FOR ONLINE PUBLICATION

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 $\mu g/m^3$. 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 inter-

national 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 μg/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

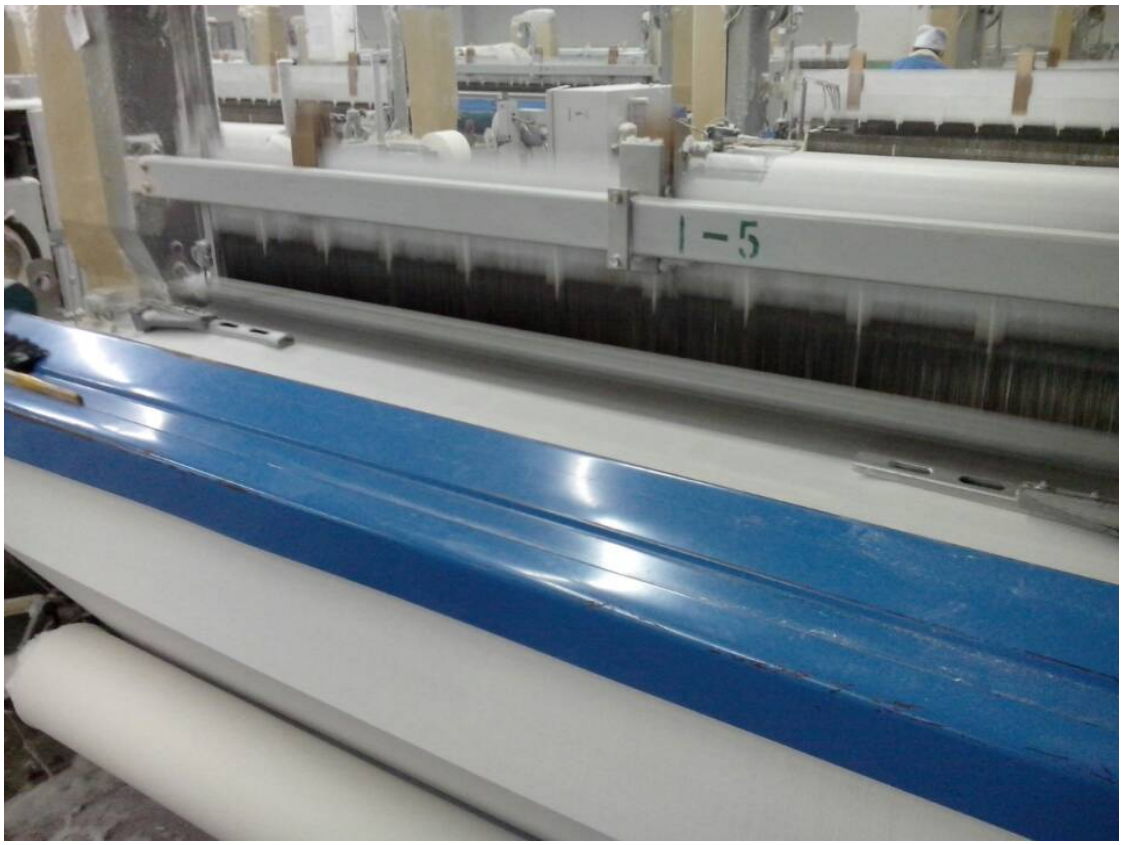


Figure A.4: Machine 5 (a Toyota air-jet loom), workstation 1

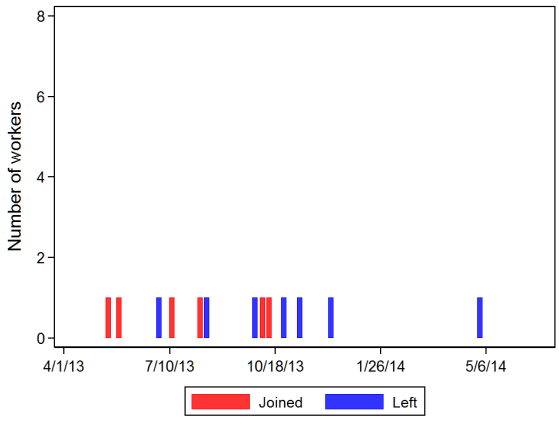


(a) Inspecting defects

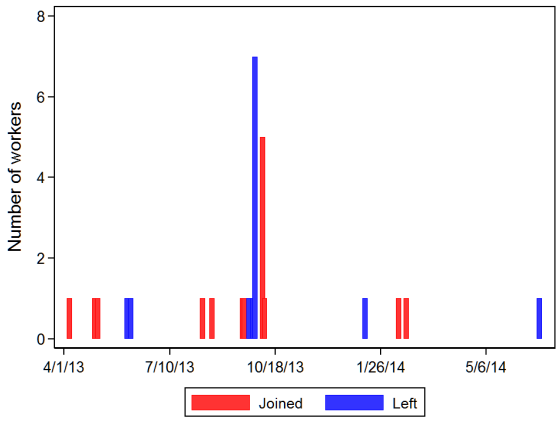


(b) An example of output records (on paper) for three workers over one month

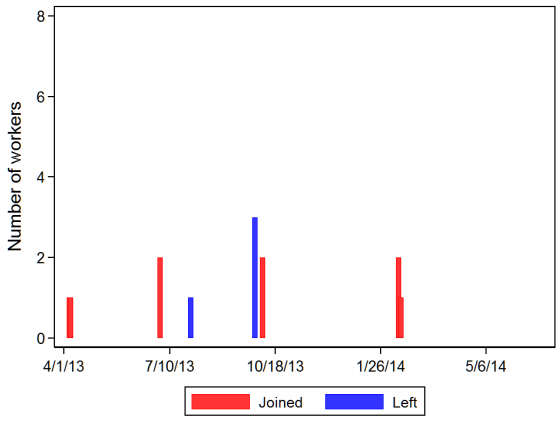
Figure A.5: Inspecting and recording individual output



