Discussion Papers in Economics

The Impact of Temperature on Productivity and Labor Supply: Evidence from Indian Manufacturing

E. Somanathan, Rohini Somanathan, Anant Sudarshan, Meenu Tewari

November 2014

Discussion Paper 14-10

Indian Statistical Institute, Delhi
Economics and Planning Unit
7, S. J. S. Sansanwal Marg, New Delhi 110016, India

The Impact of Temperature on Productivity and Labor Supply: Evidence from Indian Manufacturing

E. Somanathan, Rohini Somanathan, Anant Sudarshan, Meenu Tewari*

November 10, 2014

*Sudarshan (corresponding author): University of Chicago and IFMR, anants@uchicago.edu; Tewari: University of North Carolina, Chapel Hill and ICRIER, mtewari@unc.edu; E. Somanathan: Indian Statistical Institute, som@isid.ac.in, Rohini Somanathan: Delhi School of Economics, rohini@econdse.org. Acknowledgements: We thank ICRIER, the PPRU of ISI Delhi, and the Rockefeller Foundation for financial support. For comments that have improved this paper we thank Michael Greenstone, Ben Jones, Kyle Meng, Solomon Hsiang, Rohini Pande, Geoff Heal, Jisung Park, Christos Makridis, M. Mani and seminar participants at the NBER Summer Institute 2014, NEUDC 2013, the Indian School of Business and the Indian Statistical Institute. We also thank Kamlesh Yagnik, the South Gujarat Chamber of Commerce and Industry and especially Anant Ahuja. Mehul Patel provided important field assistance.

Abstract

This paper shows that high temperatures may reduce manufacturing output by lowering worker productivity via heat stress. Using an annual panel of manufacturing plants in India, and daily primary micro-data from case-study firms, we find that (i) output in labor-intensive settings decreases at high temperatures by about 3 percent per degree Celsius (ii) workplace climate control may enable adaptation and (iii) sustained heat may reduce worker attendance. These mechanisms might contribute to the negative correlation between temperature and aggregate output observed in poor countries. Failing to account for reduced labor productivity may underestimate the costs of climate change.
Keywords: temperature, heat stress, worker productivity, manufacturing, climate change.
JEL: Q54, Q56

1 Introduction

Extreme events excepted, the economic impact of global warming has been thought to operate mostly through its effect on agricultural output. The Fifth Assessment Report of the Intergovernmental Panel on Climate Change (Field et al., 2014) acknowledges that "Few studies have evaluated the possible impacts of climate change on mining, manufacturing or services (apart from health, insurance, or tourism)".[1] Understanding the impact of high temperatures on human behavior and economic performance is especially important for poor and middle income countries where populations are less protected, warmer temperatures more extreme and rapid urbanization has created heat islands in which a rising share of the population resides (Mohan et al., 2012; Zhao et al., 2014).

[1]The previous Fourth Assessment Report (Working Group II) on Impacts, Adaptation and Vulnerability, made the even stronger claim that "*Climate-change vulnerabilities of industry, settlement and society are mainly related to extreme weather events rather than to gradual climate change (very high confidence).*"

The effects of climate on agricultural output have been extensively studied. High temperatures are associated with lower yields of specific crops (Lobell, Schlenker and Costa-Roberts, 2011; Schlenker and Roberts, 2009; Mendelsohn and Dinar, 1999; Auffhammer, Ramanathan and Vincent, 2006). Yet agriculture alone seems insufficient to explain the intriguing negative correlation between temperature and aggregate economic output, which is observed in countries with both large and small agricultural sectors. Recent studies have documented that temperature-output associations are observed using country-level aggregates for non-agricultural sectors as well. Dell, Jones and Olken (2012) using a cross-country panel, find reductions in both agricultural and non-agricultural output for poor countries in years with higher than average temperatures. Similarly, for countries in Central America and the Caribbean, Hsiang (2010) finds negative temperature effects in the services sector that exceed impacts on agricultural output.

Isolating specific mechanisms through which these temperature effects operate has been a challenge. While Hsiang (2010) points out that thermal stress on workers is consistent with his empirical results, it is hard to rule out alternative mechanisms. For example, higher temperatures may discourage tourism in the region and falling demand may be part of the story. Other proposed channels for the temperature-output correlation include effects on mortality and political conflict (Dell, Jones and Olken, 2012). In agriculture, temperatures are likely to affect both farmer productivity and have a direct influence on crop yields.

This paper uses a nationally representative panel of manufacturing plants in India and high frequency production data collected from firms in specific industries to provide the first direct evidence that high temperatures do reduce industrial output, and that heat stress on the job is an important source of declines in worker productivity. We begin by constructing a panel dataset of manufacturing plants in India and estimate the impacts of annual temperature shocks on annual factory output. We show that these impacts are economically significant with an output decline of one to three percent per degree Celsius on hot days. An advantage of our plant-level data is that this relationship is estimated using variation within plants over time and is therefore much less affected by changes in the composition of output and technology than estimates using aggregate data. The magnitude and the non-linear relationship to temperature that we observe is consistent with physiological studies of heat stress. We also find that temperature effects are most acute in plants with a high labor share and low electricity intensity, which we use as a proxy for the likelihood of climate control.

We augment this nationwide panel with high frequency data on temperature and output from three different manufacturing processes: cloth weaving, garment manufacture and steel rolling. Firms within these industries reflect the wide variation in automation, climate control and labor intensity within the manufacturing sector. We find that high temperatures are associated with significant reductions in worker output except in workplaces that are mechanized or use climate control. These 'no-effect' cases are consistent with our hypothesized mechanism of heat stress and also suggest that climate control technologies can provide effective adaptation in the workplace. We also use a fourth case-study of diamond cutting and polishing firms to study climate control investment decisions. These firms have both high labor intensity and extremely high value addition. In these firms, we find that climate control is more likely to be adopted for processes that are labor-intensive or critical to output quality.

Even with workplace climate control, high temperatures may have health effects because they are experienced outside the working day. We explore this possibility using daily worker attendance records for different firms. We find that sustained high temperatures are associated with increased absenteeism if occasional absences are not penalized by the wage contract. For such workers, an additional day of elevated temperatures is associated with a 1 to 2 percent increase in absenteeism. For daily wage workers, where the cost of every absence is high, we find no significant correlation between temperature and absenteeism.

Quantifying the link between environmental factors and human welfare is a central part of the research agenda of modern environmental economics (Greenstone and Jack, 2013). Estimates of temperature-productivity interactions and an understanding of the mechanisms through which they operate can help us evaluate the costs of global and local warming and formulate appropriate policies for mitigation and adaptation. Our paper studies manufacturing but the mechanism of heat stress we investigate is more widely relevant. For example, our results suggest that there may be significant productivity benefits from investments in technologies that limit local temperatures through better urban planning and innovations in lighting and building design (Adhvaryu, Kala and Nyshadham, 2014).

The remainder of this paper is organized as follows. Section 2 summarizes the physiological evidence on heat stress and provide a production framework that relates temperature to economic output. Section 3 describes our data sources and Section 4 presents results from from the national panel of plants and our four case studies. Section 5 quantifies the importance of these effects in the context of climate model predictions for India and Section 6 concludes.

2 Theory and Mechanisms

The physics of how temperature affects human beings is well known. Heat generated while working must be dissipated to maintain body temperatures and avoid the adverse health effects of heat stress. The efficiency of such dissipation depends primarily on ambient temperature but also on humidity and wind speed. If body temperatures cannot be maintained at a given activity level, it becomes necessary to reduce the intensity of work (Kjellstrom, Holmer and Lemke, 2009; ISO, 1989). The elevated temperatures and high humidity present for much of the year in many Indian manufacturing plants could therefore plausibly reduce productivity well before becoming a health hazard.

Several indices of ambient weather parameters have been constructed to measure heat stress. The most widely used is the Wet Bulb Globe Temperature, or WBGT, which in indoor conditions, is determined largely by temperature and humidity (Parsons, 1993; ISO, 1989).[2] Its direct measurement requires specialized instruments but Lemke and Kjellstrom (2012) provide a formula using the air temperature in degrees Celsius, T_A, and water vapour pressure, ρ, calculated from humidity levels. This is reproduced in Equation (1) and we use this in our analysis instead of the ordinary dry bulb temperature whenever data on humidity is available.

$$\begin{aligned} WBGT &= 0.567T_A + 0.216\rho + 3.38, \\ \rho &= (RH/100) \times 6.105 \exp\left(\frac{17.27T_A}{237.7 + T_A}\right). \end{aligned} \tag{1}$$

We would expect human responses to temperature to be non-linear. Comfort increases with temperature at low levels, there may be little impact at moderate levels and heat stress should become progressively more severe at high levels. Laboratory studies suggest that the efficiency with which human beings carry out ergonomic and cognitive tasks falls by approximately 1-2 percent for every degree rise in wet bulb temperatures above 25 degrees Celsius, that is, even at levels that are not considered unsafe from the point of view of occupational safety (Hsiang, 2010). While laboratory estimates cannot directly inform us about temperature-productivity relationships in the workplace (monetary incentives for productivity and the nature of tasks performed may be quite different), they do provide a useful benchmark.

Exposure to high temperatures generates responses on short and long time scales. Heat stress can be expected to have visible effects within minutes or hours and these may be compounded when high temperatures are sustained over longer periods. Our identification strategy exploits this short-run response that sets temperature apart from many other environmental stressors such as air pollutants, and from economic factors such as spillovers from agricultural

[2] Outdoor WBGT levels may also vary with solar radiation and wind speeds.

output. We are able to use high frequency daily data on temperatures and output from a number of industrial case studies to separate short-run responses from longer-term effects, which might be the combined result of multiple physical and economic interactions.

Even with a reasonably good understanding of temperature effects on human physiology, it is not obvious how this might influence economic performance. Workplace activity in the manufacturing and service sectors (in contrast to construction and mining) does not typically require exertion nearing physical limits and takes place indoors or in shielded conditions.[3] Furthermore, the economic implications of reductions in worker productivity also depend on the value added by the tasks being carried out.

A simple production model helps clarify our approach. Consider a plant in which output, Y, is given by the following Cobb-Douglas production function:

$$Y = AL(T_I, L_o)^{\alpha} E^{\beta} K^{\gamma} \tag{2}$$

L, E, K represent labor, energy and capital inputs and A is the total factor productivity. L is a function of labor input L_o. T_I, the indoor or WBGT temperature, depends on ambient temperature T_A via the function $T_I = a + bT_A$. Adaptive technologies such as air conditioning might drive b towards zero, breaking the link between T_I and T_A. Threshold effects of temperature on productivity are captured by the following specification:

$$L(T_I, L_o) = \begin{cases} L_o & \text{if } T_I \text{ is less than } T_C \\ L_o e^{-\theta T_I} & \text{if } T_I \text{ is greater than } T_C \end{cases}$$

[3]The mining sector, where temperature and humidity exposures can be high enough to create serious health hazards, has been an important setting for research on heat stress and for designing occupational safety regulation (Wyndham, 1969).

Differentiating $Z = log(Y)$ with respect to T_A then leaves us with

$$\frac{dZ}{dT_A} = \begin{cases} 0 & \text{if } a + bT_A \text{ is less than } T_C \\ -\alpha\theta b & \text{if } a + bT_A \text{ is greater than } T_C \end{cases}$$

At high temperatures ($T_A > (T_C - a)/b$), Z declines with temperature. The extent of this decline is higher for firms in which the value added by labor (α) is high and where workplace temperatures closely follow ambient temperatures.

