

THE EFFECT OF POLLUTION ON WORKER PRODUCTIVITY: EVIDENCE FROM CALL CENTER WORKERS IN CHINA^{*}

Tom Chang
University of Southern
California

Joshua Graff Zivin
University of
California at San Diego
and NBER

Tal Gross
Columbia University
and NBER

Matthew Neidell
Columbia University
and NBER

June 2016

Abstract

We investigate the effect of pollution on worker productivity in the service sector by focusing on two call centers in China. Using precise measures of each worker's daily output linked to daily measures of pollution and meteorology, we find that higher levels of air pollution decrease worker productivity. These results manifest themselves at levels of pollution commonly found in large cities throughout the developing and developed world. This is the first study to demonstrate that the negative impacts of pollution on productivity extend beyond physically demanding tasks to indoor, white-collar work.

* We are grateful to Nick Bloom, Jennifer Cao, John Roberts, Jenny Zhichun Ying for help with the Ctrip data, with special thanks to James Liang, CEO of Ctrip. We are also grateful to Junjie Zhang and Shuang Zhang for help with the environmental data. We thank conference and seminar participants at Goethe University, Hong Kong University of Science and Technology IZA Conference on Labor Market Effects of Environmental Policies, the NBER Health Economics program meeting, Pampeu Fabra University, the University of Rotterdam, the University of Southern California, and the University of Stockholm for useful feedback.

1. Introduction

A growing body of evidence finds that pollution can reduce the productivity of workers in physically demanding occupations (Graff Zivin and Neidell, 2012; Chang et al., 2014; Heyes and Saberian, 2015).¹ Nevertheless, the importance of these occupations to aggregate economic performance in modern economies, where the service and knowledge sectors account for the majority of economic output, is quite modest. Thus, understanding whether pollution limits productivity in higher skilled, cognitively demanding professions is a question of tremendous economic importance. In this paper, we provide the first examination of the effect of pollution on workers whose primary tasks require cognitive rather than physical performance.

In particular, we investigate the effect of pollution on call-center workers using a unique panel dataset on the daily productivity of employees at a firm with offices in Shanghai and Nantong, China. This setting is important for at least three reasons. First, worker output at the call center is routinely monitored, providing precise measures of each worker's daily output (Bloom et al., 2014). Second, particulate matter (PM) pollution, which easily penetrates indoors (Thatcher and Layton, 1995; Vette et al., 2001), is both prevalent and highly variable in our study areas. Exposure could impair productivity through changes in cardiovascular and lung functioning (Seaton et al., 1995), irritation of the ear, nose, throat, and lungs (Pope, 2000), as well as through direct impacts on cognitive performance (Lavy et al., 2014).

Third, China is a large economy whose dramatic growth over the past several decades has been accompanied by equally dramatic increases in pollution. By nearly every available metric, China's environmental quality ranks among the lowest in the world, with unprecedented levels of air pollution. An effect on worker productivity would suggest that China's prioritization of industrial

¹ The Heyes and Saberian (2015) study of baseball umpires is particularly interesting in the context of our study. While it is a physically demanding occupation, cognitive performance (e.g. calling balls and strikes) plays a significant role in worker productivity.

expansion over environmental protection throughout the past few decades may have undermined some of the economic growth its policies were designed to achieve. The importance of strong environmental institutions to foster economic prosperity has implications for a wide range of countries that have yet to manage the pollution problems associated with urbanization and industrialization.

The call center we study is Ctrip, China's largest travel agency. Several aspects of the firm's operations allow us to credibly isolate the causal effect of pollution on the marginal product of labor. First, the workers have little discretion over their labor supply, a claim we confirm directly. This is important because such discretion could bias our productivity estimates due to changes in labor composition. If only the most productive workers come to work on polluted days, for example, we would underestimate the productivity effects of pollution. Second, daily variation in pollution in Shanghai and Nantong is plausibly unrelated to the firm's output. The firm serves clients throughout China, and, thus, demand for the firm's services on any given day is likely unrelated to *local* pollution on that particular day. Moreover, because calls are routed at random to the two locations, we can test directly whether such aggregate-level correlation is driving our results.

Finally, a brief pilot program at Ctrip allows us to probe the influence of one potential confounder: traffic. Traffic is a potentially important confounder because it generates pollution and can directly reduce productivity by both creating emotional stress and making employees late for work. We study a randomized experiment at Ctrip that allowed some employees to work from home for part of our study period. That experiment enables us to examine directly the importance of traffic in explaining our results.

