

NBER WORKING PAPER SERIES

THE LIGHT AND THE HEAT:
PRODUCTIVITY CO-BENEFITS OF ENERGY-SAVING TECHNOLOGY

Achyuta Adhvaryu
Namrata Kala
Anant Nyshadham

Working Paper 24314
<http://www.nber.org/papers/w24314>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2018

We are incredibly thankful to Anant Ahuja, Chitra Ramdas, Shridatta Veera, Manju Rajesh, Raghuram Nayaka, Sudhakar Bheemaroo, Paul Ouseph, and Subhash Tiwari for their coordination, enthusiasm, support, and guidance. Thanks to Prashant Bharadwaj, Michael Boozer, Rahul Deb, Josh Graff Zivin, Catherine Hausman, Tom Lyon, Robyn Meeks, Nick Ryan, Antoinette Schoar, Tavneet Suri, and Joseph Shapiro as well as seminar participants at Yale, MIT, Michigan, Carnegie Mellon, CEGA, the NBER, the IGC, the World Bank, PacDev, and NEUDC for helpful comments and suggestions. Robert Fletcher and Karry Lu provided excellent research assistance. Adhvaryu gratefully acknowledges funding from the NIH/NICHD (5K01HD071949), and all authors acknowledge funding from PEDL, a CEPR-sponsored grant initiative. All errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Achyuta Adhvaryu, Namrata Kala, and Anant Nyshadham. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Light and the Heat: Productivity Co-benefits of Energy-saving Technology
Achyuta Adhvaryu, Namrata Kala, and Anant Nyshadham
NBER Working Paper No. 24314
February 2018
JEL No. J24,O14,Q56

ABSTRACT

Measurement of the full costs and benefits of energy-saving technologies is often difficult, confounding adoption decisions. We study consequences of the adoption of energy-efficient LED lighting in garment factories around Bangalore, India. We combine daily production line-level data with weather data and estimate a negative, nonlinear productivity-temperature gradient. We find that LED lighting, which emits less heat than conventional bulbs, decreases the temperature on factory floors, and thus raises productivity, particularly on hot days. Using the firm's costing data, we estimate the pay-back period for LED adoption is nearly one-sixth the length after accounting for productivity co-benefits.

Achyuta Adhvaryu
Ross School of Business
University of Michigan
701 Tappan Street
Ann Arbor, MI 48109
and NBER
adhvaryu@umich.edu

Anant Nyshadham
Department of Economics
Boston College
Maloney Hall, 324
Chestnut Hill, MA 02467
and NBER
nyshadha@bc.edu

Namrata Kala
Harvard University
27 Hillhouse Ave
New Haven, CT 06511
namrata.kala@gmail.com

1 Introduction

Innovations in energy efficiency have been cited as a primary means to curb the acceleration of climate change (Granade et al., 2009). Despite this promise, energy efficient technologies are consistently adopted at low rates (Allcott and Taubinsky, 2015).¹ Given the repercussions of rising global temperatures due to climate change (IPCC, 2013), and the startling rate of growth of global energy demand (Wolfram et al., 2012), achieving high adoption rates of technologies that mitigate climate change is a key policy priority.² In this study, we estimate the productivity consequences of the adoption of energy-saving technology, using daily production line-level data from a large garment firm operating factory units in and around Bangalore, India. We show that the introduction of light-emitting diode (LED) technology on factory floors substantially attenuates the negative relationship between temperature and productivity, measured here as production line efficiency (realized output over target output).³

LED lighting modulates the temperature-productivity gradient through reduced heat dissipation: LED technology, in addition to being 7 times more energy-efficient than standard fluorescent lighting in our setting, also emits about one-seventh the heat. We study the impacts of the staggered roll-out of LEDs over more than three years on the sewing floors of 26 garment factories.⁴ The switch to LED lighting was largely driven by changes in international buyers' recommendations regarding environmental sustainability for their suppliers. We demonstrate in a variety of checks that the timing of the roll-out across factory units was not systematically related to temperature, nor to a variety of business processes.

We estimate the extent to which the introduction of LED lighting, through the reduced dissipation

¹Recent studies point to information frictions, or a lack of salience of information, as key determinants of this "efficiency gap": if individuals and firms knew the true returns to investment in energy efficiency, or if information were made more salient, widespread adoption of these technologies would occur more quickly (Allcott and Greenstone, 2012). It may also be that low adoption is simply a result of the fact that returns are smaller, or costs higher, in practice than engineering projections predict (Fowlie et al., 2015; Ryan, 2017).

²Economic productivity is projected to suffer, not only due to the increased frequency of extreme weather events (see, e.g., Dell et al. (2012); Deschênes and Greenstone (2007); Guiteras (2009); Hsiang (2010); Kala et al. (2012); Kurukulasuriya et al. (2006); Lobell et al. (2011); Parker (2000)), but also because excessive heat increases health risks (Burgess et al., 2011, 2014; 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 et al., 2015).

³Impacts of temperature are highly nonlinear: for outdoor wet-bulb temperatures below 19°C (the dry-bulb equivalent at average humidity levels in our sample is 27-28°C), temperature has a very small impact on efficiency. But for mean daily temperatures above this cutoff (about one quarter of production days), there is a large negative impact on efficiency of approximately 2 efficiency points per degree Celsius increase in temperature. This nonlinear gradient is remarkably consistent with the physiology of temperature effects: at high ambient temperatures, the body loses the ability to dissipate heat, and begins the process of heat absorption, which negatively affects performance (Hancock et al., 2007).

⁴Our data include 30 factories (all owned by the same garment firm), four of which did not receive LED lighting.

of heat on factory floors, flattens the temperature-efficiency gradient. Specifically, LED installation has no impact on the gradient for wet-bulb temperatures below the 19°C wet bulb globe temperature (WBGT) cutoff, but attenuates the negative slope of the gradient by more than 80 percent for temperatures above this threshold. The reason that LED installation flattens only the top of the temperature-productivity gradient has to do with the nonlinear nature of the gradient itself. An engineering study we commissioned found that bulb replacement with LEDs likely led to a reduction of indoor temperature by about 2.4°C (which is about 1.42 °C in wet bulb globe temperature (WBGT), the measure of temperature we use), and that this reduction was approximately constant across the temperature distribution. The reason mitigation through LED installation was larger where the temperature-productivity gradient was steeper is then made clear: the introduction of LEDs constituted a movement leftward along the gradient, and this movement generates large increases in efficiency in high temperature ranges, and small efficiency increases elsewhere.

We present these results in a variety of ways. Our baseline specification uses linear splines with a knot at the wet-bulb temperature of 19°C. We also estimate semi-parametric models that allow for flexibility in the impact of temperature on efficiency before v. after LED installation. We then difference across these estimated functions within 0.1°C bins to calculate the gradient difference at each point along the temperature distribution (along with bin-specific standard errors), which yields the impact of LED at each 1/10th degree. These impact estimates are quite consistent with the linear spline results, showing larger LED impacts at higher temperatures. We then combine these estimates with the distribution of degree-days over a one-year period in our data to construct a probability-weighted average semi-parametric treatment effect across the temperature distribution. This estimate, approximately 0.723 points (and statistically significant), tells us the *average increase* in efficiency after the LED introduction.

Last, we conduct an event study analysis in which we compute this weighted average treatment effect in the months immediately preceding and immediately following LED installation. The event study results corroborate our main findings: prior to LED introduction, the average efficiency difference across LED and non-LED factory units is very small, but starting immediately on the month of installation, there is a large, sustained efficiency difference.

Finally, we perform cost-benefit calculations for LED adoption, combining the above estimates with the firm's actual costing data for LED replacement and actual energy savings over CFL lighting. The results of this analysis show that the productivity co-benefits of LED adoption are substantially larger

than the energy savings. Indeed, accounting for productivity increases dramatically shifts the break-even point for the firm, from over three and half years to less than eight months. The firm was unaware of these potential productivity effects of LED adoption at the time of the switch, but has subsequently adopted a policy of installing LEDs and other efficient lighting technologies in all new factories in part due to these large results.

Our study contributes to the literature on the returns to climate change mitigation. A related literature has established patterns of adaptation to climate change and the returns to this adaptation (e.g. Barreca et al. (2016)). The few recent studies that examine “co-benefits,” or additional gains, of mitigation focus largely on the indirect public returns (see IPCC (2013) for a review). Our study examines a novel, *private* co-benefit of climate change mitigation. This distinction is important because the success of most mitigation strategies rests 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, ignoring the productivity benefits of LEDs would seriously underestimate the private returns to adoption (by about five-fold). Moreover, while engineering estimates of the heat dissipation of LED (vis-a-vis CFL) bulbs exist (indeed we commissioned such an estimate for this paper), these estimates are not always perfectly predictive of reality; Fowlie et al. (2015) is a recent prime example of this. Our study evaluates a context in which we are able to test directly whether productivity impacts exist, as engineering estimates, coupled with existing evidence of the temperature-productivity gradient, would predict.

We also contribute to the understanding of the effects of environmental and infrastructural factors (which are often related to the environment) on productivity in developing countries (Adhvaryu et al., 2016; Allcott et al., 2014; Hsiang, 2010; Sudarshan et al., 2015).⁵⁶ Indeed, the impacts of temperature on productivity appear to hold quite consistently across countries and time (Burke et al., 2015; Dell et al., 2012). Our results corroborate what these studies have found, and highlight an interaction between high temperatures and the co-benefits of energy-efficient technologies.

The remainder of the paper is organized as follows. Section 2 describes contextual details regarding

⁵Several recent studies document this relationship in more developed settings (Chang et al., 2014; Costinot et al., 2016; Graff Zivin and Neidell, 2012; Hanna and Oliva, 2016).

⁶More broadly, our paper fits into the literature on 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), inter- and intra-firm networks (Bandiera et al., 2009, 2010; Cai and Szeidl, 2016), incentive structures (Bandiera et al., 2007), and ethnic boundaries (Hjort, 2014) all significantly impact productivity.

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 reviews the cost-benefit analysis and concludes.

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 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 stages of garment production: cutting, sewing, and finishing. First, 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 the second stage, 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.

Finally, in the finishing stage, garments are checked, ironed, and packed. The majority of quality checking is done “in-line” on the sewing floor, but final checking occurs during the finishing stage. Garments with quality issues are sent back to the sewing floor for re-work or, in rare cases, 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; Seppanen et al., 2006), 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 psychomotor ability and degraded perceptual task performance (Hancock et al., 2007). The impact on individual subjects 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).

One key finding from this literature is that there is a non-monotonic relationship between ambient temperature and human performance. The overall shape of the relationship is an inverse-U: performance suffers at excessively cold and excessively warm temperatures (Parsons, 2010). Moreover, one meta-analysis highlights the dry-bulb threshold of 85°F (29.4°C) as particularly important (Hancock et al., 2007). This threshold value represents the temperature above which the body performs obligatory heat storage. As Hancock et al. (2007) put it, “[in] this circumstance, although the individual is dissipating heat at the maximal rate, he or she experiences a dynamic increase in core body temperature” (p. 860). In line with this physiology, measured effects on performance are larger for temperatures above the 85°F threshold. To our knowledge, mitigation of temperature effects is not explored extensively in this literature, beyond an emphasis on the substantial variation in effect size across stud-

⁷In 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 varies slightly by style or factory.

ies (Hancock et al., 2007). More recently, Sudarshan et al. (2015) study temperature mitigation in an industrial setting using air washers. Our contribution is to highlight the productivity effects of energy efficient technology (LED lighting) which are driven by the temperature-productivity gradient.⁸

2.4 LED v. Fluorescent Lighting

The LED light bulbs that replaced the fluorescent bulbs in the factories in our data are approximately 7 times as energy-efficient (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).⁹ Heat emissions for LEDs are substantially lower than fluorescent bulbs: the average LED bulb emits 3.4 Btus, as compared to 23.8 Btus for the fluorescent lighting in the setting we study.¹⁰

2.5 LED Roll-out: Summary and Timeline

The factories began installing LED lighting in October 2009 and completed the installations by February 2013. According to senior management at the firm, 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.¹¹

⁸Air conditioning, which would mitigate the temperature-productivity gradient altogether is extremely rare in the Indian manufacturing sector (ISHRAE, 2015).

⁹Note 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. Also, it should be noted that many varieties of LED and fluorescent bulbs exist. The energy and lighting specifications and calculations presented and discussed in this paper are specific to the bulbs involved in the factory replacements in our data and will not represent universal comparisons. Accordingly, generalizing our findings would require an understanding of how bulb specifics might differ from those used in this empirical context.

¹⁰Changing factory lighting may have consequences for productivity through mechanisms other than temperature changes, as highlighted by the results of the original Hawthorne lighting experiment (Mayo et al., 1939; Snow, 1927), as well as new analysis by Levitt and List (2011). Our analysis allows for this possibility by including the main effect of LED installation, but we find limited evidence for productivity changes through mechanisms other than temperature changes. This is not altogether surprising given the degree of care and attention placed on lighting conditions in the garment production setting. Senior management emphasized that the lighting replacement was designed such that light quantity and quality at the point of production operation would remain within the strict industry and buyer guidelines before and after the replacement.

¹¹This process, of course, still leaves room for endogeneity in the timing of LED adoption across factory units. We check explicitly for this endogeneity in Table 7, and find little evidence that LED adoption at the unit level was correlated with a

The replacement took the form of substituting 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 roughly 1,000 fluorescent lights consuming 7 W each with LED lights of 1W each.¹² 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 30 units from which we have productivity data, LED replacements occurred in 26 units during the observation period. Since our productivity data ranges from April 2010 to June 2013, some units already have LEDs at the beginning of our productivity data, and all but four units have LED by the end of our sample period.¹³

3 Data

Here we provide an overview of data sources, describe the variables of interest, and present summary statistics.

3.1 Weather Data

We use daily temperature, precipitation and relative humidity data from The National Centers for Environmental Prediction Climate Forecast System Reanalysis (CFSR) (Saha et al., 2010). The CFSR data is a re-analysis dataset that uses historical station-level and satellite data combined with climate models to produce a consistent record of gridded weather variables from 1979 to the present. It has a spatial resolution of about 38 km, and each factory is matched to the nearest data grid point. These data provide daily weather data at a fine spatial level, and are therefore preferable to station level data in India which is not always consistently available at the daily level.

We present our results using a variety of temperature indices, two that incorporate temperature
variety of business operations and outcomes.

¹²The 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 about 1,000. Replacing overhead CFL lights with machine mounted LED lights implies that while the light dissipates less heat, it is also now closer to the worker than before, which might partially offset the impacts of heat dissipation - our results however indicate that the effects of lower heat dissipation dominate the impacts on productivity.

¹³Regression results that omit units that had LED lighting at the start of the sample period or did not receive LED lighting by the end of the sample period yield very similar estimates, and are reported in appendix tables A10-A12.

and humidity into an index and the third using (dry-bulb) temperature controlling for humidity. We present our results with Wet Bulb Globe Temperature in the body of the paper, and the results with the other two measures of temperature in the appendix. We incorporate relative humidity into the temperature measure because the effect of relative humidity on thermal comfort may vary with temperature, by affecting evaporative heat loss from the human body (Jing et al., 2013), but also show that our results hold with dry bulb temperature. With mean daily temperature and relative humidity data, we construct the Wet Bulb Globe Temperature measure that is suitable for indoor exposure (that does not take into account wind or sunlight exposure, since that is not applicable in this context). The formula is taken from Lemke and Kjellstrom (2012), and is given by:

$$WBGT = 0.567T_d + 0.216 \left(\frac{rh}{100} * 6.105 \exp \left(\frac{17.27T_d}{237.7 + T_d} \right) \right) + 3.38. \quad (1)$$

where T_d = dry bulb temperature in Fahrenheit and rh = relative humidity (%).¹⁴

Note that the weather data we are using include only daily outdoor temperature measures. Of course, indoor temperature in the factory would likely be the most impactful for worker productivity; however, we do not have data on indoor temperature from the time period over which the LED roll-out occurred. Accordingly, we use outdoor ambient temperature as discussed above as a proxy for indoor conditions. In order for outdoor temperature to represent a valid proxy, we would like to verify that fluctuations in outdoor temperature pass through to indoor temperature. Although we do not have indoor temperature data from the study period, we did collect roughly a year's worth of indoor and outdoor temperature from two factories and six months of data from a third factory after the study period.¹⁵

¹⁴We also calculate an alternative measure, the Heat Index (HI), that is calculated based on the formula:

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

The formula for the calculation is derived from the Rothfus regression that replicates the HI values from Steadman (1979).

