HIGH TEMPERATURES IMPEDE PRODUCTIVITY GROWTH ON THE JOB*

Isadora Frankenthal[⊕†] Heather Schofield[⊕‡] Pedro Bessone[§]

We study how temperature affects productivity growth among Indian data-entry workers new to the job. Higher temperatures impede on-the-job learning, reducing the rate of productivity growth by 3.0 percentage points per degree Celsius over the past week. This productivity loss persists through the second half of the study, with no evidence of catch-up among those exposed to high initial temperatures. These effects are beyond the contemporaneous impact of heat: each additional degree Celsius lowers same-day productivity by 1.1%, driven by increased breaks, with no change in days worked or hours in the office. The effect of temperature on both the growth rate and level of output occurs despite strong incentives for production. The findings suggest that rising temperatures slow the accumulation of productive human capital, with consequences for both workers' earnings trajectories and firms' productivity and training costs. These impacts are likely to be more pronounced in low-income countries, where workers face greater heat exposure.

JEL Codes: J24, Q56

*Research reported in this publication was also supported by the National Institute on Aging of the National Institutes of Health under award number K01AG055691. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. We thank Achyuta Adhvaryu, Kelsey Jack, Maulik Jagnani, Jacob Moscona, Anant Nyshadham, Ben Olken, Frank Schilbach, and Nick Swanson for their insightful comments. We thank Alosias A, Srinivas Balasubramanian, Thomas Escande, Juliette Finetti, Kannan Kumar, Gunjita Gupta, Erik Hausen, Andrew Locke, Madeline McKelway, Adrien Pawlik, Jane von Rabeau, Sangeetha Ramanathan, and Krishna Prasad Srinivasan for excellent research assistance on the study which generated the data used here. We thank Shuning Ge, Aden Halpern, Simon Taye, and Haoran Wang for research assistance on this paper. The original experiment underlying this study received IRB approval from Harvard University, protocol no. 14-2294, the Institute for Financial Management and Research in India, and the Indian Council of Medical Research. All remaining errors are our own.

[⊕]Shared first-authorship
[†]MIT (isadoraf@mit.edu)
[‡]Cornell University (hws44@cornell.edu). Corresponding author. Sage Hall 375, Ithaca, NY 14850.
[§]Uber (tepedino@uber.com)

1. Introduction

On-the-job learning constitutes a substantial component of skill acquisition and productivity growth throughout workers' careers (Becker, 1962; Mincer, 1962; Arrow, 1962; Ben-Porath, 1967). Learning and productivity growth are common even for workers performing seemingly routine tasks, as in auto manufacturing or garment production, where the exact task changes frequently — for instance, the particular garment being sewn (Levitt et al., 2013; Adhvaryu et al., 2023). The importance of on-the-job learning is likely to grow further as economies shift toward work that requires greater training and learning, particularly in developing countries where knowledge-intensive jobs are rapidly expanding (Lo Bello et al., 2019). Understanding the factors that enhance or impede on-the-job learning is therefore increasingly important.

One such factor that has the potential to impede growth in labor productivity is high temperatures. Previous research has established that high temperatures reduce contemporaneous productivity (Adhvaryu et al., 2020; Somanathan et al., 2021; LoPalo, 2023; Masuda et al., 2021; Huppertz, 2024; Garg et al., 2025). However, there is little empirical evidence on whether elevated temperatures not only reduce productivity in the moment but also slow the rate of labor productivity growth. These dynamic effects may be particularly damaging for low-income countries that are disproportionately impacted by increasing temperatures with little climate control (Liu et al., 2017; Pittock, 2017; Mendelsohn et al., 2006; Ibarrarán et al., 2009).

We study the effect of higher temperatures on productivity and productivity growth in the context of data-entry jobs in Chennai, India.[1] The work arrangement is particularly well suited to this purpose for two reasons. First, the data-entry task was purely cognitive in nature, compensated primarily on a piece rate basis (leading to a close mapping between productivity and wages), and fully flexible in terms of labor supply. This allows us to test whether the negative effect of temperature on productivity — well-documented in more physically demanding work — also extends to a setting with no physical exertion, strong performance incentives, and opportunities for workers to adapt. Second, all workers were entirely new to data entry, enabling us to track individual productivity growth as they learn on the job and to assess whether elevated temperatures affect their trajectories.

[1] The jobs were offered as part of a randomized controlled trial studying the impact of sleep on labor outcomes.

The study's location, Chennai, India, is representative of typical tropical environments with typical daytime highs ranging from approximately 27°C in December to nearly 40°C in May (Weather and Climate, 2024; Wilkinson, 2014).

We estimate both contemporaneous and dynamic effects of temperature on output by exploiting variation in temperature, controlling for month by year and worker fixed effects. We find that higher temperatures generate economically meaningful reductions in hourly productivity. A one degree Celsius increase in the average daily temperature reduces hourly productivity by 1.1%. This result is consistent with previous literature that documents the effect of temperature on productivity and work quality in more physically intensive work such as door to door surveying or garment manufacturing, despite the more cognitive nature of data-entry (Adhvaryu et al., 2020; Somanathan et al., 2021; LoPalo, 2023; Huppertz, 2024; Walter and Moneke, 2024). Turning to mechanisms, we find that this reduction in productivity is driven by reductions in typing time, predominantly via increased breaks, with no impact of heat on total time spent at the office, or absenteeism. Mistakes rise by 0.9% per degree, but with fewer than one mistake per 100 entries, this increase has little effect on overall productivity or earnings.

