The Economic Impacts of Temperature on Industrial Productivity: Evidence from Indian Manufacturing

By Anant Sudarshan and Meenu Tewari*

Draft: October 24, 2013

An issue of critical importance in climate and development policy involves understanding the economic impact of changes in surface temperatures. Most of what we know about the impacts of climate change is restricted to evidence on agriculture, human health and natural disasters. This paper suggests that surface temperatures might influence economic output more broadly through their impact on worker productivity. We provide empirical evidence in support of this hypothesis in the context of manufacturing sector output in India. We exploit a national level dataset of individual manufacturing plants in India to show that (i) plant level manufacturing output responds negatively to high temperatures; (ii) the magnitude of losses is economically significant (of the order of 2 percent of daily output per degree celsius); (iii) output response is non-linear in temperature and primarily associated with high temperature days and (iv) temperature response is greatest in plants where the worker value-added is high. This response pattern agrees with the physiological response of human beings to temperature documented from heat stress studies in the lab. We strengthen our conclusions by also directly collecting daily worker productivity measures from selected case study sites. Our results suggest that integrated assessment models may underestimate the global economic costs of climate change by neglecting to account for reduced worker productivity during high temperatures.

JEL: Q54, Q56

* 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. Acknowledgements: We thank ICRIER and the Rockefeller Foundation for financial support. We thank Michael Greenstone for helpful comments. We thank Mehul Patel for important field assistance and Kamlesh Yagnik, President South Gujarat Chamber of Commerce and Industry for advice and encouragement to this project. All errors are the responsibility of the authors.

Governments across the world are attempting to determine an appropriate set of policy actions to address climate change concerns. In order to do so, it is critical to understand both the costs of a climate related policy intervention and its associated benefits. Unfortunately, while the costs of taking action are immediately salient in policy debate (albeit not necessarily well quantified), incomplete evidence exists on benefits because of the difficulty in both quantifying avoided damages and uncovering the mechanisms involved. For this reason, quantifying the link between environmental factors and economic performance is a central part of the research agenda within modern environmental economics (Greenstone and Jack, 2013). In particular, climate change policy cannot be sensibly framed without estimates of both the magnitude of the economic costs of environmental change and the mechanisms through which these costs are created. Adaptation strategies for example, cannot be formulated without knowing not just whether climate shocks matter but also why they matter.

Climate policy thus provides an immediate motivation for understanding the link between environmental parameters such as temperature and economic growth. Nevertheless this question is independently related to a larger research agenda in the development economics literature which has sought to uncover the reasons for differences in the economic performance of different countries. The growth and development literature has often examined the relationship between average temperature and aggregate economic variables across countries (Gallup et al., 1999, Nordhaus, 2006) and noted that poorer countries tend to also be warmer. A question of fundamental importance is therefore *why* this association exists?

One view is that this correlation does not hint at causality and is largely driven by spurious associations of temperature with other national characteristics such as institutional quality where the latter is causally linked to economic performance (Acemoglu et al., 2002). Yet while institutions certainly impact economic output, if there also exists evidence on direct causal channels through which temperature may have wide ranging impacts on productivity then appealing to non-environmental factors alone may be insufficient to explain the totality of cross country variation in economic performance.

It is in this context that this paper seeks to provide empirical evidence suggesting that differences in surface temperature (specifically an increase in high temperature days) might impose sustained and significant economic costs through an additional independent channel. This mechanism involves the direct impact of temperature on industrial output through its effect on labour productivity.

In doing so we complement two recent studies that also provide some evidence that is at least suggestive of the importance of this channel. First, Dell et al. (2012) examine a 50 year panel dataset of country GDP and find that at a national level, total economic output in poor countries responds very strongly to temperature (about a 1 percent decline in annual growth rates per degree temperature rise). In addition the authors provide evidence that output in sectors other than agriculture also responds to temperature. The size of the temperature response is at least suggestive that economic sectors other than agriculture might respond to temperature shocks although the authors cannot rule out the possibility of aggregate economic growth rates and output responding through other well known channels (including spillovers from agriculture shocks to other sectors).

Second, Hsiang (2010) examines economic output across different sectors for a set of countries in the Caribbean and Central America and estimates that a 1 degree C increase in temperature is associated with a 2.4 percent decrease in GDP and that this decrease is primarily associated with warm season temperature shocks. While some concerns remain in interpreting the mechanisms behind these results[1], nevertheless, this work makes a strong case for taking more seriously the direct impact of climate on non-agricultural sectors.

The basis for hypothesizing that temperature might negatively impact the economic performance of non-agriculture sectors, does not have to be made only with recourse to studies that examine how national level economic output covaries with temperature. Laboratory and experimental evidence has documented how human performance begins to deteriorate at higher temperatures and temperature effects have been seen in commercial office settings as well (see for example Tanabe et al. (2007)). Useful summaries of the literature on heat stress as relevant to climate change are available in Kjellstrom et al. (2013) and Hsiang (2010).

For the most part the literature on heat stress has focused primarily on documenting the physiology of human response to heat from laboratory or experimental studies and studying specific real world settings where temperature and humidity exposures are high enough to create occupational health concerns (Wyndham, 1969). Thus it is not obvious whether this mechanism is economically important in a wider range of real world settings where temperatures are more moderate, work requirements need not be physically strenuous and the workplace is indoors. Nevertheless, this evidence provides an independent basis for hypothesizing that temperature impacts on human physiology could be one channel through which climate change might impose economic costs on national output.

Furthermore, to the extent these impacts might be important, they may be both understudied and imperfectly accounted for in climate policy today. For the most part, environmental economists

[1]The setting for this study is a region of the world heavily dependent on tourism and therefore it is plausible that demand shifts coincident with temperature shocks might explain the economic effects found here. This is an alternative explanation that may be especially hard to rule out given that the paper reports statistically significant effects only in the service sector (which is directly affected by tourism) but not in manufacturing or other labour intensive sectors.

and scientists have largely focused on three channels[2] through which even moderate climate change might affect human welfare. The first channel is through the impact of climate change on agriculture (Mendelsohn and Dinar, 1999, Auffhammer et al., 2006). The second is through the direct effect of climate on human health (Barreca et al., 2013). The third channel is through an increased potential for extreme climate events (droughts, hurricanes, heat waves) resulting in large one-time economic losses (Mendelsohn and Dinar, 1999) and as recent evidence suggests, occasionally longer lasting spillover impacts (Antilla-Hughes and Hsiang, 2013). Indeed the Intergovernmental Panel on Climate Change in its Fourth Assessment Report, Working Group II on Impacts, Adaptation and Vulnerability states 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)*". It could be argued that this confidence in the relative immunity of non-agricultural sectors to gradual environmental change owes more to an absence of evidence on other mechanisms, than to evidence of their absence.

This paper seeks to provide empirical evidence on the economic significance of temperature shocks in non-agricultural sectors, especially manufacturing, and to address some of the gaps in the literature to date. We exploit a rich, nationwide survey of manufacturing units (factories) in India to provide rigorous evidence of the impacts of temperature on economic output in the manufacturing sector. To the best of our knowledge this study is the first to directly examine the impacts of surface temperatures on output using plant level micro-data. We show that (i) these impacts are economically significant (an output decline of about 2 percent per degree Celsius), (ii) have a specific non-linear response to temperature that is consistent with laboratory studies of human response to temperature and (iii) temperature impacts on economic performance seem most acute in sectors where the value added per worker is high. In addition, we collect detailed daily worker level output data from weaving plants in Surat and estimate a set of highly restrictive empirical specifications to show that worker output measured directly does seem to drop at high wet bulb globe temperatures.

Our emphasis on manufacturing is driven in part by an effort to provide some evidence on the mechanisms involved. By focusing on manufacturing plants we isolate both a crucial engine of economic growth in developing countries and also concentrate on a setting plausibly less directly affected by first order demand shocks correlated with temperature (such as travel and tourism). Further, by using micro-data at the level of individual manufacturing units we are able to more

[2]Recently an important literature has also reinforced the possibility that climate might affect conflict events (Hsiang et al., 2013). Nevertheless much remains to be learned about the mechanisms through which this might occur and it is difficult to assess the economic costs of such a link.

precisely control for unobservables that might otherwise affect studies using more aggregate data. We also avoid some of risks inherent in drawing conclusions about climate effects on non-agricultural sectors of the economy on the basis of national GDP aggregates that may be influenced strongly by cross-sector economic links. Lastly we use not just national level plant micro-data but also daily data from a set of case study sites. This allows us to validate the mechanism hypothesized because we are able to directly measure worker output at the daily level and can confirm that output does seem to reduce when temperatures rise.

The remainder of this paper is organized as follows. In Section I we briefly summarize the underlying evidence from heat stress studies on human productivity and provide a simple framework for thinking about how these physiological effects might impact economic output in a manufacturing plant.In Section II we describe data and methods used to analyze daily worker output at the factory level from a set of textile weaving plants in Surat. In Section III we introduce the national level dataset we use to broaden our analysis to a range of manufacturing plants and utilizing annual variation in temperature. In Section IV we present results from this national level analysis. In Section V we provide some anecdotal evidence on adaptive strategies. In Section VI we utilize modeled climate change impacts on temperature distributions for India, in conjunction with our estimates of their impact on manufacturing to provide an approximate quantification of the first order expected economic costs imposed by climate changes on the manufacturing sector. Overall we suggest that it is plausible that climate change costs owing to temperature impacts on non-agricultural sectors may even dwarf the economic costs of climate change impacts on agriculture, especially for regions of the world where agriculture is a small share of the gross domestic product. We conclude in Section VII.

