The Light and the Heat: Productivity Co-benefits of Energy-saving Technology*

Achyuta Adhvaryu†
Namrata Kala‡
Anant Nyshadham§

February 24, 2014

Abstract

Using detailed data from garment factories in Bangalore, India, we estimate large, negative impacts of temperature on productive efficiency (actual over target output). The introduction of LED lighting, which emits less heat than standard lighting, attenuates the temperature-productivity gradient by 75 percent. Our estimates suggest that productivity returns to LEDs are more than four times larger than the energy savings, significantly shifting firms' cost-benefit calculations in favor of adoption. We find no evidence of contemporaneous impacts of temperature on worker attendance, indicating labor supply is not a primary mechanism of impact between temperature and productivity.

Keywords: climate change, mitigation, co-benefits, temperature, energy-saving technology, LED, firm productivity
JEL Codes: J24, L20, O13, O14, Q54, Q56

*We are incredibly thankful to Anant Ahuja, Chitra Ramdas, Shridatta Veera, Manju Rajesh, Raghuram Nayaka, Sudhakar Bheemarao, Paul Ouseph, and Subhash Tiwari for their coordination, enthusiasm, support, and guidance. Special thanks to Prashant Bharadwaj, Josh Graff Zivin, Catherine Hausman, Tom Lyon, Robyn Meeks, Antoinette Schoar, and Joseph Shapiro as well as seminar participants at Michigan for helpful comments and suggestions. Robert Fletcher and Karry Lu provided excellent research assistance. Adhvaryu gratefully acknowledges funding from the NIH/NICHD (5K01HD071949). This research was partially funded by a Private Enterprise Development in Low-income Countries grant initiative. All errors are our own.

†University of Michigan; adhvaryu@umich.edu; achadhvaryu.com
‡Yale University; namrata.kala@yale.edu; namratakala.com
§University of Southern California; nyshadha@usc.edu; anantnyshadham.com

1 Introduction

Mean global temperatures are projected to rise by at least 1.5 degrees Celsius by 2100 (IPCC, 2013). Developing countries, whose adaptive and protective capacities are low and who are on average hotter than developed countries, will bear the lion's share of the ensuing negative impacts (Mendelsohn et al., 2006). Agricultural and manufacturing productivity may suffer, not only due to the increased frequency of extreme weather events[1] but also because excessive heat increases health risks (Burgess et al., 2011, 2013; Danet et al., 1999; Deschênes and Greenstone, 2011; Kudamatsu et al., 2012) and decreases the body's capacity for exertion (Kjellstrom et al., 2009; Lemke and Kjellstrom, 2012; Sudarshan and Tewari, 2013).

Given these repercussions, there is great academic and policymaking interest in quantifying the impacts of temperature (and of environmental factors more broadly) on economic outcomes. While the effect of rising temperatures on agriculture has been studied in depth by recent work (Deschênes and Greenstone, 2007; Guiteras, 2009; Kala et al., 2012; Kurukulasuriya et al., 2006; Lobell et al., 2011), impacts on industrial productivity remain relatively unexplored.[2] As economies in many low-income countries undergo major structural transformations away from agriculture and into manufacturing and services sectors (Roberts and Tybout, 1996), understanding these effects becomes even more important.

In this study, we estimate the impacts of temperature anomalies using detailed production data from garment factories in Bangalore, India. We begin by estimating the effect of changes in temperature on worker efficiency (realized output over target output) at the production line by day level. We find that efficiency decreases substantially on hotter days: a 1 degree Celsius increase in mean daily temperature lowers production efficiency by .23 points (average efficiency is 53.4 percent). Lagged temperature (averaged over the week prior) also matters for current efficiency: a 1 degree Celsius increase in the lagged average temperature decreases efficiency by .34 points. Contemporaneous effects of temperature on attendance are small and tightly bounded around 0, suggesting that the temperature-efficiency gradient derives from the direct physiological effect of temperature on workers' productive capacity, rather than an indirect mechanism via worker attendance.[3] However, lagged temperature does significantly affect attendance in most specifications, suggesting that labor supply might be one mechanism for cumulative, persistent effects of temperature on efficiency.

Having documented substantial impacts of temperature on industrial productivity and absenteeism, the natural next question is whether it is possible to mitigate these impacts. Finding effective mitigation strategies is not enough; even the most tempered by the willingness of indi-

[1] See, e.g., Deschênes and Greenstone (2007); Guiteras (2009); Hsiang (2010); Kala et al. (2012); Kurukulasuriya et al. (2006); Lobell et al. (2011).

[2] Though the idea that heat stress in the workplace could hamper productivity, particularly in low-income countries, has been written about for some time (see, e.g., Lemke and Kjellstrom (2012) and Kjellstrom et al. (2009), to our knowledge only one recent study in economics makes strides toward quantifying this impact (Sudarshan and Tewari, 2013).

[3] It bears mentioning that absenteeism is a significant problem in these factories, so there is substantial day-to-day variation in attendance.

viduals and firms to adopt them. This process is not easy: though the public benefits of mitigation may be high, the immediate private returns are generally assumed to be low or even negative (Knittel and Sandler, 2011). Achieving widespread adoption, then, might require costly sustained subsidies or taxation. These questions are obviously important from a policy perspective, and matter, too, for our understanding of firms' technology adoption decisions.[4]

We make strides toward answering this question by estimating the productivity consequences of the adoption of energy-saving technology in the garment factories under study. We show that the introduction of light-emitting diode (LED) technology on factory floors substantially attenuates the negative relationship between temperature and worker efficiency. LED lighting is 7 times more energy-efficient than standard fluorescent lighting in our setting, and emits about 1/7th the heat. We study the impacts of a staggered roll-out of LEDs over four years on the stitching floors of 25 factories operated by a large garment export firm in India.[5] The switch to LED lighting was driven in large part by changes in international buyers' recommendations for "green" policies for their suppliers. We demonstrate in a variety of checks that the roll-out across factory units and over time was not systematically related to temperature or productivity pre-trends.

We estimate the extent to which the introduction of LED lighting, through the reduced intensification of temperature on factory floors, flattens the temperature-efficiency gradient. Our estimated magnitudes are startling: on average, the introduction of LED lighting eliminates roughly 75 percent of the negative impact of temperature on efficiency. These results are robust to changes in empirical specification and to a variety of temperature measures. Interestingly, though we find that LED introduction generates significant attenuation of the temperature-productivity gradient across the temperature distribution, this differential effect is nearly twice as large below median temperature as compared to above. Finally, we explore the degree to which LED adoption affects productivity directly or beyond its interaction with temperature by estimating the main effect of LED on productive efficiency. We find no evidence of a main effect.

Our study contributes to the understanding of the effects of environmental factors on economic productivity. Recent work has documented significant labor supply and productivity impacts of air pollution (Graff Zivin and Neidell, 2012; Hanna and Oliva, 2011) and temperature (Graff Zivin and Neidell, 2010; Heal and Park, 2013; Hsiang, 2010; Sudarshan and Tewari, 2013). Our detailed productivity measures, relatively long time span, and high-frequency temperature data allow us to quantify impacts with precision. Our findings are quite consistent with the previous studies: deviations in temperature strongly impact labor productivity. These studies make important inroads, but they stop short of evaluating policies that actually lessen the negative impact of heat on productivity. We provide the first evidence to our knowledge of the impact of technology that flattens the temperature-productivity gradient, while also having important environmental benefits via climate change mitigation.

[4]In an interesting and related study, Allcott and Taubinsky (2013) test for behavioral hinderances as an explanation for low adoption of energy-efficient lightbulbs among consumers. We show among firms that non-adoption is even more suboptimal once production benefits are taken into account.

[5]Our data sample includes 29 factories, four of which did not receive LED lighting.

We also add to the literature on the returns to climate change mitigation.[6] The few recent studies that examine "co-benefits," or additional gains, of mitigation focus largely on the indirect public returns (Bollen et al., 2009; Knittel and Sandler, 2011). For example, a carbon tax aimed primarily at reducing CO_2 emissions may also reduce emissions of accompanying pollutants such as nitrous oxides or carbon monoxide, thus reducing substantially the tax's welfare costs. Our study examines a novel, *private* co-benefit of climate change mitigation. This distinction is important because the success of most mitigation strategies relies on individuals' and firms' willingness to adopt them, and this willingness is largely driven by private returns. If energy-saving technologies like LEDs do have substantial private co-benefits, this should meaningfully alter firms' benefit-cost calculations. Indeed, by our estimation, the benefits of LEDs in terms of productivity are substantially larger than energy savings (in fact, roughly 4 times as large), indicating that ignoring the former source of benefit would seriously underestimate the private returns to adoption.

Finally, we contribute to the study of the determinants of firm and worker productivity in low-income contexts. Recent work has demonstrated that management quality (Bloom et al., 2013; Bloom and Van Reenen, 2010), intra-firm networks (Bandiera et al., 2009, 2010), incentive structures (Bandiera et al., 2007), ethnic boundaries (Hjort, 2013), and environmental disasters (De Mel et al., 2012) all significantly impact productivity. We show that temperature, and interventions that affect ambient temperature on factory floors, play an importnant role in worker efficiency as well.

The remainder of the paper is organized as follows. Section 2 describes contextual details regarding garment production in India and LED technology. Section 3 provides details on the temperature and production data. Section 4 describes our empirical strategy. Section 5 describes the results, and section 6 has a concluding discussion.