Our analysis reveals a statistically significant, negative impact of pollution on the productivity of workers at the firm. A 10-unit increase in the air pollution index (API) decreases the number of

daily calls handled by a worker by 0.35 percent on average.² Productivity declines are largely linearly increasing in pollution levels, with statistically significant results emerging at an API above 100 for some measures of productivity and 150 for all of them. These estimates are robust to the inclusion of meteorology controls, worker-specific fixed effects, and a number of robustness checks that support the contention that we are estimating causal effects. Furthermore, when we decompose this effect, we find that the decrease in calls comes from an increase in the amount of time spent on breaks rather than from changes in the duration of phone calls, a finding consistent with Bloom et al. (2014), who found that breaks were the most malleable aspect of productivity at this firm.

To our knowledge, these results are the first evidence of an effect of pollution on white-collar labor. These impacts are economically significant. A back-of-the-envelope calculation suggests that even a very modest drop in air pollution could increase productivity in the Chinese service sector by billions of dollars per year. That consistently significant effects manifest themselves at an API of 150 also underscores that these impacts are not isolated to the most polluted cities in the developing world. Major metropolitan areas around the world, most of which employ considerably more non-manual labor, exceed that level with varying degrees of frequency. For example, Los Angeles, California, experienced 13 days with API greater than 150 in 2014, and Phoenix, Arizona, experienced 33 such days, with nearly half of those exceeding an air quality index of 200.³

The remainder of our paper is organized as follows. In the next section, we discuss the air quality index used to measure pollution in China and the potential mechanisms through which particulate matter—the principal driver of that index—can generate impacts on labor productivity.

² As detailed in the next section, particulate matter pollution is the driver of the API 98% of the time in our study setting.

³ The index uses a nonlinear pollutant-specific formula to convert each pollutant into a metric meant to capture its relative health risk to the public. The reported index is based on the highest value amongst all indexed pollutants on a given day. EPA Air Quality Index Report: http://www3.epa.gov/airdata/ad_rep_aqi.html

In Section 3, we describe the remainder of the data used for our analysis, and, in Section 4, we present our empirical strategy. Our results are presented in Section 5, and Section 6 concludes.

2. Background on Pollution

China, like most countries, releases a daily API, which is a composite measure of pollution that ranks air quality based on its associated health risks as a means to facilitate comprehensibility by the public. The index is cardinal, with higher values indicating poorer air quality. Index values less than 100 are generally deemed acceptable, those between 100 and 150 are labeled unhealthy for sensitive groups, and impacts on the general population begin to emerge at levels greater than 150. The API in China converts concentrations of three criteria air pollutants into a single index, using an algorithm developed by the U.S. Environmental Protection Agency (EPA, 2006). The pollutant that has the highest index, referred to as the primary pollutant, determines the API on a given day. In our dataset, for API levels of greater than 50, particulate matter smaller than $10 \mu\text{g}/\text{m}^3$ in diameter (PM_{10}) is the primary pollutant over 98 percent of the time.⁴ As such, we devote the rest of this section to providing a basic scientific background on particulate matter pollution.

Particulate matter consists of airborne solid and liquid particles that range considerably in size. PM is typically divided into fine PM, which includes all particles with a diameter of less than $2.5 \mu\text{g}/\text{m}^3$, and coarse PM, which includes all particles with a diameter between 2.5 and $10 \mu\text{g}/\text{m}^3$.⁵ Coarse PM penetrate into the lungs, while fine PM passes beyond the lung barrier to enter the bloodstream. Particles originate from a variety of sources and are largely the result of fossil fuel combustion, particularly when gases from power plants, industries, and automobiles interact. Given its diminutive size, PM can remain suspended in the air for extended periods of time and travel

⁴ A primary pollutant is not reported when air quality is measured as “good” (i.e., API levels of 0 to 50).

⁵ To put this in context, $10 \mu\text{g}/\text{m}^3$ is one-seventh the width of a human hair. Particulate matter greater than $10 \mu\text{g}/\text{m}^3$ is generally not believed to be harmful to human health.

lengthy distances.

A large body of epidemiological and toxicological literature suggests that exposure to PM harms health (see EPA [2004] for a comprehensive review). These risks arise primarily from changes in pulmonary and cardiovascular functioning (Seaton et al., 1995). Although extreme exposures (or more modest levels for sensitive populations) can be debilitating, even relatively modest levels of exposure may have an impact on worker productivity due to changes in blood pressure; irritation in the ear, nose, throat, and lungs; and mild headaches (Ghio et al., 2000; Pope, 2000; Sorenson et al., 2003). Symptoms can arise in as little as a few hours after exposure, but effects may also be triggered after several days of elevated exposure, particularly for people with existing cardiovascular and respiratory conditions.