I. Theory and Mechanisms

The physics of how temperature affects human beings is well known. The physical exchange of heat between the human body and the 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. When this balance cannot be maintained during normal activity levels then it becomes necessary for safety reasons 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).

The central mechanism the body uses to dissipate heat is through the evaporation of sweat and the efficiency of such dissipation depends on temperature, humidity and wind speed (the movement of air over the skin). Together these can be encapsulated in various ways for ease of use (Parsons, 1993), but perhaps the most internationally accepted index is the Wet Bulb Globe Temperature (ISO, 1989). Ignoring the local wind-speed (which indoors is largely determined by access to fans) the Wet Bulb Globe Temperature (WBGT) is largely determined by two ambient factors, temperature and humidity. At elevated temperatures or high humidity, heat stress can begin to reduce productivity. In this paper we estimate WBGT from temperature and relative humidity using a formula reported in Lemke and Kjellstrom (2012), who discuss and compare several ways of arriving at this measure using different types of meteorological data. This equation serves well as an estimate of WBGT indoors (outdoor levels may be different on accounting for solar radiation).

$$WBGT_{id} = 0.567T_a + 0.216\rho + 3.38 \tag{1}$$

where $WBGT_{id}$ is measured in $\circ C$, T_a is the air temperature and ρ is the water vapour pressure and can be calculated from the relative humidity using the physical relationship below.

$$\rho = (RH/100) \times 6.105 \exp \frac{17.27Ta}{237.7 + T_a} \tag{2}$$

The literature on heat stress also suggests that the response of human beings to temperature (or more precisely, WBGT) is not uniformly linear. Intuitively we might expect that at very cold temperatures, productivity (or at least comfort) might increase in temperature and at moderate levels, temperature variations might have no impact. At higher levels however, heat stress may become progressively more severe as the temperature (or humidity) rises. While the precise shape of the dose response relationship is not well known or even necessarily deterministic, empirical evidence (as well as theory) is consistent with this pattern. Figure 1 reproduces a graph from Hsiang (2010) based on a meta-analysis of over 150 ergonomic studies. Table 1 provides results at higher temperatures from Kjellstrom et al. (2013).

WBGT (°C)	<26	26	27	28	29	30	31	32	33	34	35	36	37	38
300W, % loss	0	0	0	3	9	17	25	35	45	55	64	74	81	85
400W, % loss	0	0	9	17	25	35	45	55	64	74	81	85	88	90

Table 1—Human productivity loss at high temperature. Source: Kjellstrom et al. (2013)

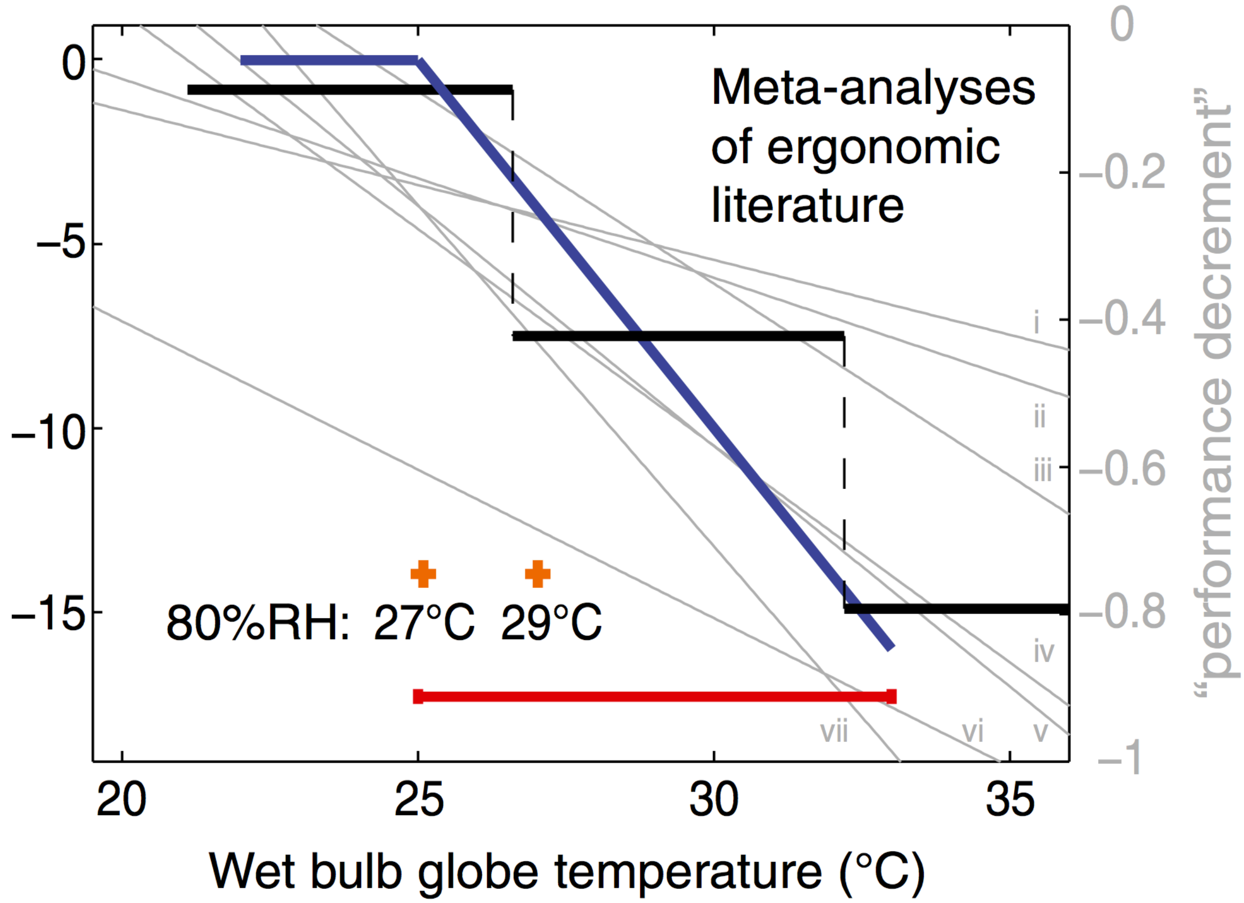

Figure 1. Reduction in productivity measured in ergonomic studies with increasing wet bulb globe temperature. Productivity declines are marked after 25 degrees Celsius. Figure reproduced from Hsiang (2010)

A. *Worker productivity and economic output*

In the preceding section we briefly reviewed the physical mechanisms through which temperature and humidity affect human beings. These mechanisms are well known and well studied but it is nevertheless not obvious how important they might be as a factor influencing productive economic output, especially in sectors such as manufacturing.

There are several reasons for this. First, day to day productive worker activity need not require energy expenditures and physical activity at or near physical limits. This is particularly so in formal, skilled work in the manufacturing and service sectors, as distinct from purely manual and unskilled labour that might play a significant part in the construction or mining sectors.

Secondly, unlike construction and mining, most labour in manufacturing (or services) can be expected to take place indoors or in shielded conditions. These work conditions may provide some protection from ambient temperatures even where air conditioning is rare.

Thirdly, the *economic* impact of reductions in worker productivity may be very different from the physiological impact. In particular, the marginal economic 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 or zero in the case of low value added tasks. As an example, manufacturing units in the diamond sector typically employ some workers to undertake a sorting activity which involves separating raw mined stones into quality grades. Even a small reduction in the ability of the workers to carry out this sorting process without errors may result in extremely large economic losses to the firm, since errors will increase the fraction of high quality stones that are discarded. Conversely, while a worker involved in a low value add activity (such as loading coal into a boiler or overseeing non-critical processes) might see physical productive potential decline in high temperature conditions, this need not necessarily translate to significant economic costs.

This distinction can be easily understood in the context of a simple production function model. Consider a plant with output following a simple Cobb-Douglas production equation as below, where Y is total output, L, E, K represent labour, energy and capital inputs and A is the total factor productivity. L in turn is written as a function of input labour L_o and T, temperature (or wet bulb globe temperature) to denote the dependence of effective labour hours on ambient conditions.

$$Y = AL(T, L_o)^{\alpha} E^{\beta} K^{\gamma} \tag{3}$$

Let

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

Of course, the functional forms here are not meant to be taken seriously and are chosen more to maximize mathematical convenience than anything else. The effective labour available will not depend on T exponentially in practice but the definition here suffices to capture a situation where available labour reduces as temperature increases provided the temperature is ‘high’.

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

$$\frac{dZ}{dT} = \begin{cases} 0 & \text{if } T \text{ is less than } T_c \\ -\alpha\theta & \text{if } T \text{ is less than } T_c \end{cases}$$

In other words, temperature shocks do not affect productivity when temperatures are moderate.

At higher temperatures ($T > T_c$), Z declines with temperature. This decline is higher when α is large, which might represent a firm where the value added by labour is high.

Taken together what we know about the physiological response of human beings to temperature and what we know about the economics of production suggest two additional tests that we can take to the data that might provide some evidence on whether this mechanism is indeed important.

Test 1 If manufacturing output responds negatively to temperature due to its effects on worker productivity then this response should occur primarily at higher temperatures (above approximately $25°C$) and should be zero at moderate (or even positive at low temperatures).

Test 2 Temperature impacts should be greatest in sectors where value added by labour is high

Taking these tests to data is important and useful in this context because we would ideally like to understand not just whether temperature matters to economic output in non-agricultural settings but also why it matters. Absent evidence on mechanisms, it is impossible to determine what adaptation strategies are most likely to be effective or how policy should be framed to deal with climate impacts.