This model leads us to expect three empirical regularities that we test for in this paper:

1. Temperature effects on worker productivity should occur mostly at high temperatures (above $25°C$).
2. Temperature impacts on output should be higher where the share of value added by labor is high.
3. Temperature impacts on output should be higher where climate control is limited.

Temperature effects on manufacturing output that do not stem from heat stress are unlikely to conform to all three of these patterns. For example, winter is one of the main agricultural growing seasons in India so agriculture-related spillovers on manufacturing should be linked with the cooler temperatures found in the growing season. Similarly, the presence of climate control within a plant might not matter if output effects were mainly through demand shocks arising in other sectors of the economy.

We now turn to the data used in our two-part empirical strategy. We first describe our annual panel of manufacturing plants and then the micro data collected from our four industry case-studies. Our objective will be to verify whether these multiple contexts and data sets support the theory and model we have just described.

3 Data Sources

3.1 Panel of Manufacturing Plants

We construct a plant-level panel using data from the Annual Survey of Industry (ASI) by the Government of India. The Survey is a census of large plants (employing over 100 workers) and a random sample of about one-fifth of the smaller plants registered under the Indian Factories Act. This design results in an unbalanced panel with large firms appearing every year and smaller firms appearing in multiple years only if they are surveyed. The ASI provides annual information on output, working capital and expenditures in broad categories (fixed capital, energy, labor, etc.) as well as numbers of skilled and unskilled workers employed. The format is similar to census data on manufacturing in many other countries (Berman, Somanathan and Tan, 2005).

In most of our analysis, our dependent variable is some function of the value of output, observed by plant and year. For survey years between 1998-99 to 2007-08 two versions of the survey data were available: (i) a panel dataset containing plant identifiers without district identifiers and (ii) a repeated cross-section containing district codes without plant identifiers. We purchased both versions and matched observations across them to generate a panel with district locations for each plant.[4] This allowed us to match each plant to weather data that is available at the level of a district (see Section 3.3). We drop units that appear less than three times during the study period and perform some data cleaning operations (described in the Appendix) to transparently eliminate outliers. Our final sample has 21,509 manufacturing units (Figure A.1 in the Appendix shows their distribution across districts).

One drawback of the ASI is that many Indian manufacturing firms are not registered under the Factories Act and so are excluded from the survey. This informal and small scale manufacturing sector plays an important role in In-

[4]Districts are the primary administrative sub-division of Indian states with an average area of about 4000 sq.km.

dian manufacturing and may have more limited means to adapt to temperature change. Plants surveyed in the ASI may therefore primarily inform us about temperature sensitivity within larger firms that have greater adaptive capacity.

3.2 Micro-data from Case Studies

To isolate the short-run productivity response resulting from heat stress, we supplement our analysis of the ASI with high-frequency micro-data from plants in different manufacturing settings. We investigate whether the physical productivity of workers is related to daily variations in temperature. Differences in the degree of mechanization, climate control, labor intensity and value addition across these plants allows for a finer test of the heat stress hypotheses. Our four industries and their locations are shown in Figure 1 and the workplace environments are described below.

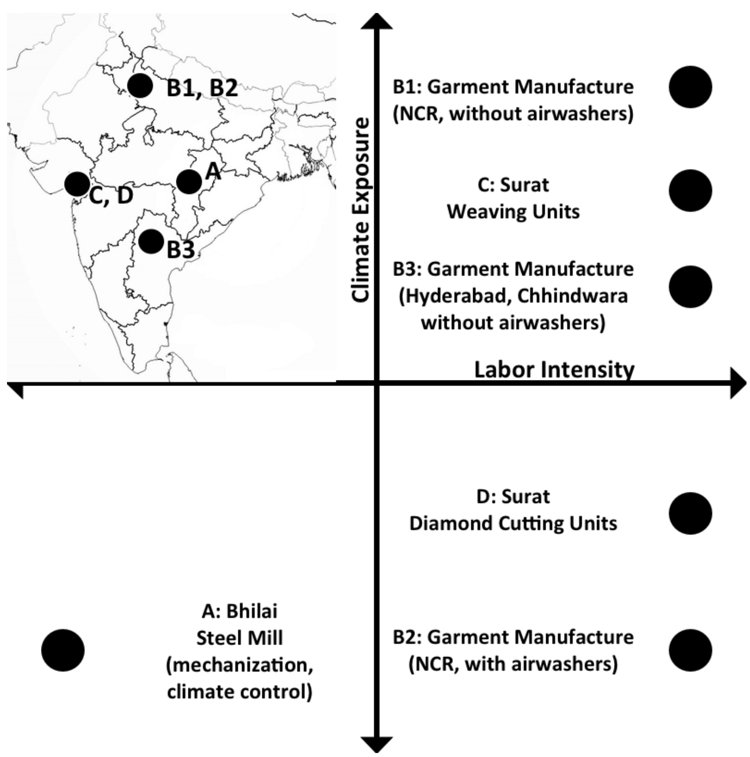

Figure 1: Case study sites span a variety of operating conditions.

Weaving Units in Surat: We use high frequency production data from three cloth weaving units located in the city of Surat in the state of Gujarat in western India. Figure 2, Panel C shows a photograph of the production floor in one of these units. Over 5 million people across the country are employed in weaving units using power looms such as the ones shown here.[5] Labor is semi-skilled and each worker is typically responsible for operating between 6 to 12 mechanized looms producing woven cloth. Workers walk up and down between looms, occasionally adjusting alignment, restarting feeds when interrupted and making other necessary corrections. In this type of small-scale manufacturing setting, temperature control is often limited to the use of windows and some fans.

Workers in these units are only paid for days present and payments are based on the meters of cloth woven (the per meter payment was about INR 2.00). The cloth produced is also the final output for the firm and is sold in wholesale markets or to dying and printing firms. Thus physical worker output directly corresponds to plant revenue. Using administrative firm records, we assemble a data set of daily output and attendance for 147 workers over the financial year 2012-2013 and examine the relationship between daily temperatures, productivity and attendance.[6]

The Bhilai Steel Plant: Our second production setting is a rail mill in one of the largest integrated steel plants in India. The steel plant is located in the town of Bhilai in the state of Chattisgarh in central India and manufactures a variety of steel products. The Bhilai rail mill is the exclusive producer of steel rails for the Indian Railways. To produce rails, steel is first formed into rectangular blocks called *blooms*. Each bloom goes through a furnace and is then rolled (shaped), cut to the desired specifications and then cooled. When a bloom is successfully produced it is said to have been *rolled*. When faults occur, the bloom is referred to as *cobbled* and is discarded.

[5] Government of India, Ministry of Textiles, Annual Report 2011.

[6] Indian minimum wage laws are not legally binding on small firms and are not enforced in these units. Incentive effects on output due to payment non-linearities from minimum wages (Zivin and Neidell, 2012) can therefore be ignored.

This process is heavily mechanized and capital-intensive. It is mostly automated and runs continuously with breaks for correcting malfunctioning machinery. Figure 2, Panel A shows part of the production line where steel blooms are being cast. Workers who manipulate the machinery for shaping blooms sit in cooled cabins. A key output variable tracked by the plant management is the number of blooms rolled into rails for each of three shifts in a day. We also have available measures of cobbled blooms per shift, line delays and worker attendance records. We use daily aggregates of output and labor for the period 1999-2008 and combine these with local temperature data. The Bhilai rail mill is a good example of a capital intensive and mechanized production process with some use of climate control. Therefore workplace heat stress due to ambient temperature shocks should be limited.[7]

Garment Manufacturing: Our third case study is of a firm which produces international brands of garments, largely for export. The firm owns a number of plants across India and we obtained micro-data from six factories in the National Capital Region (NCR) surrounding Delhi, one in Hyderabad in South India and one in Chhindwara in Central India. Garment manufacture is an important part of Indian manufacturing. The textile sector as a whole is estimated to make up about 14 percent of India's industrial production and contributes about 27 percent of export earnings.

Production involves cloth cutting, sewing, embroidery, finishing and washing. We focus on sewing lines in these plants, each of which consists of a group of 10-20 workers who together create part or all of a clothing item. Lines are highly stable in their composition of workers, although the garment manufactured by a given line changes based on production orders. We collected line-level data on the hourly productivity of each line over a two year period from April 2012 to March 2014. Figure 2, Panel B shows a picture of a typical sewing line.

Measuring productivity in this context is more difficult than for the weaving units because garment output rates depend on the complexity of operations

[7]Portions of this dataset are also analyzed and made available as supplementary material with Das et al. (2013).

carried out by a line. The garment export sector is however characterized by significant competitive pressures and large garment manufacturers track worker output in sophisticated ways. We use two variables defined by the management of the firm to encapsulate output: *Budgeted Efficiency* and *Actual Efficiency*. The first of these is an hourly production target set by the management after having the desired operations completed by a special line of 'master craftsmen'. The Actual Efficiency is the per unit rate of throughput actually attained by the line every hour. We use the Actual Efficiency as a measure of the combined productivity of each line of workers, and use Budgeted Efficiency as a control in our regression models.

During the period covered by our data, the firm was in the process of installing cooling systems in its plants. In five manufacturing units in the NCR, production floors were equipped with at least one air-washing system. Air washers enable temperature control and dehumidification and therefore help manage wet bulb globe temperatures effectively. One manufacturing unit in the NCR did not have air-washing installed until 2014 and workers on this site only had access to fans or evaporative coolers.[8] The two plants in Hyderabad and Chhindwara were also without climate control. A comparison of units within this firm therefore provides us with an exceptional opportunity to test for the influence of workplace climate control on the link between temperature and productivity.

Results on productivity responses to temperature in plants with and without climate control are in Section 4.2. The differential assignment of cooling to plants is admittedly not random. Nevertheless, these comparisons of the temperature sensitivity of output in otherwise similar units with identical management suggest that firms can mitigate the impact of temperature on production by investing in workplace cooling even when workers continue to be exposed to uncomfortable temperatures at home.

Diamond Polishing: Our final case study is of diamond polishing firms

[8] The latter may actually *increase* humidity and decrease comfort under high humidity conditions

Figure 2: Production floor images from A: Rail mill, B: Garment manufacture plants, C: Weaving units

in the same city of Surat where our weaving units are located. The city is a hub for various industrial sectors but diamonds are of special importance. An astonishing ninety percent of global diamond output is estimated to pass through Surat for initial cutting and polishing. Diamond firms are particularly interesting as a counterpart to weaving units. Like weaving plants, many diamond units are small and labor-intensive. Unlike weaving however, worker value-addition is very high. Perhaps for this reason, diamond firms in Surat exhibit significant investments in air-conditioning. They also exhibit variation in the use of cooling across production tasks within the same firm.

Diamond polishing can be broadly classified into five distinct operations: (i) sorting and grading, (ii) planning and marking, (iii) bruting, (iv) cutting, (v) polishing. The importance of each of these varies across firms. For example, smaller firms often sort and cut stones and then transfer them to larger firms for final polishing. There is some mechanization at each stage, but trained labor remains critical throughout.

To understand the factors underlying the adoption of climate control technologies, in August 2014, we surveyed a random sample of 150 firms in the city from about 500 manufacturing units formally registered with the Surat Diamond Association. Each firm was asked to provide information on the use of air-conditioning or cooling technologies in each of the five operations listed above (if they took place within the firm) as well as the number of workers and machines used at each stage of production. They were also asked to rate, on a scale of 1-5, the importance of each of these processes to the quality of final output.

This survey was used to estimate the probability of observing air cooling investments at different stages of the production process. We find that climate control is significantly more likely to be present in labor-intensive and high value production stages. Although this is survey and not production data, this evidence supports the idea that firms selectively adopt climate control only when the benefits from reduced heat stress justify the additional costs.