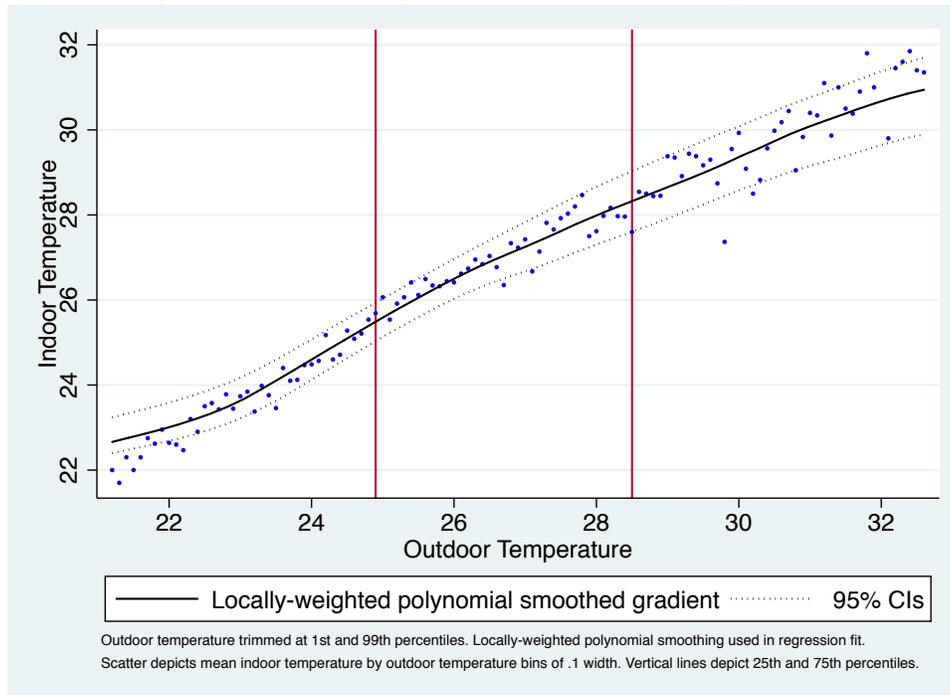
For about 0.6% of our data, the relative humidity is greater than 85% and daily temperature ranges between 80 and 87 degrees Fahrenheit, and the following adjustment is applied:

$$HI = HI + [(rh - 85)/10] * [(87 - T_d)/5] \quad (3)$$

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. For all our results, we report the main effect of WBGT. In the appendix, we report the results corresponding to specifications using dry bulb temperature controlling for relative humidity, as well as the heat index as evidence of robustness.

¹⁵We collected data from 22nd September 2014 to 11th August 2015 in one factory, from 27th September 2014 to 10th August 2015 in a second factory, and from 28th January 2015 to 10th August 2015 in a third factory.

Figure 1: Indoor Temperature vs. Outdoor Temperature



In Figure 1, we plot mean indoor temperature values for each .1 degree bin of outdoor temperature along with a local polynomial regression fit curve and 95 percent confidence intervals. Indoor temperature appears to be a linear function of outdoor temperature with a slope of roughly 0.79. That is, there appears to be large but not perfect pass through of outdoor temperature fluctuations to indoor temperature and this relationship appears to be constant for all levels of outdoor temperature. A positive intercept indicates that at lower outdoor temperature levels (e.g., 22 degrees Celsius wet bulb globe) the indoor temperature is slightly higher than the outdoor temperature reflecting a flow source of heat inside the factory independent of outdoor temperature (e.g., lighting and machinery). Furthermore, a regression of indoor temperature on outdoor temperature has an r-squared of about 0.84, implying that a very large amount of the variation in indoor temperature is explained by the variation in outdoor temperature.

Note that these data were collected after the introduction of LED in the factories, and therefore, depict the *ex post* relationship between indoor and outdoor temperature. Engineering calculations based on building and lighting specifications provided by an industry consultant suggest that LED introduction would have dropped the intercept of this relationship by about 2.4 degrees Celsius on average across factories with some variation due to the quantity of lights replaced and building size

and materials for each factory.¹⁶ This impact should be generally constant across the distribution of outdoor temperature depending on such factors as ventilation.¹⁷ Nevertheless, in what follows we do not impose a functional form on the impact of LED introduction on the relationship between indoor and outdoor temperature, but rather allow the data to determine the change in shape of the observed productivity-temperature relationships before and after LED introduction.

3.2 Factory Data

We use data on line-level daily production from 30 garment factories in and around Bangalore, India. Identifiers include factory unit number and production 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 largely derives from a central database of styles, with potential amendments by the factory’s IE department during “sampling.”¹⁸

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, with all factory units in our data operating a single day-time production shift. Accordingly, a line producing a style with a SAM of .5 will have a target of 120 garments per hour, or 960 garments per day. Most importantly, the target quantity is almost always fixed across days (and 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. Depending on the order size (or “scheduled quantity”) for a style, multiple lines may produce the same style at one time.¹⁹ Variations in expected average efficiency over the life of a particular garment order due to order size are reflected in the budgeted efficiency. Budgeted efficiency remains fixed for a given line over the life of a particular order. However, actual efficiency of a given style will vary systematically across lines and within line over time due to absenteeism, machine failures, working conditions, etc. We are, of course, interested in

¹⁶This implies a 0.8 standard deviation drop in temperature, since as indicated in Table 1, the standard deviation of temperature is 2.96 degree Celsius.

¹⁷Consultant report available upon request.

¹⁸Sampling is the process by which a style that is ordered by a buyer is costed in terms of labor and production time. So-called sampling tailors (highly trained) make a garment of a particular style entirely and recommend any alterations to the SAM for that style to the IE department.

¹⁹Indeed, in our data, lines produce styles for between 1 and 268 days.

these deviations of actual efficiency from expected or budgeted efficiency due to transitory temperature. We will accordingly control for budgeted efficiency and include line fixed effects in the regression analysis below.

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. Without normalizing production observations to target quantity and accounting for budgeted efficiency, 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 order size varied by temperature due to seasonal buying of winter garments at certain times in the fall months, resulting variations in efficiency could be attributed to temperature spuriously. Accordingly, we argue that actual efficiency, controlling for budgeted efficiency, is the 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. We check this explicitly in the empirical analysis below. 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, predictable variation in expected efficiency over the life of an order arises due to the interaction of order size and learning by doing, with lines producing larger orders of the same garment style achieving higher maximum (and therefore, average) efficiency than those producing smaller orders. True daily fluctuations in productivity are, therefore, best measured by actual efficiency controlling for budgeted efficiency.²⁰

3.3 Summary Statistics

We present means and standard deviations of variables used in the analysis in Table 1 below. Our sample consists of 523 production lines across 30 factory units. The range of dates over which we have production data spans 1,001 days in total. However, we do not observe all factory units, nor all lines within a unit, for all dates.²¹ Altogether, our data includes nearly 240,000 line x day observations. Roughly, one-third of the observations correspond to days in factory units prior to the introduction of

²⁰We check the sensitivity of all main results to alternate definitions of the productivity outcome and find the results to be robust.

²¹Appendix table A3 tests that production line-day observations for which data is missing are not correlated with either temperature or the LED installation decisions.

Table 1
Summary Statistics: Weather, Production, and LED Introduction

Number of line-day observations	239,680	
Number of lines	523	
Number of days	1,001	
Number of units	30	
	Mean	SD
<i>Weather</i>		
Temperature (Celsius)	24.353	2.966
Relative Humidity (%)	0.647	0.174
Heat Index (Celsius)	23.128	2.871
Wet Bulb Globe Temperature (Celsius)	17.230	1.683
<i>Production</i>		
Actual Efficiency	55.234	26.233
Budgeted Efficiency	61.981	11.545
Standard Allowable Minutes (SAM)	0.724	2.445
<i>Attendance</i>		
1(Present for Full Work Day)	0.843	0.363

LED lighting and the remainder are post-LED observations.

4 Empirical Strategy

In this section, we provide preliminary graphical evidence on the shape of the temperature-productivity gradient, the effects of LED introduction, and the persistence of this evidence after accounting for various unobservables. We then leverage these motivating facts in developing a two stage empirical strategy to flexibly estimate the impact of LED introduction on productivity as mediated through ambient temperature.

4.1 Preliminary Graphical Evidence

We begin by motivating the empirical specifications and techniques with descriptive plots of production and temperature data.

4.1.1 Productivity-Temperature Gradient

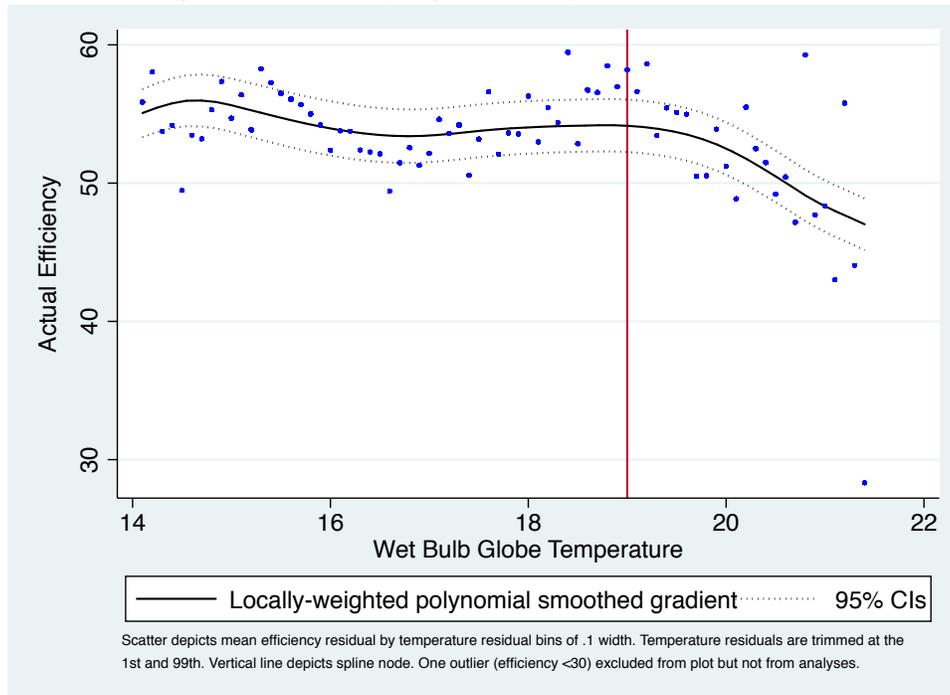
Since we intend to estimate how LED introduction impacts the relationship between efficiency and temperature, we must first understand the nature of this relationship. Accordingly, we first investigate the raw relationship between efficiency and wet bulb temperature in the data prior to LED introduction. Figure 2 presents a scatter plot of the average efficiency for each 0.1 degree bin of wet bulb temperature observed in the data. We also include in the figure a local polynomial smoothed fit and 95 percent confidence intervals like those depicted in Figure 1. Figure 2 shows that, in the absence of LED lighting, indeed efficiency appears to be a decreasing function of temperature, and this relationship is quite nonlinear with the strongest declines in efficiency occurring at the highest wet bulb temperatures. Specifically, the gradient goes from modestly decreasing to strongly decreasing to the right of the vertical line in Figure 2. This vertical line, denoting 19 degrees Celsius in wet bulb temperature, represents a strong break in the slope. Accordingly, in the parametric regression analysis proposed below, we specify a linear spline with a node at 19 to capture this dichotomous slope in the gradient.

Notably, a wet bulb globe outdoor temperature of 19 degrees Celsius corresponds in our data to an outdoor ambient dry bulb temperature of roughly 27.5 degrees Celsius and is likely equivalent to an indoor dry bulb temperature of roughly 29.5 degrees before LED introduction.²² This 29.5 degree dry bulb temperature is remarkably consistent with estimates from previous studies on the physiological threshold for the absorption of heat into the body above which temperature is more impactful for human functioning (Hancock et al., 2007).²³

²²This approximate relationship is derived from the indoor-outdoor temperature we collected and the engineering study of LED installation we commissioned.

²³Of course, this threshold applies to ambient indoor temperatures. We collected a small sample of indoor temperatures to calculate an indoor-outdoor temperature gradient (presented in Figure 1), and found that at 27 degrees Celsius, post-LED installation, the temperature indoors is roughly the same as outdoors. Prior to LED installation, according to estimates from the engineering study we commissioned, this differential would have been about 2.4 degrees larger. Thus at outdoor dry-bulb temperatures of roughly 27 degrees C, prior to LED installation the temperature indoors would have been about 29.4 degrees C, which is squarely in the range of the physiological threshold value.

Figure 2: Efficiency Against Temperature (Pre-LED)



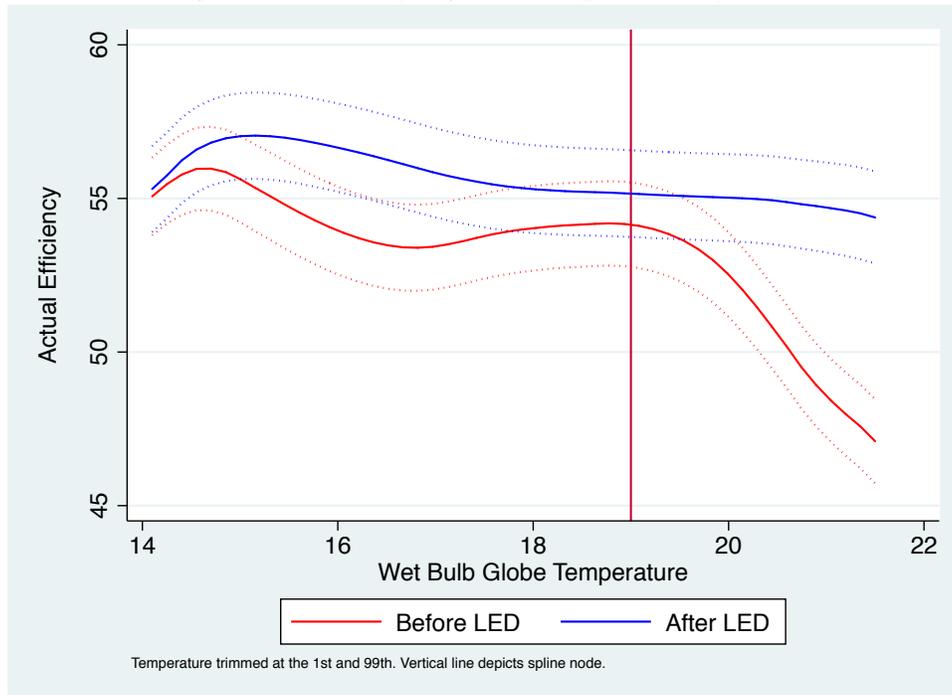
4.1.2 Impacts of LED Introduction

Having established the shape of the temperature-productivity gradient for the garment factories in our data before the introduction of LED, 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 3.²⁴ The evidence suggests that factories are more efficient at all temperatures after the LED introduction, but this gain (or attenuation) is increasing at high temperatures. That is, the pre-LED gradient (red line) in Figure 3, of course, replicates the non-linear shape depicted in Figure 2, but the post-led gradient exhibits a flatter slope to the right of the 19 degree vertical line allowing the gap between the before and after LED gradients to widen at higher temperatures and indicating a persistently significant treatment effect above 19 degrees.

As mentioned above in section 3, the engineering calculations for the impact of LED introduction on indoor temperature indicated that post-LED indoor temperatures should be around 2.4 degrees lower than would have prevailed at a given outdoor temperature before the introduction of LED. This

²⁴Note in each figure from here onwards in the paper with both pre- and post-LED plots, we show 83% confidence intervals, which allow the reader to visually assess the hypothesis of a difference between the two curves – if the confidence intervals do not overlap at a given point, then the two curves are significantly different at the 5% level at that point.

Figure 3: Efficiency Against Temperature by LED



would translate into a shift to the right of the efficiency-temperature gradient in Figure 2 after the introduction of LED, as each outdoor temperature on the x-axis corresponds to a lower indoor temperature. Notably, the difference between the before and after LED gradients could be explained by a shifting of the pre-LED gradient a few degrees to the right and truncating the right tail at around 21.5 degrees wet bulb which is the boundary of the support of the underlying temperature distribution. In any case, since we are interested in the average or total impact on efficiency of LED introduction given the observed temperature distribution, the exact structural relationship between outdoor temperature, LED, indoor temperature and subsequently efficiency is not required, nor is it feasible for us to estimate due to data limitations. Rather, we can measure empirically the difference in slopes of the efficiency-temperature gradients before and after LED, allowing for slope changes at 19 degrees, by estimating the parametric spline regressions proposed below. We can also calculate average impacts of LED introduction on efficiency from semiparametric estimated impacts at each point along the temperature distribution and weighted by the probability that each temperature value prevails. These two strategies are described in detail below.

