

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

1 Introduction

Mean global temperatures are projected to rise by at least 1.5 degrees Celsius by 2100 (25). Developing countries, whose adaptive and protective capacities are low and who are on average hotter than developed countries, will bear the lion's share of the ensuing negative impacts (33). Agricultural and manufacturing productivity may suffer, not only due to the increased frequency of extreme weather events¹ but also because excessive heat increases health risks (10; 11; 13; 16; 29) and decreases the body's capacity for exertion (27; 31; 40).

Given these repercussions, there is great academic and policymaking interest in quantifying the impacts of temperature (and of environmental factors more broadly) on economic outcomes. While the effect of rising temperatures on agriculture has been studied in depth by recent work (15; 19; 26; 30; 32), impacts on industrial productivity remain relatively unexplored.² As economies in many low-income countries undergo major structural transformations away from agriculture and into manufacturing and services sectors (38), understanding these effects becomes even more important.

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

Having documented substantial impacts of temperature on industrial productivity and absenteeism, the natural next question is whether it is possible to mitigate these impacts. Finding effective mitigation strategies is not enough; even the most tempered by the willingness of individuals and firms to adopt them. This process is not easy: though the public benefits of mitigation may be high, the immediate private returns are generally assumed to be low or even negative (28). Achieving widespread adoption, then, might require costly sustained subsidies or taxation. These

¹See, e.g., Deschênes and Greenstone (15); Guiteras (19); Hsiang (24); Kala et al. (26); Kurukulasuriya et al. (30); Lobell et al. (32).

²Though the idea that heat stress in the workplace could hamper productivity, particularly in low-income countries, has been written about for some time (see, e.g., Lemke and Kjellstrom (31) and Kjellstrom et al. (27), to our knowledge only one recent study in economics makes strides toward quantifying this impact (40).

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

questions are obviously important from a policy perspective, and matter, too, for our understanding of firms' technology adoption decisions.⁴

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

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

Our study contributes to the understanding of the effects of environmental factors on economic productivity. Recent work has documented significant labor supply and productivity impacts of air pollution (18; 21) and temperature (22; 24; 40; gra). Our detailed productivity measures, relatively long time span, and high-frequency temperature data allow us to quantify impacts with precision. Our findings are quite consistent with the previous studies: deviations in temperature strongly impact labor productivity. These studies make important inroads, but they stop short of evaluating policies that actually lessen the negative impact of heat on productivity. We provide the first evidence to our knowledge of the impact of technology that flattens the temperature-productivity gradient, while also having important environmental benefits via climate change mitigation.

We also add to the literature on the returns to climate change mitigation.⁶ The few recent studies that examine "co-benefits," or additional gains, of mitigation focus largely on the indirect public returns (9; 28). For example, a carbon tax aimed primarily at reducing CO_2 emissions may

⁴In an interesting and related study, Allcott and Taubinsky (2) test for behavioral hinderances as an explanation for low adoption of energy-efficient lightbulbs among consumers. We show among firms that non-adoption is even more suboptimal once production benefits are taken into account.

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

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

style will be the same for every garment produced by that line until the order for that garment is met). Lines consist of 20-100 sewing machine operators (depending on the complexity of the style) arranged in sequence and grouped in terms of segments of the garment (e.g. sleeve, collar, placket).⁷ Completed sections of garments pass between these groups, are attached to each other in additional operations along the way, and emerge at the end of the line as a completed garment. These completed garments are then transferred to the finishing floor.

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

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

2.3 Physiology of the Temperature-Productivity Gradient

The physical impact of temperature on human beings is a very well-studied area (17; 35), 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 (42). Higher temperatures and consequent thermal stress can impact human beings not only physically, but also through lower cognition and psychomotor ability (20). For instance, Ramsey et al. (37) find increases in unsafe behavior by workers at temperatures greater than 35 degrees Celsius WBGT (Wet Bulb Globe Temperature). The individual impact on a person varies based on factors such as the type of task and its complexity, duration of exposure, as well as the worker-level skill and acclimatization level (36), which contributes to the issues in setting a particular limit in working environments (20).