Next we leverage the rapid increase in productivity over time — a roughly four-fold increase in hourly productivity over a month, with roughly 90% of the gain in the first two weeks — to test whether heat reduces the rate of productivity growth. We build on Dell et al. (2012)'s method for estimating the effect of temperature on GDP growth by estimating an analogous worker-level growth-rate specification with individual productivity growth as the outcome. We find that higher temperatures significantly reduce the rate of productivity growth during the first half of the study period, when productivity is growing: a one degree Celsius increase in the average temperature during the previous six days reduces productivity growth by 3.0 percentage points (p=0.002). The causal interpretation of this result is reinforced by two placebo tests. First, we show that temperature in the immediate *future* is not associated with slower productivity growth, suggesting the effect we find for past temperatures is not driven by omitted variables, reverse causality, or anticipatory effects. Second, we show that there is no relationship between temperature and productivity growth in the second half of the study, when productivity growth is essentially flat. Beyond serving as a placebo test, this second fact suggests that those who are initially exposed to greater heat, fail to "catch up" in the second half of the study. This persistence is further supported by an additional test: workers exposed to a one degree Celsius higher average temperature during the first half of the study period remain 6.2% less productive and earn 5.7% less in the final week of the study controlling for temperatures in the intervening period. Although we can not rule out catch-up in the long-run, together these results suggest that productivity losses from early heat exposure persist for the duration of the study period.

This study contributes to three growing literatures on the economic effects of temperature and the potential consequences of rising temperatures due to climate change. First, building on evidence of humans' physiological responses to heat (Hancock et al., 2007), economists have documented harmful contemporaneous effects of higher temperatures on individual productivity in manufacturing, agriculture, and field surveying, as well as reductions in the productivity of teams (Zhang et al., 2018; Adhvaryu et al., 2020; Somanathan et al., 2021; LoPalo, 2023; Masuda et al., 2021; Huppertz, 2024; Garg et al., 2025). We document detrimental effects in a novel context of desk-based office work that is exclusively cognitive and individual. In contrast to most environments studied in this literature where workers typically have flat wages, we focus on a context in which workers face strong incentives for productivity: 73% of their total payment was a function of the time they spent actively typing and their output, yet workers still experience heat-related deficits in production. Further, our detailed data allow us to explore and document the margins of adjustment through which effects operate.

Second, this paper speaks to the literature on human capital accumulation at work and on-the-job learning. Following foundational work by Arrow (1962), Becker (1962), Mincer (1962), and Ben-Porath (1967), an extensive body of research has documented that workers acquire human capital throughout their careers, with learning on the job representing a major component of productivity growth (e.g., Lagakos et al. (2018); Levitt et al. (2013)). Past work has documented steep learning curves across diverse settings and in seemingly routine tasks (Bessen, 1997; Levitt et al., 2013; Adhvaryu et al., 2023), and that learning is directly influenced by training as well as by institutional factors such as the quality of peers and managers (e.g., Ma et al. (2024); Mas and Moretti (2009);

Lazear et al. (2015)). However, despite considerable research studying institutional and technological factors that influence human capital accumulation at work, little is known about environmental constraints on learning-on-the-job. We provide the first evidence that ambient temperature systematically impedes workers' ability to acquire skills and grow productivity.[2] Given that firms in low-income settings face high turnover and real-world employment typically involves continuous exposure to new challenges — e.g., even in routine work like data entry, over time workers regularly encounter updated software, modified procedures, or different content types — these findings suggest that rising temperatures may represent an underappreciated barrier to the growth of human capital.

Finally, this paper is related to the literature studying temperature effects among school-aged children, which finds negative effects of high temperatures on the test day itself on test performance (Park, 2022) and mixed, though very suggestive impacts of higher temperatures throughout the school year on test performance (Park et al., 2020; Graff Zivin et al., 2018; Garg et al., 2020). Our study complements this work by examining temperature's effects on labor productivity growth — which is likely driven at least in part by on-the-job learning, a related but distinct construct from academic learning — under strong economic incentives and among adults. Our results suggest that temperature meaningfully impacts on-the-job productivity growth, in spite of the high powered incentives and opportunities to adapt to heat through work flexibility that are present in our setting.

2. Background

2.1. Data-entry Work

This study draws on data from workers in a previously conducted experiment studying the impact of increased sleep on productivity, cognition, and well-being (Bessone et al., 2021). The workers were men and women, 25 to 55 years old, in Chennai, India. Workers were recruited by advertising the job in low-income neighborhoods near the study office: the median distance between the participants' neighborhood and the study office was 3.8 kilometers. Given screening criteria and

[2] There is a related macro-economics literature studying the effects of temperature on GDP and TFP growth; see Deschênes and Greenstone (2007) and Deschenes (2023) for reviews. Our micro-level evidence on individual productivity growth complements this literature, though connecting the two requires assumptions about aggregation that are beyond the scope of this paper.

this recruitment procedure, the workers were literate but typically from lower socio-economic strata. While each participant was enrolled for one month, recruitment was rolling and occurred over 19 months, providing substantial temperature variation with mean highs ranging from roughly 27°C in December to just under 40°C in May (Figure Ic). Roughly 30% of workers had prior computer experience and none had previous data-entry experience.

The experiment was designed to replicate a real data-entry job. Participants worked in an office equipped with computers and the data-entry software. Each worker had a fixed desk. Workers were asked to transcribe what appeared to be handwritten text following a prescribed template from the left portion of the screen to the right portion of the screen (Figure A.1). Because each worker was employed for one month, the task was held constant over time. The data-entry software records the total number of characters entered and the number of mistakes per worker every five minutes. Time worked within the day was also measured precisely: after two minutes without active typing the software would note that the work was paused. Workers could pause and resume at any time.