In the present instance one could think of other competing hypothesis on why we might see temperature effect a variety of non-agricultural sectors including spillovers from agricultural productivity shocks (Burgess et al. (2011) suggest that observed health impacts of temperature may partially owe to agricultural productivity shocks). Yet spillovers from agriculture would suggest temperature response patterns that would not necessarily match those described in Tests 1 and 2.

Agricultural growing seasons in India for instance take place during a time of the year where temperatures are relatively moderate and one of the two primary growing seasons is in the winter. Thus it would seem plausible that if non-agricultural sectors respond to temperature shocks primarily through agriculture related economic spillovers then these impacts should be highest when temperature shocks occur at the cooler temperatures found in the growing season. Furthermore by looking at physical output at the level of an individual plant rather than aggregate sector output we should minimize the degree to which we need to be concerned about cross-sector spillovers. Similarly if temperature influences manufacturing output through some channel not involving labour, there should be no particular reason we would expect to see climate sensitivity being higher where value added by labour is high or to follow a non-linear pattern consistent with heat stress.

One final point that is worth making relates to the time-scales on which temperature may affect human beings. The mechanism described above is one that can be expected to create fairly immediate performance changes in human beings contemporaneously with short term changes in temperature. In other words, unlike other environmental stressors (such as certain air pollutants for example), the effects of exposures to high temperatures can be expected to be visible on short time scales. At the same time these impacts are unlikely to disappear when temperature changes are sustained (absent adaptive actions taken to reduce exposures) because the efficiency of heat transfer from the human body depends in straightforward ways on ambient temperatures. Thus it is plausible that sustained temperature differences between populations might lead to sustained differences in the productivity of labor and this mechanism is perhaps worth considering as an environmental variable that might impose long run constraints on growth. Consistent with these ideas, we next investigate temperature effects on productivity at two time-scales using both short term, high frequency variation (Section II) and longer run annual variation (Section IV).

II. Evidence from Weaving Units in Surat

As a starting point we begin by asking whether evidence of reduced worker productivity with rising temperatures can be detected during normal working conditions in a manufacturing environment. Normal working conditions naturally vary by context but at a minimum workers in a formal manufacturing plant typically carry out activities requiring a range of skills, deploying a combination of cognitive and physical faculties, and have breaks woven into their work day. While air conditioning is still the exception rather than the norm in Indian manufacturing [3] nevertheless workers (especially those who work with machines) are typically in indoor environments shielded from direct solar radiation.

In order to understand how temperatures affect labour output in a factory we begin by assembling a unique dataset of *daily* worker productivity and attendance using data hand collected from three small weaving units located in the city of Surat in the state of Gujarat in India. Our analysis in this section bears some similarity to recently published work in Zivin and Neidell (2012) who examines the affect of ozone on worker output under outdoor working conditions in a farm by analyzing daily worker output.

Our choice of the weaving sector is motivated by a number of factors that make it well suited for our purposes. First, the textile sector as a whole (NIC code 1713) and the weaving sector in particular are particularly important manufacturing sectors in developing countries and are highly labour intensive. Although the weaving sector has become increasingly mechanized (excepting the handloom sector), a weaving unit remains a labour intensive enterprise. The workforce in mechanized weaving units is largely engaged in semi-skilled or skilled but relatively low value added

[3] Some manufacturing contexts make even fans difficult to use especially where air flow might interfere with machinery

activity [4]. According to the numbers from the Confederation of Indian Industry 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. This sector accounts for about 21 percent of total employment of which a significant share can be attributed to weaving units (12.5 million people, over a third of total textile sector employment, according to Raaja, 2011).

Second weaving workers in the sector (including those in the enterprises we study) are largely temporary and paid on a piece rate basis linked directly to their output. This makes it possible to collect high frequency output measures at the worker level which are in general not available or reliable in settings where fixed monthly or longer term payments and contracts are used. In the weaving units of Surat, employment is not contracted for even on a monthly basis. Instead workers are paid for the days when they show up to work. Payments are made on the basis of a simple measure of worker level physical output, namely metres of cloth produced, multiplied by a per meter payment (about INR 2.00 per metre in Surat). In the firms we study, this output is also essentially the final output for the plant which is then sold at wholesale markets to dying and printing units. Thus the worker output directly corresponds to plant revenue. Given the non-legally binding nature of Indian minimum wage laws, the low absolute level of minimum wages and the lack of enforcement for small firms, we find that workers are paid for exactly what they produce. We may therefore ignore complications introduced by payment non-linearities at a minimum wage lower bound (as in Zivin and Neidell (2012)).

Third, the work involved is *not* physically demanding and allows for frequent pauses. A weaving worker is primarily responsible for overseeing mechanized looms (each loom can be regarded as a work station). A worker must walk up and down between work stations[5], occasionally adjusting alignment, restarting feeds when interrupted and making occasional corrections as needed. The fact that this work is physically low intensity is important because this is very different from settings where heat stress is known to be a significant health concern (mining for instance, see Wyndham (1969)). Finding temperature impacts on worker output in this setting is therefore more likely to imply that this is may be an important mechanism affecting productivity across a range of industries. Conversely, while intensive manual labour in non-mechanized outdoor settings might be more likely to be temperature sensitive, this type of work is also less likely to be representative of manufacturing or to tell us much about productivity across a range of settings.

[4] Weaving units precede both dying and printing firms, and apparel manufacture shops (including embroidery) in the textile industry supply chain

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

Lastly, because weaving workers are mobile and can choose whether to show up to work on any given day we can examine the impacts of temperature shocks on attendance as distinct from productivity on the job (both of which may affect final output). There exists some limited evidence on this second channel from the United States, based on empirical work using the American Time Use Surveys (Zivin and Neidell, 2010). Of course it could be argued that it is ex-ante unlikely that worker absenteeism is a major concern for these firms for a couple of reasons. First, workers are paid only when they arrive at work and therefore face a clear cost of absenteeism. Since incomes are low, the relative income effect of absenteeism is therefore substantial. Secondly, absenteeism may not hurt factory output if regularly absent workers can be quickly and easily replaced (as seems plausible in this context, where labour is substitutable and there are no long term contracts).

A. Data and Results

In order to examine labour productivity and how it changes with temperature, we collected daily worker output data from three partnering weaving units in the city of Surat over a period spanning a one financial year (365 days) from April 2012 (the start of a new financial year) until the end of March 2013. Unfortunately data preceding April 2012 was generally unavailable (to avoid tax audits, small firms in India often keep only a minimal record of historic financial expenditures). Our dataset comprises 151 workers for whom we observe daily output when they arrive at work. These records are likely to be fairly reliable since they are constructed from the payment slips generated at the end of the day by the firms and because in this simple but mechanized setting, workstation output is easily measured. This dataset is then merged with weather data from a station in Surat (weather station CWOP ID: 42840) providing daily temperature, precipitation and humidity records for the entire period of time for which worker records are available. We use these records in conjunction with Equation 1 to calculate a daily wet bulb globe temperature (WBGT) measure.

In order to estimate the impact of temperature on productivity, we utilize the quasi-random day to day variation in WBGT associated with day to day variation in worker output. More precisely, we estimate coefficients of the linear model below through ordinary least squares.

$$log(\mathrm{Y}_{i,d}) = \alpha_i + \gamma_M + \omega_W + \beta_k WBGT_d \times \mathrm{D_k} + \theta f(R_{i,d}) + \epsilon_{i,d} \tag{4}$$

Here $Y_{i,d}$ is worker output for worker i on day d measured in metres of cloth. A log specification allows us to interpret regression coefficients more easily and reduces sensitivity to outliers. α_i is a worker specific fixed effect allowing an idiosyncratic daily output level for each worker. 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 a two degree polynomial in rainfall to control flexibly for precipitation. $WBGT_d$ 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 quartiles of the temperature distribution.This allows us to estimate the marginal effect on output of a change in temperature within different quartiles of the distribution.

Table 2 summarizes our results (omitting all fixed effects for clarity). The first four rows report the impact on the dependent variable (production) for a one degree increase in wet bulb globe temperature, estimated for four different ranges of WBGT with cut-offs corresponding to the quartiles of the distribution observed over the year. Because WBGT is an index that combines various parameters, for ease of interpretation the temperature ranges presented in the table are all normalized to 65 percent relative humidity.

We use a range of models including the baseline model of Equation 4 with both month fixed effects and day of week fixed effects (column 2) as well as more restrictive models allowing for firm specific month shocks and firm specific weekday fixed effects. We use both the log of output as our dependent variable (our preferred specification) and a level model with meters of cloth produced by a worker as the dependent variable. Our basic results remain similar across models, although in the most restrictive specification, where we introduce *week* fixed effects, our estimated coefficients are less precise since this model leaves very little residual variation in the data.

Of primary interest here is in the impact of a change in WBGT, within quartiles of dew point temperatures. As we have discussed, human physiology leads us to expect a non-linear impact of temperature on productivity with a slope going from positive or zero at low and moderate temperatures to negative impacts as the underlying WBGT increases (Test 1 in Section I). This pattern is borne out quite clearly in Table 2.