4.2 Parametric Spline Regression Analysis

Motivated by the preliminary graphical evidence above, we set forth a more rigorous regression analysis below to causally identify both the effect of temperature on production efficiency at various points along the temperature distribution 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 along with the non-linearities depicted in Figures 2 and 3 above.

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

$$E_{ludmy} = \alpha_0 + \beta^L T_{dgm_y}^L + \beta^H T_{dgm_y}^H + \phi B_{ludmy} + \alpha_l + \gamma_{uy} + \eta_{um} + \delta_d + \varepsilon_{ludmy}. \quad (4)$$

Here, E is actual 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^L is daily wet bulb globe temperature from grid point g in degrees Celsius up to the spline node of 19, above which it records a constant 19; T^H is daily wet bulb temperature minus 19 degrees Celsius from grid point g above the spline node, below which it records a constant 0; α_l are production line fixed effects; γ_{uy} are unit x year fixed effects; η_{um} are unit x month fixed effects; δ_d are day-of-week fixed effects; and α_0 is an intercept. β^L and β^H are the coefficients of interest, giving the impact of a 1-degree Celsius increase in wet bulb globe temperature on line-level efficiency for temperatures below and above 19 degrees, respectively.

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

$$\begin{aligned} E_{ludmy} = & \alpha_0 + \beta_1^L (T_{dgm_y}^L \times LED_{umy}) + \beta_1^H (T_{dgm_y}^H \times LED_{umy}) + \beta_2 LED_{umy} \\ & + \beta_3^L T_{dgm_y}^L + \beta_3^H T_{dgm_y}^H + \phi B_{ludmy} + \alpha_l + \gamma_{uy} + \eta_{um} + \delta_d + \varepsilon_{ludmy}. \end{aligned} \quad (5)$$

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 of LED introduction in a particular factory unit. The coefficients of interest in the above specification are β_1^L , β_1^H , β_3^L and β_3^H . β_3^L and β_3^H indicate the effect of temperature on productivity below and above the 19 degree spline node, respectively, *before* LED introduction. β_1^L and

β_1^H are the extent of attenuation of the temperature-productivity gradient below and above the 19 degree spline node, respectively, once LED lighting is introduced. The sums $\beta_1^L + \beta_3^L$ and $\beta_1^H + \beta_3^H$ gives the net effect of temperature on productivity below and above the spline node, respectively, following LED introduction. Note that we choose this spline specification with a single node at 19 degrees WBGT for two reasons: 1) the raw data plots in Figures 2 and 3 clearly show that the relationship between temperature and efficiency (and the difference in this relationship across LED) changes at this point in the temperature distribution and does not vary much on either side of this cutoff; and 2) this point corresponds remarkably well to previous studies of the physiology of heat stress (Hancock et al., 2007).²⁵

In order to account for common error distributions at the factory level over time, standard errors are clustered at the unit level. This cluster structure is appropriate given that LED introduction occurs at the unit level. However, given the relatively small number of clusters (30), we employ wild cluster bootstrap inference and report calculated p -values in parentheses in all tables unless otherwise noted.²⁶

4.2.1 Attendance

We also estimate the same specifications presented in equations 4 and 5, but replacing the efficiency outcome on the left hand side with mean attendance (or probability of each worker being present in the factory) at the line-daily level. These regressions are intended to investigate the degree to which temperature impacts on efficiency and the corresponding attenuation from LED introduction might be working through impacts on worker attendance. We also estimate the original efficiency specifications from equations 4 and 5, but with the inclusion of mean line-daily worker attendance as an additional control. The combination of these two sets of results allow us to investigate whether temperature and LED introduction indeed have impacts on worker attendance and whether controlling for any impacts on attendance changes the estimated impacts of temperature and LED on the primary outcome of interest (efficiency).

²⁵Nevertheless, we explored more flexible spline specifications with more nodes and found the results to be qualitatively identical with less precision.

²⁶See Cameron et al. (2008) for a thorough treatment of clustering approaches with few clusters and a discussion of their relative performance, which highlights that wild cluster bootstrap inference works best in a setting with few clusters. We report p -values in all regressions estimated via the wild cluster bootstrap since the estimation in Stata reports p -values.

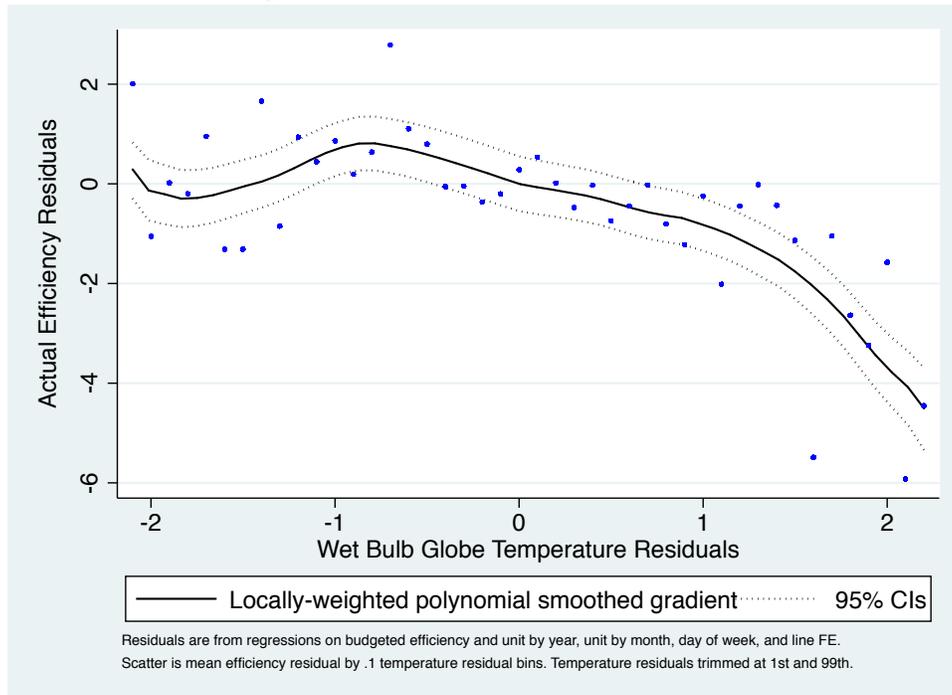
4.2.2 Distributed Lags

It should be noted that daily temperature could potentially reflect short-term serial correlation that might pose a challenge for identifying the impacts of contemporaneous exposure to temperature. Following previous studies, we augment both equations 4 and 5 to include 7-day distributed lag spline terms and their interactions with LED in addition to the contemporaneous spline and LED interaction terms of primary interest. In the distributed lag models, we interpret the coefficients on contemporaneous spline and interaction terms as the *incremental* impacts of contemporaneous temperature exposure after controlling for any persistent impacts of lagged exposure. This allows for the isolation of the impact of contemporaneous exposure from any persistent impacts of lagged exposure. If the coefficient on the contemporaneous temperature terms are similar with and without the inclusion of the 7-day distributed lag terms, we interpret the results as indicative of a minimal role for serial correlation and persistence in impacts of lagged exposures. On the other hand, we can recover the composite impact of both the incremental innovation in contemporaneous temperature exposure and the persistent impacts of lagged exposures by summing up the coefficients from contemporaneous temperature and the full set of lagged exposures, but this composite impact will be nearly identical to that estimated from the original specification presented in equation 4 and 5 as the set of relevant lagged temperatures included grows.

4.2.3 Controls and Unobservables

Note that all specifications include as controls budgeted efficiency, line fixed effects, year fixed effects, factory unit x calendar month fixed effects, and day-of-the-week fixed effects. As mentioned in section 3, budgeted efficiency accounts for expected variation in achievable efficiency due to order size and learning by doing on the line, but the remaining controls are meant to account for various unobservable determinants of efficiency that might correlate with temperature. 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

Figure 4: Residual Gradient (Pre-LED)

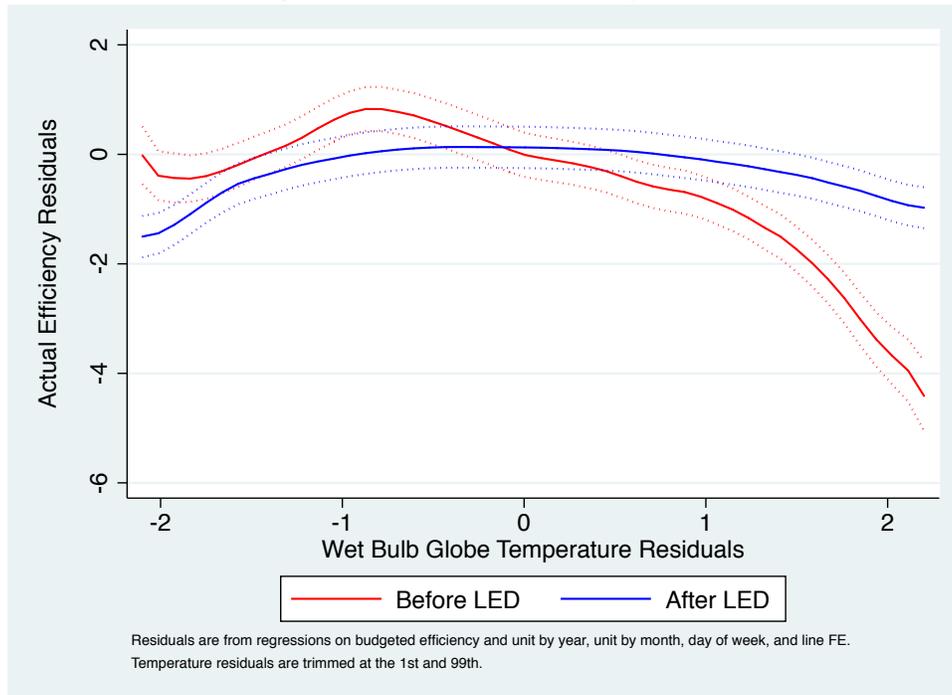


factors driving unit-specific non-linear trends such as differential rates of expansion across factories or primary buyers. Unit by month fixed effects control for unit-specific seasonality due to, for example, garment demand and labor supply patterns; day of week fixed effects control for fluctuations in efficiency across work days due to for example fatigue or weekend salience.

Finally, to check that the patterns depicted in Figures 2 and 3 above persist even after controlling for all of these unobservables, we can repeat the exercises depicted in those figures but using residuals from regressions of efficiency, temperature, and LED on all of these controls. Figure 4 shows that the residual efficiency-temperature gradient after controlling for the full set of covariates listed above is still negative and non-linear with a more steeply negative slope at higher temperatures. Figure 5 plots the residual efficiency-temperature gradient before and after LED, respectively. It shows that the difference between the with and without LED gradients grows at high temperatures (residualized) as the pre-LED gradient becomes more steeply negative and the post-LED gradient remains relatively flat. Note that in fact the two gradients in Figure 5 are not statistically significantly different at low temperatures, but the low LED residual gradient (red line) falls statistically significantly below the high LED residual gradient (blue line) just below wet bulb residual values of 0.²⁷ The comparison depicted

²⁷For figures representing the difference in gradients exactly as depicted in both Figures 3 and 5, please refer to the appendix.

Figure 5: Residual Gradient by LED

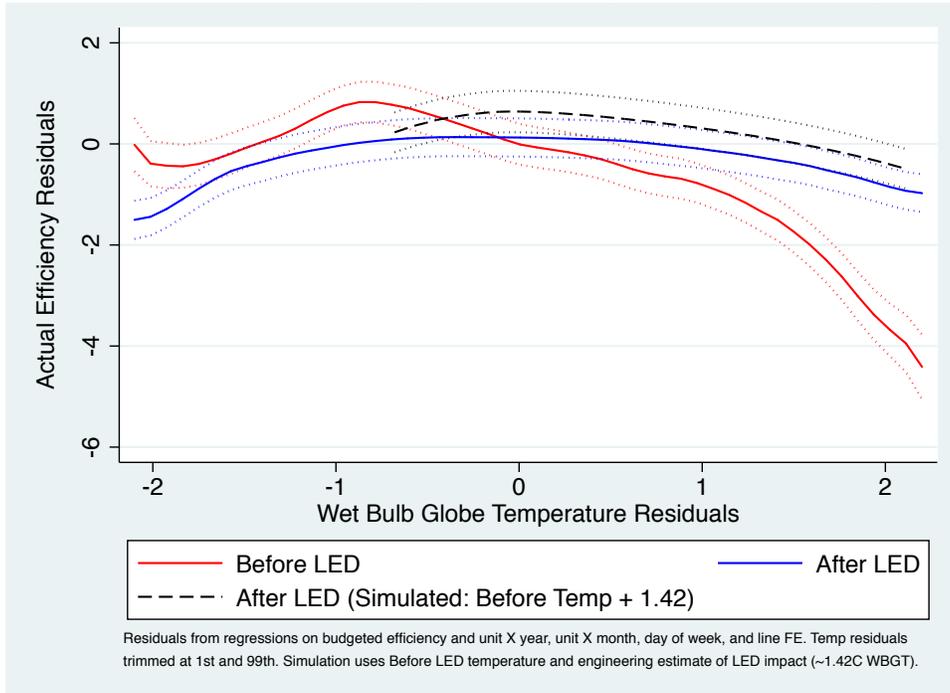


in Figure 5 illustrates the intuition of the flexible semiparametric estimation strategy we propose in the next section.

As discussed above, we commissioned an engineering report to provide estimates of the effect of LED on indoor temperature for any given outdoor temperature for the building and lighting specifications of a representative factory in our sample. The report stated that the LED lighting change should have reduced the indoor temperature by around 2.4 degrees Celsius dry-bulb temperature or roughly 1.42 degrees Celsius wet-bulb globe temperature (at mean levels of relative humidity observed in our data) at each outdoor temperature. That is, the report predicts that LED installation would have shifted the intercept of the relationship in Figure 1 down. This would translate precisely into a shift to the right of the pre-LED curve in Figure 5. To illustrate the result of this exact impact as predicted by the engineering calculations, we show in Figure 6 the same estimated gradients presented in Figure 5 for pre- and post-LED, but with the addition of the simulated post-LED gradient. This simulated curve is precisely the pre-LED data shifted to the right by the increment estimated in the engineering report (1.42 degrees WBGT), with the support of the simulated temperature distribution restricted to be common with the support the observed temperature distribution.²⁸

²⁸Note of course that the simulated temperature distribution will be truncated to the left at the point 1.42 degrees to the right of the left limit of the observed temperature distribution.

Figure 6: Residual Gradient by LED Including Simulated Impact of LED from Engineering Estimates



This simulation matches the observed post-LED gradient remarkably well, validating the interpretation that LED adoption impacted efficiency precisely by way of a shift downward in the intercept of the indoor-outdoor temperature relationship as depicted in Figure 1. Note, however, that if one takes a non-linear curve as given by the pre-LED gradient and shifts it to the right, and then attempts to measure precisely the difference in the slopes between the two curves, a simple parametric functional form specification in OLS will not perfectly fit the difference in these curves. Rather the best way to measure the difference between these non-linear curves is to fit each non-parametrically (or, more accurately, semi-parametrically given that Figure 5 presents residuals from OLS regressions on the full set of controls), and measure the difference between these non-parametric curves. We, accordingly, undertake this exact exercise as described below.

4.3 Semiparametric Treatment Effect Estimation (Mediation Analysis)

The above parametric spline regression analysis approximates the estimation of the change in efficiency-temperature gradients due to the introduction of LED lighting. However, the parametric spline specification embodies functional form assumptions based on visual inspection of the gradients in Figures 2 and 3. A more flexible and agnostic empirical approach would express efficiency as some general func-

tion of temperature after accounting for all of the relevant covariates and allow this function to differ before and after the introduction of LED. Specifically, this would amount to attempting to estimate the total impact of LED from the following equation:

$$E_{ludmy} = \alpha_0 + f[T_{dgm_y}](1 - LED) + g[T_{dgm_y}](LED) + \phi B_{ludmy} + \alpha_l + \gamma_{uy} + \eta_{um} + \delta_d + \varepsilon_{ludmy}. \quad (6)$$

Here $f[T_{dgm_y}]$ is a general function of temperature which explains efficiency when $LED = 0$, after controlling for the full set of covariates; and $g[T_{dgm_y}]$ is the analogous general function of temperature explaining efficiency when $LED = 1$.