3.3 Meteorological Data

We use meteorological data from two sources. The first is a $1° \times 1°$ gridded data product released by the Indian Meteorological Department (IMD) which provides daily temperature and rainfall based on measurements from the IMD's monitoring stations across the country. A strength of this dataset is that it is based on quality controlled ground-level monitors and not simulated measures from reanalysis models.[9]

Weighted averages of the gridded data provide us with district-level measures of temperature and rainfall for the 609 districts in the country.[10] Although comprehensive in terms of spatial and temporal coverage, this data set cannot be used to estimate WBGT because it does not contain measures of relative humidity. When using the national-level ASI panel, we therefore rely on dynamic variation in temperature alone to estimate the effect of heat on industrial output.[11]

Our second source of data is from weather stations in the vicinity of our case study sites. For all but one of these sites, we obtain local measures of temperature and humidity and use these to compute daily WBGT using (1). The steel plant at Bhilai has no public weather station data available for the period for which we obtain production data. We therefore use the IMD gridded dataset in combination with humidity measures from climate models (NCEP/ NCAR reanalysis datasets) to estimate WBGT for this site.

4 Results

4.1 Annual Manufacturing Output

We begin by examining the response of annual manufacturing output to temperatures experienced during the year. We focus on testing for non-linearities in output response since our hypothesized mechanism of heat stress should

[9] See Auffhammer et al. (2013) for a discussion of some of the concerns that arise when using temporal variation in climate parameters generated from reanalysis data.

[10] The value of temperature or rainfall that we assign to a district is the weighted average from all grid points within a 200km radius of the district centroid with weights inversely proportional to the squares of distances between grid points and centroid. The average district area is about 4000 square km while the grid spacing is about 110 km.

[11] Table A.2 in the Appendix provides results from a robustness check using humidity estimates from climate models (NCEP/ NCAR reanalysis datasets) to approximate WBGT for all districts.

depend mostly on exposure to high temperatures.

Let $V(T_d)$, be the daily output of a manufacturing unit as a function of daily temperature, T_d. We approximate the non-linear response to temperature using a stepwise linear function of production in temperature similar to Hsiang (2010) and Burgess et al. (2011)):

$$\bar{V}(T_d) = \bar{V}(T_0) + \sum_{k=1}^{N} \beta_k D_k(T_d). \tag{3}$$

$D_k(T_d)$ is the number of degree days within a given temperature bin and its coefficient measures the linear effect of a one degree change in temperature on output, within the kth temperature bin.[12] Splitting the annual average temperature into degree days allows us to approximate the true temperature response curve by a piecewise linear function. Since our plant panel from the ASI has annual measures of output and inputs, we use daily district temperatures to compute the number of degree days within each year and district that fall in specified temperature bins, and then estimate the following model:

$$V_{it} = \alpha_i + \gamma_t + \omega K_{it} + \sum_{k=1}^{N} \beta_k D_{itk} + \phi W_{it} + \theta R_{it} + \epsilon_{it}, \tag{4}$$

where V_{it} is the value of output produced by plant i during financial year t, α_i is a plant fixed effect, γ_t are time fixed effects capturing aggregate influences on manufacturing in year t, K_{it} is total working capital at the start of year t, W_{it} is the number of workers and R_{it} is rainfall in millimeters. D_{itk} is the number of degree days in year t that lie in temperature bin k, calculated for the district in which plant i is located. β_k is the output effect of a one degree rise in temperature within bin k. If heat stress causes output declines, we would expect β_k to be close to zero for moderate temperatures (or even positive for

[12] Degree days are commonly used to summarize the annual temperature distribution and carry units of temperature (Jones and Olken, 2010). By this measure a day with a mean temperature of 23 degrees contributes 3 degrees within a temperature bin with bounds of [20,25] degrees.

low temperatures) while for higher degree-day bins we should see negative coefficients. We use daily mean temperatures in all our specifications because they are less noisy than daily maximum temperatures. This is important to keep in mind when interpreting our results. Maximum temperatures are on average 6 °C higher than daily mean temperatures so, for example, with a mean temperature of 25 °C, a substantial part of the day may be spent working at temperatures above 30 °C.

We use working capital available to the plant at the start of the financial year as an input control because it determines resources available for purchasing inputs and is also plausibly exogenous to temperatures experienced during the year and to realized labor productivity. This would not be true of actual labor, energy or raw material expenditures during the year because lower labor productivity due to temperature changes may also reduce the wage bill under piece rate contracts and be accompanied by lower raw material use.

We estimate (4) using both absolute output as well as log output as outcome variables. When using the former, coefficients are expressed as proportions of the average output level. Results are in Table 1. Columns (1) and (3) contain estimates from our base specification. Columns (2) and (4) control for the reported total number of workers W_{it} on the right hand side. These are not our preferred estimates because employment data is both less complete and may contain measurement errors.[13]

The results provide clear evidence of a non-linear effect of temperature on output. Output declined by between 3 and 7 per cent per degree above 25°C, depending on the specification used. This non-linear response suggests the degree day specification in (4) is likely to be more appropriate than simpler formulations using a single annual average temperature. Climate models for India also predict a significant increase specifically in the number of extreme temperature days and not a secular increase in temperatures over the year

[13] Firms in the ASI sometimes do not provide employment numbers. They may also under-report labor to avoid legal and tax implications associated with hiring more workers. This is also why we do not directly use output per unit of labor as a dependent variable.

(Section 5). However for comparison with the literature, we also estimate a linear model and report results in the Appendix in Table A.1. For the most conservative specification, with both capital and worker controls, we estimate a 2.8 percent decrease in output for a one degree change in average annual temperature. Dell, Jones and Olken (2012) find a 1.3% decrease in GDP per degree change in annual temperature in countries that were below the global median GDP in 1960, while Hsiang (2010) finds the corresponding number to be 2.4% in the Caribbean and Central America.

Heterogeneity in Temperature Response

We argued in Section 2 that heat stress should generate the highest production declines in manufacturing plants with high labor share of output and limited climate control. To investigate whether temperature has heterogeneous effects on productivity based on these characteristics, we calculate for each plant in our dataset the ratio of wages paid over every year to output in that year and also the ratio of electricity expenditures to total cash on hand at the start of the year (our measure of capital). Electricity consumption is used as a proxy for air conditioning, which is electricity intensive, because we do not observe climate control investments directly in the annual survey data.[14] We then classify our plants by the quartile to which they belong on each of these measures, interact these quartile dummies (Q_i) with mean temperature and estimate Equation (5) separately for labor shares and electricity quartiles to examine whether temperature effects are heterogeneous in the manner we expect.

$$V_{it} = \alpha_i + \gamma_t + \omega K_{it} + \beta T_{it} \times Q_i + \theta R_{it} + \epsilon_{it} \tag{5}$$

We find that output from plants with higher labor shares is indeed more

[14] We examine the effect of air cooling more directly in Section 4.2 where we observe the climate control technologies that are actually adopted.

Table 1: Non-Linear Effect of Temperature on Manufacturing Industry Output

	Dependent variable:			
	Plant Output Value		Log Plant Output	
	(1)	(2)	(3)	(4)
Below 20°C	0.015	0.013	−0.004	−0.006
	(0.025)	(0.024)	(0.024)	(0.023)
20°C to 25°C	−0.039	−0.034	−0.056***	−0.046*
	(0.026)	(0.026)	(0.022)	(0.022)
Above 25°C	−0.069***	−0.056***	−0.039***	−0.031**
	(0.016)	(0.015)	(0.014)	(0.013)
rainfall	0.008***	0.006**	0.002	0.001
	(0.003)	(0.003)	(0.002)	(0.002)
capital	0.382***	0.343***		
	(0.010)	(0.009)		
log(capital)			0.383***	0.303***
			(0.006)	(0.006)
workers		0.002***		
		(0.0001)		
log(workers)				0.415***
				(0.008)
Plant FE	*Y*	*Y*	*Y*	*Y*
Year FE	*Y*	*Y*	*Y*	*Y*
Worker Controls	*N*	*Y*	*N*	*Y*
Units	21,509	21,509	21,509	21,509
R^2	0.249	0.291	0.196	0.272

Note: 1. Robust standard errors correcting for serial correlation and heteroskedasticity
2. *p<0.1; **p<0.05; ***p<0.01

strongly affected by temperature and that those with greater electricity consumption appear less vulnerable (Table 2).

Robustness Checks: Price Shocks and Power Outages

We have so far focussed on the heat-stress model of Section 2 and not considered other pathways by which temperature may affect output.

For example, temperature shocks might change the prices of plant inputs, especially those coming from agriculture. Although global shocks that affect the population of manufacturing plants are captured by time fixed-effects, there may be local price changes that vary with local temperatures. The ASI surveys allow us to investigate this because plants are asked to report their most common input materials and the per unit price for these inputs each year. We create a price index defined as the log of the average price across the three most common inputs used by each plant. We use this index as the dependent variable in a fixed-effects model similar to Equation (4). We find no evidence that input prices change in high temperature years. These results are in Appendix Table A.3.

A second confounding factor we consider is the regularity of power supply. It is possible that power supply to a plant might be influenced by local temperature shocks. To examine whether power outages might drive the observed temperature-output associations, we control for outages using a detailed measure of state-year outage probabilities for India estimated in Allcott, Collard-Wexler and O'Connell (2014). We find our point estimates across temperature bins remain very similar (Appendix Table A.3).

We turn now to data from our case-studies. Because we are able to observe daily worker output as well as the use of climate control, these sites allow an independent and more precise test of the heat-stress hypothesis.

Table 2: Heterogeneity in the association of output with temperature (by wage share and electricity intensity)

	A: Wage Share Quartiles	*B: Electricity Expenditure Quartiles*
	plant output	plant output
	(1)	(2)
Mean Temperature	-0.038^{***}	-0.064^{***}
	(0.013)	(0.013)
Mean Temperature X		
Quartile 2	-0.007	0.018^{***}
	(0.006)	(0.005)
Quartile 3	-0.023^{***}	0.031^{***}
	(0.007)	(0.007)
Quartile 4	-0.031^{***}	0.032^{***}
	(0.008)	(0.008)
Number of Units	21,509	21,509
Mean Obs. per Unit	4.8	4.8
R^2	0.302	0.260

Note:
1. All models include plant, year fixed effects and capital controls.
3. Robust standard errors
4. *p<0.1; **p<0.05; ***p<0.01

4.2 Daily Worker Output

The physiological basis of heat stress suggests that temperature effects on productivity should become apparent over fairly short periods of exposure. This makes high frequency data especially valuable in isolating heat stress from other mechanisms (such as inter-sectoral spillovers and demand shocks) that might be correlated with temperature but operate over longer time scales.

As discussed in Section 3.2, we obtain daily measures of output from all our case study plants. For the rail mill at Bhilai, we use team-level output measured in rolled blooms. For garment manufacturing units, output is measured as the daily efficiency for sewing teams (lines). For weaving units in Surat, output is measured in daily meters of cloth produced per worker.

For garment and weaving units we estimate a linear model in daily output (Y_{id}) for worker or team i on day d and relate this output to daily WBGT using the following model:

$$log(Y_{id}) = \alpha_i + \gamma_M + \gamma_Y + \omega_W + \beta_k WBGT_{id} \times \mathrm{D_k} + \theta R_{id} + \epsilon_{id}. \qquad (6)$$

In Equation 6 worker or team fixed effects are denoted by α_i while $\gamma_M, \gamma_Y, \omega_W$ denote month, year and day of the week fixed-effects. Together, these capture the idiosyncratic productivity levels of specific workers and teams and control for temporal shocks. For example, there may be seasonal changes in demand and productivity might vary across weekdays. R_{id} and $WBGT_{id}$ are daily measures of precipitation and WBGT for the city in which the unit i is located.

