

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

E. Somanathan

Indian Statistical Institute

Rohini Somanathan

Delhi School of Economics

Anant Sudarshan

University of Chicago

Meenu Tewari

University of North Carolina at Chapel Hill

Hotter years are associated with lower economic output in developing countries. We show that the effect of temperature on labor is an important part of the explanation. Using microdata from selected firms in India, we estimate reduced worker productivity and increased absenteeism on hot days. Climate control significantly mitigates productivity losses. In a national panel of Indian factories, annual plant output falls by about 2%

This work was made possible by funding from the Rockefeller Foundation, the Indian Council for Research on International Economic Relations, and the Indian Statistical Institute. We are grateful to Mehul Patel for field assistance and to Sheekha Verma for research assistance. This paper has benefited from helpful comments from four anonymous referees, Michael Greenstone, and participants at several conferences and seminars. Data are provided as supplementary material online.

Electronically published April 30, 2021

[*Journal of Political Economy*, 2021, vol. 129, no. 6]

© 2021 by The University of Chicago. This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0), which permits non-commercial reuse of the work with attribution. For commercial use, contact journalpermissions@press.uchicago.edu. 0022-3808/2021/12906-0004\$10.00

per degree Celsius. This response appears to be driven by a reduction in the output elasticity of labor. Our estimates are large enough to explain previously observed output losses in cross-country panels.

I. Introduction

Recent research has uncovered a systematic negative correlation between temperature and aggregate national output, especially in tropical developing countries (Dell, Jones, and Olken 2012; Burke, Hsiang, and Miguel 2015). High temperatures are associated with reduced crop yields as well as lower output in nonagricultural sectors.¹ Explanations for this relationship include heat stress on workers and temperature-related increases in mortality, conflict, and natural disasters.² Establishing and quantifying the relative importance of these mechanisms is crucial for identifying possibilities of adapting to a hotter world.

In this paper, we focus on understanding and quantifying the role of heat stress in mediating the temperature-output relation. Our knowledge of human physiology suggests that workers should respond fairly quickly when made to work in uncomfortable temperatures. Heat impacts on labor can therefore be identified both in daily or weekly output and in data at higher levels of aggregation. This distinguishes heat stress from many alternative mechanisms. We use several microdata sets and a nationally representative panel of manufacturing plants to estimate the effects of high temperatures on labor. Although we focus on Indian manufacturing, since heat stress is a universal physiological mechanism, the implications of our results may extend to other sectors and countries.

There are two channels through which high temperatures might affect factory workers. They may produce less while at work and be absent more often. We assemble high-frequency data on workers in three different manufacturing settings—cloth weaving, garment sewing, and steel products—and separately identify these two effects. We find that the output of individual workers and worker teams declines on hot days as well as in weeks with more hot days. Absenteeism is increasing in both contemporaneous temperatures and temperatures experienced over the preceding week. Stronger effects are visible for paid leave, with a weaker temperature-absenteeism relationship for unpaid leave. Climate control in the workplace eliminates productivity declines but not absenteeism, presumably because workers remain exposed to high temperatures at home and outside.

¹ For evidence on yields, see Mendelsohn and Dinar (1999), Auffhammer, Ramanathan, and Vincent (2006), Schlenker and Roberts (2009), Lobell, Schlenker, and Costa-Roberts (2011), and Gupta, Somanathan, and Dey (2017).

² Hsiang (2010) discusses heat stress, Burke, and Miguel (2013) identify a temperature-conflict relationship, and Burgess et al. (2017) study effects on mortality.

To examine whether the temperature effects for workers in these firms are more generally reflected in India's factory sector, we use a 15-year nationally representative panel of manufacturing plants. We find that the value of plant output declines in years with more hot days. Annual output is predicted to fall by 2.1% if every day warms by 1°C. We use a Cobb-Douglas specification to show that temperature-induced reductions in the output elasticity of labor, rather than capital or other factors, drive this response. This is not surprising, given that industrial air-conditioning was rare in India even in 2012, the last year covered by our data. The demand for large commercial units was a small fraction of the demand in both China and the United States, in spite of India being the warmest of these three countries.

After presenting our main results, we consider some alternatives to the heat stress channel, including natural disasters, power outages, and conflict. For the years covered by our plant panel, we collect data on instances of flooding, power shortages, and workdays lost in all recorded industrial disputes. We find that these variables cannot account for the estimated effect of temperature on output. Other possible explanations for the negative effect of high temperatures on manufacturing plant output include temperature effects acting through input prices and via linkages with agriculture. However, we find no effect of temperatures on input prices after controlling for state and year fixed effects, so this cannot account for our results. Also, we find that output declines occur across manufacturing sectors, so agricultural linkages (which vary greatly across sectors) are unlikely to be an important part of the explanation.

Our final set of results are at a yet higher level of aggregation, the Indian district. Official data on gross domestic product (GDP) in the manufacturing sector is available for Indian districts for the period between 1998 and 2009. We use a panel of 438 districts with unchanged boundaries over this period to directly estimate the impact of a 1°C increase in temperature on district output. We estimate declines of 3% per degree Celsius. This is comparable to the plant response.

To situate these findings within the context of the country-level relationships that motivate this paper, it is helpful to compare the temperature-output relationship estimated at several different levels of aggregation. Putting together our results from worker, plant, and district data, we find that effect sizes in all three cases are similar. Strikingly, these effects are large enough to account for the country-level response to temperature observed in the literature. Although this does not imply that heat stress is the sole reason for country-level decreases in manufacturing-sector output during hot years, it does indicate that this may be a much more important mechanism than previously believed.

Notwithstanding the importance of these temperature effects, adaptation through climate control is limited. For example, the cloth-weaving

firms we study are labor intensive but do not use climate control. Given the costs of electricity, value added per worker may be too low to justify these investments. In the garment firms, value addition by workers is greater, and we see partial climate control. In our national plant panel, we find that temperature effects on output fall over time, perhaps the result of investments in adaptation.

If heat stress plays an important role in reducing output, then firms that do make costly climate-control investments should strategically allocate these resources toward tasks that are labor intensive and add significant value. We surveyed the management of 150 plants in the diamond-processing industry to test these hypotheses. We find that air-conditioning is selectively used in rooms with activities that are both labor intensive and critical in determining diamond quality.

The remainder of this paper is organized as follows. Section II summarizes the physiological evidence on heat stress. Section III describes our data sources. Our main results are in section IV. In section V, we compare effect sizes from our worker, plant, and district-level data and show that these are of similar magnitude and consistent with country-level estimates in the literature. Section VI examines the adoption of climate-control investments within firms. Section VII discusses alternative explanations and the robustness of our main results. Section VIII concludes.

II. Prior Literature

The science of how temperature affects human beings is straightforward. Heat generated while working must be dissipated to maintain body temperatures and avoid heat stress. If body temperatures cannot be maintained at a given activity level, it becomes necessary to reduce the intensity of work (ISO 1989; Kjellstrom, Holmer, and Lemke 2009). The efficiency of this process depends primarily on ambient temperature but is also influenced by humidity and wind speed (ISO 1989; Parsons 1993). Laboratory studies often use an adjusted measure of heat that accounts for these factors—the wet bulb temperature (WBT; Lemke and Kjellstrom 2012). Unfortunately, outside the lab, data on humidity are often unavailable. For this reason, and to enable comparisons with prior work, we use daily maximum temperatures as our measure of heat throughout this paper.³

There have been a number of studies in the physiology and engineering literature that find that high temperatures reduce labor productivity. Mackworth (1946) conducted an early artifactual field experiment with wireless telegraph operators and found that they made more mistakes at high temperatures. Parsons (1993) and Seppanen, Fisk, and Faulkner (2003)

³ Section A1.3 in the appendix provides estimates using WBT for our factory sites.

summarize important findings in this area. Hsiang (2010) presents a meta analysis of recent laboratory evidence that shows that once WBTs rise above 25°C, task efficiency appears to fall by approximately 1%–2% per degree. A WBT of 25°C at 65% relative humidity is roughly equivalent to a temperature of 31°C in dry conditions.⁴ These temperatures are not considered unsafe from the point of view of occupational safety and commonly occur in many countries.⁵

Controlled experiments in the laboratory or workplace provide a useful benchmark but do not fully capture real manufacturing environments. Workers and management generally operate well within physical limits and have room to increase effort in response to incentives. The output-temperature relationship therefore depends on the physical as well as behavioral aspects of employment, such as the wage contract, particularities of production, management techniques, and mechanization. This makes data from nonexperimental settings particularly valuable. As early as 1915, Huntington exploited daily variations in temperatures experienced by workers and students performing various tasks and found that high temperatures appeared to reduce output (Huntington 1915).⁶ More recently, Adhvaryu, Kala, and Nyshadham (2020) exploit variation in workplace temperatures induced by low-heat LED lighting and conclude that worker productivity increases when temperatures are reduced.