2 Context

In this section, we 1) discuss the garment sector in India and key elements of the garment production process; 2) review the physiology of the relationship between temperature and worker productivity; 3) provide an overview of energy usage and heat emissions in LED v. fluorescent lighting; and 4) describe the roll-out of LED lighting across the garment factories in our data.

2.1 The Indian Garment Sector

Global apparel is one of the largest export sectors in the world, and vitally important for economic growth in developing countries (Staritz, 2010). India is the world's second largest producer of textile and garments, with the export value totaling $10.7 billion in 2009-2010. Women comprise the majority of the workforce (Staritz, 2010). Total employment in India's formal apparel and

[6] A related literature has established patterns of adaptation to climate change and the returns to this adaptation (e.g. Barreca et al. (2013)).

textile industry was about 2 million in 2008, of which 675,000 was in the formal apparel sector, making this a crucial component of India's industrial sector.

2.2 The Garment Production Process

There are three broad stages of garment production: cutting, sewing, and finishing. In the factories that comprise our production data, garments are sewn in production lines. Each line will produce a single style of garment at a time (i.e. color and size will vary but the design of the style will be the same for every garment produced by that line until the order for that garment is met). Lines consist of 20-100 sewing machine operators (depending on the complexity of the style) arranged in sequence and grouped in terms of segments of the garment (e.g. sleeve, collar, placket).[7] Completed sections of garments pass between these groups, are attached to each other in additional operations along the way, and emerge at the end of the line as a completed garment. These completed garments are then transferred to the finishing floor.

Before reaching the sewing floor, pieces of fabric needed for each segment of the garment are cut using patterns from a single sheet so as to match color and quality perfectly. These pieces are divided according to groups of sewing operations (e.g. sleeve construction, collar attachment) and pieces for 10-20 garments are grouped and tied into bundles. These bundles are then transported to the sewing floors where they are distributed across the line at various "feeding points" for each group of sewing operations.

In finishing, garments are checked, ironed, and packed. A great degree of quality checking is done "in-line" on the sewing floor, but final checking occurs in the finishing stage. Any garments with quality issues are sent back to the sewing floor for rework or, if irreparably ruined, are discarded before packing. Orders are then packed and sent to port.

2.3 Physiology of the Temperature-Productivity Gradient

The physical impact of temperature on human beings is a very well-studied area (Enander, 1989; Parsons, 2010), and has traditionally been important in order to establish occupational safety standards for workers exposed to very high or low temperatures for continued periods of time (Vanhoorne et al., 2006). Higher temperatures and consequent thermal stress can impact human beings not only physically, but also through lower cognition and psychomotor ability (Hancock et al., 2007). For instance, Ramsey et al. (1983) find increases in unsafe behavior by workers at temperatures greater than 35 degrees Celsius WBGT (Wet Bulb Globe Temperature). The individual impact on a person varies based on factors such as the type of task and its complexity, duration of exposure, as well as the worker-level skill and acclimatization level (Pilcher et al., 2002), which contributes to the issues in setting a particular limit in working environments (Hancock et al., 2007).

[7] In general, we describe here the process for woven garments; however, the steps are quite similar for knits and even pants, with varying number and complexity of operations. Even within wovens, the production process can vary a bit by style or factory.

2.4 LED v. Fluorescent Lighting

LED light bulbs are approximately 7 times as energy-efficient as fluorescent bulbs (requiring about 3 as opposed to 21 KWh/year in electricity in our setting), and thus operate at about 1/7 the cost of fluorescent lighting. In addition, they generate a tenth of the CO_2 emissions (5.01 pounds of CO_2 per year per bulb, as compared to 35.11 pounds for fluorescent lighting).[8] Heat emissions for LEDs are substantially lower than fluorescent bulbs: the average LED bulb emits 3.4 Btus, as against 23.8 Btus for fluorescent lighting in the setting we study.

2.5 LED Roll-out: Summary and Timeline

The factories began installing LED lighting in October 2009 and completed the installations by February 2013. There was no formal documentation of the reasons for LED adoption in each factory, but according to senior management at the firm, the reasons were twofold. First, over the last decade, buyers have become more stringent in their regulation of their suppliers' production standards and environmental policies. This generated a staggered roll-out of LEDs across factories within the firm because some factories were more heavily involved in the production of orders from particular buyers than others. So, for example, if buyer A's environmental regulations or production guidelines become more stringent, then the supplier might choose to convert to LED lighting in factories processing many orders from buyer A. When buyer B's regulation change, the firm will prioritize factories servicing buyer B, and so on. Second, over the study period, the firm itself began a variety of "green" initiatives firm-wide, and thus scaled up LED introduction across its factories. These two factors combine to generate wide variation in the timing of LED take-up across the factories we study. To test whether this timing was exogenous to line-daily efficiency, we perform some checks, which are described in detail in section 5.

The replacement took the form of substituting all fluorescent lights targeted at individual operations with an equivalent number of small LED lights mounted on individual workers' machines. The replacements were designed to maintain the original level of illumination. On average, each unit replaced about 1,000 fluorescent lights of 7W each with 1,000 LED lights of 1W each.[9] Based on the factories' operating time cost calculation, this meant an energy saving of 18KWh per light per year. In the conclusion, we discuss the magnitude of the environmental benefits from the installation.

A particular factory received the installation within a single month. 8% of the LED rollout (2 units) was completed in 2009, 48% (12 units) in 2010, 16% (4 units) in 2011, about 24% (6 units) in 2012 and the rest (1 unit) in 2013. Of the 29 units from which we have productivity data, LED replacements occurred in 25 units. Since our productivity data ranges from April 2010 to June

[8] Note that while both fluorescent and LED lighting are much more efficient than incandescent bulbs, the factories in our sample did not have any incandescent lighting on the production floor. For details on emissions calculations, please refer to section 6.

[9] The number of lights installed unit by unit is a function of the number of machines in the unit, and varies from about 100 to 2,550 with a mean of 1,000.

2013, some units already have LEDs at the beginning of our productivity data, and all but four units have LED for the last four months of our productivity data range. This is why our results report the impact of temperature not only for the whole sample, but also separately for the units that did not have LED at a particular time.

3 Data

Here we provide an overview of data sources, describe our data via summary statistics, and provide preliminary graphical evidence on the temperature-productivity gradient and the effects of LED introduction.

3.1 Data Sources

3.1.1 Temperature Data

Our temperature data is from the National Climatic Data Center (NCDC) at the National Oceanic and Atmospheric Administration (NOAA). The Center compiles global station-level weather data at the hourly level, although most stations outside the US report data only at 3-hour intervals. We have temperature data from three stations in Bangalore over the period spanned by our productivity data (April 2010 to June 2013). Over 80% of the observations are at the 3 hourly level (midnight IST onwards), and the rest are across somewhat arbitrary time measurements. We consider all observations across the three available stations between 9am and 6pm for each day as our daily temperature measure, which constitutes the working hours at the factories we consider.[10] This is dry bulb temperature, and at the daily level ranges between 18.8 and 34.4 degrees Celsius, with an average of about 27.6 degrees Celsius.

We have relative humidity data at the monthly level for three stations in Bangalore from January 2010 to November 2013, from the NOAA's National Data Center (NNDC). We were unable to find relative humidity data at a daily level for Bangalore, and are consequently unable to include daily relative humidity. Accordingly, the specifications using dry bulb temperature alone at the daily level are our preferred specifications in the results below. Nevertheless, at the monthly level, relative humidity ranges between about 44% and 85.7%.[11].

With the daily temperature and monthly relative humidity data, we construct two measures that incorporate both temperature and humidity. The first is the Heat Index (HI) that is calculated

[10]Note that since not all observations are available for each of the three stations, assigning factories to the nearest station is not feasible. About 60% of our observations come from the first station, and about 20% each from the second and third. The correlation between the observations from the first and second station as well as the first and third station is between 0.92 and 0.93, which is very high (the second and third stations do not overlap in terms of data availability).

[11]Analogous to the temperature calculations, we average over the station observations - the correlation between station-level observations ranges from 0.92 to 0.99

based on the formula:

$$\begin{aligned} HI = & -42.379 + 2.04901523 * T_d + 10.14333127 * rh - .22475541 * T_d * rh \\ & - .00683783 * T_d^2 - .05481717 * rh^2 + .00122874 * T_d^2 rh \\ & + .00085282 * T_d * rh^2 - .00000199 * T_d^2 * rh^2. \end{aligned} \tag{1}$$

where T_d = dry bulb temperature in Fahrenheit and rh = relative humidity (%). The formula for the calculation is derived from the Rothfusz regression that replicates the HI values from Steadman (1979).