The job was flexible but workers were incentivized for production. They could choose whether or not to come to the office, their arrival and departure time between 9:30 a.m. and 8 p.m., and how much work to do when they were present.[3] A surveyor recorded the time they entered and left the office each day. Workers were paid a fixed amount for every hour spent typing, and a variable amount based on performance: a piece rate for every correct entry, and a penalty for every incorrect entry. The incentives for production were designed to be strong: roughly two-thirds of pay was performance-based (a piece rate on output) and the remaining third was time-based (an hourly rate for time spent typing).

2.2. Climate in Chennai

Chennai, India has a tropical climate, though with substantial variation over the course of the study duration given the rolling-enrollment. The coolest month is December, with a mean daytime temperature of 25.2°C, and the warmest month is May, with a mean daytime temperature of 35.3°C (Figure I, panel C). The mean temperature during working hours ("daytime") over the 19 month

[3] A small subset of days was limited to shorter fixed hours to provide a productivity measure independent of the choice of hours worked. All specifications include day-in-study fixed effects to control for these days.

FIGURE I: Temperature in Chennai

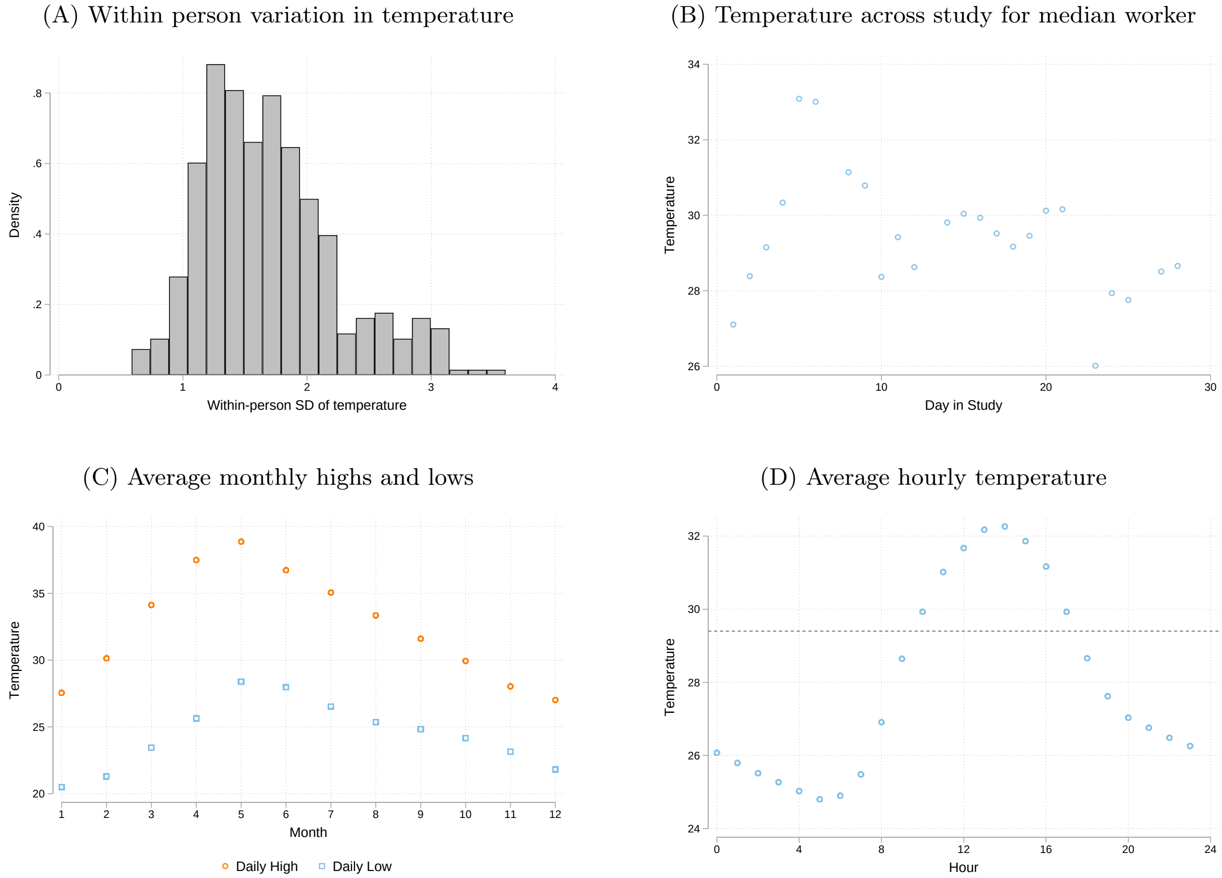

Notes: This figure shows temperature variation for study participants and during the years of the study. All measures are from gridded temperature data for the location of study office in Chennai, India. Panel A displays a histogram of the standard deviation of average daily temperature for each worker during the study period. Panel B plots the average daily temperature for each day in the study for a "representative worker" who had a median standard deviation of temperature across the study (1.6 SD) and was within 1 degree Celsius of the mean temperature. Panel C graphs the average daily maximum and daily minimum temperature for each month. Panel D graphs the average temperature at each hour of the day.

period in which workers were enrolled was 30.3°C, ranging from 21.2°C to 40.8°C. Temperatures typically peak around 2 p.m. (Figure I, panel D), though stay high until late afternoon or early evening. Maximal temperatures were over 40°C, however only 4.1% of days are above this cutoff. These average temperatures are representative of the tropics, which house about 40% of the world's population (Wilkinson, 2014).

3. Data

The analysis combines detailed data from three datasets: 1) productivity of 452 workers over 4 weeks of work, 2) weather data, and 3) pollution data. This section describes the sources of data and how the raw data is cleaned and compiled for the analysis.