Particles at the finer end of the spectrum are a particularly important concern for two reasons. First, the diminutive size of fine PM allows it to easily penetrate buildings (Ozkaynak et al., 1996; Vette et al., 2001), making exposure difficult to avoid, even for office workers. Second, fine PM is small enough to be absorbed into the bloodstream and even travels along the axons of the olfactory and trigeminal nerves into the central nervous system (CNS), where it can become embedded deep within the brain stem (Oberdörster et al., 2004). This, in turn, can cause inflammation of the CNS, cortical stress, and cerebrovascular damage (Peters et al., 2006). Greater exposure to fine particles is associated with lower intelligence and diminished performance over a range of cognitive domains (Suglia et al., 2008; Power et al., 2011; Weuve et al., 2012). Consistent with this epidemiological evidence, a recent study of Israeli teenagers found that students perform worse on high-stakes exams on days with higher PM levels (Lavy et al., 2014).

While the more dramatic health effects due to pollution exposure may lead to changes in labor supply, milder impairment of respiratory, cardiovascular, and cognitive function may reduce productivity on the intensive margin. Focus, concentration, and critical thinking are all essential

components of office-based job performance and depend heavily on a well-functioning brain and CNS. The goal of our analysis is to estimate the effect of pollution on the marginal product of labor, independent of any possible effects on the extensive margin of labor supply. Given the ubiquity of office work and the value it adds to global economic output, the welfare implications of any link between pollution and productivity in this setting are potentially enormous.

3. Data

Our data on worker productivity come from Ctrip International, China's largest travel agency. Ctrip's primary line of business involves making travel arrangements for clients. It earns revenue through commissions from hotels, airlines, and tour operators. In contrast to U.S. and European-based agencies that operate in markets with deep internet penetration, Ctrip conducted much of its business on the telephone during our study period. The firm was listed on the NASDAQ stock exchange in 2003 and its market capitalization was more than \$15 billion by the end of 2015.

Our analysis is focused on Ctrip's call-center workers in Shanghai and Nantong who book travel for clients throughout China.⁶ The offices are located in large multistory buildings, and are filled with cubicles and modern telecommunications and computing hardware. Equipment and staffing practices are identical at both offices, and both sets of workers follow the same procedural guidelines. A central server automatically routes customer calls and assignments to workers logged into the system, based on a computerized call-queuing system.

All workers in our sample receive compensation based partly on productivity. This productivity-based pay is primarily a function of call and order volumes, with additional adjustments made for call quality. As a result, Ctrip records various measures of the daily on-the-job performance of each of its call-center workers. Most relevant for our analysis are three measures: number of

⁶ Nantong is approximately 100 kilometers north of Shanghai.

phone calls handled per shift, number of minutes spent on the phone, and number of minutes logged in to the call center's computer system (i.e., the number of minutes within the work day that workers are available to handle calls).⁷ We also use data on worker absences and hours worked to assess potential changes in labor supply. Our data span the period from January 1, 2010, to December 9, 2012, with different coverage periods for the two offices.⁸

An additional feature of our data is that, during part of our study period, Ctrip performed a controlled experiment to analyze whether working from home affected worker performance. Approximately 250 employees from the Shanghai office worked from home during part of our study period, with a comparably sized control group. The experiment ran from December 6, 2010 until August 14, 2011, during which time the home-based workers were provided with the necessary equipment to allow them to perform their usual work responsibilities in a manner identical to that at the office. In Section 4, below, we exploit this experiment to test whether pollution affects workers even when they do not commute to the office. More details on call-center operations and the work-from-home experiment can be found in Bloom et al. (2014).

Our daily pollution data are obtained from the China National Environmental Monitoring Center (CNEMC), which is affiliated with the Ministry of Environmental Protection of China. These data provide a measure of pollution that is based on the average API score across all monitors within a city. During our study period there were 10 operating national pollution monitors in Shanghai and 5 in Nantong. Daily API measures are calculated in 4 steps: 1) a 24-hour average pollution concentration is calculated at the station level; 2) a city average is calculated from the

⁷ Call quality was assessed based on a 1-percent sample of recorded call transcripts that were audited and scored by an external review team. Unfortunately, we do not have access to this data.