For about 0.6% of our data, the relative humidity is greater than 85% and daily temperature ranges between 80 and 87 degrees Fahrenheit, and the following adjustment is applied:

$$HI = HI + [(rh - 85)/10] * [(87 - T_d)/5] \tag{2}$$

The second measure is a particular method for calculation Wet Bulb Globe Temperature that is suitable for indoor exposure. The formula is from Lemke and Kjellstrom (2012), and is given by:

$$WBGT = 0.567T_d + 0.216 \left(\frac{rh}{100} * 6.105 \exp\left(\frac{17.27T_d}{237.7 + T_d} \right) \right) + 3.38. \tag{3}$$

All the three measures of temperature – dry bulb temperature, Heat Index (HI), and Wet Bulb Globe Temperature (WBGT) – are converted into Celsius to ensure interpretative ease across regression specifications. While there are numerous formulae for the calculation of varied heat indices, we chose these two since they were relatively easy to calculate and interpret. For all our results, we report the main effect of dry bulb temperature as well as the main effect of dry bulb temperature controlling for relative humidity in addition to the impact of the Heat Index and the Wet Bulb Globe Temperature. The results corresponding to specifications including dry bulb alone are preferred, but the other measures are reported as evidence of robustness.

3.1.2 Factory Data

We use data on line-level daily production from 29 garment factories in and around Bangalore, India. Identifiers include factory unit number and line number within the factory. For each line and day within each factory unit, production measures include actual quantity produced, actual efficiency, and budgeted efficiency.

Actual efficiency is actual quantity produced divided by target quantity. The target quantity is derived from an industrial engineering (IE) measure for the complexity of the garment called "Standard Allowable Minute" (SAM). This measure amounts to the estimated number of minutes required to produce a single garment of a particular style. This estimate comes from a central database of styles, but is then amended by the factory's IE department during "sampling." Sampling is the process by which a style that is ordered by a buyer is costed in terms of labor and production time. The highest skill grade tailor, called a sampling tailor, will make a garment of a particular style entirely and recommend a SAM for that style to the IE department.

This SAM is then used to calculate the target quantity for the line for each hour of production. Each line runs for 8 hours during a standard work day. Accordingly, a line producing a style with a SAM of .5 will have a target of 120 garments per hour, 960 garments per day. Most importantly, the target quantity will be largely fixed across days (in fact, across hours within the day) within a particular order of a style.

Each line will only produce a single style at any time. However, depending on the order size (or "scheduled quantity") for a style, multiple lines might be producing the same style at one time and each line could produce a style for many days.[12] Of course, a line which has been producing the same style for many days will likely be more efficient at producing that style than will a line which has been producing a style (of even the same complexity or SAM) for less days.

That is, let us say that line 1 is producing some style X. The order from the buyer for style X is for a quantity of 10,000 garments. The SAM for style X is calculated by the sampling department and IE department to be .5. Then, line 1 is estimated to make $60/.5 = 120$ garments per hour, or 960 garments per day. Then, if line 1 produces exactly 960 garments of style X each day, its actual efficiency will be $960/960 = 100\%$. At this efficiency, line 1 will complete the order in roughly 10.5 production days. If, instead, line 1 produces only 480 garments of style X on the first day, because it is still learning the production details of style X or perhaps because it is too hot to produce efficiently, the actual efficiency for line 1 and day 1 will be 480/960=50%. Of course, even if line 1 produces at 100% efficiency on all subsequent days, it will take a full 11 days to complete the order instead of 10.5 days.

Predictable variations in efficiency, due to learning of new styles or line-specific characteristics such as number of operators of each skill grade, are reflected in the budgeted efficiency. Consequently, actual efficiency of a given style will vary systematically across lines and within line over time. We are, of course, interested in unpredictable, relatively transitory fluctuations in productivity due to temperature rather than these systematic fluctuations across lines due to line-specific or operator-specific characteristics and within lines over time due to order size and style complexity. We will accordingly control for budgeted efficiency and include line fixed effects in the regression analysis below.

Most importantly, we use actual efficiency rather than produced quantity as our outcome of choice. Produced quantity would not account for systematic variation due to complexity of style or number of operations. Without normalizing production observations to target quantity, one could potentially misrepresent an association between temperature and style complexity or order size as an impact on productivity. That is, for example, if garment complexity or line length varied by temperature due to seasonal buying of winter garments at certain times in the fall months or labor supply in harvest weeks, resulting variations in efficiency could be attributed to temperature spuriously. Accordingly, we argue that actual efficiency, controlling for budgeted efficiency, is the

[12] Indeed, in our data, lines produce styles for between 1 and 199 days.

most appropriate outcome for the empirical exercise proposed in this study.

To summarize, target quantity will reflect only style by line characteristics which do not vary day to day and certainly do not vary with temperature fluctuations across days. Actual quantity will indeed vary with daily productivity, of which we hypothesize temperature is an important determinant, but must be normalized by target quantity to be compared across lines and within lines across styles. Even within styles and lines, certain predictable variations and evolutions in actual efficiency arise due to, among other things, learning by doing which is a function of garment complexity, order size, and line- and order-specific characteristics. True daily fluctuations in productivity are, therefore, best measured by actual efficiency net of budgeted efficiency.

3.2 Summary Statistics

We present means and standard deviations of variables used in the analysis below. Our sample consists of 446 lines across 29 factory units. The range of dates over which we have production data spans 941 days in total. However, we do not observe all factory units, nor all lines within a unit, for all dates. We restrict our attention to lines for which we observe production data for at least 40% of dates. Altogether, our data includes nearly 215,000 line-day observations. Roughly, one-third of the observations correspond to days in factory units prior to the introduction of LED lighting and the rest are post-LED observations.

The summary statistics indicate a great deal of variation in the measures of temperature. The means of these measures appear quite similar before and after the introduction of LED. This is to be expected given that the timing of LED introduction varied at the factory level across nearly the entire date range. That is, 6 units already have LED lighting at the beginning of the date range and 4 units still have not received LED lighting by the end of the date range. On the other hand, both actual and budgeted efficiency appear to differ on average between the before and after LED samples. Average efficiency appears higher after LED introduction, while budgeted efficiency appears lower.

3.3 Preliminary Graphical Evidence

We begin by motivating the central exercise of this study with descriptive plots of production and temperature data.

3.3.1 Temperature-Productivity Gradient

Underlying the analysis conducted below is the assertion that productivity and temperature are negatively correlated. In order to check this assertion in our empirical context, we plot both the daily time series of actual production efficiency and the dry bulb temperature. These plots, presented in Figure 1, depict a distinct negative correlation and, perhaps, a slight lag in the impact of

Figure 1: Efficiency and Temperature Time Series

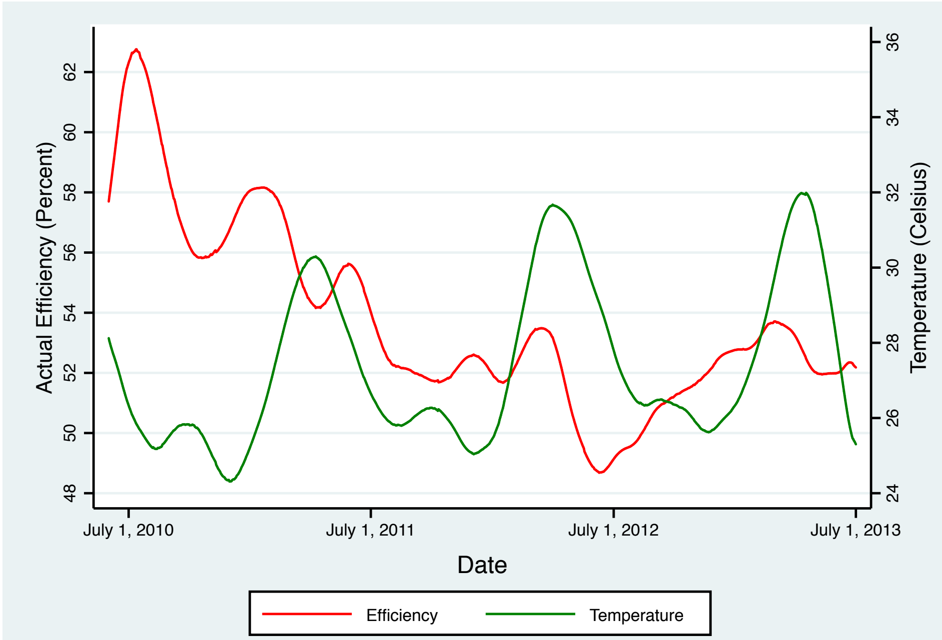

Notes: Efficiency is produced quantity divided by targeted quantity. Date trimmed before June 1, 2010 due to limited production data.

temperature on efficiency.[13]

We next collapse the time element of the data and plot actual efficiency as a function of dry bulb temperature. Figure 2 shows that indeed efficiency appears to be a downward-sloping function of temperature.

3.3.2 Impacts of LED Introduction

Having provided preliminary evidence of a negative temperature-productivity gradient for the garment factories in our data, we next check for evidence that this gradient is affected by the replacement of the ambient fluorescent lighting in factories with focused, machine-mounted LED lighting. We repeat the exercise from Figure 2 for subsets of the data from before and after the LED roll-out in each factory. These plots are presented in Figure 3. The evidence suggests that factories are more efficient at all temperatures after the LED introduction. This efficiency gap is increasing in temperature due to a less negative slope with LED, particularly in the lowest and highest ranges of temperature.

We also return to the exercise conducted in Figure 1, but plot the efficiency series separately for factory-day observations with and without LED lighting. Once again, we include the temperature

[13]Note that there appears to be a downward trend in efficiency over the time period covered by the data. This is likely due to rampant expansion across many factory units during this time that stresses resources and periodically depresses efficiency within the unit. This is another strong indication for the inclusion of unit x year FE and budgeted efficiency as a control.

