Abstract

This paper provides evidence that high temperatures impact manufacturing output by reducing worker productivity via heat stress. We collect micro-data and survey data from manufacturing plants in India and show that (i) output in labor-intensive settings decreases at high temperatures by 1-3 percent per degree celsius (ii) workplace climate control can provide effective adaptation and (iii) sustained temperature increases reduce worker attendance. This mechanism helps explain the recently documented negative correlation between temperature and GDP changes in poor countries. Climate-economy models that do not account for reduced labor productivity may underestimate costs of climate change. Keywords: temperature, worker productivity, industry, climate change, global warming.
JEL: Q54, Q56

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

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

September 8, 2014

*Sudarshan (corresponding author): Harvard Kennedy School and IFMR, anant.sudarshan@ifmr.ac.in; 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, 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.

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] This paper uses an annual nationally representative survey of manufacturing plants in India as well as primary daily data from a few industries to show that high temperatures reduce industrial output, and that one channel through which this occurs is heat stress on the job that reduces worker productivity. Our evidence also suggests that workplace climate control can provide effective adaptation in the workplace. We also provide some evidence indicating that sustained high temperatures (which will be felt outside the workplace as well) might reduce worker attendance. This relationship may be mediated by the nature of the wage contract.

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). To formulate adaptive responses to climate change it is essential that we determine not just whether climate variables are correlated with welfare but also precisely how this link is created. Similarly, to assess the value of mitigation requires a complete quantification of benefits, accounting for the different channels through which climate change might impose costs on society. Recent empirical studies have identified a robust - and plausibly causal - negative relationship between between high temperature years and developing country economic output (Dell et al., 2012; Hsiang, 2010). The challenge for researchers is to determine the mechanisms that might explain this stylized fact.

As noted above, one explanation involves appealing to reductions in agricultural output in response to temperature shocks (Lobell et al., 2011; Schlenker and Roberts, 2009; Mendelsohn and Dinar, 1999; Auffhammer et al., 2006).

[1] The Fourth Assessment Report (AR4, Working Group II) on Impacts, Adaptation and Vulnerability, stated 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).*"

Yet agriculture alone seems insufficient to explain the observed link between temperature and GDP, which remains present both in countries with an economically unimportant agriculture sector and in output from non-agricultural sectors (Dell et al., 2012). Other causal channels have therefore been suggested, including temperature impacts on mortality, conflict and worker productivity.

Hsiang (2010) examines economic output for a set of countries in the Caribbean and Central America and finds that output from the services sector decreases in high temperature years. Although it remains difficult to isolate specific mechanisms using aggregate data,[2] Hsiang points out that one mechanism consistent with his observations involves the direct impact of temperature on worker productivity, as predicted by physiological studies of heat stress in human beings.

In this paper, we provide the first ecologically valid evidence confirming that this mechanism operates in the manufacturing sector. In doing so, we provide a partial answer to the question of why increases in temperature above a threshold appear to reduce economic output in non-agricultural sectors. Our evidence indicates that the mechanism operates in labor-intensive industries without climate control, suggesting that it may account for the relation between temperature and developing-country GDP.

We begin by constructing a multi-year, nationwide panel dataset of manufacturing plants (factories) in India and directly estimate the impacts of annual temperature shocks on annual factory output. We show that these impacts are economically significant (an output decline of about one to three percent per degree Celsius), and have a magnitude and non-linear relationship to temperature that is consistent with physiological studies of heat stress when exposed to high wet bulb globe temperatures. We also find that temperature impacts on plant output seem most acute in sectors where labor's share of output is high and where electricity intensity (used as a proxy for the likelihood of climate control) is low.

[2]For instance, because the setting for this study is a region of the world heavily dependent on tourism, it is possible that demand shifts coincident with temperature shocks might explain the economic effects found here.

We then augment this nationwide panel with independently collected, daily production data from different manufacturing settings. These independent datasets allow us to directly observe high-frequency, worker-level performance outcomes in the workplace. We show that daily temperature is non-linearly associated with decreased worker output with significant reductions primarily occurring when temperatures (more accurately, wet bulb globe temperatures) are high. We also show that the temperature-output link is broken for production settings that are largely mechanized or climate controlled. These 'no-effect' cases are consistent with our hypothesized mechanism, namely, that heat stress on the job can result in productivity declines under elevated surface temperatures. Lastly, we examine worker attendance records in these firms and show that high temperatures sustained over a few days lead to a 2 to 4 percent increase in absenteeism. We also present evidence suggesting that this response may vary with the nature of labor contracts. When the cost of absenteeism is higher for workers, the link between temperature and absenteeism is weaker.

That said, global environmental change is not the only reason to be concerned about temperature-productivity interactions. An emerging strand of research in the economics literature has begun to document how environmental factors can directly influence productivity (Zivin and Neidell, 2012). Our work contributes to this strand of research. Recent evidence has also rigorously documented that local urban temperatures may be significantly elevated owing to urban heat island effects (for example Mohan et al. (2012) study heat islands in the Indian capital of New Delhi). This provides an immediate motivation to understand the impacts of such locally elevated temperatures on economic activity.

The remainder of this paper is organized as follows. In Section 1 we summarize the physiological evidence on heat stress and provide a framework describing how these physiological effects might impact economic output. In Section 2 we describe our data sources. In Section 3 we present evidence of temperature impacts on manufacturing using plant-level output data from India's Annual Survey of Industries and annual variations in temperature. In

Section 4, we present evidence relating daily temperature to daily worker productivity from one firm in each of three industries – weaving, garments, and steel. We then quantify the importance of temperature-productivity effects in the context of climate model predictions for India in Section 5. We conclude in Section 6.

1 Theory and Mechanisms

The physics of how temperature affects human beings is well known. The physical exchange of heat between the human body and surrounding air is fundamentally related to health because in order to maintain normal body temperatures, the human body must dissipate the heat it generates internally to the ambient (Parsons, 1993). When energy is expended while working, internal heat generation increases and correspondingly greater rates of heat loss become necessary. If this balance cannot be maintained at a given activity level, it becomes necessary to reduce the rate at which energy is consumed or to suffer the adverse consequences of over-heating including heat strokes (Kjellstrom et al., 2009; ISO, 1989). It is therefore plausible that at elevated temperatures or high humidity, heat stress might begin to reduce productivity even before it becomes a significant health hazard.

The primary mechanism the human body uses to dissipate heat is the evaporation of sweat. The efficiency of such dissipation depends primarily on ambient temperature but also on humidity and wind speed. These ambient parameters can be encapsulated in various ways to form indices capturing the threat of heat stress (Parsons, 1993), the most commonly accepted index being the Wet Bulb Globe Temperature (ISO, 1989).

In indoor conditions the Wet Bulb Globe Temperature (WBGT) is determined largely by two variables – temperature and humidity.[3] Direct measurement of WBGT requires instruments that are not in common use, but Lemke and Kjellstrom (2012) show that an accurate approximation can be derived from temperature and humidity using a formula reproduced in Equation 1.

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

In this paper, wherever data permits, we will use this composite measure of WBGT instead of the ordinary dry bulb temperature.

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

In Equation 1, $WBGT$ is measured in $^{\circ}C$, T_A is the air temperature and ρ is the water vapour pressure estimated from the relative humidity (RH).

The literature on heat stress also suggests that the response of human beings to temperature (or wet bulb globe temperature) is not linear. At very cold temperatures, productivity (or at least comfort) might increase with temperature and at moderate levels, temperature variations might have little impact. At higher levels however, heat stress should become progressively more severe. While the precise shape of the dose-response relationship is not well known or even necessarily deterministic, laboratory evidence is consistent with this pattern. Hsiang (2010) reports that a meta-analysis of laboratory studies suggest reductions in the ability of human beings to carry out ergonomic and cognitive tasks by approximately 1-2 percent per 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.[4]

One final point relates to the time-scales at which temperature may affect human beings. The effects of exposures to high temperatures can be expected to be visible on short time scales - within minutes or hours. At the same time these impacts are unlikely to disappear when temperature changes are sustained (absent adaptive actions taken to reduce exposure). Thus it is plausible that sustained temperature differences between populations might lead to sustained differences in the productivity of labor and also that these differences should be detectable using both short run and more sustained temperature variation. Our identification strategy exploits the fact that this short-run response sets temperature apart from many other environmental stressors (in-

[4] Although laboratory evidence cannot directly inform us about temperature-productivity relationships in the workplace where incentives and conditions can be very different from those in the lab, it does provide a benchmark.

cluding air pollutants) as well as economic factors that might be linked with temperature (such as economic spillovers from changes in agricultural output).

1.1 Worker Performance and Economic Output

The physiology behind heat stress is well known. Nevertheless it is not obvious how significant temperature might be as a factor influencing productive economic output. Daily workplace activity does not normally require exertion nearing physical limits. This is particularly so in formal, skilled work in the manufacturing and service sectors, as distinct from purely manual and unskilled labor that might play a significant part in the construction or mining sectors.[5] Secondly, most labor in manufacturing (or services) can be expected to take place indoors or in shielded conditions. These work conditions provide some protection from ambient temperatures even absent air conditioning. Lastly, the *economic* impact of reductions in worker productivity may be very different from the physiological impact. The marginal costs of a reduction in the physical or cognitive effectiveness of workers engaged in high value-added activities may be very high. Conversely the marginal cost of decreased effectiveness may be minimal in the case of low value-added tasks.

These distinctions can be easily understood in the context of a simple production model. Consider a plant with output given by a Cobb-Douglas production function as below, where Y is total output, L, E, K represent labor, energy and capital inputs and A is the total factor productivity. L in turn is written as a function of input labor L_o and T_I, the indoor or workplace temperature (or wet bulb globe temperature). Further let $T_I = a + bT_A$ denote the dependence of workplace temperature on the ambient T_A. Adaptive technologies such as air conditioning for instance, might drive b towards zero, breaking the link between T_I and T_A. Then,

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

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

Let

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

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

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

In other words, temperature shocks may not affect productivity when temperatures are moderate. At higher temperatures ($T_A > (T_C - a)/b$), Z declines with temperature. This decline is greater when α is large, which might represent a firm where the value added by labor is high. It is also larger when b is large, i.e when the relationship between the ambient temperature T_A and workplace temperature T_I is strong.

Taken together this suggests three empirical features we might expect to see in the response of manufacturing plant output to temperature.

Test 1 If manufacturing output responds negatively to temperature primarily because of temperature effects on worker productivity then this response should occur mostly at high temperatures (above $25°C$).

Test 2 Temperature impacts on output should be higher where the share of value added by labor is high.

Test 3 Temperature impacts on output should be higher where climate control and cooling is likely to be limited.

These tests also help distinguish between different mechanisms through which temperature might influence manufacturing output. For instance, one might hypothesize that temperature could be correlated with industrial output due to some form of spillover from the agricultural sector.[6] However spillovers from agriculture might suggest temperature response patterns that would not

[6] Burgess et al. (2011) suggest that some of the observed health impacts of temperature may partially owe to agricultural productivity shocks

necessarily match those described in Tests 1-3. Agricultural growing seasons in India take place during periods where temperatures are moderate and one of the two primary growing seasons is in the winter. Thus if non-agricultural sectors respond to temperature shocks primarily through agriculture related economic spillovers then these impacts should be highest when temperature shocks occur during the cooler temperatures found in the growing season. Similarly, climate control in a plant would not influence the temperature sensitivity of output shocks if this were operating only through agricultural spillovers.

With this background, we follow a two-part empirical strategy to determine whether temperature matters for manufacturing productivity. First we examine plant-level data from the Annual Survey of Industries in India. Next we zoom in to examine worker-level output at a daily frequency, using detailed micro-data that we collected from a number of sites located in different regions of India and in different industries. Our objective will be to verify whether the predictions of the simple model we have outlined in this section are indeed reflected in multiple independent datasets.

