, weather, and co-pollutant con-

[23] The quartiles of the distributions of 8-hour means—which we use throughout—are 37, 56 and 89 μg/m^3 for SO_2 and 1.1, 1.6 and 2.5 μg/m^3 for CO.

centrations. One possible interpretation of this relationship is as follows. Starting at relatively low levels, a moderate increase in ambient PM2.5 makes outdoor activities less attractive, so non-attendance declines with PM2.5 concentration. On the other hand, once PM2.5 exceeds a high threshold, further increases might affect a worker's health or the health of her family, preventing her from working. If PM2.5 drives non-attendance through its impact on the attractiveness of substitute activities when concentration levels are relatively low, this impact is likely to vary across seasons as well as across shifts. The value of activities that compete with work is likely to be higher under mild weather and during daytime hours. For this reason, Table 4 reports estimates of OLS regressions on separate samples: (i) by season—"cold," from October to March, versus "warm," from April to September—in columns (1) and (2); and (ii) by shift—again, starting at 0 am, 8 am or 4 pm—in columns (3) to (5).

Reassuringly, estimates suggest that rising PM2.5 is more precisely associated with lower non-attendance in the warm months, and this effect is more pronounced in the bottom quartile, up to a concentration of 62 μg/m^3, conditions in which outdoor activity is presumably (still) attractive. During the cold season, or when PM2.5 levels cross into the second quartile, the coefficient on PM2.5 is imprecisely estimated (and may reverse sign). We also obtain intuitive estimates on the time-of-day subsample regressions, for example, rising PM2.5 concentrations in the bottom quartile being more significant, both statistically and economically, for the 4 pm shift, when the opportunity cost of work might be higher (e.g., joint family consumption), compared with the 0 am shift. Column (6) indicates that the negative association between non-attendance and PM2.5, starting at low levels, is present among mothers with young children, as one might expect from our possible interpretation.

Another subsample analysis we conduct is for workers with (mean) productivity above the median of the worker productivity distribution separately from workers below this median, i.e., high productivity versus low productivity workers (recall panel (a) of Figure 1). We find that, starting at relatively low levels, rising PM2.5 concentrations reduce non-attendance more among less productive (lower ability) workers than among their more productive (able) counterparts. We take this result to be consistent with the simple model we developed in Section 3. While the model suggests that the output response to increased pollution is larger among workers of higher ability, the gross surplus from working remains higher for these workers compared to workers of lower ability—such differing utility can be seen in the area between marginal product and marginal cost schedules in Figure 5, panel (b). Thus, if the fall in the value of leisure $\phi(\Omega^O, \Lambda, F)$ is similar across workers of differing ability, the decline in non-attendance with rising pollution may be more pronounced among lower ability workers compared to higher ability ones, as indeed we find.

5 The impact of PM2.5 pollution on output

Table 5 reports OLS estimates for variations of the output equation (5), estimated on the full sample of 27,776 worker-date pairs between April 2013 and June 2014 for which a worker attended work.[24] The dependent variable is a worker's total output, in meters, over an 8-hour work shift; recall the mean is 509 m/shift. Variable definitions are as in the preceding section, unless noted otherwise. Column (1) specifies a linear spline of contemporaneously observed 8-hour mean PM2.5 levels. In addition to the knots specified earlier, corresponding to the first, second and third quartiles of the distribution of 8-hour means, we specify an additional knot corresponding to the 90th percentile, namely 230 μg/m^3, to allow the estimated impact to vary within the fourth quartile, where PM2.5 ranges from 149 to a maximum of 773 μg/m^3.

Column (2) adds year-month, day-of-week (again, including dummies for any work days that were on the national holiday calendar), and (quarter-specific) time-of-day fixed effects, to account for seasonal/cyclical patterns. We also control for: days that immediately precede or come after plant holidays; and a worker's last day of work prior to initiating a non-attendance spell, as well as first day of work after returning from one. Column (3) adds worker characteristics, as in Table 3, column (2). As the workspace is

[24]Some 8-hour pollution means are missing, so in practice there are slightly fewer observations than 27,776. We show below that our findings are robust to taking, as the dependent variable, output net of defects, $q = \tilde{q}\xi$, for the subsample of 20,930 worker-dates for which we observe defective output.

temperature controlled and sheltered from precipitation and wind, meteorology should not shift output other than through a possible selection effect or through (as an instrument for) pollution—we consider both channels below and do not, thus, include meteorology, providing an exclusion restriction.[25] Relative to column (1), estimated standard errors in column (2) are lower. Replacing worker characteristics by a full set of worker fixed effects, in column (4), improves precision further, and the R^2 almost triples, from 0.10 to 0.28, yet the estimated non-linear relationship between PM2.5 and worker output changes only slightly. Similarly, estimated PM2.5 effects hardly change on adding, in column (5), linear spline functions of contemporaneous SO_2 and CO concentrations, and, in column (6), a selection correction based on the inverse Mills ratio from probit model (3) in Table 3. We further comment on these specifications below.

Estimates in column (5) suggest that, at the lower fine-particle levels in ambient air, a 10 μg/m^3 increase in (contemporaneous 8-hour mean) PM2.5 concentration reduces a worker's output by 4.3 meters of fabric, equivalent to about 0.9% of mean output in the sample. This is a large effect. Estimated marginal effects are similar or somewhat lower over the interquartile range of the PM2.5 distribution, up to 149 μg/m^3, and then halve in magnitude (while remaining statistically significant) between the 75th and 90th percentiles, namely -2.0m of fabric per 10 μg/m^3 increase. Beyond the 90th percentile of the PM2.5 distribution, namely 230 μg/m^3, the estimated marginal effect is a low and marginally significant -0.5m per 10 μg/m^3 increase. Figure 4 illustrates the fitted spline, based on column (5) estimates, as well as the cubic polynomial discussed below.

We can integrate over the first 100 μg/m^3 increase in ambient PM2.5 concentrations, starting at the sample minimum of 10 μg/m^3, to obtain a prediction for the output shortfall over this range, namely 42m of fabric, equivalent to 8.5% of the sample mean, with estimated standard error (s.e.) of 6m.[26] The subsequent 100 μg/m^3 increase in PM2.5 levels, starting at 110 μg/m^3, leads to further output loss of 29m (s.e. 3m). In

[25] We check that the estimates on PM2.5 reported in Table 5 are robust to including either set of meteorological covariates used in the non-attendance analysis, e.g., in columns (4) and (5) of Table 3. Appendix Figure A.13 plots output against temperature—as well as output against PM2.5 by way of contrast—once calendar-month, day-of-week and time-of-day effects have been partialled out.

[26] For clarity, the point estimate is evaluated as -.43×(62-10)-.41×(94-62)-.43×(110-94). The first 100 μg/m^3 increase corresponds to a shift from the sample minimum into the third quartile.

total, taking the sample minimum of 10 μg/m^3 as the point of departure and raising PM2.5 concentrations by 200 μg/m^3 leads to an output shortfall of 71m (s.e. 7m), equivalent to 14% of mean output (or 13% of mean output when PM2.5 concentrations fall below the 25th percentile).

Returning to the inclusion of individual worker fixed effects in column (4), compared to observed worker characteristics included in column (3), the increase in explanatory power is evidence of the role of time-invariant unobservable productivity differences across workers. However, the finding that PM2.5's estimated impact on output changes little when adding this rich set of controls suggests that selection effects driven by unobserved heterogeneity likely plays a small role in our setting. We note that on adding time fixed effects in column (2), compared to column (1), the point estimate for the bottom quartile of the PM2.5 distribution grows more negative. This is consistent with changes in workforce composition being correlated with pollution, e.g., the workers who joined the department in the fall, just before the winter when PM2.5 concentrations tend to be higher, being more productive than those who left in the fall, as we report in the Appendix. Failing to correct for such a workforce composition effect, in this case through time fixed effects, would seem to understate the effect of pollution on output.

Columns (4) and (5) thus include a full set of individual fixed effects (respectively, without and with controls for co-pollutants). We go further and, in column (6), we control for self-selection following Wooldridge (1995). The variables that are included in the non-attendance equation but not in the output equation are meteorology variables and their interactions with quarter. The rationale, as discussed above, is that given the temperature controlled work environment (and rain and wind), once workers show up to work their productivity should not depend on contemporaneous meteorology. As we might expect following the inclusion of worker fixed effects in column (4), the coefficient on the inverse Mills ratio in column (6) is not statistically significant even at the 10% level and the estimated impacts are not sensitive to whether we include the inverse Mills ratio as an additional control.

Table 6 performs a series of subsample analyses, based on the "baseline" specification in column (5), Table 5. Since the chemical composition and thus toxicity of PM2.5 (itself a multi-pollutant) may vary across season (Bell, 2012), we estimate the model separately on two subsamples, "cold" months (again, October to March) and "warm" (the remainder). Point estimates (columns (1) and (2)) suggest a larger adverse effect of PM2.5 in the cold months, particularly up to the median of the PM2.5 distribution. Columns (3) to (5) report estimates when the model is implemented separately by time of day (shift): while the estimated slopes over the different quartiles vary across time-of-day subsample, the combined output effect over a 10 to 149 μg/m^3 increase in PM2.5 is similar, namely -59m, -56m and -56m for the 0 am, 8 am and 4 pm shifts, respectively. Our model predicted that the negative impact of PM2.5 on output is likely to be stronger for higher ability workers compared to those of less ability. This is supported by the empirical results in columns (6) and (7). Except for days when PM2.5 is at the bottom quartile, an increase in PM2.5 tends to have a more negative impact on workers with mean productivity above the sample mean.

Sensitivity analysis. We provide some robustness checks in Table 7.[27] First, in column (1), we take as dependent variable the quantity of output net of defective output for the subsample of worker-dates for which we observe defective output (recall defective records were not available for Team 3). Second, in column (2), we drop observations for which a worker's output during an 8-hour shift might appear either too high or too low. Specifically, we drop observations with output above the 95th or below the 5th percentiles, 660 m and 365 m respectively.[28] Point estimates, while still highly significant both economically and statistically, shrink somewhat relative to baseline estimates.

[27]Among other robustness tests that we do not report here, we add a spline function of mean PM2.5 levels over the 24 hours that immediately precede each work shift, on top of contemporaneous PM2.5, to the baseline specification. All contemporaneous and lagged point estimates are negative and most are highly statistically significant. Specifically, for the bottom quartile the estimated contemporaneous effect falls slightly from, -43 to -37 (s.e. 12), and the lagged effect is -9 (s.e. 13). For the second quartile the contemporaneous effect falls more, from -41 to -25 (s.e. 11) and the lagged effect is -46 (s.e. 11). One may speculate from such results whether higher levels of PM2.5 might have a more lasting effect compared to lower levels. PM2.5 estimates are plotted in Appendix Figure A.14, panel (b).

[28]See Appendix Figure A.14, panel (a). Our working hypothesis for a few observations with output above the nominal capacity of a workstation, as discussed in the Appendix, is that the worker obtained approval to continue working into the subsequent shift, noting that such extended hours are presumably unusual and are not recorded. Our working hypothesis for some low output rates is that the worker felt sick during the shift and went home, or was asked to perform an alternative task.

Column (3) replaces the spline function of our baseline specification by a cubic function in PM2.5 concentrations. For comparability, the fitted cubic is plotted in Figure 4 alongside the estimated spline from column (5), Table 5. Finally, in column (4), Table 7 the cubic specification is estimated by two-stage least squares, using meteorology as an instrument for PM2.5 pollution. Rather than rely on the entire variation in PM2.5 concentrations, the first-stage project PM2.5 concentrations on the set of meteorology by quarter controls (and included exogenous variables). The identifying assumption is that meteorology affects worker output in the climate controlled work environment only through its effect on PM2.5 pollution (Table 2). The objective is to gauge the extent to which measurement error in pollution or unobserved determinants of output that correlate with pollution may be affecting our estimates. Estimates by 2SLS are similar to those by OLS, as can be seen comparing columns (4) against (3), as well as in Figure 4.

6 Implications and concluding remarks

We take our estimated output model, namely the estimates reported in column (5) of Table 5, and conduct a back-of-the-envelope calculation, with a view to drawing implications from these results. We focus on output effects, rather than the combined attendance-output model, as we are less confident that attendance choices at our specific labor institution, i.e., the department of 98 individual workers and their characteristics, generalize to other labor settings; also, we do not attempt here to value activities that may substitute for work as air pollution varies.[29]

We predict each of the 98 individual workers' output over the course of one year (to capture all seasons) were PM2.5 pollution to follow the time path observed in each of 190 major cities in China today. (See the Appendix for data sources and calculation details.) Label this baseline scenario "actual pollution." We then truncate values in these actual 190 city-specific hourly PM2.5 series at 25 μg/m^3, and again predict output over the 12 months under this "counterfactual pollution" scenario. For example, if the actual

[29]Were we to predict work attendance across China, we would also need meteorology data (besides air pollution, explained below).

PM2.5 level in Guangzhou at 2 pm on October 1, 2013 was 49 μg/m^3, we take it to be $min(25, 49) = 25$ μg/m^3 in the counterfactual. By comparing output predicted under the two scenarios, and using GDP to aggregate across space and seasons, we obtain a rough measure of the labor productivity gains that China could experience were PM2.5 pollution to be abated to a concentration no greater than 25 μg/m^3 at any point in time and in space—without, of course, considering abatement costs (or other benefits such as health), and assuming that other workplaces in China were similar, in terms of dose-response, to the one we study.[30] Table 8 presents estimates of output gains by region of China and by season. Across the nation and over the year, the predicted output gain amounts to 3.8%. The largest proportionate output effect, +7.2%, is predicted for North China (which hosts both the sampled city and Beijing) in the coldest quarter of the year; the lowest predicted effects correspond to the mild months of July to September. Crudely multiplying the aggregate output effect by an annual GDP for China of US$ 16.1 trillion (2013, purchasing-power parity, per the IMF) as well as a labor share of GDP of 40% (Karabarbounis and Neiman, 2014), yields US$ 240 billion.