(a) Team 1's turnover

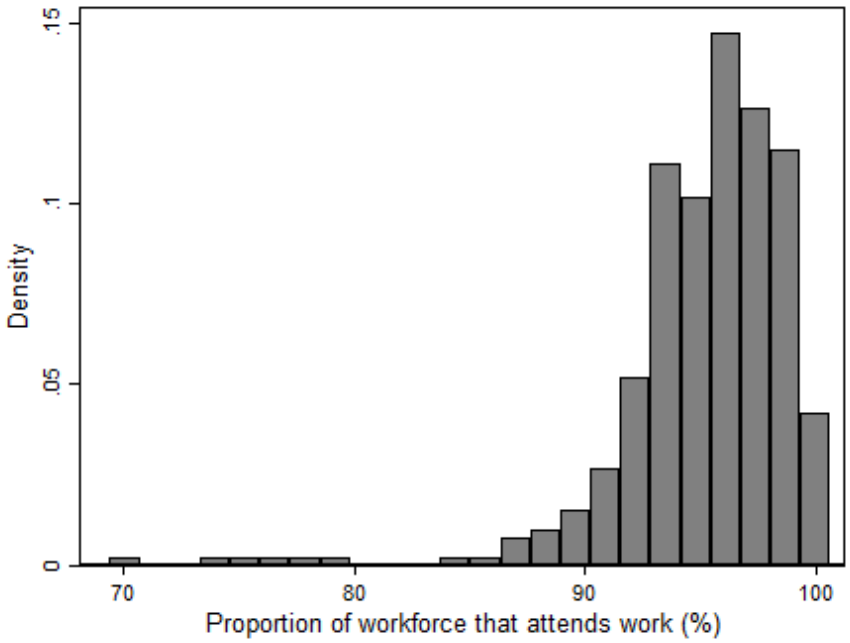


(b) Team 2's turnover

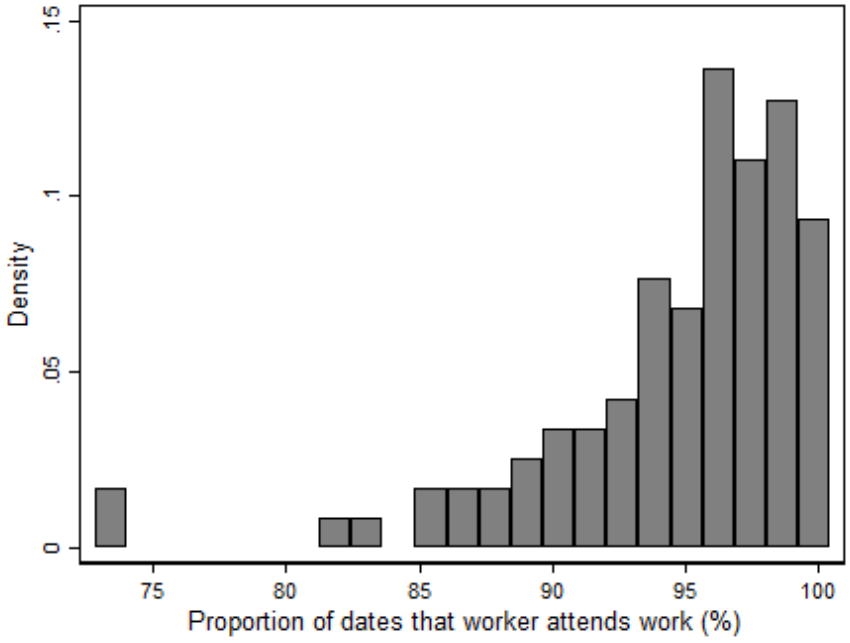


(c) Team 3's turnover

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