Workplace productivity aside, high temperatures may also reduce our willingness and ability to even be present at work. Much less prior evidence exists on absenteeism, although Zivin and Neidell (2014) find that people in the United States allocate less time to work in exposed industries when temperatures are very high.

III. Data Sources

Our labor and output data are at three levels of aggregation: the worker or worker team, the plant, and the district. For each data set, we match output to measures of temperature. We also conduct a survey of diamond firms to study the selective use of climate control. Official data in India are typically available for financial years, which run from April 1 through March 31. When referring to a financial year, we use the initial calendar year. Our data sets are described below and summarized in table 1.

⁴ The WBT scale is compressed relative to temperature, so a 1°C change in WBT corresponds to a higher than 1°C change in temperature

⁵ Temperature exposure in sectors such as mining can be high enough to create serious health hazards. These settings have long been used for research on heat stress and occupational safety (Wyndham 1969).

⁶ We are grateful to an anonymous reviewer for pointing us to some of this literature.

TABLE 1
SUMMARY OF WORKER AND FIRM DATA SETS

Source	Location	Unit (N)	Dependent Variables	Time	Climate Control
Cloth-weaving firms	Surat	Worker (147)	Meters of cloth, worker attendance	365 days	No
Garment-sewing plants	NCR, Hyderabad, Chhindwara	Sewing line (103)	Operations completed	730 days (varies by line)	Partial (74 lines)
Garment-sewing plants	NCR, Hyderabad, Chhindwara	Sewing line (266)	Absences	730 days (varies by line)	Partial (224 lines)
Steel mill	Bhilai	Shift team (9)	Blooms rolled, team absences	3,337 days (production), 857 days (attendance)	Yes
Association of Diamond Firms	Surat	Plants $\times$ operations (150 $\times$ 5)	Air-conditioning indicator	Cross-section	Partial
Annual Survey of Industry	National	Plant (58,377)	Value of output	15 years	NA
Planning Commission of India	National	District (438)	Manufacturing GDP	12 years	NA

A. *Worker Data*

We collected worker output and attendance data from selected firms in three industries: cloth weaving, garment sewing, and the production of large infrastructural steel products. Figure A.1 (figs. A.1–A.8 are available online) includes photographs of production lines in each of these industries. Our three cloth-weaving factories are all located in the industrial city of Surat in the state of Gujarat, in western India. Our garment factories are managed by a single firm, with six plants located in the National Capital Region (NCR) in North India and two others in the cities of Hyderabad and Chhindwara in south and central India. Our steel-production data are from the rail and structural mill of a large public-sector steel plant in the town of Bhilai in central India. Each of these sites is part of an important manufacturing sector in the Indian and global economies. The textile sector (which includes spinning, weaving, and dyeing) employs about 12% of factory workers in India. The garment sector employs about 7% of factory workers, and the Bhilai steel mill is the largest producer of steel rails in the world.⁷

For the three cloth-weaving factories, we gathered daily data on meters of cloth woven and attendance of 147 workers employed during the financial year starting April 2012. A worker in each of these factories operates about six mechanized looms producing woven cloth. Workers are engaged in monitoring looms, adjusting alignment, restarting feeds when interrupted, and making other necessary corrections. The cloth produced is sold in wholesale markets or to dyeing and printing firms. Workers are paid based on the meters of cloth woven by these looms, and no payments are made for days absent. Protection from heat is limited to the use of windows and some fans. We obtained payment slips for each day and digitized these to generate a worker-level data set of daily output and attendance. For most types of cloth, workers were paid Rs 2 per meter.

For garment sewing, we have production data from eight factories owned by a single firm producing garments for foreign apparel brands. Unlike in the cloth-weaving firms described above, these workers are paid monthly wages that do not directly penalize workers for small variations in productivity or occasional absences. In each plant, production is organized in sewing lines of 10–20 workers, with each line creating part or all of a clothing item. Lines are usually stable in their composition of workers, while the garment manufactured by a given line changes based on production orders. Our productivity measure relates to the entire sewing line. The garment sector is highly competitive, and firms track worker output in sophisticated ways. In our case, the firm used an hourly production target

⁷ For employment shares, see ASI 2009–10, volume 1. A description of the steel plant at Bhilai is available from the Steel Authority of India. The steel rails from Bhilai are used for the entire network of public railroads in the country.

for each line, based on the time taken to complete the desired garment by an experienced line of “master craftsmen.” The actual hourly output, when controlled for the target, provides a measure of the line productivity. The target is not revised each day so it is not sensitive to daily temperatures. The firm management provided us with daily production from 103 sewing lines over a period of 730 days during the calendar years 2012 and 2013. They also gave us attendance records over the same period, allowing us to construct a daily count of absences within sewing lines in their factories.⁸

These garment factories also provide us an opportunity to study the effects of climate-control investments on productivity. During the period for which we have data, the firm was in the process of installing cooling equipment on its shopfloors. This installation of climate control had been completed in five of the manufacturing units in the NCR before 2012, but the sixth unit did not get this until 2014. Of the 103 sewing lines, 84 lines were located in the NCR, of which 74 had climate control. Two factories in Hyderabad and Chhindwara (19 sewing lines) were also without climate control, but average temperatures in these areas are lower than in the NCR. This phased rollout allows us to compare temperature effects in colocated factories with and without climate control.

The rail and structural mill in Bhilai is the primary supplier of rails to the Indian Railways and produces steel products used for large infrastructural projects. Rectangular blocks of steel called “blooms” form the basic input for all these products. They enter a furnace and are then shaped into rails, or “structurals,” to meet ordered specifications.⁹ When a bloom is successfully shaped, it is said to have been “rolled.” The number of blooms rolled in an 8-hour shift is our measure of output.

There are three shifts on most days, starting at 6 a.m., and workers are assigned to one of three teams that rotate across these shifts. The median number of workers on the factory floor is 66. Our production data record the team and the number of blooms rolled for each working shift during the 1999–2008 period. We observe a total of 9,172 shifts over 3,337 working days. In addition to the team output in each shift, we also have team-level absences over a shorter period of 857 working days between February 2000 and March 2003.¹⁰

⁸ Not all sewing lines are operational for all days during these 2 years. The number of observations over the time span of 730 days, therefore, varies by sewing line. Our attendance data cover more workers than our output data, e.g., employees engaged in cloth cutting but not sewing activities in the same factories. Since output data do not identify individual workers and lines are labeled differently in the two data sets, we separately analyze productivity and absenteeism and do not investigate interactions.

⁹ “Structurals” refer to a miscellaneous set of steel products used mostly in construction projects such as roads and bridges.

¹⁰ These data were first used by Das et al. (2013), who provide a detailed account of the production process in the mill.

Unlike the weaving and garment units, the production of rails is highly mechanized, and the mill runs continuously with breaks only for repair, maintenance, and adjustment for different products. Workers who manipulate the machinery used to shape rails sit in air-conditioned cabins. Others perform operations on the factory floor. This is the most capital intensive of our case-study sites with both automation and climate control.

B. Panel of Manufacturing Plants

We purchased secondary data from the Annual Survey of Industries (ASI) covering the financial years 1998–99 to 2012–13. The ASI is a Government of India census of large plants and a random sample of about one-fifth of smaller plants registered under the Indian Factories Act. Large plants are defined as those employing more than 100 workers.¹¹ The ASI provides annual data on output, the value of fixed assets, debt, cash on hand, inventories, input expenditures, and the employment of workers and management. The format is similar to census data on manufacturing in many other countries.¹²

The ASI provides plant identifiers for the 2000–2010 period but not in other years. To create a longer panel requires matching observations across different years using time-invariant plant characteristics. Following a procedure similar to that of Allcott, Collard-Wexler, and O’Connell (2016), we create an unbalanced panel of 58,377 plants from 1998 to 2012.¹³ We match plants to temperature and rainfall at the level of the district.¹⁴

C. District Panel of Manufacturing GDP

The Planning Commission of India has published data on district-level manufacturing-sector GDP over a 12-year period from 1998 to 2009. These figures include ASI plants as well as estimates from unregistered manufacturing and smaller factories not covered by the ASI. We use these statistics to directly estimate the effect of temperature on economic output, aggregated at the level of districts. Unfortunately, after 2009, this information has not been systematically compiled. Data for some districts were either not available in this data set or not reliable because of changes in boundaries over this period. Kumar and Somanathan (2009) provide a review of these boundary modifications. Therefore, our estimates are based on a

¹¹ For regions with very little manufacturing, the ASI covers all plants, irrespective of their size.

¹² Berman, Somanathan, and Tan (2005) discuss the measurement of variables in the ASI and its comparability with manufacturing data in other countries.

¹³ Appendix sec. A1.4 provides details on panel construction.