We end by comparing the marginal effect found in Chang et al. (2014)—a 6% drop in indoor worker productivity for every 10 μg/m^3 increase in outdoor PM2.5 concentrations—to what we find if we restrict our sample to the same range of fine-particle pollution, namely up to 60 μg/m^3. We re-estimate our baseline specification on the sub-sample of observations in which 8-hour mean PM2.5 concentrations lie below this threshold. This drops 78% of our sample, but we can still count on 6178 observations. (We also restrict the effect of PM2.5 to be linear, given the reduced range of variation.) We find that every 10 μg/m^3 increase in PM2.5 levels is associated with a 4.6 m output loss, significant at the 1% level and equivalent to a 0.9% drop relative to mean output in this sub-sample. That we find highly significant effects of PM2.5 pollution on labor productivity over a very wide range, yet only one-sixth the magnitude found in Chang et al. (2014) for a comparatively narrow range, suggests that examining further labor

[30] For example, the exercise assumes the following to be "similar": office and farm labor, the nature of manufacturing tasks, outdoor versus indoor exposure to ambient air, the distribution of worker characteristics such as age, health and unobserved heterogeneity, the composition of PM2.5 (a multi-pollutant), and so on.

settings is an important direction for future research.

References

Auffhammer, M. and Kellogg, R. (2011). Clearing the air? The effects of gasoline content regulation on air quality. *American Economic Review*, 101(6):2687–2722.

Beirle, S., Platt, U., Wenig, M., and Wagner, T. (2003). Weekly cycle of NO_2 by GOME measurements: A signature of anthropogenic sources. *Atmospheric Chemistry and Physics*, 3(6):2225–2232.

Bell, M. (2012). Assessment of the health impacts of particulate matter characteristics. Technical Report 161, Health Effects Institute.

Chang, T., Graff Zivin, J. S., Gross, T., and Neidell, M. J. (2014). Particulate pollution and the productivity of pear packers. Working Paper 19944, National Bureau of Economic Research.

Chen, Y., Ebenstein, A., Greenstone, M., and Li, H. (2013). Evidence on the impact of sustained exposure to air pollution on life expectancy from China's Huai River policy. *Proceedings of the National Academy of Sciences*, 110(32):12936–12941.

Christiani, D. C., Ye, T.-t., Zhang, S., Wegman, D. H., Eisen, E. A., Ryan, L. A., Olenchock, S. A., Pothier, L., and Dai, H.-l. (1999). Cotton dust and endotoxin exposure and long-term decline in lung function: Results of a longitudinal study. *American Journal of Industrial Medicine*, 35(4):321–331.

Crocker, T. D. and Horst, R. L., J. (1981). Hours of work, labor productivity, and environmental conditions: A case study. *The Review of Economics and Statistics*, 63(3):361–368.

Cyrys, J., Pitz, M., Bischof, W., Wichmann, H.-E., and Heinrich, J. (2004). Relationship between indoor and outdoor levels of fine particle mass, particle number concentrations and black smoke under different ventilation conditions. *Journal of Exposure Science and Environmental Epidemiology*, 14(4):275–283.

Davis, L. W. (2008). The effect of driving restrictions on air quality in Mexico City. *Journal of Political Economy*, 116(1):38–81.

Diette, G. B., Markson, L., Skinner, E. A., Nguyen, T. T., Algatt-Bergstrom, P., and Wu, A. W. (2000). Nocturnal asthma in children affects school attendance, school performance, and parents' work attendance. *Archives of Pediatrics & Adolescent Medicine*, 154(9):923–928.

Dustmann, C. and Rochina-Barrachina, M. E. (2007). Selection correction in panel data models: An application to the estimation of females' wage equations. *The Econometrics Journal*, 10(2):263–293.

Ebenstein, A., Fan, M., Greenstone, M., He, G., and Zhou, M. (2015). Growth, pollution, and life expectancy: China from 1991-2012. *American Economic Review Papers & Proceedings*, (forthcoming).

Fann, N., Lamson, A. D., Anenberg, S. C., Wesson, K., Risley, D., and Hubbell, B. J. (2012). Estimating the national public health burden associated with exposure to ambient PM2.5 and ozone. *Risk Analysis*, 32(1):81–95.

Graff Zivin, J. and Neidell, M. (2012). The impact of pollution on worker productivity. *American Economic Review*, 102(7):3652–73.

Graff Zivin, J. and Neidell, M. (2014). Temperature and the allocation of time: Implications for climate change. *Journal of Labor Economics*, 321:1–26.

Gupta, A. and Cheong, K. (2007). Physical characterization of particulate matter and ambient meteorological parameters at different indoor–outdoor locations in Singapore. *Building and Environment*, 42(1):237–245.

Hanna, R. and Oliva, P. (2014). The effect of pollution on labor supply: Evidence from a natural experiment in Mexico City. *Journal of Public Economics*, (forthcoming).

Karabarbounis, L. and Neiman, B. (2014). The global decline of the labor share. *Quarterly Journal of Economics*, 129(1):61–103.

Madronich, S. (2014). Atmospheric chemistry: Ethanol and ozone. *Nature Geoscience*, 7(6):395–397.

Morawska, L., He, C., Hitchins, J., Gilbert, D., and Parappukkaran, S. (2001). The relationship between indoor and outdoor airborne particles in the residential environment. *Atmospheric Environment*, 35(20):3463–3473.

Ostro, B. D. (1983). The effects of air pollution on work loss and morbidity. *Journal of Environmental Economics and Management*, 10(4):371–382.

Salvo, A. and Geiger, F. M. (2014). Reduction in local ozone levels in urban Sao Paulo due to a shift from ethanol to gasoline use. *Nature Geoscience*, 7(6):450–458.

Salvo, A., Geiger, F. M., Brito, J., and Artaxo, P. (2015). Fine and ultra-fine particles and shifting gasoline and ethanol use in urban Sao Paulo. Mimeo.

Wooldridge, J. M. (1995). Selection corrections for panel data models under conditional mean independence assumptions. *Journal of Econometrics*, 68(1):115–132.

World Health Organization (1999). Hazard prevention and control in the work environment: Airborne dust. Technical Report WHO/SDE/OEH/99.14, World Health Organization.

World Health Organization (2006). WHO air quality guidelines for particulate matter, ozone, nitrogen dioxide and sulfur dioxide: Global update 2005. Technical Report WHO/SDE/PHE/OEH/06.02, World Health Organization.

Table 1: Descriptive statistics

Variables	(1) N	(2) Mean	(3) Std.Dev.	(4) Minimum	(5) Maximum
Total output (meters per worker per shift)	27,776	508.98	96.02	128.30	1,137.00
Defective output (meters per worker per shift)	20,930	3.79	7.84	0.00	82.00
Worker chose not to work (non-attendance, yes=1)	29,200	0.05	0.22	0.00	1.00
First day of non-attendance event (yes=1)	28,642	0.03	0.17	0.00	1.00
Age of worker (years, in 2014)	98	40.79	9.31	17.00	58.00
Worker is male (yes = 1)	98	0.18	0.39	0.00	1.00
Worker's schooling (years)	98	9.73	1.62	6.00	12.00
Worker is local (yes = 1)	98	0.81	0.40	0.00	1.00
Worker's number of children	98	0.97	0.51	0.00	2.00
Age of worker's youngest child (years, in 2014)	84	16.23	7.40	2.00	34.00
Worker has child aged ≤ 12 (yes=1)	98	0.28	0.45	0.00	1.00
Worker's tenure at firm (years, in 2014)	98	18.52	12.07	1.00	35.00
Worker's contract duration (years)	98	7.07	3.03	1.00	10.00
PM2.5, within-shift mean at closest monitor ($\mu g/m^3$)	1,201	123.65	98.62	12.67	773.13
PM2.5, within-shift mean in Beijing ($\mu g/m^3$)	1,204	95.02	82.27	4.63	609.88
CO, within-shift mean at closest monitor ($\mu g/m^3$)	1,198	2.15	1.66	0.24	13.82
SO_2, within-shift mean of mean across city's 4 monitors ($\mu g/m^3$)	1,200	71.48	47.75	6.54	320.67
O_3, within-shift mean of median across city's 4 monitors ($\mu g/m^3$)	1,187	52.10	44.71	1.93	216.13
NO_2, within-shift mean of median across city's 4 monitors ($\mu g/m^3$)	1,200	60.26	29.35	7.69	241.81
Atmospheric pressure, within-shift mean in city (mb)	1,206	1,014.83	9.79	994.63	1,036.75
Temperature, within-shift mean in city (^{0}C)	1,206	18.50	10.77	-9.63	41.75
Humidity, within-shift mean in city (%)	1,206	45.28	19.87	7.50	98.38
Wind speed, within-shift mean in city (mph)	1,206	6.38	2.92	0.63	20.00
Precipitation, within-shift mean in city (mm/hour)	1,206	0.04	0.27	0.00	6.88

Notes: Statistics for pollution and meteorological variables reported here restrict the date-shift sample to dates that were not plant holidays (rather than considering the entire 456-day period between April 1, 2013 and June 30, 2014). See the Appendix for data sources.

Table 2: Observed determinants of PM2.5 mass concentrations

	(1)	(2)	(3)	(4)	(5)
Maximum temperature (^{o}C, 8-h period)		2.835**			0.014*
		(1.185)			(0.008)
Minimum temperature (^{o}C, 8-h period)		1.756			0.019**
		(1.131)			(0.008)
Maximum humidity (%, 8-h period)		1.352***			0.011***
		(0.259)			(0.002)
Minimum humidity (%, 8-h period)		0.304			0.001
		(0.356)			(0.003)
Precipitation >0 but ≤0.5 mm/hour (8-h period)		-18.802***			-0.163***
		(6.615)			(0.055)
Precipitation >0.5 mm/hour (8-h period)		-21.315**			-0.195*
		(10.558)			(0.109)
Atmospheric pressure (mb, 8-h period)		2.134***			-0.002
		(0.675)			(0.004)
Wind speed > 5 but ≤10 mph (8-h period)		-6.977			-0.042
		(4.725)			(0.033)
Wind speed > 10 mph (8-h period)		-16.041**			-0.169***
		(7.106)			(0.059)
North wind (yes = 1, on date)		2.071			-0.016
		(13.454)			(0.077)
South wind (yes = 1, on date)		7.585			0.035
		(13.718)			(0.080)
North-South wind (yes = 1, on date)		13.620			0.072
		(15.073)			(0.087)
South-North wind (yes = 1, on date)		29.824**			0.207**
		(14.944)			(0.088)
Year-month fixed effects	Yes	Yes	Yes	Yes	Yes
Day-of-week fixed effects	Yes	Yes	Yes	Yes	Yes
Public holiday indicators	Yes	Yes	Yes	Yes	Yes
Time-of-day × Quarter	Yes	Yes	Yes	Yes	Yes
Meteorology × Quarter	-	-	Yes	Yes	-
Contemporaneous and lagged meteorology	-	-	-	Yes	-
Observations	1,363	1,363	1,363	1,362	1,363
R-squared	0.381	0.453	0.499	0.577	0.464
Mean value of dependent variable	121.464	121.464	121.464	121.442	4.559

Notes: An observation is a date-shift pair between April 1, 2013 and June 30, 2014. The dependent variable is PM2.5 mass concentration (μg/m^3, mean over the 8 hours of each shift) in columns (1) to (4), and its logarithm in column (5). Public holiday indicators separately control for New Year, Chinese New Year, Qingming Festival, International Workers' Day, Dragon Boat Festival, Mid-Autumn Day, National Day, as well as Saturday and Sundays which were set to be working days in the official holiday calendar. Time-of-day are indicators for the 0 am, 8 am and 4 pm shifts. Lagged meteorology comprise meteorological conditions in the preceding 8-hour period. Ordinary least squares estimates. Standard errors in parentheses, clustered by date. ***Significant(ly different from zero) at (the) 1% (level), **significant at 5%, *significant at 10%.

Table 3: The impact of PM2.5 pollution on non-attendance

	(1)	(2)	(3)	(4)	(5)
PM2.5 10-62 μg/m^3 (per 100)	-0.063*	-0.067***	-0.068***	-0.065***	-0.071***
	(0.032)	(0.020)	(0.017)	(0.021)	(0.024)
PM2.5 62-94 μg/m^3 (per 100)	0.030	0.023	0.024	0.024	0.028
	(0.020)	(0.018)	(0.016)	(0.018)	(0.018)
PM2.5 94-149 μg/m^3 (per 100)	-0.021**	-0.025***	-0.023***	-0.025***	-0.031***
	(0.009)	(0.009)	(0.008)	(0.009)	(0.009)
PM2.5 149-773 μg/m^3 (per 100)	0.004**	0.005**	0.005***	0.005**	-0.002
	(0.002)	(0.002)	(0.002)	(0.002)	(0.004)
Worker is male (yes = 1)		0.004	0.004		
		(0.003)	(0.003)		
Worker's schooling (years)		-0.017***	-0.014**		
		(0.007)	(0.006)		
Worker's tenure at firm (years, in 2014)		0.000	0.000		
		(0.001)	(0.001)		
Worker has child aged $\leq$ 12 (yes=1)		0.018***	0.018***		
		(0.003)	(0.002)		
Worker is male and has child aged $\leq$ 12 (yes = 1)		-0.021***	-0.020***		
		(0.005)	(0.006)		
Worker is currently married (yes = 1)		0.009***	0.011***		
		(0.002)	(0.003)		
Worker is local (yes = 1)		-0.017***	-0.016***		
		(0.003)	(0.003)		
Year-month fixed effects	No	Yes	Yes	Yes	Yes
Day-of-week fixed effects	No	Yes	Yes	Yes	Yes
Time-of-day $\times$ Quarter	No	Yes	Yes	Yes	Yes
Days before or after plant holiday	No	Yes	Yes	Yes	Yes
Worker fixed effects	No	No	No	Yes	Yes
Meteorology $\times$ Quarter	No	Yes	Yes	Yes	No
Flexible meteorology	No	No	No	No	Yes
Co-pollutant concentrations	No	No	No	No	Yes
Observations	28,579	28,579	28,549	28,579	28,579
Number of regressors	5	101	101	97	165
R-squared	0.001	0.019		0.038	0.042
Mean value of dependent variable	0.028	0.028	0.028	0.028	0.028

Notes: An observation is a worker-date pair in the sample of work and non-work choices, i.e., dates with non-zero worker output and first date of non-attendance events (for events not exceeding three days). The dependent variable is 1 if the worker chose to initiate a non-attendance spell and 0 if the worker chose to work. Worker's schooling and tenure are in logs. Day-of-week fixed effects include public holiday indicators (see notes to the previous table). Days before or after plant holiday are separate indicators for 1, 2 or 3 before or after a plant holiday. See the preceding table for meteorology controls and the text for more flexible meteorology controls. Co-pollutant concentrations are spline functions of SO_2 and CO concentrations, each with three knots set at the first, second and third quartiles of the respective distribution (as for PM2.5). All pollutant concentrations are means for the 24 hours immediately preceding a shift. Ordinary least squares estimates except in column (3), which reports marginal effects from a probit model evaluated at the sample means. Standard errors in parentheses, clustered by date. ***Significant at 1%, **significant at 5%, *significant at 10%.