Productivity For each worker, we construct hourly and daily measures of performance at work. The main measure of productivity, which was pre-registered for the original experiment, is average quality-adjusted output: correct entries minus eight times incorrect entries. In addition, we examine the total number of characters entered, the number of mistakes, how long participants spent actively typing, and their typing earnings. We use the hourly data to compute average hourly output. We use the daily measures to examine whether higher temperatures reduce total daily output, beyond their effect on average hourly productivity by shifting work across time.

Beyond productivity, the data also record whether the worker was present at work and, if so, exactly what time they arrived and left. This information is used to investigate the effect of temperature on the extensive margin of labor supply. In addition, three tests were administered to measure cognition. First, a psychomotor vigilance test — capturing alertness — was administered daily, second, a hearts and flowers test — measuring inhibitory control — was administered on alternate days; and third, a Corsi task — testing short term visual and spatial memory — was also administered on alternate days. We use pre-specified formulas which combine each worker's accuracy and reaction time on each task to create a comprehensive measure of performance. The three metrics are also combined into an Anderson (2008) index.

Temperature We use gridded hourly temperature data for the location of the study office.[4] These data are collected by the NCEP Climate Forecast System Reanalysis (CFSR) provided by the NOAA National Centers for Environmental Prediction (NCEP), and are described in detail in Saha et al. (2010). The data records the temperature (and relative humidity, which we use to compute heat index for robustness) at the beginning of each hour in degrees Celsius, accurate to the nearest 0.1

[4]Because participants were recruited from nearby neighborhoods, these temperature measures are also strongly correlated with workers' temperature exposure outside the office, including at night.

degree.[5] The temperature data are matched to the productivity data at the hour-day level. From the hourly temperature data we also create a daily-level dataset which records the average temperature for each workday. We compute daily temperature as the average temperature during each worker's office hours, providing an accurate measure of the temperature experienced while at work.[6]

Pollution $PM_{2.5}$, SO_2, NO_2, and O_3 pollution data is sourced from Sharma and Mauzerall (2021). PM_{10} is sourced from Copernicus Atmosphere Monitoring Service (2020). For each pollution measure, we source data captured every 3 hours from the closest of the 3 measurement stations in Chennai that recorded data (the stations are approximately 9.7 kilometers, 14.2 kilometers and 14.5 kilometers away from the study site) and aggregate to the daily level.

4. Empirical Strategy

Effect on Productivity There are two ways in which contemporaneous temperature can affect daily productivity. First, higher temperatures can make the worker less productive per unit of time, either via reduced speed or by increasing mistakes. Second, higher temperatures might reduce time spent at work, or time spent actively working while at the office. To investigate the direct effects of temperature on contemporaneous productivity via these channels, we estimate the following equation:

$$\text{Y}_{id} = \beta \cdot \text{Temperature}_{t(i,d)} + \alpha_i + \delta_d + \gamma_m + \tau_{i,d} + \mathbf{L}_{t(i,d)} + \epsilon_{id} \tag{1}$$

Where i indexes workers, d indexes days since enrollment in the study (day-in-study), $t(i,d)$ indexes the calendar date for worker i on study day d, and m indexes calendar time using month-year bins. Y_{id} is a measure of productivity for a worker-day. $\text{Temperature}_{t(i,d)}$ is the average temperature in degrees Celsius for worker i on day d. α_i, δ_d, and γ_m are worker, day-in-study, and month-by-

[5] NOAA data captures external temperatures. While indoor measurements are not available, the office was not air conditioned and was connected to an open air courtyard via open windows and an open door, so indoor temperatures tracked outdoor conditions relatively closely.

[6] When estimating effects on attendance and arrival and departure times, we instead measure temperature over the standard working day (9 a.m. to 8 p.m.) to avoid defining temperature using workers' endogenous labor supply choices.

year fixed effects, respectively. Worker fixed effects hold constant any differences across people in inherent ability or anything about them that is fixed over time (e.g. their neighborhood). Day-in-study fixed effects account for differences in output over time, which is particularly important given the steep learning curve, with productivity quadrupling over the first 16 days on average.[7] We include month-by-year fixed effects, which are particularly important given the rolling enrollment of workers in recruitment waves throughout the study period. By comparing workers exposed to different temperatures within the same month and year, these fixed effects absorb differences in average worker ability across recruitment periods that could otherwise be correlated with temperature. Month-by-year fixed effects also absorb idiosyncratic month-specific factors, such as festivals, that may affect productivity.[8] $\tau_{i,d}$ is a vector of binary variables capturing the six experimental arms in the original experiment, which are implemented following a baseline period on day eight. Following previous work that estimates the effect of temperature on output, the main specification also controls for three lagged values of temperature, $\mathbf{L}_{t(i,d)}$, to minimize bias due to serial correlation (Adhvaryu et al., 2020; Somanathan et al., 2021; LoPalo, 2023). Standard errors are clustered at the worker level.

The coefficient of interest is β, which captures the effect of average daily temperature on measures of output. We estimate the effect separately on quality adjusted output, the total number of entries, the time spent typing, and the number of mistakes. Further, to investigate the margins of adjustment underlying the total time spent typing, we estimate whether workers alter arrival or departure times in response to higher temperatures. Separately, we also examine whether temperature impacts attendance at the office.

The advantage of this worker-by-day level specification is that it takes into account an important margin of adaptation: on warmer days, workers can choose to shift effort from hotter hours of the day to cooler hours of the day. If they did so systematically, the immediate effect of temperature on output in a given hour might be large, but the effect of temperature on *average* hourly productivity would be much smaller. We can interpret β from the equation above as the average effect of temperature on hourly productivity inclusive of this margin of adaptation.

[7]Day-in-study fixed effects also address day-specific activities administered as part of the original RCT.