2.4 LED v. Fluorescent Lighting

LED light bulbs are approximately 7 times as energy-efficient as fluorescent bulbs (requiring about 3 as opposed to 21 KWh/year in electricity in our setting), and thus operate at about 1/7 the cost of fluorescent lighting. In addition, they generate a tenth of the CO_2 emissions (5.01 pounds of CO_2 per year per bulb, as compared to 35.11 pounds for fluorescent lighting).⁸ Heat emissions

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 can vary a bit by style or factory.

⁸Note 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.

for LEDs are substantially lower than fluorescent bulbs: the average LED bulb emits 3.4 Btus, as against 23.8 Btus for fluorescent lighting in the setting we study.

2.5 LED Roll-out: Summary and Timeline

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

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

A particular factory received the installation within a single month. 8% of the LED rollout (2 units) was completed in 2009, 48% (12 units) in 2010, 16% (4 units) in 2011, about 24% (6 units) in 2012 and the rest (1 unit) in 2013. Of the 29 units from which we have productivity data, LED replacements occurred in 25 units. Since our productivity data ranges from April 2010 to June 2013, some units already have LEDs at the beginning of our productivity data, and all but four units have LED for the last four months of our productivity data range. This is why our results report the impact of temperature not only for the whole sample, but also separately for the units that did not have LED at a particular time.

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

3 Data

the calculation is derived from the Rothfusz regression that replicates the HI values from Steadman (1979).

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

$$HI = HI + [(rh - 85) = 10] [(87 - T_d) = 5] \quad (2)$$

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

$$WBGT = 0.567T_d + 0.216 \frac{rh}{100} 6.105 \exp \frac{17.27T_d}{237.7 + T_d} + 3.38 \quad (3)$$

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

3.1.2 Factory Data

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

Actual efficiency is actual quantity produced divided by target quantity. The target quantity is derived from an industrial engineering (IE) measure for the complexity of the garment called “Standard Allowable Minute” (SAM). This measure amounts to the estimated number of minutes required to produce a single garment of a particular style. This estimate comes from a central

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

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

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

Most importantly, we use actual efficiency rather than produced quantity as our outcome of choice. Produced quantity would not account for systematic variation due to complexity of style or number of operations. Without normalizing production observations to target quantity, one could potentially misrepresent an association between temperature and style complexity or order size as an impact on productivity. That is, for example, if garment complexity or line length varied by temperature, directly the per a 299stheper1(theperahebna)sofal(ef)17ohstheper 0 -16.435 Td [(W)92(elab7(linn)-su

garment complexity, order size, and line- and order-specific characteristics. True daily fluctuations in productivity are, therefore, best measured by actual efficiency net of budgeted efficiency.

3.2 Summary Statistics

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

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

3.3 Preliminary Graphical Evidence

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

3.3.1 Temperature-Productivity Gradient

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

FIGURE 1: EFFICIENCY AND TEMPERATURE TIME SERIES

FIGURE 2: EFFICIENCY AGAINST TEMPERATURE

FIGURE 3: EFFICIENCY AGAINST TEMPERATURE BY LED

3.3.2 Impacts of LED Introduction

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

We also return to the exercise conducted in Figure 1, but plot the efficiency series separately for factory-day observations with and without LED lighting. Once again, we include the temperature time series for comparison and present the plots in Figure 4. The evidence in Figure 4 also suggests that LED lighting improves efficiency on all days, but particularly smooths the fluctuations in efficiency due to temperature.

Finally, we explore graphically the main effects of LED lighting on efficiency. That is, there might be many pathways by which LED introduction might affect efficiency directly or outside of its interaction with the temperature-efficiency gradient. For example, the quality and quantity of light might have effects on attentiveness and sight, both of which are important in garment manufacturing. While, the previous figures indicate that LED might have a mitigative impact on

FIGURE 4: EFFICIENCY AND TEMPERATURE TIME SERIES BY LED

the temperature-efficiency gradient, estimating the mean composite effect of LED introduction on efficiency is also of interest in what follows. We present preliminary evidence of this main effect by plotting actual efficiency over time relative to the date of LED introduction. However, Figure 5 shows little evidence of a mean composite effect.