⁸ The difference in data coverage periods is due to a combination of pollution data availability and the time coverage of the productivity data provided by the firm. The Shanghai data sample runs from January 1, 2010, to August 14, 2011. The Nantong sample runs from September 1, 2011, to December 9, 2012.

station averages; 3) the city average is converted to API based on pollutant specific formulas; and finally 4) the overall API is defined as the max of individual pollutant APIs.⁹

Although discrepancies between CNEMC and data from the U.S. embassy in Beijing have called into question the reliability of the CNEMC, recent work finds no evidence of manipulation in our study regions (Ghanem and Zhang, 2014). To the extent that it may still exist, such manipulation should be unrelated to demand for travel services. As can be seen in Figure 1, the distribution of pollution in both Shanghai and Nantong are continuously distributed, with no evidence of displacement around important cutoffs, such as when the API reaches 100.

Table 1 presents simple descriptive statistics for our sample. The average API in the merged sample is 66, with an average value of 65 in Shanghai and 69 in Nantong. In comparison, the median air quality index in New York City was 55 in 2014.¹⁰ Our study includes 393 workers from Shanghai and 4,499 workers from Nantong, with this difference due to distinct sampling frames for workers across facilities as well as differences in the availability of environmental data. The Shanghai sample follows the subset of workers who participated in the work-from-home experiment from January 1, 2010, to August 14, 2011, while the Nantong sample includes all worker observations from September 1, 2011, to December 9, 2012. Importantly, one of the inclusion criteria for the work-from-home experiment was having at least six months tenure with the firm. Thus, the Shanghai sample, but not the Nantong sample, leaves out workers with short tenure at the firm. As a result, the number of worker-day observations between the two locations is more balanced than the number of workers at each office might suggest. Similarly the difference in coverage periods for the two cities explains the weather differences between the two samples: the worker-days covered in Nantong are more likely to have occurred during winter than those in Shanghai.

⁹ Since citywide measures of pollution are likely to be a noisy measure of the exposure of individual workers in our sample, our estimates may be biased downward.

¹⁰ Figures are derived from the EPA Air Quality Index Report, which can be found at: http://www3.epa.gov/airdata/ad_rep_aqi.html. Last accessed on February 9, 2016.

Figure 2 describes the distribution of worker productivity. Panel A plots the distribution of calls per day, which appears to be normal distributed. On a typical day, workers handle an average of 66 calls. Panel B describes the average productivity across workers. There are two peaks in the data corresponding to two distinct types of worker tasks.¹¹ As we describe below, Figure 4 shows that the impacts across productivity quantiles is essentially flat, suggesting that all worker types are impacted similarly by pollution.

4. Econometric model

Our goal is to estimate the effect of pollution on worker productivity. We estimate the following hybrid production function:

$$\log(y_{ijt}) = \beta \times \text{API}_{jt} + X_{jt}'\gamma + \delta_t + \alpha_i + \varepsilon_{ijt}.$$

The outcome y_{ijt} is the measure of daily productivity for each worker, either total number of calls per shift or the number of minutes logged in to the call center’s computer system. Because workers clock out of the system when they are not available to field calls, this latter measure captures the length of a worker’s breaks taken during the workday. The variable API is a daily average measure of the air pollution index in city j , where β captures the effect of API on our two measures of productivity. Although we specify API linearly here, we also include specifications with a series of API indicator variables to capture potential nonlinear effects of pollution on the outcome of interest. The vector X_t includes temperature, the only other covariate we consistently observe at both locations. The parameter δ includes day-of-week and year-month indicator variables to account for trends within the week and over time. The variable α_i is a worker-specific fixed effect for those

¹¹There are two types of workers: workers that call customers directly, and workers that take calls from customers. The workers that take calls from customers make relatively few calls per day, whereas the workers that make calls to customers make many calls per day. Those two groups lead to the “bimodal” distribution in Panel B of Figure 2. Unfortunately, we only observe worker type for a subset of Shanghai workers. For those workers, we have confirmed that the overall productivity effects of pollution are similar.

specifications that rely only on within-worker variation in pollution exposure to identify impacts on productivity. Because the error term ε likely exhibits auto-correlation between observations based on the same worker or for all workers on the same day of PM measurement, we allow for two-way clustering (Cameron et al., 2011) along those dimensions.

There are several challenges to uncovering a causal effect of pollution on worker productivity. First, individuals can sort into locations based on the amount of pollution in that area, leading to non-random assignment of pollution. Moreover, within a location, an individual's labor supply might change in response to pollution. Although the high-frequency-panel nature of our data allows us to overcome the issue of geographic selection—workers are unlikely to sort on daily fluctuations in pollution levels—sample selection due to pollution-induced changes in labor supply remains a potential concern. For example, workers who are most susceptible to pollution may be those least likely to work on more polluted days, thus biasing our estimates of worker productivity if the most susceptible workers also differ in their average level of productivity. Similar concerns regarding worker composition apply if the most susceptible workers shorten their workday in response to pollution. Because workers have very few sick days and limited discretion over their work hours, we suspect that extensive-margin responses of this sort are unlikely in our sample. We explicitly test this claim by examining whether API relates to the probability of working and the number of hours worked (using the regression equation above) and find no effect of API levels on either measure of labor supply.

