



PEDL Research Papers

This research was partly or entirely supported by funding from the research initiative Private Enterprise Development in Low-Income Countries (PEDL), a Foreign, Commonwealth & Development Office (FCDO) funded programme run by the Centre for Economic Policy Research (CEPR).

This is a PEDL Research Paper which emanates from a PEDL funded project. Any views expressed here are those of the author(s) and not those of the programme nor of the affiliated organizations. Although research disseminated by PEDL may include views on policy, the programme itself takes no institutional policy positions.

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

Achyuta Adhvaryu, Namrata Kala, and Anant Nyshadham*

Abstract—We study the adoption of energy-efficient LED lighting in garment factories around Bangalore, India. Combining daily production line-level data with weather data, we estimate a negative, nonlinear productivity-temperature gradient. We find that LED lighting raises productivity on hot days. Using the firm’s costs data, we estimate that the payback period for LED adoption is less than one-third the length after accounting for productivity co-benefits. The average factory in our data gains about \$2,880 in power consumption savings and about \$7,500 in productivity gains.

I. Introduction

INNOVATIONS in energy efficiency and regulation-driven adoption of energy-efficient technologies have been cited as a primary means of curbing the acceleration of climate change (Granade et al., 2009). Despite this promise, energy-efficient technologies are usually adopted at low rates (Allcott & Taubinsky, 2015). Recent studies point to several explanations for this “energy-efficiency gap.” The first is market failures such as information frictions or credit constraints that drive a wedge between socially and privately optimal adoption (Allcott & Greenstone, 2012). The second is behavioral factors such as consumer inattention to energy costs (Allcott, Mullainathan, & Taubinsky, 2014). The third possible explanation is that returns are smaller, or costs higher, in practice than engineering projections predict (Burlig et al., 2017; Fowlie, Greenstone, & Wolfram, 2013; Ryan, 2017). Furthermore, behavioral responses to energy-efficiency (such as increased consumption) may offset returns to energy efficiency investments. Thus, estimating the true returns to energy efficiency requires testing for mechanisms that may drive a wedge between engineering and economic returns, including imperfect maintenance of the investments and rebound effects.

In this study, we estimate the productivity consequences of the adoption of energy-saving technology, using daily pro-

duction line data from a large garment firm operating factories in and around Bangalore, India. First, we show that days with higher outside temperatures have lower productivity, measured as production line efficiency (realized output over target output). We then show that the replacement of compact fluorescent lamps (CFLs) with light-emitting-diode (LED) lighting on factory floors attenuates the negative relationship between mean daily outdoor temperature and efficiency. Driven by buyers’ environmental standards, factories replaced a substantial fraction of CFL bulbs with LED bulbs. LED lighting reduces ambient temperature on the factory floor because less electricity is converted to waste heat, relative to CFL lighting. This lower ambient temperature reduces the effect of higher outside temperature on efficiency. We study the impacts of the staggered rollout of LEDs over more than three years on the sewing floors of 26 garment factories.¹ We use rich administrative data on worker attendance, working hours, and productivity to test for mechanisms that would mitigate or offset the returns to energy-efficient lighting. We also demonstrate in a variety of checks that the timing of the rollout across factories was not systematically related to business processes or working conditions, such as time of the start or end of the workday, total working hours, wages, or the composition of hiring patterns by worker skill levels.

Our measure of mean daily temperature exposure, wet bulb globe temperature (WBGT), takes into account both temperature and humidity, since the impact of temperature on thermal regulation varies by humidity levels. Impacts of outdoor temperature on productive efficiency, estimated using a spline regression (controlling for factory by year, factory by month, production line, and day of the week fixed effects), are quite nonlinear: for mean daily WBGT of below 19°C (the temperature equivalent at average humidity levels in our sample is 27° to 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.² We then estimate the extent to which the introduction of LED lighting, likely through the reduced dissipation of heat on factory floors, flattens the temperature-productivity gradient. LED installation has no impact on the gradient below the 19°C WBGT cutoff but attenuates the negative slope of the gradient by more than 80% for temperatures above this threshold. Our results are robust to the inclusion of a variety of fixed effects and controls,

¹Our data examine thirty factories (all owned by the same garment firm), four of which did not receive LED lighting.

²This nonlinear gradient is remarkably consistent with the physiology of temperature effects: at high ambient temperatures, the body loses the ability to dissipate heat, which negatively affects performance (Hancock et al., 2007).

Received for publication January 31, 2018. Revision accepted for publication April 5, 2019. Editor: Rohini Pande.

*Adhvaryu: University of Michigan, NBER, BREAD, and Good Business Lab; Kala: MIT Sloan School of Management, BREAD, NBER, and JPAL; Nyshadham: Boston College, NBER, University of Michigan, and Good Business Lab.

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, Nicholas 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. A.A. gratefully acknowledges funding from the NIH/NICHHD (5K01HD071949), and all of us acknowledge funding from PEDL, a CEPR-sponsored grant initiative. All errors are our own.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/rest_a_00886.

including factory by year by quarter fixed effects, as well as alternative specifications such as semiparametric estimation. 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 and is likely due to a leftward movement along the gradient. This movement would generate large increases in efficiency in high temperature ranges and small efficiency increases elsewhere.³

While engineering estimates of the heat dissipation of LED (versus those of CFL) bulbs exist, those estimates are not always reflective of economic returns, as Fowlie et al. (2013) and Burlig et al. (2017) showed recently. In our setting too, a field study has several advantages in estimating the true productivity returns to energy efficiency. First, if factories respond to energy savings by increasing working hours, then the cobenefits to these investments may change: they may be higher if workers respond to the more comfortable environment on hotter days by continuing to be more productive for extra hours, and they may be lower or 0 if workers respond to longer hours by slowing their productivity per hour. Using data on working hours, we can directly test for this response by the factory managers. Second, if the temperature-productivity relationship is driven by lower attendance on hotter days, and not by workers responding to a less comfortable work environment, then LED lighting may not mitigate this relationship (e.g., temperatures outside of working hours may affect workers' health, and therefore their propensity to attend work). Using data on worker attendance, we can rule out that this is the case. Third, if workers respond to the lighting by changing their attendance (either because they are now more comfortable or because they are uncomfortable with the new lighting), the productivity cobenefits may be higher or lower. Finally, our results indicate that energy-efficient lighting can generate these co-benefits in settings where workers are exposed to heat generated by conventional bulbs, and air-conditioning is not cost-effective (which is typical of manufacturing workplaces in low-income countries).