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

4 Empirical Strategy

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

$$E_{ludmy} = \alpha_0 + \alpha_1(\text{concerns}) - 245(\text{r}) + \dots + \alpha_6(\text{year60}) + \dots$$

year60

FIGURE 5: EFFICIENCY BEFORE AND AFTER LED

line fixed effects; u_{iy} are unit x year fixed effects; m are month fixed effects; d are day-of-week fixed effects; and α_0 is an intercept. α_1 is the coefficient of interest, giving the impact of a 1-degree Celsius increase in temperature on line-level efficiency.

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

$$E_{ludmy} = \alpha_0 + \alpha_1(Q_1 \times T_{dmy}) + \alpha_2(Q_2 \times T_{dmy}) + \alpha_3Q_1 + \alpha_4Q_2 + \alpha_5 + u_{iy} + m + d + \mu_{ludmy}; \quad (5)$$

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

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

$$E_{ludmy} = \alpha_0 + \alpha_1 T_{dmy} \times LED_{umy} + \alpha_2 LED_{umy} + \alpha_3 T_{dmy} + B_{ludmy} + \alpha_4 + u_{iy} + m + d + \mu_{ludmy}; \quad (6)$$

Here LED_{umy} is a dummy variable for the presence of LED lighting in unit

from 0 to 1 in the month following LED introduction in a particular factory unit. The coefficients of interest in the above specification are β_1 and β_3 . β_3 indicates the effect of temperature on productivity *before* LED introduction. β_3 is the extent of attenuation of the temperature-productivity gradient once LED lighting is introduced. The sum of these two, $\beta_1 + \beta_3$, gives the net effect of temperature on productivity following LED introduction.

Finally, we test if the attenuation impact of LED is heterogenous for temperatures above and below the mean temperature (by using the interaction term $(\text{temp} - \text{mean_temp}) \times \text{led}$).

5 Results

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

5.0.1 Average Impact of Temperature

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

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

might be contributing to the persistent, cumulative effects of temperature on efficiency depicted in Table 3. We do not push the interpretation of lagged impacts further due to concerns regarding serial correlation in temperature, and undertake specification robustness checks in Table A.3 to partially alleviate concerns regarding serial correlation in temperature.

5.2 LED and the Temperature-Efficiency Gradient

and above the median. The main effect of the below median dummy is omitted to preserve the constant.