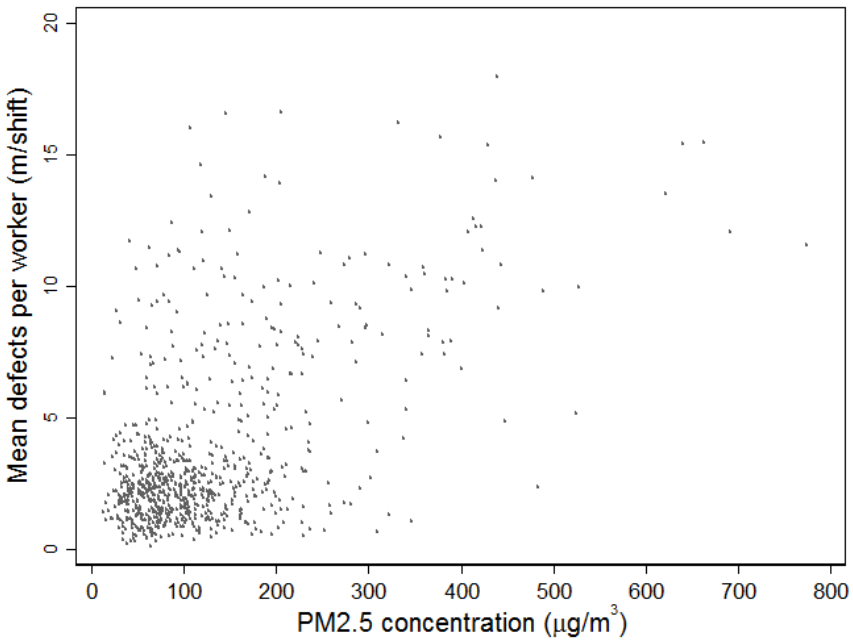


(a) Across days

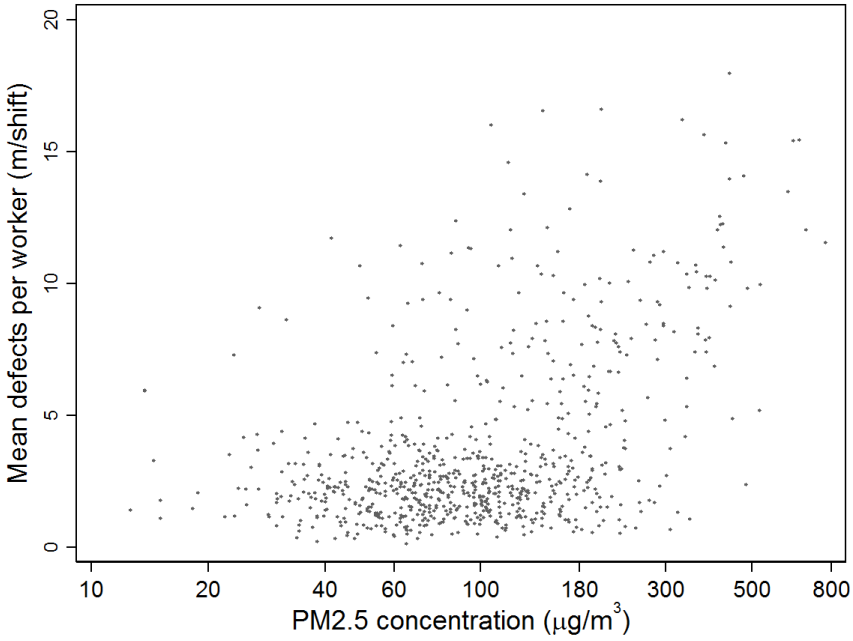


(b) Across workers

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)