2 Data Sources

We use three sources of data for our empirical work. The first is a nationally representative, annual panel data set of individual manufacturing plants in India. The second is a collection of independent datasets, containing daily production measures over multiple years, that we put together from different manufacturing environments. The third is meteorological data, including surface temperature and rainfall.

2.1 Annual Manufacturing Plant Survey Data

Our data on plant level output comes from the Annual Survey of Industry (ASI) carried out by the Government of India. The ASI is a survey of individual manufacturing plants in every state within India. The population eligible to be surveyed consists of all industrial plants registered under India's

Factories Act. Each annual cross-section is a census of registered plants that employ over 100 workers (not including short term contract labor) together with a random sample covering 18 percent of the plants below this threshold.

The survey is intended to capture critical variables relating to production inputs and outputs (in both physical and monetary units and including energy inputs), annual income and expenditures under various heads, labor utilization, wages and annual man-days worked. The measure of output that we use is calculated by multiplying the reported market price of products manufactured with the production quantity in each year.

The ASI has historically taken the form of repeated cross-sections. However for survey years between 1998-99 and 2007-08, survey micro-data may be purchased with a panel identifier so that it is possible to identify repeated observations on plants across years.[7] The panel is unbalanced since only large firms with over 100 employees are surveyed every year, with smaller firms appearing in multiple years only if they are surveyed.

We carry out a few data-cleaning operations before analysis primarily to transparently eliminate outliers. These steps are detailed in Appendix Section A. Overall we have data from 21,525 manufacturing units that are observed at least three times within the panel. Figure .1 in the Appendix shows the distribution of the number of surveyed plants in different districts in the country.

One limitation 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 Indian 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 with greater capital investment and a certain minimal level of adaptive capacity.

[7] The panel dataset does not provide the geographical location of a factory. However for these years an alternative version of the dataset was made available containing district identifiers without the panel identifier. We purchased both versions and then matched observations across the two data views to generate a dataset with both a plant identifier and a location identifier for each observation.

2.2 Worker Daily Output Data

The ASI provides a wide breadth of sector and regional variability in manufacturing plants. However, all measures of economic output are observed only annually and for a manufacturing unit as a whole. As a consequence it is impossible to directly observe worker output, fundamentally the quantity that most interests us here.

We therefore complemented the ASI by collecting *daily* worker and plant level output data from three different (and independent) manufacturing settings. These units are chosen partly because they belong to industrial settings where high-frequency output data can be obtained and also because they represent very different workplace contexts. Figure 1 shows that the case studies include production settings where ambient temperature might matter a lot as well as contexts where it should matter less.

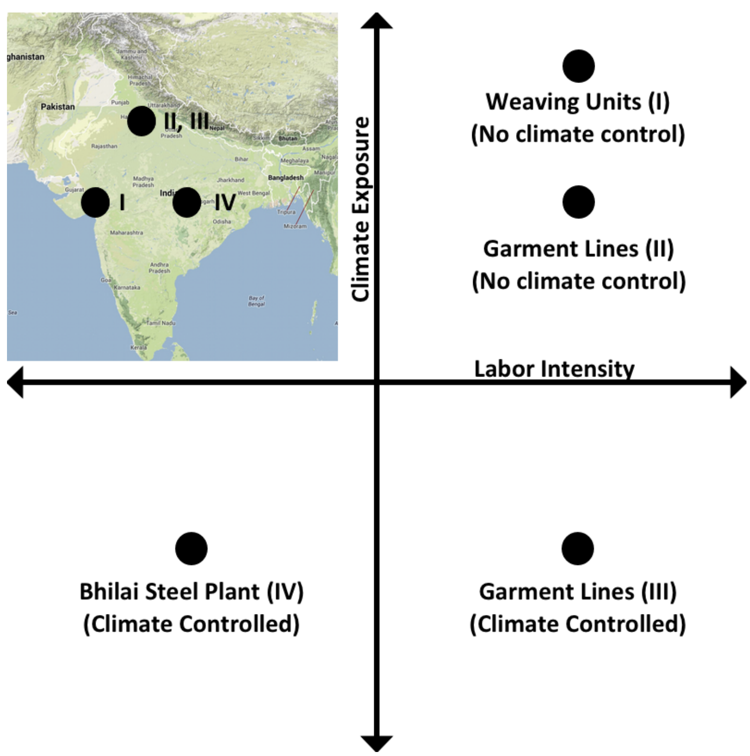

Figure 1: Case study sites span a variety of operating conditions.

2.2.1 Weaving Units in Surat

The first setting involves three weaving units located in the city of Surat in the state of Gujarat in western India. We assembled daily worker-level output and attendance data for the financial year 2012-2013 that tracks 147 weaving workers. These weaving firms use a labor-intensive manufacturing process, where temperature control is limited to the use of windows and some fans. Figure 2, Panel B shows a picture of the production floor.

Our choice of the textile sector (more precisely, weaving) is motivated by a number of factors that make it well suited for our purposes. The textile sector is estimated to make up about 14 percent of India's total industrial production (and about 3 percent of GDP) and to contribute to about 27 percent of foreign exchange from exports. Weaving in particular is a major source of industrial employment - The 2011 Ministry of Textiles Annual Report estimates that the power-loom sector employs 5.7 million people.

The workforce in mechanized weaving units consists of semi-skilled labor. Labor is often non-contractual and workers are paid piece rates. This makes it possible to collect high frequency output measures at the worker level. In the weaving units of Surat for instance, workers are paid only for days when they show up to work. Payments are made on the basis of a simple measure of worker-level physical output, namely meters of cloth produced, multiplied by a per meter payment (a little over INR 2.00 per metre in Surat)[8]. In the firms we study, this cloth output is also essentially the final output for the plant and is then sold in wholesale markets or to dying and printing firms. Thus worker output directly corresponds to plant revenue.

It is also interesting that while weaving is labor intensive it is not physically strenuous. A weaving worker is primarily responsible for operating mechanized looms (each loom can be regarded as a work station). A worker must walk up and down between work stations,[9] occasionally adjusting alignment, restarting

[8]Indian minimum wage laws are both poorly enforced and not legally binding on small firms. We can therefore ignore complications introduced by payment non-linearities at a minimum wage lower bound as in Zivin and Neidell (2012).

[9]A single worker typically works on about 6-12 looms.

feeds when interrupted and making occasional corrections as needed. The fact that this work is not physically demanding is important because this is very different from settings where heat stress is known to be an occupational health concern (e.g. mining, see Wyndham (1969)). Finding temperature impacts on worker output in this setting is therefore more likely to imply that this may be an important mechanism affecting productivity across a range of industries.

2.2.2 Bhilai Steel Plant

The second setting we investigate is one of the largest integrated steel plants in India, located at Bhilai in the state of Chattisgarh. The Bhilai steel plant manufactures a variety of steel products. For all final products, steel is first formed into rectangular blocks called blooms that can then be shaped further.

The bloom production line consists of a furnace section, a steel milling section, a mechanized hot saw area and finally a cooling yard. When a bloom is successfully produced it is said to have been 'rolled'. When faults occur, the product must be discarded (the steel may be later reused) and this is referred to as a 'cobbled' bloom.

A key output variable tracked by the plant management is the number of blooms rolled. We also collect measures of cobbled blooms per shift (there are three shifts in the day) and line delays. Lastly we use data on worker attendance. We obtained this data for every day for the period 1999-2008.[10]

The bloom production line we study is interesting because the production process is heavily mechanized and capital-intensive. Figure 2, Panel A shows part of the production line where steel blooms are being cast. Many workers on this line work out of air-conditioned glass cabins allowing for remote operation of production line machinery. In other words, referring back to Section 1, the Bhilai Steel Plant is a good example of a capital intensive and mechanized production process with some use of climate control. Therefore, if a temperature-output relationship exists because of the performance of labor at high temperatures, then this factor should be very limited in this

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

setting. Conversely, if other mechanisms drive a relationship between temperature and output, then the Bhilai Steel Plant might nevertheless see output shocks coincident with temperature shocks.

2.2.3 Garment Manufacturing Units

The third setting is a cluster of large garment manufacturing plants, all operated by the same firm. Six of these units are located in the National Capital Region (NCR) of North India, one in Hyderabad and one in Chhindwara. These plants produce finished clothes for multiple international and national brands, largely for export. Production activities and management systems are similar across all units and involve cloth cutting, sewing, embroidery, finishing and washing.

For our analysis here, we focus on sewing lines, each of which consists of a group of 10-20 workers, all carrying out specific operations that combine to create a finished clothing item (or sub-item). Lines are highly stable groups of workers who remain assigned to the same line as a general practice, although the garment manufactured by a given line changes periodically. We collect line-level data on the hourly productivity of each line over a two year period from April 2012 to March 2014. Figure 2, Panel C shows a picture of the production floor.

Productivity in this context requires a little bit of discussion. Unlike weaving firms where a single output variable - meters of cloth produced - fully encapsulates worker productivity, the garment setting we study is more complicated because output rates will depend on the type of garment being manufactured and the complexity of operations carried out by each line. To track productivity, we use two variables defined by the management of the firm in question: Actual Efficiency and Budgeted Efficiency.[11]

Budgeted Efficiency is an hourly production target that is set for each line by the management of the firm. It depends on the product being manufactured

[11] The export garment sector is characterized by both significant competitive pressures and a production function where labor productivity is extremely important. As a consequence large garment manufacturers track worker output in sophisticated ways, providing an excellent testbed for research into determinants of productivity.

and is calculated based on an analysis of the time taken to complete each individual operation going into the line product[12]. The Actual Efficiency is the per unit rate of throughput actually attained by the line every hour. We use the Actual Efficiency, controlling for Budgeted efficiency, as a measure of line productivity (productivity for the cluster of 10-20 workers taken together).

During the period for which we obtain data, the firm we study was in the process of installing central cooling systems in its different plants. We thus have an exceptional opportunity to test whether workplace climate control can reduce the link between temperature changes and productivity shocks. This test is important both to validate the hypothesis that workplace heat stress may cause productivity declines and as important evidence on the effectiveness of adaptation. The cross-country temperature-output literature as in Dell et al. (2012) finds that country GDP is sensitive to temperature shocks only in developing countries. To the extent that climate control technologies could mitigate the effects of high temperatures on labor, this pattern might reflect the relative prevalence of air-conditioning in developed vs. developing countries.

In our dataset five manufacturing units (located in the NCR) had production floors that had been equipped with at least one air-washing system. Air washers enable temperature control and dehumidification and therefore the ability to manage wet bulb globe temperatures effectively. One manufacturing unit in the NCR did not have air-washing installed until 2014 and workers in this setting had access only to fans or evaporative coolers (the latter may actually *increase* humidity and decrease comfort under high humidity conditions). We observe two units outside the NCR (in Hyderabad and Chhindwara in South India), that also lacked climate control.

We take advantage of this variation by comparing the responsiveness of worker output to temperature shocks in settings with and without effective climate control. The differential assignment of cooling to plants is admittedly not random. Nevertheless, comparisons of the temperature sensitivity of

[12]The best case performance rate is computed by having the desired operations completed by a special line of 'master craftsmen'.

output in otherwise highly similar units with the same outputs and identical management is arguably informative in the present context. One use of these comparisons is to determine whether firms could mitigate the impact of temperature on production by investing in workplace cooling. Since workers might continue to be exposed to uncomfortable temperatures at home the answer to this question is not ex-ante obvious but does carry important implications for the possibility of adaptation.