In order to recover the average impact of LED on efficiency from equation 6, we first partition the regression to isolate the terms containing temperature from the remaining covariates, both with and without LED. We do this by regressing efficiency and temperature on budgeted efficiency and the full set of fixed effects and calculating the residuals from each regression, separately for the sample with and without LED.²⁹

We then non-parametrically estimate using kernel-weighted local polynomial smoothing $f[T_{dgm_y}]$ and $g[T_{dgm_y}]$ for each 0.1 width bin in wet bulb globe temperature residuals using the subsample of data with and without LED, respectively. We also recover standard errors for each bin from both curves using the non-parametric estimation procedure. Next, we subtract estimated values of $f[T_{dgm_y}]$ from $g[T_{dgm_y}]$ for each 0.1 width bin of the wet bulb residual and calculate the appropriate two-sample standard error for the difference.

Note that this amounts to estimating the difference at each temperature point between the non-parametric residual gradients depicted in Figure 5, and recovers the estimated treatment effect of LED on efficiency at each point along the observed temperature distribution, after accounting for any endogeneity in unobservables as discussed above. Figures depicting these point for point differences between the residual gradients and their statistical significance are presented and discussed in section 5 below. It should also be noted that this semiparametric procedure is identical in intuition to the degree and decile bin temperature effects specifications estimated in previous studies (Barreca et al., 2016),

²⁹Note that this assumes conditional mean independence of LED, which is supported by the empirical tests shown in Figure 8 indicating that after accounting for the full set of covariates and fixed effects LED and temperature are indeed orthogonal. Instead of using the LED binary variable, we can approximate the residualized $(1 - LED)$ and (LED) terms with a dummy that takes the value 1 if the LED residual (residual from regressing $1(LED)$ on budgeted efficiency on all the fixed effects) ≥ 0 and value 0 if the LED residual < 0 . We have conducted the analysis under this assumption as well and find the results to be qualitatively similar to the preferred approach reported in the paper. These alternate results are available upon request.

but extends and generalizes previous approaches in two ways. First, we leverage the additional granularity and quantity of data in our setting to estimate effects for each 0.1 degree temperature residual bin rather than degree or decile bins of greater width. Second, we combine non-parametric estimation techniques with fixed effects specifications to allow temperature effects to vary as flexibly as possible across bins while preserving causal interpretation of the estimates. Beyond this added granularity and flexibility, the intuition behind previous approaches to estimating non-linear impacts of temperature across the distribution is preserved.

Finally, we calculate the temperature weighted average treatment effect of LED by multiplying the difference between the gradients at each temperature point at the 0.1 degree level by the probability that temperature occurs and then adding the full set of these products. The temperature probability distribution is calculated from the data. This procedure provides us with an estimate of the total impact of LED on efficiency as mediated by temperature, which is necessary for the cost-benefit calculations we conduct below. To this end, the semiparametric procedure developed here represents a novel approach to mediation analysis in which a continuous covariate is believed to mediate the impact of a regressor of interest on an outcome, but the functional form of the relationship between the regressor and the outcome and the structure of the mediating mechanism are either unknown or not easily or parsimoniously parametrized. In particular, when the relationship between the regressor of interest and the outcome is believed to be (or assumed to be) linear and the impact of the regressor on the mediating factor can be easily estimated, simpler parametric approaches to mediation analysis can be used.³⁰

5 Results

5.1 Parametric Spline Regression Analysis

We begin by reporting results from the estimation of the parametric spline specifications presented in equations 4 and 5. Columns 1 and 2 of Table 2 reports estimates of β^L and β^H from equation 4 with column 2 estimates corresponding to a specification with an additional control for precipitation. The additional precipitation control ensures that impacts are indeed being driven by temperature exposure alone and are not composite effects reflecting the impacts of other correlated weather conditions.

³⁰A long literature develops and implements analyses of this type. For a recent application of this more traditional parametric approach to mediation analysis see Heckman et al. (2013).

Columns 3 and 4 report estimates of β_1^L , β_1^H , β_2 , β_3^L and β_3^H from equation 5, once again with column 4 reporting results after the inclusion of an additional control for precipitation.

The spline regression estimates from columns 1 and 2 reflect the pattern shown in Figures 2 and 4 with the slope of the efficiency-temperature gradient below 19 degrees Celsius of wet bulb globe temperature being slightly negative (statistically indistinguishable from 0) and the slope above 19 degrees being strongly negative and statistically significant at the 1 percent level. Point estimates indicate that, at wet bulb globe temperatures above 19 degrees Celsius, a one degree increase in temperature leads to a reduction of more than 2.1 percentage points in actual efficiency. A comparison of estimates across columns 1 and 2 show that the inclusion of an additional control for precipitation has minimal impact on results.

The results in columns 3 and 4 are consistent with the pattern reflected in Figures 3 and 5 with the introduction of LED having no significant impact on the slope of the efficiency-temperature gradient below 19 degrees Celsius, but strong attenuating impact on the negative slope of the gradient above 19 degrees. That is, the estimates indicate that the introduction of LED offsets the negative impacts of temperature on efficiency by roughly 85%, attenuating the magnitude of the negative slope above 19 degrees from around -2 to roughly -0.3. LED shows no significant impact below 19 degrees Celsius which is consistent with the evidence from ergonomics and physiology literatures suggesting that temperature is most impactful on human functioning at temperatures above this level. The estimate of the main effect of LED is positive and large, consistent with the pattern shown in Figures 3, but is imprecisely estimated and statistically indistinguishable from 0.

The results reported in Table 3 correspond to the regression of mean line-daily worker attendance on the identical specifications to those in Table 2 as described in section 4.2.1. The estimates from Table 3 suggest a negative impact of temperature on attendance at temperatures below 19 degrees Celsius; however, the magnitudes of the point estimates are extremely small (less than 1% of the mean). All other estimates of coefficients, including those reflecting the impacts of LED, are statistically indistinguishable from 0. In general, we interpret the results in Table 3 as indicative of no real impacts of temperature on worker attendance. These results imply that it is unlikely that impacts of temperature on worker attendance are contributing to the estimated impacts of temperature and LED installation on efficiency.

To further verify that worker attendance is not a primary mediating mechanism of the impacts of temperature and LED installation on efficiency, we repeat the analysis reported in Table 2 with mean

Table 2

Impact of Temperature on Production Efficiency and Mitigative Impact of LED Lighting

	(1)	(2)	(3)	(4)
	Actual Efficiency (Actual Production / Targeted Production)*100			
Wet Bulb Globe Temperature <19	-0.3 (0.5)	-0.32 (0.47)	-.094 (0.73)	-0.105 (0.7)
Wet Bulb Globe Temperature >=19	-2.135*** (0.002)	-2.169*** (0.002)	-1.95*** (0.002)	-1.98*** (0.002)
1(LED)*(Wet Bulb Globe Temperature <19)			-.106 (0.79)	-0.103 (0.70)
1(LED)*(Wet Bulb Globe Temperature >=19)			1.67*** (0.006)	1.68*** (0.004)
1(LED)			3.45 (0.68)	3.39 (0.68)
Fixed Effects	Factory x Year, Factory x Calendar Month, Production Line, Day of the Week			
Precipitation Control	N	Y	N	Y
Observations	74,939	74,939	239,680	239,680
Mean of Dependent Variable	53.73	53.73	55.234	55.234

Notes: Wild-cluster bootstrap p-values in parentheses (***) denotes significance at the 1% level, ** denotes significance at the 5% level, * denotes significance at the 10% level). Clustering is done at the factory level. All measures of temperature are in degree Celsius. All regressions include daily budgeted efficiency as a control variable.

Table 3

Impact of Temperature on Attendance and Mitigative Impact of LED Lighting

	(1)	(2)	(3)	(4)
Worker Presence (Line-Level Mean Daily Probability)				
Wet Bulb Globe Temperature <19	-0.0061*** (0.01)	-0.0059** (0.02)	-0.0011 (0.53)	-0.0007 (0.7)
Wet Bulb Globe Temperature >=19	0.0003 (0.90)	0.0007 (0.82)	0.0056 (0.27)	0.0064 (0.19)
1(LED)*(Wet Bulb Globe Temperature <19)			0.0003 (0.89)	0.0002 (0.94)
1(LED)*(Wet Bulb Globe Temperature >=19)			-0.0051 (0.38)	-0.0053 (0.37)
1(LED)			-0.0065 (0.83)	-0.0054 (0.86)
Fixed Effects	Factory x Year, Factory x Calendar Month, Production Line, Day of the Week			
Precipitation Control	N	Y	N	Y
Observations	392,601	392,601	392,601	392,601
Mean of Dependent Variable	0.829	0.829	0.829	0.829

Notes: Wild-cluster bootstrap p-values in parentheses (***) denotes significance at the 1% level, ** denotes significance at the 5% level, * denotes significance at the 10% level). Clustering is done at the factory level. All measures of temperature are in degree Celsius.

Table 4

Impact of Temperature on Production Efficiency and Mitigative Impact of LED Lighting
(Controlling for Mean Line-Daily Attendance)

	(1)	(2)	(3)	(4)
	Actual Efficiency (Actual Production / Targeted Production)*100			
Wet Bulb Globe Temperature <19	-0.443 (0.39)	-0.47 (0.36)	-0.137 (0.63)	-0.148 (0.60)
Wet Bulb Globe Temperature >=19	-2.498*** (0.004)	-2.55*** (0.004)	-2.164*** (0.008)	-2.196*** (0.008)
1(LED)*(Wet Bulb Globe Temperature <19)			-0.016 (1.00)	-0.013 (1.00)
1(LED)*(Wet Bulb Globe Temperature >=19)			1.605** (0.02)	1.617** (0.02)
1(LED)			1.617 (0.89)	1.565 (0.89)
Line-Level Mean Daily Worker Presence	1.884 (0.48)	1.884 (0.48)	2.188** (0.03)	2.195** (0.03)
Fixed Effects	Factory x Year, Factory x Calendar Month, Production Line, Day of the Week			
Precipitation Control	N	Y	N	Y
Observations	61,782	61,782	203,554	203,554
Mean of Dependent Variable	53.05	53.05	55.09	55.09

Notes: Wild-cluster bootstrap p-values in parentheses (***) denotes significance at the 1% level, ** denotes significance at the 5% level, * denotes significance at the 10% level). Clustering is done at the factory level. All measures of temperature are in degree Celsius. All regressions include daily budgeted efficiency as a control variable. Line-Level Mean Daily Probability of Worker Presence is the average probability that a worker is present for a production line on a given day.

line-daily worker attendance as an additional control. To the degree that estimates remain largely unchanged after including this additional control for attendance, we conclude that attendance is not a primary mediator of the impacts of temperature on efficiency nor of the attenuation of impacts caused by LED installation. Indeed, the results from these regressions reported in Table 4 are remarkably similar to those presented in Table 2. Overall, we interpret the results in Tables 3 and 4 as strong evidence against the importance of attendance as a primary mediator of the impacts of temperature and LED installation on efficiency. That is, we find that exposure to higher temperatures impacts the intensive margin of productivity per unit labor supplied, but does not impact strongly the extensive margin of the quantity of labor units supplied. Similarly, the introduction of LED attenuates greatly the impacts of temperature on the intensive margin of efficiency, but has no perceptible impact on the extensive margin.

Next, we investigate whether the impacts we are estimating of contemporaneous temperature exposure on efficiency are indeed reflecting contemporaneous exposure alone rather than a composite estimate of contemporaneous exposure impacts and persistent impacts of lagged exposure. Similarly, we check that the estimated attenuation from LED installation is working through contemporaneous temperature exposure. Although persistent impacts of lagged exposures and serial correlation in temperature would not invalidate in any way the analysis conducted above, the interpretation of the point estimates will change based on the underlying sources of variation. As discussed in section 5, we repeat the analysis reported in Table 2 but include 7-day distributed lag temperature spline terms and, where appropriate, their interactions with LED installation. The results from the estimation of these augmented specifications are reported in Table 5. Specifically, all results reported in Table 5 correspond to specifications including 7-day distributed lag temperature spline terms and results in columns 3 and 4 correspond to specifications also including interactions of distributed lag spline terms with the LED installation dummy. Overall, the results in Table 5 are qualitatively identical to the main results reported in Table 2, but with larger magnitudes for coefficients on the above 19 degree temperature spline and the corresponding LED interaction terms. These results indicate that indeed estimates of temperature impacts and attenuation from LED installation are being driven by contemporaneous exposures and that a more rigorous isolation of contemporaneous temperature variation leads to even more pronounced impacts of temperature and LED installation. While daily temperature is generally believed to reflect some degree of serial correlation, the similarity in results with and without distributed lags is not altogether surprising in our study. In particular, we should note that the baseline specifications already include a large set of heterogeneous non-linear trends (e.g., unit by month FE) to soak up a great deal of this less transitory variation in temperature. Indeed, the correlations between contemporaneous temperature and lagged temperature values after partialling out the full set of controls are quite small (never more than .25 and mostly below .1).

5.2 Semiparametric Treatment Effect Estimation

After estimating the non-linear relationship between temperature and efficiency and the attenuating impact of LED installation on this relationship, we turn to a more flexible empirical approach in order to recover the overall impact of LED installation on efficiency as mediated by the full distribution of temperature exposures. Specifically, in order to fully capture the impact of LED on efficiency at

Table 5

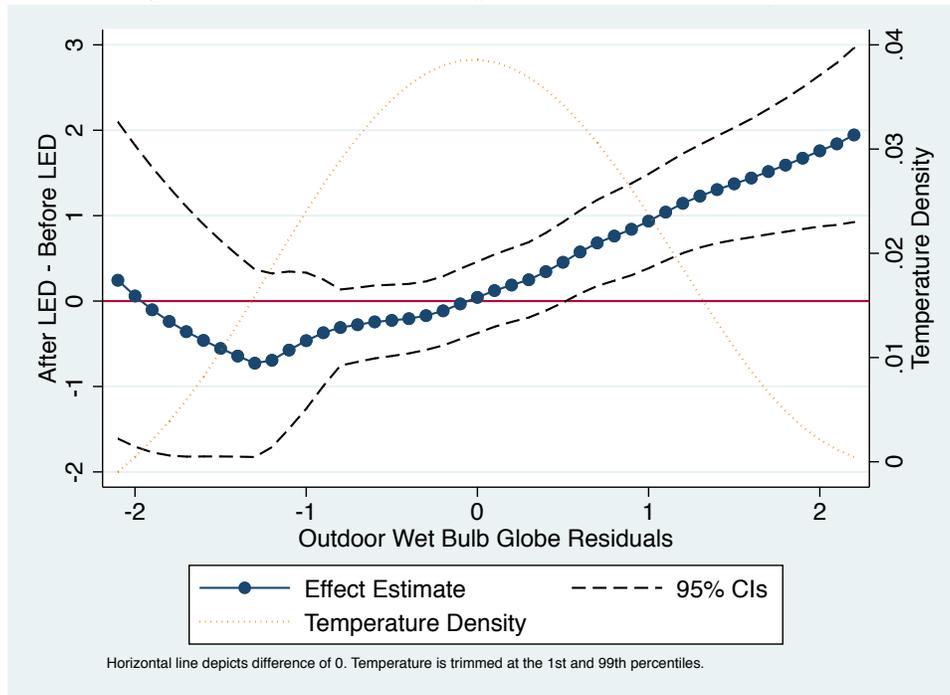
Impact of Temperature on Production Efficiency and Mitigative Impact of LED Lighting (Distributed Lag Specification)				
	(1)	(2)	(3)	(4)
Actual Efficiency (Actual Production / Targeted Production)*100				
Wet Bulb Globe Temperature <19	-0.44 (0.29)	-0.467 (0.26)	-0.227 (0.41)	-0.245 (0.35)
Wet Bulb Globe Temperature >=19	-2.236*** (0.002)	-2.271*** (0.002)	-2.295*** (0.002)	-2.32*** (0.002)
1(LED)*Wet Bulb Globe Temperature <19			0.072 (0.8)	0.078 (0.79)
1(LED)*Wet Bulb Globe Temperature >=19			2.375*** (0.0)	2.384*** (0.0)
1(LED)			-9.199 (0.40)	-9.165 (0.40)
7-day Distributed Lag Temperature Splines	Y	Y	Y	Y
7-day Distributed Lag Spline Interactions with LED	N	N	Y	Y
Fixed Effects	Factory x Year, Factory x Calendar Month, Production Line, Day of the Week			
Precipitation Control	N	Y	N	Y
Observations	74,939	74,939	239,680	239,680
Mean of Dependent Variable	53.732	53.732	55.23	55.23

Notes: Wild-cluster bootstrap p-values in parentheses (***) denotes significance at the 1% level, ** denotes significance at the 5% level, * denotes significance at the 10% level). Clustering is done at the factory level. All measures of temperature are in degree Celsius. All regressions include daily budgeted efficiency as a control variable. Full table reporting coefficients on 7 day distributed lag temperature splines and their interactions with LED is presented in the Appendix Table A1.

each point along the temperature distribution, we must relax the parametric restrictions imposed in the above analysis. Although we wish to maintain the importance of covariates and fixed effects in isolating the causal relationships between temperature, LED installation, and actual efficiency, we do not want to impose functional forms and parametric restrictions on the relationships depicted in Figure 5. The most flexible and agnostic approach, as described in section 4.3, would be to estimate the observed distance between the residual temperature-efficiency gradients (of general shape) with and without LED at each point along the common support of residual temperature distribution for pre-LED and post-LED observations. This amounts to calculating the distance between the low LED residual (red line) gradient and high LED residual (blue line) gradient in Figure 5 for each .1 degree temperature residual bin along the x-axis.