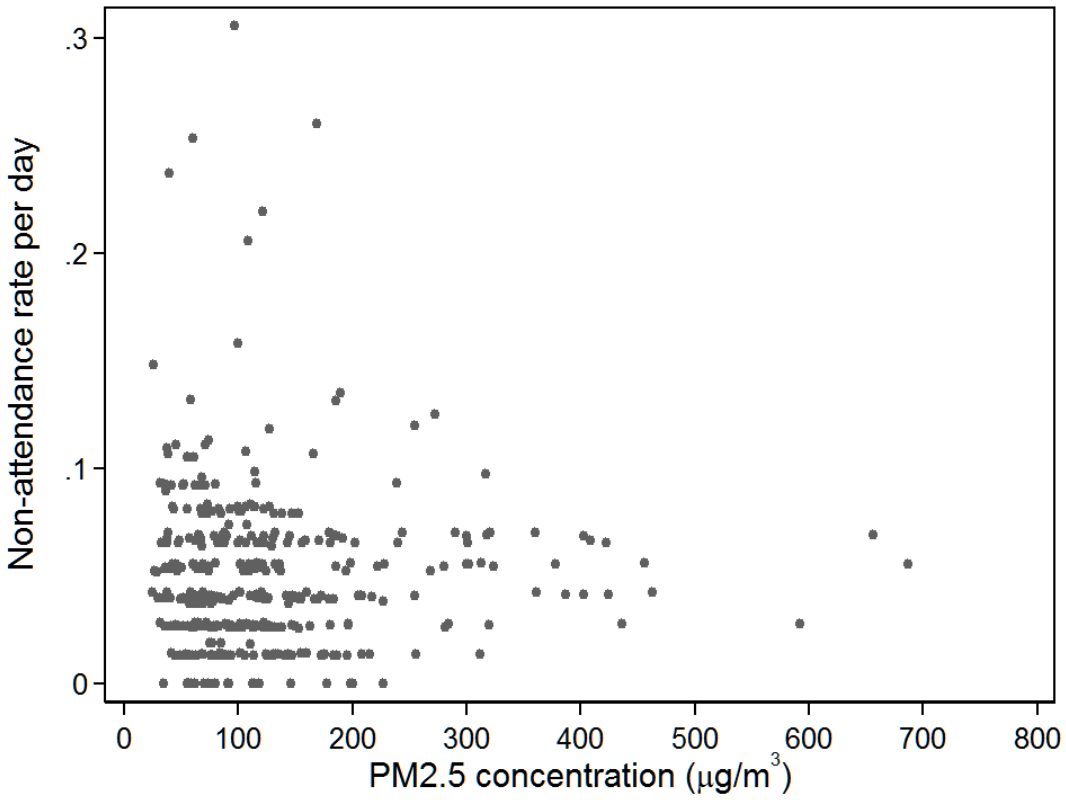
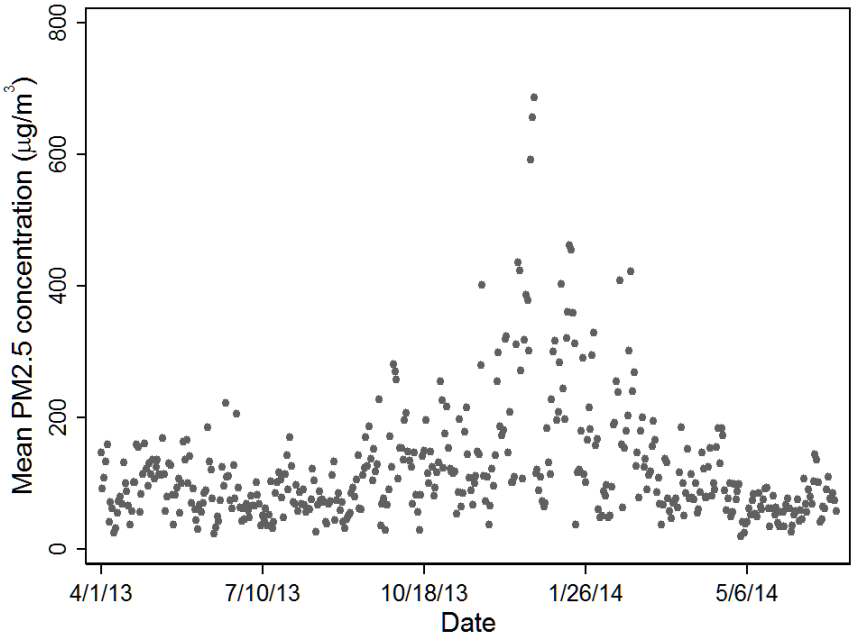
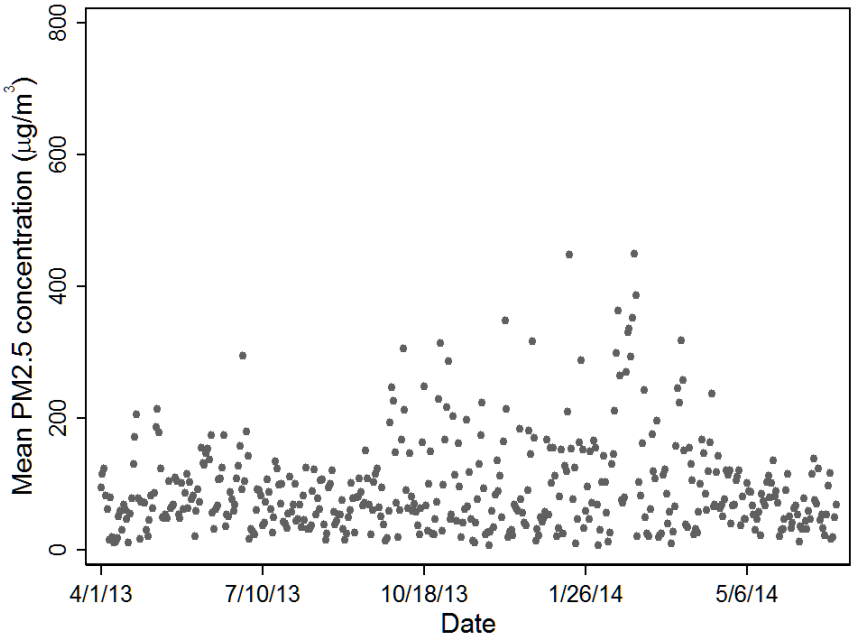


Figure A.9: PM2.5 concentration (μg/m^3, daily means) and the proportion of the workforce that does not attend work