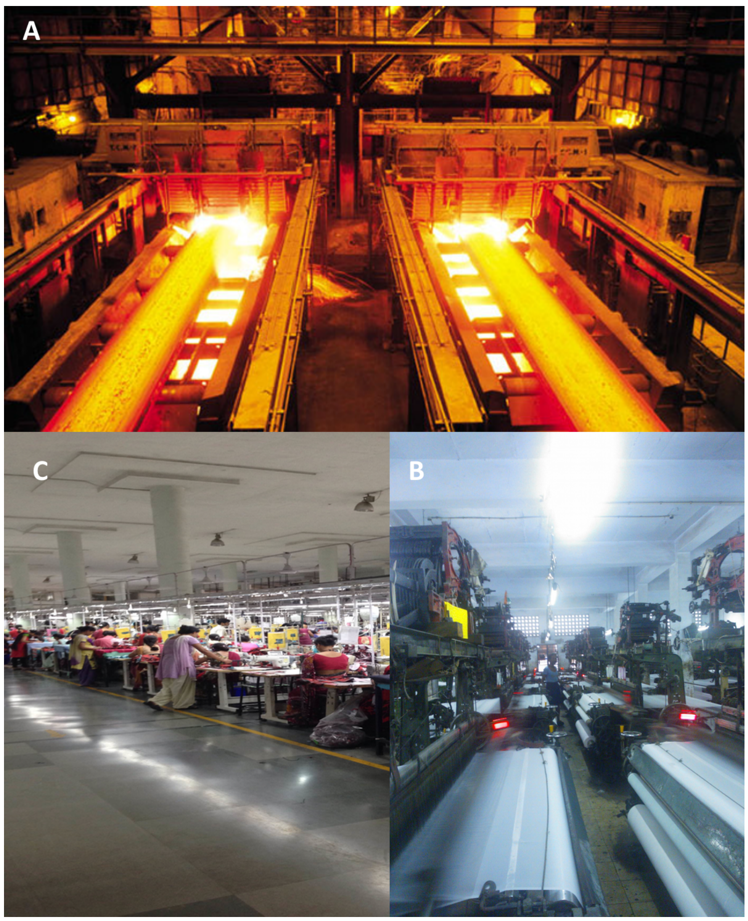

Figure 2: Production floor images from A: Steel mill, B: Weaving units, C: Garment manufacture plant

2.2.4 Meteorological Data

The meteorological data used in this paper comes from two sources. The first is a $1^{\circ} \times 1^{\circ}$ gridded data product released by the Indian Meteorological Department (IMD) which provides daily temperature and rainfall measures interpolated from the IMD's monitoring stations across the country. Weighted averages of the gridded data were used to get district-level measures of these variables for the 609 districts in the country.[13] A key strength of this dataset is that it is based on data from quality controlled ground-level monitors and not sub-sampled measures from regional climate models or reanalysis data (see Auffhammer et al. (2013) for a discussion of some of the concerns that arise when using temporal variation generated from climate models).

The second weather data source is local weather station readings of temperature, humidity and rainfall. We use this data to construct WBGT from Equation 1 and to control for rainfall in the case study sites, each of which is matched to a quality controlled weather station located in the same city.[14]

Unfortunately creating a nationwide WBGT measure using Equation 1 is difficult because reliable time series data on relative humidity across India is not easily available. For this reason in Section 3 we use dynamic variation in temperature alone to estimate the effect of heat on industrial output in our nationwide analysis.[15]

[13] The value of x for a district is the weighted average of x from all grid points within a 200km radius of the district centroid with weights inversely proportional to the squares of the (great-circle) distances between grid points and centroid. The average district area is about 4000 square km while the grid spacing is about 110 km.

[14] The one exception to this is our temperature measure for the steel plant at Bhilai. The absence of a quality-controlled public weather station near Bhilai, means that the only available option to create a consistent temperature record for Bhilai is to use the district measure for Durg district from the IMD gridded dataset.

[15] Table 5 in the appendix provides results from a robustness check using humidity data from reanalysis models to approximate WBGT for all districts.

3 Temperature and Manufacturing Output

We identify the impact of temperature on India's manufacturing output through a comparison of year-to-year variation in a single plant's output with year-to-year variation in temperature. This ensures that we can isolate the effect of temperature, independent of other variables that might be associated with cross-sectional temperature differences between units but might affect independently affect output (such as altitude for example). We control for price shocks or other macro-economic variables that might influence the sector as a whole through the use of time fixed effects.

Implementing this strategy using average annual temperatures as an explanatory variable is straightforward. However to determine whether heat effects of workers might cause temperature impacts on manufacturing output it is useful to estimate a more general non-linear specification. To do so we exploit the fact that although plant output is available to us only on an annual basis, our temperature records are available for every day in the year.

We let $V(T_d)$ represent the daily output of a manufacturing unit as a function of daily temperature, T_d. In general $V(T_d)$ may be represented as follows

$$V(T_d) = V(T_o) + \int_{T_o}^{T_d} \frac{\partial V}{\partial T} dT \tag{3}$$

We may approximate the general non-linear response to temperature by specifying a stepwise linear function of production in temperature similar to Hsiang (2010) and Burgess et al. (2011)). Thus we obtain,

$$\bar{V}(T_d) = \bar{V}(T_0) + \sum_{k=1}^{N} \beta_k D_k(T_d) \tag{4}$$

Here

$$D_k(T_d) = \int_{xl_k}^{xu_k} \mathbf{1}[\mathbf{T_d} \leq \mathbf{x}] dx \tag{5}$$

where $\mathbf{1}[...]$ represents an indicator function which is 1 when the statement in brackets is true and 0 otherwise. In other words $D_k(T_d)$ measures the degree days within the year that fall in a given temperature bin. Provided we assume that $V(T_d)$ does not vary with the time of year, the formulation above is equivalent to estimating annual production as a piecewise linear function of degree days in different temperature bins where the coefficient associated with each degree day bin represents the change in production caused by an increase of one degree-day within that bin. We can therefore write annual output V_t as a function of degree days D_k as follows

$$V_t = V_0 + \sum_{k=1}^{N} \beta_k D_k \tag{6}$$

Because we observe district temperatures at a daily level throughout the years of our study it is possible to calculate a degree day measure associated with each year. We may then estimate a regression of the form

$$V_{i,t} = \alpha_i + \gamma_t + \omega K_{i,t} + \sum_{k=1}^{N} \beta_k D_k + +\phi W_{i,t} + \theta R_{i,t} + \epsilon_{i,t} \tag{7}$$

where $V_{i,t}$ is the value of output produced by plant i during financial year t, α_i is a fixed effect representing average level of output for each manufacturing unit, γ_t are time fixed effects capturing national changes in manufacturing output year to year, $K_{i,t}$ is total working capital, and $R_{i,t}$ is a control for rainfall in mm.

Equation 7 allows us to see if β_k, the effect on output of a 1° rise in WGBT in the kth temperature bin, varies over temperature bins. If temperature reduces industrial productivity through its impact on workers, we should expect to find the hypothesis $\beta_k = 0$ true for low temperatures and to see negative values of β_k for higher degree-day bins (see Section 1 for details).

$K_{i,t}$ is the total working capital available to the plant at the start of the financial year (a measure that includes cash generated from the previous years output less expenditures). Cash on hand at the start of the financial year is converted by the plant into labor wages, raw material purchases or energy inputs and these in turn are transformed via the factory production function into outputs. Thus being able to explicitly control for liquid capital at the start of the year helps us to cleanly identify the impact of temperature realizations on output produced in the year, controlling for a fundamental measure of inputs available at the start of the year. Working capital at the start of the financial year is also plausibly exogenous to temperatures experienced during the year and to realized labor productivity. This would not be true of labor, energy or raw material expenditures actually realized during the year. For instance, in our case study of weaving workers in Surat (Section 4 we note that workers appear to produce smaller amounts of woven cloth on high temperature days. These productivity declines can be expected to translate to lower labor expenditures (since wages are linked to output) and to lower raw material use (since finished cloth is mechanically correlated with raw cloth inputs).

Equation 7 is estimated and the results reported in Table 1 column 2. Multiple specifications are presented as a robustness check. We estimate models that include the reported total number of workers W_{it} on the right hand side. We use both the absolute output as well as logged output as outcome variables. We also estimate a model using the log of output per worker as a dependent variable (this outcome may be very noisy since ASI employment numbers are sometimes missing and may under-report contract labor). Overall we find fairly strong evidence that output is non-linearly impacted by temperature shocks. Furthermore, this link seems strongest for changes in the highest temperature bins, with output falling by 3-4% per °C above 25 °C.

The literature examining the impact of temperature on country level output has normally used average annual temperatures as an explanatory variable (Dell et al., 2012). However the non-linear response implied by Table 1 suggests that it is the degree-day model of Equation 7 that should be of primary interest rather than simply the average estimate of the impact of temperature on productivity. One reason for this is that historic variations in temperature do not necessarily correspond to the forward looking predictions of climate models. For India, these models predict a significant increase in the number of extreme temperature days and not a secular increase in temperatures over the year (see Section 5 for more details).

Nevertheless it is straightforward to estimate the average impact of annual temperature on output and we present results in Appendix Table 4. We find that across all models the coefficient on temperature remains negative and significant. In the most conservative specification, with logged output and both capital and worker controls we obtain a point estimate suggesting a 2.8 percent decrease in output for a one degree change in average annual temperature (aggregated over all days in the year). It is interesting to compare this estimate with those for aggregate GDP from the recent studies on this issue. Dell et al. (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.

3.1 Heterogeneity in Impact: Labor Value Added

In Section 1.1 we argued that if temperature shocks reduce worker productivity we should expect that this effect should result in percentage declines in production that are highest in manufacturing sectors with a high value added per worker.

In order to test this hypothesis we require a measure of the value added by labor within a particular sector. The approach we use is to calculate for each plant in our dataset the ratio of wages paid over every year to output in that year. This quantity is not the same as the marginal value of an additional unit of labor but we use this to proxy for firms where labor costs are a significant share of output value and presumably therefore, labor adds a significant amount of value. We discretize this variable creating a dummy variable for every plant identifying the quartile of the wage ratio distribution to which the unit belongs. Next we regress the log of factory output on mean temperature interacted with the wage share dummies as in Equation 8. This allows us to flexibly examine whether there exists a relationship between temperature effects on output and the importance of labor, in particular whether plants in

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

	Dependent variable:						
	Plant Output Value			Log Plant Output			Log (Output/Worker)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Below 20°C	0.015	0.013	0.053*	−0.004	−0.006	0.004	−0.007
	(0.025)	(0.024)	(0.027)	(0.024)	(0.023)	(0.026)	(0.025)
20°C to 25°C	−0.039	−0.034	0.038	−0.056***	−0.046*	−0.039	−0.038*
	(0.026)	(0.026)	(0.032)	(0.022)	(0.022)	(0.026)	(0.023)
Above 25°C	−0.069***	−0.056***	−0.048**	−0.039***	−0.031**	−0.034*	−0.021
	(0.016)	(0.015)	(0.021)	(0.014)	(0.013)	(0.018)	(0.014)
rainfall	0.008***	0.006**	0.009***	0.002	0.001	0.002	0.000
	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)
capital	0.382***	0.343***	0.394***				
	(0.010)	(0.009)	(0.010)				
log(capital)				0.383***	0.303***	0.395***	0.195***
				(0.006)	(0.006)	(0.007)	(0.006)
workers		0.002***					
		(0.0001)					
log(workers)					0.415***		
					(0.008)		
shortages			−0.113			−0.067	
			(0.108)			(0.087)	
Plant FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Year FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Worker Controls	*N*	*Y*	*N*	*N*	*Y*	*N*	*N*
Units	21,509	21,509	21,509	21,509	21,509	21,509	21,509
R^2	0.249	0.291	0.257	0.196	0.272	0.202	0.092

Note:
1. *shortages* is a proxy for state outages from Allcott et al. (2014)
2. Coefficients for models 1-3 are expressed as proportions of average output level
3. Daily maximum temperature is on average $6°C$ above the daily mean temperature
4. Robust standard errors correcting for serial correlation and heteroskedasticity
5. *p<0.1; **p<0.05; ***p<0.01

the highest quartile of labor wage shares show greater output changes during high temperature years.

$$log(V_{i,t}) = \alpha_i + \gamma_t + \beta T_{i,t} \times VA_i + f(R_{i,t}) + \epsilon_{i,t} \tag{8}$$

Here VA_i is the constructed dummy variable. We are interested in the coefficients of the interaction between temperature $T_{i,t}$ and VA_i. Table 2, Panel A summarizes our estimates. While not a perfect implementation of Test 2 in Section 1.1, we find that plants with a higher wage share of output seem significantly more negatively impacted by temperature shocks.