Finally, we perform cost-benefit calculations for LED adoption, combining the above estimates with the firm's actual cost data for LED replacement and projected energy savings. 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 significantly shifts the break-even point for the firm, from over three and half years to less than eight months. With some assumptions on how worker productivity translates into profits (detailed in section VII), we estimated that the average factory gained about \$2,880 in power consumption savings and about \$7,500 in productivity gains.

³One major drawback of our study is that we do not have indoor temperature data in the factories before and after LED installation. Thus, other aspects of LED lighting that affect the productivity-temperature gradient such as unmeasured light quality changes may contribute to the aggregate effect of LED lighting mitigating the productivity-temperature relationship, as long as these unmeasured changes affect productivity only on hotter days.

Our study contributes to the literature on the returns to climate change mitigation and energy efficiency. Recent studies have indicated that energy-efficient lighting can both reduce electricity consumption (Burlig et al., 2017) and generate positive externality co-benefits such as greater electricity grid reliability (Carranza & Meeks, 2020). Other studies that examine co-benefits, or additional gains, of climate change mitigation broadly speaking, such as carbon taxes, also focus largely on the indirect public returns (Knittel & Sandler, 2011; see IPCC, 2013, for a review). We study 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. By our estimation, ignoring the productivity benefits of LEDs would significantly underestimate the private returns to adoption.

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, Kala, & Nyshadham, 2016; Allcott, Collard-Wexler, & O'Connell, 2014; Hsiang, 2010; Sudarshan et al., 2015) and adaptation to higher temperatures.⁴ The impacts of temperature on productivity appear to hold quite consistently across countries and time (Burke, Hsiang, & Miguel, 2015; Dell, Jones, & Olken, 2012). A related literature has established patterns of adaptation to climate change and the returns to this adaptation (Barreca et al., 2016). Our results indicate that energy-efficient lighting can be a form of adaptation to higher temperatures in settings characterized by low air-conditioning adoption and significant indoor heat exposure from conventional lighting. Our results thus highlight an interaction between high temperatures and the co-benefits of energy-efficient technologies.

The remainder of the paper is organized as follows. Section II describes contextual details regarding garment production in India and the LED installation. Section III provides details on the temperature and production data. Section IV describes our empirical strategy. Section V describes the results, section VI offers additional robustness checks, and section VII reviews the cost-benefit analysis and concludes.

II. Context

A. *Physiology of the Temperature-Productivity Gradient*

The physical impact of temperature on human beings is well studied (Enander, 1989; Parsons, 2010; Seppanen, Fisk, & Lei, 2006) and has been important for establishing occupational safety standards for workers exposed to very high or

⁴Several recent studies document this relationship in more developed settings (Chang et al., 2014; Costinot, Donaldson, & Smith, 2016; Graff Zivin & Neidell, 2012; Hanna & Oliva, 2015).

low temperatures for extended periods of time (Vanhoorne, Vanachter, & De Ridder, 2006). Thermal stress can affect human beings physically and through lower psychomotor ability and degraded perceptual task performance (Hancock, Ross, & Szalma, 2007). The impact on individual subjects varies based on factors such as the type of task and its complexity, duration of exposure, and the worker-level skill and acclimatization level (Pilcher, Nadler, & Busch, 2002). This contributes to the difficulty of setting a specific 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 U: performance suffers at excessively cold and excessively warm temperatures (Parsons, 2010). Moreover, one meta-analysis highlights the dry-bulb threshold of 29.4°C (85°F) as particularly important (Hancock et al., 2007). This threshold value represents the temperature above which the body starts to store heat. 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 29.4°C threshold.

B. *Measuring Garment Productivity and Overview of the LED Installation*

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). Garments are usually sewn in production lines in manufacturing plants. Each line produces a single style of garment at a time (possibly with varying colors or sizes) until the order for that garment is met. Lines consist of sixty to seventy sewing machine operators (depending on the complexity of the style) arranged in sequence and grouped in terms of parts of the garment (e.g., sleeve, collar).⁵ 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.

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 past decade, buyers have become more stringent in their regulation of their suppliers’ production and environmental standards. This prompted a staggered rollout 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 become more stringent, then the supplier might choose

⁵In general, we describe here the process for woven garments; however, the steps are quite similar for knits and even pants, with a varying number and complexity of operations. Even within wovens, the production process varies slightly by style or factory.

to upgrade to LED lighting in factories processing many orders from buyer A. When buyer B’s regulations change, the firm will prioritize factories servicing buyer B, and so on.⁶ One thing to note is that there are still CFL bulbs in all factories after the change. That is, only about half the bulbs were replaced, with each fixture now containing one CFL bulb instead of two.

The replacement took the form of substituting a portion of CFLs 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 factory replaced about 1,200 CFLs consuming 7 W each with LED lights of 1 W each.⁷ The LED light bulbs that replaced the CFLs in the factories in our data require about 3 as opposed to 21 kWh/year in electricity in our setting, and thus operate at about one-seventh the cost of CFL lighting.⁸ Based on the factories’ operating time cost calculation, this meant an energy saving of 18 kWh per bulb per year. Heat emissions for a single LED bulb are 3.4 Btus, compared to 23.8 Btus for a single CFL lighting bulb.⁹ In section VII, we discuss the magnitude of the environmental benefits from the installation.

Each factory received the installation within a single month. Eight percent of the LED rollout (2 factories) was completed in 2009, 48% (12 factories) in 2010, 16% (4 factories) in 2011, about 24% (6 factories) in 2012, and the rest (1 factory in 2013). Of the 30 factories from which we have productivity data, LED replacements occurred in 26 factories during the observation period. Since our productivity data range from April 2010 to June 2013, some factories already had LEDs at the beginning of our productivity data, and all

⁶We check for the endogeneity of LED adoption in tables 5 and 6 and conduct other robustness checks. We find little evidence that LED adoption at the factory level was correlated with a variety of business operations and outcomes.

⁷The number of lights installed is a function of the number of machines in the factory and varies from about 100 to 2,550, with a mean of about 1,200.

⁸It should be noted that there are many varieties of LED and CFL bulbs. The energy and lighting specifications and calculations presented and discussed in this paper are specific to the bulbs used in the factory replacements in our data and do 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.