Another threat to identification lies in the existence of a local factor that both creates local pollution and has a direct effect on worker productivity. Indeed, if firms are directly responsible for much of the local air pollution, a naïve analysis may erroneously conclude that pollution increases productivity. Although this may be a substantial concern for pollution-intensive firms, such as those

engaged in heavy manufacturing, this type of direct linkage between firm-level pollution and productivity is unlikely at the call center that we study or the service sector, more generally.

Of course, indirect linkages also can confound inference in this type of analysis. While it is impossible to rule out all potential factors that may generate a spurious relationship between pollution and productivity, the most salient concern relates to vehicular traffic. The transportation sector is a major contributor to urban air pollution (Fenger, 1999), and increases in roadway congestion can simultaneously increase air pollution and commute time for workers.

As described earlier, our analysis of extensive-margin changes in hours worked does not find any evidence for commute-related changes in the length of the workday. Productivity changes on the intensive margin remain a concern if the stress and aggravation from an arduous commute alter performance on the job. To test whether such a traffic channel confounds the results, we exploit a unique feature of our data. For part of the period for which we have data, a subset of workers was randomly assigned to work from home (Bloom et al., 2014). Because employees who worked at home had no commute, their productivity should be unaffected by traffic, thus allowing us to isolate the pollution impacts of interest.

A final concern regarding identification relates to demand-mediated effects. If higher levels of pollution are the result of increased economic activity and travel is a normal good, higher pollution levels could lead to greater demand for Ctrip's services. To the extent that worker slack exists, such changes in demand could conceivably lead to changes in worker productivity. As the largest travel agency in China, Ctrip draws its clients from a broad geographic base both across mainland China and internationally. This fact greatly minimizes concerns that our measured impacts of local air pollution on productivity are driven by demand-side factors. To further bolster this claim, we also test whether a common demand-level shock is driving our results by regressing

productivity in Shanghai on Nantong pollution and vice versa, as well as the effect of Beijing pollution, the most polluted city in China, on productivity in Shanghai and Nantong.

5. Empirical Results

This section presents our empirical results. We begin by exploring the impacts of pollution on the extensive margin through changes in labor supply. This is followed by an analysis of the effect of pollution on the intensive margin through changes in productivity as well as a decomposition of this effect. Finally, we present a series of robustness checks.

5A. The Extensive Margin: Labor Supply

Our first task is to assess whether pollution affects labor supply in this setting. Table 2 presents these results. The first column of Panel A presents results from a regression in which the dependent variable is an indicator variable for whether a worker comes to work on a given day. To address planned vacations, we drop absence spells of more than 5 consecutive days from the sample. We find a small, statistically insignificant relationship between pollution and the probability of working, shown in column (1). The third column of Panel A repeats the analysis with a focus on shift length for the subsample of workers on which we have detailed timestamp records.¹² We find a similarly insignificant effect of pollution on hours worked. Both coefficients remain largely unchanged by the addition of worker-specific fixed effects, presented in columns 2 and 4, respectively.

We can also rule out very small effects based on these estimates. For example, the lower 95-percent confidence interval for the probability of working is 0.17 percent, while the estimate for hours worked (in the fixed-effects specification) implies a 0.5-minute decrease in time at the office.

¹² We also drop worker-day observations for which the implied shift length is implausible: less than 15 minutes, longer than 12 hours, or shorter than the recorded amount of time spent on the phone for that shift.

Both estimates are sufficiently small that we are confident that we can conclude that labor supply is unaffected by air pollution.

Panel B of Table 1 repeats this exercise with indicator variables for API between 50 and 100, 100 and 150, 150 and 200, and 200 plus (with API less than 50 as the reference group). That specification explores potential non-linear effects of pollution on labor supply. As in the linear case, we find no evidence of pollution impacts on the extensive margin.

This absence of an effect on labor supply is consistent with the institutional realities of the Ctrip workplace, as described earlier. It also implies that our estimates of the effect of pollution on productivity that follow are not affected by sample-selection bias.