¹⁴ There are 529 districts with at least one plant in the data set. Figure A.4 shows the geographic distribution of ASI plants and locations of our microdata sites.

subsample of 438 districts with static boundaries and at least two nonmissing observations over this period.

D. Weather Data

Our weather data come from two sources. We use recordings from public weather stations within the cities where our cloth-weaving and garment-sewing factories are located. We also use a $1^\circ \times 1^\circ$ gridded data product sold by the India Meteorological Department (IMD), which provides daily historical temperature and rainfall measurements interpolated over the IMD's network of monitoring stations across the country. The first of these provides a more precise measure for locations near a weather station. The second is best suited to averaging over larger areas.¹⁵

In the case of our worker data, we know the precise factory locations and can use data from nearby public weather stations wherever available. We characterize the temperature of a day using the daily maximum temperature, which occurs during working hours and is therefore a useful proxy for heat exposure at the workplace. There were no public weather stations in the proximity of the Bhilai Steel Plant over the period for which we have data. For this plant, we instead rely on the IMD gridded data set and use an inverse distance weighted average of grid points within 50 km of the plant to assign daily maximum temperature values.

For our annual panel of manufacturing plants, we use daily maximum temperatures from the IMD gridded data sets as well as daily precipitation. Since we do not have precise location coordinates from the ASI, we assign to each plant the temperature and rainfall corresponding to the district in which it is situated. These numbers are obtained by spatially averaging grid temperatures over the geographical boundaries of each district. Additional details are in appendix section A1.4 (appendix is available online).

When using the ASI data, in our main specification, we aggregate daily temperatures up to the annual level using counts of the number of days in the year falling within different temperature bins. We use temperature bins defined as $\{(0,20], (20,25], (25,30], (30,35], (35,50]\}$. To summarize the temperature distribution over the year, we construct a vector $\mathbf{T} = (T^1, T^2, T^3, T^4, T^5)$, with counts of the number of days in each of these bins. This is calculated for every district and each year. Taken together, these bins are nonoverlapping and span the observed range of temperatures in the data, so that any given day is assigned to exactly one bin. We also estimate additional specifications using alternative functions of daily

¹⁵ The physiology literature often uses WBTs to study heat stress. This measure combines temperature and humidity. We are not aware of a good source of time-varying measures of WBT for the whole country. For this reason, and to ease comparison with previous work, we use maximum temperatures throughout the main paper.

maximum temperatures over the year, including a degree-day measure. These are described in section IV.B.

When using worker-level data, we also use similar binned specifications. The cutoffs and widths of these bins vary, reflecting differences in the distribution of weather in different sites. Bin definitions for workers are discussed in section IV.A and shown in figure 1.

E. Climate Control within Diamond Firms

In August 2014, we surveyed 150 diamond-cutting plants, randomly sampled from more than 500 units formally registered with the industry association of the city of Surat (the same location as our cloth-weaving units). Each plant carries out five operations: (i) sorting and grading, (ii) planning and marking, (iii) bruting (rounding a diamond), (iv) cutting, and (v) polishing. Although these factories are small and labor intensive like the cloth-weaving plants, the value added in production is much greater, and these units commonly deploy air-conditioning in at least some parts of the plant.

We asked the management of each firm about the number of workers and machines and the use of air-conditioning in each of the five operations. They were also asked to rate, on a scale of 1–5, the importance of each of these processes to the quality of final output. We use these responses to study the selective deployment of climate control.

IV. Results

A. Temperature Effects on Worker Output

Temperature can influence worker output through different channels. People may be more likely to miss work on very hot days. They may also be less productive at the workplace because of heat stress. Both contemporaneous and lagged temperatures potentially matter.

We begin by estimating the effects of temperature on the output of workers at the weekly level. These estimates reflect the combined effects of absenteeism and reduced productivity at work. We then use daily data to separately examine the nonlinear effects of contemporaneous and lagged temperatures on productivity and attendance.

Output is related to temperature using the following binned specification:

$$y_{iw} = \alpha_i + \gamma_M + \gamma_t + \sum_{j=2}^J \beta_j T_{iw}^j + \theta R_{iw} + \lambda X_{iw} + \epsilon_{iw}. \quad (1)$$

Our output measure is in physical units in each of the three types of firms that we study. For cloth weaving, y_{iw} is the inverse hyperbolic sine transformation of the daily meters of cloth produced by worker i averaged over the course of week w . If a worker is absent, we set output for that day




FIG. 1.—Effect of temperature on worker output. Estimates are percentage changes in daily output (averaged over a week) for a day in the week moving to a hotter temperature bin from the coolest (omitted) bin. Shaded areas represent 90% confidence intervals using robust standard errors clustered at the worker level. The number of bins varies across locations, reflecting differences in observed temperatures. *A*, Garment sewing lines in the National Capital Region (NCR). *B*, Garment sewing lines in Hyderabad and Chhindwara. *C*, Cloth weaving in Surat. *D*, Steel mill in Bhilai. The output variable for garment plants (*A*, *B*) is defined as the logarithm of the efficiency measure of each sewing line. In panel *C*, the output variable is defined as the inverse hyperbolic sine transformation of the meters of cloth woven by a worker. In panel *D*, the output variable is the logarithm of blooms rolled by a team of workers.

at zero. We use this transformation instead of logarithms since our output indicator can take zero values. For the steel mill, y_{iaw} is the logarithm of the average number of rectangular blooms rolled in shift i during week w . As described in section III, a bloom is an intermediate steel product that

is used in the manufacture of railway tracks. There are three shifts in the workday, each manned by a different worker team. For garment plants, y_{iw} is the logarithm of the efficiency of each sewing line (a team of workers). “Efficiency” is a performance metric used by the garment firm based on the number of operations completed every hour by the sewing line. We also control for a line-specific target efficiency that is set by the firm, as described in section III. We do this because the lines carry out operations of varying complexity over time, and the target helps to control for this. Note that the target itself is not updated daily and is therefore independent of temperature.

We include a range of fixed effects to control for idiosyncratic worker productivity and temporal and seasonal shocks. Fixed effects for the i th unit are denoted by α_i . A unit is an individual worker in the cloth-weaving firms, a sewing line in garment firms, and a team shift for the steel mill. As mentioned in section III, for the steel mill, there are three shifts a day, and three teams of workers rotating across shifts, producing a total of nine indicator variables.

Output is likely to respond to (possibly seasonal) demand, so we also include month and year fixed effects (γ_M, γ_Y). We use R_{iw} to indicate the weekly average of daily rainfall and X_{iw} other controls including the number of working days in the week and, for garment workers, the target efficiency. The variable T^j is a count of the number of days in the reference week that fall in a given temperature bin j . We use the following temperature bins: (0,19], (19,21], (21,23], (23,25], (25,27], (27,29], (29,31], (31,33], (33,35], (35,50]. Taken together, these capture the nonlinear relationship between output and temperature.

The temperature range we observe for each unit depends on its location. For units in the NCR around Delhi, we use all 10 temperature bins. For each of the other factory locations, we combine some of the lower temperature bins because observed temperatures span a smaller range. To facilitate comparisons, the highest bin is pegged at maximum temperatures above 35°C. The cloth-weaving workers in Surat face warmer temperatures, so our first bin ends at 29°C. This produces five bins: (0,29], (29,31], (31,33], (33,35], and (35,50]. For the steel plant and the garment-sewing lines outside the NCR, our first bin ends at 27°C. Because the sum of all bin counts is a constant, we omit the lowest bin in our regressions. The estimate of the coefficient of T^j should be interpreted as the effect of a single day in the week moving from the lowest (coldest) temperature bin, T^1 , to a warmer temperature range corresponding to bin j .

Figure 1 presents coefficient estimates β_j for all worker sites, with 90% confidence intervals. In the absence of climate control, output falls in weeks with more hot days.¹⁶ In climate-controlled garment plants in the

¹⁶ Large shop floors are not cooled by typical air-conditioning units. Thus, when we refer to climate control, we mean a plant that has a centralized cooling system such as an air washer installed.

NCR (fig. 1A), we see no negative effects of temperature on output. For the steel mill, which is largely automated and has climate control, if anything, output rises slightly at higher temperatures (fig. 1D). This might occur if climate control is turned on only on hot days, making workplace conditions on those days actually more comfortable. It is also possible that foundry operations are negatively affected by cold weather because metal may set too quickly, causing faults in the final output (Fiorese et al. 2015). We return to the question of interactions of capital equipment with temperature in section IV.B.¹⁷

Our estimates are heterogeneous across workplace settings. For garment plants in the NCR without climate control, the effect of an additional day in a week moving from the lowest to highest temperature bin is to reduce average daily efficiency by as much as 8%. The estimate for the garment plants in Hyderabad and Chhindwara is about half of this. For weaving workers, it is as low as 2%. These differences are not surprising because the omitted bin is not the same across sites—in warmer regions, the omitted bin spans higher temperatures than for sites in the NCR. That said, workplaces vary along many other dimensions such as worker health and income, the nature of physical or cognitive tasks they perform, differences in the output measure, financial incentives, and the nature of employment contracts. These factors may lead to heterogeneous effects of heat, even if the observed temperature ranges are the same.