⁹Changing factory lighting may have consequences for productivity through mechanisms other than temperature changes, as highlighted by the results of the original Hawthorne lighting experiment (Snow, 1927; Mayo et al., 1939), as well as new analysis by Levitt & 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. However, any unmeasured light quality changes that affect the temperature gradient in addition to indoor temperature would form part of our estimates of LED lighting to mitigate the efficiency-temperature gradient. We cannot distinguish the two, since we do not have measurements of indoor temperature before and after the study.

but four factories had LEDs by the end of our sample period. Figure A1 in the appendix presents the cumulative proportion of factories adopting LED against mean temperature.¹⁰

III. Data

A. Weather Data

We use mean 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 reanalysis data set that uses historical station-level and satellite data combined with climate models to produce a consistent record of gridded weather variables from 1979 to 2014. It has a spatial resolution of about 38 km; each factory in our sample is matched to the nearest data grid point.¹¹

We use a temperature index that incorporates temperature and humidity. 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 we 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 (WBGT) 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 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 (%). Both measures of temperature—dry bulb temperature and WBGT—are converted into Celsius to ensure interpretative ease across regression specifications.

Note that the weather data we use are mean daily outdoor temperature measures. While indoor temperature in the factory is what would affect worker productivity, we do not have data on indoor temperature from the period of the LED rollout. Accordingly, we use outdoor ambient temperature as

¹⁰Regression results that omit factories 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.

¹¹There are eight temperature grid points in our sample. The factories are located in and around Bangalore city, so while they are not clustered in a particular part of the city, the identification is largely coming from the time series variation in temperature. The reanalysis data allow us to exploit this cross-sectional relationship slightly better. There are eight reanalysis data points and only one station in Bangalore that regularly report weather data across our sample period that we found in the Global Historical Climate Network (GHCN) data. If we compare the time series of the mean daily temperatures from our eight reanalysis points (averaged over each day) with the mean daily temperature from the Bangalore weather station, the correlation in daily temperatures is about 0.8, which seems to suggest that the reanalysis data correlate reasonably well with the station-level data.

discussed above as a proxy for indoor conditions. 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 about 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.¹²

In figure 1, we plot mean indoor temperature values for each 0.1 degree bin of outdoor temperature along with a local polynomial regression fit curve and 95% confidence intervals.¹³ 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 temperatures (e.g., 22°C 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, in addition to heat generated by workers' presence on the factory floor). Furthermore, a regression of indoor temperature on outdoor temperature has an R^2 of about 0.84, implying that a very large amount of the variation in indoor temperature is explained by the variation in outdoor temperature. However, it is important to 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.

B. Factory Data

We use daily data at the production line level from thirty garment factories in and around Bangalore. Identifiers include factory number and production line number within the factory. For each line and day within each factory, production measures include actual quantity of garments produced and target quantities of the line on that day.


Actual efficiency is actual quantity produced divided by target quantity. The target quantity is derived from an industrial engineering measure for the complexity of the garment—the “Standard allowable minute (SAM), which is 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 adjustments by the factory's industrial engineering (IE) department during “sampling.”¹⁴

¹²We collected data from September 22, 2014, to August 11, 2015, in one factory, from September 27, 2014, to August 10, 2015, in a second factory, and from January 28, 2015, to August 10, 2015, in a third factory.

¹³Fit reflects kernel-weighted local mean smoothing, using the Epanechnikov kernel.

¹⁴Sampling is the process by which a cost estimate is generated for a buyer when ordering a garment style. Sampling tailors make a garment of

FIGURE 1.—INDOOR TEMPERATURE VERSUS OUTDOOR TEMPERATURE



The SAM measure is used to calculate the target quantity for the line for each hour of production. Each line runs for eight hours during a standard workday from 9 a.m. to 5 p.m., with all factories in our sample operating a single daytime production shift. Accordingly, a line producing a style with a SAM of 0.5 will have a target of 120 garments per hour, or 960 garments per day. Most important, 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 produces only a single style at a time.¹⁵ Variations in expected achievable efficiency over the life of a particular garment order due to order size are reflected in a measure that incorporates learning by doing, budgeted efficiency. Budgeted efficiency remains fixed for a given line over the life of a particular order and reflects the efficiency that management believes a line might be able to achieve given the expected length of time the line will be producing the order. Actual efficiency of a given order will vary systematically across lines and within a line over time due to, for example, absenteeism, machine failures, or working conditions. We are interested in variation in actual efficiency due to transitory temperature. We therefore control for budgeted efficiency to account for systematic variation in efficiency deriving from order size and include line fixed effects in the regression analysis that follows. In the robustness checks, we show that our results are not affected by excluding this control variable.

a particular style and recommend any alterations to the SAM for that style to the IE department.

¹⁵Indeed, in our data, lines produce styles for between 1 and 268 days.

TABLE 1.—SUMMARY STATISTICS: WEATHER, PRODUCTION, AND LED INTRODUCTION


	Mean	Standard Deviation
Number of line-day observations	239,680	
Number of lines	523	
Number of days	1,001	
Number of factories	30	
Weather		
Temperature (Celsius)	24.353	2.966
Relative humidity (%)	0.647	0.174
Wet bulb globe temperature (Celsius)	17.230	1.683
Production		
Actual efficiency	55.234	26.233
Budgeted efficiency	61.981	11.545
Standard allowable minutes (SAM)	0.724	2.445
Attendance		
1(Present for Full Work day)	0.843	0.363

C. Summary Statistics

We present means and standard deviations of variables used in the analysis in table 1. Our sample consists of 523 production lines across thirty factories. The range of dates over which we have production data spans 1,001 days. However, we do not observe all factories for all dates.¹⁶ Altogether, our data include nearly 240,000 line by day observations. About one-third of the observations correspond to days in factories prior to the introduction of LED lighting, and the remainder are post-LED observations.

¹⁶Once a factory starts reporting data, it continues to do so until the end of the sample period. In the appendix, we restrict the analysis of the main productivity specifications to only production lines that have a proportion of missing data less than or equal to 30% of observations.

FIGURE 2.—EFFICIENCY AGAINST TEMPERATURE (PRE-LED)



IV. Empirical Strategy

In this section, we provide preliminary graphs 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 to develop an empirical strategy to flexibly estimate the impact of LED introduction on productivity as moderated through ambient temperature.

A. Descriptive Evidence

We begin by motivating the empirical specifications and techniques with descriptive plots of production and temperature data.¹⁷

Productivity-temperature gradient. To estimate how LED lights have an impact on the relationship between efficiency and temperature, 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% confidence intervals like those depicted in figure 1.¹⁸ Figure 2 shows that in the absence of LED lighting, efficiency appears to be a decreasing function of temperature, and this relationship is quite non-linear, with the largest 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°C 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°C corresponds in our data to an outdoor ambient dry bulb temperature of about 27°C and is likely equivalent to an indoor dry bulb temperature of about 29.5°C before LED introduction.¹⁹ This 29.5°C dry bulb temperature is quite consistent with estimates from previous studies on the physiological threshold for the absorption of heat into the body, above which temperature affects human functioning (Hancock et al., 2007).