5B. The Intensive Margin: Worker Productivity

Table 3 presents estimates of our main regression specification.¹³ The first column of Panel A suggests that a 10-unit increase in API decreases the number of calls per shift by 0.35 percent, an effect statistically significant at the 1-percent level. Adding worker-specific fixed effects, shown in column 2, does not appreciably change our estimate. Columns 3 and 4 present estimates of the same regression, but with the logarithm of minutes-on-the-phone per shift as the dependent variable. Similar to the effect on the number of phone calls, we find that minutes logged-in decreases in response to pollution. Column 3 implies that a 10-unit change in API reduces minutes logged in by 0.25 percent, a finding that is statistically significant and robust to the inclusion of worker-specific fixed effects.

Panel B of Table 3 repeats this analysis using the same flexible specification of API as in the previous table. Across columns we find a similar pattern of coefficients. Specifically, we find that

¹³ The decrease in sample size in Table 3 relative to the extensive-margin regressions in columns 1 and 2 of Table 2 is due to the fact that the extensive-margin regressions include worker-day observations for days on which the worker does not work (i.e., the average worker works just under 5 days a week).

productivity monotonically decreases in pollution levels, with statistical significance emerging for minutes logged in at API levels exceeding 100, and near statistical significance for impacts on the number of calls. All results are statistically and economically significant for pollution levels exceeding 150. To compare to the linear estimates, we plot the nonlinear coefficients alongside the linear effects calculated for the midpoint of each bin, shown in Figure 3, based on the cross-sectional estimates for the number of calls. As evident in the figure, the more-flexible estimates align quite closely with the linear estimates. The divergence in the highest bin reflects that the linear estimate is based on an API of 225, while the bin includes all API measures above 200. In general, these estimates suggest an approximately linear effect of pollution on worker productivity.

Based on the results in Table 3, it is unclear whether the decline in the number of calls handled by workers over the course of their shift is driven exclusively by the observed reductions in the amount of time logged into the call system or the duration of those calls. Table 4 presents results that decompose these two channels. The first column examines the impact of pollution on the total amount of minutes spent on the phone per shift. Column 3 examines the impact of pollution on the amount of time spent on each call. The effect of pollution on time on the phone appears to be driven entirely by time logged in and thus the availability of workers to handle calls; time spent on each call is unaffected by pollution. These results are robust to the inclusion of worker-specific fixed effects (columns 2 and 4) and follow a similarly non-linear pattern as those presented in Table 3. Thus, it appears that the productivity effects of pollution revealed in Table 3 are driven by workers taking more frequent or longer breaks. This finding is consistent with the results of Bloom et al. (2014), who found that breaks were the most affected aspect of productivity in this study population.

To test whether the effect of pollution is heterogeneous in baseline worker productivity, we estimate our equation for the two main productivity outcomes using quantile regressions. The

coefficients from these regressions along with their 95% confidence intervals are shown in Figure 4. The figure clearly shows that, with the exception of the lowest decile, which has by far the largest confidence interval, the coefficients are nearly identical across quantiles for both our main measures of productivity. These results indicate that the effect of air pollution on productivity is remarkably stable across workers, equally affecting both low and high productivity types and thus workers performing different tasks for the firm (e.g. accepting vs. placing calls) who vary in their average levels of productivity.

While all of the analyses thus far have assumed a contemporaneous effect of pollution on worker productivity, the effect of pollution may accrue over time. As such, we explore a more dynamic regression specification in which we also include lagged measures of pollution. We also include a lead of pollution in this specification as a falsification test; future pollution should not affect current productivity. The results from this exercise are summarized in Figure 5, which presents estimates from a regression of calls per day (Panel A) and total minutes logged in (Panel B) on contemporaneous pollution, three lags, and one lead, while controlling for all of the covariates included in our main regressions above. The effect from contemporaneous pollution is statistically significant and larger than the effect from the lags, with only the three-day lag on total minutes on the phone significant at the 5-percent level. Since all other lagged effects are small and statistically insignificant, we view this evidence as supportive of the notion that pollution effects are rather immediate. Reassuringly, the coefficient on future pollution is statistically insignificant and smaller in magnitude than the contemporaneous coefficient.

5C. Robustness Checks

As discussed in the previous section, two remaining concerns could threaten the internal validity of the analysis. First, traffic may be a source of both air pollution and a factor that affects worker

productivity. Second, the impacts that we find on worker productivity may be the result of changes in the demand for Ctrip services. Since Ctrip serves clients across China, this is tantamount to a concern that local pollution near the call center is correlated with national pollution levels that shape aggregate demand for travel. We discuss each of these in turn.