3.2 Heterogeneity in Impact: Electricity Inputs

Air cooling is one obvious adaptive measure that a manufacturing plant could use to mitigate against any temperature effects on workers. We investigate the effect of air cooling more directly in Section 4 but attempt an indirect investigation using our annual survey data. The ASI surveys do not report whether or not a plant uses climate control. However we do observe reported expenditures of electricity. We create a new variable for each plant defined as the ratio of electricity expenditures to total cash on hand at the start of the year (capital). Because air cooling is an extremely electricity intensive technology, plants with high spending on electricity during the year (as a fraction of cash on hand at the start of the year) are arguably more likely to be using climate control.

We look for differences in temperature sensitivity interacted with dummy variables for each quartile of electricity intensity. Assuming this quantity is correlated with the use of air cooling, we might expect to see temperature sensitivity decrease for plants with higher electricity expenditures. Indeed we do see this pattern in the data, albeit somewhat imprecisely estimated for the model using log of plant output as the dependent variable. Our results are reported in Table 2, Panel B.

Table 2: Impact of temperature on plant output (variation by wage share and electricity intensity)

	A: Temp. Effect on Output by Wage Share Quartiles	
	plant output	log(plant output)
	(1)	(2)
meant	−0.038***	−0.042***
	(0.013)	(0.010)
wage share quartile X meant		
Q2	−0.007	0.0001
	(0.006)	(0.004)
Q3	−0.023***	−0.004
	(0.007)	(0.005)
Q4	−0.031***	−0.019**
	(0.008)	(0.009)
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	0.302	0.364

Note: 1. Fixed effects and covariates including quartile dummies omitted for clarity
2. Robust standard errors
3. *p<0.1; **p<0.05; ***p<0.01

	B: Temp. Effect on Output by Elec. Share Quartiles	
	plant output	log(plant output)
	(1)	(2)
meant	−0.064***	−0.043***
	(0.013)	(0.012)
elec. share quartile X meant		
Q2	0.018***	0.009*
	(0.005)	(0.005)
Q3	0.031***	0.011*
	(0.007)	(0.007)
Q4	0.032***	0.006
	(0.008)	(0.007)
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	0.260	0.233

Note: 1. Fixed effects and covariates including quartile dummies omitted for clarity
2. Robust standard errors
3. *p<0.1; **p<0.05; ***p<0.01

3.3 Local Input Price Shocks

Equation 2 provides a simple way to think about how temperature might influence output through its impact on labor. However it is possible that temperature shocks might change the price of other inputs - especially inputs linked to agricultural output - and that the output shocks we see might reflect input price changes rather than changes in the effectiveness of labor. Although any price shocks that affect the broader population of manufacturing plants will already be captured by time fixed effects, these controls might not account for local input price shocks that vary with local average temperatures.

The ASI surveys allow us to directly test this proposition 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 mean reported per unit price for the three most common inputs used by a plant. We then use fixed effect regressions similar to Equation 10 to test whether the price index for a given plant changes in years when local temperatures rise. We find no evidence that the price index we estimate changes (increases) in high temperature years or an increase in high temperature days within the year. These results do suggest that local input price shocks are not an important explanatory factor in this setting. We report our results in Appendix Table 6.

3.4 Power Outages

A possible confounding factor in these panel regressions is the impact of power outages - which may also be correlated with temperature - on productivity. In Section 4 we use daily worker output from plants with power back-up to estimate temperature impacts in a setting without this concern. Nevertheless, there are a few tests that can be carried out using the ASI data which suggest that outages are unlikely to be a major factor driving our results.

First 2 suggests that more electricity intensive plants seem less - not more - sensitive to temperature. Second, we introduce the power outages measure derived in Allcott et al. (2014) as a control in some of our specifications (Table

1, Columns 3,6) and find our point estimates remain similar[16].

4 Temperature and Daily Worker Output

In Section 3 we showed that output from individual manufacturing plants decreases with an increase in the number of high temperature days in the year. We also find that output associations with temperature are reduced in plants which use greater electricity (a proxy for the use of climate control) and in plants where the labor share of output is low. These patterns are consistent with a causal mechanism that involves the impact of temperature on worker productivity, through the physiological mechanism of heat stress.

However, ASI data does not allow us to directly observe productivity measures at the level of an individual worker or groups of workers. In this section therefore, we examine how daily worker or plant level output changes with daily temperature shocks. We compare the sensitivity of worker output to temperature in the presence of climate control, to the productivity response in plants without climate control.

The use of daily output measures also serves as an important test of the mechanism we propose here. Heat stress has a physiological basis which would predict that temperature effects should become apparent over fairly short periods of exposure. This very quick response is characteristic of this mechanism, since other proposed explanations for the impact of temperature on non-agricultural sectors (conflict, economic spillovers, demand shocks) are unlikely to be detectable at very short time scales.

We begin by linking daily output data from each of our case study sites with a measure of local ambient wet bulb globe temperature (WBGT) from nearby climate stations or from historical weather data products made available by the IMD (the latter only in the case of the Bhilai steel plant). Equation 1 is used to calculate a daily WBGT measure from temperature and humidity

[16]While useful as a robustness check, measures of power shortages may be 'bad controls' in general because power outages are also plausibly correlated with both temperature and the use of climate control. Therefore adding these controls may not only control for the direct impact of power shortages on output, but also the impact of temperature on output.

records. We utilize day to day variation in WBGT to estimate the impact of changing temperature on worker output. By obtaining high frequency output data we are able to control for individual fixed effects at the level of the worker (weaving firms) or of a small group of workers (garment units).

For weaving units and garment manufacturers (I and II in Figure 1) we estimate coefficients of the linear model below through ordinary least squares. Note that the model below is estimated on daily output measures, and in those cases where hourly data is available (garment manufacture units), we first aggregate up to daily measures.

$$log(Y_{i,d}) = \alpha_i + \gamma_M + \omega_W + \beta_k WBGT_d \times \mathrm{D_k} + \theta R_{i,d} + \epsilon_{i,d} \tag{9}$$

Here $Y_{i,d}$ is worker output for worker i on day d (or in the case of garment manufacturing sites, output for worker line i on day d). For weaving workers, output is measured in meters of cloth. For garment manufacture lines, output is measured as the ratio of actual efficiency to budgeted efficiency for the product being manufactured.

We use both output and log output as the dependent variable (the latter being less sensitive to outliers). α_i is a worker (or line) specific fixed effect allowing an idiosyncratic daily output level for each worker (or line). γ_M is a month fixed effect allowing for general shocks to daily productivity affecting all workers each month (M). This captures seasonalities and market effects of all kinds that might influence output during the year. ω_W is a *day of week* fixed effect that captures unobserved shifts in production associated with specific days of the week (for example there may be lower production on weekends). $R_{i,d}$ is rainfall in mm on day d. $WBGT_d$ represents the average daily wet bulb temperature on day d. We interact the effect of daily wet bulb temperature, $WBGT_d$ with a dummy variable D_k for different bins of the temperature distribution.This allows us to separately estimate the marginal effect on output of a change in temperature within different regions of the distribution.

The production data from the Bhilai steel plant is also used in a similar specification with the primary difference being that high frequency output data is available only at the plant level so worker fixed effects α_i are absent.

Output is measured in rolled blooms, and we also run models using the fault rate (the fraction of cobbled blooms) as the outcome variable.

Table 3 summarizes our results for all case study sites (omitting all fixed effects for clarity). Columns 1-2 report results from the steel mill in Bhilai, columns 3-5 report results from the garment manufacturing firm, and columns 6-7 report outcomes observed in weaving workers in Surat. The shaded columns (1-3) represent climate-controlled plants and the other columns plants without climate control. Columns 3 and 4 are similar garment units operated by the same firm, all located in the National Capital Region (NCR), with and without climate control respectively. We report coefficients associated with a one degree change in wet bulb globe temperature on the output variable (or log output variable), conditional on the value of wet bulb globe temperature. We split the response curve into four wet bulb globe temperature bins[17]: $< 21°C$, $< 21°C - 25°C$, $< 25°C - 27°C$ and $\geq 27°C$. Using the same temperature bins for all case study sites allows for a comparison of effect sizes across datasets.

In Table 3, Rows 3-6 provide the incremental change in output for a one degree change in wet bulb globe temperature within a given WBGT bin. Across all models, identification comes from correlations between dynamic variation within a unit's output (worker or line) with dynamic variation in temperature, controlling for rainfall, time invariant worker or line fixed effects and month and year fixed effects.

In addition to the binned piecewise linear models in the table we also fit a set of highly flexible non-linear estimates of wet bulb globe temperatures on output by modeling this relationship using a set of restricted cubic splines with four knots. These splines also provide useful robustness checks against the bin cut-offs of Table 3. Figure 3 shows the predicted impacts of temperature on output measures for all four case study sites.

Table 3 and Figure 3 identify a few clear patterns that seem strongly supportive of the hypothesis that heat stress is a factor influencing manufacturing

[17]Break points of 25 and 27 degrees allow a comparison to the breakpoints used in Hsiang (2010)

output.

1. In the highly mechanized steel mill where some workers are located in air-conditioned cabins there is no clear relationship between temperature and output, faults or delays. We identify fairly tight zeros or near zero effects and there is no drop-off in output at high temperatures (Panel A of Figure 3 and Columns 1,2 in Table 3).

2. Worker output in garment plants in the National Capital Region that had air-washers installed show a very different temperature response relative to lines located on production floors in without air-washers (Panel B of Figure 3 and Columns 3 and 4 in Table 3). Output levels are similar at cool temperatures but lines without access to air-washers shows a clear drop in output with increasing wet bulb globe temperatures. Note that all plants in this comparison were managed by the same firm, are located in the same region and produce similar products with efficiency measured identically in both cases.

3. Garment manufacture lines on production floors located in Hyderabad and Chhindwara - where air-washers were not installed - also show a drop in efficiency with increasing wet bulb temperatures (Panel C of Figure 3 and Column 5 of Table 3). Note that temperatures are moderate most of the year in these areas that in the NCR.

4. In small weaving units in Surat, another setting without climate control, a similar non-linear pattern of temperature impact on worker output is observed with no temperature effect at lower temperatures and negative estimates in the highest degree day bins (Panel D of Figure 3 and Columns 6 and 7 of Table 3).

The data in Panel B and C of Figure 3 comes from garment units with power backups. In the case of Panel B, we compare co-located plants for whom the frequency of any power outages would be similar. This provides fairly strong evidence that the differences in temperature sensitivity we identify here are not being caused by power supply variations. For weaving units (Panel D) backup power was not present. Nevertheless, Gujarat is a power surplus state with claimed 24x7 supply in urban areas. Weaving plant managers in interviews conducted while collecting data also indicated that unscheduled power outages were not a serious concern.