[8]We also show that workers recruited in different seasons are comparable in English literacy, prior computer experience, years of education, and math ability (Appendix Table A.1).

Effect on Productivity Growth Beyond direct effects in the moment, higher temperatures can also slow productivity growth over time if heat makes it harder for workers to learn skills on the job. The key challenge in studying how temperature affects growth rates is distinguishing between level effects, a temporary shock that reduces output in the short run but lasts only for the duration of the shock, and growth effects, a shock that permanently slows the rate of productivity accumulation. Dell et al. (2012) tackle this distinction by exploiting year-to-year variation in temperature within countries and estimating distributed-lag models of growth. A pure level effect shows up as a short-run drop in output that is offset once conditions normalize, so that the sum of current and lagged temperature coefficients on growth converges to zero.[9] By contrast, if temperature shocks leave the economy permanently behind, the negative impacts persist and the cumulative sum of coefficients across lags is negative. The intuition for our setting is similar: temporary heat shocks may reduce workers' productivity on a given day but should not alter their trajectory of improvement. In contrast, impacts on growth would suggest that exposure to heat in the past impairs the ability to learn on the job, resulting in persistently lower productivity growth over time. In this case, productivity growth remains lower across subsequent periods, and the cumulative sum of temperature effects provides a direct test of whether heat alters the pace of learning on the job, not just contemporaneous performance.

We estimate the following equation:

$$\text{(2)} \qquad \text{Productivity Growth Rate}_{ip} = \sum_{j=0}^{J} \beta_j \cdot \text{Temperature}_{t(i,p-j)} + \alpha_i + \delta_p + \gamma_m + \tau_{i,p} + \epsilon_{ip}$$

Here, i indexes workers, $t(i,p)$ indexes the dates during which worker i is in period p, m indexes the month-year. j indexes the lag (i.e., the number of two-day blocks prior to the current period, p). To reduce noise due to random fluctuations in output across days, productivity growth is calculated by comparing differences in productivity for each worker across two-day blocks. The definition of a time period, p, is therefore a two-day block. Temperature is calculated as the simple average of the temperature during work hours across each period. The Productivity Growth Rate$_{ip}$

[9] Note that individual coefficients on each lag are not directly interpretable.

FIGURE II: Increase in Average Hourly Productivity Over Study Period

(A) Productivity growth

(B) Workers split by exposure to high temperature

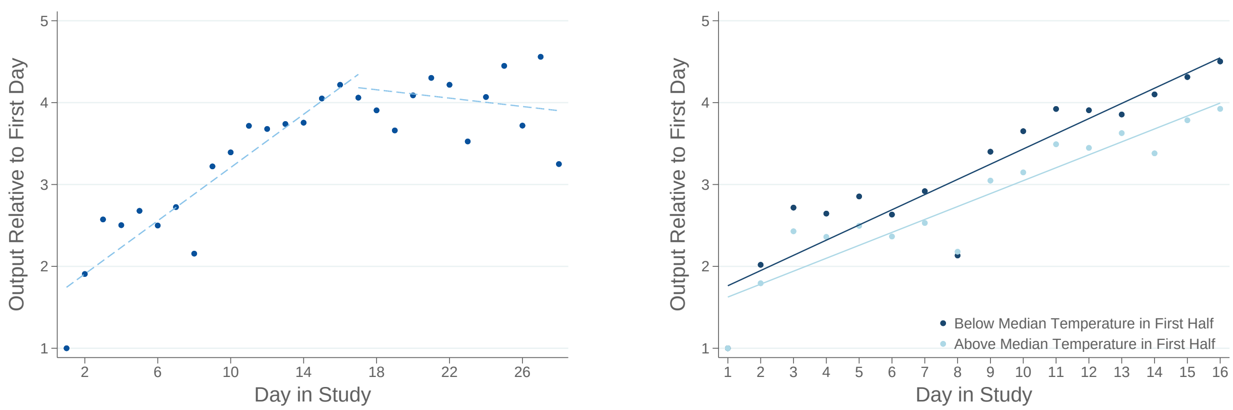

Notes: This figure shows how worker productivity evolves over the course of the study. Panel A plots average hourly productivity across the 28 study days, normalized by first-day output. The dotted lines represent a piecewise linear fit to the data, with a break after day 16. Panel B examines the same outcome variable, but separates participants who are exposed to above-median (light blue) and below-median (dark blue) average temperatures over the study. Points are generated via a binscatter, with a best-fit line for each temperature group. The sample in Panel B is restricted to the first 16 days to focus on the period of productivity growth.

is calculated as productivity in period p minus productivity in $p-1$, divided by productivity in $p-1$. Temperature$_{t(i,p-j)}$ is the average temperature in degrees Celsius for worker i in period $t(i,p-j)$. We estimate the equation with one, two, and three lags. Given that productivity only grows meaningfully for the first 16 days we are limited in the number of potential lags that can be included. Three two-day lags provides nearly a week of lagged data while also leaving ten days over which to evaluate growth. α_i, δ_p, and γ_m are worker, time period, and month-by-year fixed effects, respectively. $\tau_{i,p}$ are dummies for the participant's treatment assignment from the original RCT and are as defined in Equation 1. Standard errors are clustered at the worker level.

The estimate of interest is the sum of the coefficients on the measures of temperature, $\sum_{j=0}^{J} \beta_j$, which captures the effect of temperature on the growth rate of productivity (Dell et al., 2012). A negative and significant sum would indicate that cumulative exposure to higher temperatures negatively affects productivity growth over time. The main results are estimated using the first 16 days of the study, when individual productivity is growing rapidly (Figure II, Panel A).[10] The sample consists of the 89.2% of workers who were present at the office more than 75% of the days

[10]This break point is data-driven. We estimate a piece-wise linear function to the raw productivity data and find a break point at day 16. From day 17 onward, growth is flat.