FIGURE 2: EFFICIENCY AGAINST TEMPERATURE

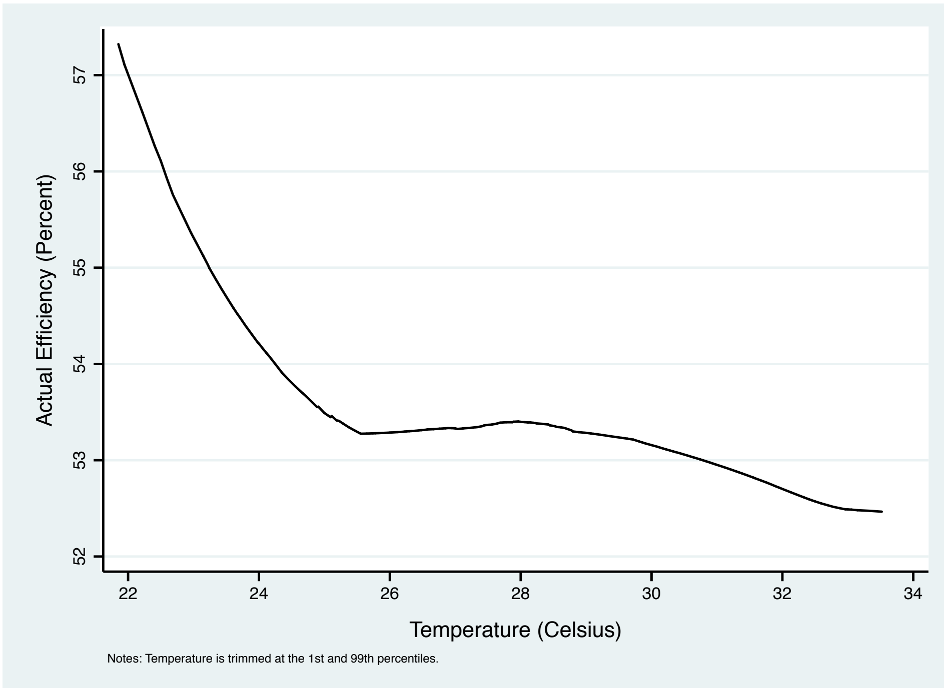

FIGURE 3: EFFICIENCY AGAINST TEMPERATURE BY LED

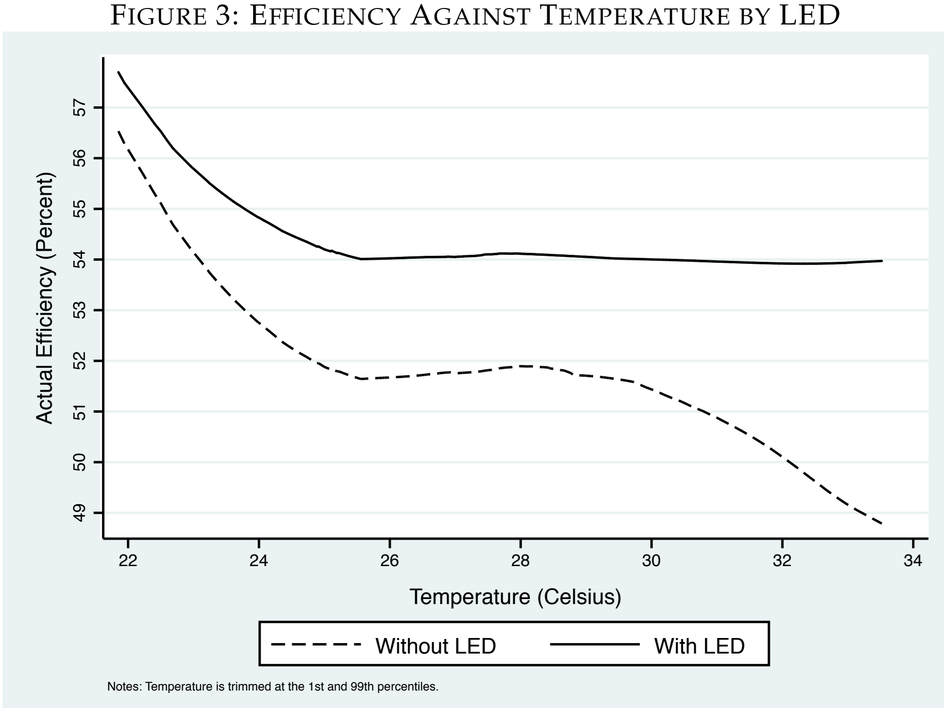

FIGURE 4: EFFICIENCY AND TEMPERATURE TIME SERIES BY LED

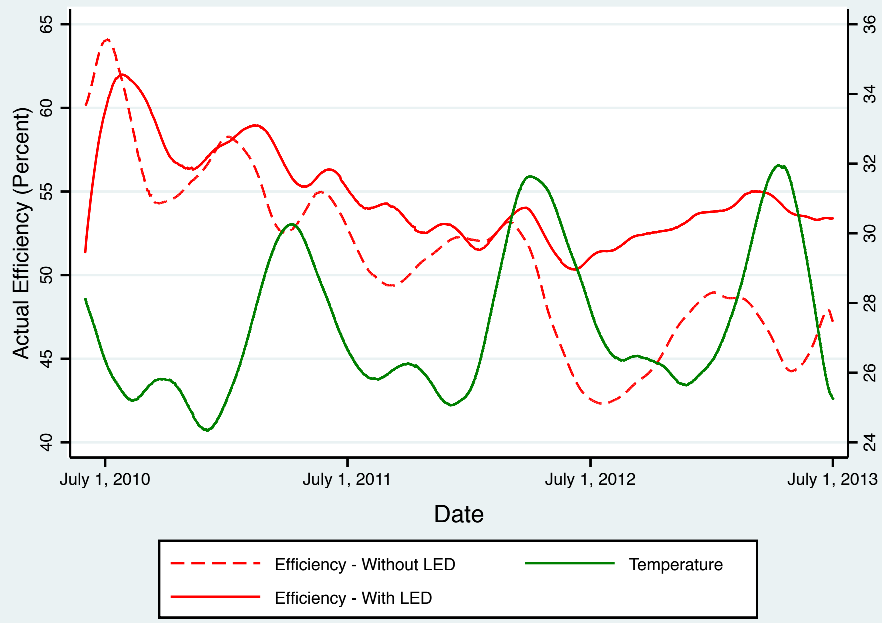

Notes: Efficiency is produced quantity divided by targeted quantity. Date trimmed before June 1, 2010 due to limited production data.

time series for comparison and present the plots in Figure 4. The evidence in Figure 4 also suggests that LED lighting improves efficiency on all days, but particularly smooths the fluctuations in efficiency due to temperature.

Finally, we explore graphically the main effects of LED lighting on efficiency. That is, there might be many pathways by which LED introduction might affect efficiency directly or outside of its interaction with the temperature-efficiency gradient. For example, the quality and quantity of light might have effects on attentiveness and sight, both of which are important in garment manufacturing. While, the previous figures indicate that LED might have a mitigative impact on the temperature-efficiency gradient, estimating the mean composite effect of LED introduction on efficiency is also of interest in what follows. We present preliminary evidence of this main effect by plotting actual efficiency over time relative to the date of LED introduction. However, Figure 5 shows little evidence of a mean composite effect.

Motivated by this preliminary evidence we set forth a more rigorous regression analysis below to causally identify both the effect of temperature on production efficiency and the attenuation of this impact driven by the replacement of traditional fluorescent lighting with LED technology. In particular, we address concerns regarding unit-level trends in efficiency, line-level unobservables, seasonality in efficiency, and the exogeneity of the LED introduction.

FIGURE 5: EFFICIENCY BEFORE AND AFTER LED

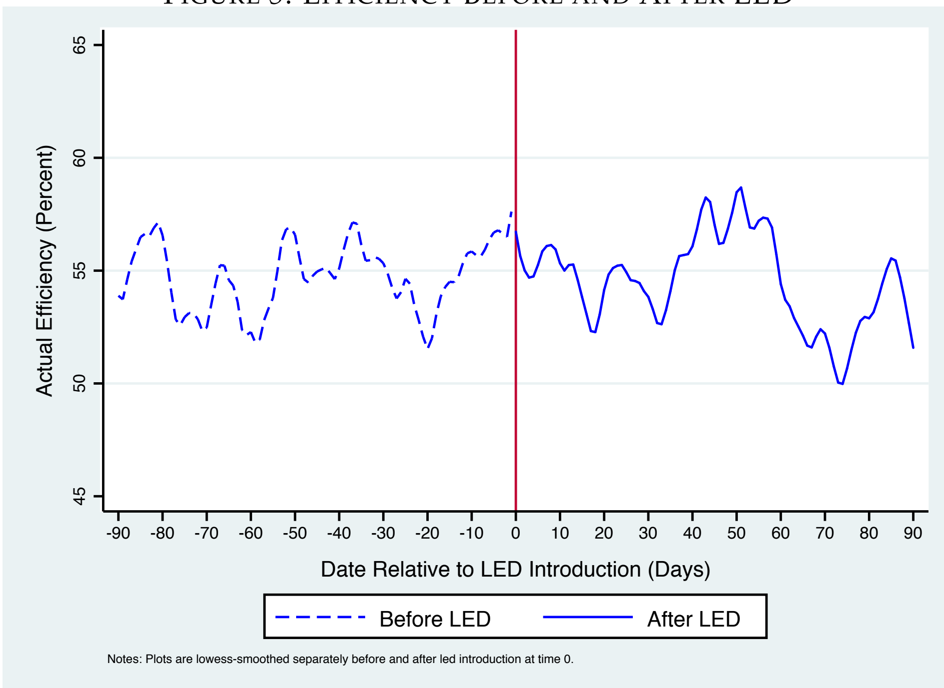

4 Empirical Strategy

'

First, we estimate the following empirical specification of the relationship between worker efficiency and temperature:

$$E_{ludmy} = \alpha_0 + \beta T_{dmy} + \phi B_{ludmy} + \alpha_l + \gamma_{uy} + \eta_m + \delta_d + \varepsilon_{ludmy}. \tag{4}$$

Here, E for efficiency of line l of unit u on day d in month m and year y; B is budgeted efficiency for line l of unit u on day d in month m and year y; T is daily temperature in degrees Celsius; α_l are line fixed effects; γ_{uy} are unit x year fixed effects; η_m are month fixed effects; δ_d are day-of-week fixed effects; and α_0 is an intercept. β is the coefficient of interest, giving the impact of a 1-degree Celsius increase in temperature on line-level efficiency.