4.1 Worker Absenteeism

Temperature exposures may plausibly impact worker attendance in various ways and with some lag time. A very hot day might immediately reduce desire to work. Sustained high temperatures may eventually lead to fatigue or illness. Longer term seasonal variations could create differences in disease burden and modify preferences over working environments. 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 - while also linked to heat stress - could affect output (or labor input costs) independently of actual workplace performance.

Although the ASI survey data provides no good measure of worker attendance, for our case study sites we were able to collect detailed histories worker attendance that can be used to investigate this question. For the steel plant in Bhilai we observe a daily count of total unplanned worker absences over an approximately three year period between Feb 2000 and March 2003. For the cluster of garment workers in the NCR we observe worker attendance histories in 2012 and 2013. For weaving workers in Surat, we observe worker attendance histories through the financial year 2012-2013. These records can be used to create a time-series measure of total daily worker attendance (or total absences) in each case.

Importantly, the different worker populations we observe vary in the nature of their wage contracts. Workers in the steel mill and garment manufacture units are full-time contracted employees paid a monthly wage. However weaving workers are not, and earn day wages based on output produced. This means the cost of each day of absence for weaving workers is relatively high,

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

	Dependent variable:						
	Steel Mill		Garment Manufacture Plants			Weaving Plants	
	log(blooms)	faults	log(efficiency)			log(meters)	meters
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(1) rainfall	0.001***	−0.00003**	0.0002	0.002	−0.003	0.006	1.512
	(0.0004)	(0.00002)	(0.0004)	(0.001)	(0.002)	(0.008)	(0.958)
(2) log(budgetedeff)			0.795***	0.495***	0.539***		
			(0.034)	(0.157)	(0.041)		
(3) WBGT:[<20]	−0.001	−0.0001	0.014***	−0.015**	−0.15	0.001	0.462
	(0.006)	(0.0002)	(0.004)	(0.007)	(0.009)	(0.009)	(0.998)
(4) WBGT:[20-25)	−0.004	0.00004	−0.015**	−0.095***	−0.004	0.006	1.627*
	(0.006)	(0.0003)	(0.007)	(0.020)	(0.009)	(0.009)	(0.835)
(5) WBGT:[25-27)	0.013	−0.0003	0.024*	−0.024	0.004	−0.014	−0.492
	(0.014)	(0.001)	(0.014)	(0.025)	(0.020)	(0.014)	(1.125)
(6) WBGT:[≥27]	0.009	−0.001*	−0.001	−0.092***	−0.03**	−0.085**	−7.131**
	(0.020)	(0.001)	(0.007)	(0.026)	(0.016)	(0.038)	(2.923)
Number of Plants	1	1	5	1	2	3	3
Number of Units	1	1	74	10	19	147	147
Climate Control	*Y*	*Y*	*Y*	*N*	*N*	*N*	*N*
Worker or Line FE	*N*	*N*	*Y*	*Y*	*Y*	*Y*	*Y*
Month FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Year FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Weekday FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*

Note:

1. Shaded columns represent observations from plants with climate control (use of AC or airwashers)
2. Output measure for steel mill (1-2): Log of rolled blooms and share of cobbled blooms (faults)
3. Output measure for garment units (3-5): Line efficiency (controlling for budgeted efficiency)
4. Output measure for weaving units (6-7) Cloth woven (meters and log meters)
5. Number of units refers to the number of distinct workers or groups of workers identified in data
6. Robust standard errors correcting for serial correlation and heteroskedasticity
7. *p<0.1; **p<0.05; ***p<0.01

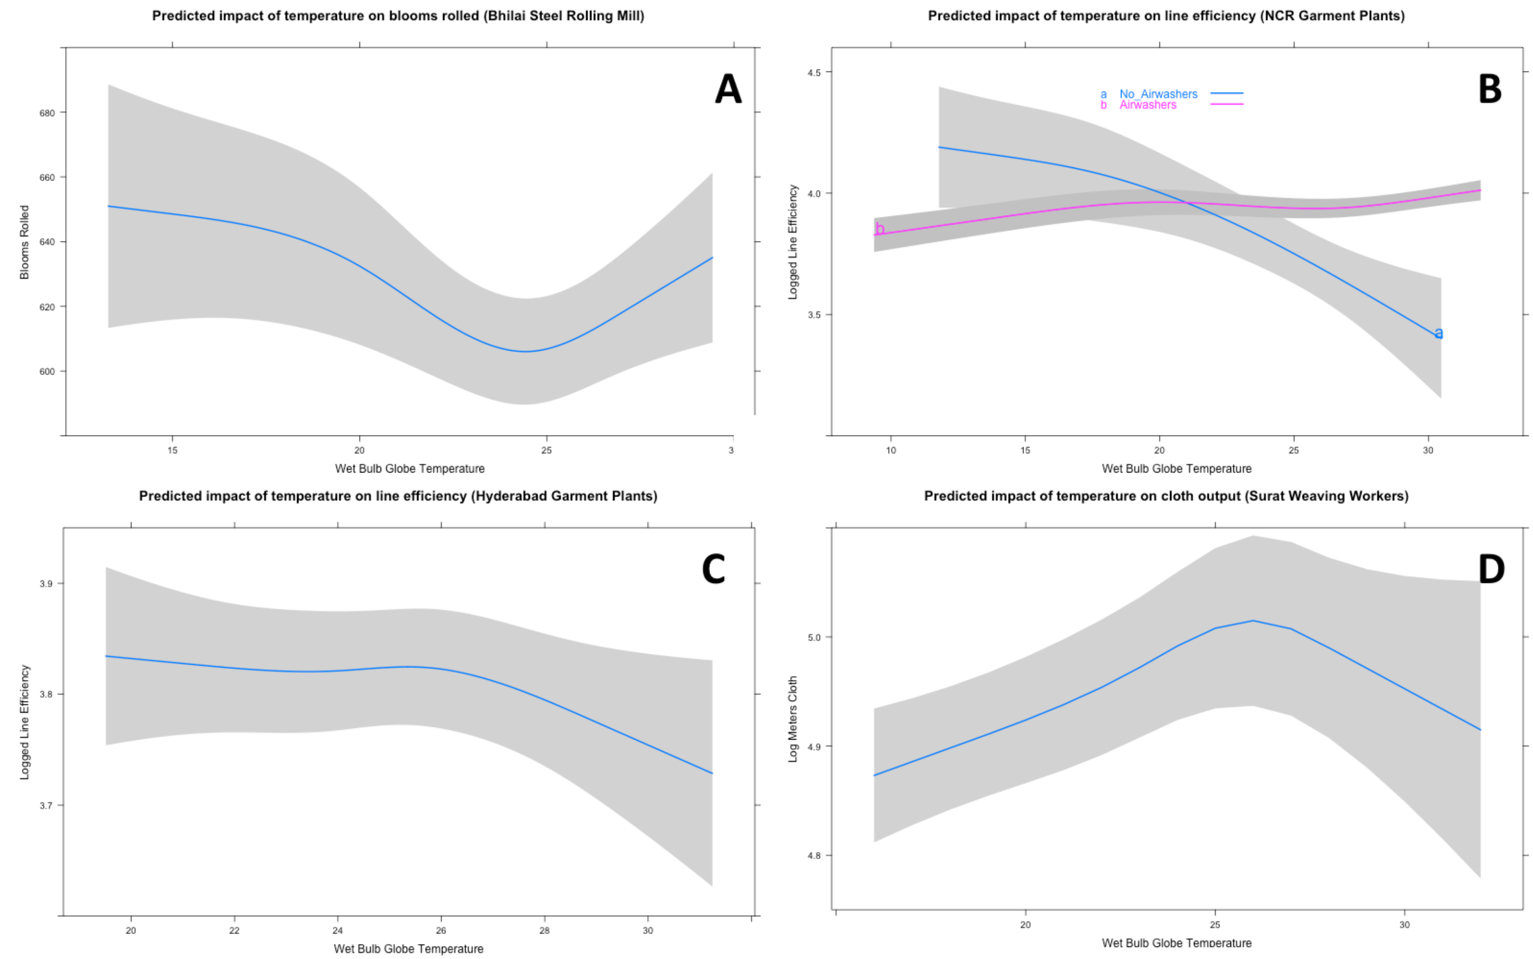

Figure 3: Restricted cubic spline models of the impact of temperature on output measures (with 90 percent bootstrapped confidence intervals). Panel A: Rolled blooms against temperature (Bhilai Steel Mill). Panel B: Logged efficiency in garment plants in NCR both with airwashers (5 plants) and without 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

while it may be small or zero in the other two cases.

To flexibly model the impact of temperature on attendance we proceed as follows. Let the probability of worker attendance p on day t_o be $p(t_o)$. Then we could write

$$p(t_o|X) = \sum_{k=0...K} f_X(W_{t_o-k}, k)$$

Here f is some non-linear function of the lag period, k, as well as wet bulb globe temperatures W_{t_o-k} experienced at time $t_o - k$, for all $k \in [0, K]$. The relationship between lagged temperatures and attendance (f_X) might also change depending on characteristics of the underlying economic environment (captured by X). In particular, the costs of absenteeism might vary depending on the nature of labor contracts, a setting of particular economic interest that also varies across our datasets.

This type of exposure-response relationship can be flexibly modeled using non-linear distributed lag models. These models are based on estimating the parameters of a flexible functional form specification describing the relationship between the response variable (worker attendance) and exposure to temperature W_{t_o-k} experienced at lag period k[18].

In the present context we estimate a model using two third order polynomials to describe the marginal effect of both temperature levels and the lag period of exposure on worker attendance measures (controlling for month fixed effects and rainfall)[19]. With these parameters estimated we can then simulate how worker attendance responds to any given exposure history of temperatures.

Figure 4 shows how attendance (or absence) measures from the three case study sites are predicted to change for a one degree increase in WBGT (on a base of 25 degrees celsius) sustained over varying periods of time[20].

For both steel workers in Bhilai and garment workers in the NCR we find

[18]See (Gasparrini et al., 2010; Gasparrini, 2013) for more details on empirical estimation

[19]For garment workers we restrict the cohort of interest to a subset of about 2700 workers employed for at least 80 percent of the two year period. This ensures we capture only workers with long term contracts.

[20]See Appendix Table 8 for linear models of the relationship between contemporaneous and week-averaged temperature shocks on worker attendance probabilities.

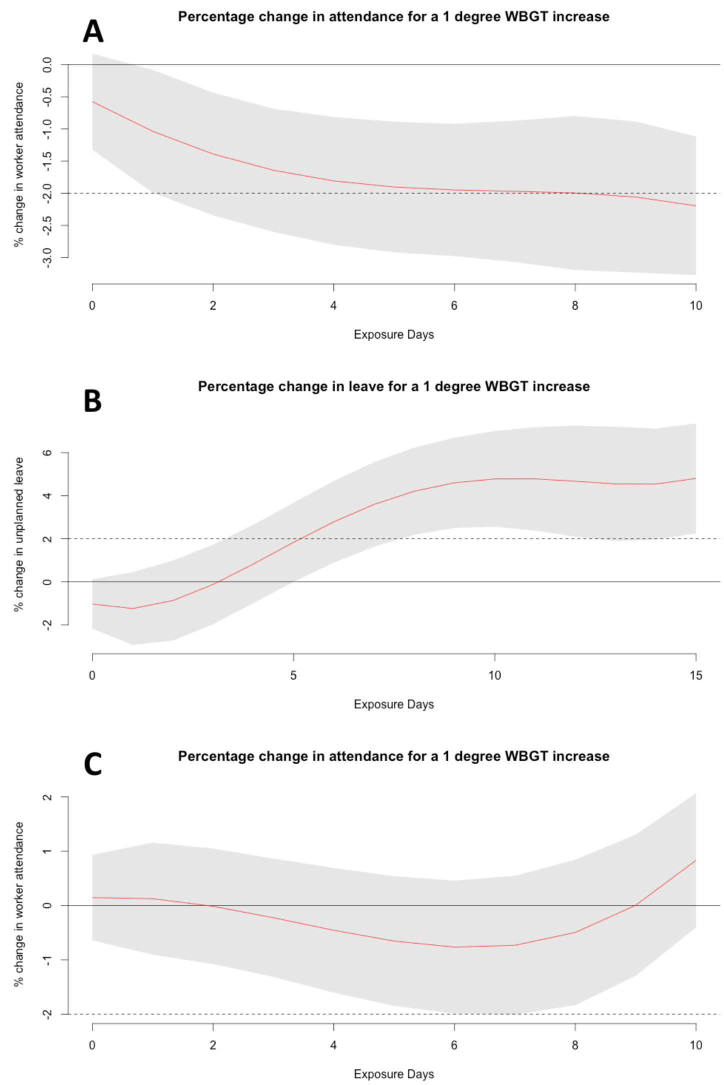

Figure 4: Predicted impact of wet bulb globe temperature on attendance measures for contracted garment workers (Panel A), contracted steel workers (Panel B) and daily wage workers in weaving firms (Panel C). Notes: (i) Effects reported for a base of 25 degrees celsius. (ii) Third order polynomial fits for both temperature and lag dependence. (iii) 90 percent confidence intervals

no significant effect of contemporaneous temperature shocks on same day attendance. However as cumulative exposure increases, we see absenteeism rates first increase and then stabilize at a level about 2 to 4 percent per degree celsius higher than baseline. It is interesting that these effect sizes are similar to those reported by (Zivin and Neidell, 2014).