In order to assess the role of traffic, we analyze the effect of air pollution on workers that participated in a randomized-control trial to test the effects of working from home. That experiment allows us to measure the impact of pollution on workers who worked from home (and therefore did not have to commute to work) compared to the control group (those who volunteered for the experiment but worked from the office) for the duration of the experimental period (December 6, 2010 through August 14, 2011). The first four columns of Table 5 present the results of this analysis. Columns 1 and 2 present the effect of pollution on the logarithm of the number of phone calls taken by workers who worked from the office and at home, respectively. In both cases, the coefficient is negative and statistically significant at conventional levels, and comparable to estimates in Table 2. Columns 3 and 4 present the effect of pollution on the logarithm of minutes logged in to the system for workers who worked from the office and their homes, respectively. The coefficient is again negative but in this case neither is statistically significant at conventional levels, likely a reflection of our greatly reduced sample during this experimental period.¹⁴ That the point estimates for the workers in the office are not statistically significantly different from those that work at home suggests that traffic is unlikely to be a source of confounding for the productivity effects described in the previous section.

Table 6 explores the possibility that the observed effect of pollution is driven by demand-side factors. To test for such a possibility, we repeat our main analysis with controls for the API in

¹⁴ Note that the increase in standard errors for the analysis on this subset of workers is likely due to a combination of the much smaller sample size and an increase in classical measurement error. Given the long commutes of many of the workers, their amount of pollution exposure may be measured with more error (see Bloom et al. (2014) for details regarding the characteristics of the workers in the WFH experiment).

the other city: we add data on the pollution in Nantong for Shanghai workers and pollution in Shanghai for Nantong workers. Thus, the coefficient on other API will reflect the degree to which Nantong pollution affects productivity in Shanghai, and vice-versa. A priori, we would not expect to see a statistically significant effect from pollution in a different location unless they were driven by a common factor correlated with aggregate demand. As an additional test of demand side factors, we also include data on the API from Beijing. In all cases, we find that the coefficient for pollution from a different location is much smaller and statistically insignificant, while the coefficient for own-location pollution remains statistically significant and largely unchanged.

As a final robustness check, Table 7 repeats our core analysis separately by city. Since calls are randomly routed to each location, we should not expect to see significant differences by location. Reassuringly, the table suggests qualitatively similar results: no effect of pollution on labor supply, and a negative relationship between pollution and productivity. While the point estimates on productivity appear slightly larger for the Nantong sample, they remain statistically indistinguishable from those in Shanghai, providing additional support for our pooling strategy.

6. Conclusion

In this paper, we analyze the relationship between air pollution and the productivity of individual workers at a large call center in China. We find that a 10-unit increase in the air pollution index decreases the number of daily calls handled by a worker by 0.35 percent. Our analysis also suggests that these productivity losses are largely linearly increasing in pollution levels. To our knowledge, this is the first evidence that the negative impacts of pollution on worker productivity extend to labor markets beyond those centered on physically demanding labor.

To place these results in context, it is instructive to compare them to productivity estimates found in prior research that focused on other economic sectors. Graff Zivin and Neidell (2012)

found an elasticity of 0.26 in an agricultural setting, and Chang et al. (2016) found an elasticity of 0.08 in a manufacturing setting. Our estimates from column 1 of Table 3 imply an elasticity of 0.003. The smaller estimates for call center workers is consistent with laboratory evidence that suggests that the increased respiratory rates associated with physical activity can exacerbate the effects of pollution (Lippmann et al., 2003). Measurement error may also be a factor that is biasing our estimates downward, since China only releases citywide measures of pollution rather than those from individual pollution monitors that may be more proximate to the firm.

While our point estimates are smaller than those for other sectors, the impacts for China and other rapidly industrializing countries with sizable service sectors and persistent pollution problems are profound. A simple back-of-the-envelope calculation may be useful to fix ideas. The coefficient from our simple linear specification in Column 1 of Table 3 implies that a 10-unit change in the API translates into a 0.35 percent change in daily productivity. If we assume that this effect applies to all service-sector workers in China, an across the board 10-unit reduction in national pollution levels would increase the monetized value of worker productivity by more than \$2.2 billion US per year.¹⁵

That statistically significant effects emerge in some dimensions when the API exceeds 100 and for all outcome measures at an API of 150 suggests that this is not simply an issue for the world's most-polluted cities, since such levels obtain with some frequency in urban environments around the world. Given the size of the service and knowledge sectors in the developed world, and the relatively high levels of labor productivity within them, even very small impacts from pollution could aggregate to rather substantial economic damages. The case of Los Angeles is illustrative. In 2014, the air quality index exceeded the EPA standard on 90 days. If all of those days were brought into regulatory compliance, service sector productivity in the county of Los Angeles would have

¹⁵ The service sector is defined to include: Hotels, IT, Financial Intermediation, Real Estate, Business Services, Science, Household Services, Education, Health, Sports, and Public Management. According to the National Bureau of Statistics of China (2014), the total urban wage bill in these sectors totaled 3,958,960 million Yuan in 2013. At 6.19 Yuan to the dollar in 2013, that translates into \$639,243,040,000. Multiplying that figure by 0.35 percent yields \$2,429,123,552.

been \$374 million larger.¹⁶ The sum of these impacts across all major metropolitan areas in the US would be substantially higher.