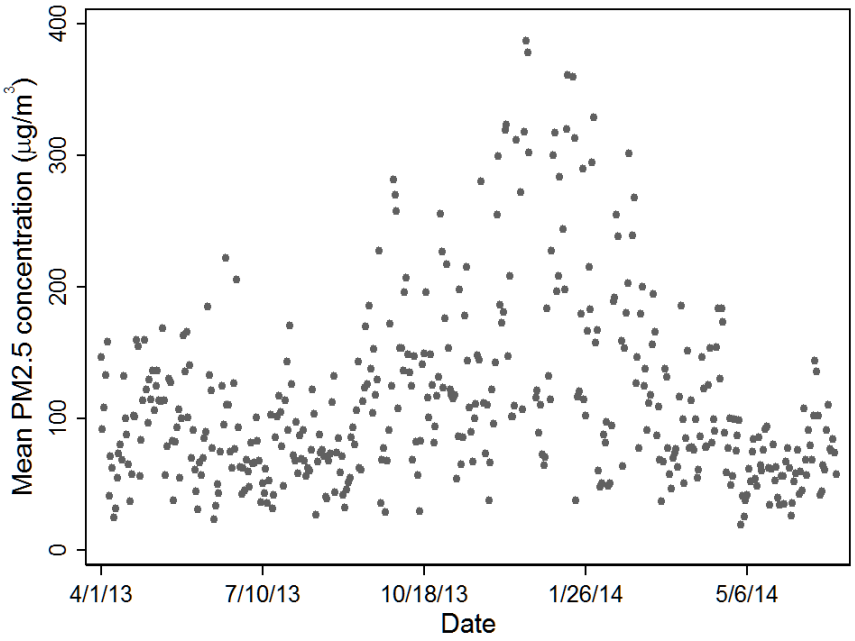


(a) 2 km from the plant

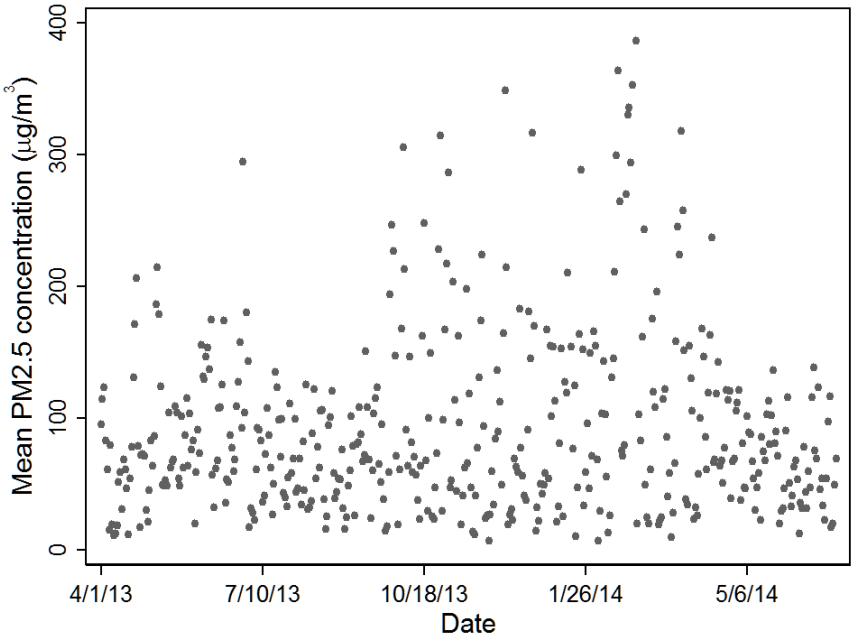


(b) In Beijing

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