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

5.3 Main Effect and Exogeneity of LED Introduction

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

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

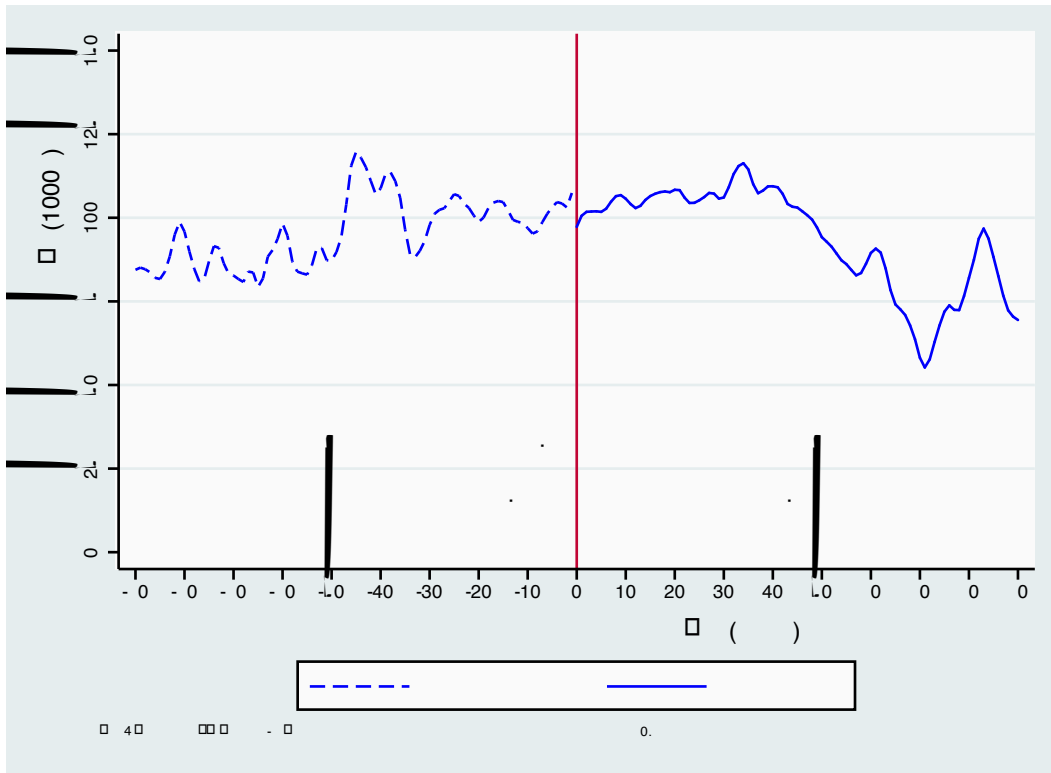
Nevertheless, for the sake of interpretation and external validity, we investigate determinants of the timing of LED introduction. Specifically, we investigate the degree to which standard allowable minutes, budgeted efficiency, scheduled quantity, and target quantity correlate with LED introduction.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

We first regress SAM on LED introduction to check whether the timing of LED roll-out coincided with the arrival of orders for more or less complex garments. The results of this regression

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

FIGURE 6: EFFICIENCY BY DATE RELATIVE TO

FIGURE 7: SCHEDULED QUANTITY BY DATE RELATIVE TO LED INTRODUCTION



are reported in column 2 of Table 7 and show no evidence of a relationship between SAM and LED roll-out.

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

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

6 Discussion

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

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

ingly, at the current estimates of carbon prices, these benefits are relatively small in comparison to the annual private benefits.²⁰

We believe our work is an important first step in quantifying private co-benefits of climate change mitigation strategies, but that much more needs to be done to quantify the full returns to the variety of mitigation strategies. For example, as Knittel and Sandler (28) suggest, carbon taxes likely have health benefits due to decreases in local air pollution. If consumers internalize these benefits, the effective costs of the tax will be substantially lower. Whether similar co-benefits exist for other types of mitigation – e.g., renewable energy investments, public transport systems, energy-efficient built environments, etc. – is an open and vital question.

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

References

[gra]

- [2] Allcott, H. and Taubinsky, D. (2013). The lightbulb paradox: Evidence from two randomized control trials. *working paper*.
- [3] 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.
- [4] Bandiera, O., Barankay, I., and Rasul, I. (2009). Social connections and incentives in the workplace: Evidence from personnel data. *Econometrica*, 77(4):1047–1094.
- [5] Bandiera, O., Barankay, I., and Rasul, I. (2010). Social incentives in the workplace. *The Review of Economic Studies*, 77(2):417–458.
- [6] Barreca, A., Clay, K., Deschenes, O., Greenstone, M., and Shapiro, J. S. (2013). Adapting to climate change: The remarkable decline in the u.s temperature-mortality relationship over the 20th century.
- [7] 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.
- [8] Bloom, N. and Van Reenen, J. (2010). Why do management practices differ across firms and countries? *The Journal of Economic Perspectives*, pages 203–224.
- [9] Bollen, J., Guay, B., Jamet, S., and Corfee-Morlot, J. (2009). Co-benefits of climate change