These calculated differences, representing treatment effect estimates at each temperature value, are depicted in Figure 7 along with the observed probability density of temperature residuals. Figure 7 shows that, as indicated in the preliminary graphical evidence and the parametric spline estimates presented above, estimates of the treatment effect of LED on efficiency are small at low temperatures but rise monotonically with higher temperatures, ultimately plateauing at around the 90th (value of

Figure 7: Difference in Semiparametric Gradients by LED



.98) percentile of the residual temperature distribution. Gains in efficiency due to LED installation range from 1.2 to 1.4 percentage points for the top 25 percent of temperature values.

These semiparametric treatment effect estimates for each .1 degree bin along the temperature residual distribution corroborate with empirical flexibility and rigor the pattern of impacts shown in the parametric spline results above. However, the primary value to conducting the semiparametric analysis is the ability to calculate the total impacts of LED installation on efficiency by way of temperature-probability-weighted averages of treatment effects along the entire temperature distribution. As described in section 4.3, we multiply the value represented by each solid blue dot in the connected line of treatment effects depicted in Figure 7 by the corresponding density value shown in the underlying temperature distribution (faint dotted line) and then sum across this full set of probability-weighted treatment effects. This is the computational equivalent to integrating the distance between the curves in Figure 5 over the temperature residual distribution.

The results of this exercise are reported in Table 6. The first row of column 1 in Table 6 reports that the temperature-probability-weighted average treatment effect of LED installation on actual efficiency is just over .7 percentage points and is significant at the 1 percent level. This estimate of the overall impact of LED installation on efficiency allows us to do cost-benefit calculations on the adoption of

Table 6

Treatment Effects from Differencing of Semiparametric Efficiency-Temperature Gradients Across LED Status

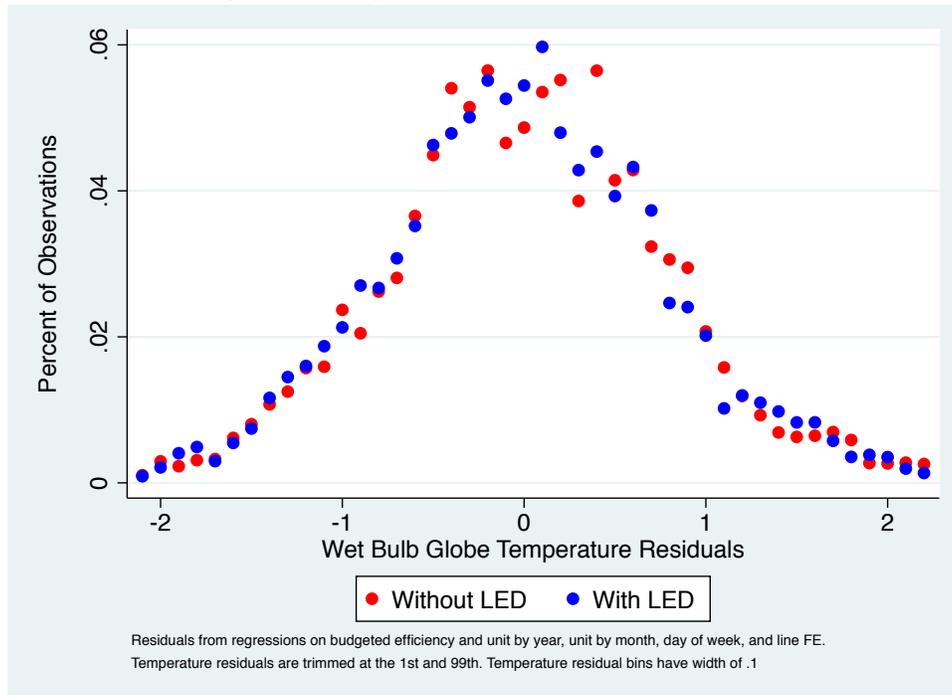
	(1)	(2)	(3)	(4)
	Actual Efficiency (Actual Production / Targeted Production)*100		Worker Presence (Line-Level Mean Daily Probability)	
Temp Prob Weighted Average Treatment Effect	0.7231807***	0.7043784***	0.0006151	-0.0010373
Temp Prob Weighted Average Standard Error	(0.2845352)	(0.2423353)	(0.0053535)	(0.0065272)
Temp Prob Weighted Average T-Stat	2.541621	2.906628	0.1149049	-0.1589111
Uniform Weighted Average Treatment Effect	0.6151872**	0.6390683**	0.0000539	0.0017291
Uniform Weighted Average Standard Error	0.4017829	0.3105415	0.0062784	0.0080576
Uniform Weighted Average T-Stat	2.170351	2.163561	0.1426833	0.0701154
Fixed Effects	Factory x Year, Factory x Calendar Month, Production Line, Day of the Week			
Precipitation Control	N	Y	N	Y
Observations	234888	234888	384749	384749

Notes: Treatment effect estimates are from locally weighted polynomial smoothing functions relating residuals of the outcome variable to residuals of temperature. Smoothed values of the outcome residual are calculated at 50 points along the temperature residual distribution. These sets of 50 smoothed values are calculated separately for led residual values less than zero and greater than zero. Residuals are taken from regressions of outcome, led and temperature variables on all controls and fixed effects. The two curves are then differenced point by point along the temperature residual distribution, and the weighted average of this difference is calculated using the probability that temperature residuals fall within bins corresponding to the 50 points as the weights. Estimated standard errors in parentheses are calculated as the square root of the estimated conditional variance from a higher order local polynomial fit within a bandwidth of 1.5 times the smoothing bandwidth. Reported t-statistics are the corresponding weighted averages of treatment effects at each smoothed point divided by the estimated standard error at each smoothed point. P-values are calculated by comparing t-statistics to conventional asymptotic student t distributions (*** p<0.01, ** p<0.05, * p<0.1). All regressions include daily budgeted efficiency as a control variable.

LED lighting. These calculations are presented in Table 8 and discussed in section 6. In column 2, we report the analogous estimate from the same exercise but with the inclusion of an additional control for precipitation. As in the parametric spline results, the additional control does not meaningfully affect the estimate. Below these estimates, we report the average treatment effect estimate obtained using a uniform weight rather than the underlying temperature probability density. These estimates are qualitatively similar, but are slightly smaller in magnitude indicating that in our data higher temperatures at which LED installation has a larger impact on efficiency are more frequent than the lower temperatures at which LED has little impact. Accordingly, without accounting for the underlying distribution of temperature as the mediator of LED impacts on efficiency, we would underestimate the total impact of LED on efficiency. In columns 3 and 4, we report estimates from the identical exercise with line-level mean daily presence probability as the outcome. The results show that even when adjusting for the probability distribution of temperature, LED has no measurable impact on worker attendance.

One underlying assumption for the validity of the semiparametric exercise conducted here is the equivalence of the observed temperature distributions before and after LED installation. While the method will by construction not reflect differences in the underlying support of the temperature dis-

Figure 8: Temperature Residual Distribution

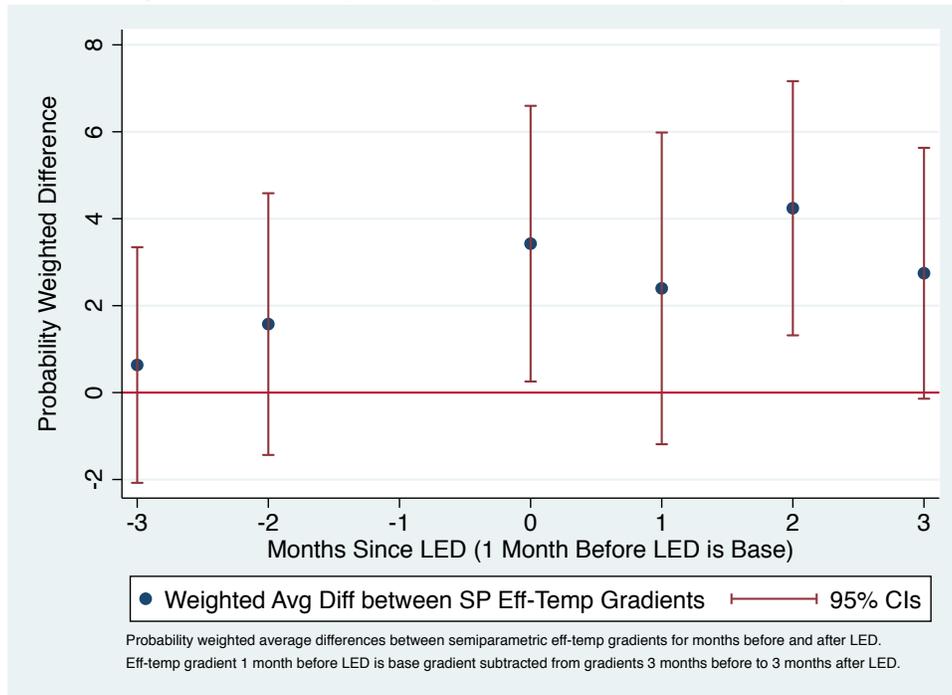


tributions, differences in frequency of particular temperature ranges before and after LED installation will convolute the analysis. That is, while the method explicitly calculates treatment effects from only those temperature values that exist in both pre-LED and post-LED samples and therefore will not reflect issues of uncommon support, a higher likelihood of high temperatures before LED installation as compared to after LED, for example, would impact the estimates adversely. Accordingly, we check that the residual temperature distributions, after controlling for the full set of covariates and fixed effects, for low LED residual and high LED residual samples are statistically equivalent. Figure 8 plots the two distributions and visually the distributions appear equivalent. We also conduct a Kolmogorov-Smirnov nonparametric test of the equivalence of the distributions and cannot reject that they are equivalent.³¹

Finally, now that we have developed and implemented a method for recovering total impacts of LED on efficiency, we can use this method to present an event study in support of the sharp timing of the impact. That is, with LED having a highly non-linear impact on efficiency dependent on which temperatures prevail at a given time, a simple event study using coefficients from linear regressions

³¹The results of this test are available upon request. This empirically verified orthogonality between LED and temperature residuals allows us to omit temperature from the LED residual regression and LED from the temperature residual regression discussed in section 4.3.

Figure 9: Efficiency Semiparametric Estimate Event Study

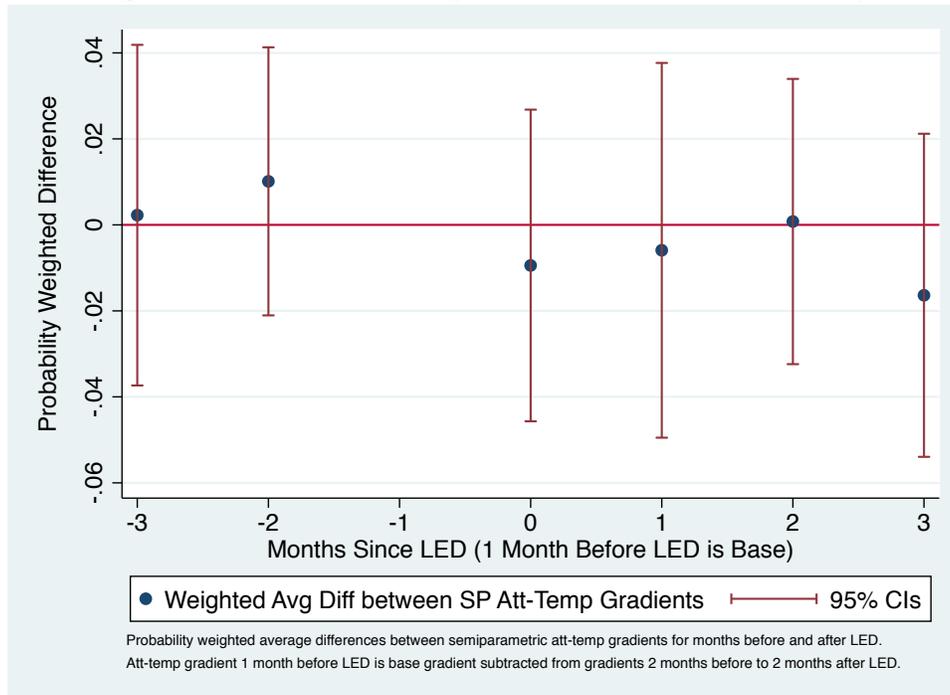


or other parametric specifications would be subject to the issues of commonality in support month by month which are precisely avoided by the semiparametric approach developed here. Accordingly, we utilize our semiparametric temperature-probability-weighted treatment effects estimator to calculate treatment effects month by month relative to the month of LED installation. That is, we draw the semiparametric temperature-efficiency residual gradient using only the data from the month directly before LED installation in each factory.³² This represents our baseline gradient. We then draw the analogous gradient using only data from the month in which LED lighting was introduced in each factory. Then, following the steps set forth in section 4.3, we calculate the temperature-probability-weighted difference between these gradients and the corresponding standard errors and t-statistics. These estimates are plotted at time 0 in Figure 9. We repeat this exercise to calculate the difference between the gradient for 1 month after LED installation and the base gradient of 1 month before LED installation, as well as the difference for the gradients 2 and 3 months after LED installation. Also, as falsification checks, we calculate the differences between the base gradient and gradients 2 and 3 months *before* LED installation.

All of these probability-weighted treatment effect estimates and standard errors and corresponding

³²Note that since timing of LED installation is central to this exercise, factory units that already had LED lighting at the beginning of our data and those which still did not have LED lighting by the end of our data are excluded.

Figure 10: Attendance Semiparametric Estimate Event Study

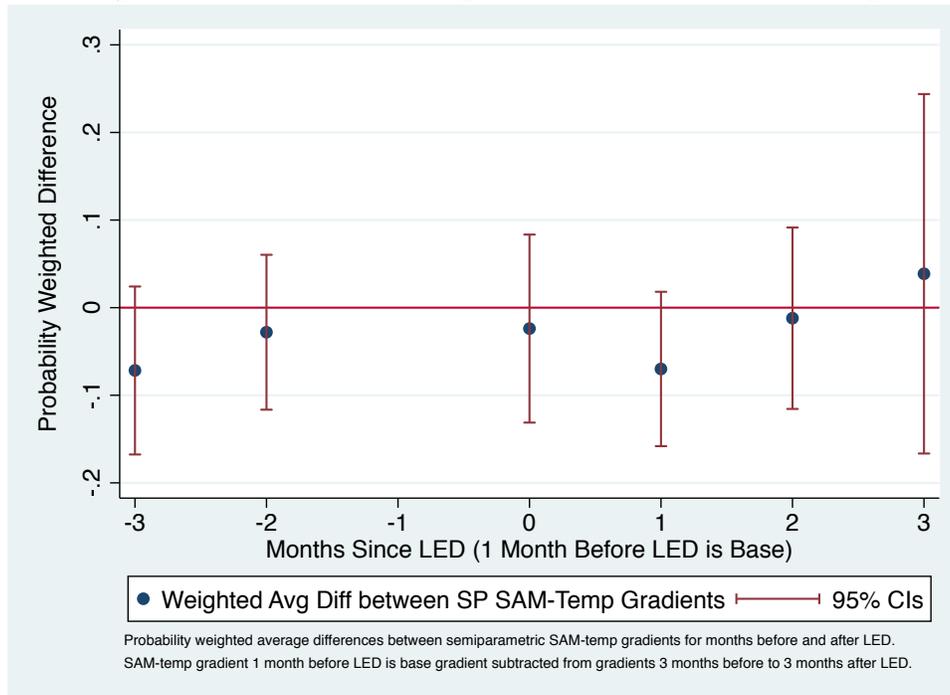


t-statistics are plotted in Figure 9. We interpret this event as remarkably strong evidence of the sharpness of the timing of impacts around LED installation. Indeed, prior to LED installation efficiency-temperature gradients are statistically indistinguishable, but as soon as LED lighting is introduced efficiency-temperature gradients reflect large, positive, and statistically significant differences from the base gradient of one month prior to LED installation. We present the analogous event study for attendance in Figure 10. Consistent with the previous estimates on attendance, we find no evidence of even a transitory impact of LED installation on attendance in Figure 10.