We also note the absence of any short-run impact of temperature on attendance for weaving workers in Surat. One important difference between weaving workers and our other settings is that the costs of absenteeism are high for weaving workers because wages are paid daily. Recognizing that there may also be other unobserved differences at play here, this difference in attendance response at least suggests that economic incentives might be able to mitigate absenteeism.

The analysis here is restricted to short-run responses of attendance to temperature shocks (our models include month fixed effects). It is possible that temperature might be causally associated with long-term seasonal effects on attendance or changes in employment patterns. Such longer run responses might actually be more acute in the case of non-contracted workers who do have greater employment mobility.

In interviews with weaving firm managers in Surat a frequent complaint related to the difficulty of getting daily wage workers in summer months even though these periods correspond with the agricultural off-season. Figure 5 suggests there may indeed be seasonal reductions in the availability of non-contracted workers (Panel A) that are concentrated in high temperature months. These seasonal patterns are absent for garment workers with long term contracts (Panel B).

While our data does not allow us to draw any strong conclusions about longer-run patterns, this is an area that might benefit from further research.

5 Climate Model Projections

We begin by noting two stylized facts that underline why temperature effects on productivity might be a significant concern both in India and elsewhere. Panel A of Figure 6 reproduces a map of annual wet bulb temperature maximums from (Sherwood and Huber, 2010). It is striking that India lies in the right tail of the spatial distribution, as does much of the tropical belt.

Panel B provides projections of future changes in temperature from two

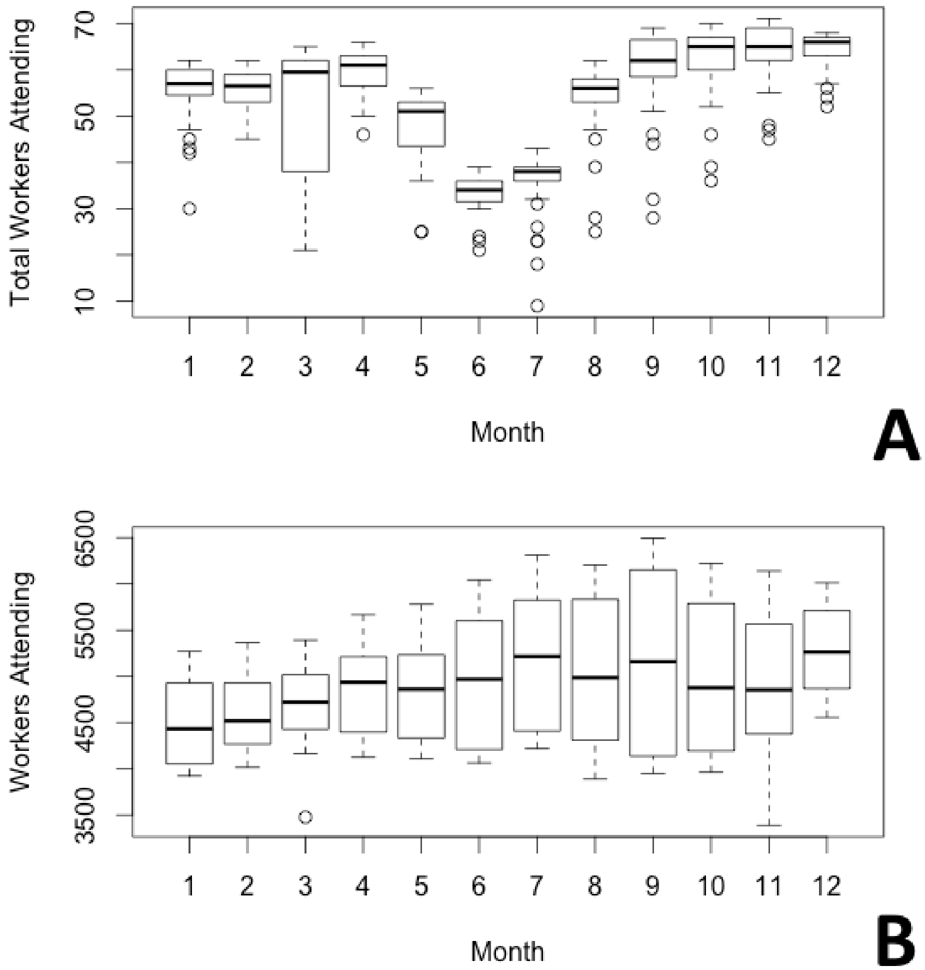

Figure 5: Worker attendance by month for daily wage workers in weaving units (Panel A) vs contract workers in garment manufacture units (Panel B)

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. These projections suggest that India is likely to see significant shifts in annual high temperature degree days and a corresponding reduction in cooler days, for a net increase in annual average temperatures. We overlay on these projections our estimate of the non-linear effects of temperature on manufacturing output from Table 1 (column 3).

Figure 6 shows why significant adaptive measures may be necessary to mitigate heat impacts on workers. These adaptive measures might include air conditioning investments, reallocation of manufacturing to cooler regions, urban planning measures (green cover, water bodies) designed to lower local temperatures, building design changes (cool roofs) and so on. Adaptation could also include measures to reduce the intensity of work, implementing economic incentives to encourage worker effort and low cost cooling investments (installing fans, shading windows etc). Recent work also suggests adaptive possibilities from the use of LED lighting (Adhvaryu et al.).

Ignoring adaptation, these projections, combined with our estimates of the impact of temperature on output (Column 2 of Table 1) allow us to compute an estimate of the upper bound impact of climate change on manufacturing output in India due to heat effects on workers.[21] The predicted changes in daily average degrees in the three bins are (-1.79, -0.64, 3.34) for ($\leq 20°C$, $20°C - 25°C$, $> 25°C$) respectively in the Hadley model projections. For the CCSM model predicted changes in the highest degree day bin are lower but still significant (-1.17, -0.55, 1.32). Even assuming the lower projection is a more reasonable estimate, our empirical estimates suggest that absent adaptation, the estimated upper bound impact on manufacturing could be as high as -6.99 percent (95 percent CI = [-2.77,-10.69]).

Although this exercise cannot be interpreted as a prediction of long run impacts, it does place the potential importance of this mechanism into context. Also while climate projections are necessarily imprecise, heat island effects in urban areas have already led to temperature hotspots that are more than five degrees warmer than surrounding areas (Mohan et al., 2012).

[21] To the extent that some adaptive measures may already be widely adopted, our results could be interpreted as being net of low cost adaptation.

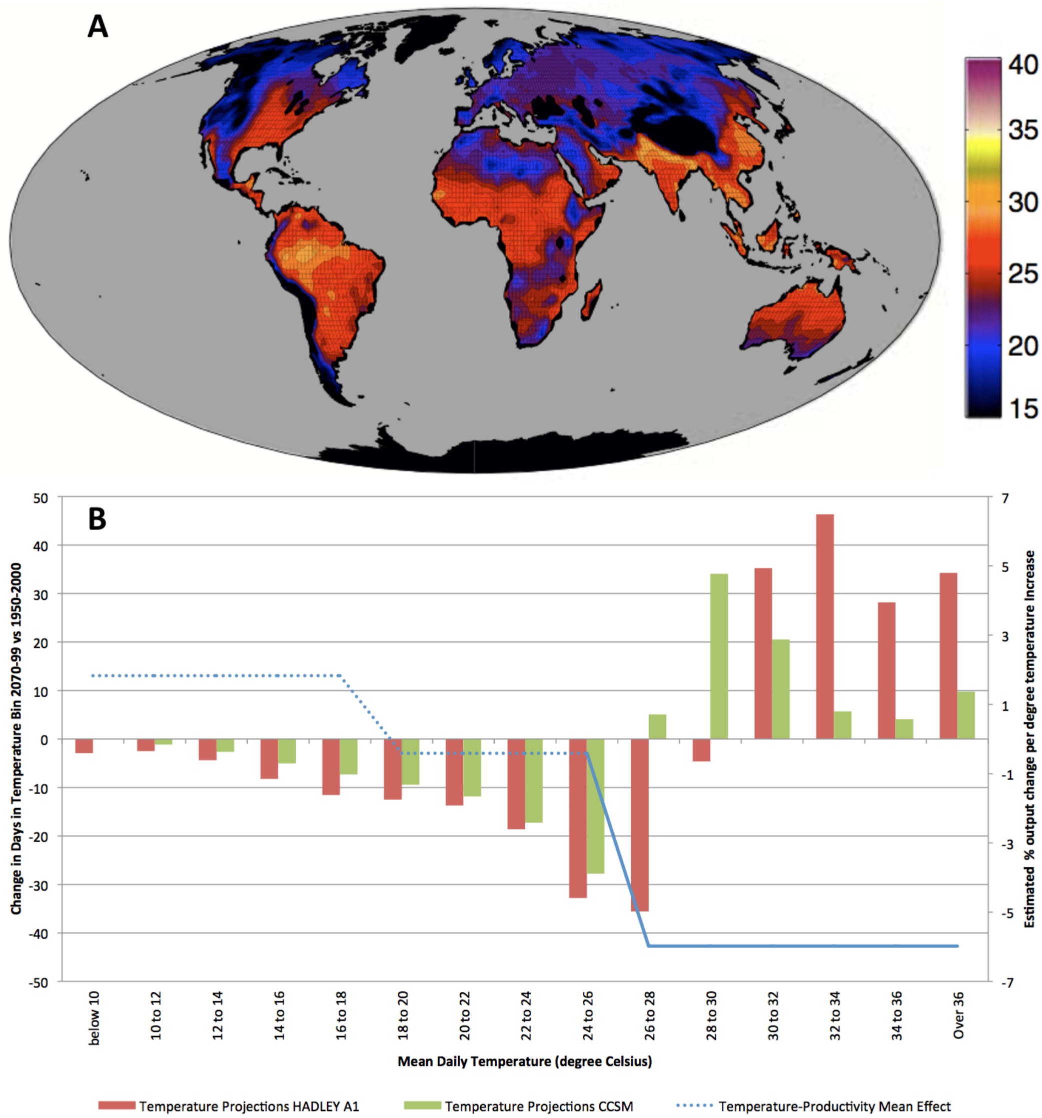

Figure 6: Panel A: Estimated annual wet bulb globe temperature maxima, 1999-2008. Source: Sherwood and Huber (2010). Panel B: Historical and projected temperatures under a business as usual climate change scenario for India. Source: Burgess et al. (2011). Overplotted lines denote estimated productivity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

6 Conclusions

This paper has sought to make three contributions. We provide new evidence to show that high temperatures reduce worker productivity and attendance and thereby may reduce economic output in the absence of climate control. In so doing, we provide a micro-foundation and evidence for a specific mechanism that could help explain previously observed correlations between surface temperatures and the economic output of poor and developing countries (Dell et al., 2012).