For worker sites, we are also able to obtain data on temperature and humidity and can estimate WBTs that are commonly used in the physiology literature to measure heat stress. In figure A.3 (figs. A.1–A.8 are in the appendix) we replicate the results in figure 1 using bins in WBT instead of maximum temperatures. We find the same patterns of output response as we do when using maximum temperatures to proxy for heat. If anything, standard errors are smaller and effect sizes slightly larger.

Lagged effects on output and absenteeism.—To examine the effect of contemporaneous and lagged temperatures on workplace productivity and absenteeism, we turn to our disaggregated daily data. Exposure to very hot days may generate fatigue and illness, lowering output and increasing absenteeism. Strokes, fatigue, and even cases of organ damage have been directly linked to heat stress, and continued exposure may increase overall vulnerability (Kovats and Hajat 2008). Other illnesses may be influenced by sustained warm weather through different mechanisms, for example, the increased breeding of pathogens and disease vectors.

¹⁷ High temperatures could directly reduce productivity if they are associated with power outages. All the factories in our data set have a power backup, so this is not a concern. Also, if outages were driving our results, we should expect to see this effect in plants with and without climate control.

We modify (1) to include lagged temperature bins. Here, L_{id}^j is a count of the number of days falling in bin j in the six days preceding day d . Our output and other variables are as before, except now at the daily rather than weekly level. In the case of weaving workers, we include only those present at work on day d . We estimate

$$y_{id} = \alpha_i + \gamma_M + \gamma_t + \sum_j \beta_j T_{id}^j + \sum_j \omega_j L_{id}^j + \theta R_{id} + \lambda X_{id} + \epsilon_{id}. \quad (2)$$

We now use T^j as an indicator for the day falling in temperature bin j ; R_{id} is daily rainfall; and X_{id} now includes a fixed effect for the day of the week, and as before, for sewing lines, it also includes the target efficiency for the line. Our estimates from weekly data in figure 1 suggest that most of the temperature effects occur in the two highest bins. We focus on these temperatures by aggregating over cooler bins. Therefore, there are a total of three bins in both T and L .¹⁸

Our results are in table 2. Declines in daily output on hotter days are seen only in sites without climate control.¹⁹ Lagged temperatures reduce output for some sites. The clearest effects are found for weaving workers, where an additional day above 35°C in the six preceding days causes a 2.7% decrease in contemporaneous daily output. Notice that lagged temperatures seem to matter even in climate-controlled garment plants. This may reflect exposure outside the workplace. This is related to our findings on absenteeism, which we turn to next.

We have a daily indicator of absenteeism for our cloth-weaving workers. In the case of garment and steel plants, we have daily counts of the number of absences in the worker team. Using these measures of absenteeism as the dependent variable, we estimate (2). From table 3, we see absenteeism effects in settings with and without climate control. Lagged high temperatures increase the likelihood of missed work in climate-controlled garment factories, the steel plant, and the weaving plants. For garment plants with no climate control, our coefficients are imprecisely estimated.

The garment workers in our sample provide us with some insight into how workers respond to incentives. These workers are allocated a certain amount of paid leave, and our data distinguish paid and unpaid absences for each worker. In climate-controlled garment plants in the NCR (cols. 1 and 2), we find that the number of paid absences increases with both contemporaneous and lagged temperatures but that the probability of unpaid leave does not change with temperature. This suggests that monetary

¹⁸ Including lagged variables for all temperature bins increases the number of coefficients being estimated and reduces the precision of our estimates.

¹⁹ As before, we see positive effects on output in the case of climate-controlled sites. Standard errors are high for the garment plants in central and south India, and we are unable to draw clear conclusions.

TABLE 2
EFFECT OF HOT DAYS ON WORKER OUTPUT

	CLIMATE CONTROL		NO CLIMATE CONTROL		
	Garments (Log Efficiency)	Steel (Log Blooms Rolled)	Weaving (IHS Meters)	Garments (Log Efficiency)	
	(1)	(2)	(3)	(4)	(5)
<i>T</i> (33°–35°C)	.025** (.010)	.028* (.017)	−.040** (.019)	−.129*** (.042)	−.007 (.037)
<i>T</i> (above 35°C)	.035*** (.014)	.020** (.009)	.011 (.022)	−.154*** (.041)	.008 (.046)
<i>L</i> (33°–35°C)	−.004 (.005)	.005 (.004)	−.033*** (.011)	−.009 (.012)	.004 (.010)
<i>L</i> (above 35°C)	−.011** (.005)	−.002 (.005)	−.027*** (.009)	−.019 (.027)	.015 (.018)
Climate control	Yes	Yes	No	No	No
Number of units	74 lines	9 teams	147 workers	10 lines	19 lines
Time span (days)	730	3,337	365	730	730

NOTE.—Robust standard errors are clustered at the worker level. *T* is an indicator for a day falling in the specified temperature bin. *L* is a count of the number of days falling in the specified temperature bins in the six preceding days. Models include unit-level fixed effects (individuals for weaving and teams for garments and steel) and fixed effects for the month, year, and day of the week. Columns 1 and 2 have estimates from climate-controlled garment plants in the National Capital Region (NCR) and the steel mill in Bhilai. Columns 3–5 are for settings without climate control: weaving workers in Surat and garment-sewing lines in the NCR and south and central India. Output for weaving workers is an inverse hyperbolic sine (IHS) transformation of meters of cloth woven. The output variable for garment workers is the logarithm of the efficiency measure. Rainfall is included as a control, but estimates are not presented.

* $p < .01$.
** $p < .05$.
*** $p < .10$.

disincentives could weaken the temperature-absenteeism link.²⁰ For non-climate-controlled garment plants (cols. 5 and 6), our point estimates are too noisy to draw any conclusions.

Absenteeism driven by contemporaneous high temperatures may be partially due to time-allocation decisions and labor-leisure trade-offs (Zivin and Neidell 2014). Lagged effects may also reflect the effects of morbidity. Although workplace climate control may reduce the effects of temperature on worker productivity on the shop floor, it may not remove negative output effects caused by absenteeism. Absenteeism might also result in costs we do not measure, such as firms hiring redundant workers. The presence

²⁰ We focus here on daily absenteeism. The incentives generated by employment contracts may affect other types of absences and the duration of employment. Section A1.2 in the appendix provides data on monthly absences for these two types of workers. Those without paid leave are much more likely to leave during summer months. This is also borne out by interviews with factory owners in the city of Surat, where our cloth-weaving plants are located.

TABLE 3
EFFECT OF HOT DAYS ON WORKER ABSENTEEISM

	CLIMATE CONTROL			NO CLIMATE CONTROL		
	Garments		Steel (All) (3)	Weaving (All) (4)	Garments	
	Paid (1)	Unpaid (2)			Paid (5)	Unpaid (6)
<i>T</i> (33°–35°C)	.082*** (.022)	–.083 (.065)	–.011 (.048)	.003 (.004)	–.001 (.128)	.796 (.678)
<i>T</i> (above 35°C)	.115*** (.027)	.031 (.049)	.051 (.068)	–.004 (.004)	–.034 (.117)	1.001 (.862)
<i>L</i> (33°–35°C)	–.018 (.011)	–.047 (.032)	.044*** (.014)	.006*** (.002)	.017 (.077)	.772 (.686)
<i>L</i> (above 35°C)	.021** (.010)	–.001 (.022)	.045** (.020)	.005*** (.002)	.078 (.083)	.567 (.426)
Number of units	224 lines		9 teams	147 workers	42 lines	
Time span (days)	730		3,337	365	730	

NOTE.—Robust standard errors are clustered at the worker level. *T* is an indicator for a day falling in the specified temperature bin. *L* is a count of the number of days falling in the specified temperature bins in the six preceding days. Models include unit-level fixed effects (individuals for weaving and teams for garments and steel) and fixed effects for the month, year, and day of the week. Columns 1 and 2 present estimates of the effect of temperature on the number of paid and unpaid leaves for sewing lines in climate-controlled garment plants, col. 3 reports coefficients for absences in climate-controlled steelworker teams, col. 4 reports the probability of a weaving worker being absent, and cols. 5 and 6 give estimates of temperature effects on paid and unpaid leaves for sewing lines in non-climate-controlled garment plants.

** $p < .05$.
*** $p < .10$.

of redundant labor has been documented for the steel plant we study (Parry 1999), and this might explain why we do not see output effects in climate-controlled plants in spite of increased absenteeism.