5.3 Checks for Exogeneity of LED Roll-Out

In this section, we present several additional checks on the exogeneity of the timing of LED installation. We begin by conducting an event study like that presented in Figure 9, but for a placebo outcome that should not be impacted by the introduction of LED. Specifically, we present in Figure 11 the event study for SAM. If the introduction of LED was timed around peak production cycles or seasonal buying patterns, the event study for SAM, which measures differences in the types and complexity of garments, should show systematic fluctuations relative to the timing of LED installation. Figure 11 shows no evidence of endogeneity in the timing of LED installation with respect to production charac-

Figure 11: Placebo SAM Semiparametric Estimate Event Study



teristics.

In Table 7, we report additional regression results in further support of the exogeneity of LED roll-out. In column 1, we report estimates of the coefficients on the temperature spline terms from the regression of the LED introduction dummy on the main specification in equation 4. We find no evidence that LED installation was timed around particular temperature realizations. These results are consistent with the evidence of the equivalence of pre-LED and post-LED temperature residual distributions presented in Figure 8. In columns 2 and 3 of Table 7, we report results from the regression of SAM and budgeted efficiency, respectively, on the LED installation dummy, the date relative to LED installation, and their interaction with the remaining specification identical to that depicted in equation 4. These regressions are meant to check whether garment style and complexity (SAM) and order size (budgeted efficiency) varied systematically in the lead up to LED installation or immediately after. Significant coefficient estimates in columns 2 and 3 would reflect evidence that the timing of LED introduction is endogenous with respect to these production factors; however, we find no such evidence. Finally, in columns 4 through 6, we check that LED installation was not accompanied by other forms of upgrading. Specifically, we regress the proportion of each of the three skill levels of tailors - A, B, and C grade - hired on each day in each factory unit on the same specification reported

Table 7

Checks for Exogeneity of LED Roll-Out

	(1)	(2)	(3)	(4)	(5)	(6)
	1(LED)	Standard Allowable Minutes (SAM)	Budgeted Efficiency	Proportion of A grade tailors hired	Proportion of B grade tailors hired	Proportion of C grade tailors hired
Wet Bulb Globe Temperature <19	0.0017 (0.72)					
Wet Bulb Globe Temperature >=19	-0.0078 (0.24)					
1(LED)*Date Relative to LED Installation		0.0010 (0.18)	-0.000004 (0.82)	0.00001 (1)	0.00008 (0.46)	-0.00005 (0.78)
1(LED)		2.994 (0.33)	-0.0388 (0.478)	-0.0211 (0.40)	0.0439 (0.19)	-0.019 (0.38)
Date Relative to LED Installation		-0.0007 (0.17)	0.000003 (0.81)	0.00006 (0.94)	0.00129* (0.06)	-0.0012 (0.19)
Fixed Effects	Factory x Year, Factory x Month, Production Line, Day of the Week			Factory x Year, Factory x Month, Day of the Week		
Precipitation Controls	Y	Y	Y	Y	Y	Y
Temperature Controls	N	Y	Y	Y	Y	Y
Observations	239,680	134,326	134,326	8,595	8,561	8,562
Mean of the dependent variable	0.69	0.751	61.64	0.44	0.29	0.27

Notes: Wild-cluster bootstrap p-values in parentheses (***) denotes significance at the 1% level, ** denotes significance at the 5% level, * denotes significance at the 10% level). Clustering is done at the factory level. Since columns 2 through 6 consider the date relative the LED installation, units that had LED lighting at the beginning of the sample period or did not have LED lighting by the end of the sample period are omitted. The first 3 columns are at the production line-date level, and the last 3 columns are defined at the unit-date level.

in columns 2 and 3.³³ We do not find any evidence that hiring patterns changed in the lead up to LED installation or immediately after.

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 which accrue privately to the adopting firm.

Specifically, we find that the introduction of LEDs eliminates roughly 85% percent of the negative impact of temperature on worker efficiency during relatively hot days. Using the semiparametric treat-

³³A grade tailors are the most skilled, followed by B grade tailors, and C grade tailors are the least skilled.

Table 8

Cost-Benefit Calculations for LED Adoption

<i>Cost of Implementation (one-time)</i>	
Investment per bulb (bulb, wiring, etc.)	\$8.53
Number of bulbs replaced per factory unit	1193
Total Cost of Implementation	\$10,180.27
<i>Energy Savings (per year)</i>	
Power consumption savings per bulb	\$2.40
Number of bulbs replaced	1193
Total Energy Savings	\$2,863.20
<i>Efficiency Gains (per year)</i>	
Average efficiency gain in percentage point from LED-caused temperature reductions (Table 7)	0.7044
Efficiency percentage point gain to profit percentage point gain translation	0.1875
Profit margin at baseline	5%
Average revenue in USD per factory unit per year	\$10,000,000.00
Average profit in USD per factory unit per year	\$500,000.00
Total Efficiency Gains	\$13,207.10
Total Net Savings from LED Adoption in the first year	\$5,890.03
Total Net Savings from LED Adoption in the second year	\$16,070.30
Carbon (Public) Benefits from LED adoption (at \$27/tC)	\$101.23
Carbon (Public) Benefits from LED adoption (at \$93/tC)	\$345.68
Notes: Profit margin at base was taken from accounting department of the firm. Calculation for translating efficiency gains to profit come from accounting estimates of the proportion of "Cut to Make" (non-material) costs that can be recovered via efficiency gains. "Cut to Make" costs make up 25% of total costs, and 75% of these costs are recoverable via improved efficiency. Average revenue in USD per factory unit per year is obtained by taking balance sheet revenues for the firm and dividing by number of plants. For additional details on the calculation of carbon benefits, please refer to the Discussion section.	

ment effect estimator developed in this study to calculate the average total impact of LED installation given the specific treatment effect at each observed temperature and the underlying probability distribution of temperature, we estimate an average total increase in production efficiency of roughly .7 percentage points (or more than 1% improved productivity from the mean), as presented in Table 6.

6.1 Private Benefits (Firm Cost-Benefit Calculations)

We combine our estimate of average total efficiency gains with actual production and costing data from the firm to calculate annual costs and benefits of LED installation and report these calculations in Table 8. Senior management at the firm, with whom we worked closely on this study, estimated

that the profit gains for each percentage point gain in efficiency were 0.1875 percentage points.³⁴ Thus, a 0.7 percentage point gain in efficiency from LED installation translates to a .13 percentage point gain in profits (or a 2.6% increase in profitability from the 5% baseline profit margin of the firm). At an approximate profit value per factory per year of 500,000 USD, the introduction of LED results in increased profits of 13,207 USD per factory per year from gains in production efficiency.

How does this estimate change the cost-benefit calculations of LED adoption for the firm? To start, we obtained the energy cost savings calculations the firm used when making its LED adoption choices. Management estimated that the total energy cost savings per year per factory unit of LEDs (as compared with CFL bulbs, which were being used before LED introduction) were approximately 2.40 USD per bulb replaced or roughly 2,863 USD in total for an average replacement of 1193 bulbs per factory in our data.³⁵ The added profits from efficiency gains per year we computed are nearly 5 times this amount. The cost of replacing the average factory's bulbs to LEDs is roughly 10,180 USD. Thus, if only energy savings were taken into account, it would take more than 3 and a half years to break even. However, when the productivity benefits are included, the firm breaks even in less than 8 months after LED installation. After this initial payback period, the firm benefits from an on-going combined increase in profitability from energy savings and efficiency gains of roughly 3.2% or an increase in their profit margin of .16 percentage points.

6.2 Public Benefits (Emissions Calculations)

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.³⁶ 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

³⁴This calculation comes from management identifying what proportion of total non-material "Cut to Make" costs are recoverable by increasing labor efficiency. "Cut to Make" costs make 25% of the total cost of the garment and the accounting department of the firm estimated that 75% of "Cut to Make" costs were recoverable via efficiency improvements.

³⁵For these calculations, we use the average number of bulbs replaced in the 14 factories we observe before and after LED installation in the production data as these factories best represent the treatment effects estimates.

³⁶The 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.

carbon prices, these benefits are relatively small in comparison to the annual private benefits.³⁷

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. 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. Our findings highlight the potential importance of information dissemination regarding co-benefits, and point the way to government policies designed to subsidize the wide distribution of such information in the private sector, thus increasing economic output and generating environmental benefits, as well.

³⁷ 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

- Adhvaryu, A., Kala, N., and Nyshadham, A. (2016). Management and shocks to worker productivity.
- Allcott, H., Collard-Wexler, A., and O'Connell, S. D. (2014). How do electricity shortages affect productivity? evidence from india. *NBER Working Paper*, (w19977).
- Allcott, H. and Greenstone, M. (2012). Is there an energy efficiency gap? *Journal of Economic Perspectives*, 26(1):3–28.
- Allcott, H. and Taubinsky, D. (2015). Evaluating behaviorally-motivated policy: Experimental evidence from the lightbulb market. *Sloan Foundation Economics Research Paper*, (2640628).
- 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. (2016). Adapting to climate change: The remarkable decline in the us temperature-mortality relationship over the twentieth century. *Journal of Political Economy*, 124(1):105–159.
- Billor, N., Hadi, A. S., and Velleman, P. F. (2000). Bacon: blocked adaptive computationally efficient outlier nominators. *Computational Statistics & Data Analysis*, 34(3):279–298.
- 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.
- 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. (2014). The unequal effects of weather and climate change: Evidence from mortality in india.
- Burke, M., Hsiang, S. M., and Miguel, E. (2015). Global non-linear effect of temperature on economic production. *Nature*.
- Cai, J. and Szeidl, A. (2016). Interfirm relationships and business performance. Technical report, National Bureau of Economic Research.
- Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *The Review of Economics and Statistics*, 90(3):414–427.
- Chang, T., Graff Zivin, J., Gross, T., and Neidell, M. (2014). Particulate pollution and the productivity of pear packers. *NBER Working Paper*, (w19944).
- Costinot, A., Donaldson, D., and Smith, C. (2016). Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world. *Journal of Political Economy*, 124(1):205–248.
- 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.
- Dell, M., Jones, B. F., and Olken, B. A. (2012). Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics*, pages 66–95.
- 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.

- Fowlie, M., Greenstone, M., and Wolfram, C. (2015). Are the non-monetary costs of energy efficiency investments large? understanding low take-up of a free energy efficiency program. *The American Economic Review*, 105(5):201–204.
- Graff Zivin, J. and Neidell, M. (2012). The impact of pollution on worker productivity. *American Economic Review*, 102:3652–3673.
- Granade, H. C., Creyts, J., Derkach, A., Farese, P., Nyquist, S., and Ostrowski, K. (2009). Unlocking energy efficiency in the us economy.
- Guiteras, R. (2009). The impact of climate change on indian agriculture.
- 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. (2016). The effect of pollution on labor supply: Evidence from a natural experiment in mexico city. Technical report.
- Heckman, J., Pinto, R., and Savelyev, P. (2013). Understanding the mechanisms through which an influential early childhood program boosted adult outcomes. *The American Economic Review*, 103(6):1–35.
- Hjort, J. (2014). Ethnic divisions and production in firms. *The Quarterly Journal of Economics*, 129(4):1899–1946.
- 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.
- ISHRAE (2015). The Indian HVAC&R Growth Story. Technical report, The Indian Society of Heating, Refrigerating and Air Conditioning Engineers.
- Jing, S., Li, B., Tan, M., and Liu, H. (2013). Impact of relative humidity on thermal comfort in a warm environment. *Indoor and Built Environment*, 22(4):598–607.

- 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.
- Kudamatsu, M., Persson, T., Strömberg, D., et al. (2012). Weather and infant mortality in africa. Technical report, CEPR Discussion Papers.
- 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.
- Levitt, S. D. and List, J. A. (2011). Was there really a hawthorne effect at the hawthorne plant? an analysis of the original illumination experiments. *American Economic Journal: Applied Economics*, 3(1):224–38.
- Lobell, D. B., Schlenker, W., and Costa-Roberts, J. (2011). Climate trends and global crop production since 1980. *Science*, 333(6042):616–620.
- Mayo, E., Roethlisberger, F., and Dickson, W. (1939). *Management and the worker*.
- Nordhaus, W. (2008). *A question of balance*.
- Parker, P. M. (2000). *Physioeconomics: The basis for long-run economic growth*. MIT Press.
- 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.
- Ryan, N. (2017). *Is there an energy-efficiency gap? experimental evidence from indian manufacturing plants*. Technical report, Working Paper. Yale University.

- Saha, S., Moorthi, S., Pan, H.-L., Wu, X., Wang, J., Nadiga, S., Tripp, P., Kistler, R., Woollen, J., Behringer, D., et al. (2010). The ncep climate forecast system reanalysis. *Bulletin of the American Meteorological Society*, 91(8):1015–1057.
- Seppanen, O., Fisk, W. J., and Lei, Q. (2006). Effect of temperature on task performance in office environment. *Lawrence Berkeley National Laboratory*.
- Snow, C. (1927). Research on industrial illumination. *The Tech Engineering News*, 8(6):257.
- 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., Somanathan, R., Somanathanan, E., and Tewari, M. (2015). 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.
- Wolfram, C., Shelef, O., and Gertler, P. (2012). How will energy demand develop in the developing world? *Journal of Economic Perspectives*, 26(1):119–38.

A Additional Tables

In Figure A1, we present the difference between the productivity-temperature for the with and without LED sample. Essentially, we are presenting the difference between the two gradients presented in Figure 3. Analogously, in Figure A2, we present the difference between the productivity-temperature gradient net of all the fixed effects and budgeted efficiency for the with and without LED sample, which is the difference between the two gradients presented in Figure 5.

In appendix table A1, we present the full results of the distributed lag specification with efficiency as the dependent variable, including lagged coefficients for each of the two splines for seven-day lagged temperatures, and in columns 3 and 4, the interactions of these with the LED installation dummy variable, as well as the LED installation dummy variable. We find that the estimated coefficients for the lagged temperature realizations and those of the lagged temperature realizations interacted with the LED installation dummy variable are largely not statistically significant at conventional levels, and for the coefficients that are statistically significantly different from zero, the point estimates are very small relative to the coefficients on contemporaneous temperature, especially at higher contemporaneous temperatures which is where the impacts on productivity seem most important. Furthermore, as mentioned in the text, both the point estimates and the level of statistical significance of higher contemporaneous temperatures above the 19 degree wet bulb globe temperature cutoff, as well as the mitigative impact of LED lighting of the productivity-temperature gradient at this range of temperatures, remains unchanged. This further supports the conclusion that the results are not driven by omitting lagged temperature variables which might be correlated with contemporaneous temperature.

Appendix table A2 presents the results from the distributed lag specification with mean daily line-level worker attendance as the dependent variable. As with table A1, we include coefficients for each of the two splines for seven-day lagged temperatures, and in the third and fourth columns, interactions of each of these fourteen variables with the LED installation dummy, in addition to the LED installation dummy. The results mirror those of the regression where these lags (and their interactions with LED) are omitted - the estimated coefficients are very small relative to the mean level of attendance, and not statistically significantly different from zero for higher ($WBGT \geq 19$ degrees Celsius) temperatures, which is where the productivity impacts are concentrated. In addition, the contemporaneous point estimates do not change much relative to the specification that omits lagged temperature. Thus, this further supports the conclusion that attendance is unlikely to be the mediating mechanism

for the impacts of temperature on productivity, or the mitigative impact of LED on these effects of temperature.

Appendix table A3 tests that production line-day observations for which data is missing are not correlated with either temperature or the LED installation decisions, and therefore for the purpose of this analysis, can be assumed to be missing randomly. About 20% of all line-day observations are missing efficiency data, since the administrative production data we use are recorded by hand in most factory units, and then inputted daily into the firms data repository servers. This digitization process is imperfect, however, and some data were lost in the process of collating and uploading. The obvious worry here is if temperature is correlated with the incidence of missing production data. We study this in a selection-type regression, in which we create a dummy for missing production line x day observations, and put this in our baseline specification: temperature splines with a node at wet-bulb globe temperature of 19 degrees C, and all baseline fixed effects. We also regress the dummy for missing production line-day observations on all the fixed effects and the LED installation dummy variable. The results of this analysis are presented in Table A3. The coefficients on the temperature splines are very small and statistically insignificant, indicating that conditional on baseline fixed effects, temperature or LED installation does not predict missing observations.

We then test the robustness of our results to alternative temperature measures, namely dry bulb temperature (controlling for daily relative humidity) and the Heat Index (HI). Tables A4-A6 present the results for dry bulb temperature, and tables A7-A9 for the heat index. The spline nodes for these alternate temperature variables are chosen to correspond to roughly the same proportion of observations as the spline for wet bulb globe temperature, and the productivity-temperature gradient exhibits a significant negative gradient at these chosen nodes. This is reflected in the regression results. Controlling for precipitation does not affect any of the results.