We estimate a similar model for the Bhilai rail mill except that we have three output measures per day corresponding to three daily shifts. Three groups or *brigades* of workers rotate through the shifts. Since productivity varies across night and day shifts, we use a shift-day as our unit of observation and control for nine brigade-shift fixed effects, α_{bs}. We do not observe hourly temperatures so all shifts in a particular day are assigned the average daily temperature.

To capture non-linearities in the effects of heat-stress, we interact daily wet bulb temperature, $WBGT_{id}$ with a dummy variable D_k for different temperature ranges. This allows us to separately estimate the marginal effect on output for a degree change in temperature within different temperature bins. We split the response curve into four wet bulb globe temperature bins: $< 21°C$, $< 21°C - 25°C$, $< 25°C - 27°C$ and $\geq 27°C$. This allows for an easy comparison of our results with Hsiang (2010) who uses breakpoints of 25 and 27 degrees.

Table 3 summarizes our results for all case study sites. Column 1 is based on the rail mill data, columns 2-4 on garment manufacturing lines and columns 5-6 on cloth output from weaving units. The shaded columns (1 and 2) represent climate-controlled plants. Columns 2 and 3 are similar garment units operated by the same firm and located in the National Capital Region (NCR). Column 4 presents data from garment plants located in the milder climate of South-Central India (Hyderabad and Chhindwara). Rows 3-6 provide the incremental change in output for a one degree change in wet bulb globe temperature within a given WBGT bin.

In addition to this binned piecewise linear model, we estimate this relationship much more flexibly by modeling the impact of WBGT on output using cubic splines with four knots. Figure 3 shows the predicted impact of temperature on output measures using these spline fits. Table 3 and Figure 3 together identify patterns that are strongly supportive of the hypothesis that heat stress is a factor influencing manufacturing output.

The clearest evidence in support of our hypothesis is obtained from a comparison of garment manufacture units, located in the capital region around Delhi, with identical management but different levels of climate control (Panel B of Figure 3 and Columns 3 and 4 in Table 3). Production lines on floors without access to air-washers show a clear drop in output with increasing wet bulb globe temperatures especially in the highest temperature bins. This link is almost completely broken in settings with climate control. The width of confidence bands for these two splines varies since we have more data available from teams on production floors with air-washers (Table 3). Garment lines located in Hyderabad and Chhindwara - where air-washers were not installed - also show a drop in efficiency with increasing wet bulb temperatures but the estimated response is smaller, most likely due to the more moderate ambient temperatures in these areas relative to Delhi (Panel C of Figure 3 and Column 5 of Table 3).

In small weaving units of Surat, another setting without climate control, a similar non-linear pattern of temperature impact on worker output is observed with negative impacts on days with high wet bulb temperatures (Panel D of Figure 3 and Columns 6 and 7 of Table 3). In contrast, in the highly mechanized rail mill, where many workers are also located in air-conditioned cabins, there is no evidence that output is negatively affected by very high temperatures and our point estimates are very small and often not statistically significant from zero. The production of rails involves the heating and casting of steel which may be directly influenced by ambient temperatures even if there is no effect on workers. This may be one reason for the more complicated response function seen in spline graphs in Panel A of Figure 3).

The micro data from our case studies also helps assuage concerns about power outages underlying the estimated relationship between productivity and temperature. The data in all panels of Figure 3 comes from manufacturing settings with power backups. In the case of Panel B, we additionally compare co-located plants for whom the incidence of power outages should be similar. For weaving units (Panel D) we were able to confirm that the electricity utility in Surat occasionally scheduled pre-announced weekly power holidays on Mondays. Any effect of such power outages (notwithstanding power backups) is controlled for in our estimates by including day of week fixed effects. While we were not able to observe and control for the plant level power outages in the more aggregate ASI data, these case studies suggest that observed differences in temperature sensitivity are unlikely to be driven by power supply variations.

Table 3: Effect of Wet Bulb Globe Temperature on Daily Worker Output

	Dependent variable:					
	Rail Mill	Garment Manufacture Plants			Weaving Plants	
	log(blooms)	log(efficiency)	log(efficiency)	log(efficiency)	log(meters)	meters
	(1)	(2)	(3)	(4)	(5)	(6)
(1) rainfall	0.001***	0.083**	0.044	−0.067	0.006	1.512
	(0.0002)	(0.030)	(0.192)	(0.035)	(0.008)	(0.958)
(2) log(budgeted efficiency)		0.796***	0.421***	0.525***		
		(0.034)	(0.126)	(0.044)		
(3) WBGT:[<20]	−0.008*	0.014***	−0.026***	−0.15	0.001	0.462
	(0.003)	(0.004)	(0.007)	(0.097)	(0.009)	(0.998)
(4) WBGT:[20-25)	−0.0002	−0.014**	−0.064***	−0.004	0.006	1.627*
	(0.005)	(0.007)	(0.020)	(0.009)	(0.009)	(0.835)
(5) WBGT:[25-27)	0.011*	0.029**	−0.149***	0.004	−0.014	−0.492
	(0.006)	(0.014)	(0.026)	(0.020)	(0.014)	(1.125)
(6) WBGT:[≥27]	0.016	0.001	−0.087***	−0.037**	−0.085**	−7.131**
	(0.011)	(0.007)	(0.024)	(0.016)	(0.038)	(2.923)
Number of Plants	1	5	1	2	3	3
Number of Units	9	74	10	19	147	147
Climate Control	*Y*	*Y*	*N*	*N*	*N*	*N*
Worker or Line FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Month FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Year FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Weekday FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*

Note:
1. Shaded columns represent sites with climate control (use of AC or airwashers)
2. Number of units refers to the number of distinct workers (weaving) or teams (steel and garments)
3. Robust standard errors correcting for serial correlation and heteroskedasticity
4. *p<0.1; **p<0.05; ***p<0.01

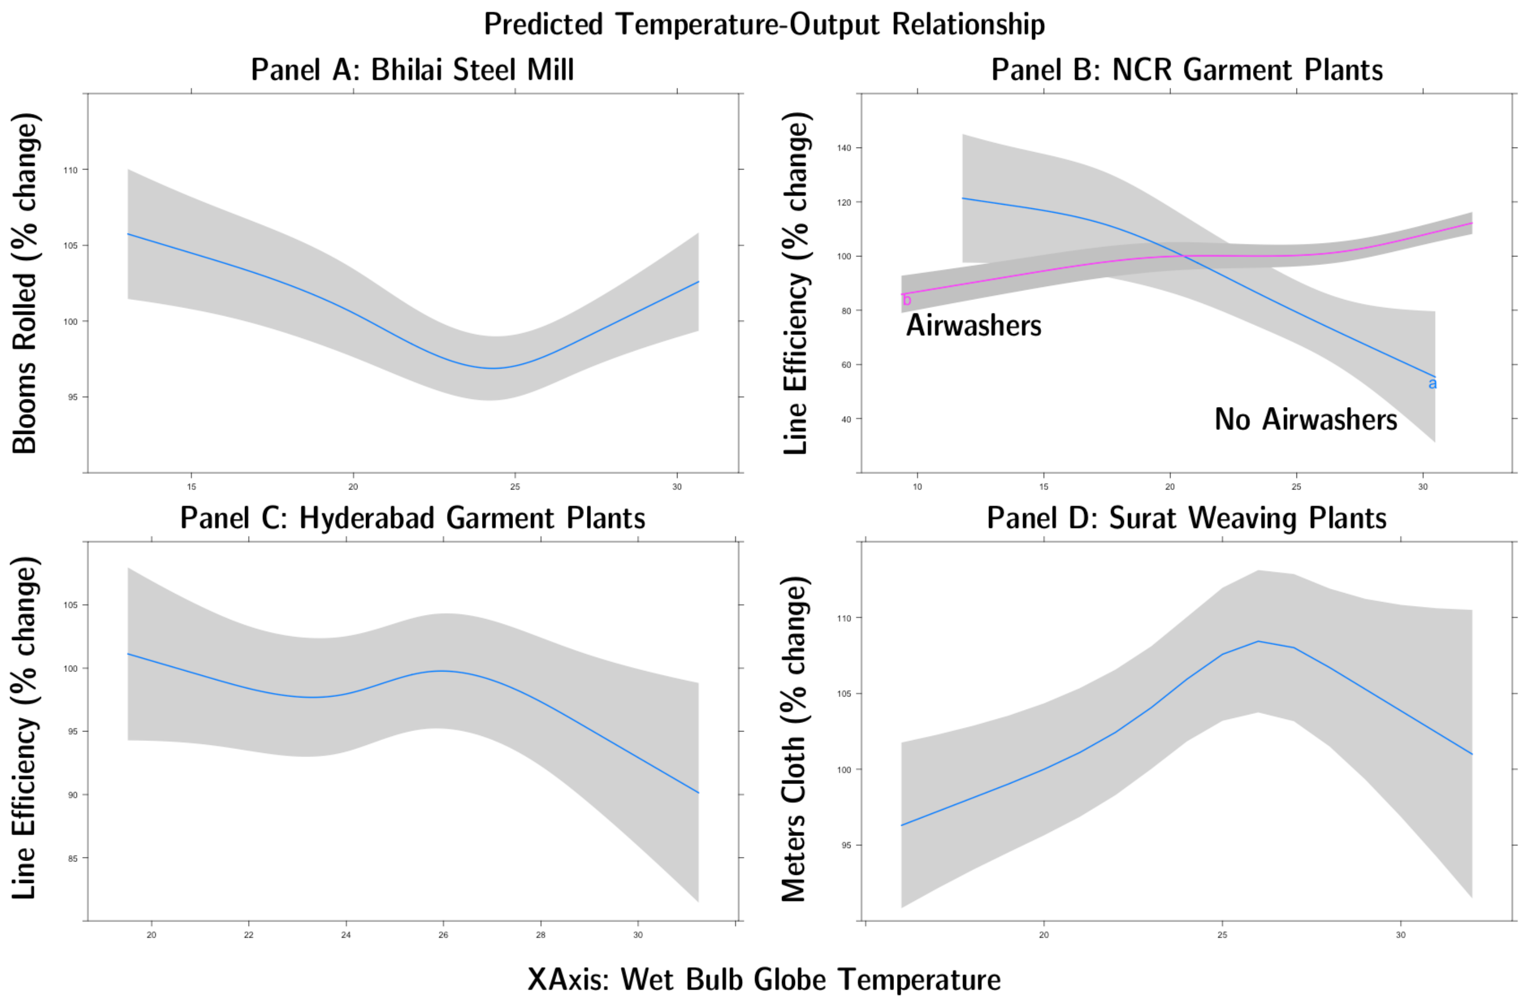

Figure 3: Restricted cubic spline models of the impact of temperature on output measures (90 percent bootstrapped confidence intervals, output at 25 degrees normalized to 100 percent). Panel A: Rolled blooms against temperature (Bhilai Rail Mill). Panel B: Logged efficiency in garment plants in NCR both with airwashers (5 plants) and without airwashers (1 plant). Panel C: Logged efficiency for garment plant in Hyderabad and Chhindwara without airwashers. Panel D: Logged meters of cloth produced by weaving workers in Surat

4.3 Worker Absenteeism

Ambient temperatures may impact worker attendance in various ways. A very hot day might reduce the desire to go to work. Sustained high temperatures may eventually lead to fatigue or illness. Longer term seasonal variations could create differences in disease burden and change the choice of occupation. Recent evidence from the United States suggests people may allocate time away from work on hot days (Zivin and Neidell, 2014). Such changes in attendance could affect labor input costs and output independently of actual workplace performance.