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

¹⁷Residualized graphs with fixed effects and controls mentioned in section IVB look very similar and are available on request.

¹⁸Fit reflects kernel-weighted local mean smoothing, using the default Epanechnikov kernel and bandwidth of 1.

¹⁹This approximate relationship is derived from the indoor-outdoor temperature we collected and back-of-the-envelope calculations about how LED affected internal temperature. While a full engineering projection of heat dissipation is beyond the scope of this study, we present a simple heat gain calculation. The difference in energy consumption is 18 kWh per bulb per year, which translates into 0.058 kWh per bulb per day (assuming a six-day workweek). For the average factory, which received 1,000 LED bulbs, that implies a lowered electricity consumption of 58 kWh/day. Taking the heat capacity of air as 1 joule/(g °C) and the density of air as 1.18 kg/m³, 58 kWh would heat 73,700 m³ of air (or, for instance, a factory of 192 by 192 m square with a height of 2 m) by 2.4°C. This temperature difference is a significant ambient temperature difference that would explain our results, a calculation based on the fact that 1 kWh is 3.6 million joules, and heating 1 m³ of air requires $1 \times 1.18 \times 1,000 \times 2.4 = 2832$ joules = 0.00079 kWh. (We thank one of our referees for suggesting this back-of-the-envelope calculation.)

FIGURE 3.—EFFICIENCY AGAINST TEMPERATURE BY LED



check for evidence that this gradient is affected by the partial replacement of the CFLs 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 rollout in each factory. These plots are presented in figure 3.²⁰ The evidence suggests that factories are somewhat more efficient at all temperatures after the LED introduction, but this gain (or attenuation) increases at high temperatures. That is, the pre-LED gradient (red line) in figure 3 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.

B. Parametric Spline Regression Analysis

Motivated by the graphical evidence we estimate the regression equations 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 LED replacement. In particular, we address concerns regarding factory-level trends in efficiency, line-level unobservables, seasonality in efficiency, and the exogeneity of the LED introduction along with the nonlinearities depicted in figures 2 and 3.

²⁰Fit reflects kernel-weighted local mean smoothing, using the Epanechnikov kernel and bandwidth of 1. Note in each figure from here onward 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.

First, we estimate the following empirical specification of the relationship between production line efficiency and temperature using only observations prior to LED installation:

$$E_{ludmy} = \alpha_0 + \beta^L T_{dgm}^L + \beta^H T_{dgm}^H + \phi B_{ludmy} + \alpha_l + \gamma_{uy} + \eta_{um} + \delta_d + \varepsilon_{ludmy}. \quad (2)$$

Here, E is the 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°C 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 by year fixed effects; η_{um} are unit by 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°C Celsius increase in wet bulb globe temperature on line-level efficiency for temperatures below and above 19°C, respectively.²¹

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

$$E_{ludmy} = \alpha_0 + \beta_1^L (T_{dgm}^L \times LED_{umy}) + \beta_1^H (T_{dgm}^H \times LED_{umy}) + \beta_2 LED_{umy} + \beta_3^L T_{dgm}^L + \beta_3^H T_{dgm}^H + \phi B_{ludmy} + \alpha_l + \gamma_{uy} + \eta_{um} + \delta_d + \varepsilon_{ludmy}. \quad (3)$$

²¹While the effect of temperature on productivity may vary within the day, this is not testable given our data, since we only observe mean productivity and outdoor temperature for a production line each day.

Here LED_{umy} is a dummy for the 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°C 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°C spline node, respectively, once LED lighting is introduced. The sums $\beta_1^L + \beta_3^L$ and $\beta_1^H + \beta_3^H$ give 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°C WBGT for two reasons: (a) 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 (b) this point corresponds remarkably well to previous studies of the physiology of heat stress (Hancock et al., 2007).²²

To account for common error distributions within a factory over time, standard errors are clustered at the factory level. This cluster structure is appropriate given that LED introduction occurs at the unit level. However, given the relatively small number of clusters (thirty), we employ wild cluster bootstrap inference and report 95% confidence intervals in parentheses in all tables unless otherwise noted.²³

Attendance. We also estimate the same specifications presented in equations (2) and (3), 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 affects efficiency, and the corresponding attenuation from LED introduction might be working through effects on worker attendance. In robustness checks, we also estimate the original efficiency specifications from equations (2) and (3), including mean line-daily worker attendance as an additional control. The combination of these two sets of results allows us to investigate whether temperature and LED introduction affect worker attendance and whether controlling for attendance changes the estimated impacts of temperature and LED on the primary outcome of interest (efficiency).

Distributed lags. Daily temperature could reflect short-term serial correlation, which would make it difficult to identify the impacts of contemporaneous exposure to temperature. Following previous studies, we augment equations 2

and 3 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 lagged exposure. This isolates the impact of contemporaneous exposure from that of lagged exposure. If the coefficients on the contemporaneous temperature terms are similar with and without the inclusion of the 7-day distributed lag terms, we interpret the results as indicating a minimal role for serial correlation and persistence in impacts of lagged exposures. We can recover the composite impact of both the contemporaneous temperature exposure and 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 equations (2) and (3).

V. Results

A. Main Results

We report results from the estimation of the parametric spline specifications presented in equations (2) and (3) in table 2. Columns 1 and 2 report estimates of β^L and β^H from equation (2), with column 2 estimates corresponding to a specification with an additional control for precipitation. The precipitation control ensures that impacts are being driven by temperature exposure alone and are not composite effects reflecting the impacts of other correlated weather conditions. Columns 3 and 4 report estimates of β_1^L , β_1^H , β_2 , β_3^L , and β_3^H from equation (3), once again with column 4 reporting results after controlling for precipitation.

The spline regression estimates from columns 1 and 2 reflect the pattern shown in figure 2 with the slope of the efficiency-temperature gradient below 19°C of wet bulb globe temperature being slightly negative (statistically indistinguishable from 0) and the slope above 19°C being strongly negative and statistically significant at the 1% level. Point estimates indicate that at wet bulb globe temperatures above 19°C, a 1°C 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 shows that including an additional control for precipitation has a minimal impact on results.

The results in columns 3 and 4 are consistent with the pattern reflected in figure 3, with the introduction of LED having no significant impact on the slope of the efficiency-temperature gradient below 19°C, but a significant attenuating impact on the negative slope of the gradient above 19°C. That is, the introduction of LED offsets the negative impacts of temperature on efficiency by about 85%, attenuating the magnitude of the negative slope above 19°C from about -2 to about -0.3 . LED shows no significant impact below 19°C, which is consistent with the ergonomics and physiology

²²Nevertheless, we explored more flexible spline specifications with more nodes and found the results to be qualitatively identical with less precision.