Table 4: The impact of PM2.5 on non-attendance, by season, time of day and worker productivity

	Cold months (1)	Warm months (2)	0 am shift (3)	8 am shift (4)	4 pm shift (5)	Women with young child (6)	More productive (7)	Less productive (8)
PM2.5 10-62 μg/m^3 (per 100)	-0.070	-0.075***	-0.035	-0.042	-0.119***	-0.113***	-0.034	-0.108***
	(0.048)	(0.024)	(0.032)	(0.040)	(0.036)	(0.042)	(0.028)	(0.028)
PM2.5 62-94 μg/m^3 (per 100)	0.014	0.024	0.004	0.033	0.061**	0.061*	0.007	0.032
	(0.026)	(0.024)	(0.031)	(0.030)	(0.027)	(0.031)	(0.024)	(0.024)
PM2.5 94-149 μg/m^3 (per 100)	-0.008	-0.033**	-0.021	-0.046***	-0.031**	-0.043**	-0.034***	-0.019
	(0.012)	(0.014)	(0.014)	(0.016)	(0.013)	(0.018)	(0.012)	(0.013)
PM2.5 149-773 μg/m^3 (per 100)	0.003	-0.010	0.014**	-0.002	-0.005	0.002	0.003	0.000
	(0.004)	(0.017)	(0.007)	(0.006)	(0.005)	(0.009)	(0.005)	(0.005)
Year-month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day-of-week fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time-of-day × Quarter	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Days before or after plant holiday	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Worker fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Meteorology × Quarter	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Co-pollutant concentrations	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	11,212	17,367	9,519	9,564	9,496	7,332	14,246	14,333
Number of regressors	62	68	95	97	96	105	105	105
R-squared	0.038	0.048	0.051	0.065	0.051	0.037	0.033	0.053
Mean value of dependent variable	0.026	0.030	0.028	0.028	0.028	0.039	0.026	0.030

Notes: An observation is a worker-date pair in the sample of work and non-work choices. The dependent variable is 1 if the worker chose to initiate a non-attendance spell and 0 if the worker chose to work. *More productive* workers refer to those whose mean productivity is above the median and *less productive* workers are those with below-median productivity. Ordinary least squares estimates. Standard errors in parentheses, clustered by date. ***Significant at 1%, **significant at 5%, *significant at 10%.

Table 5: The impact of PM2.5 pollution on output

	(1)	(2)	(3)	(4)	(5)	(6)
PM2.5 10-62 μg/m^3 (per 100)	-24.52	-45.02***	-44.41***	-40.13***	-43.49***	-39.86***
	(19.22)	(14.34)	(14.18)	(12.28)	(12.56)	(12.30)
PM2.5 62-94 μg/m^3 (per 100)	-36.01**	-33.71***	-34.21***	-39.08***	-40.73***	-39.28***
	(15.68)	(12.44)	(12.33)	(10.97)	(11.09)	(11.00)
PM2.5 94-149 μg/m^3 (per 100)	-43.83***	-48.35***	-47.90***	-44.38***	-43.18***	-44.35***
	(10.12)	(8.69)	(8.63)	(7.92)	(8.23)	(7.94)
PM2.5 149-230 μg/m^3 (per 100)	-20.61***	-16.67***	-16.91***	-18.07***	-19.84***	-18.02***
	(7.58)	(6.40)	(6.37)	(5.58)	(6.06)	(5.54)
PM2.5 230-773 μg/m^3 (per 100)	-1.76	-0.70	-0.70	-1.34	-4.96*	-1.36
	(2.12)	(1.99)	(2.00)	(2.06)	(2.79)	(2.06)
Last work day before non-attend.		9.34*	10.55**	11.94***	12.11***	11.92***
		(5.17)	(5.18)	(4.44)	(4.46)	(4.43)
First work day after non-attend.		-4.72	-3.55	-3.26	-3.22	-3.27
		(4.65)	(4.67)	(4.15)	(4.15)	(4.15)
Inverse Mills ratio						0.34
						(1.59)
Year-month fixed effects	No	Yes	Yes	Yes	Yes	Yes
Day-of-week fixed effects	No	Yes	Yes	Yes	Yes	Yes
Time-of-day × Quarter	No	Yes	Yes	Yes	Yes	Yes
Days before or after plant holiday	No	Yes	Yes	Yes	Yes	Yes
Worker characteristics	No	No	Yes	No	No	No
Worker fixed effects	No	No	No	Yes	Yes	Yes
Co-pollutant concentrations	No	No	No	No	Yes	Yes
Observations	27,673	27,673	27,673	27,673	27,585	27,673
Number of regressors	6	48	55	49	59	50
R-squared	0.045	0.099	0.103	0.281	0.282	0.281
Mean value of dependent variable	509.04	509.04	509.04	509.04	508.94	509.04

Notes: An observation is a worker-date pair with non-zero output in the sample. The dependent variable is a worker's total output during her shift. See notes to the previous tables for fixed effects. Co-pollutant concentrations are spline functions of SO_2 and CO concentrations. All pollutant concentrations are means for the 8-hour work shift. Ordinary least squares estimates in columns (1) to (5). In column (6), variables excluded from the output equation but included in the selection equation are meteorological conditions that are allowed to vary by quarter of year. We implement the selection equation at the monthly rather than daily level due to the lack of variation in non-attendance. Standard errors in parentheses, clustered by date. ***Significant at 1%, **significant at 5%, *significant at 10%.

Table 6: The impact of PM2.5 on output, by season, time of day and worker productivity

	Cold months (1)	Warm months (2)	0 am shift (3)	8 am shift (4)	4 pm shift (5)	More productive (6)	Less productive (7)
PM2.5 10-62 μg/m^3 (per 100)	-93.38**	-34.66***	-52.06*	-9.23	-55.52***	-43.13***	-42.26***
	(45.41)	(10.42)	(27.40)	(19.53)	(14.74)	(16.11)	(14.31)
PM2.5 62-94 μg/m^3 (per 100)	-60.68**	-34.38***	-19.55	-66.73***	-48.30***	-54.34***	-25.21*
	(28.66)	(10.97)	(21.13)	(18.15)	(17.04)	(13.19)	(13.00)
PM2.5 94-149 μg/m^3 (per 100)	-40.07***	-41.33***	-46.39***	-54.15***	-21.44	-47.37***	-40.73***
	(14.81)	(9.65)	(13.42)	(12.03)	(13.47)	(9.28)	(9.85)
PM2.5 149-230 μg/m^3 (per 100)	-22.12**	-22.69**	-23.40**	-22.03***	-16.87	-20.76***	-19.15**
	(8.74)	(10.90)	(10.09)	(7.63)	(13.65)	(6.77)	(7.55)
PM2.5 230-773 μg/m^3 (per 100)	-4.78*	23.05	-0.30	-5.16	-14.74*	-6.19	-3.51
	(2.82)	(16.29)	(4.82)	(3.94)	(7.91)	(3.92)	(2.78)
Last work day before non-attend.	17.33***	8.29	17.36***	4.71	15.10**	2.39	21.27***
	(6.10)	(6.08)	(6.53)	(6.14)	(6.67)	(5.19)	(5.81)
First work day after non-attend.	1.18	-5.87	-5.13	-7.67	0.76	-5.27	-1.18
	(5.84)	(5.58)	(6.46)	(6.80)	(7.21)	(5.45)	(5.33)
Year-month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day-of-week fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time-of-day × Quarter	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Days before or after plant holiday	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Worker characteristics	No	No	No	No	No	No	No
Worker fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Co-pollutant concentrations	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10,826	16,759	9,075	9,295	9,215	13,770	13,815
Number of regressors	42	47	51	51	51	59	59
R-squared	0.291	0.292	0.286	0.293	0.310	0.192	0.186
Mean value of dependent variable	502.72	512.95	509.26	508.69	508.87	542.54	475.44

Notes: An observation is a worker-date pair with non-zero output in the sample. The dependent variable is a worker's total output during her shift. *More productive* workers refer to those whose mean productivity is above the median and *less productive* workers are those with below-median productivity. Ordinary least squares estimates. Standard errors in parentheses, clustered by date. ***Significant at 1%, **significant at 5%, *significant at 10%.

Table 7: The impact of PM2.5 on output: Further sensitivity analysis

	Net output (1)	No extremes (2)	OLS, Cubic (3)	2SLS, Cubic (4)
PM2.5 10-62 μg/m^3 (per 100)	-46.48***	-27.95***		
	(13.35)	(7.99)		
PM2.5 62-94 μg/m^3 (per 100)	-43.26***	-36.34***		
	(12.62)	(7.99)		
PM2.5 94-149 μg/m^3 (per 100)	-42.28***	-32.89***		
	(9.23)	(6.16)		
PM2.5 149-230 μg/m^3 (per 100)	-29.88***	-15.39***		
	(7.06)	(4.89)		
PM2.5 230-773 μg/m^3 (per 100)	-2.96	-4.09**		
	(2.96)	(1.86)		
PM2.5, within-shift mean (μg/m^3)			-65.48***	-69.32***
			(6.63)	(22.77)
PM2.5 squared			14.86***	13.97
			(2.46)	(11.98)
PM2.5 cubed			-1.15***	-0.95
			(0.26)	(1.63)
Year-month fixed effects	Yes	Yes	Yes	Yes
Day-of-week fixed effects	Yes	Yes	Yes	Yes
Time-of-day $\times$ Quarter	Yes	Yes	Yes	Yes
Days before or after plant holiday	Yes	Yes	Yes	Yes
Worker characteristics	No	No	No	No
Worker fixed effects	Yes	Yes	Yes	Yes
Co-pollutant concentrations	Yes	Yes	Yes	Yes
Observations	20,849	24,862	27,585	27,585
Number of regressors	59	59	53	53
R-squared	0.317	0.263	0.281	0.277
Mean value of dependent variable	500.49	507.44	508.94	508.94

Notes: An observation is a worker-date pair with non-zero output in the sample. The dependent variable is a worker's total output during her shift, except for column (1), where the dependent variable is a worker's output net of defects. Column (2) drops observations with very high or very low output. Column (3) replaces the splines by cubic functions of each pollutant's (8-hour mean) concentrations. Ordinary least squares estimates in columns (1) to (3). Column (4) reports 2SLS estimates, where PM2.5 and co-pollutant concentrations are instrumented with meteorology. Standard errors in parentheses, clustered by date. ***Significant at 1%, **significant at 5%, *significant at 10%.

Table 8: Proportionate increase in output from counterfactual particle pollution abatement, by region of China and by quarter

Region of China	Annual PM2.5	Quarter 1 % change	Quarter 2 % change	Quarter 3 % change	Quarter 4 % change	Annual % change
North China	91.1	7.2	4.5	4.1	5.5	5.3
Northeast China	61.9	4.3	1.9	1.4	4.7	3.1
East China	65.8	4.7	3.0	1.6	4.3	3.4
South Central China	69.8	4.9	3.4	2.1	4.4	3.7
Southwest China	69.2	5.7	2.3	2.0	4.2	3.5
Northwest China	77.6	6.6	2.8	2.5	5.1	4.2
Entire country	71.2	5.3	3.2	2.2	4.5	3.8

Notes: Predictions from the estimated worker-day level model. See the text for details and the Appendix for data sources, including aggregation weights. PM2.5 mass concentrations are annual means in $\mu g/m^3$.