Table I: The Contemporaneous Effect of Temperature on Labor Outcomes

	Average Hourly Outcomes					Daily Outcomes			
	Quality Adjusted Output	**Total Number of Entries**	**Active Typing Time**	**Mistakes** (per 100 entries)	**Typing Earnings**	**Present** (=1)	**Check-in Time**	**Check-out Time**	**Hours of Work**
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Temperature ($°C$)	-20.891***	-21.069***	-0.288***	0.007**	-0.456***	-0.000	0.004	-0.008	-0.011
	(2.828)	(2.921)	(0.038)	(0.003)	(0.076)	(0.002)	(0.004)	(0.007)	(0.008)
Dep. Var. Mean	1820.3	1948.4	29.23	0.827	35.27	0.893	10.59	18.33	7.740
Observations	7268	7268	7268	7267	7268	12655	10884	10884	10884

Notes: The unit of observation is a worker-day. The dependent variable is the average hourly quality adjusted output, calculated as was pre-registered, for each worker and day in column 1. In column 2 the dependent variable is the average total number of entries (correct and incorrect) per hour. Active typing time (column 3) is the average number of minutes per hour the worker spends actively typing. Mistakes (column 4) is the number of incorrect entries per 100 entries. Typing earnings (column 5) are the average earnings from data-entry in Rupees per hour. Present (column 6) is a dummy for whether the participant was present in the office at any time that day. Check-in time (column 7) and Check-out time (column 8) are measured in base 10 hours, such that a coefficient of 0.1 corresponds to 0.1 hours or 6 minutes. Hours of work (column 9) is calculated at check-out time minus check in time to capture the total time spent in the office. The sample includes all hours of the workday when workers are in the office. All columns control for worker, day-in-study, and month-by-year fixed effects as well as three lags of daily temperature and binary variables for the six underlying experimental arms. Standard errors are clustered at the worker level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels respectively.

to improve precision. We also remove the $<$0.01% of days in which participants were present in the office for less than 1 hour to minimize noise.

Recent work in econometrics has shown that bias known as "dynamic bias" can arise if the dependent variable is serially correlated and the specification does not control for the lag of the dependent variable (Klosin, 2024). This source of bias is often much greater than the bias that arises from including the lag as a control, known as "Nickell" bias (Nickell, 1981; Klosin, 2024). We therefore also report estimates of the main specification controlling for the lag of the dependent variable as well as lagged temperature to examine robustness.

5. Results

We document two sets of results: first, higher temperatures meaningfully reduce the contemporaneous productivity of workers in an individual desk-based job, and second, higher temperatures have a negative, statistically significant, economically large, and persistent effect on workers' productivity growth.

5.1. Heat Reduces Productivity

Table I displays the effect of temperature on average hourly productivity and its components. The dependent variable in column 1 is average hourly quality-adjusted output, which is calculated as the total quality-adjusted entries produced by a worker in a day divided by the total number of hours they were present at the office. The effect of temperature is negative and statistically significant, suggesting a one degree Celsius increase in average daily temperature reduces average hourly quality-adjusted output by 20.89 entries, or 1.1%. Columns 2 through 4 decompose this effect into the effect on the total number of entries (i.e. both correct and incorrect entries), the number of minutes spent actively typing, and the number of mistakes, respectively. The effect of temperature on the total number of entries is very similar to the effect on quality-adjusted output, suggesting that the effect is driven primarily by fewer characters entered rather than a substantial increase in mistakes. Columns 3 and 4 confirm this intuition. A one degree Celsius increase in average daily temperature reduces the time spent actively typing by 1.0%, accounting for the strong majority of the decline in entries. Mistakes also increase by 0.9%. However, with less than 1 mistake per 100 entries on average, this increase in errors has little influence on overall productivity as it results in roughly 0.14 additional errors per hour. This decline in time spent typing and entries leads to a corresponding decrease in pay, which is a function of both time spent typing and output: average typing earnings per hour falls by 1.3% (column 5). Further, consistent with an increase in breaks being the primary driver of productivity declines, we find no relationship between temperature and several measures of cognitive performance which capture sustained attention (PVT), working memory (Corsi), and inhibitory control (Hearts and Flowers) (Table A.2).[11] However, we cannot rule out that other measures of cognition beyond the three tested here – sustained attention, inhibitory control, and working memory – could be impacted.

The results described above show that higher temperatures affect contemporaneous productivity by reducing the amount of time spent actively working conditional on coming to work. In addition,

[11]While previous work has found impacts of heat on cognition, the results are sensitive to the domain being tested, the stakes of the test, and the baseline ability of the group such that null results appear with some regularity (Graff Zivin et al., 2018, 2020; Park, 2022; Deschênes and Greenstone, 2007). These null results are unlikely to reflect poor measurement: the same measures detected differences in cognition due to napping, suggesting they are sensitive enough to pick up differences in cognition in this population.

higher temperatures could also affect the extensive margin of labor supply or the time of arrival and departure (Graff Zivin and Neidell, 2014). To investigate this margin of adjustment, columns 6-8 document whether temperature has an effect on worker absenteeism and arrival and departure times. In column 6, the dependent variable is a binary variable that equals one if the worker is present at the office on a given day, and zero otherwise. In columns 7 and 8, the dependent variables are the time the worker checked into and out of the office, respectively. Temperature has no significant effect on absenteeism (column 6) or on arrival and departure times (columns 7–8): all three effects are small (less than 30 seconds for check-in/check-out) and statistically indistinguishable from zero, consistent with (Cai et al., 2018) but in contrast to the absenteeism effects found in (Somanathan et al., 2021). Workers in this context thus do not adjust labor supply in response to higher temperatures, either by altering the number of days worked or by altering when they choose to work.[12]