While call-center work is mental rather than physical work, it is important to emphasize that it remains a semi-skilled occupation. If our measured productivity impacts are the result of diminished cognitive function, the negative impacts of pollution may well be larger for high-skilled occupations that form the backbone of the service and information economy. The development of suitable measures of productivity in those occupations and assessing its relationship with environmental quality represents a fruitful area for future research.

¹⁶ The EPA air quality standard is 100. Los Angeles County exceeded this figure on 90 days for a total of 1895 excess API points relative to the standard. The total service sector wage bill in Los Angeles County in 2014 was \$205,620,109,131 (US Bureau of Labor Statistics, 2014), which translates into a daily service sector wage of \$563,342,765. Multiplying the excess API points by the 0.35 percent coefficient per 10 API point change and the daily wage for the 90 days in violation implies a total productivity effect of \$373,637,089.

References

Bloom, Nicholas, James Liang, John Roberts and Zhichun Jenny Ying (2014). "Does working from home work? Evidence from a Chinese experiment." *Quarterly Journal of Economics*, forthcoming.

Calderon-Garciduenas, Lilian, et al. "DNA damage in nasal and brain tissues of canines exposed to air pollutants is associated with evidence of chronic brain inflammation and neurodegeneration." *Toxicologic pathology* 31.5 (2003): 524-538.

Cameron, Colin, Jonah Gelbach, and Douglas Miller (2011). "Robust inference with Multi-Way Clustering." *Journal of Business and Economic Statistics*, 29(2).

Chang, Tom, Graff Zivin, Joshua, Gross, Tal and Matthew Neidell (2014). "Particulate Pollution and the Productivity of Pear Packers," NBER Working Paper 19944.

Dalia Ghanem and Junjie Zhang (2014). "Effortless Perfection: Do Chinese Cities Manipulate Air Pollution Data?" *Journal of Environmental Economics and Management*, 68(2).

Fenger, Jes. "Urban air quality." *Atmospheric environment* 33.29 (1999): 4877-4900.

Graff Zivin, Joshua and Matthew Neidell (2012). "The Impact of Pollution on Worker Productivity," *American Economic Review*, 102(7).

Lavy, Victor, Avraham Ebenstein, and Sefi Roth (2014). "The Impact of Short Term Exposure to Ambient Air Pollution on Cognitive Performance and Human Capital Formation." NBER Working Paper 20648.

Lippmann, Morton, et al. "The US Environmental Protection Agency Particulate Matter Health Effects Research Centers Program: a midcourse report of status, progress, and plans." *Environmental Health Perspectives* 111.8 (2003): 1074.

National Bureau of Statistics of China, *China Statistical Yearbook* (2014):
<http://www.stats.gov.cn/tjsj/ndsj/2014/indexeh.htm>.

- Oberdörster, Günther, et al. "Translocation of inhaled ultrafine particles to the brain." *Inhalation toxicology* 16.6-7 (2004): 437-445.
- Peters, Annette, et al. "Translocation and potential neurological effects of fine and ultrafine particles a critical update." *Part Fibre Toxicol* 3.13 (2006): 1-13.
- Power, Melinda C., et al. "Traffic-related air pollution and cognitive function in a cohort of older men." *Environmental health perspectives* 119.5 (2011): 682.
- Suglia, S. Franco, et al. "Association of black carbon with cognition among children in a prospective birth cohort study." *American journal of epidemiology* 167.3 (2008): 280-286.
- Thatcher, T. L., & Layton, D. W. (1995). "Deposition, resuspension, and penetration of particles within a residence." *Atmospheric Environment*, 29(13), 1487-1497.
- U.S. Bureau of Labor Statistics, Quarterly Census of Employment and Wages, 2014.
- Vette, A. F., Rea, A. W., Lawless, P. A., Rodes, C. E., Evans, G., Highsmith, V. R., & Sheldon, L. (2001). "Characterization of indoor-outdoor aerosol concentration relationships during the Fresno PM exposure studies." *Aerosol Science & Technology*, 34(1), 118-126.
- Weuve, Jennifer, et al. "Exposure to particulate air pollution and cognitive decline in older women." *Archives of Internal Medicine* 172.3 (2012): 219-227.