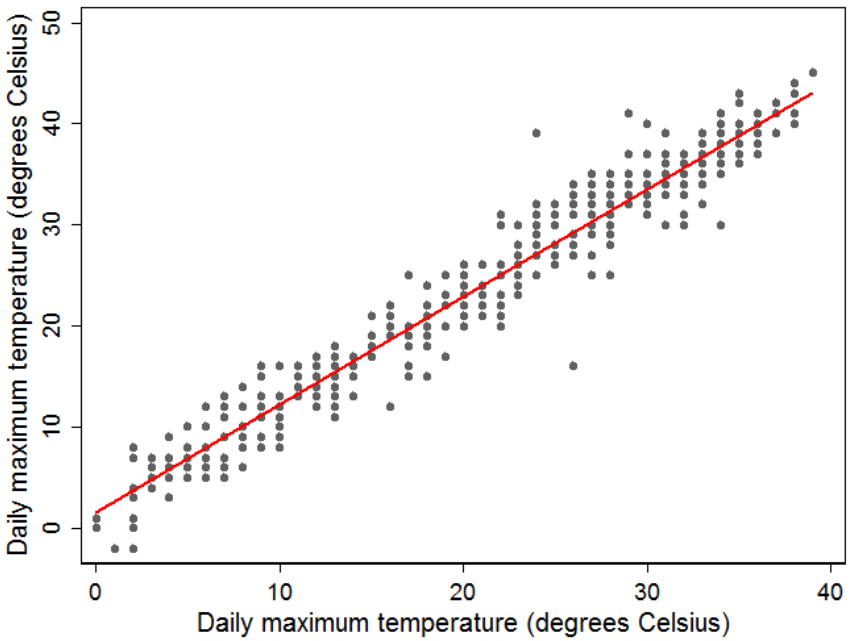


(a) 2 km from the plant

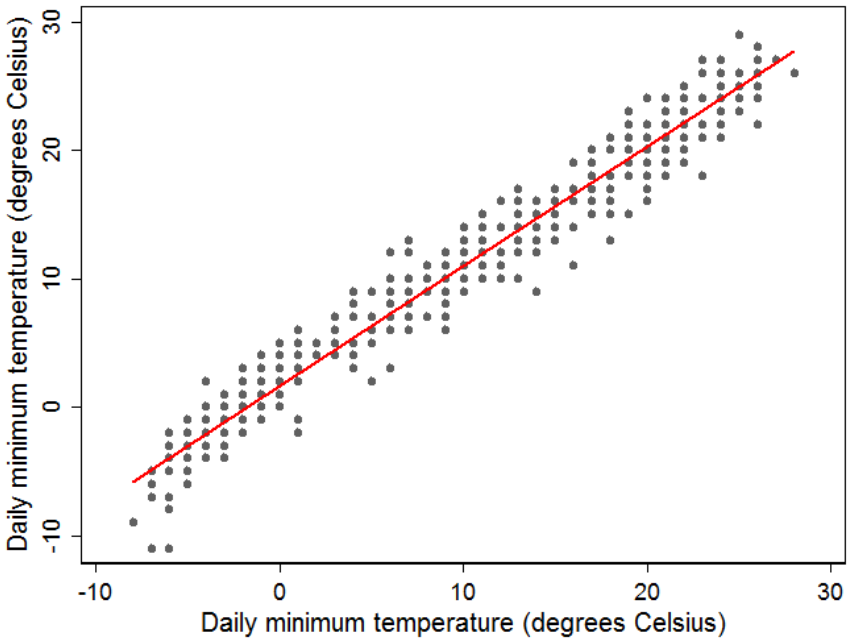


(b) In Beijing

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)