The ASI does not provide a good measure of worker attendance. However we were able to collect detailed histories of daily worker attendance using administrative records from garment units located in the NCR and from the rail mill at Bhilai. For weaving plants in Surat we obtain information on worker attendance using worker level payment records. For all three cases we construct a time series measure of the total number of worker absences every day.[15] These absence records span two years (2012 and 2013) for garment plants, three years (Feb 2000 to March 2003) for the rail mill and one year (April 2012 to March 2013) for Surat weaving units.

This micro-data can be used to investigate the relationship between absenteeism and temperature. To begin, we note that the probability of worker absence (equivalently, the number of absences in a cohort) on any given day may depend on both contemporaneous and lagged temperatures.

Denoting the number of absences in a cohort of workers observed on day t_0 by A_{t_0}, we might model

$$A_{t_0} = \alpha + \beta E_{t_0} + \gamma X_{t_0} + \epsilon_{t_0}$$

[15] In the case of the rail mill and garment plants an absence is defined as a recorded leave day. In the case of daily wage weaving workers an absence is defined as any day when no payments were recorded for a worker. Absences for garment workers are calculated for workers observed for at least 600 days over the two year period.

Here $E_{t_0} = f(W_{t_0}, \ldots, W_{t_0-K})$ is the accumulated heat exposure at time t_0 that depends on the history of all wet bulb globe temperatures experienced over the previous K days. γX_{t_0} denotes other covariates (such as festival seasons) that might change A_{t_0}.

In general E_{t_0} could vary non-linearly with both the level of wet bulb globe temperatures, W_{t_0-k}, as well as the lag period k. This general model can be simplified (subject to assumptions) and parameters estimated from the data.

One common specification is to set $E = \beta W_{t_0}$. This assumes that exposures are a function only of contemporaneous temperatures. A second natural specification involves setting $E = \beta \mathbf{W^K_{t_0}}$ where $\mathbf{W^K_{t_0}} = \frac{\sum_{k=0}^{K} W_{t_0-k}}{K}$ is the mean wet bulb globe temperature experienced over the previous K days. This assumes that all temperatures W, experienced over the K days preceding t_0, contribute equally and linearly to exposure E_{t_0}. We might also be interested in testing for the presence of non-linearities in the response as a function of the level of average temperatures experienced over the lag duration K. One way to do this is to separately estimate β for different quartiles of observed $\mathbf{W^K_{t_0}}$.

We estimate models for absenteeism using both contemporaneous WBGT ($E_{t_0} = \beta W_{t_0}$) as well as a model relating absenteeism to the average WBGT for the previous week ($K = 7$), interacted with dummies for different quartiles (i.e $E_{t_0} = \beta_j \mathbf{W^K_{t_0}} \times \mathrm{D_j}$, where $\mathrm{D_j}$ is a dummy for quartiles of weekly average WBGT). We additionally control for month fixed effects, year fixed effects, day of week fixed effects and rainfall.

Table 4 presents the results. We find evidence that high temperatures are associated with an increase in absenteeism for workers in the rail mill and garment plants. For the highest quartile of lagged weekly temperatures, a $1°C$ increase in the average weekly WBGT is associated with a 10 percentage point increase in absences for rail mill workers and a 6 percentage point increase for garment workers. In contrast, we do not see absenteeism effects for weaving workers, perhaps because of their very different wage contracts. In garment and rail plants, workers are full-time employees paid a monthly wage, while in the weaving units they are daily wage workers who are not paid when they do not come to work. This means that the marginal cost of an additional absence is relatively high for weaving workers, while it may be small or zero in the other two cases.

A little more insight can be had by estimating a less restrictive model. The exposure-response relationship can be flexibly modeled using distributed lag models (DLMs) or non-linear DLMs (Gasparrini (2013) provide details on empirical estimation). These models represent E_{t_0} as a weighted sum of lagged wet bulb globe temperatures so that $E_{t_0} = \tau_0 W_{t_0} + \tau_1 W_{t_0 - 1} + ... + \tau_K W_{t_0 - K}$ with weights τ related to each other by some flexible function whose parameters can be estimated from the data. A non-linear DLM extends this idea to allow exposure weights to vary across both lag-space, K and temperature levels, N. A completely unrestricted model would require a full set of NxK parameters to be estimated but by assuming that the variation of weights in lag-space and levels can be described by two polynomials, a fairly parsimonious yet flexible model can be estimated.[16] We use two third order polynomials to describe how cumulative exposure E_{t_0} at time t_0 varies with both the level and lag period of ambient WBGT. We then use this model to simulate predicted changes in absenteeism under any specified history of WBGT exposures.

Figure 4 displays two cross-sections. The left column shows the predicted change in the logarithm of daily absences for a $1°C$ increase in WBGT, over a $25°C$ reference, sustained for K days (K ranging from 1 to 10). For workers with long term contracts - rail mill (Panel A) and garment firms (Panel B) - absences increase approximately linearly with every additional day of elevated temperatures at the rate of approximately 1 to 2 percent per day. We see no effect on daily wage workers. In the right column, we simulate the variation in absenteeism for a fixed exposure duration (10 days) at varying levels of temperature. We see clear evidence that high temperatures drive the absenteeism response as suggested by the binned weekly WBGT models in Table 4.

[16] We use the dlnm package in R (Gasparrini, 2011) to estimate these models.

Table 4: Effect of Temperature on Worker Absenteeism

	Dependent variable: log(Absences)					
	Rail Mill		Garment Manufacture		Weaving	
	(1)	(2)	(3)	(4)	(5)	(6)
WBGT	0.032***		0.014*		0.012	
	(0.010)		(0.008)		(0.012)	
Weekly WBGT						
x Q1		0.051*		0.025		0.014**
		(0.030)		(0.018)		(0.006)
x Q2		0.042		−0.044*		−0.009
		(0.037)		(0.024)		(0.013)
x Q3		0.073		0.006		0.017
		(0.048)		(0.026)		(0.025)
x Q4		0.097***		0.059***		0.015
		(0.034)		(0.009)		(0.035)
rainfall	0.002	0.001	0.45***	0.51***	0.027	−0.001
	(0.002)	(0.002)	(0.17)	(0.17)	(0.029)	(0.007)
Days	857	857	662	662	365	365
Workers	198	198	2700	2700	147	147
Month FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Year FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Weekday FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*

Note:
1. *p<0.1; **p<0.05; ***p<0.01
2. Robust standard errors correcting for serial correlation and heteroskedasticity
3. Q1-Q4 refer to quartiles of weekly WBGT
4. Rainfall in 1,2,5,6 measured in mm.

Rainfall in 3,4 measured in fraction of hours with recorded precipitation event

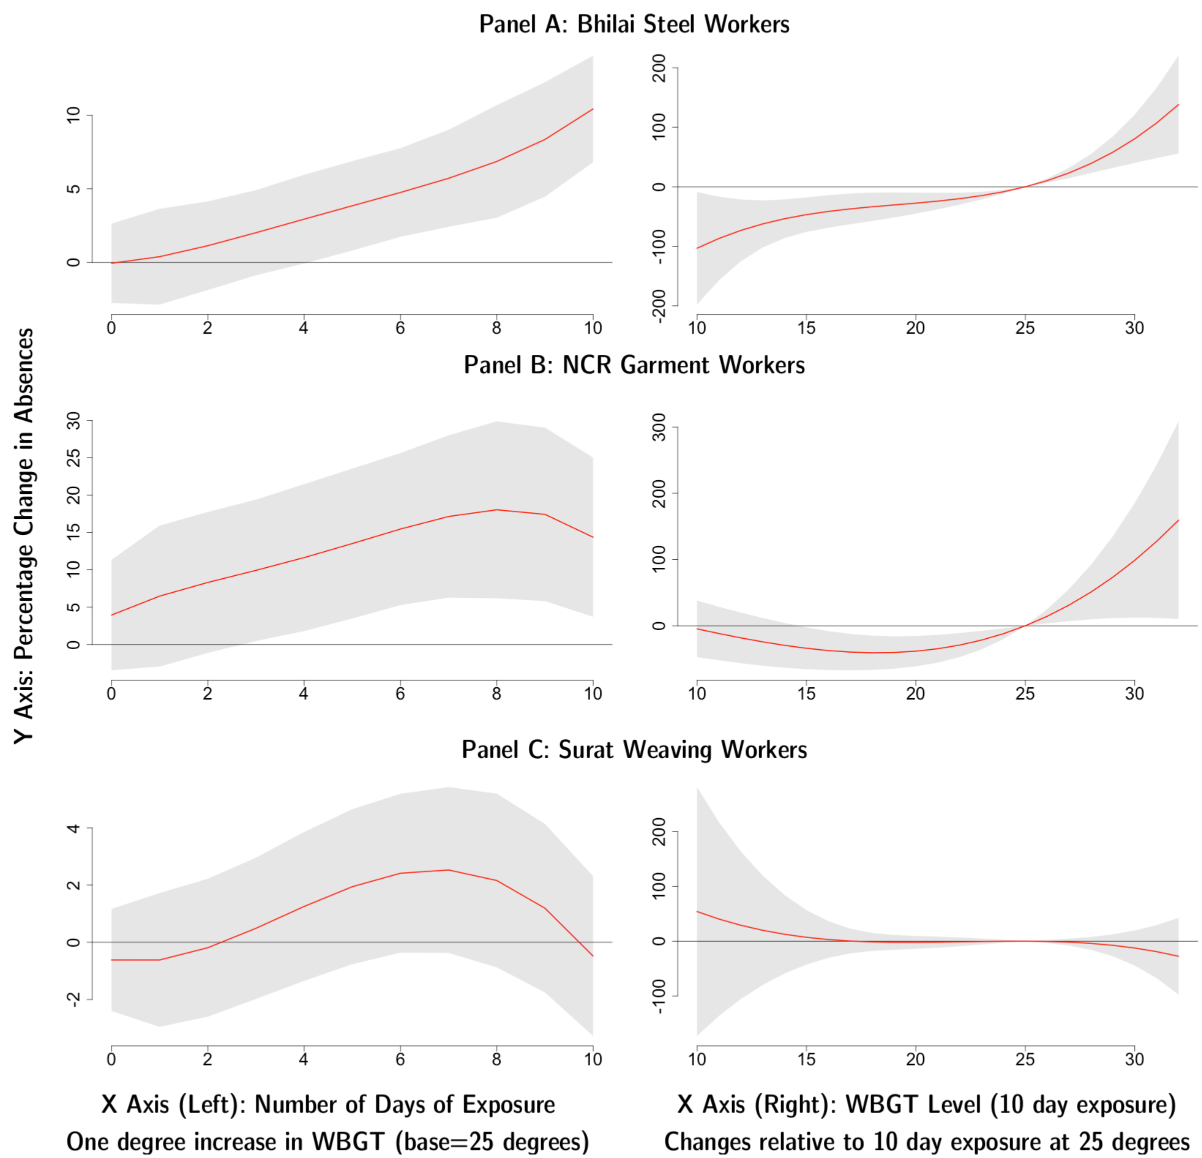

Figure 4: Predicted impact of wet bulb globe temperature on attendance measures for rail mill workers (Panel A), garment workers (Panel B) and workers in weaving firms (Panel C). Left column plots predicted percentage change in absenteeism under a one degree temperature increase sustained for varying periods of time. Right column plots predicted percentage change in absenteeism for different temperature levels sustained for 10 days (relative to the exposure level corresponding to 25 degrees WBGT experienced for 10 days). Bootstrapped 90 percent confidence intervals

Our analysis is restricted to short-run (10-day) responses of attendance to temperature shocks. Although our data does not support a detailed investigation of longer run responses, Figure A.2 in the Appendix suggests there are seasonal reductions in the availability of daily wage workers (but not full-time contracted workers) during high temperature months. Daily wage workers have greater flexibility to shift occupations relative to workers on full time contracts.