In addition to the average effect of temperature on efficiency, we are also interested in testing whether this effect, and its corresponding attenuation by LED lighting, is heterogenous across the distribution of temperature. To implement this, we estimate equation 4 allowing for varying slopes of temperature above and below the median of the temperature distribution. We use the following empirical specification:

$$E_{ludmy} = \alpha_0 + \zeta_1 \left(Q_1 \text{ x } T_{dmy}\right) + \zeta_2 \left(Q_2 \text{ x } T_{dmy}\right) + \zeta_5 Q_1 + \zeta_6 Q_2 + \alpha_l + \gamma_{uy} + \eta_m + \delta_d + \varepsilon_{ludmy}. \tag{5}$$

where Q_1 is a dummy variable that is 1 if temperature is above median temperature and zero otherwise, and Q_2 is a dummy variable that is 1 if temperature is below median temperature and zero otherwise.

We then estimate the extent to which the introduction of LED lighting attenuates the temperature-productivity relationship via the following specification:

$$E_{ludmy} = \alpha_0 + \beta_1 \left(T_{dmy} \text{ x } LED_{umy}\right) + \beta_2 LED_{my} + \beta_3 T_{dmy} + \phi B_{ludmy} + \alpha_l + \gamma_{uy} + \eta_m + \delta_d + \varepsilon_{ludmy}. \quad (6)$$

Here LED_{umy} is a dummy for presence of LED lighting in unit u in month m and year y. It changes from 0 to 1 in the month following LED introduction in a particular factory unit. The coefficients of interest in the above specification are β_1 and β_3. β_3 indicates the effect of temperature on productivity *before* LED introduction. β_3 is the extent of attenuation of the temperature-productivity gradient once LED lighting is introduced. The sum of these two, $\beta_1 + \beta_3$, gives the net effect of temperature on productivity following LED introduction.

Finally, we test if the attenuation impact of LED is heterogenous for temperatures above and below the median temperature by estimating the following specification:

$$\begin{aligned} E_{ludmy} = \alpha_0 &+ \beta_1 Q_1 \text{ x } T_{dmy} \text{ x } LED_{umy} + \beta_2 Q_2 \text{ x } T_{dmy} \text{ x } LED_{umy} + \psi_1 Q_1 \text{ x } T_{dmy} + \psi_2 Q_2 \text{ x } T_{dmy} \\ &+ \zeta_1 Q_1 \text{ x } LED_{umy} + \zeta_2 Q_2 \text{ x } LED_{umy} + \theta Q_1 + \phi B_{ludmy} + \alpha_l + \gamma_{uy} + \eta_m + \delta_d + \varepsilon_{ludmy}. \quad (7) \end{aligned}$$

Following Graff Zivin and Neidell (2012), we use two-way clustering of standard errors in all regressions (Cameron et al., 2011). Standard errors are clustered 1) at the production line level, because our cross-sectional variation is the line level, and we would expect inter-temporal correlation in productivity within lines; and 2) at the date level, since this is the level at which temperature varies.

4.1 Controlling for Unobservables

Note that all specifications include as controls budgeted efficiency, line fixed effects, unit by year fixed effects, month fixed effects, and day of the week fixed effects. Line fixed effects are meant to control for unobservable determinants of efficiency at the line level that are static over time such as line supervisor characteristics (e.g. management style, experience, rapport and relationship with workers), type of garment usually produced by the line (e.g. shirt vs. pant, denim vs. twill), and position in the factory (e.g. higher floor where it is hotter, closer to the window where there is better light and ventilation). Unit by year fixed effects are meant to control not only for static unobservables at the unit level such as characteristics of factory management and factory location, but also for unobservable factors driving unit-specific non-linear trends such as differential rates of expansion across factories or primary buyers. Month fixed effects control for seasonality due to, for example, garment demand and labor supply patterns; day of week fixed effects control for fluctuations in efficiency across work days due to for example fatigue or weekend salience.

Budgeted efficiency controls for variations in actual efficiency due to unobserved determinants at the line by garment style level that vary over time such as garment style complexity, order size and time since garment style production began in that line, number of operations and operators in the line, and distribution of skill grade among operators in the line. Specifically, any fluctuations in efficiency which are predictable from the perspective of factory management and line supervisors and, therefore, might factor into endogenous production decisions, are well-summarized in budgeted efficiency.

5 Results

In this section, we present and discuss the results of the estimation strategy proposed in section 4 above.

5.0.1 Average Impact of Temperature

We begin by verifying the apparent negative temperature-productivity gradient depicted in Figure 2. In Table 2, we present results from the regression of actual efficiency on various measures of temperature. Unless otherwise noted, all specifications, as discussed in section 4, include budgeted efficiency as a control as well as month, day-of-week, line, and unit by year fixed effects.

Column 1 in Panel A shows that a one degree Celsius increase in the dry bulb temperature leads to a .23 percentage point reduction in efficiency (as compared to a mean efficiency of roughly 53.4 percent). This estimate is significant at the 5 percent level and is robust across various temperature measures. Column 2 presents the impact of a one degree rise in dry bulb temperature controlling for relative humidity; while columns 3 and 4 present the impact of a one degree rise in heat index and wet bulb globe temperature, respectively. Across all measures estimates are negative, statistically significant and of nearly identical magnitude. Estimates range from .16 to .26 percentage points across the other measures of temperature.

The regressions reported in Panel A are estimated on the sample of observations prior to the introduction of LED lighting in order to isolate the unmitigated temperature-efficiency gradient. Panel B of Table 2 shows estimates from regressions identical to those reported in Panel A, but including observations from both before and after the introduction of LED lighting. Point estimates from Panel B are consistently smaller than those from Panel A ranging from .1 to .17 percentage points. All estimates in Panel B, as in Panel A, are significant at the 1 to 5 percent level.

5.1 Lagged Impact of Temperature

Having shown the significant impact of contemporaneous temperature on efficiency, we next explore the degree to which this impact persists and/or compounds over time. We test for persistent, cumulative effects of temperature by re-estimating the regressions from Table 2 including the mean temperature over the prior week as an additional regressor. The results of these regressions are reported in Table 3. Once again, the point estimates in panel A correspond to the pre-LED sample; while Panel B presents estimates from the whole sample.

In Panel A, the point estimates on the one week lag are precisely estimated, negative, and large in magnitude across all measures of temperature. Point estimates are more than 50 percent larger than the contemporaneous effect of temperature, presented in panel A of Table 2. These results suggest that the effect of temperature on efficiency does in fact persist or even compound. Estimates in Panel B are less precisely estimated and the magnitudes are less than half of those from Panel A.[14]

Note that when the 1 week lagged temperature is included, estimates of the effect of contemporaneous daily temperature on efficiency become smaller and less significant. This is most likely due to the strong serial correlation in temperature. The pairwise correlations between yesterday's temperature, today's temperature, and tomorrow's are upwards of .8. Therefore, it is unfortunately impossible to determine which day's temperature in the past week is driving the effect of the 1 week lagged temperature, or if all the days in the past week are impactful. In Table A.2, we repeat the analysis from Table 3, but disaggregate the 1 week lagged temperature into the temperatures from each of the 7 days in the last week and include these days 1 at a time. The results indicate that indeed all the days in the past week negatively impact current efficiency, but due to the serial correlation only some of the coefficients are statistically significant.

Next, we explore the impact of temperature on worker attendance. Considering that when the work attendance decision is made in the morning the temperature is still quite mild, we do not expect that contemporaneous temperature will have much of an effect on attendance. Of course, information on temperature forecasts might be available and the high serial correlation in temperature indicates that the temperature from the day before provides a reasonable forecast; however, empirical estimates of the effects of contemporaneous temperature alone on attendance show little evidence of an effect (not reported).

However, we suspect that lagged temperature might impact future attendance, perhaps by way of exhaustion. We estimate regressions identical to those reported in Table 3, but with a binary for whether the worker attended a full day of work as the outcome. While partial attendance, particularly leaving early due to heat exhaustion, might also be of interest, we focus here on full day attendance because partial absences are extremely rare in the data (less than 1 percent). We present results from these regressions in Table 4.