While worker heat stress may not be the only factor explaining these macro-level correlations, the effect sizes we identify in different and 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.

Climate change aside, the link between environmental variables (especially temperature) and economic growth has long been studied by economists (Gallup et al., 1999; Nordhaus, 2006). The evidence in this paper therefore relates directly to research on the environmental determinants of long-term growth.

Our second contribution is to link an emerging economics literature on the environmental determinants of productivity to the scientific literature on urban temperature changes. Urban heat island effects have been extensively studied in the scientific literature (Arnfield, 2003) but relatively little attention has been paid to them by economists. The evidence we present 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).

Thirdly, we show that adaptation against temperature productivity impacts is possible through the use of workplace climate control.[22] We also find that attendance reductions are not observed in workers who face high opportunity costs of absenteeism. This suggests that economic incentives could also 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.

Appendix

A Annual Survey of Industry Data Cleaning

We carry out a few data-cleaning operations before analysis in order to transparently eliminate outliers (units with implausibly large output values or zero and negative output) and instances with missing and possibly incorrect data.

Manufacturing Restriction We restrict the sample to surveyed units that report NIC codes belonging to the manufacturing sector.

Trimming 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 a number of outliers - firms with implausibly high reported output (or other variables) as well as a long tail of plants producing near zero reported output.

Missing or Incorrect Data 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

[22] Of course in some small-scale manufacturing settings the high energy costs of such climate control may not be worth the productivity benefits.

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.

Panel Information We drop units that appear twice or fewer times from our panel. All remaining observations form part of our dataset.

B Impact of Temperature on Manufacturing Output

Formally we estimate the following regression equation,

$$V_{i,t} = \alpha_i + \gamma_t + \omega K_{i,t} + \phi W_{i,t} + \beta T_{i,t} + R_{i,t} + \epsilon_{i,t} \tag{10}$$

Here $V_{i,t}$ is the recorded value of output produced by a specific industrial unit i during financial year t. This quantity is the product of physical output with average prices per unit product (aggregated over all outputs). α_i is a fixed effect representing average level of output for each manufacturing unit. γ_t are time fixed effects capturing national changes in manufacturing output year to year. $T_{i,t}$ is our primary variable of interest, namely the average temperature during the financial year t (so that a year is calculated from April 1 through March 31). $R_{i,t}$ is a control for rainfall. K_{it} is a control variable that measures the total working capital available to the plant at the start of the financial year (a measure that includes cash generated from the previous years output less expenditures). $W_{i,t}$ is a control for contract workers employed.

Results for various specifications are presented in Table 4. Results using WBGT values based on long run average measures of daily relative humidity between 1981-2010 from the NCEP/ NCAR reanalysis datasets are presented in Table 5.

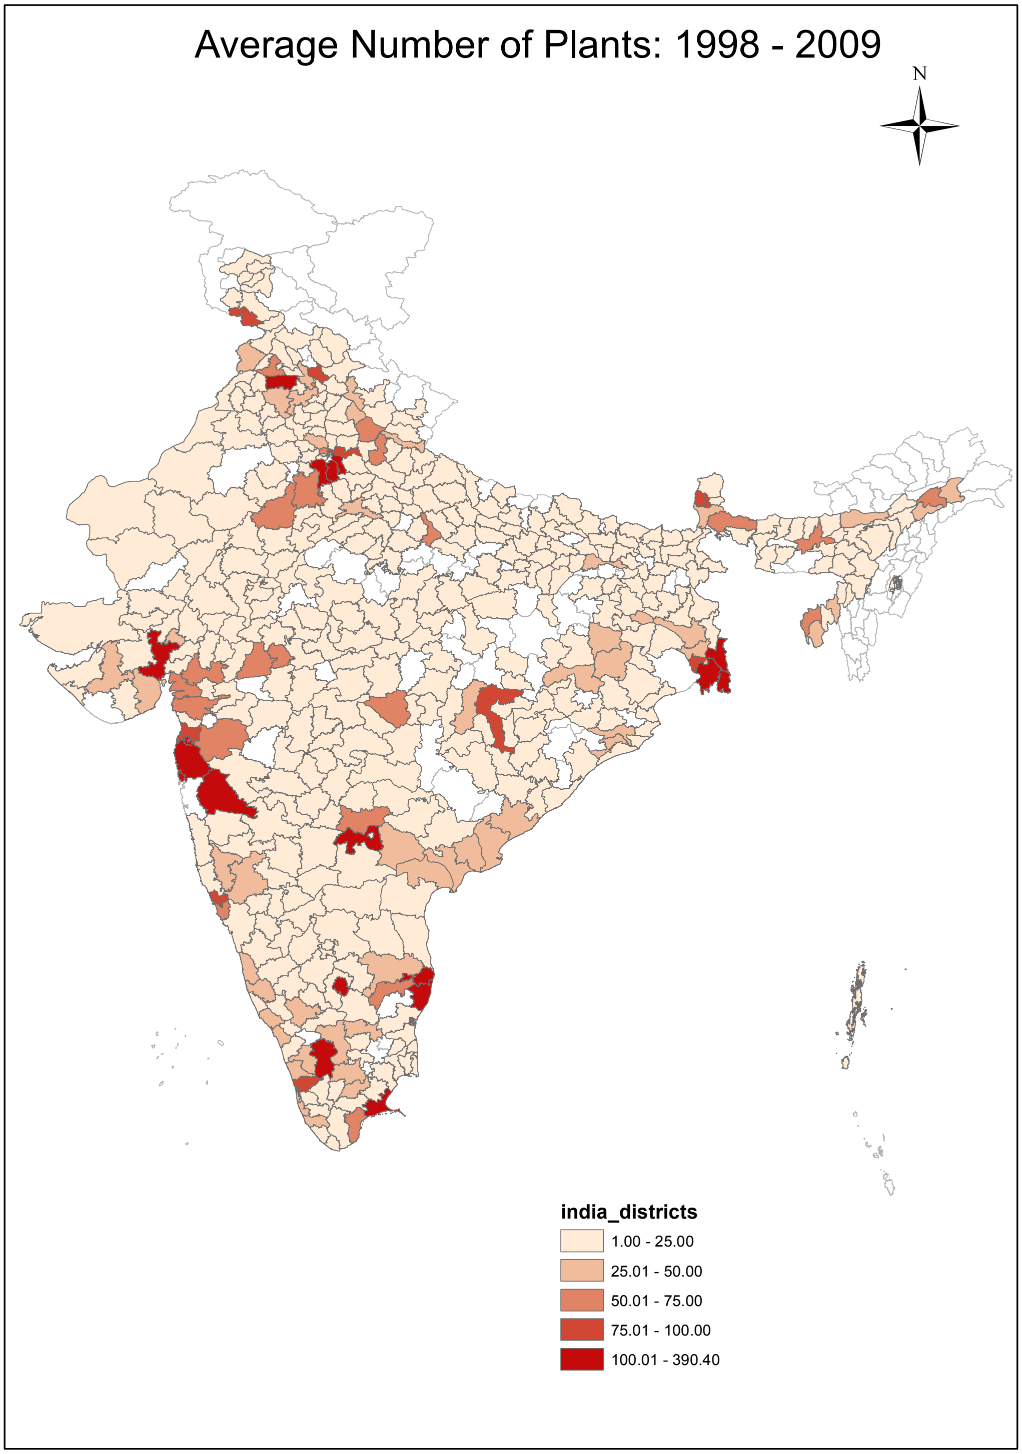

Figure .1: Distribution of annual ASI survey observations over India districts

Table 4: Effect of Temperature on Manufacturing Industry Output

	Dependent variable:					
	Plant Output Value			Log Plant Output Value		Log (Output/Worker)
	(1)	(2)	(3)	(4)	(5)	(6)
mean temp	−.043***	−0.042***	−0.036***	−0.032***	−0.028***	−0.022**
	(0.013)	(0.012)	(0.012)	(0.010)	(0.010)	(0.012)
rainfall	0.013***	0.009***	0.006***	0.003	0.001	0.00
	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)	(0.002)
capital		0.386***	0.346***	0.384***	0.339***	0.197***
		(0.010)	(0.009)	(0.003)	(0.006)	(0.006)
Plant FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Year FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Capital Controls	*N*	*Y*	*Y*	*Y*	*Y*	*Y*
Worker Controls	*N*	*N*	*Y*	*N*	*Y*	*N*
Units	21,525	21,525	21,525	21,525	21,525	21,525
R^2	0.0076	0.4615	0.4876	0.6705	0.6595	0.2930

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.

C Estimates using Wet Bulb Globe Temperature

As described in Section 1 the environmental quantity of most direct relevance to heat stress is not temperature but rather the wet bulb globe temperature, an index that also accounts for ambient humidity. Unfortunately creating a nationwide WBGT measure across the time period of the ASI panel is difficult because reliable time series data on relative humidity across India is not available. For this reason our analysis of national level ASI data uses local dry bulb temperatures[23].

However one way of approximating wet bulb temperatures is to use water vapour pressure or humidity measures produced by re-analysis models such as the NCEP/ NCAR reanalysis datasets. Wet bulb temperatures computed this way are likely to be noisy since reanalysis models are not generally calibrated to provide accurate estimates of temporal variation in humidity. Nevertheless as a robustness check we also repeat our estimate of the average temperature effect on output using a WBGT measure obtained by combining temperature with long run average measures of daily relative humidity between 1981-2010 from the NCEP/ NCAR reanalysis datasets. Appendix Table 5 summarizes our results which look very similar to those in Appendix Table 4.

D Local Input Price Shocks

An input price index is created for all 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. This information is missing in about 28 percent of responses. In addition to dropping plants with missing data 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.

Formally we estimate the model below where $P_{i,t}$ is the log of the plant

[23] In our analysis of high frequency worker output data we are able to link production output to local weather stations and can therefore compute wet bulb globe temperature in these cases.

Table 5: Effect of Wet Bulb Globe Temperature on Manufacturing Industry Output

	Dependent variable:					
	Plant Output Value			Log Plant Output Value		Log (Output/Worker)
	(1)	(2)	(3)	(4)	(5)	(6)
wbgt	−.042***	−0.044***	−0.036**	−0.036***	−0.030**	−0.022*
	(0.015)	(0.014)	(0.014)	(0.013)	(0.012)	(0.013)
rainfall	0.013***	0.009***	0.007***	0.003	0.001	0.00
	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)	(0.002)
capital		0.386***	0.346***	0.390***	0.339***	0.197***
		(0.010)	(0.009)	(0.006)	(0.006)	(0.006)
Plant FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Year FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Capital Controls	*N*	*Y*	*Y*	*Y*	*Y*	*Y*
Worker Controls	*N*	*N*	*Y*	*N*	*Y*	*N*
Units	21,525	21,525	21,525	21,525	21,525	21,525
R^2	0.0076	0.4615	0.4876	0.6705	0.6595	0.2930

Note:
1. *p<0.1; **p<0.05; ***p<0.01
2. Robust standard errors correcting for serial correlation and heteroskedasticity
3. Maximum temperature is on average $6°C$ above the mean temperature
4. Coefficients for models 1-3 are expressed as percentages of average output level.

input price index and other variables are as reported in Appendix Section B. Results are reported in Table 6.

$$P_{i,t} = \alpha_i + \gamma_t + \omega K_{i,t} + \sum_{k=1}^{N} \beta_k D_k + +\phi W_{i,t} + R_{i,t} + \epsilon_{i,t} \qquad (11)$$

Table 6: Impact of temperature on input price index

	Dependent variable: Input Price Index	
	(1)	(2)
meant	0.063	
	(0.040)	
DD1		0.023
		(0.087)
DD2		0.121
		(0.081)
DD3		0.050
		(0.051)
rainfall	0.002	0.002
	(0.007)	(0.007)
Plant FE	*Y*	*Y*
Year FE	*Y*	*Y*
Capital Controls	*Y*	*Y*
Number of Units	21,525	21,525
Mean Obs. per Unit	4.8	4.8
R^2	0.480	0.685
Note:	1. *p<0.1; **p<0.05; ***p<0.01	

E Intra-day Temperature Shocks and Output Variation

In Table 3 in the main paper we report models relating daily output to daily temperatures. However for some of our case study sites, namely garment manufacture units located in the NCR as well as Hyderabad, we observe not just daily but *hourly* data on line output efficiencies (where a line is a small group of workers). This hourly output data can be matched to hourly temperature and humidity data from nearby weather stations.