4.4 Adaptation and Investments in Climate Control

In Sections 4.1 and 4.2 we directly investigated the relationship between temperature and worker output. An indirect way of testing the heat-stress mechanism is by observing the way in which plants make investments in climate control technologies. We would expect that firms that are concerned about heat impacts on workers would preferentially invest in cooling for production activities that are high value and labor intensive.

We conducted a short survey of over 150 diamond polishing firms in the city of Surat in Gujarat to collect information on the presence of air conditioning. Diamond polishing units are well suited for this purpose because, like the weaving units we study, many diamond units are small-scale and labor-intensive operations. Unlike weaving however, workers engage in activities with high value-addition. Diamond firms in Surat exhibit significant investments in air-conditioning and also substantial variation in the use of cooling across different production tasks within the same firm.

We estimate a logit model of the probability of firms reporting air-cooling at different stages of the diamond production process as a function of the share of workers employed in the process (worker intensity), the share of machines used within a process (mechanization intensity) and the self-reported importance of the process in determining stone quality (a proxy for value addition), controlling for total number of workers (a proxy for firm size) and the years since the first climate control investment.

Figure 5 summarizes our results. We find that diamond firms in Surat are significantly more likely to use climate control for production tasks they consider important in determining product quality and for tasks that are labor-intensive. These patterns are consistent with a model of adaptive investments where firms choose to preferentially cool high value and labor intensive processes.

It is also possible that investing in air-conditioning reflects a form of compensation to attract higher quality workers rather than an effort to offset negative temperature impacts. This explanation seems unlikely because wages are low and workplace activities are not physically taxing. Workers would therefore probably prefer wage increases to equivalent expenditures on air-conditioning. Workers also move between different parts of the production process. Air-conditioning is therefore better regarded as being associated with a production activity rather than a form of compensation.

5 Climate Model Projections

The economic impact of productivity losses from global warming will depend on how the distribution of temperature shifts. Since losses from heat stress arise only on hot days, the economic impact of warming will be modest if summers are sufficiently cool to begin with, or if warming consists mainly of warmer winters rather than hotter summers. Unfortunately, as we shall now see, neither of these is true for India.

Panel A of Figure 6 reproduces a map of annual wet bulb temperature maximums from (Sherwood and Huber, 2010). It is seen that Indian summers are among the hottest on the planet, along with those in the tropical belt and the eastern United States. The areas in red in Figure 6 all experience maximum wet bulb temperatures that are above $25°C$. This suggests that - absent adap-

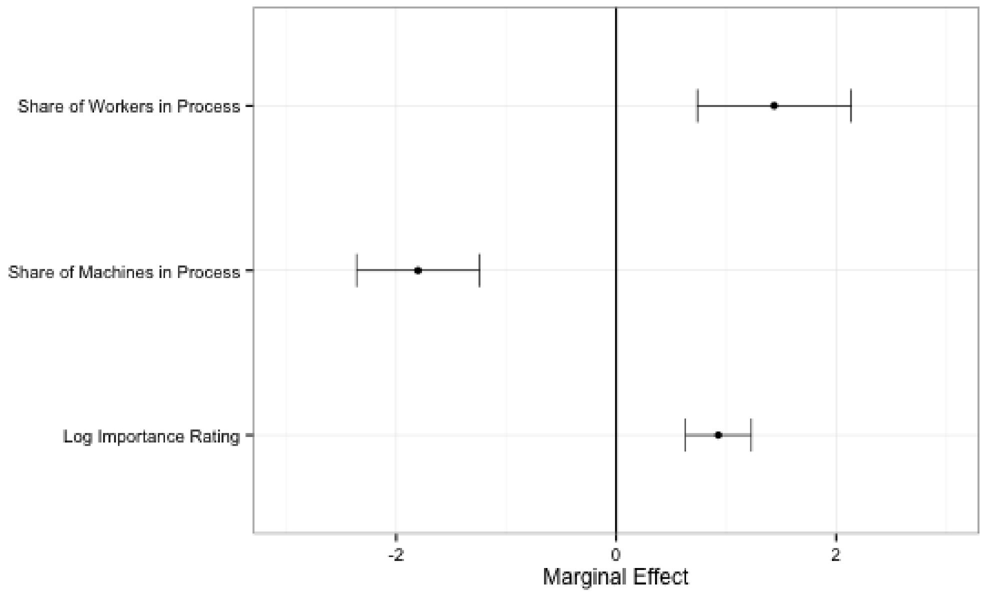

Figure 5: Marginal effect of covariates on probability of seeing climate control for a single process within the diamond production line. Bootstrapped robust standard errors.

tation - an increase in the frequency or severity of high WBGT days might rapidly impose large productivity costs in these regions.

Panel B of Figure 6, (left axis), plots projections of the long run change in the annual temperature distribution for India from two commonly cited climate models: (i) the A1F1 "business-as-usual" scenario of the Hadley Centre Global Environmental Model (HadGEM1) from the British Atmospheric Data Centre and (ii) the A2 scenario of the Community Climate System Model (CCSM) 3, from the National Center for Atmospheric Research. As is evident, the predicted increase in degree days is concentrated in the highest temperature bins.

We then overlay (right axis of Panel B of Figure 6) our estimated marginal effects of temperature on manufacturing output using the ASI data from Table 1 (column 2). The temperature range where we estimate significant negative productivity impacts from an additional degree day is precisely the range where the largest increases in degree days are predicted by the climate models.

We can use the climate models to obtain the projected change in degree-days per year within the temperature bins of ($\leq 20°C$, $20°C - 25°C$, $> 25°C$). For the Hadley model projections these are (-1.79, -0.64, 3.34) degrees respectively. For the CCSM projections they are (-1.17, -0.55, 1.32) degrees. Assuming the lower projection is a more reasonable estimate - and multiplying these numbers by our empirical estimate of the impact of a degree day on output - suggests that absent adaptation, the estimated impact on manufacturing could be as high as -7% (95 percent CI = [-2.77,-10.69]).

This estimate is only indicative because it holds the labor-intensity of manufacturing constant, ignores adaptation that will mitigate large impacts, and cannot include effects of temperatures outside the range of what has been observed. Adaptive actions might include air conditioning, shifting manufacturing to cooler regions, urban planning measures designed to lower local temperatures (green cover, water bodies), building design modifications (cool roofs) and so on. Adaptation could also include techniques to reduce the intensity of work, or the use of economic incentives to encourage worker effort. Recent work also suggests adaptive possibilities from the use of LED lighting (Adhvaryu, Kala and Nyshadham, 2014). Many of these measures are neither easy or costless and research into affordable technologies to reduce heat exposures is likely be worthwhile. Heat island effects in urban areas have already led to temperature hotspots that can be more than five degrees warmer than surrounding areas (Mohan et al., 2012; Zhao et al., 2014).

6 Conclusions

This paper has provided new evidence to show that heat stress may be an important mechanism underlying previously observed correlations between surface temperatures and the economic output of poor and developing countries (Dell, Jones and Olken, 2012).[17] While this is not the only factor explaining

[17]To the extent that climate control technologies could mitigate the effects of high temperatures on labor, this result might reflect the relative prevalence of these technologies in

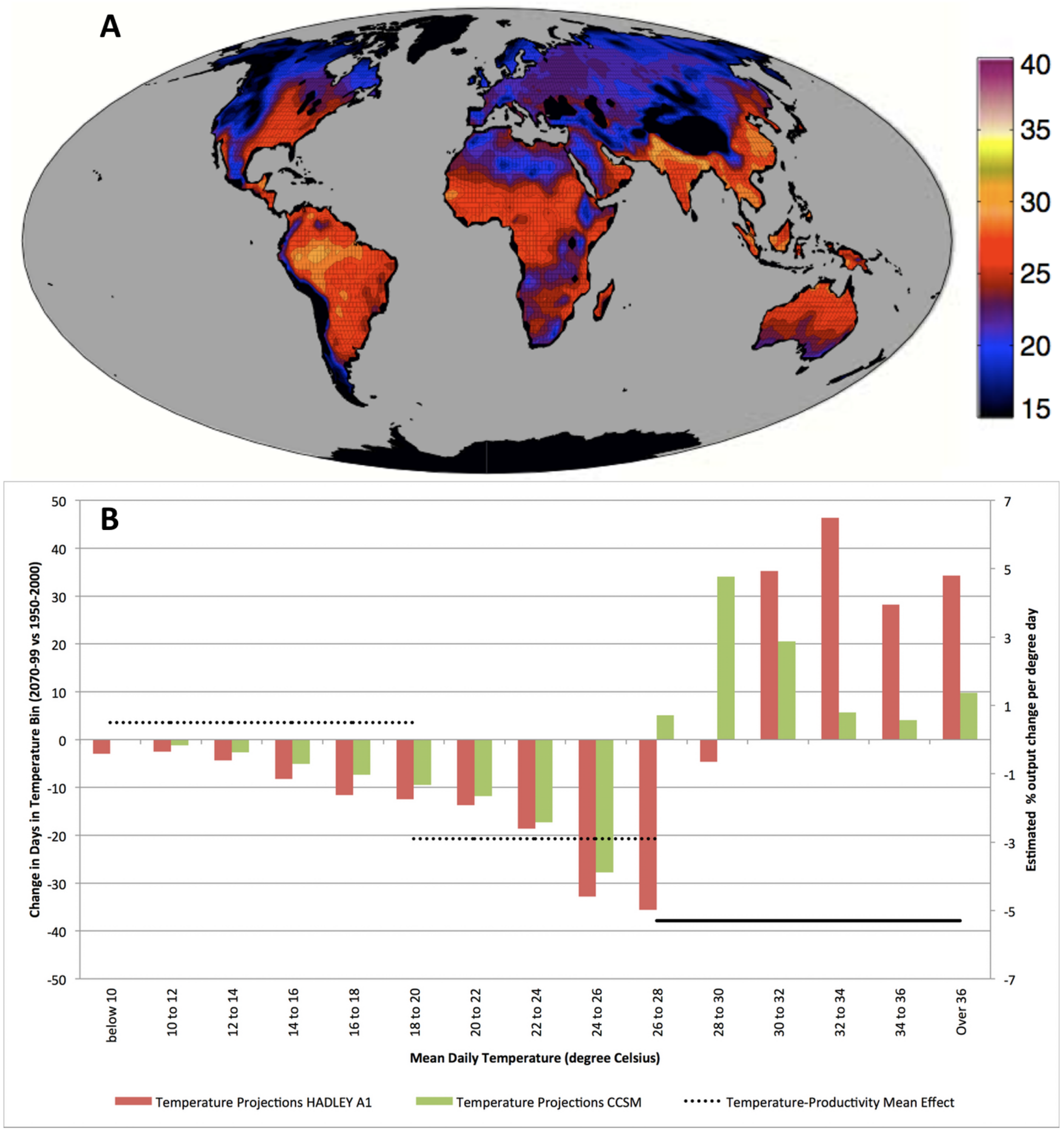

Figure 6: Panel A: Estimated annual wet bulb globe temperature maxima, 1999-2008. Source: Sherwood and Huber (2010). Panel B: Projected temperatures under a business as usual climate change scenario for India. Source: Burgess et al. (2011). Overplotted lines denote estimated productivity impacts of temperature from Table 1, Column 3. Solid segments imply statistically significant effects

macro-level correlations, the effect sizes we identify in independent datasets from India's manufacturing sector are similar in magnitude both to laboratory studies and to evidence from country level panel studies. Taken altogether, we argue that there is a compelling case for being concerned about temperature impacts on worker productivity and therefore the direct economic costs of gradual climate change.