²³See Cameron, Gelbach, and Miller (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.

TABLE 2.—IMPACT OF TEMPERATURE ON PRODUCTION EFFICIENCY AND MITIGATIVE IMPACT OF LED LIGHTING

	1	2	3	4
	Efficiency (Actual Production / Targeted Production) × 100			
Wet bulb globe temperature <19	-0.299 [-1.803, 0.532]	-0.318 [-1.813, 0.510]	-0.0940 [-1.017, 0.421]	-0.105 [-1.008, 0.404]
Wet bulb globe temperature ≥19	-2.135*** [-3.312, -1.395]	-2.169*** [-3.369, -1.399]	-1.953** [-3.00, -1.206]	-1.981*** [-3.020, -1.230]
1(LED) × (Wet bulb globe temperature <19)			-0.106 [-0.847, 0.852]	-0.103 [-0.843, 0.853]
1(LED) × (Wet bulb globe temperature ≥19)			1.671*** [0.718, 2.787]	1.681*** [0.725, 2.809]
1(LED)			3.447 [-18.34, 16.85]	3.393 [-18.39, 16.85]
Fixed effects	Factory × Year, Factory × 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

Wild-cluster bootstrap 95% CIs in brackets: significant at ***1%, **5%, and *10%. Clustering is done at the factory level. All measures of temperature are in degrees 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** [-0.0166, -0.000630]	-0.0059** [-0.0167, -0.000305]	-0.0011 [-0.00534, 0.00162]	-0.0007 [-0.00506, 0.00204]
Wet bulb globe temperature ≥19	0.0003 [-0.00585, 0.00973]	0.0007 [-0.00533, 0.00989]	0.0056 [-0.00409, 0.0186]	0.0064 [-0.00341, 0.0193]
1(LED) × (Wet bulb globe temperature <19)			0.0003 [-0.00283, 0.00442]	0.0002 [-0.00298, 0.00428]
1(LED) × (Wet bulb globe temperature ≥19)			-0.0051 [-0.0197, 0.00771]	-0.0053 [-0.0199, 0.00743]
1(LED)			-0.0065 [-0.0708, 0.0427]	-0.0054 [-0.0701, 0.0441]
Fixed effects	Factory × Year, Factory × 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

Wild-cluster bootstrap 95% CIs in brackets: significant at ***1%, **5%, and *10%. Clustering is done at the factory level. All measures of temperature are in degrees Celsius.

literature, suggesting that temperature has the highest impact on human functioning at temperatures above this level. The estimate of the main effect of LED is positive and large, but it 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 IVB. The estimates from table 3 suggest a negative impact of temperature on attendance at temperatures below 19°C, 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 indicating no real impacts of temperature on worker attendance. These results imply that it is unlikely that the impact of temperature on worker attendance contributes

to the estimated impacts of temperature and LED installation on efficiency.

Next, we investigate whether the estimated impacts of contemporaneous temperature exposure on efficiency reflect contemporaneous exposure alone rather than a composite of contemporaneous exposure and 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 our analysis, the interpretation of the point estimates will change based on the underlying sources of variation. As discussed in section IV, we repeat the analysis reported in table 2 but include seven-day distributed lag temperature spline terms and, where appropriate, their interactions with LED installation. The results are reported in table 4. All results in table 4 correspond to specifications including seven-day distributed lag

TABLE 4.—IMPACT OF TEMPERATURE ON PRODUCTION EFFICIENCY AND MITIGATIVE IMPACT OF LED LIGHTING (DISTRIBUTED LAG SPECIFICATION)

	1	2	3	4
		Efficiency (Actual Production / Targeted Production) × 100		
Wet bulb globe temperature < 19	−0.440 [−1.895, 0.415]	−0.467 [−1.900, 0.380]	−0.227 [−1.191, 0.365]	−0.245 [−1.198, 0.342]
Wet bulb globe temperature ≥ 19	−2.236*** [−3.462, −1.359]	−2.271*** [−3.479, −1.367]	−2.295*** [−3.678, −1.229]	−2.320*** [−3.686, −1.244]
1(LED) × wet bulb globe temperature < 19			0.0724 [−0.510, 0.908]	0.0776 [−0.504, 0.902]
1(LED) × wet bulb globe temperature ≥ 19			2.375*** [1.234, 3.726]	2.384*** [1.233, 3.740]
1(LED)			−9.199 [−42.54, 11.93]	−9.165 [−42.55, 11.92]
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 × Year, Factory × 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

Wild-cluster bootstrap 95% CIs in brackets: significant at ***1%, **5%, and *10%. Clustering is done at the factory level. All measures of temperature are in degrees Celsius. All regressions include daily budgeted efficiency as a control variable. Full table reporting coefficients on seven-day distributed lag temperature splines and their interactions with LED are presented in appendix table A1.

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 4 are qualitatively identical to the main results reported in table 2, but with slightly larger magnitudes for coefficients on the above 19°C temperature spline and the corresponding LED interaction terms. These results indicate that the estimates of temperature impacts and attenuation from LED installation are being driven by contemporaneous exposures and not the correlation of contemporaneous and lagged temperature. Daily temperature is generally believed to reflect some degree of serial correlation, so the similarity in results with and without distributed lags is not altogether surprising in our study. The baseline specifications already include a large set of heterogeneous nonlinear trends (e.g., unit by month FE) to control for this less transitory variation in temperature. Indeed, the correlations between contemporaneous temperature and lagged temperature values after partialing out the full set of controls are not very high (never more than .25 and mostly below .1).²⁴

B. Checks for Exogeneity of LED Roll-Out

In this section, we check for the exogeneity of the timing of LED installation.

In column 1 of table 5, 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 (2). We find no evidence that LED installation was timed around particular temperature realizations. In columns 2 and 3 of table 5, we report results from the regression of SAM (a proxy for the complexity of the garments being pro-

duced) and budgeted efficiency (a proxy for scope for learning by doing due to order size), 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 (2). 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 suggest that the timing of LED introduction is endogenous with respect to these production factors; however, we find no such evidence. 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 in columns 2 and 3.²⁵ We do not find any evidence that hiring patterns changed in the lead-up to LED installation or immediately after.

Finally, we test whether there were changes to working hours or wage contracts in the lead-up to LED installation. We use two other data sources for this purpose. The first data source is daily data at the worker level showing when individuals clocked in and out of work, which we use to construct the two measures of daily line-level working hours. The first is the average time at the line level that we observe a worker clock in for a given line on a given day (measured in terms of elapsed minutes since midnight). The second is the average time that we observe a worker on a line clock out on a given day (also measured in terms of elapsed minutes since midnight). These measures are at the production-line daily

²⁴The coefficients on the lag terms are reported in appendix table A1.