The impact of a one degree change in WBGT[6] when it is relatively cold (temperatures between $17°C$ and $24.34°C$ at 65% relative humidity) is small but positive and significant across all speci-

[6]While the estimates in Table 2 tell us the impact of WBGT on temperature, it is straightforward to convert these to temperature units (at a given relative humidity). At the reference humidity of 65% a one degree rise in temperature corresponds to a *less than* one degree increase in WBGT. Thus all estimates in Table 2 must be scaled down slightly to reflect the effect of a temperature change alone (the scaling factor ranges from 0.78 for the lowest WBGT quartile to 0.93 for the highest quartile). The required multiplier is easily obtained by differentiating Equation 1 with respect to temperature.

fications (Row 1). This is the baseline effect of WBGT in the regression specifications (with Rows 2,3,4 telling us how this baseline effect changes at higher temperatures). Row 2 suggests that the incremental change from this baseline level for moderate temperatures is small and *negative*, albeit not statistically significant. Rows 3 and 4 tell us the change from the baseline effect at higher levels of WBGT. Across specifications, these estimates are negative, statistically significant and large. As an example, for our baseline model (column 2) the net effect of a one degree change in WBGT for temperature ranges in the fourth quartile can be obtained by adding the baseline impact (0.018) to the change in impact at high temperatures (-0.067), giving us a mean impact of -0.049 (a reduction of 4.9% of output on average).

Table 2 provides some useful insight into the effects that temperature might have on worker productivity even for tasks that are not physically demanding and even in indoor settings. Another possible mechanism through which temperature might affect labour is through its impact on absenteeism. We are able to observe whether or not a worker is present for work at the firm he normally works for and can therefore test to see if absences are associated with high temperature days. We therefore estimate a simple linear probability model as follows

$$IsPresent = \alpha_i + \gamma_M + \omega_W + \beta_k WBGT_d \times \mathrm{D_k} + \theta f(R_{i,d}) + \epsilon_{i,d} \tag{5}$$

where $IsPresent$ is a binary variable that takes the value 0 when the worker does not report for work and 1 otherwise[7]. The covariates on the right have the same definitions as in Equation 4. We find no evidence in this setting that workers are more likely to be absent when temperatures are higher (Table 3) although as remarked earlier, this may not be surprising in the context we study.

III. National Level Evidence from the Annual Survey of Industry

In Section II we provide evidence suggesting that even in industrial settings where worker activities are not very physically demanding, temperature increases are associated with decreases in labor output and consequently reduced productivity. This case study evidence is intended to provide direct empirical evidence on the plausibility of the channel through which we have hypothesized that the environment might influence economic output, namely the direct effect of temperature on labor. Nevertheless by itself the evidence in Section II is of limited utility in determining whether temperature has a significant and direct causal impact on economic output more broadly. There

[7]Of course a worker who does not report for work at the firms we observe may work elsewhere so a measure of absence is not the same as observing a day off. Nevertheless the two are likely strongly correlated.

Table 2—Effect of Wet Bulb Globe Temperature on Weaving Worker Output

	Dependent variable:							
At 65% relative humidity	log(Meters)						Meters	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
WBGT:[$\leq 24.34°C$]	0.015**	0.018***	0.017***	0.018***	0.018***	0.020***	1.761**	1.753**
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.007)	(0.827)	(0.806)
X [$24.34 - 29.00°C$]	−0.014	−0.014	−0.016	−0.016	−0.016	−0.003	−1.929	−2.044*
	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.014)	(1.412)	(1.240)
X [$29.00 - 30.62°C$]	−0.063***	−0.057***	−0.056***	−0.053***	−0.053***	−0.046*	−2.015	−1.605
	(0.020)	(0.020)	(0.020)	(0.020)	(0.020)	(.026)	(1.985)	(1.998)
X [$\geq 30.62°C$]	−0.055	−0.067**	−0.064**	−0.066**	−0.066**	−0.036	−5.368*	−5.363**
	(0.034)	(0.034)	(0.032)	(0.032)	(0.032)	(.026)	(2.775)	(2.616)
RAIN	0.012	0.010	0.008	0.009	0.008	0.008	1.761*	1.859*
	(0.008)	(0.008)	(0.008)	(0.008)	(0.021)	(.009)	(.978)	(0.986)
RAINSQ					0.0004			
					(0.005)			
Worker FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Month FE	*Y*	*Y*	*N*	*N*	*N*	*N*	*Y*	*N*
Weekday FE	*N*	*Y*	*Y*	*N*	*N*	*Y*	*Y*	*N*
Month x Firm FE	*N*	*N*	*Y*	*Y*	*Y*	*N*	*N*	*Y*
Weekday x Firm FE	*N*	*N*	*N*	*Y*	*Y*	*N*	*N*	*Y*
Week FE	*N*	*N*	*N*	*N*	*N*	*Y*	*N*	*N*
Adjusted R^2	0.008	0.013	0.036	0.043	0.043	0.025	0.009	0.040
Mean Daily Output (mts)	124.77	124.77	124.77	124.77	124.77	124.77	124.77	124.77

Note:

1. *p<0.1; **p<0.05; ***p<0.01
2. Temperature ranges corresponding to WBGT quartiles expressed at 65 percent relative humidity
3. Cluster robust standard errors (Arellano-Bond) correcting for serial correlation and heteroskedasticity
4. Coefficients for models 1-6 may be interpreted as percentages after multiplying by 100.
5. Coefficients for columns 7, 8 are in units of daily meters of cloth per worker per day.
6. WBGT is expressed in units of temperature at a reference humidity of 65% (the annual average humidity in Surat).

Table 3—Effect of Temperature on Weaving Worker Presence

	Dependent variable: Binary Attendance Indicator
	IsPresent
WBGT:[$\leq 25.6°C$]	0.002
	(0.003)
X [$25.6 - 30.58°C$]	0.0002
	(0.005)
X [$30.58 - 32.29°C$]	−0.015
	(0.010)
X [$\geq 32.29°C$]	−0.003
	(0.012)
RAIN	0.001
	(0.004)
Constant	0.681***
	(0.056)
R^2	0.422
Adjusted R^2	0.420
Note:	*p<0.1; **p<0.05; ***p<0.01; Fixed effects omitted for clarity

are two reasons for this.

First, as the discussion in I.A makes clear, there are varying degrees to which the physiological impacts of temperature on workers can be expected to translate into economic impacts. How these two quantities are linked depends on the production function of the firm in question. For this reason the evidence from workers in Surat suggests that temperature might matter to a limited degree for small weaving units but does not in itself tell us very much about whether temperature is an economically significant determinant of productivity on a national scale. In this section therefore we consider how temperature changes affect average output aggregated over a nationwide sample of manufacturing plants of different sizes and spanning a wide range of sectors.

Secondly, we use high frequency daily variation to identify temperature impacts on weaving worker output in Section I.A. Although the physics of heat stress on human beings is dependent on contemporary exposure to elevated temperature (and should therefore be identifiable using quasi-random short term variation as well as long run changes), it may nevertheless be the case that firms invest in some degree of adaptation when temperatures increase over longer time periods. Such adaptation could reduce the economic importance of temperature shocks by limiting the degree to which workers are exposed to ambient temperatures. For this reason it is useful to examine the sensitivity of manufacturing output to temperature using variation aggregated over a longer time period than a day. In this section therefore we examine how annual variation in temperature changes the output of individual manufacturing plants.

A. Manufacturing Plant Data

We begin by constructing a panel dataset of annual observations at the level of a discrete manufacturing unit made up of factories sampled from across India using the Annual Survey of Industry (ASI). The ASI is a rich survey of individual plants carried out on an annual basis in every state within India. The survey is carried out every year the by Field Operation Division (FOD) of National Sample Survey Organization (NSSO), Government of India through its network of Zonal/ Regional/ Sub- Regional offices. The population eligible to be surveyed consists of industrial plants registered under India's Factories Act. Registration normally implies that a plant has at least 10 employees in total. Each annual cross-section includes every unit from the population of registered firms who employ over 100 workers (not including short term contract labour) as well as a random sub-sample of units smaller than this cut-off. This sub-sample is sized to cover 18 percent of total eligible units every year. The ASI has been carried out annually since 1960, but we focus for this study on the survey years between 1998-99 and 2008-09. Between these years, survey micro-data may be purchased with a panel identifier so that it is possible to identify repeated observations on plants across years.[8] 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. Entry and exit further reduce the number of observations per unit.

The survey is intended to capture critical variables relating to factory level production inputs and outputs (in both physical and monetary units and including energy inputs), annual income and expenditures under various heads, labour utilization, wages and annual man-days worked. In addition the survey reports basic characteristics of the surveyed factory including a full NIC classification (sector). Of primary importance to us in this study is that the survey provides a measure of the total value of output produced at the end of the financial year for every plant. This is calculated by multiplying the market price of all products manufactured with the production quantity (the ten most important categories are separately tabulated alongside the total output). It is this quantity that we will use as our primary measure of economic output.

Before we use the data from the Annual Survey of Industry, we carry out a few simple data-cleaning operations. First we remove firms that are surveyed two times or less over the entire period 1998-2008. Since our identification strategy relies on within manufacturing unit variations in output correlated with temperature, firms observed twice or fewer times cannot contribute to

[8]The panel dataset as made available does not provide the geographical location of a factory precisely. However an alternative version of the same data is made available by the NSSO with the district in which a plant is located, but without a panel identifier for each plant. We match observations aross both versions to generate our final dataset which contains both a plant identifier and a district identifier for each observation.

our results. Secondly, before estimating our various econometric specifications we trim the top 2.5 percent and bottom 2.5 percent of the distribution of observations by output value and workforce size. This is done to transparently eliminate outliers (units with implausibly large output values or zero and negative output) that are likely due to data entry errors or otherwise represent special cases. Lastly we remove a small number of firms who report having less than 10 workers employed because this represents a discrepancy between firm data on file with the government (and used to select the survey sample) and what is actually reported by the firm. This discrepancy is also plausibly associated with false reporting since firms with less than 10 workers are subject to very different labor laws and taxation regimes under Indian law. In addition, firms with a very small total workforce are unlikely to be relevant to the questions we ask in this paper, since our interests are explicitly in the impacts of temperature on workers. All remaining observations form part of our dataset. Overall we are left with about 28076 manufacturing units that are observed at least three times over the nine years of our panel.