Our findings also relate to the scientific literature on urban temperature changes. Urban heat island effects have been extensively studied by scientists but relatively little attention has been paid to them by economists. Our results suggest that urban heat islands may have direct and economically significant economic effects in developing country settings where climate control is limited. Satellite based heat island studies in Delhi for instance show that urban hotspots can experience temperature elevations of greater than five degrees celsius (Mohan et al., 2012).

Finally, we show that adaptation to high temperatures is possible through the use of workplace climate control. We also find that attendance reductions are not observed in workers who face high opportunity costs of absenteeism. This suggests that economic incentives could be used to mitigate some behavioral responses driven by environmental change.

While our study has examined only the manufacturing sector in India, the mechanism that we identify of heat stress reducing worker productivity may be even more pronounced in agriculture and other sectors involving outdoor activity. Observed productivity losses in agriculture that have been attributed by default to plant growth responses to high temperatures may in fact be partly driven by lower labor productivity. These possibilities are yet to be researched.

richer countries.

References

Adhvaryu, Achyuta, Namrata Kala, and Anant Nyshadham. 2014. "The Light and the Heat: Productivity Co-benefits of Energy-saving Technology."

Allcott, Hunt, Allan Collard-Wexler, and Stephen O'Connell. 2014. "How Do Electricity Shortages Affect Productivity? Evidence from India."

Auffhammer, M., S. M. Hsiang, W. Schlenker, and A. Sobel. 2013. "Using Weather Data and Climate Model Output in Economic Analyses of Climate Change." *Review of Environmental Economics and Policy*, 7(2): 181–198.

Auffhammer, M, V Ramanathan, and J R Vincent. 2006. "From the Cover: Integrated model shows that atmospheric brown clouds and greenhouse gases have reduced rice harvests in India." *Proceedings of the National Academy of Sciences*, 103(52): 19668–19672.

Berman, Eli, Rohini Somanathan, and Hong W Tan. 2005. "Is Skill-biased Technological Change Here Yet?: Evidence from Indian Manufacturing in the 1990's." *Annals of Economics and Statistics*, 79-80(299-321).

Burgess, Robin, Olivier Deschenes, Dave Donaldson, and Michael Greenstone. 2011. "Weather and Death in India."

Das, Sanghamitra, Kala Krishna, Sergey Lychagin, and Rohini Somanathan. 2013. "Back on the Rails: Competition and Productivity in State-Owned Industry." *American Economic Journal: Applied Economics*, 5(1): 136–162.

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken. 2012. "Temperature Shocks and Economic Growth: Evidence from the Last Half Century." *American Economic Journal: Macroeconomics*, 4(3): 66–95.

Field, CB, V Barros, K Mach, and M Mastrandrea. 2014. "Climate Change 2014: Impacts, Adaptation, and Vulnerability."

Gasparrini, Antonio. 2011. "Distributed Lag Linear and Non-Linear Models in R: The Package dlnm." *Journal of Statistical Software*, 43(8): 1–20.

Gasparrini, Antonio. 2013. "Modeling exposure-lag-response associations with distributed lag non-linear models." *Statistics in Medicine*, 33(5): 881–899.

Greenstone, Michael, and B. Kelsey Jack. 2013. "Envirodevonomics: A Research Agenda for a Young Field." National Bureau of Economic Research NBER Working Paper: 19426.

Hsiang, Solomon M. 2010. "Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America." *Proceedings of the National Academy of Sciences of the United States of America*, 107(35): 15367–72.

ISO. 1989. "Hot environments – Estimation of the heat stress on working man, based on the WBGT-index (wet bulb globe temperature)." International Standards Organization.

Jones, Benjamin F, and Benjamin A Olken. 2010. "Climate Shocks and Exports." *American Economic Review*, 100(2): 454–459.

Kjellstrom, Tord, Ingvar Holmer, and Bruno Lemke. 2009. "Workplace heat stress, health and productivity: An increasing challenge for low and middle-income countries during climate change." *Global Health Action*, 2: Special Volume: 46–51.

Lemke, Bruno, Bruno, and Tord Kjellstrom, Tord. 2012. "Calculating Workplace WBGT from Meteorological Data: A Tool for Climate Change Assessment." *Industrial Health*, 50(4): 267–278.

Lobell, David B, Wolfram Schlenker, and Justin Costa-Roberts. 2011. "Climate trends and global crop production since 1980." *Science*, 333(6042): 616–620.

Mendelsohn, R, and A Dinar. 1999. "Climate Change, Agriculture, and Developing Countries: Does Adaptation Matter?" *The World Bank Research Observer*, 14(2): 277–293.

Mohan, Manju, Yukihiro Kikegawa, B. R. Gurjar, Shweta Bhati, and Narendra Reddy Kolli. 2012. "Assessment of urban heat island effect for different land use-land cover from micrometeorological measurements and remote sensing data for megacity Delhi." *Theoretical and Applied Climatology*, 112(3-4): 647–658.

Parsons, K C. 1993. *Human Thermal Environments.* Informa UK (Routledge).

Schlenker, W., and Michael Roberts. 2009. "Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change." *Proceedings of the National Academy of Sciences*, 106(37): 15594–15598.

Sherwood, Steven C., and Matthew Huber. 2010. "An adaptability limit to climate change due to heat stress." *Proceedings of the National Academy of Sciences*, 1–4.

Wyndham, C.H. 1969. "Adaptation to heat and cold." *Environmental Research*, 2: 442–469.

Zhao, Lei, Xuhui Lee, Ronald B. Smith, and Keith Oleson. 2014. "Strong contributions of local background climate to urban heat islands." *Nature*, 511(7508): 216–219.

Zivin, Joshua Graff, and Matthew Neidell. 2012. "The Impact of Pollution on Worker Productivity." *American Economic Review*, 102(7): 3652–3673.

Zivin, Joshua Graff, and Matthew Neidell. 2014. "Temperature and the Allocation of Time: Implications for Climate Change." *Journal of Labor Economics*, 32(1): 1–26.

Appendix: For Online Publication

A.1 Annual Survey of Industry Data Cleaning

The following data-cleaning operations are performed to arrive at the sample used in our analysis:

1. We restrict the sample to surveyed units that report NIC codes belonging to the manufacturing sector.
2. We trim the top 2.5 percent and bottom 2.5 percent of the distribution of observations by output value, total workers, cash on hand at the opening of the year and electricity expenditures. This is done to transparently eliminate outliers. There are some firms with implausibly high reported values of these variables and also a long tail of plants with near zero reported output.
3. We remove a small number of manufacturing units that report having less than 10 workers employed because this represents a discrepancy between the criterion used to select the survey sample and reported data. Such discrepancies may be associated with false reporting since firms with less than 10 workers are subject to very different labor laws and taxation regimes under Indian law. We mark as missing all plants with zero or negative values of output, capital, workers or raw materials used.
4. We drop units that appear less than three times during our study period.

All remaining plants are in our panel.

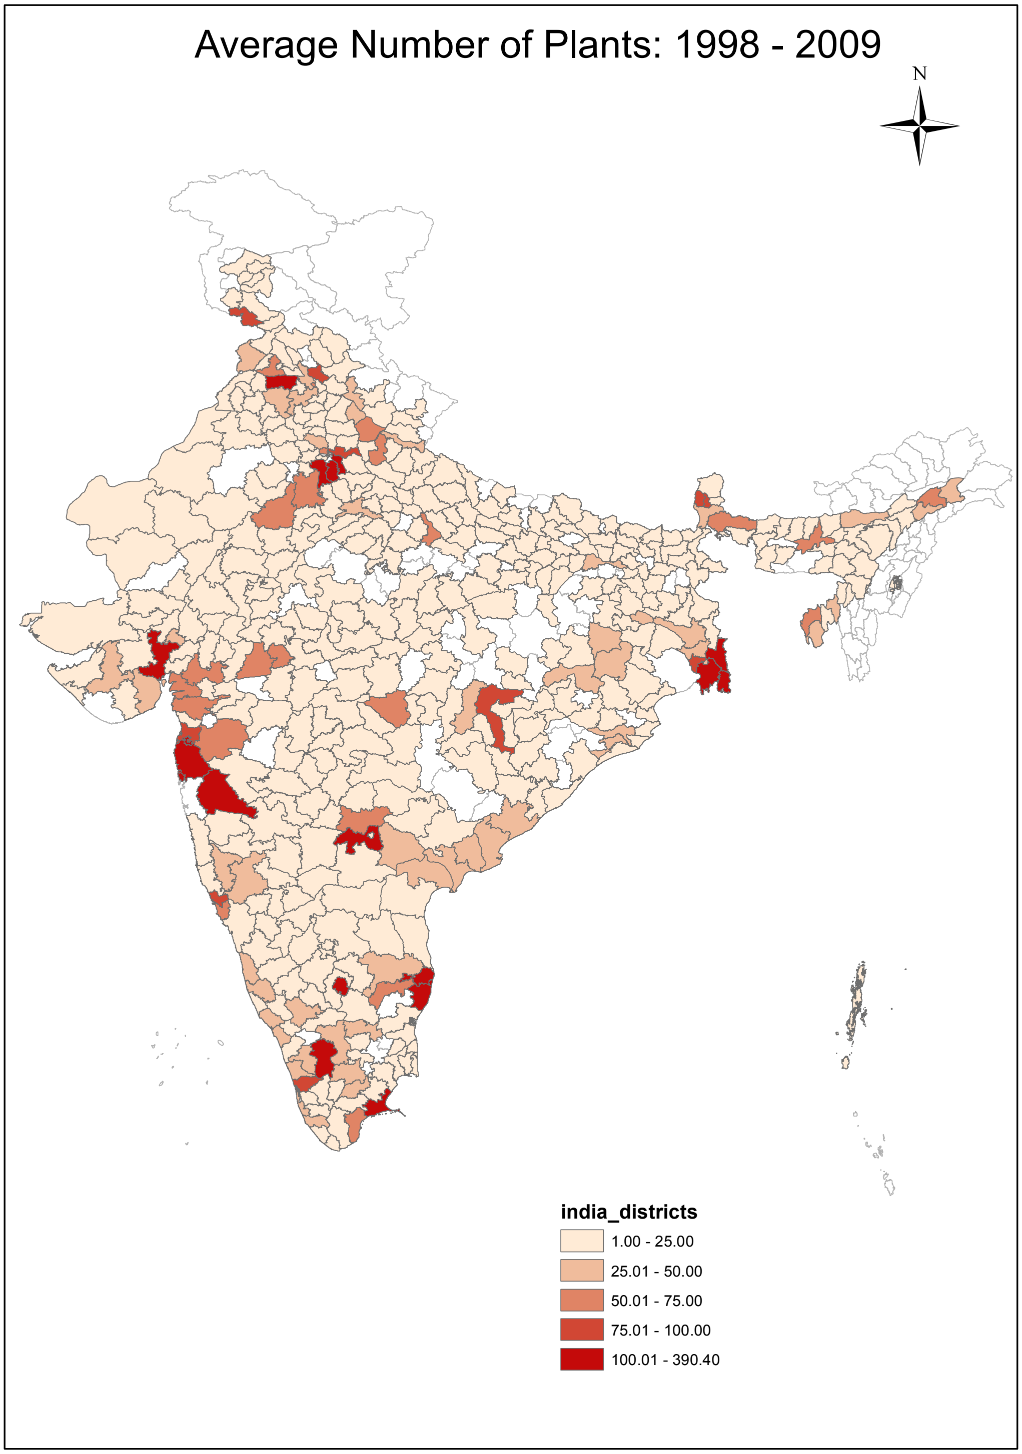

Figure A.1: Distribution of annual ASI survey observations over Indian districts

A.2 Additional Results

Annual Average Temperature and Manufacturing Output

The model in Equation 4 allowed for a non-linear (or piece-wise linear) output response to temperature using four temperature bins. Here we present results from the simpler linear specification. Much of the country-level literature estimates a linear model because degree days cannot be computed for all countries. The estimates in this section facilitate a comparison of our findings with other studies. We estimate the following model:

$$V_{it} = \alpha_i + \gamma_t + \omega K_{it} + \phi W_{it} + \beta T_{it} + \theta R_{it} + \epsilon_{it} \tag{7}$$

where T_{it} is the average temperature during the financial year t (so that a year is calculated from April 1 through March 31) and the other variables are as in (4). Estimates are in Table A.1.