²⁵A-grade tailors are the most skilled, followed by B-grade tailors, and C-grade tailors are the least skilled.

TABLE 5.—CHECKS FOR EXOGENEITY OF LED ROLLOUT: PRODUCTION AND HIRING

	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
WBGT <19	0.00172 [−0.00797, 0.0103]					
WBGT ≥19	−0.00776 [−0.0210, 0.00569]					
1(LED) × Date relative to LED installation		−0.0001 [−0.00098, 0.00159]	0.0153 [−0.0284, 0.0518]	0.00001 [−0.00062, 0.00045]	0.00008 [−0.00017, 0.0003]	−0.00005 [−0.0004, 0.00043]
1(LED)		−0.0386 [−0.171, 0.0894]	2.931* [−4.637, 8.661]	−0.0205 [−0.0808, 0.0375]	0.0436 [−0.0232, 0.119]	−0.0194 [−0.0635, 0.0311]
Date relative to LED installation		0.0001 [−0.0005, 0.0008]	−0.0208 [−0.0421, 0.0048]	0.00009 [−0.0013, 0.00161]	0.00129* [−0.00003, 0.00284]	−0.00123 [−0.0032, 0.0004]
Fixed effects		Factory × Year, Factory × Calendar Month Production Line, Day of the Week		Factory × Year, Factory × Calendar Month		
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

Wild-cluster bootstrap 95% CIs in brackets: significant at ***1%, **5%, and *10%. Clustering is done at the factory level. Since columns 2 through 6 consider the date relative to 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 three columns are at the production line-date level, and the last three columns are defined at the factory-date level.

TABLE 6.—CHECKS FOR EXOGENEITY OF LED ROLLOUT: WORKING HOURS, WAGES, AND PAY DAYS

	1	2	3	4	5	6
	Mean Time In	Mean Time Out	Earned Wages	Total Deductions	Unpaid Leave	Number of Paydays
1(LED) × Date relative to LED installation	−0.0120 [−0.0631, 0.0131]	−0.0248 [−0.161, 0.0196]	−23.02 [−52.12, 13.33]	−4.190 [−9.149, 4.116]	4.449 [−26.20, 25.63]	−0.0526 [−0.170, 0.108]
Date relative to LED installation	0.00575 [−0.00438, 0.0151]	0.0189 [−0.0993, 0.188]	−6.604 [−707733, 592734]	1,294 [−99108, 158182]	3,141 [−33032, 501873]	583.6 [−91.99, 2455]
1(LED)	2.133 [−2.924, 7.739]	−2.459 [−11.16, 10.26]	−42.68 [−141, 146.9]	−6.833 [−32.27, 18.26]	86.03 [−53.07, 167.3]	−0.317 [−0.724, 0.313]
Fixed effects		Factory × Year, Factory × Calendar Month Production Line, Day of the Week		Factory × Year, Factory × Calendar Month Production Line		
Precipitation controls	Y	Y	Y	Y	Y	Y
Temperature controls	Y	Y	Y	Y	Y	Y
Observations	93,198	93,279	566,761	566,761	566,761	585,401
Mean of the dependent variable	537.4	1,058	5,629	602.3	982.3	25.34

95% wild cluster bootstrap confidence intervals in brackets: significant at ***1%, **5%, and *10%. Clustering is done at the factory level. Since this table considers the date relative to 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. All times are in minutes elapsed since midnight. The unit of observation for columns 1 and 2 is the production line-daily level, and for columns 3 through 6 is at the worker-monthly level. Earned wages are total wages accruing to the employee. Total deductions include contributions to the provident fund, taxes, and employee social security. Unpaid leave refers to wages that were unpaid because the employee took leave without pay.

level.²⁶ The second data source is at the monthly level for each individual worker and measures different aspects of the wage contract as well as the number of paydays each month for the worker. The four variables we consider are wages earned (total wages accruing to the employee, including deductions for taxes, social security, and employees' provident fund), total deductions (contributions to taxes, social security, and employees' provident fund), value of unpaid leave, and the number of days each month that the employee was present and accrued wages (the number of pay days).

Table 6 presents the results. Overall, we do not find that changes in working hours or compensation to workers change with LED adoption; there is no statistically significant change in any of the six variables leading up to LED adoption.

VI. Additional Robustness Checks

We conduct a variety of additional robustness checks.

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 line-daily worker attendance as an additional control. The results from these regressions are reported in table A6 and are very similar to those presented in table 2.

²⁶Results using the earliest and latest times that a worker on a line clocks in and out give very similar results.

Overall, we interpret the results in tables 3 and A6 as evidence against the importance of attendance as a primary mediator of the has an impact on 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 strongly affect 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 of labor supply.

We also present all of our main results with additional fixed effects. We replace factory by year fixed effects with factory by year by quarter fixed effects (and still include factory by calendar month, production line, and day of the week fixed effects, along with daily precipitation as a control variable). Table A3 presents the results using this specification. Columns 1 and 2 present the impact of temperature on productivity and the mitigating impact of LED lighting. Columns 5 and 6 present the impact of temperature on productivity and the mitigating impact of LED lighting controlling for line-level attendance, and columns 3 and 4 do so while including the same splines for seven-day distributed lag specification of temperature as for same-day temperature (as well as with interactions of the LED adoption dummy variable). The results are robust to the inclusion of these fixed effects, though the magnitude reduces a little in some specifications. These results suggest that any unmeasured changes to the working environment that we are not able to pick up in the previous robustness checks but that happened within the quarter are not driving our results.

In table A2, we present our main results without including budgeted efficiency as a control variable, to show that our results do not depend on using this control variable. In table A4, we restrict the analysis of the main productivity specifications to only production lines that have a proportion of missing data that is lower than or equal to 30% (this includes data from 344 of about 500 production lines).²⁷ The results are nearly identical to our original results, suggesting that the missing data are not substantially affecting our main results.²⁸

In table A5, we replicate results from table 2, except that we use dry bulb temperature instead of WBGT and control for relative humidity separately. Finally, in the online appendix, we present our main results using a more flexible semiparametric estimation rather than parametric spline regressions, and obtain very similar estimates.

²⁷We divide the number of days that a factory reported productivity data by the total number of days between when a factory began reporting data until the end of the sample period, excluding Sundays, to compute the proportion of missing data. We chose this proportion because once a factory starts reporting data, on average the probability that a production line is missing productivity data on a given day is about 32%.

²⁸We also check that the probability of missing productivity data is not affected by temperature or LED adoption.