The ASI does have some drawbacks that limit the inferences we are able to draw from it. First, a time series measure of individual firm output is available only over a limited period of time (1998-2008). Because our estimation strategy relies only on within firm variation in output coincident with variation in temperatures across years, a short time series limits the temperature variation that is empirically observed and therefore reduces our ability to identify impacts as precisely as a longer panel dataset might enable.

Secondly, the survey is not a good representation of small manufacturing units in India. This includes many firms that are not registered under the Factories Act because full time, contractual employees fall below the registration cut-off. This informal and small scale manufacturing sector plays an important role in Indian manufacturing and is plausibly more vulnerable to climate shocks of all kinds. To the extent that we cannot observe this population our results may under-estimate the vulnerability of the manufacturing sector in India to temperature. On the other hand, focusing on registered firms may provide results that are more generalizable and to represent climate sensitivity for firms that have a certain minimal size and capital investment and therefore presumably, a certain minimal level of adaptive capacity. Also while developing countries tend to have a large informal and small scale manufacturing sector, there are important economic and technological advantages to operating at a larger scale and therefore over time one might expect manufacturing to agglomerate and the left tail of small units to shrink. For instance, there is evidence that this process has been occurring in China (Liu and Li, 2012, Wen, 2004).

B. Meteorological Data

The Annual Survey of Industry enables us to locate every surveyed unit down to the level of a district in the country. Districts in India are an administrative subdivision of a state and for most purposes are the most granular unit at which administrative and economic data is reported. There are 609 districts in the country, with an approximate size of the order of 50 square kilometers. Knowing the geographical location of an industry with some precision is important because the temperature of interest to us is the local temperature where a manufacturing plant is located. We therefore match observations in the ASI data with temperature and precipitation observations at the district level. To do this we use a 1.0 degree 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. This dataset represents the highest quality temperature record we are aware of for India.

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). This gridded temperature and rainfall data is then mapped to a district measure by locating the district centroid and then assigning to each district the weighted average of temperature and rainfall measures from all grid points within a 200km radius of the centroid. Weights are taken inversely proportional to the square of the (great-circle) distance between grid points and the district centroid. Our original temperature and precipitation records are reported daily over the entire nine year period of interest, but we aggregate these up to annual measures as needed (since factory output is reported annually). The availability of daily data is critical however in order to accurately generate degree day measures for each year and each district (we discuss this in more detail in Section IV).

One last point is worth noting. As we discussed in Section II, strictly speaking, the environmental quantity of most direct relevance to heat stress on workers is not simply temperature but rather the wet bulb globe temperature, an index that also accounts for ambient humidity. 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. To be specific, although water vapour pressure or humidity measures are available as part of reanalysis datasets, these models were not necessarily designed to provide reasonable estimates of temporal variation in humidity and related parameters.[9] In addition little is known about how relative humidity might change over time due

[9] Auffhammer et al. (2013) discuss some of the considerations involved in using intra-annual temperature variation from

to climate change effects.

For this reason in Section IV we use dynamic variation in temperature alone to estimate the effect of heat on industrial output in our nationwide analysis. This ensures that our results are driven only by variation in temperature. However, as a robustness check we also repeat our analysis using a wet bulb globe measure obtained by combining temperature with *long run average* measures of daily relative humidity between 1981-2010 from the NCEP/ NCAR reanalysis datasets[10]. Dynamic variation in this index is however driven only by variation in temperatures over grid points in our dataset. This approach therefore is equivalent to re-weighting temperature observations based on average relative humidity measures.

IV. Methods and Results

To identify the impact of temperature on industrial productivity we exploit the presumably quasi-random variation in year to year local average temperatures and estimate the response of the industrial units to this variation, relative to their own average production level and controlling for capital inputs available at the start of the financial year. In doing so we identify the impacts of temperature on output primarily through dynamic local variation and its effect on a single manufacturing unit. This ensures that we can isolate the effect of temperature, independent of other variables that might be associated with sustained temperature differences between units but might affect independently affect output (such as altitude for example). Formally we estimate the following regression equation,

$$V_{i,t} = \alpha_i + \gamma_t + \omega K_{i,t} + \beta T_{i,t} + R_{i,t} + \epsilon_{i,t} \tag{6}$$

where $V_{i,t}$ is the recorded value of output produced by a specific industrial unit i during financial year t. This quantity is essentially the product of physical output with average prices per unit product produced (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. In our most restrictive specifications we allow separate fixed effects for each manufacturing sector at the two digit NIC level (46 sectors are represented in the data). $T_{i,t}$ 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.

reanalysis data. Because relative humidity is not a primary parameter against which these models are calibrated, these concerns are likely to be significantly more serious when using humidity output from climate models

[10]The NCEP/NCAR outputs are unfortunately available only over a relatively coarse 2.5 by 2.5 degree grid which we interpolate using a procedure similar to that followed for temperature

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). Capital on hand at the start of the financial year is converted by the plant into labour 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 capital stocks at the start of the year enables us to cleanly identify the impact of temperature realizations during the year 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 is not true of labor, energy or raw material expenditures actually realized during the year. For instance, in our case study of weaving workers in Surat 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 6 is estimated and the results reported in Table 4 column 3. Columns 1,2, 6 and 7 all present slightly different specifications as a robustness check. Models 6 and 7 for example use the log of output as the independent variable (with the log of working capital among the controls). Model 1 is the same as in Equation 6 but without a control for capital available at the start of the year (the estimated temperature effect is similar to Model 3 but unsurprisingly more noisy). Models 2 and 6 are less restrictive models using a single year fixed effect across all plants instead of sector specific year fixed effects. We also estimate (but do not report for brevity) models with an additional quadratic control for rainfall and find our results substantively unchanged.

Note that the coefficient β on mean temperature is reported as a percentage of average output in columns 1-5. For the logged models the coefficient can be interpreted as a percentage change after multiplying by 100. For our preferred specification (Model 3) we estimate an average of a 3.2 percent decrease in output for a 1 degree rise in average daily temperature in the year. In the logged specification the estimate is slightly lower at 1.8 percent (of course the two models are not exactly comparable). It is interesting to note that this estimate of a 1.8 percent output loss per degree rise in temperature at the level of individual manufacturing plants is close to the reported percent change in economy wide output for a one degree change in temperature in Hsiang (2010) (2.4 percent) as well as the percentage decline in labour supply reported by Zivin and Neidell (2010) (1.8 percent).

TABLE 4—EFFECT OF TEMPERATURE ON MANUFACTURING INDUSTRY OUTPUT

	Dependent variable:							
	Plant Output Value					Log Plant Output Value		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mean Temperature	−2.78***	−2.79***	−3.22***			−0.016**	−0.018**	
	(1.02)	(0.93)	(0.94)			(0.008)	(0.008)	
$0°C - 20°C$				3.20	1.83			0.022
				(2.48)	(2.48)			(0.021)
$20°C - 25°C$				−1.61	−0.41			−0.035*
				(2.28)	(2.28)			(0.019)
Above $25°C$				−5.07***	−5.98***			−0.022**
				(1.27)	(1.27)			(0.011)
rainfall	1.33***	1.03***	1.02***	1.06***	1.06***	0.005***	0.003*	0.004**
	(0.24)	(0.22)	(0.22)	(0.22)	(0.22)	(0.002)	(0.002)	(0.002)
Plant FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Year FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Sector-Year FE	*N*	*N*	*Y*	*N*	*Y*	*N*	*Y*	*Y*
Capital Controls	*N*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Number of Plants	28706	28706	28706	28706	28706	28706	28706	28706
Adjusted R^2	0.077	0.180	0.180	0.180	0.180	0.149	0.149	0.149

Note:

1. *p<0.1; **p<0.05; ***p<0.01
2. Cluster robust standard errors (Arellano-Bond) correcting for serial correlation and heteroskedasticity
3. Maximum temperature in afternoon is on average $6°C$ above the mean temperature
4. Coefficients for models 1-5 are expressed as percentages of average output level.
5. Coefficients for columns 6-8 can be multiplied by 100 to obtain approximate percentage changes.

A. Non-linearities in Temperature Response

The results of associating variation in average temperature against output suggest that temperature matters for manufacturing productivity. However it is possible to gain much greater insight into the mechanism involved, and to carry out an important robustness check by considering what we know from the heat stress literature about the nature of human physiological response to temperature.

The literature on heat stress makes a strong prediction of a non-linear response of human beings to temperature. Figure 1 reproduces a graph from Hsiang (2010) based on a meta-analysis of over 150 ergonomic studies and Table 1 provides results at higher temperatures from Kjellstrom et al. (2013).

A clear prediction from this type of relationship, and one that corresponds to intuition, is that the impact of a one degree change in temperature should matter most when temperature is relatively high. Figure 1 and Table 1 suggests a fairly sharp cut-off at around 25 degrees Celsius after which human performance begins to deteriorate sharply. Conversely, at lower temperatures, worker performance should remain largely unaffected by changes in temperature.Consequently, if the mechanism behind the results in Table 4 is in fact a reflection of temperature impacts on worker productivity we should expect to see a similar pattern of industrial output response as well. This suggests a simple but powerful test that we can take to the data.