Using estimated WBGT with the ASI panel

The impact of temperature degree days on output in Table 1 used temperature data rather than WBGT because measures of relative humidity are not available across all districts and over the ten year period covered by our manufacturing plant panel. An alternative is to approximate WBGT using estimates of average daily relative humidity from reanalysis models. This is not our preferred approach since reanalysis datasets are not normally calibrated to accurately estimate relative humidity - certainly not on a daily basis - and therefore this approach may increase rather than decrease measurement error, particularly since our estimation relies on temporal variation rather than cross-sectional comparisons.

Nevertheless these results make for a useful robustness check. Table A.2 presents results from models similar to those in Table A.1 using estimated

Table A.1: Effect of Temperature on Manufacturing Industry Output

	Dependent variable:				
	Plant Output Value			Log Plant Output Value	
	(1)	(2)	(3)	(4)	(5)
Annual Average Temperature	−.043***	−0.042***	−0.036***	−0.032***	−0.028***
	(0.013)	(0.012)	(0.012)	(0.010)	(0.010)
rainfall	0.013***	0.009***	0.006***	0.003	0.001
	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)
capital		0.386***	0.346***	0.384***	0.339***
		(0.010)	(0.009)	(0.003)	(0.006)
Plant FE	*Y*	*Y*	*Y*	*Y*	*Y*
Year FE	*Y*	*Y*	*Y*	*Y*	*Y*
Capital Controls	*N*	*Y*	*Y*	*Y*	*Y*
Worker Controls	*N*	*N*	*Y*	*N*	*Y*
Units	21,509	21,509	21,509	21,509	21,509
R^2	0.0076	0.4615	0.4876	0.6705	0.6595

Note:
1. *p<0.1; **p<0.05; ***p<0.01
2. Robust standard errors correcting for serial correlation and heteroskedasticity
3. Coefficients for models 1-3 are expressed as percentages of average output level.

WBGT measures calculated using Equation 1 and using daily long run average measures of relative humidity from the NCEP/ NCAR reanalysis datasets. Note that this output provides an average measure for each day but not temporal variation from year to year. This may be preferable in our context since this means temporal variation is still driven by the better measured temperature parameters. At the same time absolute temperatures are re-weighted across days of the year and across spatial locations to account for varying relative humidity levels.

Price Shocks and Power Outages

In this section we report results investigating the robustness of the non-linear response of output to temperature (reported in Table 1) to the inclusion of controls for power outage probabilities. We also test to see whether local input prices can be shown to respond to local temperature shocks to any significant degree. Table A.3 reports both results.

Column 1 provides results for a regression of a price index computed for each plant on temperature (controlling for plant fixed effects). Formally we estimate the model below where $P_{i,t}$ is the log of the plant input price index and other variables are the same as in Equation 7.

$$P_{it} = \alpha_i + \gamma_t + \omega K_{it} + \sum_{k=1}^{N} \beta_k D_k + \phi W_{it} + R_{it} + \epsilon_{it} \tag{8}$$

Note that the price index P_{it} is created only for ASI plants where input price data was reported. The price index is computed by averaging reported prices for the three most important reported inputs for each plant in each year and taking the log of the resulting price. Input price information is missing in about 28 percent of survey responses. In addition we also drop the top 2.5 percent and bottom 2.5 percent of plants within the computed input price distribution to remove outliers with very low or high reported input prices.

Table A.2: Effect of Wet Bulb Globe Temperature on Manufacturing Industry Output

	Dependent variable:				
	Plant Output Value			Log Plant Output Value	
	(1)	(2)	(3)	(4)	(5)
Annual Wet Bulb Globe Temperature	$-.042^{***}$	-0.044^{***}	-0.036^{**}	-0.036^{***}	-0.030^{**}
	(0.015)	(0.014)	(0.014)	(0.013)	(0.012)
rainfall	0.013^{***}	0.009^{***}	0.007^{***}	0.003	0.001
	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)
capital		0.386^{***}	0.346^{***}	0.390^{***}	0.339^{***}
		(0.010)	(0.009)	(0.006)	(0.006)
Plant FE	*Y*	*Y*	*Y*	*Y*	*Y*
Year FE	*Y*	*Y*	*Y*	*Y*	*Y*
Capital Controls	*N*	*Y*	*Y*	*Y*	*Y*
Worker Controls	*N*	*N*	*Y*	*N*	*Y*
Units	21,509	21,509	21,509	21,509	21,509
R^2	0.0076	0.4615	0.4876	0.6705	0.6595

Note:

1. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$
2. Robust standard errors correcting for serial correlation and heteroskedasticity
3. Coefficients for models 1-3 are expressed as percentages of average output level.

To control for power outages we download data made publicly available by (Allcott, Collard-Wexler and O'Connell, 2014) and reproduce their measure of state-year power outages that they construct from panel data on state-wise assessed demand and actual generation reported. We use this as a control for the intensity of power outages that might be experienced by all plants in a state and introduce this as an additional control in a specification similar to Equation 4. As Table A.3, Column 2 makes clear, our temperature response estimates seem robust to the addition of the outages control.

Table A.3: Testing for price shocks and robustness to power outages

	Dependent variable	
	Input Price Index	*Log Plant Output*
	(1)	(2)
Below 20°C	0.023	0.004
	(0.087)	(0.026)
20°C to 25°C	0.121	−0.039
	(0.081)	(0.026)
Above 25°C	0.050	−0.034*
	(0.051)	(0.018)
rainfall	0.002	0.002
	(0.007)	(0.002)
power outages		−0.067
		(0.087)
Plant FE	*Y*	*Y*
Year FE	*Y*	*Y*
Capital Controls	*Y*	*Y*
Number of Units	21,509	21,509
Mean Obs. per Unit	4.8	4.8
R^2	0.685	0.202

Note: 1. *p<0.1; **p<0.05; ***p<0.01
2. For details on the calculation of state power outages see Allcott, Collard-Wexler and O'C

Seasonal Patterns in Absenteeism

In interviews with weaving firm managers in Surat a frequent complaint related to the difficulty of hiring daily wage workers for industrial work during the summer months. Managers claimed that during the hottest months, daily wage workers preferred to go home to their villages and rely on income from the National Rural Employment Guarantee Scheme rather than work under the much more strenuous conditions at the factory. Some owners reported that they were actively considering the possibility of combating this preference for less taxing work by temporarily raising wages through a summer attendance bonus. However small scale weaving units operate on very tight profit margins and do not necessarily have the ability to raise wages very easily.

Figure A.2 in the Appendix suggests there may be some truth to this narrative. We see seasonal reductions in the attendance of daily wage weaving workers (Panel A), concentrated in high temperature months. These seasonal patterns are absent for the garment workers who have long term employment contracts (Panel B). It is possible that formal employment contracts - while reducing the costs to taking an occasional day of leave - significantly increase the opportunity cost of switching occupations for extended periods of time. Thus, when accounting for possible longer term responses to temperature, formal employment contracts might do better at retaining labour than daily wage arrangements. This is an area that would benefit from further research.

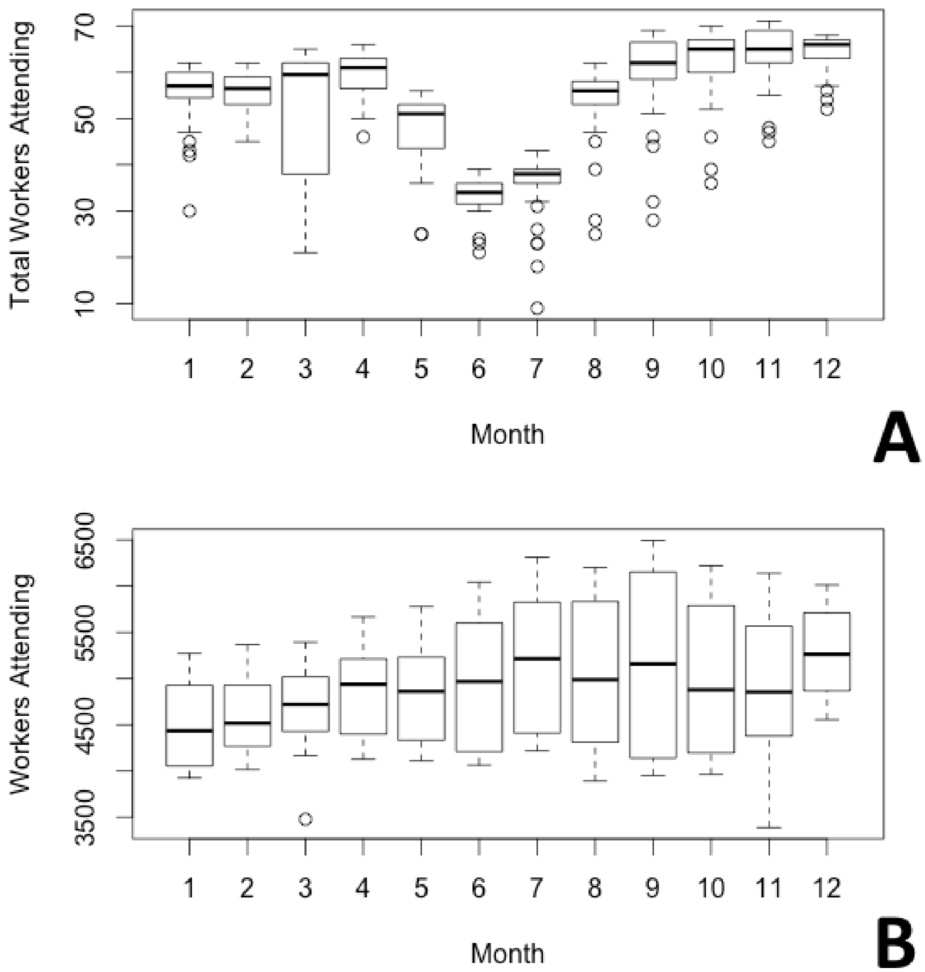

Figure A.2: Boxplots of worker attendance by month for daily wage workers in weaving units (Panel A) and regular workers in garment manufacture units (Panel B)