VII. Discussion

The promise of climate change mitigation is tempered by the willingness of individuals and firms to adopt beneficial technologies on a large scale. This willingness in turn, is a function of the private returns to adoption. In this study, we show that the introduction of energy-efficient LED lighting in Indian garment factories had substantial productivity co-be-nefits that accrue privately to the adopting firm.

Specifically, we find that the introduction of LEDs eliminates about 85% percent of the negative impact of temperature on worker efficiency on relatively hot days. Using the probability that mean daily outdoor temperature reaches or exceeds 19°C WBGT (20% of all days), we estimate an average total increase in production efficiency of 0.4 percentage points (0.2 times the mitigation coefficient on LED lighting, which is 1.95 percentage points of efficiency).²⁹

A. Private Benefits (Firm Cost-Benefit Calculations)

We combine our estimate of average total efficiency gains with actual production and cost data from the firm to calculate annual costs and benefits of LED installation. The calculations are shown in table 7. Senior management at the firm estimated that the profit gain for each percentage point gain in efficiency was 0.1875 percentage points.³⁰ Thus, a 0.4 percentage point gain in efficiency from LED installation translates to a .075 percentage point gain in profits (or a 1.5% increase in profitability from the 5% baseline profit margin of the firm). At an approximate profit per factory per year of \$500,000 (amounts in US dollars) the introduction of LED delivers productivity gains worth \$7,500 per factory per year.

How does this estimate change the cost-benefit calculations of LED adoption for the firm? We calculated this based on the energy cost calculations the firm used for its LED adoption choices. 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 per bulb replaced or \$2,880 in total for an average replacement of about 1,200 bulbs per factory in our sample.³¹ The additional annual profits from efficiency gains we computed are more than two and a half times this amount. The cost of

²⁹Using a semiparametric estimation strategy in the online appendix gives a higher impact of LEDs, with an average total increase in production efficiency of about 0.7 percentage points. However, we use the lower of the two numbers here to be conservative.

³⁰This calculation comes from management identifying what proportion of total nonmaterial costs are recoverable by increasing labor efficiency. These costs make 25% of the total cost of the garment, and the accounting department of the firm estimated that 75% of these costs were recoverable via efficiency improvements.

³¹For these calculations, we use the average number of bulbs replaced in the factories we observe before and after LED installation in the production data, as these factories best represent the treatment effects estimates. Ideally, we would be able to corroborate the engineering estimates of electricity savings with electricity billing data. Unfortunately, data limitations prevent this. Since working hours in the factory did not change with LED lighting adoption and the bulbs were run continuously throughout these hours, it seems plausible that rebound effects are not very large in this context.

TABLE 7.—COST-BENEFIT CALCULATIONS FOR LED ADOPTION

Cost of implementation (one-time)	
Investment per bulb (e.g., bulb, wiring)	\$8.53
Number of bulbs replaced per factory	1200
Total cost of implementation	\$10,240.00
Energy savings (per year)	
Power consumption savings per bulb (with savings of 18 kWh per year, at a cost of about \$0.1334/kWh)	\$2.40
Number of bulbs replaced	1,200
Total energy savings	\$2,880.00
Efficiency gains (per year)	
Average efficiency gain in percentage points from LED-induced temperature reductions	0.4000
Efficiency percentage point gain to profit percentage point gain translation	0.1875
Profit margin at baseline	5%
Average revenue in USD per factory per year	\$10,000,000.00
Average profit in USD per factory per year	\$500,000.00
Total efficiency gains	\$7,500.00
Total net savings from LED Adoption in the first year	\$140.00
Total net savings from LED Adoption in the second year	\$10,380.00
Carbon (public) benefits from LED adoption (at \$44/tC)	\$197.04
Carbon (public) benefits from LED adoption (at \$93/tC)	\$416.47

Profit margin at base was taken from the accounting department of the firm. Calculation for translating efficiency gains to profit comes from accounting estimates of the proportion of “cut to make” (nonmaterial) 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 per year is obtained by taking balance sheet revenues for the firm and dividing by the number of plants. For additional details on the calculation of carbon benefits, refer to section VII.

replacing the average factory’s bulbs with LEDs is \$10,240. Thus, if only electricity expenditure were taken into account, it would take about three and a half years to break even. However, when the productivity benefits are included, the firm breaks even within twelve months of LED installation. After the initial payback period, the firm benefits from an ongoing combined increase in profitability from energy savings and efficiency gains.

These results are of course generalizable only to settings where air-conditioning is not available in the workplace. However, since air-conditioning remains quite rare in factories in developing countries, our results indicate that energy-efficient lighting can have substantial co-benefits for worker productivity for a large section of the workforce.

B. Public Benefits (Emissions Calculations)

In addition to the private benefits of increased productivity and energy cost savings, the replacement of CFL lighting with LEDs has public benefits of avoided damages due to reduced carbon emissions. On average, the LED introduction saves 21,600 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.³² Valuing this reduction of carbon emissions at the Nordhaus (2011) estimate of \$44/tC (a 2005 carbon price) gives us avoided damages of \$197.04 per factory per year, and valuing this at the mean value of the

³²The 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 megawatt of electricity generated on the Southern grid causes 0.76 tCO₂ of emissions (Bhawan & Puram, 2014). Thus, 18,000 kWh causes about 16.42 tCO₂, or about 4.48 tC. Callaway, Fowlie, and McCormick (2018) use US electricity generation data to show that the average emissions factor of electricity, which we use here, may be quite different from the marginal emissions factor. While the latter is preferable to generate precise values of carbon generated by electricity, the lack of data on marginal emissions factors makes this impossible in our setting.

review by Tol (2005) of \$93/tC yields avoided damages of \$416.47 per factory per year. At the current estimates of carbon prices, these benefits are relatively small in comparison to the annual private benefits.³³