For the garment and steel plants, there is no straightforward way to translate increased absenteeism within worker teams into impacts on output. For weaving workers, an additional day above 35°C in the six preceding days causes a 0.005 increase in the probability of missing work. The mean worker output, on a day when the worker is present, is 134.3 meters of cloth. Since absenteeism takes output to zero, this is equivalent to a reduction of 0.7 meters. Weaving workers come to work intermittently, so their average daily output, net of absences, is about 51 meters of cloth per day. An additional hot day in the six preceding days, therefore, reduces output by about 1.4% through the absenteeism channel. This can be compared with a loss of 2.7% via the on-the-job productivity channel (table 2).

B. *Temperature Effects on Plant Output*

1. *Main Results*

Thus far, we have used high-frequency data to show that worker productivity declines on hot days. We now turn to our nationwide panel of manufacturing

plants to examine whether there are similar temperature effects on the value of plant output and, if so, whether they might be attributable to a decline in the productivity of labor.

We estimate a model analogous to (1):

$$y_{it} = \alpha_i + \gamma_t + \sum_{j=2}^5 \beta_j T_{it}^j + \theta R_{it} + \epsilon_{it}. \quad (3)$$

The dependent variable y is now the log of the value of annual plant output. Plant and year fixed effects are denoted by α_i and γ_t , respectively. For every plant i and year t , T_{it}^j is the number of days in the year with maximum temperature falling in bin j . We have five temperature bins: $\{(0,20], (20,25], (25,30], (30,35], (35,50]\}$. Variable R_{it} is the annual average of daily rainfall in the district containing plant i in year t .²¹ We use wider bins here than with our worker data to preserve precision. We have a shorter panel with only 15 years of data, as opposed to the worker data, where our shortest weekly panel is 52 weeks and our shortest daily panel covers 365 days. The topmost bin for both worker and plant models is identical.

Our coefficient estimates β_j are plotted in figure 2 and indicate an inverse relationship between temperature and annual plant output, akin to the relationship we see between temperature and worker productivity.²² Each β_j is the percentage change in annual plant output from a single day in the year moving from the coldest bin to bin j . Shaded areas represent 90% confidence intervals, with standard errors corrected for serial and spatial correlation following Conley (2010).²³ A day moving from the lowest to the highest temperature bin reduces annual output by 0.22%.

2. Alternative Specifications and Warming Scenarios

We examine the robustness of these results by running a set of related specifications. In each case, we predict the percentage change in the value of annual plant output for alternative warming scenarios. Our results are in table 4. The first four rows of columns 1–4 show the predicted percentage change in output when a single day in the year moves from 20°C to the specified temperature. The first column has the estimates of equation (3) already in figure 2. Column 2 adds state-specific quadratic time trends. Column 3 controls for floods and industrial conflicts, while column 4 controls for power outages. We discuss these three variables further in

²¹ Since temperature and rainfall data are available at the district level but not for individual plants, these variables have the same values for all plants in a district.

²² Two recent studies from China have similar findings (Zhang et al. 2018; Chen and Yang 2019).

²³ Conley errors are presented assuming a 150-km radius of spatial correlation.




FIG. 2.—Temperature effects on the value of annual plant output. Shown is the percentage change in the annual value of plant output resulting from the daily maximum temperature of a single day moving from below 20°C to the given temperature. Shaded areas represent 90% confidence intervals with standard errors corrected for serial and spatial correlation following Conley (2010). Data are from the Annual Survey of Industries for the years 1998–2012.

section VII. We see from table 4 that the additional controls in columns 2–4 do not substantially change the bin coefficients.

Columns 5–7 present models that do not use daily bin counts but depend on the distribution of daily temperature over the year in other ways. Column 5 presents a model that is piece-wise linear in degree days. The calculation of degree days is best explained with an example. A day with a temperature of 29°C contributes 20°C to the first bin (0–20], 5°C to the second bin (20–25], and 4°C to the third bin (25–30]. Thus, when a single day moves from 20°C to 25°C (the scenario in col. 5, row 1, of table 4) there is an increase of 5°C in the second degree-day bin and no change in other bins.

More formally, denote the endpoints of our five temperature bins by $(T^{1j}, T^{2j}]$, $j = 1, 2 \dots 5$. A daily temperature T contributes positive degree days to all those bins for which $T > T^{1j}$ and zero to all others. If $T \geq T^{2j}$, the day contributes $T^{2j} - T^{1j}$ to bin j . If $T^{1j} < T \leq T^{2j}$, it contributes $T - T^{1j}$ to bin j . As in (3), we now sum the degree days in each bin over the year to obtain D_{it}^j for each unit i and estimate the following model:

$$y_{it} = \alpha_i + \gamma_t + \sum_{j=2}^5 \beta_j D_{it}^j + \theta R_{it} + \epsilon_{it}. \quad (4)$$

The effects of moving a day from 20°C to 25°C, 30°C, 35°C, and 45°C in the degree-day model are shown in the first four rows of column 5. These predictions are similar to those from the binned specifications in the first four columns. Columns 6 and 7 provide results from models where logged output depends on polynomial functions of daily maximum temperature, summed over the year. Denoting by T_{dit} the maximum temperature for

TABLE 4
PREDICTED CHANGES IN PLANT OUTPUT UNDER DIFFERENT WARMING SCENARIOS

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
25°C	-.00139*** (.00043)	-.00093** (.00040)	-.00139*** (.00044)	-.00173*** (.00053)	-.00152*** (.00058)	-.00065*** (.00018)	-.00030*** (.00012)
30°C	-.00164*** (.00049)	-.00130*** (.00047)	-.00160*** (.00049)	-.00183*** (.00055)	-.00116** (.00049)	-.00108*** (.00030)	-.00061*** (.00024)
35°C	-.00189*** (.00051)	-.00159*** (.00050)	-.00184*** (.00051)	-.00225*** (.00050)	-.00175*** (.00056)	-.00127*** (.00039)	-.00091*** (.00036)
45°C	-.00217*** (.00056)	-.00183*** (.00057)	-.00214*** (.00057)	-.00238*** (.00066)	-.00133* (.00079)	-.00097* (.00065)	-.00152*** (.00061)
Mean shift by 1°C	-.02126*** (.00812)	-.02085*** (.00713)	-.02132*** (.00813)	-.02267*** (.00965)	-.01648* (.00940)	-.01581* (.00927)	-.02214*** (.00892)
Projected warming 2075–80	-.06558** (.02934)	-.0629** (.0260)	-.06715** (.0294)	-.06006* (.0345)	-.04489 (.04130)	-.05483 (.03815)	-.08856*** (.03570)
Plant fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State-specific trends	No	Yes	No	No	No	No	No
Conflict and flood controls	No	No	Yes	No	No	No	No
Outage controls	No	No	No	Yes	No	No	No

NOTE.—Columns correspond to different model specifications, and rows correspond to alternative warming scenarios. Column 1 provides estimates from our preferred specification in eq. (3). Column 2 adds state-specific quadratic time trends to the baseline model in col. 1. Column 3 adds controls for floods and conflict, while col. 4 adds a control for outages. Columns 5–7 use alternative representations of annual temperature distributions. Column 5 presents estimates from the degree-day model in eq. (4). Column 6 presents estimates from the model in eq. (5) where daily output depends on a quadratic polynomial of daily temperature, and col. 7 has output depend linearly on temperature. Rows 1–4 present the effect on log output of raising the temperature of a single 20°C day to 25°C, 30°C, 35°C, and 45°C, respectively (in cols. 1–4, these are just the bin coefficients). Row 5 predicts the effect of a 1°C increase in the temperature of every day in the year. Row 6 computes predicted output changes based on projections of long-term warming obtained from the RCP 8.5 scenario of the HadGEM2 climate model. Standard errors are corrected for serial and spatial correlation following Conley (2010). Data on the value of output and inputs are at the plant level from the Annual Survey of Industry.

* $p < .01$.

** $p < .05$.

*** $p < .10$.

plant i on day d of year t , column 6 has predictions based on the following model:

$$y_{it} = \alpha_i + \gamma_t + \sum_{d=1}^{365} \beta_1 T_{dit} + \sum_{d=1}^{365} \beta_2 T_{dit}^2 + \theta R_{it} + \epsilon_{it}. \quad (5)$$

Column 7 is based on a variant without the quadratic temperature terms, so output depends linearly on the sum of maximum daily temperatures over the year. The quadratic and linear models show smaller point estimates than the binned and degree-day specifications of columns 1–5, although the confidence intervals are overlapping.