Notice these data are available at the worker level, rather than the production line level. Accordingly, the sample of observations is much larger. The results indicate that the probability that a worker comes to work today is indeed reduced by the mean temperature over the prior week.

[14] The fact that the estimates include observations after the introduction of LED lighting presented in Panel A are attenuated and insignificant might indicate that the introduction of LED lighting offsets the persistent impacts of temperature on efficiency as well. However, there is insufficient residual variation in the interaction of the 1 week lagged temperature with LED after controlling for the interaction of LED with contemporaneous temperature.

A one degree celsius increase in mean temperature over the prior week leads to a .45 percentage point reduction in the probability of attending in the preferred specification (estimates range from .13 percentage points to .65). Estimates are quite similar from both the sample of observations prior to LED introduction (reported in Panel A) and the whole sample (reported in panel B).

On the other hand, consistent with the discussion above, point estimates of the contemporaneous effect of temperature on attendance are consistently insignificant and an order of magnitude smaller than those of the 1 week lagged temperature.[15] Taken together, we interpret the results for attendance as evidence that labor supply is not the primary mechanism of impact for contemporaneous temperature on efficiency. On the other hand, labor supply impacts of lagged temperature might be contributing to the persistent, cumulative effects of temperature on efficiency depicted in Table 3. We do not push the interpretation of lagged impacts further due to concerns regarding serial correlation in temperature, and undertake specification robustness checks in Table A.3 to partially alleviate concerns regarding serial correlation in temperature.

5.2 LED and the Temperature-Efficiency Gradient

Having established the impacts of temperature on production outcomes, we next investigate the degree to which LED attenuates these impacts. We regress actual efficiency on the interactions of our four measures of temperature with a dummy for the presence of LED lighting in the factory along with the main effects of temperature and LED introduction. The remainder of the specification is identical to that presented in Table 2. The results of these regressions are presented in Table 5. Estimates of the coefficients on the interaction of LED introduction with all measures of temperature are positive, statistically significant, and large in magnitude relative to the main effects of temperature. The estimates overall provide strong evidence that the partial replacement of incandescent bulbs with LED bulbs significantly offset the negative impact temperature on production efficiency.

The results in column 1 indicate that the introduction of LED lighting offsets .24 percentage points in efficiency loss for each degree of temperature. This amounts to roughly a 75% reduction in the impact of temperature on efficiency. The final row of Table 4 shows that the impact of temperature on efficiency net of attenuation from LED introduction is only weakly negative and insignificant. Estimates in columns 2 through 4 using alternate measures of temperature are very similar with LED introduction significantly offsetting the negative effects of temperature.

While we may not necessarily expect that LED has as strong a mitigative effect on the impacts of lagged temperature, we are not able to explicitly test for this effect. In the data, due to the strong serial correlation in temperature, there is insufficient independent variation in the interaction of LED with the 1 week lagged mean of temperature once we control for the interaction of the contemporaneous temperature with LED. Similarly, given the evidence of minimal effects of contemporaneous temperature on attendance, we do not report results from regressions of attendance

[15]Results from specifications regression attendance on only the contemporaneous temperature also show no evidence of an impact of LED and are available upon request.

on the interaction of temperature with LED introduction.[16]

5.2.1 Distributional Impacts of Temperature

In Table 6, we investigate the degree to which the temperature-efficiency gradient has different slopes above and below the median of the temperature distribution, and accordingly, the degree to which the mitigative impact of LED is more strongly realized at above or below median temperatures. The regressions reported in Table 6 are identical to those in the first columns of Tables 2 and 5, with added triple interactions of heat index, LED, and dummies for temperatures below and above the median. The main effect of the below median dummy is omitted to preserve the constant.

The results suggest that indeed the slope of the temperature-efficiency gradient is steeper below the median without LED, but this relationship does not appear in the whole sample . The mitigative impact of LED appears to be strongest for below median temperatures as well. These results could perhaps indicate that at sufficiently high temperatures the reduction in indoor temperature due to LED lighting replacement is less noticeable.[17]

5.3 Main Effect and Exogeneity of LED Introduction

As mentioned in section 3 above, there are several channels by which LED might affect efficiency directly, in addition to its demonstrated mitigative effect on the temperature-efficiency gradient. For example, changes in the quality and quantity of light might improve attention and sight, which in turn affect efficiency. However, the preliminary evidence presented in Figure 5 does not support a main effect. In order to verify Figure 5, we regress actual efficiency on the introduction of LED in the usual specification, but with temperature and its interaction with LED omitted. The results are reported in column 1 of Table 7 and show no evidence of a main effect of LED introduction on efficiency.

We also conduct checks of the exogeneity of the timing of the roll-out of LED bulb replacement across factory units. To the degree that temperature deviations from monthly means are plausibly exogenous, we do not necessarily need LED roll-out to also be exogenous in order to interpret the coefficient on the interaction of LED introduction and temperature in the main results as causal. That is, any correlation of LED introduction timing with unobservable determinants of efficiency across factories or over time ought to be addressed by the inclusion of the main effect of LED introduction, so long as these unobservables are orthogonal to temperature deviations within month. Furthermore, as mentioned in section 2, senior managers at the garment factories indicated that LED introduction was driven mostly by efforts to comply with changing environmental policies of the companies of specific buyers.

[16]The results show no evidence of a role for LED and are available upon request.

[17]Notice these regression results do not necessarily match the preliminary evidence depicted in Figure 3. This is likely due to the month fixed effects which restrict comparisons of efficiency to days within a month and remove the contribution of seasonal patterns to the temperature-efficiency gradient.

FIGURE 6: EFFICIENCY BY DATE RELATIVE TO LED INTRODUCTION

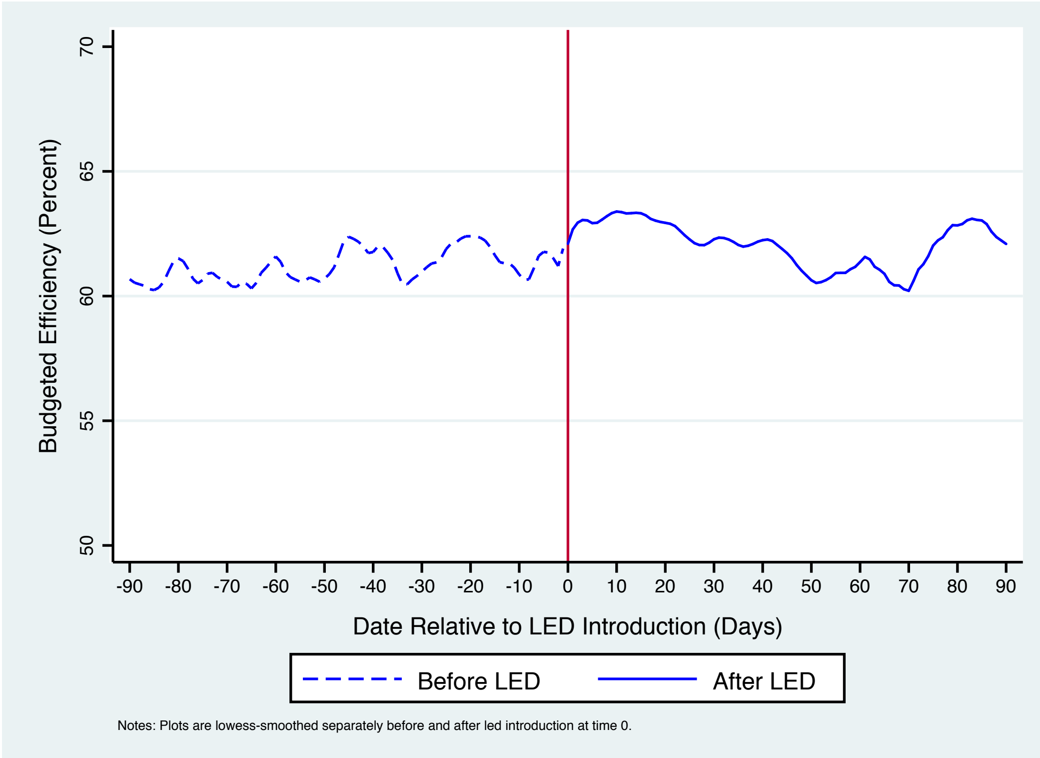

Nevertheless, for the sake of interpretation and external validity, we investigate determinants of the timing of LED introduction. Specifically, we investigate the degree to which standard allowable minutes, budgeted efficiency, scheduled quantity, and target quantity correlate with LED introduction. That is, to the extent that LED replacement is, for example, more likely during lean production times or less likely during the production of large orders from important buyers, these will be reflected in the scheduled quantity and budgeted efficiency for the days leading up to LED introduction.

We first regress SAM on LED introduction to check whether the timing of LED roll-out coincided with the arrival of orders for more or less complex garments. The results of this regression are reported in column 2 of Table 7 and show no evidence of a relationship between SAM and LED roll-out.

We next plot budgeted efficiency against the date relative to LED introduction using data from the quarters before and after LED introduction for each factory unit. This plot is presented in Figure 6 and shows no clear evidence of abnormal trends in budgeted efficiency leading up to LED introduction. Next, we perform the regression analog to Figure 6 by regressing budgeted efficiency on the date relative to LED introduction, again including the usual fixed effects. The results of this regression are reported in column 3 of Table 7 and provide no evidence of any correlation.