³³Adding 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., N. Kala, and A. Nyshadham, “Management and Shocks to Worker Productivity,” International Growth Centre working paper (2016).
- Allcott, H., A. Collard-Wexler, and S. D. O’Connell, “How Do Electricity Shortages Affect Productivity? Evidence from India,” NBER working paper 19977 (2014).
- Allcott, H., and M. Greenstone, “Is There an Energy Efficiency Gap?” *Journal of Economic Perspectives* 26:1 (2012), 3–28.
- Allcott, H., S. Mullainathan, and D. Taubinsky, “Energy Policy with Externalities and Internalities,” *Journal of Public Economics* 112 (2014), 72–88.
- Allcott, H., and D. Taubinsky, “Evaluating Behaviorally-Motivated Policy: Experimental Evidence from the Lightbulb Market,” Sloan Foundation economics research paper 2640628 (2015).
- Barreca, A., K. Clay, O. Deschenes, M. Greenstone, M., and J. S. Shapiro, “Adapting to Climate Change: The Remarkable Decline in the US Temperature-Mortality Relationship over the Twentieth Century,” *Journal of Political Economy* 124 (2016), 105–159.
- Bhawan, S., and R. Puram, “CO₂ Baseline Database for the Indian Power Sector” (New Delhi: Central Electricity Authority, Ministry of Power, Government of India, 2014).
- Burke, M., S. M. Hsiang, and E. Miguel, “Global Non-Linear Effect of Temperature on Economic Production,” *Nature* 527:7577 (2015), 39–75.
- Burlig, F., C. Knittel, D. Rapson, M. Reguant, and C. Wolfram, Machine Learning from Schools about Energy Efficiency,” NBER technical report (2017).
- Callaway, D.S., M. Fowlie, and G. McCormick, “Location, Location, Location: The Variable Value of Renewable Energy and Demand-Side Efficiency Resources,” *Journal of the Association of Environmental and Resource Economists* 5:1 (2018), 39–75.
- Cameron, A. C., J. B. Gelbach, and D. L. Miller, “Bootstrap-Based Improvements for Inference with Clustered Errors,” this REVIEW 90 (2008), 414–427.

- Carranza, E., and R. Meeks, "Energy Efficiency and Electricity Reliability," this REVIEW (2020).
- Chang, T., J. Graff Zivin, T. Gross, and M. Neidell, "Particulate Pollution and the Productivity of Pear Packers," NBER working paper 19944 (2014).
- Costinot, A., D. Donaldson, and C. Smith, "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 (2016), 205–248.
- Dell, M., B. F. Jones, and B. A. Olken, "Temperature Shocks and Economic Growth: Evidence from the Last Half Century," *American Economic Journal: Macroeconomics* 4:3 (2012), 66–95.
- Enander, A. E., "Effects of Thermal Stress on Human Performance," *Scandinavian Journal of Work, Environment and Health* 16 (1989), 44–50.
- Fowlie, M., M. Greenstone, and C. Wolfram, "Do Energy Efficiency Investments Deliver? Evidence from the Weatherization Assistance Program," *Quarterly Journal of Economics* 133 (2013), 1597–1644.
- Graff Zivin, J., and M. Neidell, "The Impact of Pollution on Worker Productivity," *American Economic Review* 102 (2012), 3652–3673.
- Granade, H. C., J. Creyts, A. Derkach, P. Farese, S. Nyquist, and K. Ostrowski, "Unlocking Energy Efficiency in the US Economy," McKinsey report (2009).
- Hancock, P. A., J. M. Ross, and J. L. Szalma, "A Meta-Analysis of Performance Response under Thermal Stressors," *Human Factors: The Journal of the Human Factors and Ergonomics Society* 49 (2007), 851–877.
- Hanna, R., and P. Oliva, "The Effect of Pollution on Labor Supply: Evidence from a Natural Experiment in Mexico City," *Journal of Public Economics* 122 (2015), 68–79.
- Hsiang, S. M., "Temperatures and Cyclones Strongly Associated with Economic Production in the Caribbean and Central America," *Proceedings of the National Academy of Sciences* 107:35 (2010), 15367–15372.
- IPCC, *Working Group I Contribution to the IPCC 5th Assessment Report Climate Change 2013: The Physical Science Basis: Final Draft Underlying Scientific-Technical Assessment* (Cambridge: Cambridge University Press, 2013).
- Jing, S., B. Li, M. Tan, and H. Liu, "Impact of Relative Humidity on Thermal Comfort in a Warm Environment," *Indoor and Built Environment* 22 (2013), 598–607.
- Knittel, C. R., and R. Sandler, "Cleaning the Bathwater with the Baby: The Health Co-Benefits of Carbon Pricing in Transportation," NBER technical report 17390 (2011).
- Lemke, B., and T. Kjellstrom, "Calculating Workplace WBGT from Meteorological Data: A Tool for Climate Change Assessment," *Industrial Health* 50 (2012), 267–278.
- Levitt, S. D., and J. A. List, "Was There Really a Hawthorne Effect at the Hawthorne Plant? An Analysis of the Original Illumination Experiments," *American Economic Journal: Applied Economics* 3 (2011), 224–238.
- Mayo, E., F. Roethlisberger, and W. Dickson, *Management and the Worker* (Cambridge, MA: Harvard University Press, 1939).
- Nordhaus, W. D., "Estimates of the Social Cost of Carbon: Background and Results from the RICE-2011 Model," NBER working paper 17540 (2011).
- Parsons, K., *Human Thermal Environments: The Effects of Hot, Moderate, and Cold Environments on Human Health, Comfort and Performance* (Boca Raton, FL: CRC Press, 2010).
- Pilcher, J. J., E. Nadler, and C. Busch, "Effects of Hot and Cold Temperature Exposure on Performance: A Meta-Analytic Review," *Ergonomics* 45 (2002), 682–698.
- Ryan, N., "Is There an Energy-Efficiency Gap? Experimental Evidence from Indian Manufacturing Plants," Yale University technical report (2017).
- Saha, S., S. Moorthi, H.-L. Pan, X. Wu, J. Wang, S. Nadiga, ... M. Goldberg, "The NCEP Climate Forecast System Reanalysis," *Bulletin of the American Meteorological Society* 91 (2010), 1015–1057.
- Seppanen, O., W. J. Fisk, and Q. Lei, "Effect of Temperature on Task Performance in Office Environment," in Lawrence Berkeley National Laboratory, *Conference Proceedings of the Fifth International Conference on Cold Climate Heating, Ventilating and Air Conditioning* (2006).
- Snow, C., "Research on Industrial Illumination," *Tech Engineering News* 8 (1927), 257.
- Staritz, C., *Making the Cut? Low-Income Countries and the Global Clothing Value Chain in a Post-Quota and Post-Crisis World* (Washington, DC: World Bank, 2010).
- Sudarshan, A., R. Somanathan, E. Somanathan, and M. Tewari, "The Economic Impacts of Temperature on Industrial Productivity: Evidence from Indian Manufacturing" (2015).
- Tol, R. S., "The Marginal Damage Costs of Carbon Dioxide Emissions: An Assessment of the Uncertainties," *Energy Policy* 33 (2005), 2064–2074.
- Vanhoorne, M. N., O. V. Vanachter, and M. P. De Ridder, "Occupational Health Care for the 21st Century: From Health at Work to Workers' Health," *International Journal of Occupational and Environmental Health* 12 (2006), 278–285.