The spline regression estimates in Table A4 from columns 1 and 2 illustrate that the slope of the efficiency-temperature gradient below 27 degrees Celsius of dry bulb temperature is slightly negative (statistically indistinguishable from 0) and the slope above 27 degrees is strongly negative and statistically significant at the 1 percent level, with a magnitude of about -1.02 percentage points of efficiency. As with the main specification with wet bulb globe temperature, the introduction of LED offsets the negative impacts of temperature on efficiency by roughly 85% attenuating the magnitude of the negative slope above 27 degrees, but has no significant impact at lower temperatures, which do not statistically significantly affect efficiency. The estimate of the main effect of LED is statistically

indistinguishable from 0.

Next, we test whether higher temperatures impact worker attendance. Table A5 reports results of the regression of mean line-daily worker attendance on the identical specifications to those in Table A4. Analogous to results that use wet bulb globe temperature as the temperature measure, we find that temperature shocks at only low temperatures (less than 27 degree Celsius) affect attendance, and the magnitudes of the point estimates are extremely small (less than 1% of the mean). All other estimates of coefficients, including those reflecting the impacts of LED, are statistically indistinguishable from 0. As with the main specifications, this suggests that the impacts of temperature on worker attendance are not contributing to the estimated impacts of temperature and LED installation on efficiency.

Table A6 confirms that attendance is not the primary mediating mechanism for the impacts of temperature on efficiency. It reports the results of regressions identical to those in Table A4, with the additional inclusion of mean line-daily worker attendance as a control variable. Once more, as with the main specifications that use wet bulb globe temperature as the temperature measure, including mean line-daily worker attendance as a control variable does not significantly change the estimated impacts of temperature, or the mitigative impact of LED. Thus, using dry bulb temperature as an alternative temperature measure supports the conclusion that while higher temperatures impact the intensive margin of productivity per unit labor supplied, but does not seem to affect much the extensive margin of the quantity of labor units supplied. Furthermore, LED lighting installation attenuates the impacts of temperature on the intensive margin of efficiency, but has no perceptible impact on the extensive margin of attendance.

Tables A7-A9 report the analogous results to Tables A4-A6, except that Heat Index is used as the measure of temperature. Once more, the results are in line with our main results (using wet bulb globe temperature as the temperature measure), and results obtained using dry bulb temperature as the temperature measure. The productivity-temperature gradient is strongly negative at higher temperatures (Heat Index ≥ 27 degrees Celsius), with a marginal impact of about -1.3 percentage points of efficiency, whereas it is statistically indistinguishable from zero at lower temperatures. The installation of LED lighting mitigates about 80% of the impact of temperature on productivity, in accordance with results that use alternate temperature measures. Furthermore, this impact does not appear to be driven by attendance, as shown by the results in Tables A8 and A9. Table A8 presents the impacts of higher temperature on attendance, and finds at most weakly statistically significant impacts only at lower temperatures (Heat Index < 27 degrees Celsius), and no impact at higher temperatures (Heat

Index ≥ 27 degrees Celsius). Table A9 shows that controlling for mean line-daily worker attendance does not impact the estimated impacts of temperature on efficiency, or the mitigative impact of LED, thereby confirming that attendance is likely not the primary mechanism driving either the estimated productivity-temperature gradient, or the mitigative impact of LED lighting of this gradient.

Finally, in tables A10-A12, we exclude factories that had LED lighting at the beginning of the sample or did not receive LED lighting by the end of the sample, and re-estimate the main regressions. Table A10 shows impacts on efficiency, A11 shows impacts on attendance, and A12 shows impacts on efficiency controlling for daily line-level attendance. We find that the results on the temperature-efficiency gradient, as well as mitigation due to LED, and the null impacts on attendance, remain unchanged across these three tables with the amended sample that excludes these factories. We see again that results are very similar (both in coefficient magnitudes and patterns of statistical significance) to results obtained by including all factories with available productivity data.

Figure A1: Difference in Gradient by LED

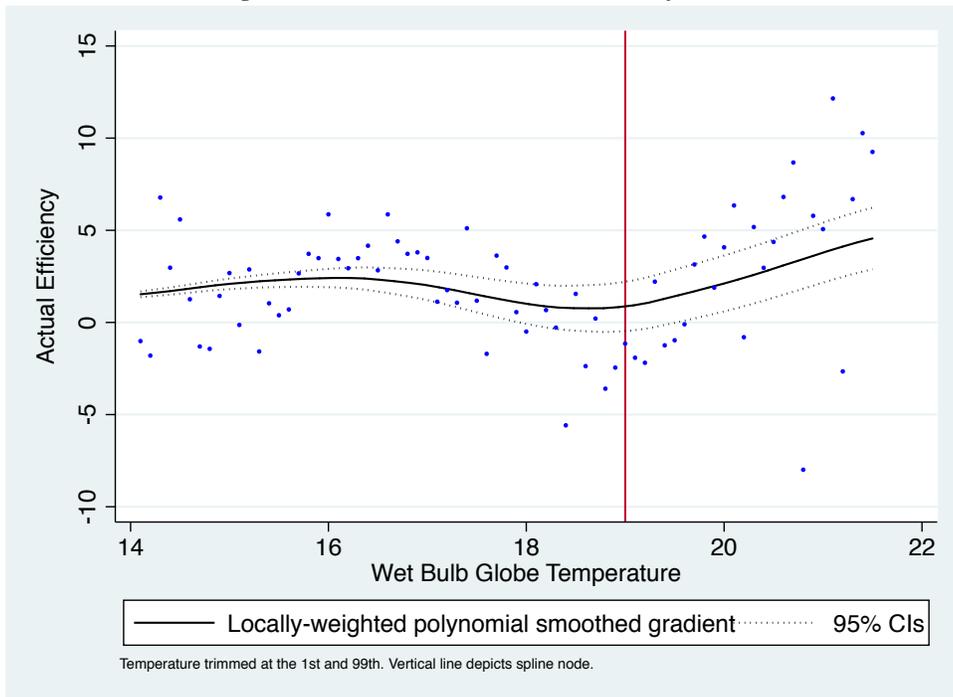


Figure A2: Difference in Residual Gradient by LED

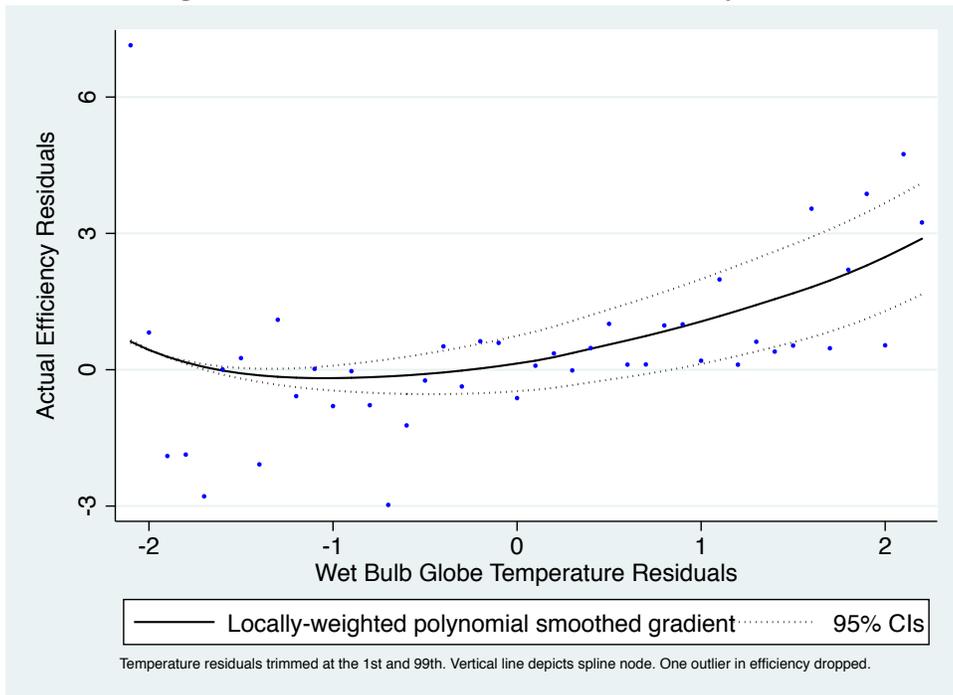


Table A1
Impact of Temperature on Production Efficiency and Mitigative Impact of LED Lighting
(Distributed Lag Specification)

	(1)	(2)	(3)	(4)
	Actual Efficiency (Actual Production / Targeted Production)*100			
Wet Bulb Globe Temperature <19	-0.44 (0.29)	-0.467 (0.26)	-0.227 (0.41)	-0.245 (0.35)
Wet Bulb Globe Temperature >=19	-2.236*** (0.002)	-2.271*** (0.002)	-2.295*** (0.002)	-2.32*** (0.002)
One-Day Lagged Wet Bulb Globe Temperature <19	-0.105 (0.3)	-0.102 (0.32)	-0.215* (0.06)	-0.212* (0.07)
One-Day Lagged Wet Bulb Globe Temperature >=19	-0.056 (0.43)	-0.05 (0.46)	0.026 (0.68)	0.029 (0.64)
Two-Day Lagged Wet Bulb Globe Temperature <19	0.136 (0.29)	0.139 (0.27)	0.059 (0.64)	0.061 (0.62)
Two-Day Lagged Wet Bulb Globe Temperature >=19	-0.045 (0.79)	-0.043 (0.79)	-0.054 (0.73)	-0.054 (0.74)
Three-Day Lagged Wet Bulb Globe Temperature <19	-0.093 (0.37)	-0.089 (0.39)	-0.15 (0.24)	-0.146 (0.25)
Three-Day Lagged Wet Bulb Globe Temperature >=19	0.25 (0.19)	0.252 (0.19)	0.302 (0.18)	0.303 (0.18)
Four-Day Lagged Wet Bulb Globe Temperature <19	0.029 (0.75)	0.029 (0.75)	-0.022 (0.69)	-0.022 (0.69)
Four-Day Lagged Wet Bulb Globe Temperature >=19	0.154 (0.43)	0.151 (0.43)	0.216 (0.36)	0.214 (0.37)
Five-Day Lagged Wet Bulb Globe Temperature <19	0.227*** (0.01)	0.226*** (0.01)	0.085 (0.45)	0.084 (0.45)
Five-Day Lagged Wet Bulb Globe Temperature >=19	0.165 (0.36)	0.165 (0.36)	0.258 (0.19)	0.257 (0.19)
Six-Day Lagged Wet Bulb Globe Temperature <19	0.053 (0.38)	0.061 (0.36)	-0.023 (0.85)	-0.019 (0.89)
Six-Day Lagged Wet Bulb Globe Temperature >=19	-0.366*** (0.006)	-0.367*** (0.006)	-0.335** (0.05)	-0.337** (0.05)
Seven-Day Lagged Wet Bulb Globe Temperature <19	0.029 (0.54)	0.036 (0.48)	-0.008 (0.92)	-0.003 (0.96)
Seven-Day Lagged Wet Bulb Globe Temperature >=19	0.194 (0.34)	0.191 (0.34)	0.253 (0.30)	0.25 (0.31)
1(LED)*Wet Bulb Globe Temperature <19			0.072 (0.8)	0.078 (0.79)
1(LED)*Wet Bulb Globe Temperature >=19			2.375*** (0.0)	2.384*** (0.0)
1(LED)*One-Day Lagged Wet Bulb Globe Temperature <19			0.23 (0.13)	0.229 (0.13)
1(LED)*One-Day Lagged Wet Bulb Globe Temperature >=19			-0.206 (0.11)	-0.206 (0.11)
1(LED)*Two-Day Lagged Wet Bulb Globe Temperature <19			0.122 (0.34)	0.122 (0.34)
1(LED)*Two-Day Lagged Wet Bulb Globe Temperature >=19			-0.044 (0.85)	-0.045 (0.84)
1(LED)*Three-Day Lagged Wet Bulb Globe Temperature <19			-0.006 (0.98)	-0.007 (0.97)
1(LED)*Three-Day Lagged Wet Bulb Globe Temperature >=19			-0.263 (0.27)	-0.262 (0.27)
1(LED)*Four-Day Lagged Wet Bulb Globe Temperature <19			0.262*** (0.01)	0.263*** (0.01)
1(LED)*Four-Day Lagged Wet Bulb Globe Temperature >=19			-0.506* (0.08)	-0.504** (0.08)
1(LED)*Five-Day Lagged Wet Bulb Globe Temperature <19			-0.156 (0.19)	-0.158 (0.18)
1(LED)*Five-Day Lagged Wet Bulb Globe Temperature >=19			-0.277 (0.19)	-0.276 (0.19)
1(LED)*Six-Day Lagged Wet Bulb Globe Temperature <19			0.064 (0.56)	0.064 (0.56)
1(LED)*Six-Day Lagged Wet Bulb Globe Temperature >=19			0.193 (0.34)	0.192 (0.34)
1(LED)*Seven-Day Lagged Wet Bulb Globe Temperature <19			0.076 (0.62)	0.072 (0.63)
1(LED)*Seven-Day Lagged Wet Bulb Globe Temperature >=19			-0.375 (0.15)	-0.375 (0.15)
1(LED)			-9.199 (0.40)	-9.165 (0.40)
Fixed Effects	Factory x Year, Factory x Calendar Month, Production Line, Day of the Week			
Precipitation Control	N	Y	N	Y
Observations	74,939	74,939	239,680	239,680
Mean of Dependent Variable	53.732	53.732	55.23	55.23

Notes: Wild-cluster bootstrap p-values in parentheses (***) denotes significance at the 1% level, ** denotes significance at the 5% level, * denotes significance at the 10% level. Clustering is done at the factory level. All measures of temperature are in degree Celsius. All regressions include daily budgeted efficiency as a control variable.

Table A2
Impact of Temperature on Attendance and Mitigative Impact of LED Lighting (Distributed Lag Specification)

	(1)	(2)	(3)	(4)
	Line-Level Mean Daily Probability of Worker Presence			
Wet Bulb Globe Temperature <19	-0.0112*** (0.002)	-0.011*** (0.002)	-0.0081*** (0.008)	-0.0075*** (0.02)
Wet Bulb Globe Temperature >=19	-0.0012 (0.61)	-0.0009 (0.68)	0.0007 (0.82)	0.0013 (0.64)
One-Day Lagged Wet Bulb Globe Temperature <19	0.002 (0.186)	0.0019 (0.19)	0.002 (0.10)	0.0019 (0.14)
One-Day Lagged Wet Bulb Globe Temperature >=19	0.0023*** (0.004)	0.0022*** (0.00)	0.003*** (0.0)	0.003*** (0.002)
Two-Day Lagged Wet Bulb Globe Temperature <19	0.0015 (0.27)	0.0015 (0.27)	0.0019 (0.14)	0.0018 (0.14)
Two-Day Lagged Wet Bulb Globe Temperature >=19	-0.0037 (0.33)	-0.0037 (0.33)	-0.0035 (0.29)	-0.0035 (0.27)
Three-Day Lagged Wet Bulb Globe Temperature <19	-0.0001 (1.00)	-0.0001 (1.00)	-0.0001 (0.95)	-0.0002 (0.93)
Three-Day Lagged Wet Bulb Globe Temperature >=19	-0.0012 (0.73)	-0.0012 (0.73)	-0.0008 (0.76)	-0.0008 (0.76)
Four-Day Lagged Wet Bulb Globe Temperature <19	-0.0006 (0.58)	-0.0007 (0.57)	-0.0004 (0.74)	-0.0004 (0.73)
Four-Day Lagged Wet Bulb Globe Temperature >=19	0.0032* (0.09)	0.0032* (0.08)	0.0037** (0.02)	0.0038** (0.01)
Five-Day Lagged Wet Bulb Globe Temperature <19	0.0071*** (0.03)	0.0071*** (0.03)	0.0069** (0.04)	0.0069** (0.04)
Five-Day Lagged Wet Bulb Globe Temperature >=19	0.0023*** (0.01)	0.0023*** (0.01)	0.0045*** (0.0)	0.0046*** (0.0)
Six-Day Lagged Wet Bulb Globe Temperature <19	0.0016 (0.47)	0.0015 (0.52)	0.0017 (0.48)	0.0015 (0.54)
Six-Day Lagged Wet Bulb Globe Temperature >=19	0.0037** (0.05)	0.0037** (0.05)	0.0048*** (0.0)	0.0049*** (0.0)
Seven-Day Lagged Wet Bulb Globe Temperature <19	-0.0003 (0.75)	-0.0004 (0.71)	-0.0007 (0.53)	-0.0008 (0.42)
Seven-Day Lagged Wet Bulb Globe Temperature >=19	-0.0005 (0.49)	-0.0005 (0.54)	0.0001 (0.98)	0.0001 (0.93)
1(LED)*Wet Bulb Globe Temperature <19			0.0017 (0.64)	0.0015 (0.69)
1(LED)*Wet Bulb Globe Temperature >=19			0.0002 (0.93)	0.0001 (0.97)
1(LED)*One-Day Lagged Wet Bulb Globe Temperature <19			0.0006 (0.67)	0.0007 (0.64)
1(LED)*One-Day Lagged Wet Bulb Globe Temperature >=19			-0.0003 (0.79)	-0.0003 (0.79)
1(LED)*Two-Day Lagged Wet Bulb Globe Temperature <19			-0.0026* (0.09)	-0.0026* (0.08)
1(LED)*Two-Day Lagged Wet Bulb Globe Temperature >=19			-0.0021 (0.47)	-0.0021 (0.46)
1(LED)*Three-Day Lagged Wet Bulb Globe Temperature <19			0.0014 (0.78)	0.0014 (0.78)
1(LED)*Three-Day Lagged Wet Bulb Globe Temperature >=19			-0.0015 (0.51)	-0.0015 (0.50)
1(LED)*Four-Day Lagged Wet Bulb Globe Temperature <19			0.0041** (0.02)	0.0041** (0.02)
1(LED)*Four-Day Lagged Wet Bulb Globe Temperature >=19			-0.0031 (0.10)	-0.0031 (0.10)
1(LED)*Five-Day Lagged Wet Bulb Globe Temperature <19			-0.0013 (0.46)	-0.0012 (0.49)
1(LED)*Five-Day Lagged Wet Bulb Globe Temperature >=19			-0.0005 (0.68)	-0.0005 (0.68)
1(LED)*Six-Day Lagged Wet Bulb Globe Temperature <19			-0.0046** (0.01)	-0.0046** (0.02)
1(LED)*Six-Day Lagged Wet Bulb Globe Temperature >=19			-0.0016 (0.32)	-0.0016 (0.32)
1(LED)*Seven-Day Lagged Wet Bulb Globe Temperature <19			0.0058*** (0.0)	0.0058*** (0.0)
1(LED)*Seven-Day Lagged Wet Bulb Globe Temperature >=19			-0.0005 (0.79)	-0.0005 (0.80)
1(LED)			-0.0846 (0.20)	-0.086 (0.20)
Fixed Effects	Factory x Year, Factory x Calendar Month, Production Line, Day of the Week			
Observations	136,062	136,062	392,601	392,601
Mean of Dependent Variable	0.85	0.85	0.83	0.83