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

$$V(T_d) = V(T_o) + \int_{T_o}^{T_d} \frac{\partial V}{\partial T} dT \tag{7}$$

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

$$\bar{V}(T_d) = \bar{V}(T_0) + \sum_{k=1}^{N} \beta_k D_k(T_d) \tag{8}$$

Here

$$D_k(T_d) = \int_{xl_k}^{xu_k} \mathbf{1}[\mathbf{T_d} \leq \mathbf{x}] dx \tag{9}$$

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 within a given temperature interval. 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. In other words we can write annual output V_t as a function of degree days D_k as follows

$$V_t = V_0 + \sum_{k=1}^{N} \beta_k D_k \tag{10}$$

Because we do in fact 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

$$V_{i,t} = \alpha_i + \gamma_{s,t} + K_{i,t} + \sum_{k=1}^{N} \beta_k D_k + f(R_{i,t}) + \epsilon_{i,t} \tag{11}$$

and observe whether β_k values associated with fluctuations in degree days vary in the way the heat stress literature suggests. In other words, if temperature impacts industrial productivity because of 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.

This response function also suggests that it is the degree day model of Equation 11 that should be of primary interest to us since the estimated impact of temperature on productivity from Equation 6 is simply an average value over degree day changes actually observed in the literature. Historic variations in temperature however do not necessarily correspond to the 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. In other words, the predicted impacts of climate change is precisely to increase the number of degree days in higher temperature buckets, shown here to be quite strongly associated with manufacturing output.

Columns 4,5 and 8 of Table 4 provides degree day results for different specifications (Model 5 is the most restrictive and our preferred specification). The effect of temperature on productivity at lower temperatures (below 25 degrees C) is small and not statistically significant. At temperatures above 25 degrees, output responds strongly to temperature increases (see also Figure 2 Even granted that temperature impacts on agriculture may have spillover impacts on other sectors, this type of relationship is difficult to reconcile with such a hypothesis because there is no strong reason to believe that agriculture is unaffected by temperature variations below 25°C. If anything the opposite might be expected to hold since the primary growing periods in India typically coincide with cooler temperatures and not the hottest summer months. It is important to keep in mind here that the temperature degree measure refers to the average temperature over a 24 hour period. This is significantly lower than the afternoon peak which may be closer to the ambient temperature during much of the working day. The average difference between maximum and mean temperatures across the entire country is about 6 degrees. Thus a day with a mean temperature of 25 degrees celsius may have a maximum temperature over 30 degrees and daytime ambient temperatures are likely to be in the high twenties.

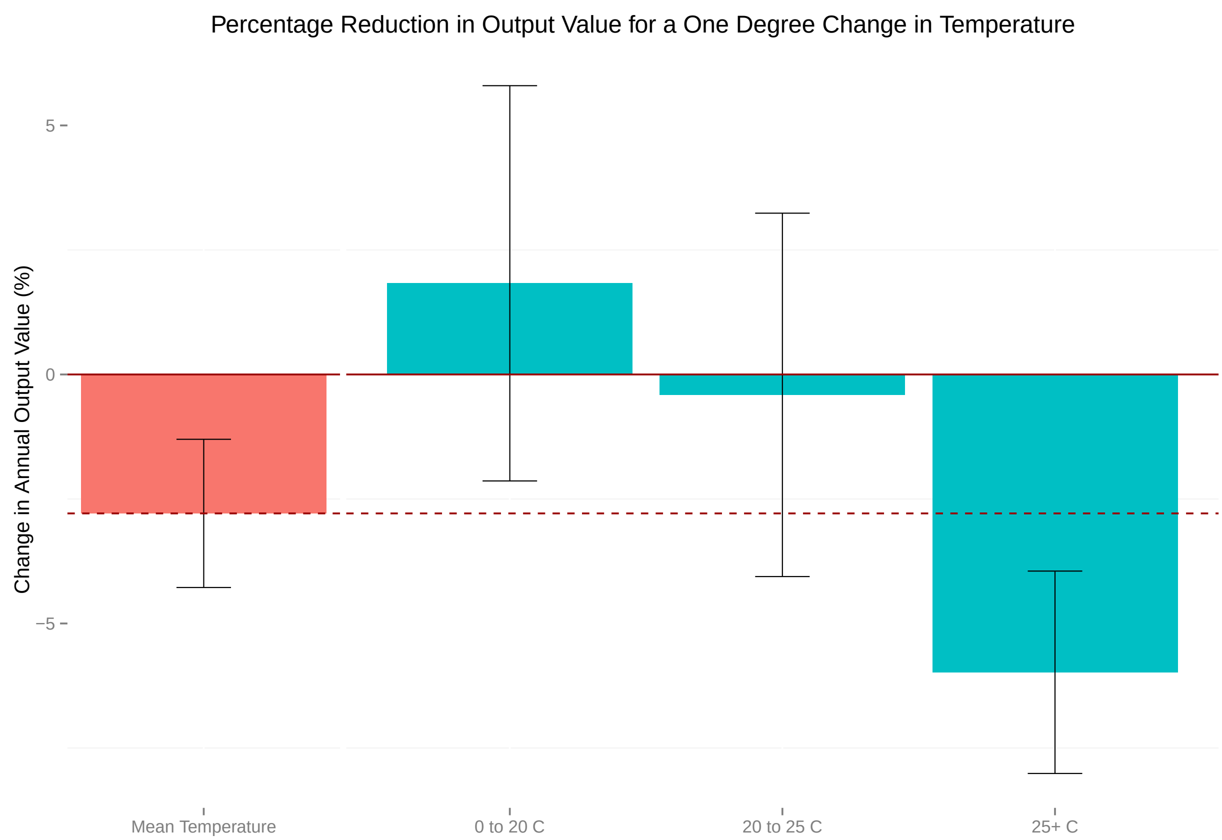

FIGURE 2. IMPACT OF A ONE DEGREE RISE IN TEMPERATURE ON DAILY PRODUCTION BY DEGREE DAY BIN. 90% CONFIDENCE INTERVALS PLOTTED

B. Results using Wet Bulb Globe Temperature

As we described in Section II the environmental quantity of most direct relevance to heat stress on workers is not simply temperature but rather the wet bulb globe temperature, an index that also accounts for ambient humidity. 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. Although water vapour pressure or humidity measures are available as part of reanalysis datasets, these models were not necessarily designed to provide reasonable estimates of temporal variation in humidity and related parameters. In addition significantly less is known about how relative humidity might change over time due to climate change effects.

For this reason we have used dynamic variation in temperature alone to estimate the effect of heat on industrial output in our analysis thus far. This ensures that our results are driven only by variation in temperature and not potentially spurious variations in humidity measures from reanalysis data. However as a robustness check we also repeat our estimation of Equation 6 and 11 using an approximate wet bulb globe measure. This is obtained by combining temperature with *long run average* measures of daily relative humidity between 1981-2010 from the NCEP/ NCAR reanalysis datasets. Dynamic variation in this index is therefore still only driven by variation in temperatures over grid points in our dataset. This approach therefore is equivalent to re-weighting temperature observations based on average relative humidity measures.

Table 5 summarizes our results which look very similar to those in Table 4. The estimate of a one degree rise in mean temperature on productivity using the WBGT measure is also negative and significant although greater in magnitude than that obtained using temperature alone. The higher magnitude using this type of model also makes sense because a one unit change in a WBGT index measure is associated with a *greater* than one unit change in temperature[11] (see Equation 1).

C. Variation in Impacts by Value Added by Labour

In Section I we argued that if temperature shocks result in reduced worker productivity we might expect that this effect should result in percentage declines in production that are highest in manufacturing sectors with a high value added per worker. Conversely, in sectors where the value added per worker is low, one might expect that plant output may be less affected by climate shocks.

The intuition here is that the impact of temperature on human performance is not the quantity of ultimate economic interest, rather it is the degree to which such performance deterioration

[11]At 25 degrees Celsius and 65% relative humidity a one unit increase in temperature corresponds to a 0.84 unit increase in wet bulb globe temperature.

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

	Dependent variable:							
	Plant Output Value					Log Plant Output Value		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
WBGT	−3.08**	−3.34***	−3.94***			−0.021**	−0.023**	
	(1.25)	(1.15)	(1.15)			(0.010)	(0.010)	
Under $20°C$				0.89	0.55			0.015
				(3.18)	(3.18)			(0.027)
$20°C - 26.5°C$				−2.91	−2.39			−0.064***
				(2.79)	(2.80)			(0.024)
Over $26.5°C$				−2.22**	−2.66**			−0.004
				(0.99)	(0.99)			(0.009)
rainfall	1.35***	1.03***	1.01***	1.05***	1.05***	0.005**	0.004*	0.005***
	(0.24)	(0.22)	(0.22)	(0.22)	(0.22)	(0.002)	(0.002)	(0.002)
Plant FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Year FE	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Sector-Year FE	*N*	*N*	*Y*	*N*	*Y*	*N*	*Y*	*N*
Capital Controls	*N*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*	*Y*
Number of Plants	28706	28706	28706	28706	28706	28706	28706	28706
Adjusted R^2	0.077	0.180	0.180	0.180	0.180	0.149	0.149	0.149

Note:
1. *p<0.1; **p<0.05; ***p<0.01
2. Cluster robust standard errors (Arellano-Bond) correcting for serial correlation and heteroskedasticity
3. WBGT levels are daily means and are expressed in Celsius at a reference humidity of 65%
4. Coefficients for models 1-5 are expressed as percentages of average output level.
5. Coefficients for columns 6-8 can be multiplied by 100 to obtain approximate percentage changes.

has consequences for productive economic output. These two quantities are linked by the degree to which workers are engaged in high value added tasks or equivalently by the importance of a marginal change in labor inputs to output. Thus while all labour becomes less productive in extreme temperatures, small changes in productivity are most likely to translate to meaningful economic outcomes when the tasks in question involve a substantial value addition.