Robustness We conduct five robustness checks. First, we aggregate these measures to estimate the total impact on daily outcomes in Table A.3 with very similar results: declines in productivity are large, statistically significant, and driven by additional breaks. Second, due to the inclusion of the lag of temperature as a control in Table I, the main specification excludes days after a worker is not present in the office because the lag of temperature is missing in that case. We address this in two ways: Table A.4 follows the specifications in Table I, but fills in these missing lags with the average temperature during work hours with similar results. The third check is displayed in Table A.5, which reproduces Table I without the control for the lag of temperature. In both of these cases, we find significant reductions in productivity driven by reduction in time spent typing and no change in hours. Fourth, we examine robustness to the heat-index, which also captures humidity in Table A.6.[13] Although the magnitudes shift slightly as is expected given the different scaling, the effects are qualitatively unchanged with similar orders of magnitude and significance (Lai et al., 2023). Lastly, to ensure that results are not driven by pollution — which is correlated with temperature and has detrimental effects on productivity — we estimate a specification with $PM_{2.5}$, PM_{10}, NO_2,

[12]This result also reduces the likelihood of any selection effects in the sample we use to estimate the effect of temperature on productivity.

[13]We choose the heat index rather than the wet-bulb temperature given that the work takes place in a shaded location with a few fans helping to create light airflow, matching the heat-index assumptions well while still capturing differences in how temperatures feel due to differences in humidity.

SO_2, and 0_3 as additional covariates. We find that the effects of temperature are once again highly robust to the inclusion of all of these pollution controls (Table A.7).

5.2. Heat Slows Productivity Growth

Next, we examine whether higher temperatures affect not only the level of productivity, but also productivity growth. The data entry task is well-suited to answer this question because none of the recruited workers have previous data-entry experience, which leads to steep productivity growth as they begin the job. Panel A of Figure II shows the average output per day over time which grows roughly four-fold during the first 16 days on the job. To study the relationship between temperature and productivity growth, we begin with a simple associational plot examining the rate of output growth over time for those exposed to above median temperatures (light blue) versus those exposed to below median temperatures (dark blue) (Figure II, Panel B). We normalize the productivity measure by day one output to account for any initial differences across individuals, either due to individual features or contemporaneous temperature. There is a clear difference in the rate of growth over time by temperature: those with below median temperatures experience a more rapid increase in productivity over time. The raw difference in productivity growth between above median average daytime temperatures and below median average daytime temperatures is roughly 15.6% on the last day of the learning period. However, these differences may underestimate or overestimate true causal effects given differences in the population enrolled over time or partial acclimatization to higher temperatures between seasons. We turn to our previously described empirical strategy to address these concerns.

Drawing on estimating Equation 2 to identify causal effects, in Table II we document the effect of temperature on productivity growth, defined as the difference in average hourly productivity between block p and $p-1$, divided by average hourly productivity in block $p-1$. The sum of all temperature coefficients is presented at the top of the column as it captures the effect of interest: the effect of higher temperatures on productivity growth (Dell et al., 2012). However, for completeness, we report both the individual coefficients associated with temperature and all its lags below as well. In columns 1 and 2, only contemporaneous temperature is included as an explanatory variable.

Table II: The Effect of Temperature on Productivity Growth

	Dependent Variable: **Productivity Growth**							
	N = 0 Lags		N = 1 Lag		N = 2 Lags		N = 3 Lags	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sum of Temperature Coefficients	-0.007*	-0.007*	-0.005	-0.008	-0.015***	-0.017***	-0.026**	-0.030***
	[0.082]	[0.078]	[0.339]	[0.110]	[0.008]	[0.003]	[0.013]	[0.002]
Temperature ($°C$)	-0.007*	-0.007*	-0.011**	-0.009*	-0.013**	-0.012**	-0.019***	-0.017***
	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)	(0.007)	(0.006)
Lag 1 of Temperature			0.006	0.002	0.005	0.001	0.002	-0.001
			(0.005)	(0.005)	(0.006)	(0.006)	(0.008)	(0.008)
Lag 2 of Temperature					-0.007	-0.006	-0.006	-0.006
					(0.006)	(0.005)	(0.006)	(0.005)
Lag 3 of Temperature							-0.003	-0.006
							(0.006)	(0.005)
Control for Lag of Dependent Variable	No	Yes	No	Yes	No	Yes	No	Yes
Dependent Variable Mean	0.0964	0.0964	0.0980	0.0980	0.0944	0.0944	0.107	0.107
Observations	2382	2382	2050	2050	1919	1919	1455	1455

Notes: The unit of observation is a worker by two-day block. The sample is the first 16 days of the study, while productivity is growing. The dependent variable is the growth rate of productivity, calculated as the difference in hourly quality-adjusted output between block p and $p-1$, divided by hourly quality-adjusted output in block $p-1$. All columns control for worker, two-day block, and month-by-year fixed effects, and binary variables for the six underlying experimental arms. The number of lags of temperature in each specification is indicated in the table at the top of the column. Even-numbered columns additionally control for the lag of the dependent variable. p-values are shown in square brackets. Standard errors are clustered at the worker level and are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels respectively.