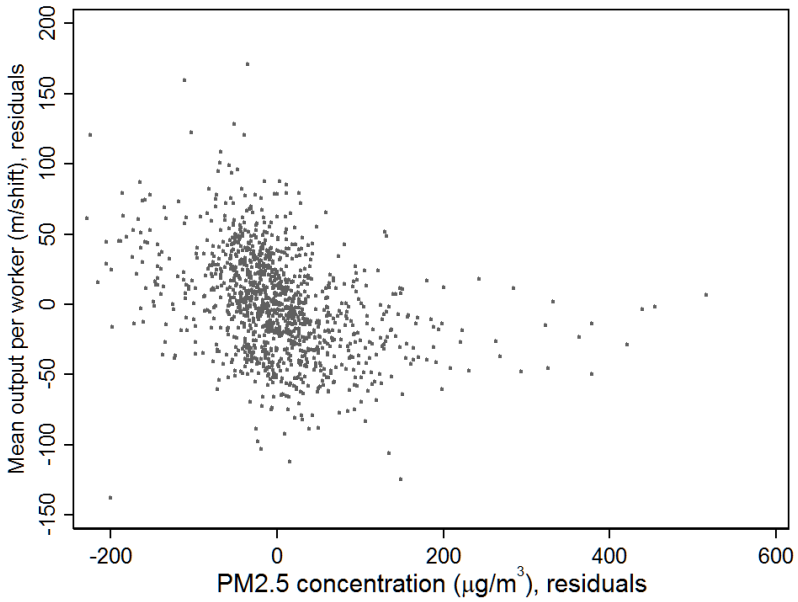


(a) Daily maximum

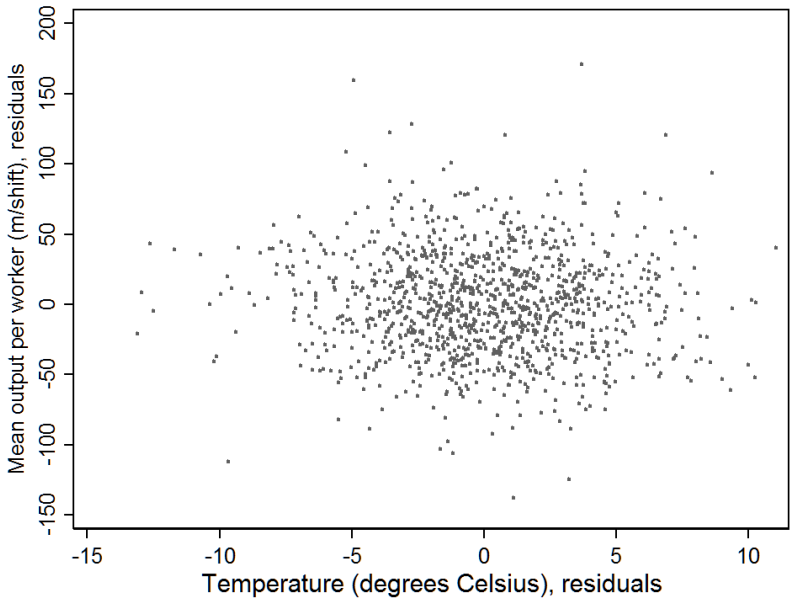


(b) Daily minimum

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

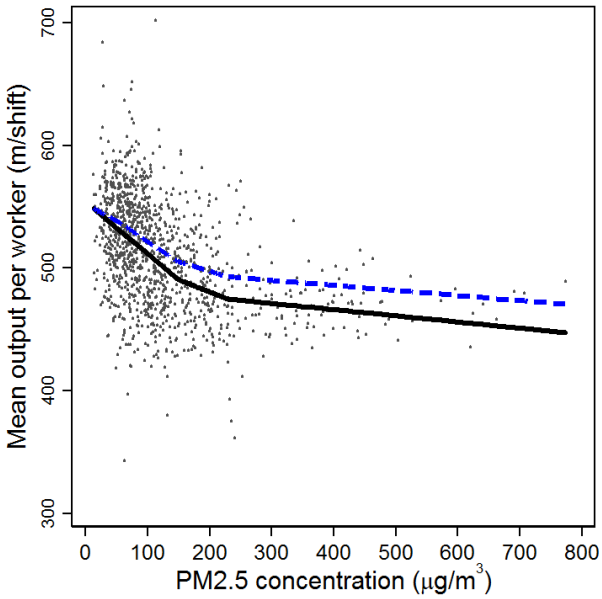


(a) Mean output (residuals) against PM2.5 level (residuals)

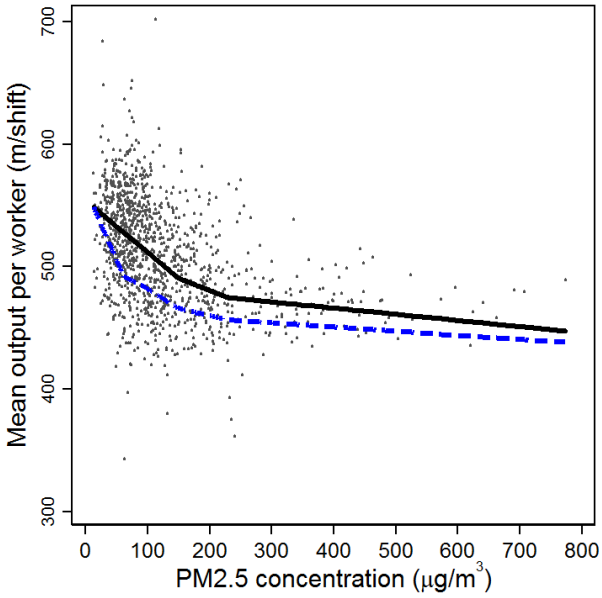


(b) Mean output (residuals) against temperature (residuals)

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).