The fifth and sixth rows of the table use our models to generate predictions from two alternative warming scenarios. We use the estimated coefficients from each of our models to compute the change in log output that would occur if the distribution of temperature changed from the one we actually observe in our data to a new warmer distribution. Row 5 shows predicted output changes when each day in the year is 1°C warmer, so that the annual average of the daily maximum temperature increases by 1°C. The estimated reduction in output ranges from 1.6% to 2.3% in the different models. Row 6 computes predicted output changes based on projections of long-term warming obtained from the RCP (Representative Concentration Pathway) 8.5 scenario of the HadGEM2 (Hadley Centre Global Environment Model version 2) climate model. For every day in the year, we compute the daily average of the 2075–80 projections and the 2005–10 projections. The difference between these two give us an estimate of the change in temperature we can expect by 2075–80 for each day of the year. We add this change in temperature to the baseline temperature distribution of average daily temperatures in our data.²⁴ Row 6 provides predictions for output changes under this warming scenario. These range between –4.5% and –8.9%.

To summarize, the inverse relationship between measures of temperature and plant output is seen across the many model specifications we consider. Results from these alternative models are broadly comparable, with some heterogeneity in effect sizes.²⁵

3. The Labor Channel

We now examine the extent to which the aggregate effects we have found in the factory panel can be explained by reductions in the productivity of labor as opposed to other factors. There was very limited deployment of climate control in the Indian factory sector during the period of our analysis. A study carried out by the Japan Refrigeration and Conditioning

²⁴ This baseline distribution averages over all plants in a year and all years in the data set so we work with a single temperature number for each day of the year.

²⁵ Table A.3 is similar to table 4 but with different bin cutoffs.

Industry Association reports that the demand for commercial scale air-conditioning units in India in 2013 was about 10% that of China and 3% that of the United States.²⁶ This, together with our results on declining labor productivity of workers in our microdata, suggests heat stress on labor may be an important explanation for the declines in the value of plant output we have presented above.

We explore this using a Cobb-Douglas production function in which the total factor productivity and the output elasticities of labor and capital are all allowed to depend on the temperature distribution as represented by the number of days in each of five temperature bins, $\mathbf{T} = (T^1, T^2, \dots, T^5)$. We assume that quantities of labor and capital within the factory are determined before the realization of $\mathbf{T}$ and so do not depend on it. While output elasticities equal input cost shares on average, they will not do so in any given year since temperature distributions are not predictable. Denoting logged values of output, capital and labor by y , k , and l , respectively, we have

$$y = \alpha(\mathbf{T}) + \omega(\mathbf{T})k + \beta(\mathbf{T})l. \quad (6)$$

We assume that total factor productivity α , output elasticity of labor β , and the output elasticity of capital ω are all linear in temperature bins indexed by j . Thus, we have

$$\alpha(\mathbf{T}) = \alpha_o + \sum_{j=2}^5 \alpha_j T^j,$$

$$\omega(\mathbf{T}) = \omega_o + \sum_{j=2}^5 \omega_j T^j,$$

and

$$\beta(\mathbf{T}) = \beta_o + \sum_{j=2}^5 \beta_j T^j.$$

Making these substitutions in (6), we obtain

$$y = \alpha_o + \sum_{j=2}^5 \alpha_j T^j + \omega_o \cdot k + \sum_{j=2}^5 \omega_j T^j k + \beta_o \cdot l + \sum_{j=2}^5 \beta_j T^j l. \quad (7)$$

We use the net value of equipment and machinery at the start of each year as our measure of capital and the number of full-time workers as

²⁶ The total sales of variable refrigerant flow air-conditioning systems, a common technology for larger commercial and industrial applications, numbered about 22,000 units in India compared to almost 600,000 in China (JRAIA 2019). Another technology used in industrial cooling, chiller systems, was even less popular, with about 4,000 units sold (USAID and BEE 2014). Low-cost technologies, such as industrial air coolers that use water rather than a refrigerant, were also uncommon. As recently as February 2019, in an interview published in the leading Indian newspaper, *Hindu BusinessLine*, the chief executive officer of India's largest manufacturer of air coolers characterized this market as "negligible," saying that "the industrial/commercial coolers segment doesn't exist in the country at present" (Vora 2019).

our measure of labor. We add controls for plant and year fixed effects as well as rainfall to (7) and estimate ω_j , β_j , and α_j .

Coefficient estimates from this model are in column 3 of table 5. Columns 1 and 2 show estimates from models that build up to this one by incrementally introducing labor and capital interactions with temperature to our base model in equation 3. We see that the temperature-labor interaction terms in column 3 are all negative and significant, while temperature effects on the output elasticity of capital are positive. Controlling for temperature interactions with labor and capital, the residual effect of temperature is also insignificant, as seen in the first four rows. These results suggest that it is temperature-induced declines in labor productivity that drive the negative effects of temperature on output.

One concern with estimating production functions of this type is potential endogeneity of labor (Levinsohn and Petrin 2003; Akerberg, Caves, and Frazer 2006). This may not be a significant concern in our setting, given India's notoriously inflexible labor market. In 2017, the World Bank ranked India as low as 130 on its global Ease of Doing Business index, citing rigid labor laws as a primary reason for the country's poor performance. Among several other weaknesses, the report draws attention to India's Industrial Disputes Act of 1947, which requires that firms with more than 100 employees obtain explicit government approval before dismissing workers. Since our measure of capital is the value of plant and machinery at the start of the year, this too is relatively inflexible and cannot be influenced by temperature shocks during the year.

Nevertheless, as a robustness check, we also estimate our production function using the Levinsohn-Petrin estimator, which allows for endogenous labor (Levinsohn and Petrin 2003). This approach assumes that labor is highly flexible and chosen by the firm in each period, after the realization of any shocks. Section A1.5 of the appendix describes the way in which we apply this method to our data, and column 4 of table 5 reports the relevant coefficient estimates. The point estimates for the labor-temperature interactions are smaller but remain negative and are statistically indistinguishable from those in column 3.²⁷

Finally, we investigate how temperature effects vary by labor and capital intensity. We measure labor intensity by the ratio of the total annual wage bill to total annual output for all plants in our sample. We measure capital intensity by the ratio of the value of capital to annual output. We classify plants into quartiles, Q^{lj} and Q^{kj} , based on their mean values of labor and capital intensity across all years and estimate the following model:

$$y_{it} = \alpha_i + \gamma_t + \beta_0 T_{it}^a + \sum_{j=2}^4 \beta_j^l T_{it}^a Q_i^{lj} + \sum_{j=2}^4 \beta_j^k T_{it}^a Q_i^{kj} + \theta R_{it} + \epsilon_{it}. \quad (8)$$

²⁷ Capital-temperature interactions and residual temperature effects are subsumed in a nonlinear control function and not separately estimated here. See app. sec. A1.5 for details.

TABLE 5
TEMPERATURE INTERACTIONS WITH FACTOR INPUTS

	(1)	(2)	(3)	(4)	(5)
T^a					.02324* (.01345)
T^2	.00256** (.00097)	.00008 (.00211)	-.00008 (.00162)		
T^3	.00147 (.00103)	-.00205 (.00229)	-.00009 (.00165)		
T^4	.00081 (.00108)	-.00094 (.00237)	-.00028 (.00170)		
T^5	.00003 (.00118)	-.00499* (.00259)	-.00171 (.00185)		
l	.8612*** (.0957)		.91426*** (.09660)	.36520*** (.05910)	
k		.20433** (.05674)	.06629 (.04114)		
$l \times T^2$	-.00098*** (.00027)		-.00134*** (.00034)	-.00056** (.00022)	
$l \times T^3$	-.00067** (.00027)		-.00104*** (.00027)	-.00038** (.00017)	
$l \times T^4$	-.00052* (.00027)		-.00077*** (.00027)	-.00030* (.00017)	
$l \times T^5$	-.00036 (.00029)		-.00075*** (.00029)	-.00039** (.00018)	
$k \times T^2$		-.00009 (.00015)	.00028* (.00015)		
$k \times T^3$		.00005 (.00016)	.00022* (.00012)		
$k \times T^4$		-.00003 (.00016)	.00016 (.00011)		
$k \times T^5$		.00022 (.00018)	.00024** (.00012)		
$T^a \times Q^{/2}$					-.04037*** (.01215)
$T^a \times Q^{/3}$					-.08313*** (.01312)
$T^a \times Q^{/4}$					-.13986*** (.01794)
$T^a \times Q^{k2}$					.04452*** (.01154)
$T^a \times Q^{k3}$					.03544*** (.01224)
$T^a \times Q^{k4}$					.00876*** (.0149)
Observations	179,107	179,107	179,107	176,620	179,107

NOTE.—Data are from the Annual Survey of Industry. Standard errors are corrected for serial and spatial correlation following Conley (2010). Models include plant and year fixed effects. Temperature bins are $\{(0,20], (20,25], (25,30], (30,35], (35,50]\}$. T^j is the number of days in the j th bin. T^j is the omitted bin. Columns 1 and 2 add interactions with labor and capital to our base model. Column 3 presents ordinary least squares estimates of the production function. Column 4 presents the first stage of a Levinsohn-Petrin estimate of the production function. Capital-temperature interactions and residual temperature effects are subsumed in a nonlinear control function and not separately reported. Column 5 interacts annual temperatures with quartiles of labor and capital intensities. Coefficients on rainfall and quartile dummies are omitted.