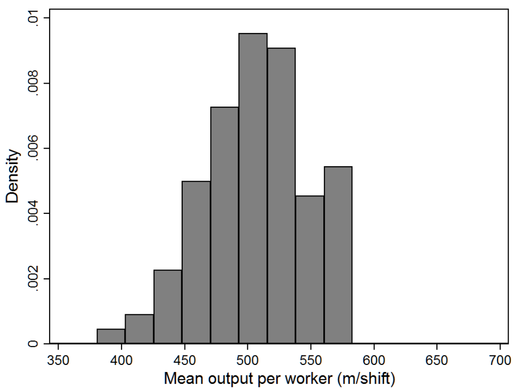

(a) Across workers

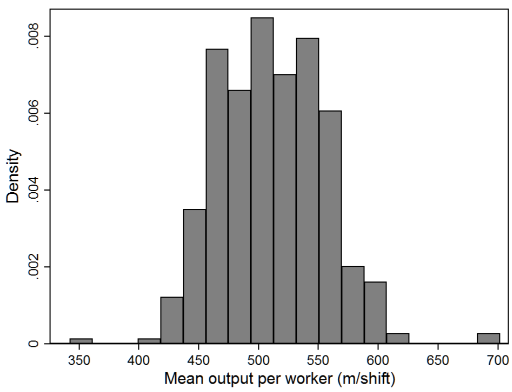

(b) Across days, 0 am shift

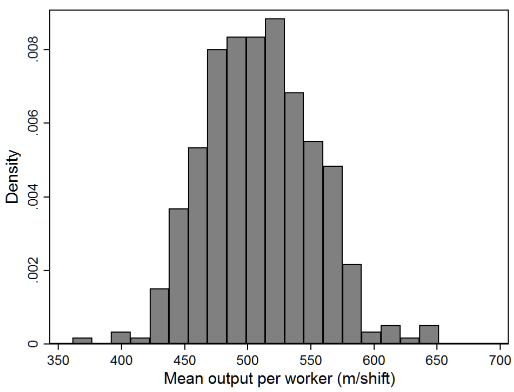

(c) Across days, 8 am shift

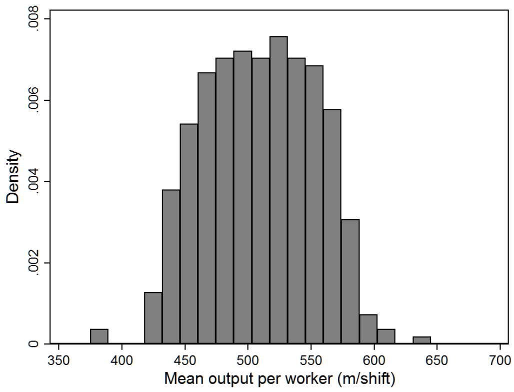

(d) Across days, 4 pm shift

Figure 1: Distributions of mean output across workers and across days in the sample (by shift)

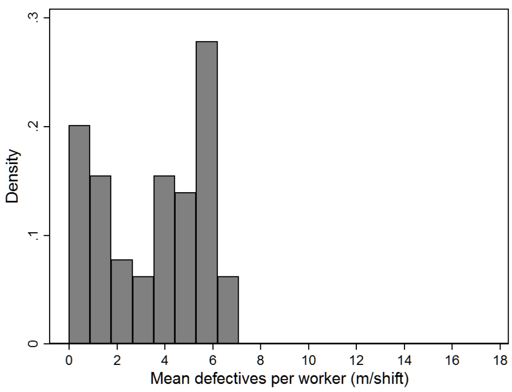

(a) Across workers

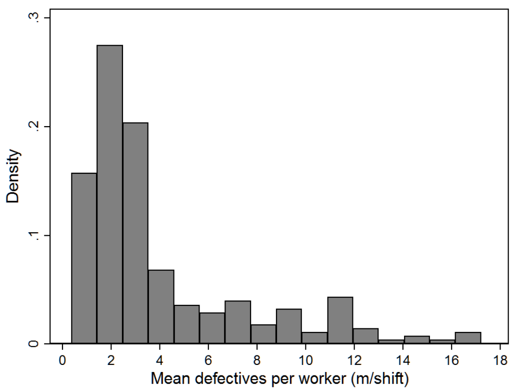

(b) Across days, 0 am shift

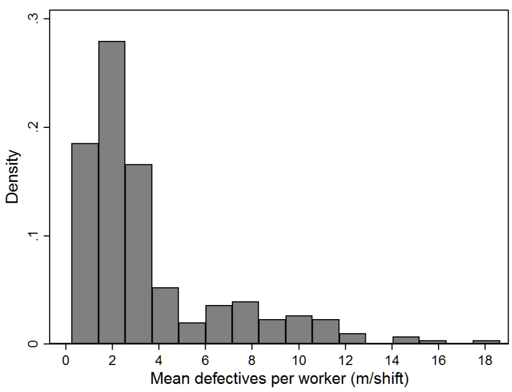

(c) Across days, 8 am shift

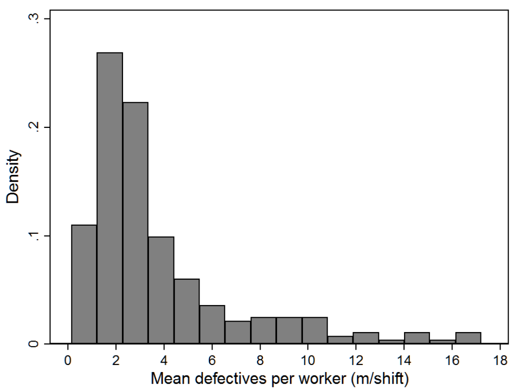

(d) Across days, 4 pm shift

Figure 2: Distribution of mean detective output across workers and across days in the sample (by shift)

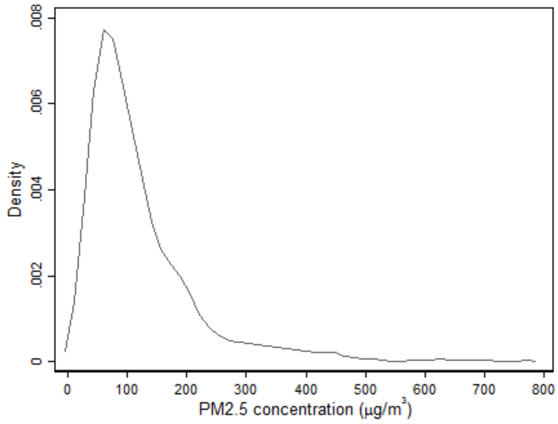

(a) PM2.5 levels 2 km from the workplace

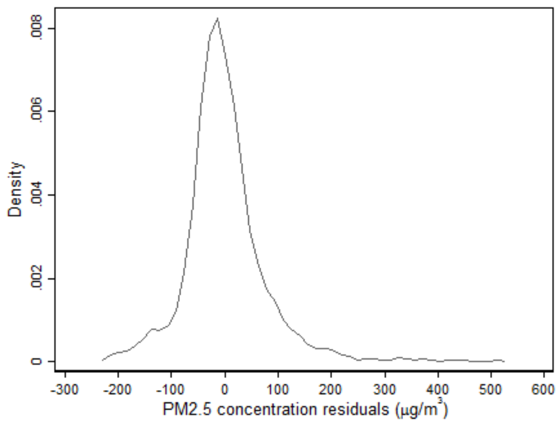

(b) Residual PM2.5 levels, 2 km away

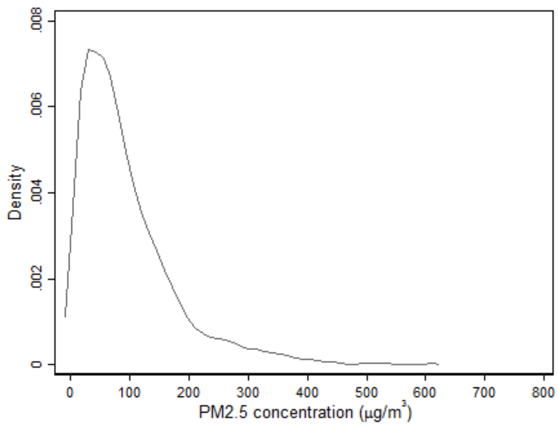

(c) PM2.5 levels in Beijing, hundreds of km away

Figure 3: Distribution of mean (8-hour) PM2.5 mass concentrations across date-shifts in the sample, 2 km from plant and in Beijing, several hundred km away

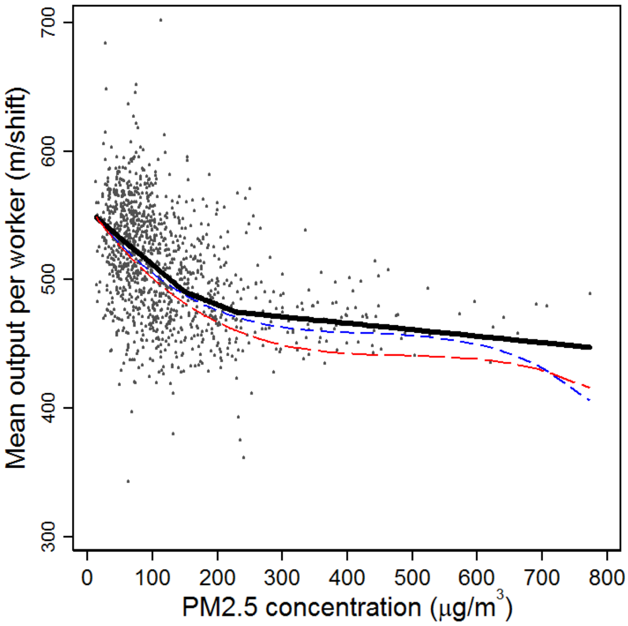

(a) Output against PM2.5, linear scale

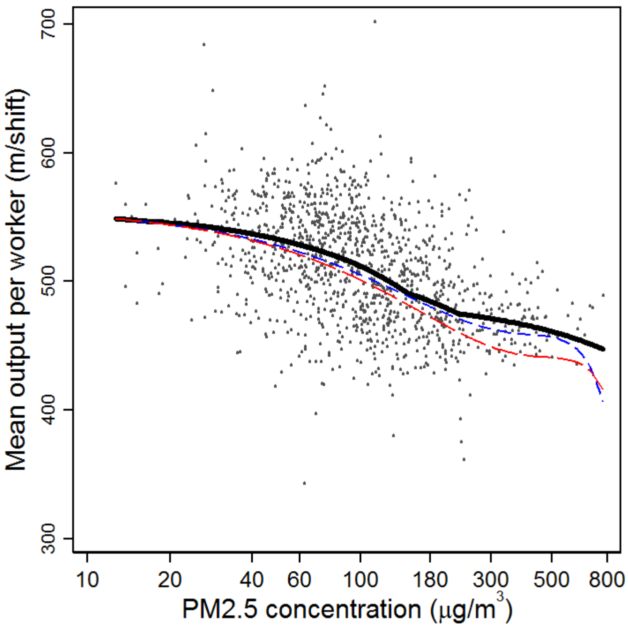

(b) Output against PM2.5, logarithmic scale

Figure 4: PM2.5 concentration (μg/m^3) and mean output per worker (m/shift). Each dot denotes a date-shift pair (mean) in the output sample. Specifications: Thick solid line shows fitted spline, short dash shows cubic polynomial fitted by OLS, long dash shows cubic polynomial fitted by 2SLS.

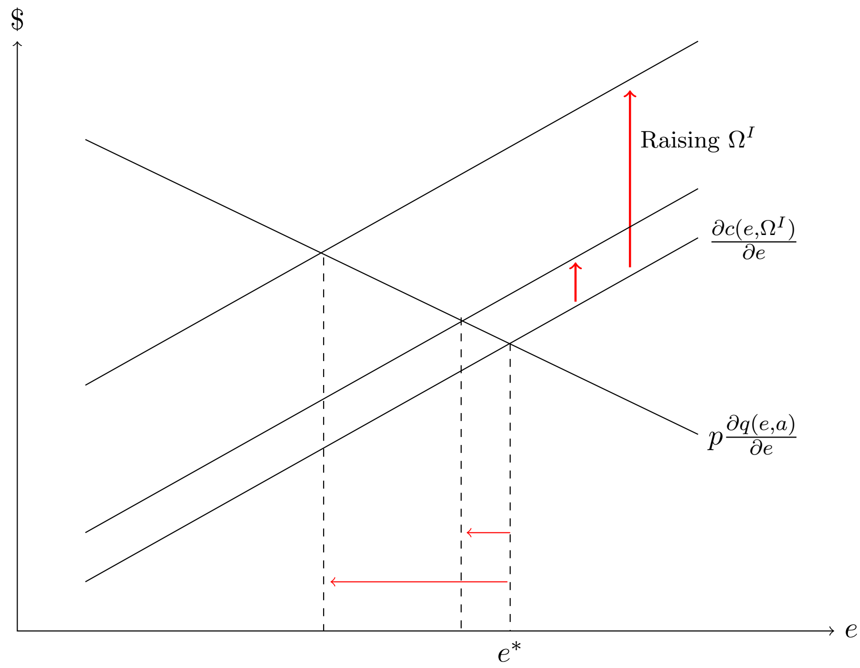

(a) Raising pollution: Less sensitive versus more sensitive workers or, more generally, non-linear effects of pollution, such as (low/moderate to high) versus (high to very high).

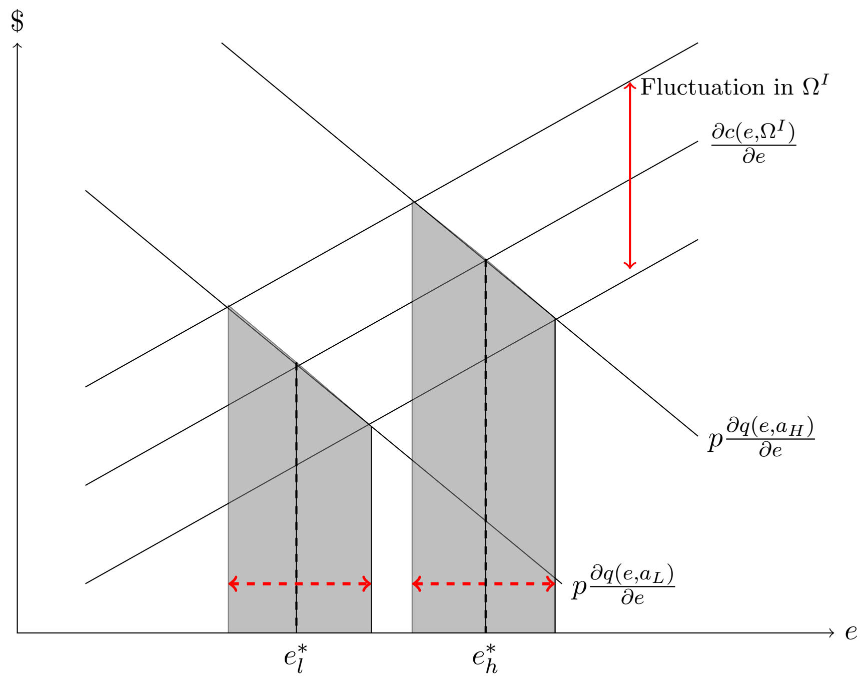

(b) Effort choice by workers of varying ability. As drawn, the optimal response of effort to pollution is similar across the two workers, implying that the output response is increasing in ability.