In order to test this hypothesis we need a measure of the whether the value added by labour within a particular sector is high or not. One simple approach is to calculate the average ratio of output value V produced to the reported number of employed workers for all units within a manufacturing sector. Naturally this quantity is not the same as the marginal value of an additional unit of labour but we regard this as an indicative proxy since sectors where the average unit generates a large amount of output value relative to the number of employed workers are plausibly likely to be most sensitive to a reduction in the productivity of their workers. We therefore calculate this quantity for each of the 46 manufacturing sectors represented in the data (at the 2 digit NIC level) and create a dummy variable which is 1 for the sectors in the top quartile by the ratio of value created to workers employed and 0 otherwise.

Next we compare the distribution of the number of employees in plants belonging to the two sectoral groups identified by our dummy variable. (sectors with a high value to labor ratio and those with a lower value to labor ratio). We find that although sectors with a high value to labor ratio do employ fewer workers, the distributions overall look reasonably similar, suggesting that it is not the case that one group of plants is so heavily mechanized that labor is not an important factor in production.

In order to gain some insight into how these two sub-samples of our data respond to temperature we therefore regress the log of factory output on temperature in a specification similar to Equation 6 as well as on binned degree days (similar to Equation 11). In addition we include interaction terms allowing temperature effects on production to vary across the two groups. We restrict attention to logged output in these specifications because our focus is on identifying whether high value added sectors are more sensitive to temperature shocks. Therefore we are interested in percentage changes relative to a firms output levels (rather than absolute reductions in output). It is possible to interpret coefficients in the logged specification (multiplied by 100) as a measure of percentage impacts.

$$log(V_{i,t}) = \alpha_i + \gamma_t + \beta T_{i,t} \times VA_i + f(R_{i,t}) + \epsilon_{i,t} \tag{12}$$

Here VA_i is the constructed dummy variable which takes a value of 1 for units belonging to sectors in the top quartile by value created per worker and 0 otherwise. We are interested in the coefficient of the interaction between temperature $T_{i,t}$ and VA_i with a negative sign indicating a greater percentage impact of temperature occurs on firms in sectors where the value added per worker is highest. Table 6 summarizes our estimates. We find that sectors with manufacturing units with the highest value added per worker are significantly more negatively impacted by temperature shocks. While not a perfect test of proposition 2 in Section I, this provides further evidence that the channel through which temperature acts in this setting is indeed linked to its effects on labor.

Table 6—Differential Effect of Temperature on Sectors with High Value Added by Labour

	Dependent variable:		
	logvalue		
	(1)	(2)	(3)
meant	−0.008	−0.010	
	(0.009)	(0.009)	
rainfall	0.005**	0.004*	0.005**
	(0.002)	(0.002)	(0.002)
meant $\times VA$	−0.036**	−0.034*	
	(0.015)	(0.018)	
$0°C - 20°C$			0.017
			(0.024)
$20°C - 25°C$			−0.030
			(0.022)
Above $25°C$			−0.009
			(0.012)
$0°C - 20°C \times VA$			0.025
			(0.052)
$20°C - 25°C \times VA$			−0.019
			(0.047)
Above $25°C \times VA$			−0.056**
			(0.025)
Plant FE	*Y*	*Y*	*Y*
Year FE	*Y*	*Y*	*Y*
Sector-Year FE	*N*	*Y*	*Y*
Capital Controls	*Y*	*Y*	*Y*
Observations	135,607	135,607	135,607
Adjusted R^2	0.149	0.149	0.149
Note:	1. *p<0.1; **p<0.05; ***p<0.01		
	2. Cluster robust standard errors (Arellano-Bond)		
Labor Size Quantiles (VA=0)	[0.0]: 10, [0.25]: 31, [0.50]: 84, [0.75]: 189, [1.0]: 824		
Labor Size Quantiles (VA=1)	[0.0]: 10, [0.25]: 28, [0.50]: 74, [0.75]: 161, [1.0]: 824		

V. Qualitative Evidence and Adaptation Strategies

To this point we have attempted to provide some evidence suggesting that temperature changes reduce manufacturing unit productivity. Further we have argued that the pattern of response, as well as an underlying literature on heat stress, suggests that reductions in labour productivity at elevated wet bulb globe temperatures may be the mechanism through which this occurs.

In a series of detailed interviews with owners and workers at weaving units in Surat we found anecdotal evidence to support this claim. Furthermore it does not seem to be the case that industry owners are unaware that temperature may influence the attention paid by workers, the speed at which they work, the number of breaks required and the faults made. We document these findings in more detail as part of a separate paper (Sudarshan and Tewari 2013) but summarize some key insights here.

First it is not obvious that additional adaptive strategies are easily found or implemented without significant costs. Weaving plants in our case studies for example already take a number of low cost actions to mitigate the impact of ambient temperatures on workers. First, breaks are built into the work day and these breaks are increased in hotter weather. Second, weaving units attempt to locate worker intensive machinery in basements or lower floor of buildings where temperatures are lower[12]. Third, while fans or air-conditioning did not seem to be commonly used, plants did seem to make a significant effort to introduce natural draft ventilation by creating windows and allowing cross ventilation to occur. In one instance the owner of one of the weaving firms we interviewed reported that he had made an order to a Chinese firm to purchase a low cost air cooling system. This planned capital investment seems to make this particular owner something of an outlier but represents an example of the type of adaptive strategy (and expense) that may grow more common if temperatures rise or labor competition increases. A key factor that interviewed weaving plant owners cited as a reason for having to live with temperature shocks (which apart from production affects also increase worker discomfort and may encourage worker attrition or decrease the attractiveness of the job) had to do with the energy costs of adaptation. Simply put, air conditioning large production floors (or investing in air cooling) may be too expensive to become a common strategy for firms in highly competitive markets such as the weaving sector. Furthermore, unlike in the case of agriculture for example, where farmers may be able to shift from climate sensitive crops such as wheat to less sensitive crops such as millets, manufacturing plants have only one species available as labor input! In addition the ability of human beings to acclimatize to elevated wet bulb globe temperates is likely to be limited (Wyndham, 1969). Thus it is plausible that adaptation can occur only through reducing exposure to elevated temperatures which may require capital investment or even plant relocation.

We also conducted a few interviews with diamond sector plants in Surat. The diamond sector

[12]Of course these location decisions are also driven by vibration effects of machinery on buildings and therefore may not occur purely for temperature related reasons.

is interesting as a counterpoint to weaving units because although both are highly labor intensive, the value added by diamond workers is significantly higher and workers in this sector are often highly skilled. We were unable to collect worker level output data from these plants but found that air-conditioning investments are significantly more common in diamond firms, even where the unit size is small and the number of workers employed is low (of the order of 10-20 workers). In larger firms all production steps were found to take place under completely climate controlled conditions. In the smallest firms this was not necessarily the case, but production steps involving very high value addition did tend to be air conditioned.

For example the first step in most diamond units is a sorting activity where raw stones are brought in, assessed for quality, and then sorted into different bins. Low value stones may be sold to other polishing units while the higher value stones are sent for polishing and cutting. Sorting is done by human beings over sorting tables and is work that requires very little physical effort but a significant degree of concentration and skill. Mistakes made during sorting have the potential to be highly expensive since they may result in mis-classifying high value stones as low value stones or vice versa. We found that even the smallest units tend to use air-conditioning at the sorting step of their production process.

VI. Temperature Impacts from Climate Model Projections

It is possible to use the estimates obtained in Section IV, in conjunction with climate change projections of anticipated changes in the distribution of annual degree days within India to obtain an estimate of the costs imposed by climate change on industrial productivity. In what follows we use the estimated impacts of temperature from our most rigorous specification, Model 5 in Table 4 which includes controls for sector-year fixed effects as well as controlling for working capital at the start of the year.

There are various model based projections available for climate change driven temperature changes over the globe and there is some inherent modeling uncertainty in all these forecasts. This uncertainty is hard to quantify so we use two different climate models to compare results.

To accomplish this we begin with a projection of changes in the annual distribution of days across temperature bins for India as documented in Burgess et al. (2011). These projections are based on a combination of historical climate data from the Reanalysis 1 project of the Climactic Research Unit of the National Center for Environmental Prediction / National Center for Atmospheric Research (NCEP) in conjunction with projected climate data from two leading climate models. The models used for climate projections are (i) the A1F1 "business-as-usual" scenario of the Hadley Centre

Global Environmental Model, version 1 (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 (NCAR 2007). These projections are combined and aggregated up from a global 2 degree grid to national boundaries to produce estimates of historical and projected temperature distributions at the national level.

Figure 3 reproduces a graph comparing historical and projected temperatures in degree day bins over India. Perhaps the most striking feature of these projections is the skewed nature of projected climate change impacts which suggest a significant increase in annual degree day bins at temperatures above 25 degrees celsius (77 degrees Fahrenheit). We overlay on this graph our empirical estimates of the impact of a change in temperature across different degree day bins. The conclusion is quite stark. The impacts of climate change for India (a pattern that is repeated for a number of other developing countries) are expected to be concentrated in precisely the region of the temperature distribution that our discussion in Section IV suggests has the greatest impact on human productivity.

These climate projections can then be combined with our empirical estimates of the impacts of temperature to obtain a back of the envelope estimate of the expected impacts of this change in temperatures on industrial productivity, assuming no adaptation in excess of what is captured within the existing status quo (and thus implicitly accounted for in our estimates). This assumption of no adaptation may not hold good in the long run of course, but adaptation is not costless and indeed is plausibly more expensive the greater the increase in ambient temperatures. Thus understanding the order of magnitude of costs imposed by heat stress effects on manufacturing is useful as a bound on total costs and indicative of the type of adaptive measures that may become cost-effective.