FIGURE 7: SCHEDULED QUANTITY BY DATE RELATIVE TO LED INTRODUCTION

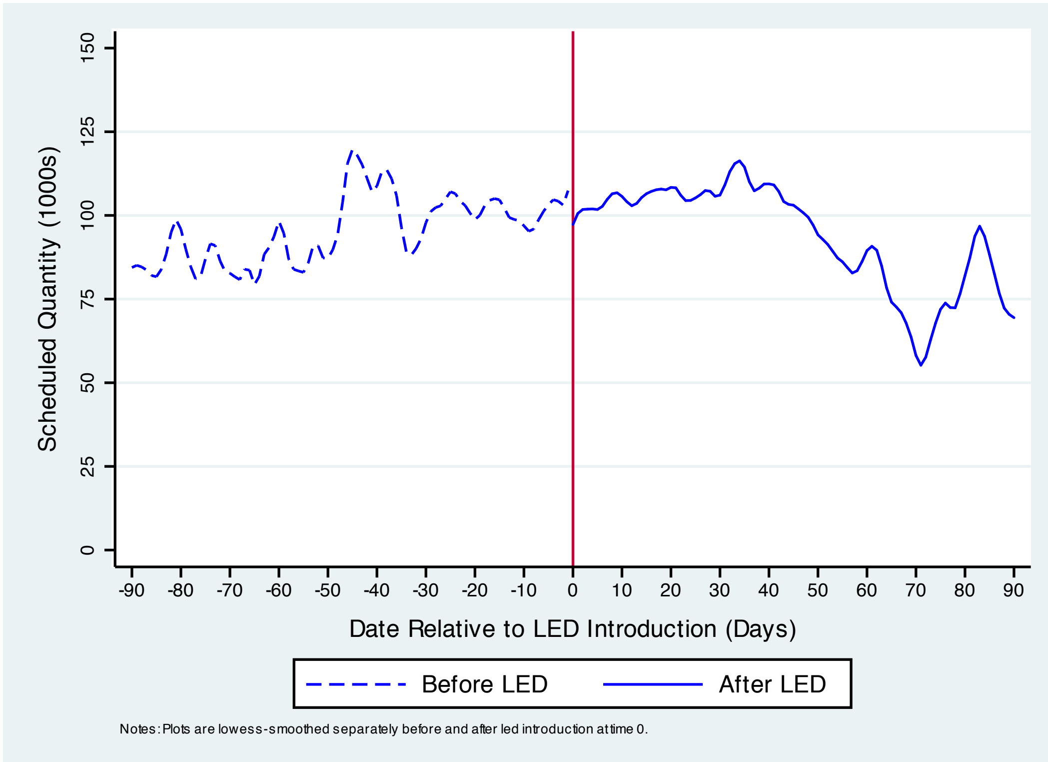

We then repeat both exercises for scheduled quantity. The plot of scheduled quantity against the date relative to LED introduction is presented in Figure 7. The results from the analogous regression are reported in columns 4 of Table 7. Both further support the exogeneity of the introduction of LED lighting.

6 Discussion

The promise of climate change mitigation is tempered by the willingness of individuals and firms to adopt these beneficial technologies on a large scale. This willingness, in turn, is a function of the private returns to adoption, which, for most mitigation strategies, are cited as low or negative even when the public benefits are large. In this study, we show that the introduction of energy-saving LED lighting in Indian garment factories has substantial productivity co-benefits. In particular, the introduction of LEDs eliminates 75% percent of the negative impact of temperature on worker efficiency.

What are the benefit-cost implications for the firm of this productivity impact? From the coefficient estimates in Table 4, we find that, at the average daily heat index of 29.669, LED introduction is associated with an increase in efficiency of -3.87 + (.15 x 29.669) = .58 percentage points.[18] What does this mean to the firm? Senior management at the firm we worked with estimated that the profit gains for each percentage point gain in efficiency were 0.2 percentage points (a fifth of every point gained in efficiency is translated to profit). Thus, at 1,067.58 dollars (USD), the approximate value of profit per factory unit per operating day, LED introduction results in a profit increase of about 41.4 USD per factory unit per operating day, or 12,922.46 USD per factory per year. This is equivalent to about 3.9 percent of daily profit per unit. Put another way, installing LEDs results in each factory unit "gaining" 12 additional days in profit per year.

How does this estimate change the benefit-cost calculations of LED adoption for the firm? To begin with, we obtained energy cost-savings calculations the firm used when making its LED adoption choices. Management estimated that the total energy and operating cost savings per year per factory unit of LEDs (as compared with CFL bulbs, which were being used before LED introduction) were approximately 1,83,000 rupees (INR), or about 3,000 USD. The productivity savings we compute are more than 4 times this amount. The cost of replacing a single factory's bulbs to LEDs is 3,84,000 INR, or about 6,300 USD. Thus, if only energy savings were taken into account, it would take more than 2 years to break even. But when the productivity benefits are included, the firm breaks even within 5 months of LED introduction.

In addition to the private benefits of increased productivity and energy cost savings, the replacement of LED lighting has public benefits of avoided damages due to reduced carbon emissions. On average, the LED replacement saves 18,000 KWh of electricity per factory unit per year, which in this case reduces electricity emissions by about 3.73 tC emissions per unit per year.[19] Valuing this reduction of carbon emissions at the Nordhaus (2008) estimate of $27/tC (a 2005 carbon price) gives us avoided damages of $101.23 per unit per year, and valuing this at the mean value of the review by Tol (2005) of $93/tC yields avoided damages of $345.68 per unit per year. Interestingly, at the current estimates of carbon prices, these benefits are relatively small in comparison to the annual private benefits.[20]

We believe our work is an important first step in quantifying private co-benefits of climate change mitigation strategies, but that much more needs to be done to quantify the full returns to the variety of mitigation strategies. For example, as Knittel and Sandler (2011) suggest, carbon taxes likely have health benefits due to decreases in local air pollution. If consumers internalize these benefits, the effective costs of the tax will be substantially lower. Whether similar co-benefits exist for other types of mitigation – e.g., renewable energy investments, public transport systems, energy-efficient built environments, etc. – is an open and vital question.

[18] We use the coefficient on the impact of heat index on efficiency here because it is the most conservative estimate of the various temperature measures we use.

[19] The conversion from electricity consumption to carbon emissions is done as follows: According to the CO2 Baseline Database for the Indian Power Sector (version 8) by the Central Electricity Authority of India, a MWh of electricity generated on the Southern grid causes 0.76 tCO2 of emissions. Thus, 18,000 KWh causes about 13.68 tCO2, or about 3.73 tC.

[20] Adding the corresponding reduction in local air pollutants would increase the valuation of public benefits, but given the sparsity of accurate data regarding marginal damages of local pollutants in Bangalore, we are unable to include this valuation in this study.

References

Allcott, H. and Taubinsky, D. (2013). The lightbulb paradox: Evidence from two randomized control trials. *working paper*.

Bandiera, O., Barankay, I., and Rasul, I. (2007). Incentives for managers and inequality among workers: evidence from a firm-level experiment. *The Quarterly Journal of Economics*, 122(2):729–773.

Bandiera, O., Barankay, I., and Rasul, I. (2009). Social connections and incentives in the workplace: Evidence from personnel data. *Econometrica*, 77(4):1047–1094.

Bandiera, O., Barankay, I., and Rasul, I. (2010). Social incentives in the workplace. *The Review of Economic Studies*, 77(2):417–458.

Barreca, A., Clay, K., Deschenes, O., Greenstone, M., and Shapiro, J. S. (2013). Adapting to climate change: The remarkable decline in the u.s temperature-mortality relationship over the 20th century.

Bloom, N., Eifert, B., Mahajan, A., McKenzie, D., and Roberts, J. (2013). Does management matter? evidence from india. *The Quarterly Journal of Economics*, 128(1):1–51.

Bloom, N. and Van Reenen, J. (2010). Why do management practices differ across firms and countries? *The Journal of Economic Perspectives*, pages 203–224.

Bollen, J., Guay, B., Jamet, S., and Corfee-Morlot, J. (2009). Co-benefits of climate change mitigation policies: literature review and new results. Technical report, OECD Publishing.

Burgess, R., Deschenes, O., Donaldson, D., and Greenstone, M. (2011). Weather and death in india. *Cambridge, United States: Massachusetts Institute of Technology, Department of Economics. Manuscript*.

Burgess, R., Deschenes, O., Donaldson, D., and Greenstone, M. (2013). The unequal effects of weather and climate change: Evidence from mortality in india.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, 29(2).

Danet, S., Richard, F., Montaye, M., Beauchant, S., Lemaire, B., Graux, C., Cottel, D., Marécaux, N., and Amouyel, P. (1999). Unhealthy effects of atmospheric temperature and pressure on the occurrence of myocardial infarction and coronary deaths a 10-year survey: The lille-world health organization monica project (monitoring trends and determinants in cardiovascular disease). *Circulation*, 100(1):e1–e7.

De Mel, S., McKenzie, D., and Woodruff, C. (2012). Enterprise recovery following natural disasters. *The Economic Journal*, 122(559):64–91.