Notes: Wild-cluster bootstrap p-values in parentheses (** p<0.01, * p<0.05, † p<0.1). Clustering is done at the factory level. All measures of temperature are in degree Celsius.

Table A3
 Partial Correlations between Missing Data and Temperature and LED Regressors

	(1)	(2)	(3)
	1(Efficiency Data is Missing)		
Wet Bulb Globe Temperature <19	0.0075 (0.26)	0.0075 (0.25)	
Wet Bulb Globe Temperature >=19	-0.00438 (0.92)	-0.00434 (0.92)	
1(LED)			-0.0165 (0.81)
Fixed Effects	Factory x Year, Factory x Calendar Month, Production Line, Day of the Week		
Precipitation Control	N	Y	N
Observations	95,526	95,526	356,924
Mean of Dependent Variable	0.21	0.21	0.33

Notes: Wild-cluster bootstrap p-values in parentheses (***) denotes significance at the 1% level, ** denotes significance at the 5% level, * denotes significance at the 10% level). Clustering is done at the factory level. All measures of temperature are in degree Celsius.

Table A4
Impact of Dry Bulb Temperature on Efficiency and Mitigative Impact of LED lighting

	(1)	(2)	(3)	(4)
	Efficiency			
Dry Bulb Temperature <27	-0.136 (0.55)	-0.1501 (0.51)	-0.0274 (0.84)	-0.0333 (0.82)
Dry Bulb Temperature >=27	-1.022*** (0.002)	-1.0279*** (0.002)	-0.9305*** (0.002)	-0.9392*** (0.002)
1(LED)*(Dry Bulb Temperature <27)			-0.0752 (0.75)	-0.0743 (0.75)
1(LED)*(Dry Bulb Temperature >=27)			0.7954*** (0.01)	0.8032*** (0.01)
1(LED)			3.412 (0.65)	3.3883 (0.66)
Fixed Effects	Factory x Year, Factory x Calendar Month, Production Line, Day of the Week			
Precipitation Control	N	Y	N	Y
Observations	74,939	74,939	239,680	239,680
Mean of Dependent Variable	53.73	53.73	55.234	55.234

Notes: Wild-cluster bootstrap p-values in parentheses (***) denotes significance at the 1% level, ** denotes significance at the 5% level, * denotes significance at the 10% level). Clustering is done at the factory level. All measures of temperature are in degree Celsius. All regressions include daily budgeted efficiency as a control variable.

Table A5

Impact of Dry Bulb Temperature on Attendance and Mitigative Impact of LED lighting

	(1)	(2)	(3)	(4)
	Worker Presence (Line-Level Mean Daily Probability)			
Dry Bulb Temperature <27	-0.0037** (0.02)	-0.0037** (0.04)	-0.0002 (0.86)	-0.00017 (0.90)
Dry Bulb Temperature >=27	0.00145 (0.38)	0.00146 (0.37)	0.0059** (0.04)	0.00595** (0.04)
1(LED)*(Dry Bulb Temperature <27)			0.0002 (0.85)	0.0002 (0.85)
1(LED)*(Dry Bulb Temperature >=27)			-0.004 (0.22)	-0.0041 (0.21)
1(LED)			-0.0083 (0.73)	-0.00876 (0.71)
Fixed Effects	Factory x Year, Factory x Calendar Month, Production Line, Day of the Week			
Precipitation Control	N	Y	N	Y
Observations	136,062	136,062	392,601	392,601
Mean of Dependent Variable	0.846	0.846	0.829	0.829

Notes: Wild-cluster bootstrap p-values in parentheses (***) denotes significance at the 1% level, ** denotes significance at the 5% level, * denotes significance at the 10% level). Clustering is done at the factory level. All measures of temperature are in degree Celsius.

Table A6
Impact of Dry Bulb Temperature on Efficiency and Mitigative Impact of LED lighting Controlling for Attendance

	(1)	(2)	(3)	(4)
	Efficiency			
Dry Bulb Temperature <27	-0.2265 (0.41)	-0.2453 (0.37)	-0.06437 (0.70)	-0.07 (0.67)
Dry Bulb Temperature >=27	-1.1928*** (0.004)	-1.1991*** (0.004)	-1.0193*** (0.004)	-1.02872*** (0.006)
1(LED)*(Dry Bulb Temperature <27)			-0.00515 (1.00)	-0.0047 (1.00)
1(LED)*(Dry Bulb Temperature >=27)			0.7265** (0.04)	0.74** (0.04)
1(LED)			1.4719 (0.89)	1.4562 (0.90)
Line-Level Mean Daily Probability of Worker Presence	1.7789 (0.49)	1.7386 (0.49)	2.1314** (0.04)	2.1197** (0.04)
Fixed Effects	Factory x Year, Factory x Calendar Month, Production Line, Day of the Week			
Precipitation Control	N	Y	N	Y
Observations	74,939	74,939	239,680	239,680
Mean of Dependent Variable	53.73	53.73	55.234	55.234

Notes: Wild-cluster bootstrap p-values in parentheses (***) denotes significance at the 1% level, ** denotes significance at the 5% level, * denotes significance at the 10% level). Clustering is done at the factory level. All measures of temperature are in degree Celsius. Line-Level Mean Daily Probability of Worker Presence is the average probability that a worker is present for a production line on a given day. All regressions include daily budgeted efficiency as a control variable.

Table A7
Impact of Heat Index on Efficiency and Mitigative Impact of LED lighting

	(1)	(2)	(3)	(4)
	Efficiency			
Heat Index <26	-0.151 (0.53)	-0.161 (0.50)	-0.0455 (0.77)	-0.051 (0.75)
Heat Index >=26	-1.324*** (0.002)	-1.347*** (0.002)	-1.184*** (0.002)	-1.202*** (0.002)
1(LED)*(Heat Index <26)			-0.051 (0.83)	-0.0497 (0.83)
1(LED)*(Heat Index >=26)			0.983*** (0.01)	0.989*** (0.01)
1(LED)			2.804 (0.70)	2.767 (0.70)
Fixed Effects	Factory x Year, Factory x Calendar Month, Production Line, Day of the Week			
Precipitation Control	N	Y	N	Y
Observations	74,939	74,939	239,680	239,680
Mean of Dependent Variable	53.73	53.73	55.234	55.234

Notes: Wild-cluster bootstrap p-values in parentheses (***) denotes significance at the 1% level, ** denotes significance at the 5% level, * denotes significance at the 10% level). Clustering is done at the factory level. All measures of temperature are in degree Celsius. All regressions include daily budgeted efficiency as a control variable.

Table A8
Impact of Heat Index on Attendance and Mitigative Impact of LED lighting

	(1)	(2)	(3)	(4)
Worker Presence (Line-Level Mean Daily Probability)				
Heat Index <26	-0.00337** (0.04)	-0.00329* (0.06)	-0.00067 (0.54)	-0.00045 (0.71)
Heat Index >=26	0.00018 (0.92)	0.0004 (0.83)	0.003987 (0.23)	0.00449 (0.18)
1(LED)*(Heat Index <26)			0.00076 (0.52)	0.0007 (0.57)
1(LED)*(Heat Index >=26)			-0.00446 (0.23)	-0.00461 (0.22)
1(LED)			-0.01725 (0.46)	-0.01673 (0.48)
Fixed Effects	Factory x Year, Factory x Calendar Month, Production Line, Day of the Week			
Precipitation Control	N	Y	N	Y
Observations	136,062	136,062	392,601	392,601
Mean of Dependent Variable	0.846	0.846	0.829	0.829
Notes: Wild-cluster bootstrap p-values in parentheses (***) denotes significance at the 1% level, ** denotes significance at the 5% level, * denotes significance at the 10% level). Clustering is done at the factory level. All measures of temperature are in degree Celsius.				

Table A9

Impact of Heat Index on Efficiency and Mitigative Impact of LED lighting Controlling for Attendance

	(1)	(2)	(3)	(4)
	Efficiency			
Heat Index <26	-0.238 (0.40)	-0.252 (0.37)	-0.079 (0.63)	-0.085 (0.61)
Heat Index >=26	-1.545*** (0.004)	-1.582*** (0.004)	-1.299*** (0.006)	-1.32*** (0.006)
1(LED)*(Heat Index <26)			0.012 (0.91)	0.014 (0.90)
1(LED)*(Heat Index >=26)			0.922** (0.042)	0.93** (0.03)
1(LED)			1.079 (0.92)	1.043 (0.93)
Line-Level Mean Daily Probability of Worker Presence	1.887 (0.47)	1.887 (0.48)	2.188** (0.03)	2.195** (0.03)
Fixed Effects	Factory x Year, Factory x Calendar Month, Production Line, Day of the Week			
Precipitation Control	N	Y	N	Y
Observations	74,939	74,939	239,680	239,680
Mean of Dependent Variable	53.73	53.73	55.234	55.234

Notes: Wild-cluster bootstrap p-values in parentheses (***) denotes significance at the 1% level, ** denotes significance at the 5% level, * denotes significance at the 10% level). Clustering is done at the factory level. All measures of temperature are in degree Celsius. Line-Level Mean Daily Probability of Worker Presence is the average probability that a worker is present for a production line on a given day. All regressions include daily budgeted efficiency as a control variable.

Table A10

Impact of Temperature on Actual Efficiency and Mitigative Impact of LED lighting
(Omitting Units that Always Had LED/Never Got LED)

	(1)	(2)	(3)	(4)
	Actual Efficiency			
Wet Bulb Globe Temperature <19	-0.526 (0.50)	-0.549 (0.47)	-0.201 (0.59)	-0.223 (0.55)
Wet Bulb Globe Temperature >=19	-2.953*** (0.004)	-3.0283*** (0.004)	-2.579*** (0.008)	-2.681*** (0.008)
1(LED)*(Wet Bulb Globe Temperature <19)			-0.222 (0.59)	-0.225 (0.59)
1(LED)*(Wet Bulb Globe Temperature >=19)			2.261*** (0.00)	2.319*** (0.00)
1(LED)			5.165 (0.58)	5.205 (0.58)
Fixed Effects	Factory x Year, Factory x Calendar Month, Production Line, Day of the Week			
Precipitation Control	N	Y	N	Y
Observations	43,291	43,291	134,326	134,326
Mean of Dependent Variable	57.372	57.372	56.34	56.34

Notes: Wild-cluster bootstrap p-values in parentheses (***) denotes significance at the 1% level, ** denotes significance at the 5% level, * denotes significance at the 10% level). Clustering is done at the factory level. All measures of temperature are in degree Celsius. All regressions include daily budgeted efficiency as a control variable.

Table A11

Impact of Temperature on Attendance and Mitigative Impact of LED lighting
(Omitting Units that Always Had LED/Never Got LED)

	(1)	(2)	(3)	(4)
Worker Presence (Line-Level Mean Daily Probability)				
Wet Bulb Globe Temperature <19	-0.00971*** (0.002)	-0.00964*** (0.002)	-0.00251 (0.16)	-0.00213 (0.26)
Wet Bulb Globe Temperature >=19	0.00452 (0.24)	0.00478 (0.20)	0.01109* (0.07)	0.01193* (0.06)
1(LED)*(Wet Bulb Globe Temperature <19)			0.00151 (0.40)	0.00142 (0.44)
1(LED)*(Wet Bulb Globe Temperature >=19)			-0.0125* (0.09)	-0.01283* (0.09)
1(LED)			-0.02482 (0.38)	-0.02437 (0.40)
Fixed Effects	Factory x Year, Factory x Calendar Month, Production Line, Day of the Week			
Precipitation Control	N	Y	N	Y
Observations	84,435	84,435	273,801	273,801
Mean of Dependent Variable	0.83	0.83	0.823	0.823

Notes: Wild-cluster bootstrap p-values in parentheses (***) denotes significance at the 1% level, ** denotes significance at the 5% level, * denotes significance at the 10% level). Clustering is done at the factory level. All measures of temperature are in degree Celsius.

Table A12

Impact of Temperature on Actual Efficiency and Mitigative Impact of LED lighting Controlling for Attendance (Omitting Units that Always Had LED/Never Got LED)

	(1)	(2)	(3)	(4)
	Actual Efficiency			
Wet Bulb Globe Temperature <19	-0.619 (0.41)	-0.646 (0.39)	-0.227 (0.53)	-0.248 (0.49)
Wet Bulb Globe Temperature >=19	-3.217*** (0.004)	-3.302*** (0.004)	-2.719*** (0.006)	-2.813*** (0.006)
1(LED)*(Wet Bulb Globe Temperature <19)			-0.171 (0.73)	-0.172 (0.73)
1(LED)*(Wet Bulb Globe Temperature >=19)			2.125*** (0.01)	2.178*** (0.01)
1(LED)			4.012 (0.67)	4.022 (0.67)
Line-Level Mean Daily Probability of Worker Presence	2.327 (0.49)	2.33 (0.50)	2.569** (0.04)	2.592** (0.04)
Fixed Effects	Factory x Year, Factory x Calendar Month, Production Line, Day of the Week			
Precipitation Control	N	Y	N	Y
Observations	41,400	41,400	121,694	121,694
Mean of Dependent Variable	57.58	57.58	56.76	56.76

Notes: Wild-cluster bootstrap p-values in parentheses (***) denotes significance at the 1% level, ** denotes significance at the 5% level, * denotes significance at the 10% level). Clustering is done at the factory level. All measures of temperature are in degree Celsius. Line-Level Mean Daily Probability of Worker Presence is the average probability that a worker is present for a production line on a given day. All regressions include daily budgeted efficiency as a control variable.

A Data Appendix

We have daily line-level data from 30 factories in Bangalore. To ensure precise 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 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 once 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 with positive production and with efficiency less than or equal to 200%. These decisions were taken following meetings with the Industrial Engineering experts at the factory regarding what constitutes feasible values of output and efficiency. However, results are extremely similar if we trim outliers using blocked adaptive computationally efficient outlier nominators (Billor et al., 2000), which uses Mahalanobis distance to compute outliers, and are available upon request.
- Finally, we remove Sundays, since these days of unrepresentative productivity (production does not usually occur on Sundays, unless factories are working over-time to finish orders).