This allows us to estimate a model of the following type:

$$log(E_{i,d}) = \alpha_i + \gamma_d + \psi_h + log(B_{i,d}) + \beta_k WBGT_{h,d} \times \mathrm{D_k} + \epsilon_{i,d} \tag{12}$$

Here $E_{i,h}$ is achieved line efficiency for line i during hour h. Similarly $B_{i,h}$ is the budgeted (target) efficiency for line i during hour h. α_i is a line specific fixed effect allowing an idiosyncratic hourly output level for each line. γ_d is a day fixed effect allowing for a idiosyncratic shock to line efficiency level for every single day. Thus introducing γ_d into the model soaks up variation in daily average temperatures. ψ_h is a fixed effect for every hour of day which flexibly controls for any consistent trends in output within the day.

$WBGT_{h,d}$ represents the wet bulb temperature on hour h of day d. We interact the effect of hourly wet bulb temperature, $WBGT_{h,d}$ with a dummy variable D_k for different bins of the temperature distribution. To make estimation more tractable we divide wet bulb globe temperatures into only two bins, below and above 25 degrees C

We estimate models with log efficiency as the dependent variable as well as absolute efficiency. We also estimate a model using the log of the ratio of actual to budgeted efficiencies as the dependent variable (output measure). All results are reported in Table 7.

F Linear Regression Models of Absenteeism

It is possible to carry out a simple examination of whether worker absenteeism is associated with high temperature days by estimating a simple linear probability model as follows

$$IsPresent = \alpha_i + \gamma_M + \omega_W + \beta_k WBGT_d \times \mathrm{D_k} + \theta R_{i,d} + \epsilon_{i,d} \tag{13}$$

where $IsPresent$ is a binary variable that takes the value 0 when a worker does not report for work and 1 otherwise. The other covariates on the right have the same definitions as in Equation 9.

A simple test of whether sustained temperatures matter can be obtained

Table 7: Effect of Hourly Wet Bulb Globe Temperature on Worker Output

	Dependent variable:		
	log(actual efficiency)	log(efficiency ratio)	actual efficiency
	(1)	(2)	(3)
(1) log(budgeted efficiency)	0.463***		
	(0.062)		
(2) budgeted efficiency			0.388***
			(0.064)
(3) wbgt:[0,25]	−0.029***	−0.028***	−1.492***
	(0.006)	(0.006)	(0.276)
(4) wbgt:[25,35]	−0.022**	−0.019*	−1.245***
	(0.010)	(0.010)	(0.465)
Observations	39,747	39,747	39,747
R^2	0.209	0.197	0.238

1. *p<0.1; **p<0.05; ***p<0.01
2. Robust standard errors correcting for serial correlation and heteroskedasticity
3. Fixed effects and bin dummies omitted

by estimating an alternative specification using the average wet bulb globe temperature over the preceding seven days as an independent variable. This specification tests for absenteeism that responds not just to a single hot day but to more sustained increases.

We report our results in Table 8. Once again we find little evidence that temperature shocks on a given day increase worker absenteeism[24]. This may reflect the fact that in developing countries the opportunity costs of absenteeism for workers may be fairly high and protection from heat through staying at home relatively low.

In examining models with weekly average temperatures on the right hand side, we do see evidence that absenteeism increases in steel mill and garment workers but not in weaving workers. Note that the models for garment and weaving workers in Table 8 are estimated within worker - with worker fixed effects - and therefore effect sizes cannot be directly compared with changes in daily attendance or absenteeism (which can be approximated by a binomial

[24]There is some evidence that high rainfall days may increase absenteeism

random variable with a number of trials equal to the size of the observed cohort). These results also include observations from all garment workers and not just those who have completed two years of work.

Table 8: Effect of Temperature on Worker Absenteeism

	Dependent variable:					
	log(extraleave)		Absent		*Absent*	*Absent*
	Steel Mill		Weaving		Garment Manufacture	
	(1)	(2)	(3)	(4)	(5)	(6)
WBGT:Q1	0.041		−0.001		0.002∗∗∗	
	(0.024)		(0.003)		(0.0004)	
WBGT:Q2	0.051		−0.001		−0.003∗∗∗	
	(0.039)		(0.003)		(0.0005)	
WBGT:Q3	0.023		−0.001		0.001	
	(0.042)		(0.005)		(0.001)	
WBGT:Q4	−0.036		−0.008		0.004∗∗∗	
	(0.034)		(0.010)		(0.0005)	
rainfall	0.001	0.001	−0.002	−0.003	0.054∗∗∗	
	(0.002)	(0.002)	(0.003)	(0.003)	(0.002)	
Mean Week WBGT		0.058***		0.003		0.0007∗∗∗
		(0.014)		(0.003)		(0.0002)
Mean Week Rain		0.006		0.006		0.049∗∗∗
		(0.005)		(0.009)		(0.002)
No. of Workers	300	300	147	147	17421	17421
Worker FE	*N*	*N*	*Y*	*Y*	*Y*	*Y*
Month-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. extraleave counts total unplanned absences each day for Bhilai steel mill
4. Absent is a binary variable with value 1 when worker is absent

References

Achyuta Adhvaryu, Namrata Kala, and Anant Nyshadham. The light and the heat: Productivity co-benefits of energy-saving technology.

Hunt Allcott, Allan Collard-Wexler, and Stephen O'Connell. How do electricity shortages affect productivity? evidence from india. 2014.

A.J. Arnfield. Two decades of urban climate research: a review of turbulence, exchanges of energy and water, and the urban heat island. *International Journal of Climatology*, 23(1):1–26, 2003.

M Auffhammer, V Ramanathan, and J R Vincent. 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, December 2006. doi: 10.1073/pnas.0609584104. URL http://dx.doi.org/10.1073/pnas.0609584104.

M. Auffhammer, S. M. Hsiang, W. Schlenker, and A. Sobel. Using weather data and climate model output in economic analyses of climate change. *Review of Environmental Economics and Policy*, 7(2):181–198, Jul 2013. doi: 10.1093/reep/ret016. URL http://dx.doi.org/10.1093/reep/ret016.

Robin Burgess, Olivier Deschenes, Dave Donaldson, and Michael Greenstone. Weather and death in india. 2011.

Sanghamitra Das, Kala Krishna, Sergey Lychagin, and Rohini Somanathan. Back on the rails: Competition and productivity in state-owned industry. *American Economic Journal: Applied Economics*, 5(1):136–162, Jan 2013. doi: 10.1257/app.5.1.136. URL http://dx.doi.org/10.1257/app.5.1.136.

Melissa Dell, Benjamin F Jones, and Benjamin A Olken. Temperature Shocks and Economic Growth: Evidence from the Last Half Century. *American Economic Journal: Macroeconomics*, 4(3):66–95, July 2012. doi: 10.1257/mac.4.3.66. URL http://dx.doi.org/10.1257/mac.4.3.66.

CB Field, V Barros, K Mach, and M Mastrandrea. Climate change 2014: Impacts, adaptation, and vulnerability. contribution of working group ii to the fifth assessment report of the intergovernmental panel on climate change, 2014.

J. L. Gallup, J. D. Sachs, and A. D. Mellinger. Geography and economic development. *International Regional Science Review*, 22 (2):179–232, Aug 1999. doi: 10.1177/016001799761012334. URL `http://dx.doi.org/10.1177/016001799761012334`.

A. Gasparrini, B. Armstrong, and M. G. Kenward. Distributed lag non-linear models. *Statistics in Medicine*, 29(21):2224–2234, Sep 2010. doi: 10.1002/sim.3940. URL `http://dx.doi.org/10.1002/sim.3940`.

Antonio Gasparrini. Modeling exposure-lag-response associations with distributed lag non-linear models. *Statist. Med.*, pages n/a–n/a, sep 2013. doi: 10.1002/sim.5963. URL `http://dx.doi.org/10.1002/sim.5963`.

Michael Greenstone and B. Kelsey Jack. Envirodevonomics: A research agenda for a young field. Technical Report NBER Working Paper: 19426, National Bureau of Economic Research, September 2013.

Solomon M Hsiang. 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, August 2010. ISSN 1091-6490. doi: 10.1073/pnas.1009510107.

ISO. Hot environments – estimation of the heat stress on working man, based on the wbgt-index (wet bulb globe temperature). Technical report, International Standards Organization, 1989.

Tord Kjellstrom, Ingvar Holmer, and Bruno Lemke. Work- place heat stress, health and productiv- ity – an increasing challenge for low and middle-income countries during climate change. *Global Health Action*, 2: Special Volume:46–51, 2009.

Bruno Lemke, Bruno and Tord Kjellstrom, Tord. Calculating workplace wbgt from meteorological data: A tool for climate change assessment. *Industrial Health*, 50(4):267–278, 2012. doi: 10.2486/indhealth.MS1352. URL `http://dx.doi.org/10.2486/indhealth.MS1352`.

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

R Mendelsohn and A Dinar. Climate Change, Agriculture, and Developing Countries: Does Adaptation Matter? *The World Bank Research Observer*, 14(2):277–293, August 1999. doi: 10.1093/wbro/14.2.277. URL `http://dx.doi.org/10.1093/wbro/14.2.277`.

Manju Mohan, Yukihiro Kikegawa, B. R. Gurjar, Shweta Bhati, and Narendra Reddy Kolli. 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, Sep 2012. doi: 10.1007/s00704-012-0758-z. URL `http://dx.doi.org/10.1007/s00704-012-0758-z`.

W. D. Nordhaus. Inaugural article: Geography and macroeconomics: New data and new findings. *Proceedings of the National Academy of Sciences*, 103(10):3510–3517, Mar 2006. doi: 10.1073/pnas.0509842103. URL `http://dx.doi.org/10.1073/pnas.0509842103`.

K C Parsons. *Human Thermal Environments*. Informa UK (Routledge), 1993. ISBN 978-0-203-35595-4. doi: 10.4324/9780203302620. URL `http://dx.doi.org/10.4324/9780203302620`.

W. Schlenker and Michael Roberts. 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, 2009.

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

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

Joshua Graff Zivin and Matthew Neidell. The impact of pollution on worker productivity. *American Economic Review*, 102 (7):3652–3673, Dec 2012. doi: 10.1257/aer.102.7.3652. URL `http://dx.doi.org/10.1257/aer.102.7.3652`.

Joshua Graff Zivin and Matthew Neidell. Temperature and the allocation of time: Implications for climate change. *Journal of Labor Economics*, 32(1):1–26, Jan 2014. doi: 10.1086/671766. URL `http://dx.doi.org/10.1086/671766`.