Deschênes, O. and Greenstone, M. (2007). The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather. *The American Economic Review*, pages 354–385.

Deschênes, O. and Greenstone, M. (2011). Climate change, mortality, and adaptation: evidence from annual fluctuations in weather in the us. *American Economic Journal: Applied Economics*, 3(4):152–185.

Enander, A. E. (1989). Effects of thermal stress on human performance. *Scandinavian journal of work, environment & health*, pages 27–33.

Graff Zivin, J. and Neidell, M. (2010). Temperature and the allocation of time: Implications for climate change. *NBER working paper*, (w15717).

Graff Zivin, J. and Neidell, M. (2012). The impact of pollution on worker productivity. *American Economic Review*, 102:3652–3673.

Guiteras, R. (2009). The impact of climate change on indian agriculture. *Manuscript, Department of Economics, University of Maryland, College Park, Maryland*.

Hancock, P. A., Ross, J. M., and Szalma, J. L. (2007). A meta-analysis of performance response under thermal stressors. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 49(5):851–877.

Hanna, R. and Oliva, P. (2011). The effect of pollution on labor supply: Evidence from a natural experiment in mexico city. Technical report, National Bureau of Economic Research.

Heal, G. and Park, J. (2013). Feeling the heat: Temperature, physiology & the wealth of nations. Technical report, National Bureau of Economic Research.

Hjort, J. (2013). Ethnic divisions and production in firms. Technical report, CESifo Working Paper.

Hsiang, S. M. (2010). Temperatures and cyclones strongly associated with economic production in the caribbean and central america. *Proceedings of the National Academy of Sciences*, 107(35):15367–15372.

IPCC (2013). *Working Group I contribution to the IPCC 5th Assessment Report "Climate Change 2013: The Physical Science Basis". Final Draft Underlying Scientific-Technical Assessment*. Cambridge University Press.

Kala, N., Kurukulasuriya, P., and Mendelsohn, R. (2012). The impact of climate change on agro-ecological zones: evidence from africa. *Environment and Development Economics*, 17(6):663.

Kjellstrom, T., Holmer, I., and Lemke, B. (2009). Workplace heat stress, health and productivity–an increasing challenge for low and middle-income countries during climate change. *Global Health Action*, 2.

Knittel, C. R. and Sandler, R. (2011). Cleaning the bathwater with the baby: The health co-benefits of carbon pricing in transportation. Technical report, National Bureau of Economic Research.

Kudamatsu, M., Persson, T., and Strömberg, D. (2012). *Weather and infant mortality in Africa*. Centre for Economic Policy Research.

Kurukulasuriya, P., Mendelsohn, R., Hassan, R., Benhin, J., Deressa, T., Diop, M., Eid, H. M., Fosu, K. Y., Gbetibouo, G., Jain, S., et al. (2006). Will african agriculture survive climate change? *The World Bank Economic Review*, 20(3):367–388.

Lemke, B. and Kjellstrom, T. (2012). Calculating workplace wbgt from meteorological data: a tool for climate change assessment. *Industrial Health*, 50(4):267–278.

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

Mendelsohn, R., Dinar, A., and Williams, L. (2006). The distributional impact of climate change on rich and poor countries. *Environment and Development Economics*, 11(02):159–178.

Nordhaus, W. (2008). A question of balance.

Parsons, K. (2010). *Human thermal environments: the effects of hot, moderate, and cold environments on human health, comfort and performance*. CRC Press.

Pilcher, J. J., Nadler, E., and Busch, C. (2002). Effects of hot and cold temperature exposure on performance: a meta-analytic review. *Ergonomics*, 45(10):682–698.

Ramsey, J. D., Burford, C. L., Beshir, M. Y., and Jensen, R. C. (1983). Effects of workplace thermal conditions on safe work behavior. *Journal of Safety Research*, 14(3):105–114.

Roberts, M. J. and Tybout, J. R. (1996). Industrial evolution in developing countries: Micro patterns of turnover, productivity, and market structure.

Staritz, C. (2010). *Making the Cut?: Low-income Countries and the Global Clothing Value Chain in a Post-quota and Post-crisis World*. World Bank Publications.

Sudarshan, A. and Tewari, M. (2013). The economic impacts of temperature on industrial productivity: Evidence from indian manufacturing.

Tol, R. S. (2005). The marginal damage costs of carbon dioxide emissions: an assessment of the uncertainties. *Energy policy*, 33(16):2064–2074.

Vanhoorne, M. N., Vanachter, O. V., and De Ridder, M. P. (2006). Occupational health care for the 21st century: from health at work to workers' health. *International Journal of Occupational and Environmental Health*, 12(3):278–285.

A Additional Tables

In appendix table A1, we test the robustness of our main results–that is, the main effect of temperature on efficiency, as well as the interaction of LED introduction and temperature–to the subtraction of a variety of fixed effects. In particular, we explore three alternate specifications. The first specification has no fixed effects, and daily budgeted efficiency is the only control variable. The second specification includes all the time fixed effects, namely year, month, and day of the week. The third specification includes unit by year, month, and day of the week fixed effects (in addition to daily budgeted efficiency as a control variable). We find that the results are very robust to these alternate specifications, and that the coefficients on temperature and the interaction of temperature and LED introduction are quite stable.

In appendix tables A2 and A3, we examine the impact of lagged daily temperature on production efficiency and attendance, respectively. We include the contemporaneous temperature as well as one of seven daily lags in turn. Thus, the first column includes contemporaneous temperature and a one-day lag, the second includes contemporaneous temperature and a two-day lag, and so on, for seven days respectively. Appendix table A2 presents the results of this specification for production efficiency, and table A3 for attendance. We notice that contemporaneous temperature negatively impacts production efficiency regardless of the lag measure used, but has a much smaller and statistically insignificant impact on attendance.

Furthermore, Table A2 shows that one-day and two day lags have a statistically insignificant impact on production efficiency, with lags further back in time, such as five or six day lags showing a more robust negative relationship with production efficiency. Part of the reason why lags closer to the contemporaneous temperatures may not show a significant relationship is that contemporaneous temperature is highly correlated with more recent lagged temperatures, and this correlation diminishes with lags further in time - e.g. the correlation between contemporaneous temperature and one-day lagged temperatures is about 0.887, whereas the relationship between contemporaneous temperature and seven-day lagged temperature is about 0.696.

Table A3 shows that one-day lagged temperature most strongly impacts attendance, likely through thermal stress as well as through its impact on expectations about the temperature the next day. Lags further back in time also have some predictive power regarding attendance, which likely works through the heat stress channel. Overall, the results from tables A2 and A3 support the evidence presented in the main paper.

In Table A4, we consider whether the mitigative impact of LED varies by whether a unit received LED earlier or later. Since we have unit by year fixed effects in all regression specifications, the estimates are not biased by the potential correlation of unit-level unobserved heterogeneity with the introduction of LED. Nevertheless, it is still instructive to examine whether units that adopted LED later exhibit differential impacts of LED on the productivity-temperature gradient. We divide our sample of 29 units into two - the first 14 units to receive LED are considered to have received LED relatively early and the other 15 relatively late.[21] Table A4 illustrates that the mitigative impact of LED does not vary by whether a unit received LED relatively late, since the triple interaction term between the LED dummy variable, temperature, and a dummy that takes the value 1 if the unit received LED relatively late and is 0 otherwise is not statistically significant. The other coefficients of interest - the interaction of the LED dummy and temperature and the main effect of temperature - are very similar to the estimates in Table 5. Temperature has a negative and statistically significant impact on production efficiency, and the introduction of LED mitigates a large portion of this negative relationship.

B Data Appendix

We have daily line-level data from 29 factories in Bangalore. To ensure accurate estimation, we remove extreme outlier values as well as unrepresentative days (such as Sundays) from the dataset. The following factors are taken into consideration when deciding our final sample.

- We create a measure of the difference between the maximum and minimum date for which each line is observed divided by the total number of days for which it is observed. This measure essentially captures the proportion of time for which a line is observed relative to its time in the data. We remove observations for which this proportion is strictly greater than 1 (10 lines) and less than 0.38, which is the 5th percentile of the observations. This is done to ensure that the sample includes lines that are consistently producing, not ad hoc lines that are sometimes set up to fulfil rush orders or orders behind schedule.

- We remove lines observed greater than twice a day, about 0.6% of our observations, since these are likely coding errors. While it is possible that a line finished a set of orders and moved onto producing a different style of garment midway through the day, it is not possible that a line finished several sets of garment orders in a single day, since orders are usually for hundreds or thousands of garments per order. For lines that are observed more than twice a day, we consider mean actual efficiency and mean budgeted efficiency across the two styles produced that day.

- We remove extreme outliers from the efficiency and quantity produced. We consider all observations between the 5th and 95th percentile of the efficiency measurements, which ranges from about 3% to about 111%. After removing these values, we also trim physical quantity produced at the 1st and 99th percentile. These decisions were taken following meetings with the Industrial Engineering experts at the factory regarding what constitutes feasible values of output and efficiency.

- Finally, we remove Sundays as well as days which have very few lines, since these are likely regional holidays or days of unrepresentatively low productivity. We consider only days for

[21] The latter measures includes four units that did not have LED throughout our sample period.

which 117 or more lines are present in the data (the 1st percentile of observations).