Figure 5: The response of effort and revenue product to PM2.5 pollution

A Data Appendix

A.1 Firm, department and workstations

The firm was founded in the 1960s.[31] It was a large state-owned enterprise (SOE), until it was privatized in the early 2000s. Currently, four manufacturing departments constitute its production line, including a preparatory department, a yarn department and two textile departments, "east" and "west," which operate in parallel and produce different types of fabric. We gained access to worker output data for the west textile department, which hereafter we refer to simply as department.

Figure A.1, panel (a) depicts one of the external walls of the building that hosts the department. The main external door to the department consists of iron bars behind plastic curtains, as can be seen in panel (b). The plastic curtains are meant to reduce the air flow between the outdoor and indoor environments. The external door leads into a warehouse that stores yarn, one of the department's main inputs. As Figure A.2, panel (a) shows, a door links this warehouse to a large indoor space that hosts the workstations, as depicted in panel (b). Inside the workplace, a long row of windows lines one of the sidewalls, shown in panel (c). Given their age, it is likely that the windows further contribute to the exchange of air between the outdoor and indoor environments. The building is mostly a single storey, with a second-storey addition on one side that hosts some offices and changing rooms.

Figure A.3 depicts the department's water-cooled air conditioning system, from the outside, in panel (a), and from the inside, in panel (b). The inside can also be seen on the ceiling shown in Figure A.2, panel (b), on top of the workstations. In addition to controlling temperature, the air conditioning system also controls indoor humidity, which is important to the textile operation (recall from Section 2 that humidity in northern China rises sharply in the summer).

The workplace consists of 348 machines—air-jet looms—that are placed in arrays and numbered. For instance, Figure A.4 depicts the fifth machine in workstation number 1,

[31] This data appendix is largely based on multiple interviews with management at the firm, in addition to the datasets described herein.

labeled 1-5. The machines were purchased from Toyota in the early 2000s, when the firm was restructured. The technical specifications under which looms operate are set by the firm's engineers/technicians. In principle, a workstation comprised of 10 looms can produce as much as 780 meters of 133-type fabric during an 8-hour shift, or about 760 m if set up to produce 134-type fabric. In practice, output will be lower because threads naturally break, requiring that the worker reconnect them while the machine lies idle.

While the median output in our sample is 506 m per worker-shift, we observe some worker-shifts with substantially higher output levels (the maximum is as high as 1,137 m). Such high levels of output exceed the nominal capacity of a workstation over an 8-hour period. On enquiring with management, we were informed that while unusual, there are cases of workers who request an extension of their work hours into the subsequent shift, subject to approval by the subsequent shift's supervisor and pending availability of workstations. While the overtime output is combined and recorded with the worker's output during her standard 8-hour shift, such extended work hours are unfortunately not recorded (likely because such occasions are exceptional). For this reason, we conduct robustness tests in which we either drop worker-shift observations with output above the 95th percentile, or replace the observed output values by 780 m.

Looms producing one type of fabric (133 or 134) can be readjusted to produce the other type by the firm's engineers/technicians (this involves dropping wares, changing operating parameters, etc). In practice, a single workstation will be set up to produce only one type of fabric. This is because a worker needs to attentively observe the number of warps, which varies across fabric type, and simultaneously operating looms producing different fabric types would be too taxing on the worker. Over our sample period, two-fifths of machines were set up to produce fabric of type 133, with the remaining producing 134.

Given the nature of materials being handled, ambient air that the workers are exposed to contains not only fine particles (i.e., less than 2.5 μm in diameter) from the outdoor environment but also, to some degree, cotton dust. Cotton dust consists of coarse, not fine, particles; however, cotton dust is a researched occupational hazard.[32] While workers in our sample are less exposed to cotton dust than those in the preparatory department, we cannot rule out the possibility that cotton dust in ambient air impacts output. Presumably, higher output levels across the department raise the concentration of cotton dust in indoor air. In the event that higher concentrations of cotton dust adversely affect output, and that this relationship holds for day-to-day variation in ambient dust concentrations, our estimates of the adverse effect of PM2.5 concentrations on output will be a lower bound.

There is no combustion inside the department, and the department does not handle chemicals (which if volatile might evaporate and contribute to the formation of particles). Thus, indoor PM2.5 pollution originates from the outdoor environment.

A.2 Workers, output, holidays and non-attendance

Worker characteristics were obtained from both the human resources department (birth year, gender, schooling, hometown, contract, etc) and the "family planning" department (number of children, their birth year, etc). We accessed such data already in digital form.

In the sample, contract duration positively correlates with a worker's tenure at the firm. In addition, workers with longer tenure (above 10 years) are eligible to purchase housing at subsidizes rates in the vicinity of the plant. Single workers with less work years at the firm typically live in dorms located inside the industrial complex.

At the end of her shift, a worker marks the point at which her production ended. The completed rolls of fabric are then inspected for defects and output is recorded. Panel (a) of Figure A.5 depicts the quality control process. Panel (b) illustrates the output records for three different workers, by date-shift, during a month. For a worker, the first column records total output in meters, whereas the second column records defective output. We obtained temporary access to these paper records (a different booklet per month per

[32] World Health Organization (1999) classifies cotton dust in the "thoracic particulate fraction," i.e., particles of diameter up to 25 μm in diameter (10 times larger than PM2.5), which can penetrate the airways of the lung and cause airway disease. Christiani et al. (1999) examine the long-term decline in lung function among cotton workers.

team), through an administrative department where they are kept, and proceeded to digitize them.

A team supervisor's main task is to oversee her workers and assist them during the shift. The supervisor rotates, along with the rest of her team, among the different shift times. A supervisor's salary is partly fixed, partly variable. The variable component of a supervisor's compensation is determined by the average output of her team's workers.

Based on output records, we observe workers who join the department after the sample period begins, as well as workers who leave the department before the sample period ends. This change in the workforce is reported in Figure A.6 by team. As mentioned in Section 2, early fall 2013 was the time with most turnover in the sample (particularly Team 2), as well as the point in which the piece rate was increased, from 0.07 to 0.10 CNY per meter of (133-equivalent) fabric. We find workers who joined the department after the wage raise to be 21 m/shift (about 4.3%) more productive than workers who left the department before the raise; a test of equality of means is rejected at the 1% significance level.[33]

Table A.1 lists dates with zero output in the department (which we label plant holidays) alongside public holidays. Though plant holidays are usually set around the official holiday calendar, they only partially overlap with public holidays as they depend on the production schedule. A robustness test we conduct is to drop observations pertaining to dates on the official holiday calendar that were not observed by the department as well as three days before and three days after public holidays.

Dates in which output records show positive output for a worker are labeled attendance dates. In addition to plant holidays, output records indicate zero output for specific workers on specific dates—we label these non-attendance dates. Figure A.7 indicates the rates of work attendance across dates and across workers, respectively. For 94% of dates in the sample (other than plant holidays), more than nine-tenths of the workforce attends work. On one isolated date, May 9, 2013, only three-quarters of the workers have positive output—we do not know what happened that day. (Our estimation results are robust to

[33] Comparing these new workers to workers who are in the sample throughout the entire period, the former group's output is slightly lower, presumably due to learning.

dropping this date from the sample.) Most workers in the sample have attendance rates exceeding 90%, noting that non-attendance includes planned leave (recall that we are unable to distinguish between planned and unplanned non-attendance). Three workers in the sample never take leave (beyond the 54 days of plant holiday). Figure A.9 plots rates of non-attendance against mean PM2.5 concentrations on dates in our sample.

A.3 Air pollution and meteorology

We obtained data from several sources. Hourly readings for pollutant mass concentrations (PM2.5, PM10, SO_2, CO, NO_2, and O_3) at four Chinese Ministry of Environmental Protection sites located in the city where the firm is based are available at `www.cnpm2d5.com/`. Hourly readings for PM2.5 mass concentrations at US Embassies in Beijing, Chengdu, Guangzhou, Shanghai and Shenyang are available from the US Department of State at `www.stateair.net/web/historical/1/1.html`. Meteorological conditions measured in the city where the firm is based come from two sources: (i) 3-hourly readings for ground temperature, humidity, atmospheric pressure, precipitation, wind speed and wind direction are available at `www.worldweatheronline.com`; and (ii) daily readings for ground temperature (maximum and minimum), precipitation, wind speed and wind direction are available at `www.lishi.tianqi.com` (these are based on readings by the National Meteorological Center).

We inspected pollutant concentrations, reading by reading, and did minimal imputing with regard to the raw data. Specifically, for the closest PM2.5 monitor, 58 date-hour observations out of a total of 10,186 non-missing observations over the sample period (April 2013 to June 2014) were recoded as missing, largely because values over adjacent hours were invariant at 3 or at 1000 $\mu g/m^3$. Our results are robust to not recoding these 0.6% of observations (a spreadsheet listing the raw data and the 58 recoded values is available from the authors). Table A.2 shows that the availability of PM2.5 measurements is high and fairly balanced over months of year and hours of day. For example, data availability ranges from 99% of possible date-hour measurements in January to 82% in November, and from 96% at 7 pm to 87% at 12 pm. Table A.3 indicates the increased data availability when using the mean concentration at the three other nearby PM2.5 monitors to impute for a missing value at the nearest monitor at a given date-hour. For example, data availability in November increases by 14 percentage points. Table A.4 reports the very high correlation between hourly PM2.5 at the four monitoring sites, all located up to 5 km from the plant. The table also reports correlation coefficients between these hourly PM2.5 concentration measurements, in the city where the firm is based, and measurements at the US Embassy in Beijing and in Shenyang. For perspective, our city of study, Beijing and Shenyang are all within several hundred km of each other. We also recoded to missing a few date-hour readings with negative concentrations in the US Department of State data.

Figure A.10 plots daily mean PM2.5 concentrations over our sample period at the monitoring site located 2 km from the plant and in Beijing, respectively. Given some very high measurements in the winter, these plots are reproduced in Figure A.11 restricting the vertical axis to no more than 400 μg/m^3. Compared to Beijing, particle pollution in the city we examine tends to be higher, while both locations follow a similar annual pattern.

We similarly inspected measured concentrations for pollutants other than PM2.5. In the case of SO_2, we recoded as missing 0.04% of all original non-missing date-hour observations. No recoding seemed called for in the case of CO. Whereas for CO our analysis uses readings at the site located 2 km from the plant, in the case of SO_2 we use the mean reading across all four city monitors, since the closest monitor exhibits some periods of low values and low variability. A spreadsheet comparing the raw individual data to the cross-monitor means is available from the authors.

Meteorological data was missing for only one date in the sample period, September 6, 2013, and only for the 24-hour dataset. Except for wind direction, where we use daily observations, our analysis uses meteorological data from the higher-frequency (3-hour) dataset, though we ensured that both data sources are in broad agreement. For example, Figure A.12 shows that mean daily temperatures as recorded in the 3-hourly data correlate tightly with those in the daily dataset.

Analysis of residuals: Worker output, PM2.5 concentration and temperature. We separately regress output (mean across workers), PM2.5 concentration, and temperature, averaged at the date-shift level, on calendar-month, day-of-week (including public holiday indicators) and time-of-day×quarter fixed effects. We note that the R^2 for the temperature regression is over double that for the PM2.5 concentration regression, respectively, R^2 of 88% and 39%, suggesting that seasonal, weekly and diurnal cycles (both natural and anthropogenic) are highly predictive of temperature, while PM2.5 concentrations show more orthogonal variation. We obtain fitted residuals for each of these three regressions and plot them in Figure A.13, namely worker output against PM2.5 in panel (a), and mean worker output (again) against temperature in panel (b). The strong negative relationship between output and fine-particle pollution—even when time fixed effects have been netted out—is evident, while no relationship between output and temperature is apparent.

A.4 Back-of-the-envelope calculations

To conduct the back-of-the-envelope calculation reported in Table 8, we obtained additional data and utilize it as explained next. First, annual mean PM2.5 mass concentrations in 2014, in μg/m^3, were obtained for 190 large Chinese cities from `www.cnpm2d5.com/`. Since this data provides cross-sectional (spatial) variation, we combine it with temporal (seasonal) variation in the US Department of State data as follows. Based on minimum (Haversine) distance, we assign every one of the 190 cities to one of the five US embassies that measure and publish PM2.5 on an hourly basis (Beijing, Chengdu, Guangzhou, Shanghai and Shenyang). For example, the nearest PM2.5-monitoring US Embassy to Tianjin, one of the 190 cities, is located in Beijing. Then, for every one of the 190 cities, we take the hourly PM2.5 levels (averages as measured between 2012 and 2014) over the 365×24 day-of-year×hour-of-day pairs at the nearest US Embassy, and scale this hourly series up or down by a scalar equal to the ratio of annual means (from the cross-sectional data) for the city relative to the US Embassy location. In the example, we take the hourly series for the US Embassy in Beijing and scale it by 86/85, the numerator and denominator being the annual means in Tianjin and Beijing, respectively.

We now have an "actual PM2.5 pollution" series for every one of the 190 cities. We take our worker-date-shift sample and, city by city: (i) replace PM2.5 values in the sample by the actual (8-hour mean) PM2.5 pollution in the city (Tianjin in the example), followed by (ii) use the modified sample to predict, based on the estimated baseline output specification (column (5) of Table 5), output by worker by day during the 12 in-sample months of May 2013 to April 2014. Aggregate predicted output by season—the first to the fourth quarter of the year—under actual PM2.5 pollution for every one of the 190 cities is then the sum of individual output across workers within the corresponding three-month quarter. (We pick May 2013 as the first month of a full 12-month period because this is when the US State Department began monitoring PM2.5 in Shenyang.)

To obtain predicted output under a counterfactual PM2.5 pollution scenario, for every one of the 190 cities, we repeat the steps described in the preceding paragraph except that we truncate actual PM2.5 concentrations (by date-hour, by city) at 25 μg/m^3. (The exercise can be repeated for other thresholds.) This yields an aggregate output prediction by city, by season under a "counterfactual PM2.5 pollution" scenario. We compute the proportionate output change, in the counterfactual relative to actual pollution, by city and by season.