* $p < .01$.

** $p < .05$.

*** $p < .01$.

Column 5 of table 5 reports coefficients β_j^l and β_j^k from this model. The negative effects of the annual average of daily maximum temperature (T^a) are greatest in plants with high wage-share output ratios. On the other hand, capital intensity is positively associated with temperature. These models include plant fixed effects, so these results cannot simply be driven by plant size.²⁸

Taken together, the evidence in this section not only suggests that temperature negatively affects manufacturing output but also that this response operates through labor productivity.

V. Comparison with Macrolevel Estimates

In this section, we show that our estimated temperature effects at worker and plant levels are consistent with each other and with estimates based on district-level manufacturing output. We also compare our results with prior country-level studies. These comparisons suggest that temperature effects on labor are large enough to account for much of the country-level response of manufacturing GDP to temperature.

Prior studies have estimated the effect of a 1°C increase in annual temperature on country GDP. To compare our estimates with these, we must report our worker and plant results in similar terms. This requires specifying how the distribution of daily temperatures across the year changes when the average annual temperature increases by 1°C. There is of course, no unique way to map changes in temperature distribution to changes in annual average temperatures. We simply assume that every day in the year warms by 1°C. Under this assumption, the change in plant output for our primary specification is -2.1%, with a 90% confidence interval of ± 1.32 . This is plotted in bar 2 of figure 3 and is from row 5, column 1, of table 4.

Our worker-level estimates in figure 1 exhibit heterogeneity across sites, depending on the type of work and the degree of protection from heat. Noting that no single setting is representative of all workers, we estimate the effect of a 1°C uniform increase in the daily temperature distribution for garment workers in the NCR who are not working in cooled environments. We use this site because it has a wide temperature range that corresponds most closely to that observed in the nationally representative plant data. The estimated percentage reduction in output is 3 ± 1.35 (bar 1 of fig. 3).²⁹

²⁸ For parsimony, this model interacts only the average daily maximum temperature with quartile dummies. We obtain similar results using days in the highest temperature bin rather than average maximum temperature. We could also interact all temperature bins with quartile dummies, but this produces a large number of imprecisely estimated coefficients.

²⁹ Since we model the relationship between temperature and worker output using a “days in temperature bins” specification, we translate a 1°C increase in the daily temperature into corresponding changes in temperature bins in order to compute this effect.




FIG. 3.—Bars 1–3 provide the marginal effect of temperature on log output at different levels of production with 90% confidence intervals as estimated in this paper. “DJO” provides the contemporaneous effect of temperature on industrial sector growth rates in poor countries in a model with no lags from Dell, Jones, and Olken’s (2012) study. “BHM” provides the contemporaneous marginal effect of temperature on all-sector country output growth, at 30°C from a similar model with no lagged effects in Burke, Hsiang, and Miguel’s (2015) study.

If the output from manufacturing plants drops in hot years, we should see corresponding changes in manufacturing GDP at the subnational level. Using the district panel described in section III, we regress manufacturing GDP on average annual maximum temperature, T^a , controlling for rainfall as well as district and year fixed effects. The coefficient on T^a gives us the effect of a 1°C increase in temperature on district output. The estimated percentage reduction in manufacturing GDP is -3.5 ± 2.6 . This is shown in bar 3 of figure 3.³⁰

The last two bars of figure 3 depict estimates from two recent country-level studies: Dell, Jones, and Olken (2012) and Burke, Hsiang, and Miguel (2015). Both studies use annual average temperatures for many countries across the world, observed over long periods of time. The specifications in

³⁰ We favor using this district panel rather than the Reserve Bank of India GDP figures for Indian states because these data are interpolated in several years and therefore unreliable. In our district panel, we have missing data in some years but no imputed estimates.

these studies are not directly comparable with ours, but their results provide a useful benchmark. In figure 3, the fourth bar, labeled DJO, provides the contemporaneous effect of temperature on industrial sector growth rates in poor countries in a model with no lags (table 5 of Dell, Jones, and Olken 2012). The last bar, labeled BHM, provides the contemporaneous marginal effect of temperature on all-sector country output growth, at 30°C from a similar model with no lagged effects (table S2 of Burke, Hsiang, and Miguel 2015). It is interesting that temperature effects on the economy as a whole are similar in magnitude to those observed for manufacturing alone and, in turn, are similar to our estimates at lower levels of aggregation. Part of the explanation might be that changes in labor productivity affect all sectors of the economy.

This exercise does not, of course, imply that the negative effects of temperature on GDP found in these cross-country studies is occurring wholly or mainly through labor. However, both studies use data going back to the 1950s, covering long periods of time when climate control was uncommon in many parts of the world. Our estimates suggest that if the effect of temperature on labor productivity in the countries and sectors studied by these authors is of the same size as what we find in Indian manufacturing, then it would be enough to explain the entire temperature effect found there.

VI. Adaptation

The loss in output caused by high temperatures encourages adaptive responses by firms. In the short term, decisions to invest in climate control depend on the costs of cooling, relative to the expected output losses resulting from heat stress.³¹ Over longer time periods, firms may increase automation, relocate plants, or change the composition of output.

Firms may also selectively invest in climate control. If labor productivity plays an important role in output losses associated with hot days, we would expect that processes that are labor intensive and add high value would be preferentially protected. To study this, we conducted a survey of 150 diamond-cutting factories located in the same city of Surat as our cloth-weaving units. These are drawn randomly from all factories registered with the local diamond industry association.

These factories handle several distinct processes, some of which are largely mechanized (such as cutting stones), while others have much greater worker input (such as sorting uncut diamonds by quality). Our survey allowed us to study the selective adoption of air-conditioning within plants. We find that climate control is indeed more likely to be used for processes

³¹ In app. sec. A1.14, we carry out a back-of-the-envelope cost-benefit analysis of climate control for weaving plants and show that electricity costs of air-conditioning are high relative to output losses.

that are labor intensive and contribute most to diamond quality. We describe our data and results in appendix section A1.15.

In our national plant panel, we find that the effects of a 1°C rise in temperature seem to be falling over a 15-year period. We modify (3) to include a full set of interactions of temperature bin counts with a continuous time variable. The negative effect on output from an additional day in the fourth and fifth temperature bins reduces by about 6%–8% per year. column 1 of table A.1 (tables A.1–A.5 are available online) provides these coefficient estimates.³² As countries grow richer, it is possible that their manufacturing sector becomes less vulnerable to output losses associated with heat.

VII. Alternative Explanations

Reduced labor productivity is not the only way in which high temperatures may reduce output. Climatic changes may increase conflict (Hsiang, Burke, and Miguel 2013) or the frequency of natural disasters (Kahn 2005). Neither of these would influence our worker-level results because they occur on timescales that are much longer than a day. They could potentially mediate the temperature effects on output that we observe in our national panel of manufacturing plants. Other factors that may influence plant output, without necessarily changing the productivity of labor, include power outages, input price changes, and agricultural spillovers.

We test some of these explanations and find that they are unable to account for our results. We have already shown in table 4 that temperature effects on output remain almost unchanged when controlling for floods, conflicts, and power outages. In appendix section A1.8, we describe the construction of these variables and provide coefficient estimates associated with them when they are included in modified versions of (3).

To examine whether input prices change with temperature, we use data on the price of the input with the largest expenditure share, as reported in the ASI (table A.2). We find no evidence of temperature effects on input prices in our data. It may be that most changes in prices are captured by the year fixed effects in our models, and price shocks from local temperature fluctuations are neutralized by storage.

Finally, to examine the role of agricultural spillovers, we provide sector-wise estimates of temperature effects by estimating a model in which we include interactions of average annual maximum temperature with indicators for two-digit manufacturing sectors. We observe negative temperature effects across sectors, even for activities with no obvious connection to agriculture (fig. A.6).

³² Since climate control requires electricity, we also look for heterogeneity in the temperature response by the electricity intensity of output. We find that plants with above-median levels of electricity intensity respond more weakly to high temperatures (table A.4, col. 2).

VIII. Conclusions

This paper estimates the impact of temperature on manufacturing output. We use selected factory settings to separately study temperature effects on the daily productivity and attendance of workers. We show that, in the absence of climate control, worker productivity declines on hot days. For absenteeism, we find effects of contemporaneous and lagged temperatures even for workers in factories with climate control, suggesting that workplace adaptation alone is insufficient to mitigate all the effects of heat. In a 15-year national panel of manufacturing plants, we find that the effect of temperature on the value of annual plant output appears to be driven in large part by its effect on the output elasticity of labor.