- [14] De Mel, S., McKenzie, D., and Woodruff, C. (2012). Enterprise recovery following natural disasters. *The Economic Journal*, 122(559):64–91.
- [15] 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.
- [16] 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.
- [17] Enander, A. E. (1989). Effects of thermal stress on human performance. *Scandinavian journal of work, environment & health*, pages 27–33.
- [18] Graff Zivin, J. and Neidell, M. (2012). The impact of pollution on worker productivity. *American Economic Review*, 102:3652–3673.
- [19] Guiteras, R. (2009). The impact of climate change on indian agriculture. *Manuscript, Department of Economics, University of Maryland, College Park, Maryland*.
- [20] 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.
- [21] Hanna, R. and Oliva, P. (2011). The effect of pollution on labor supply: Evidence from a natural experiment in mexico city. Technical report, National Bureau of Economic Research.
- [22] Heal, G. and Park, J. (2013). Feeling the heat: Temperature, physiology & the wealth of nations. Technical report, National Bureau of Economic Research.
- [23] Hjort, J. (2013). Ethnic divisions and production in firms. Technical report, CESifo Working Paper.
- [24] 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.
- [25] 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.
- [26] 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.

- [27] 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.
- [28] Knittel, C. R. and Sandler, R. (2011). Cleaning the bathwater with the baby: The health co-benefits of carbon pricing in transportation. Technical report, National Bureau of Economic Research.
- [29] Kudamatsu, M., Persson, T., and Strömberg, D. (2012). *Weather and infant mortality in Africa*. Centre for Economic Policy Research.
- [30] 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.
- [31] Lemke, B. and Kjellstrom, T. (2012). Calculating workplace wbgt from meteorological data: a tool for climate change assessment. *Industrial Health*, 50(4):267–278.
- [32] Lobell, D. B., Schlenker, W., and Costa-Roberts, J. (2011). Climate trends and global crop production since 1980. *Science*, 333(6042):616–620.
- [33] Mendelsohn, R., Dinar, A., and Williams, L. (2006). The distributional impact of climate change on rich and poor countries. *Environment and Development Economics*, 11(02):159–178.
- [34] Nordhaus, W. (2008). A question of balance.
- [35] Parsons, K. (2010). *Human thermal environments: the effects of hot, moderate, and cold environments on human health, comfort and performance*. CRC Press.
- [36] 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.
- [37] Ramsey, J. D., Burford, C. L., Beshir, M. Y., and Jensen, R. C. (1983). Effects of workplace thermal conditions on safe work behavior. *Journal of Safety Research*, 14(3):105–114.
- [38] Roberts, M. J. and Tybout, J. R. (1996). Industrial evolution in developing countries: Micro patterns of turnover, productivity, and market structure.
- [39] 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.
- [40] Sudarshan, A. and Tewari, M. (2013). The economic impacts of temperature on industrial productivity: Evidence from indian manufacturing.
- [41] Tol, R. S. (2005). The marginal damage costs of carbon dioxide emissions: an assessment of the uncertainties. *Energy policy*, 33(16):2064–2074.

A Additional Tables

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

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

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

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

In Table A4, we consider whether the mitigative impact of LED varies by whether a unit received LED earlier or later. Since we have unit by year fixed effects in all regression specifications, the estimates are not biased by the potential correlation of unit-level unobserved heterogeneity with the introduction of LED. Nevertheless, it is still instructive to examine whether units that adopted LED later exhibit differential impacts of LED on the productivity-temperature gradient. We divide our sample of 29 units into two - the first 14 units to receive LED are considered to

have received LED relatively early and the other 15 relatively late.²¹ Table A4 illustrates that the mitigative impact of LED does not vary by whether a unit received LED relatively late, since the triple interaction term between the LED dummy variable, temperature, and a dummy that takes the value 1 if the unit received LED relatively late and is 0 otherwise is not statistically significant. The other coefficients of interest - the interaction of the LED dummy and temperature and the main effect of temperature - are very similar to the estimates in Table 5. Temperature has a negative and statistically significant impact on production efficiency, and the introduction of LED mitigates a large portion of this negative relationship.

B Data Appendix

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

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

We remove lines observed greater than twice a day, about 0.6% of our observations, since these are likely coding errors. While it is possible that a line finished a set of orders and

\AM@currentdocname .png

.png