We aggregate across cities within each of 31 Chinese provinces (where the 190 cities are located) using population weights, namely the number of residents aged between 15 and 64 years of age (and thus more likely to be in the workforce). Population data are obtained from the 2010 Population Census. We further aggregate from province-quarter level to region-quarter, and similarly to country and year, using GDP by province-quarter data as weights. GDP data for 2013 was obtained from the National Bureau of Statistics of China, available at `http://data.stats.gov.cn/workspace/index?m=fsjd`.

As a second back-of-the-envelope calculation, we consider whether managers at the workplace we study would privately benefit from installing indoor pollution abatement technology. Motivating this question, personal interviews with principals at some international schools in China revealed that they had recently installed central air purifying systems; these systems are priced around US$ 1 million for the entire campus, and installation includes sealing windows and creating "positive indoor pressure" such that air that enters the building does so mostly through filters. We ignore any lagged effects of human exposure to pollution, say when the worker was at home, instead assuming that only contemporaneous PM2.5 in ambient air at the workplace affects labor productivity. We use our estimated model to predict the increase in output for the 98 workers—again holding the distribution of work days fixed—were managers at the workplace to install indoor pollution control equivalent to truncating the observed hourly (outdoor) PM2.5 concentration series at 25 $\mu g/m^3$, as in the counterfactual abatement policy considered in the first back-of-the-envelope (and again not changing SO_2 and CO levels).

With air in the city that is home to our studied workplace being among China's most polluted, the predicted increase in annual output comes out at 7.5%. (This can be compared to the predicted output changes across China's regions, reported in Table 8.) We were informed that the firm's annual revenue is about US$ 30 million, and we crudely multiply this revenue by 7.5%, the predicted output increase. We consider a few different combinations for: (i) the firm's profit margin, specifically, what proportion of the 7.5% revenue increase the firm captures; and (ii) the capital investment (the firm consists of further departments and buildings, other than the one we study); *such that the indoor pollution control project breaks even*, with a "project NPV" of US$ 0. We further assume a project horizon of 10 years, and an annual discount rate of 5%.

For the project to privately break even when the capital investment is US$ 1 million, the required margin on the additional revenue product of labor is 6%. For a capital investment of US$ 2 million, the required break-even margin on the additional revenue is 12%. This exercise, with the marginal private returns it points to, suggests to us why pollution control might not be prevalent in the workplace even in this most polluted region of China, other than at organizations that cater to concerned parents, or handle sensitive products, such as hospital services, pharmaceuticals and semiconductors.

Table A.1: Plant holidays and public holidays during the sample period

	Plant holiday	Public holiday
Qingming Festival	04/04 to 04/05/2013	04/04 to 04/06/2013
International Workers' Day	05/01 to 05/02/2013	04/29 to 05/01/2013
Dragon Boat Festival	06/09 to 06/15/2013	06/10 to 06/12/2013
Mid-Autumn Day	Not observed by the plant	09/19 to 09/21/2013
Summer Holiday	7/31 to 8/7/2013; 8/15 to 8/18/2013	-
National Day	09/30 to 10/05/2013	10/01 to 10/07/2013
New Year	Not observed by the plant	01/01/2014
Chinese New Year	01/26 to 02/11/2014	01/31 to 02/06/2014
Qingming Festival	03/28 to 03/31/2014	04/05 to 04/07/2014
International Workers' Day	05/01 to 05/04/2014	05/01 to 05/03/2014
Dragon Boat Festival	Not observed by the plant	05/31 to 06/02/2014

Table A.2: Hourly data availability for PM2.5 concentrations at the closest monitor, by month of year and hour of day

	Panel A											
By month (%)	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
	98.9	93.8	96.5	90.6	93.1	95.7	98.1	92.5	88.8	91.9	81.9	91.4
	Panel B											
By hour (%)	0 am	1 am	2 am	3 am	4 am	5 am	6 am	7 am	8 am	9 am	10 am	11 am
	92.8	93.0	93.6	93.0	92.5	92.8	93.9	92.1	92.1	93.6	93.0	93.4
	12 pm	1 pm	2 pm	3 pm	4 pm	5 pm	6 pm	7 pm	8 pm	9 pm	10 pm	11 pm
	86.8	89.3	91.4	91.9	90.4	92.8	95.0	96.1	94.7	94.3	95.4	94.7

Table A.3: Increase in hourly data availability when augmenting data at the closest PM2.5 monitor with data from three other nearby monitors

						Panel A						
By month (%)	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
	0.0	2.1	0.0	1.8	1.3	1.5	0.1	0.0	5.8	2.4	13.5	6.0
						Panel B						
By hour(%)	0 am	1 am	2 am	3 am	4 am	5 am	6 am	7 am	8 am	9 am	10 am	11 am
	1.8	2.2	2.9	2.2	2.6	2.6	2.9	3.7	3.9	3.3	3.1	3.3
	12 pm	1 pm	2 pm	3 pm	4 pm	5 pm	6 pm	7 pm	8 pm	9 pm	10 pm	11 pm
	4.4	3.5	3.5	2.2	3.9	1.5	1.5	1.3	1.8	1.5	1.5	1.3

Table A.4: Correlation coefficients for PM2.5 measurements at monitoring sites near to the plant and several hundred kilometers away

	Closest Monitor	Nearby Site 1	Nearby Site 2	Nearby Site 3	Beijing	Shenyang
	Panel A: Hourly PM2.5 concentration time series					
Closest Monitor	1.00	0.88	0.91	0.89	0.37	0.37
Nearby Site 1	0.88	1.00	0.91	0.89	0.38	0.33
Nearby Site 2	0.91	0.91	1.00	0.93	0.39	0.37
Nearby Site 3	0.89	0.89	0.93	1.00	0.40	0.35
Beijing (several hundred km away)	0.37	0.38	0.39	0.40	1.00	0.41
Shenyang (several hundred km away)	0.37	0.33	0.37	0.35	0.41	1.00
	Panel B: Daily mean PM2.5 concentration time series					
Closest Monitor	1.00	0.93	0.94	0.92	0.46	0.54
Nearby Site 1	0.93	1.00	0.96	0.95	0.50	0.52
Nearby Site 2	0.94	0.96	1.00	0.96	0.49	0.54
Nearby Site 3	0.92	0.95	0.96	1.00	0.52	0.52
Beijing (several hundred km away)	0.46	0.50	0.49	0.52	1.00	0.56
Shenyang (several hundred km away)	0.54	0.52	0.54	0.52	0.56	1.00

(a) One of the external walls of the workplace building

(b) The main door of the building, leading into the department from outside

Figure A.1: The outdoor built environment

(a) The door linking the warehouse to the workplace

(b) The workplace and its workstations

(c) A long row of windows lining one of the sidewalls

Figure A.2: The indoor built environment

(a) Viewed from outside

(b) Viewed from inside (ceiling over the workstations)

Figure A.3: The water-cooled air conditioning system

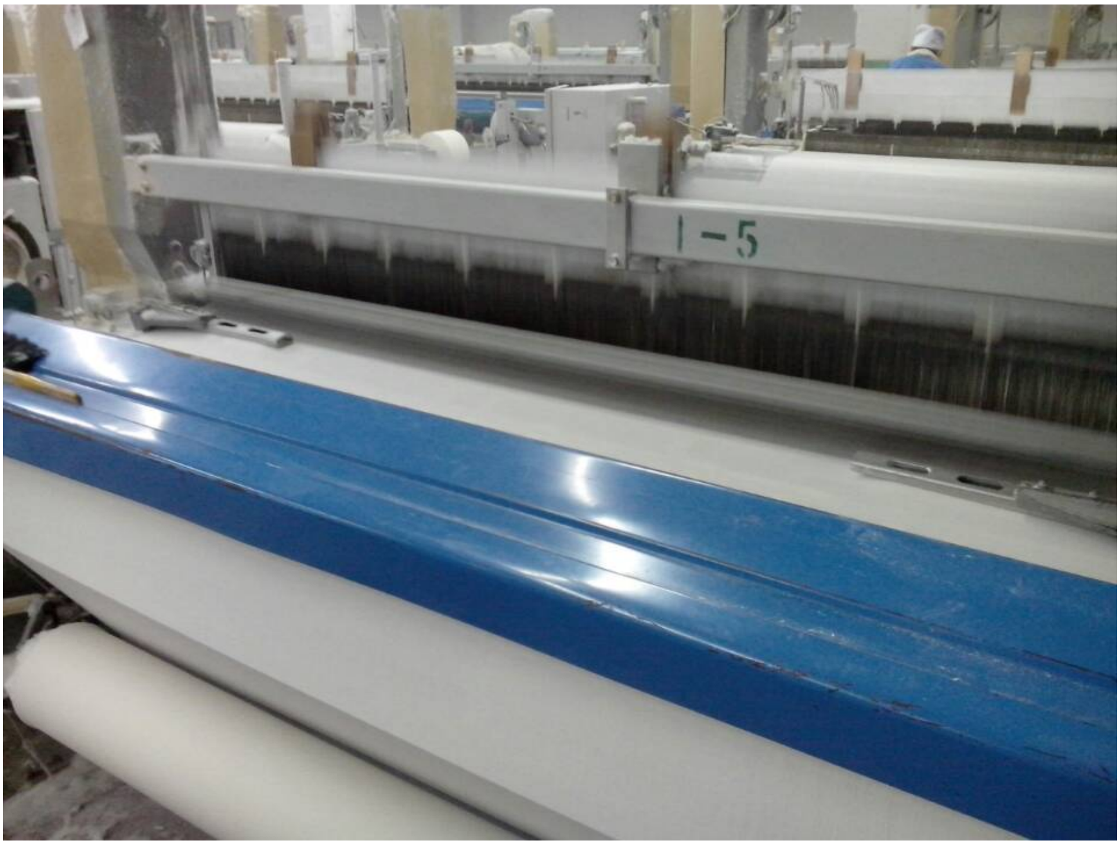

Figure A.4: Machine 5 (a Toyota air-jet loom), workstation 1

(a) Inspecting defects

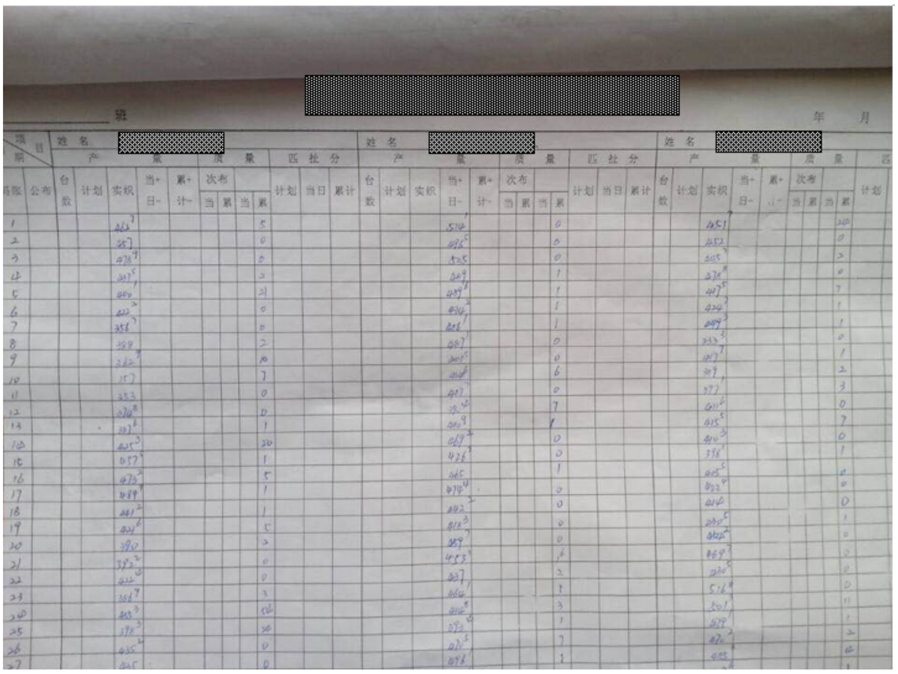

(b) An example of output records (on paper) for three workers over one month

Figure A.5: Inspecting and recording individual output

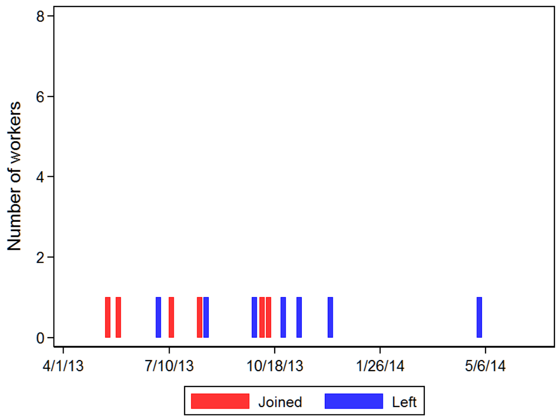

(a) Team 1's turnover

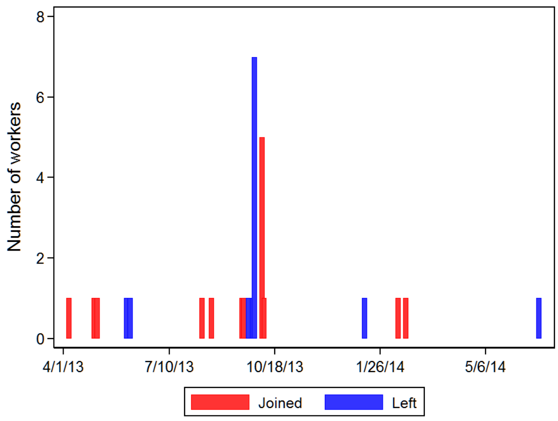

(b) Team 2's turnover

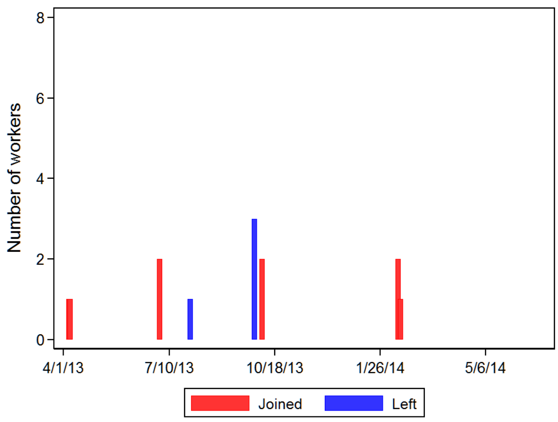

(c) Team 3's turnover

Figure A.6: Number of workers who joined and left the department during the sample period, by team