The projected changes in degree days in Figure 3 combined with our mean estimates of the impact of temperature changes on productivity from Figure 2 allow us to compute a back of the envelope estimate of the impact on manufacturing output in India of projected climate change. We collapse the projected and historical temperature distributions into the coarser temperature bins (Under 20, 20-25 and 25+ degrees celsius) over which we estimate productivity effects of temperature. Temperature appears to have insignificant effects in the first two bins but has strong negative impacts in the third bin. 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. The change in the daily average degrees above $25°C$ in this projection is extremely high and a consequence of the large number of extreme temperature days the population weighted model

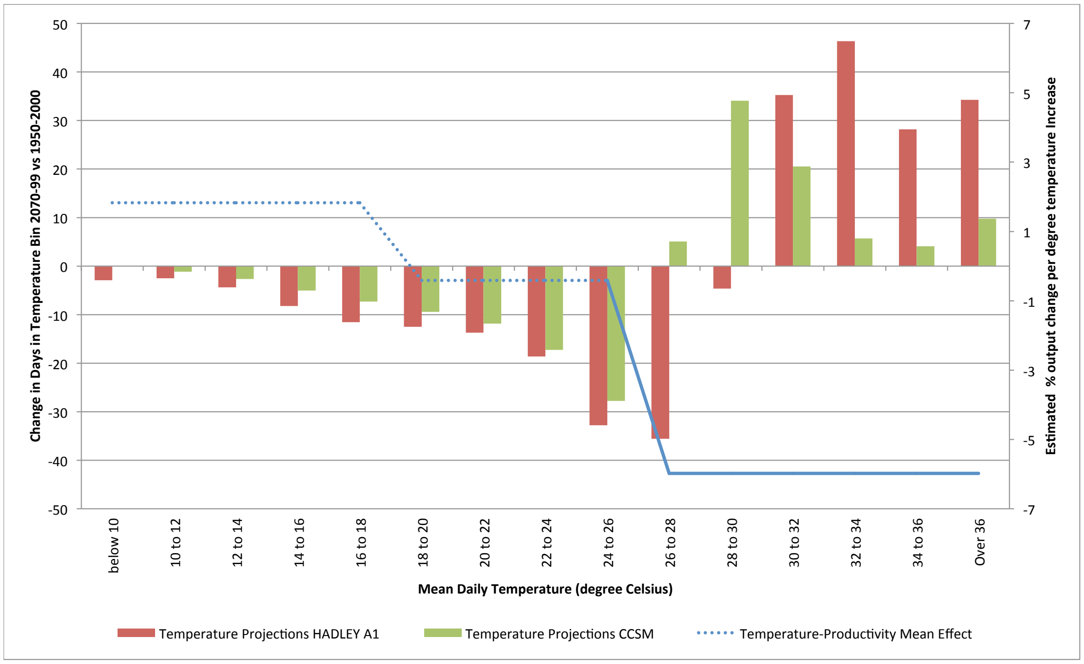

Figure 3. Historical and projected temperatures under a business as usual climate change scenario for India. See Burgess et al. (2011) for climate change projections. Lines denote estimated productivity effects of temperature with solid lines representing statistically significant effects

outputs predict for India (see Figure 3) as well as the fact that both populations and industrial activity are concentrated in relatively warmer parts of the country. 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 the more reasonable estimate our empirical estimates suggest an estimated impact on manufacturing of -7.89% (95% CI = [-4.60%,-11.21 %]).

While pinning down the magnitude of projected impacts more precisely may require further research and replication with data from other countries, we would argue that our results certainly suggest the need for much more attention to be paid toward this channel through which climate change and temperature might impact economic output. The magnitude of our back of the envelope calculation is large and economically highly significant and suggests that the quality of labour as an economic input may be significantly lower in the higher temperatures found in much of the developing world. Further, because climate change projections suggest that countries like India will primarily see an increase in the very hottest days of the year, the impact of environmental change on economic productivity and therefore growth may also be high. Mitigating these impacts is possible (and we reference some anecdotal examples in Section V) but this adaptation is likely to be expensive because it is hard to envisage adaptation options that do not involve recurring energy and capital expenditures on air cooling equipment or other infrastructure that might reduce exposure to heat.

VII. Conclusions

This paper presents empirical evidence from the manufacturing sector in India suggesting that manufacturing output decreases significantly as temperatures increase. We make reference to the literature on heat stress and the impact of temperature on human performance indices to argue that this physiological mechanism may result in labor that is less productive as ambient temperatures rise and this may impact industrial output across a range of sectors. We provide evidence for this suggestion using data at different levels of detail.

First, we collect detailed daily worker level output data from weaving plants in Surat and estimate a set of highly restrictive empirical specifications to show that worker output does indeed seem to drop at high wet bulb globe temperatures. Furthermore these negative impacts of temperature on worker output are concentrated at higher temperatures, consistent with our theoretical understanding of heat stress in human beings.

Secondly we create a panel dataset of financial and balance sheet data from an annual survey of individual manufacturing plants in India (the Annual Survey of Industry) and use dynamic variation in year to year temperatures to show that the output of individual plants seems to decrease as average temperature increases (net of controls for inputs available at the start of the year, and flexibly accounting for plant and time fixed effects). The average percentage reduction in output value associated with a degree change in temperature is of the order of about 2% per degree. These impacts are also associated with changes in the number of days with mean temperatures above $25°C$ but not below and seem concentrated in sectors where the value created per worker is high.

Finally we show that combining our empirical estimates of the economic impacts of temperature on manufacturing units with standard model estimates of temperature changes due to climate change suggests that the magnitude of the negative effect of climate change driven temperature changes, on labor productivity and thus economic output in India, may be large. We note that the physiology of human beings and the physics of heat transfer between the human body and the ambient provides a clear theoretical basis for observing impacts on labor productivity at elevated temperatures. Thus provided we take seriously the evidence provided in this paper, we should also consider that these impacts are likely to be similar in manufacturing and other sectors in countries other than India, and indeed may be a concern wherever climate change projections suggest a significant increase in high temperature days.

In conclusion, when taken in conjunction with previous work such as Hsiang (2010), we suggest that this particular channel through which the environment may affect growth may be much more important than is generally recognized. Further research using data from other countries and sectors including experimental work in industrial settings would prove extremely useful in helping us understand not just how temperature changes in the future may impact growth but also how historic temperature differences across different parts of the world may have contributed to existing differences in economic output between nations.

REFERENCES

D. Acemoglu, S. Johnson, and J. A. Robinson. Reversal of fortune: Geography and institutions in the making of the modern world income distribution. *The Quarterly Journal of Economics*, 117(4):1231–1294, Nov 2002. doi: 10.1162/003355302320935025. URL `http://dx.doi.org/10.1162/003355302320935025`.

Jesse K Antilla-Hughes and Solomon M Hsiang. Destruction, Disinvestment and Death: Economic and Human Losses Following Environmental Disaster. 2013. URL `http://ssrn.com/abstract=2220501`.

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`.

Alan Barreca, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph Shapiro. Adapting to Climate Change: The Remarkable Decline in the U.S. Temperature-Mortality Relationship over the 20th Century. Technical report, National Bureau of Economic Research, January 2013. URL `http://dx.doi.org/10.3386/w18692`.

Robin Burgess, Olivier Deschenes, Dave Donaldson, and Michael Greenstone. Weather and death in india. URL: http://faculty.chicagobooth.edu/workshops/AppliedEcon/past/pdf/AEW2011.

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.

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.

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.

S. M. Hsiang, M. Burke, and E. Miguel. Quantifying the influence of climate on human conflict. *Science*, 341(6151):1235367–1235367, Sep 2013. doi: 10.1126/science.1235367. URL http://dx.doi.org/10.1126/science.1235367.

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. URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2932627&tool=pmcentrez&rendertype=a

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 productivity – an increasing challenge for low and middle-income countries during climate change. *Global Health Action*, 2: Special Volume:46–51, 2009.

Tord Kjellstrom, Bruno Lemke, and Matthias Otto. Mapping Occupational Heat Exposure and Effects in South-East Asia: Ongoing Time Trends 1980 to 2011 and Future Estimates to 2050. *Industrial Health*, 51(1):56–67, 2013. doi: 10.2486/indhealth.2012-0174. URL http://dx.doi.org/10.2486/indhealth.2012-0174.

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.

Tung Liu and Kui-Wai Li. Analyzing Chinas productivity growth: Evidence from manufacturing industries. *Economic Systems*, 36(4):531–551, Dec 2012. doi: 10.1016/j.ecosys.2012.03.003. URL http://dx.doi.org/10.1016/j.ecosys.2012.03.003.

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.

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.

Shin-ichi Tanabe, Naoe Nishihara, and Masaoki Haneda. Indoor Temperature, Productivity, and Fatigue in Office Tasks. *HVAC&R Research*, 13 (4):623–633, July 2007. doi: 10.1080/10789669.2007.10390975. URL http://dx.doi.org/10.1080/10789669.2007.10390975.

Mei Wen. Relocation and agglomeration of chinese industry. *Journal of Development Economics*, 73(1):329–347, Feb 2004. doi: 10.1016/j.jdeveco.2003.04.001. URL http://dx.doi.org/10.1016/j.jdeveco.2003.04.001.

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

Joshua Graff Zivin and Matthew Neidell. Temperature and the Allocation of Time: Implications for Climate Change. Technical report, National Bureau of Economic Research, February 2010. URL http://dx.doi.org/10.3386/w15717.

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.