Our estimates from both worker and annual plant data are comparable to those found in studies of country-level manufacturing GDP. This suggests that heat stress, through its effects on productivity, time allocation, and morbidity, is an important underlying cause for the declines in non-agricultural GDP at high temperatures.

The evidence we provide on the effectiveness of climate control and on its limited adoption has implications for how we should think about the costs of climate change going forward. Research into low-cost technologies to protect workers from ambient temperatures may have significant social value. In the long term, there are other ways in which the industrial sector might respond to high temperatures. These include increasing automation and shifting away from labor-intensive sectors in hot parts of the world. These adaptive responses may have significant distributional implications. If directed toward more productive workers, they will tend to increase wage inequality.

Although our focus throughout this paper has been on the manufacturing sector, the potential ramifications of our findings are wider. Our conclusion that a physiological mechanism is economically important suggests that these effects may exist in labor-intensive activities across the world, such as construction and agriculture, where heat exposure is high and adaptation through climate control is expensive or infeasible. Observed productivity losses in agriculture that have been attributed by default to plant growth responses to high temperatures may in fact be partly driven by lower labor productivity. These possibilities are yet to be researched.

References

- Akerberg, Daniel, Kevin Caves, and Garth Frazer. 2006. "Structural Identification of Production Functions." MPRA Paper no. 38349. Munich Personal RePEc Archive, Munich.
- Adhvaryu, Achyuta, Namrata Kala, and Anant Nyshadham. 2020. "The Light and the Heat: Productivity Co-benefits of Energy-Saving Technology." *Rev. Econ. and Statis.* 102 (4): 779–92.

- Allcott, Hunt, Allan Collard-Wexler, and Stephen O'Connell. 2016. "How Do Electricity Shortages Affect Productivity? Evidence from India." *A.E.R.* 106 (3): 587–624.
- Auffhammer, M., V. Ramanathan, and J. R. Vincent. 2006. "From the Cover: Integrated Model Shows That Atmospheric Brown Clouds and Greenhouse Gases Have Reduced Rice Harvests in India." *Proc. Nat. Acad. Sci. USA* 103 (52): 19668–72.
- Berman, Eli, Rohini Somanathan, and Hong W. Tan. 2005. "Is Skill-Biased Technological Change Here Yet? Evidence from Indian Manufacturing in the 1990s." *Ann. Econ. Stat.* 79–80:299–321.
- Burgess, Robin, Olivier Deschenes, Dave Donaldson, and Michael Greenstone. 2017. "Weather, Climate Change and Death in India." Working paper, Climate Impact Lab. <http://www.impactlab.org/research/weather-climate-change-and-death-in-india/>.
- Burke, Marshall, Solomon M. Hsiang, and Edward Miguel. 2015. "Global Non-Linear Effect of Temperature on Economic Production." *Nature* 527 (7577): 235–39.
- Chen, Xiaoguang, and Lu Yang. 2019. "Temperature and Industrial Output: Firm-Level Evidence from China." *J. Environmental Econ. and Management* 95:257–74.
- Conley, T. G. 2010. "Spatial Econometrics." In *Microeconometrics*, edited by S. N. Durlauf and L. E. Blume, 303–13. London: Palgrave Macmillan. https://doi.org/10.1057/9780230280816_33.
- Das, Sanghamitra, Kala Krishna, Sergey Lychagin, and Rohini Somanathan. 2013. "Back on the Rails: Competition and Productivity in State-Owned Industry." *Am. Econ. J. Appl. Econ.* 5 (1): 136–62.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken. 2012. "Temperature Shocks and Economic Growth: Evidence from the Last Half Century." *Am. Econ. J. Macroeconomics* 4 (3): 66–95.
- Fiorese, Elena, Franco Bonollo, Giulio Timelli, and Lars Arnberg. 2015. "New Classification of Defects and Imperfections for Aluminum Alloy Castings." *Internat. J. Metalcasting* 9 (1): 55–66.
- Gupta, Ridhima, E. Somanathan, and Sagnik Dey. 2017. "Global Warming and Local Air Pollution Have Reduced Wheat Yields in India." *Climatic Change* 140 (3–4): 593–604.
- Hsiang, Solomon M. 2010. "Temperatures and Cyclones Strongly Associated with Economic Production in the Caribbean and Central America." *Proc. Nat. Acad. Sci. USA* 107 (35): 15367–72.
- Hsiang, S. M., M. Burke, and E. Miguel. 2013. "Quantifying the Influence of Climate on Human Conflict." *Science* 341 (6151): 1235367.
- Huntington, Ellsworth. 1915. *Civilisation and Climate*. New Haven, CT: Yale Univ. Press.
- ISO (International Organization for Standardization). 1989. "Hot Environments: Estimation of the Heat Stress on Working Man, Based on the WBGT Index (Wet Bulb Globe Temperature)." Technical report, ISO, Geneva.
- JRAIA (Japan Refrigeration and Air Conditioning Industry Association). 2019. "World Air Conditioner Demand by Region." Technical report, JRAIA, Tokyo. https://www.jraia.or.jp/english/World_AC_Demand.pdf.
- Kahn, Matthew E. 2005. "The Death Toll from Natural Disasters: The Role of Income, Geography, and Institutions." *Rev. Econ. Stat.* 87 (2):271–84.
- Kjellstrom, Tord, Ingvar Holmer, and Bruno Lemke. 2009. "Workplace Heat Stress, Health and Productivity: An Increasing Challenge for Low and Middle-Income Countries during Climate Change." *Glob. Health Action* 2, <https://doi.org/10.3402/gha.v2i0.2047>.

- Kovats, R. Sari, and Shakoor Hajat. 2008. "Heat Stress and Public Health: A Critical Review." *Ann. Rev. Public Health* 29:41–55.
- Kumar, Hemanshu, and Rohini Somanathan. 2009. "Mapping Indian Districts across Census Years, 1971–2001." *Econ. and Polit. Weekly* 44 (41–42): 69–73.
- Lemke, Bruno, and Tord Kjellstrom. 2012. "Calculating Workplace WBGT from Meteorological Data: A Tool for Climate Change Assessment." *Indus. Health* 50 (4): 267–78.
- Levinsohn, James, and Amil Petrin. 2003. "Estimating Production Functions Using Inputs to Control for Unobservables." *Rev. Econ. Studies* 70 (2): 317–41.
- Lobell, David B., Wolfram Schlenker, and Justin Costa-Roberts. 2011. "Climate Trends and Global Crop Production since 1980." *Science* 333 (6042): 616–20.
- Mackworth, N. H. 1946. "Effects of Heat on Wireless Telegraphy Operators Hearing and Recording Morse Messages." *British J. Indus. Medicine* 3:143–58.
- Mendelsohn, R., and A. Dinar. 1999. "Climate Change, Agriculture, and Developing Countries: Does Adaptation Matter?" *World Bank Res. Observer* 14 (2): 277–93.
- Parry, Jonathan. 1999. "Lords of Labour: Working and Shirking in Bhilai." *Contributions to Indian Sociology* 33 (1): 107–40.
- Parsons, K. C. 1993. *Human Thermal Environments*. Boca Raton, FL: CRC.
- Schlenker, Wolfram, and Michael J. Roberts. 2009. "Nonlinear Temperature Effects Indicate Severe Damages to US Crop Yields under Climate Change." *Proc. Nat. Acad. Sci. USA* 106 (37): 15594–98.
- Seppanen, Olli, William J. Fisk, and David Faulkner. 2003. "Cost Benefit Analysis of the Night-Time Ventilative Cooling in Office Building." Technical report, Lawrence Berkeley National Laboratory, Berkeley, CA.
- USAID (United States Agency for International Development) and BEE (Bureau of Energy Efficiency). 2014. "HVAC Market Assessment and Transformation Approach for India." Technical report, USAID and BEE, Government of India.
- Vora, Rutam. 2019. "Come Summer, Commercial Air-Coolers Will Change the Course of Industry." *Hindu BusinessLine*, Feb 20. <https://www.thehindu.com/business/Companies/come-summer-2019-commercial-air-coolers-will-change-the-course-of-industry/article26321107.ece>.
- Wyndham, C. H. 1969. "Adaptation to Heat and Cold." *Environmental Res.* 2:442–69.
- Zhang, Peng, Olivier Deschenes, Kyle Meng, and Junjie Zhang. 2018. "Temperature Effects on Productivity and Factor Reallocation: Evidence from a Half Million Chinese Manufacturing Plants." *J. Environmental Econ. and Management* 88:1–17.
- Zivin, Joshua Graff, and Matthew Neidell. 2014. "Temperature and the Allocation of Time: Implications for Climate Change." *J. Labor Econ.* 32 (1): 1–26.