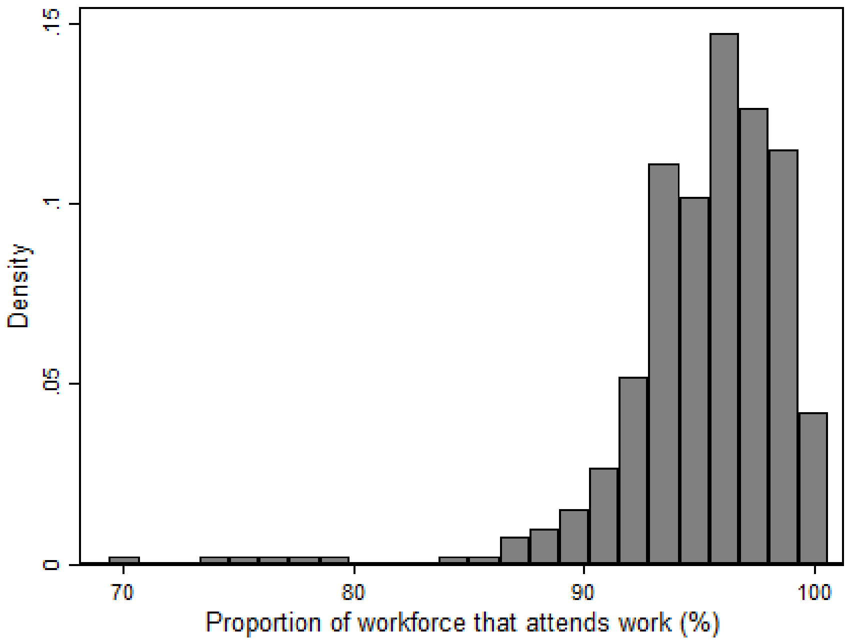

(a) Across days

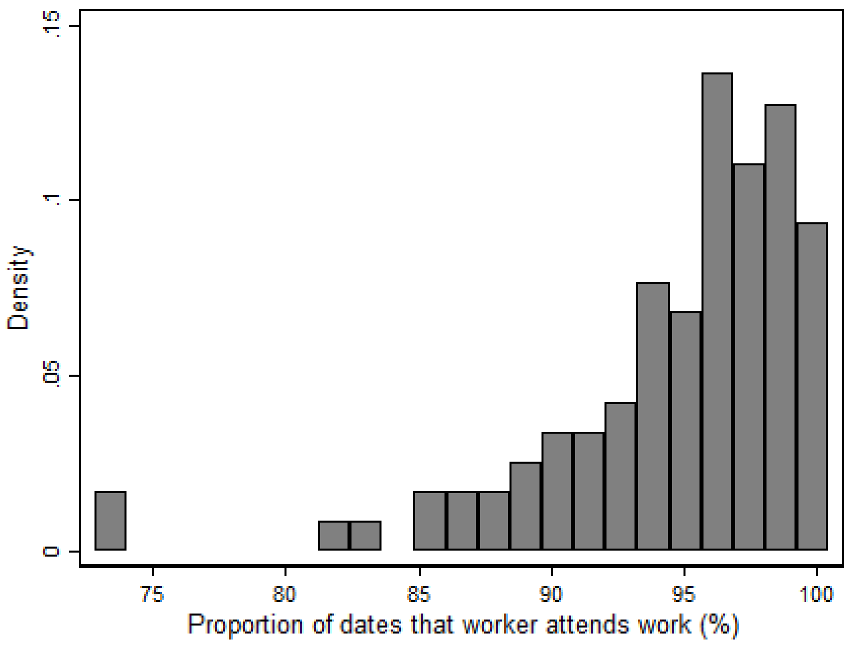

(b) Across workers

Figure A.7: Attendance rates across days and across workers in the sample

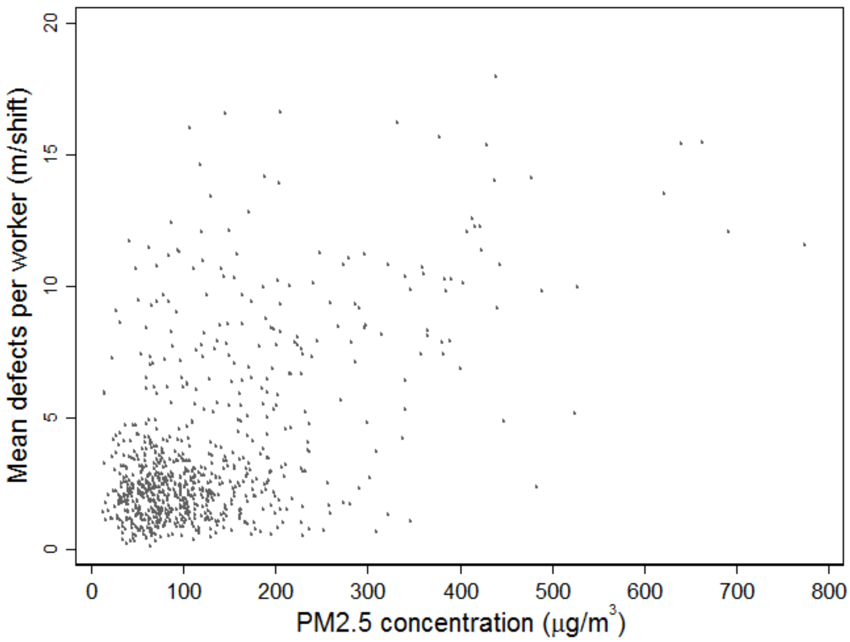

(a) Defective output against PM2.5, linear scale

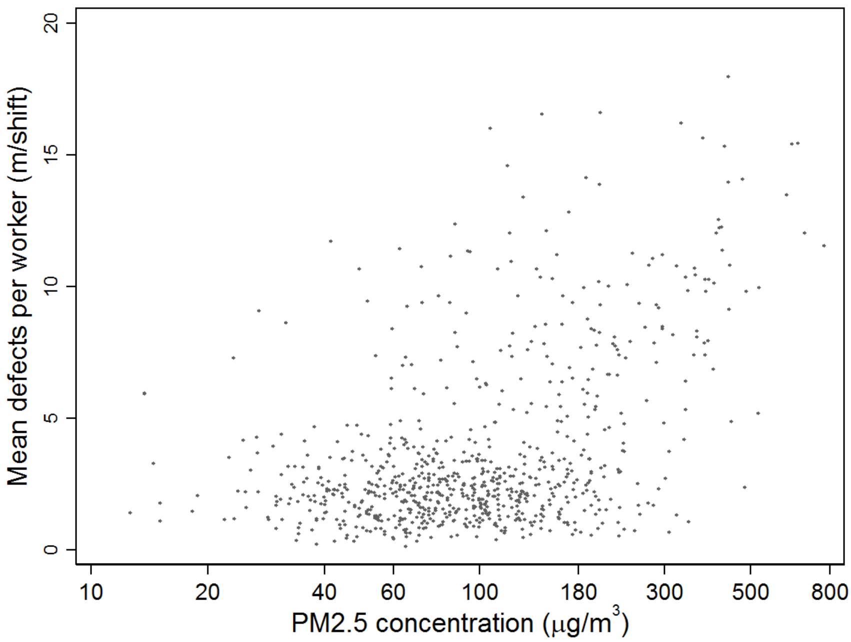

(b) Defective output against PM2.5, logarithmic scale

Figure A.8: PM2.5 concentration (μg/m^3) and mean defects per worker (m/shift)

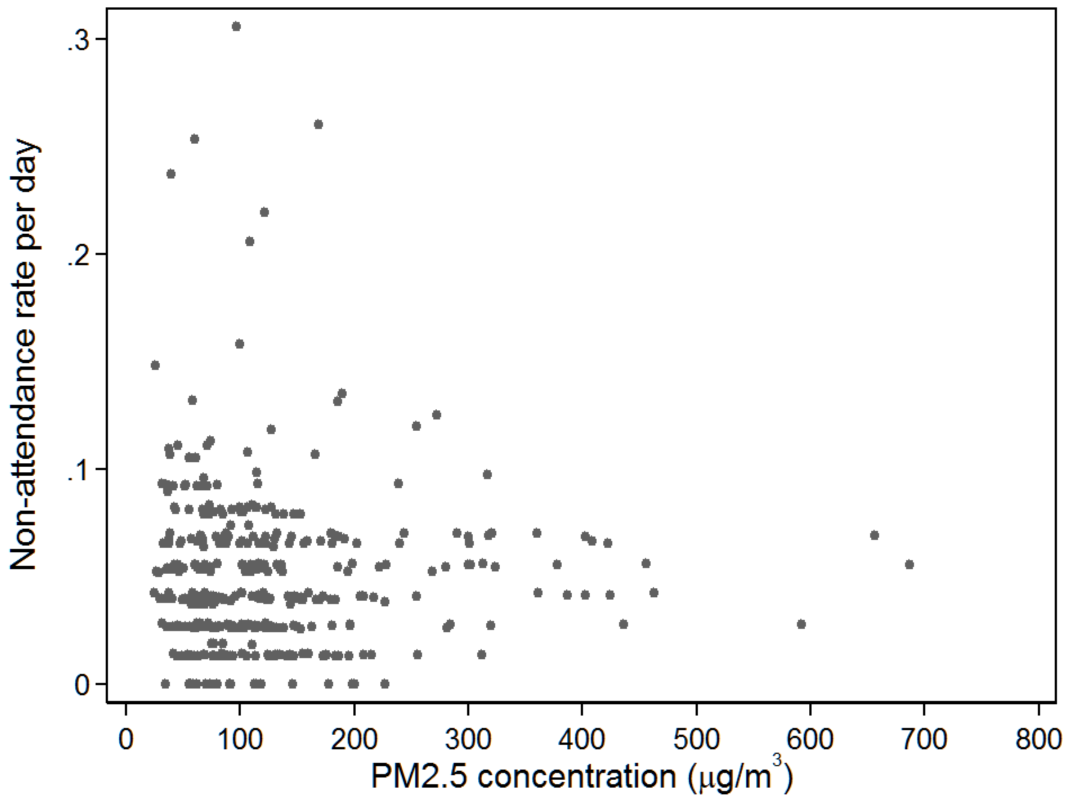

Figure A.9: PM2.5 concentration (μg/m^3, daily means) and the proportion of the workforce that does not attend work

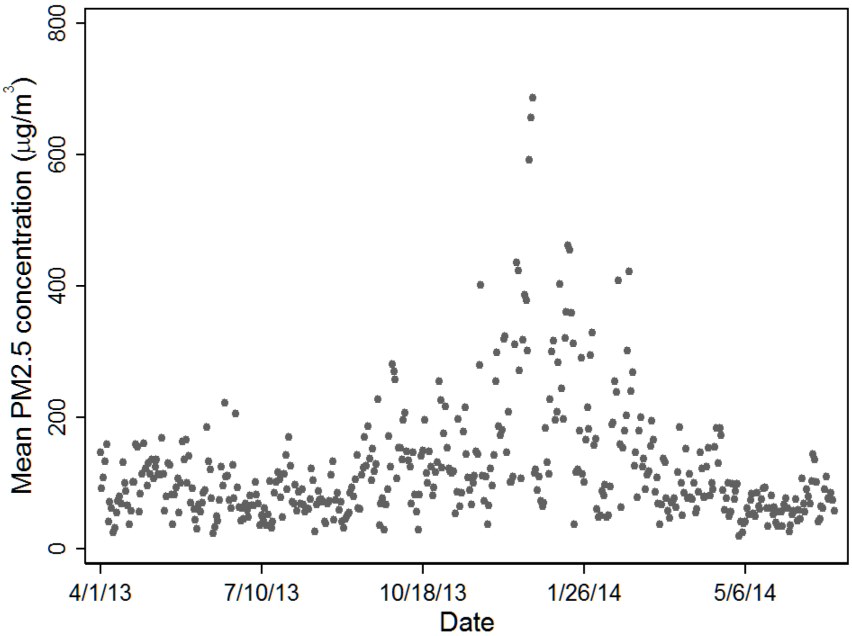

(a) 2 km from the plant

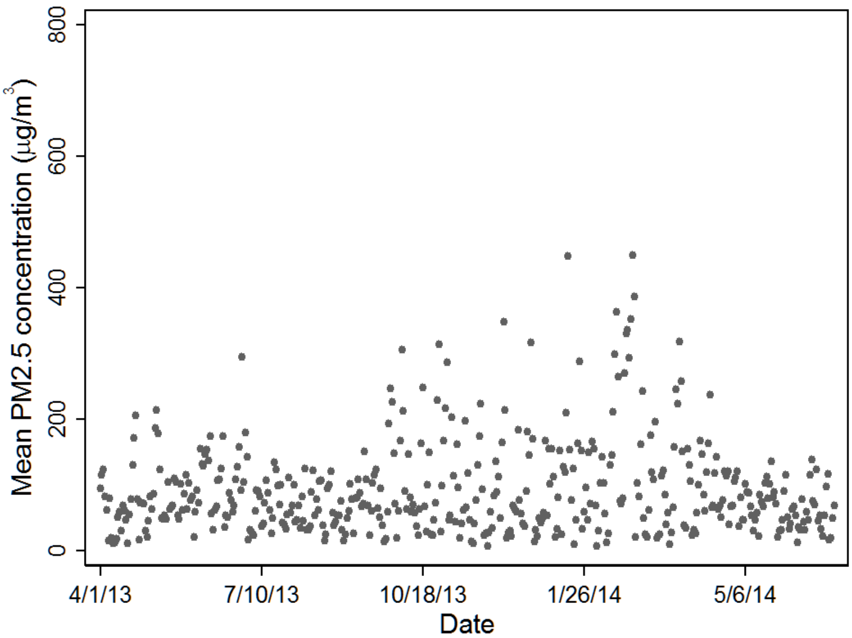

(b) In Beijing

Figure A.10: PM2.5 concentration (μg/m^3, daily means) at the closest monitoring site and in Beijing, during the sample period

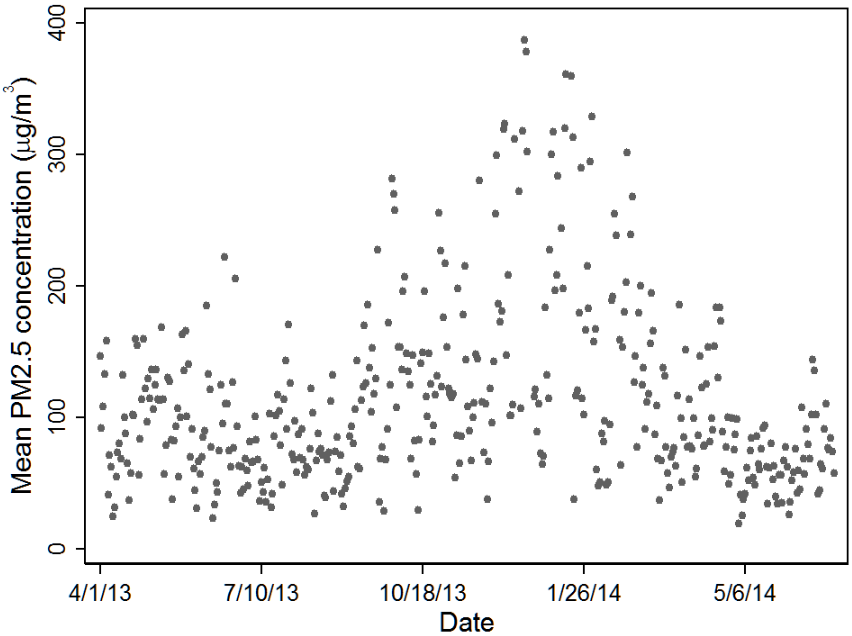

(a) 2 km from the plant

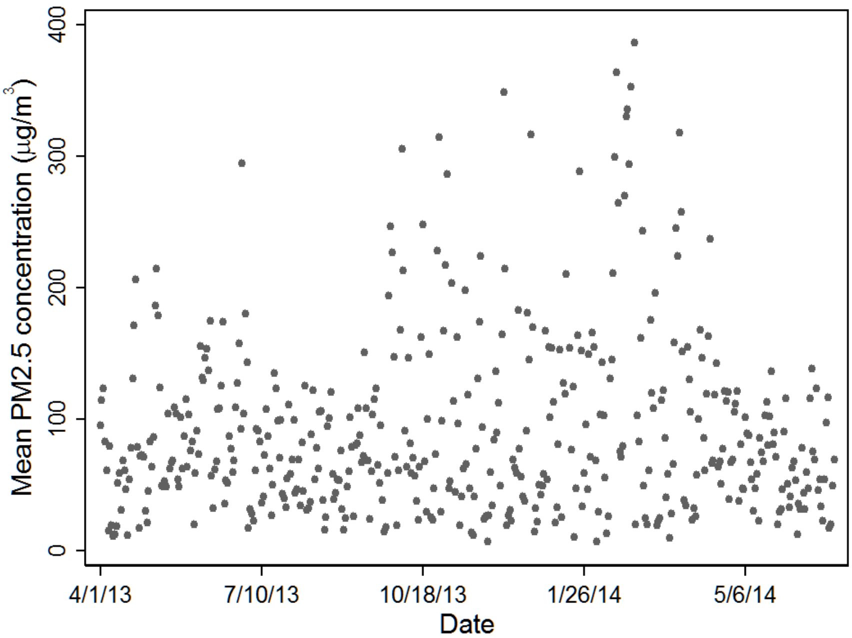

(b) In Beijing

Figure A.11: PM2.5 concentration (μg/m^3, daily means) at the closest monitoring site and in Beijing, during the sample period (excluding readings above 400 μg/m^3)

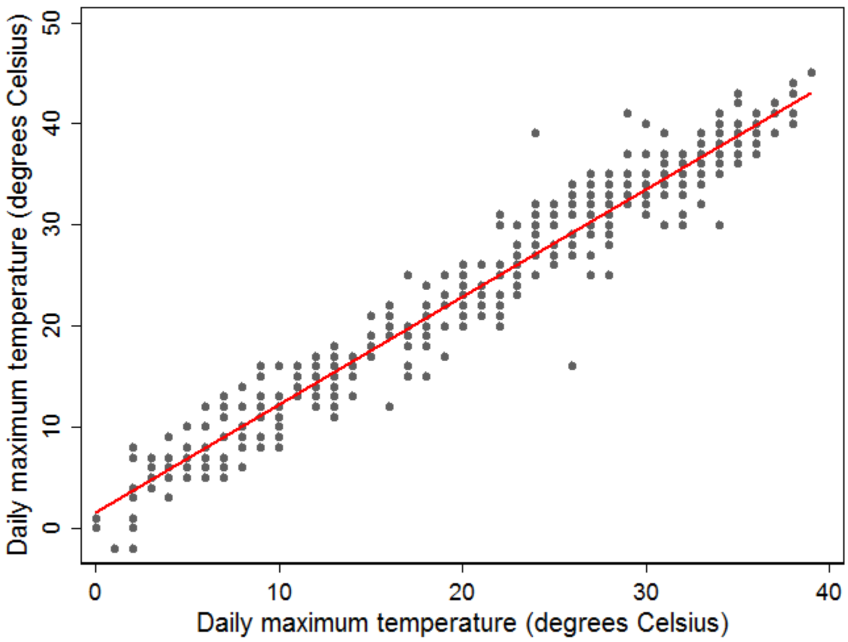

(a) Daily maximum

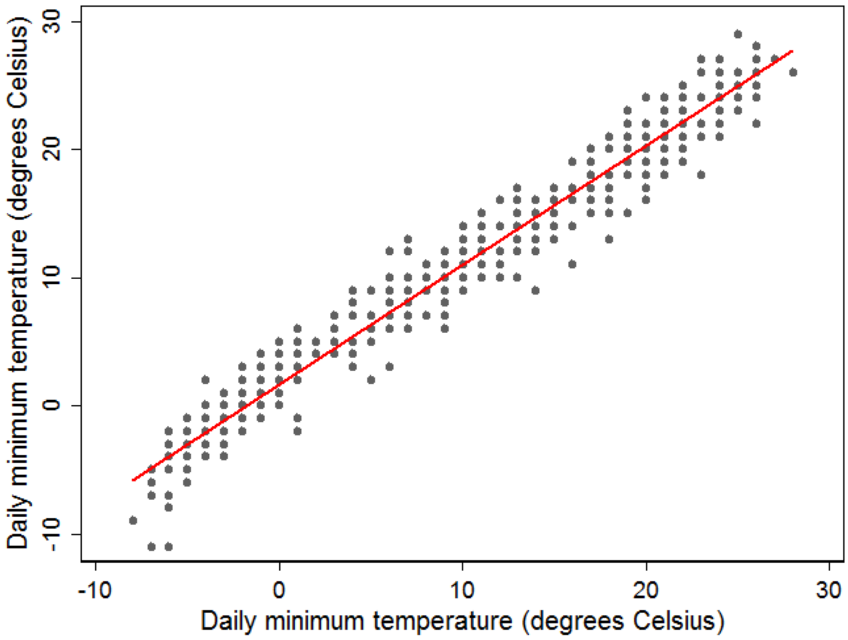

(b) Daily minimum

Figure A.12: Temperatures as observed in the two meteorological data sources

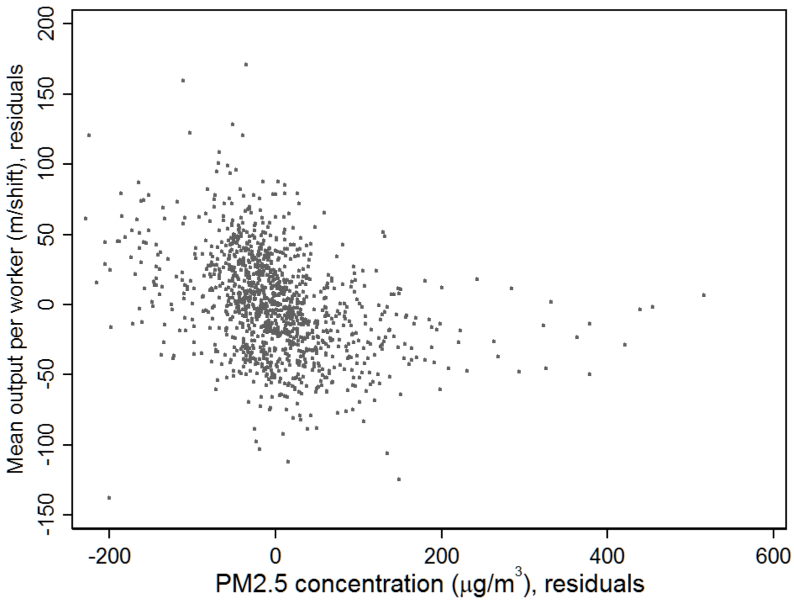

(a) Mean output (residuals) against PM2.5 level (residuals)

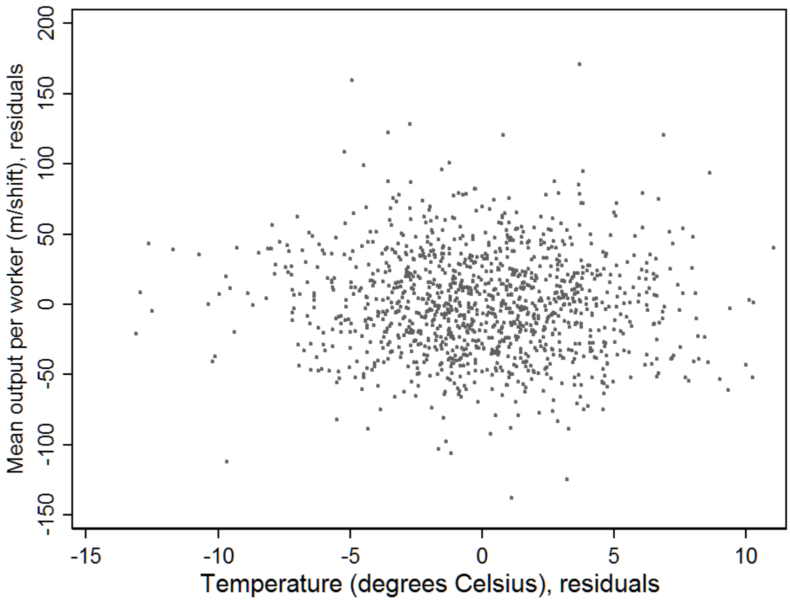

(b) Mean output (residuals) against temperature (residuals)

Figure A.13: Variation in output (mean across workers), PM2.5 concentration and temperature (outdoor), averaged at the date-shift level, when time fixed effects have been netted out

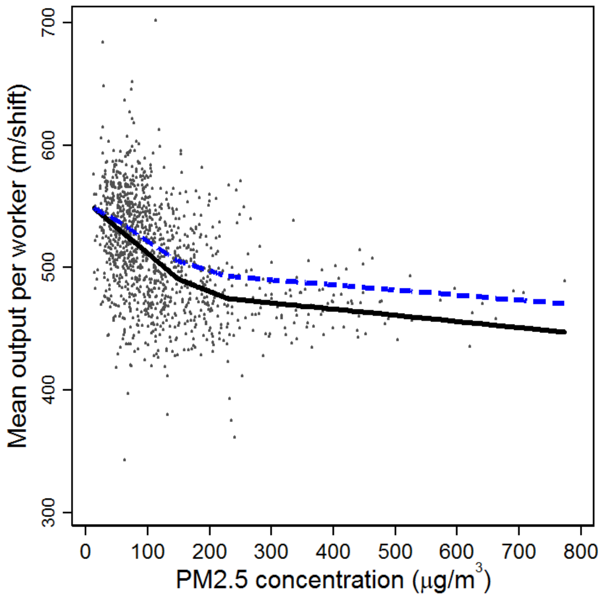

(a) Output against PM2.5, no extreme output values

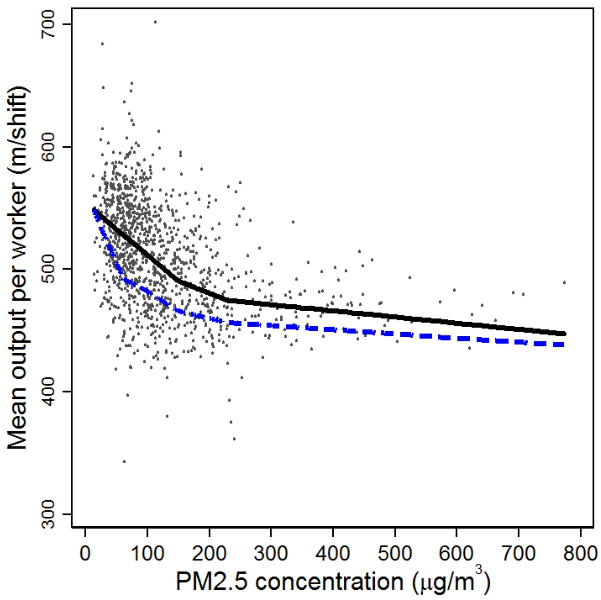

(b) Output against PM2.5, contemporaneous and lagged

Figure A.14: PM2.5 concentration (μg/m^3) and mean output per worker (m/shift). Each dot denotes a date-shift pair (mean) in the output sample. In each panel, the solid line indicates the fitted baseline specification, the dashed line shows estimates, in panel (a), when observations with very high or very low output are dropped, and, in panel (b), when 24-hour lagged PM2.5 levels are controlled for on top of contemporaneous PM2.5 (the horizontal axis denotes joint variation).