Column 2 additionally controls for the lag of the dependent variable, to ensure that dynamic bias is not driving the results, and finds very similar results with a decline in growth of 0.7 percentage points for each additional degree Celsius of average daily temperature (Klosin, 2024). The subsequent columns show the results adding one, two, and three lags of temperature, both with and without the control for the lag of the dependent variable. We include a maximum of 3 lags – or 6 days – given that the period of productivity growth is 16 days in total, resulting in a tradeoff between the number of included lags and the sample size. This choice results in between 10 and 16 days of data for the estimation of these results. We find that the total effect of temperature on productivity growth is negative and increases in magnitude as we add more lags to the specification, providing additional support for temperature's impact on growth above and beyond level effects. The estimate in column 8, with three lags as well as controls for the lag of the dependent variable is our preferred point estimate and suggests that a one degree Celsius increase in the average temperature over the last six days reduces productivity growth by 3.0 percentage points. Results are robust to using the

Table III: No Evidence of Convergence in Productivity Following Heat Exposure

	Panel A: Heterogeneous Effect of Temperature on Productivity Growth				
	Learning Half	Post-Learning		Learning Half	Post-Learning
	(1)	(2)		(3)	(4)
Sum of Temperature Coefficients	-0.026**	0.006		-0.030***	0.000
	[0.013]	[0.406]		[0.002]	[0.988]
Lag of Dependent Variable	No	No		Yes	Yes
Dependent Variable Mean	0.107	0.0182		0.107	0.0182
Observations	1455	1324		1455	1324
	Panel B: Productivity in Final Study Week				
	Quality Adjusted Output	**Total Number of Entries**	**Active Typing Time**	**Mistakes** (per 100 entries)	**Typing Earnings**
	(1)	(2)	(3)	(4)	(5)
Temperature ($°C$)	-22.213	-24.216	-0.363***	-0.000	-0.493
	(15.666)	(15.661)	(0.109)	(0.010)	(0.301)
Mean Learning Half Temperature	-127.285**	-119.597*	-0.603**	0.028	-2.222**
	(61.740)	(61.789)	(0.258)	(0.036)	(1.066)
Dep. Var. Mean	2068.8	2194.4	29.74	0.668	39.15
Observations	1495	1495	1495	1495	1495

Notes: Panel A: The unit of observation is a worker by two-day block. The dependent variable is the growth rate of productivity, calculated as the difference in hourly quality-adjusted output between block p and $p-1$, divided by hourly quality-adjusted output in block $p-1$. All columns control for worker, two-day block, and month-by-year fixed effects. Columns (1) and (3) use the first 16 days while growth is positive ("learning"); columns (2) and (4) use the remainder ("post-learning"). Columns (3)–(4) additionally include the lag of the dependent variable. p-values are shown in square brackets. Panel B: Examines productivity measures in the final week of the study, with the dependent variable as labeled in each column header. Temperature is the average daily temperature on that day. "Mean Learning Half Temperature" is the mean temperature experienced over the first 16 days of the study. The specification includes fixed effects and controls as described in Equation 1 with two exceptions. First we exclude worker fixed effects which are co-linear with the regressor of interest, though month-year fixed effects are maintained to address confounding by recruitment wave. Second, we include 2 additional lags (5 in total) of temperature to account for the gap between the first 16 days of the study and the final week, which starts on day 21. Standard errors are clustered at the worker level and are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels respectively.

difference in log productivity between t and t-1 as the dependent variable (Table A.8).

Falsification Tests Two placebo tests validate the causal interpretation of this result. First, a natural implication of a causal interpretation is that future temperature should not predict current productivity growth, even given the serial correlation in temperature. Leads thus serve as a placebo, since tomorrow's weather cannot cause today's learning. If the leads mattered, we would suspect an omitted variable, reverse causality, or anticipatory behavior. Appendix Figure A.2 explores this by reproducing columns 7 and 8 of Table II, but sums the coefficients on the leading values of temperature instead of the lag coefficients. The sum of the lead coefficients is of the opposite sign and roughly one-third as large. This suggests that the effect of temperature on workers' productivity growth is specific to temperatures experienced in the past while they are actively learning, supporting a causal interpretation of the main result.

Second, we leverage the difference in the rate of productivity growth over time: the effect of temperature on productivity growth should be present during the first half of the study, when productivity grows rapidly, but should be absent in the second half, which sees very limited productivity growth (Figure II, Panel A). To investigate this, we estimate Equation 2 separately for the first and second half of the study. Panel A of Table III confirms this pattern. Columns 1 and 3 replicate the results in Table II columns 7 and 8. Consistent with a causal effect, these coefficients become indistinguishable from zero for the second half of the study when productivity growth stalls, suggesting that the effects in the first half are not the result of simple persistence or serial correlation.

Persistence of the Effect An important question for the overall economic costs of these effects is whether the productivity deficits caused by heat during the learning period are temporary, with later convergence, or whether they persist. We examine this in two ways.

First, we ask whether workers exposed to higher temperatures during the learning period exhibit accelerated productivity growth in the post-learning period, consistent with catch-up. For full catch-up to occur, the sum of temperature coefficients in the post-learning period would need to be equal in magnitude but opposite in sign. Panel A of Table III shows no such pattern: the sum of temperature coefficients in the post-learning period is an order of magnitude smaller than in the learning period, mixed in sign, and statistically insignificant, providing no evidence of meaningful convergence.

Second, we directly test whether heat exposure during the learning period leaves a permanent mark on the level of productivity at the end of the study. Panel B examines productivity in the final week of the study as a function of mean temperature experienced during the first 16 days, controlling for contemporaneous temperature in the final week as well as temperature on all intervening days. We find that each additional degree Celsius experienced during the learning period is associated with a 6.2% reduction in productivity and a 5.7% reduction in earnings in the final week of the study — effects that are statistically significant and economically meaningful. Together, these results suggest that productivity losses from heat during the learning period persist for at least two weeks beyond that period. While we cannot rule out catch-up in the longer run, there is no structural reason to expect a delayed catch-up rather than an immediate one. To the extent these effects persist beyond the study period, the implications are compounding: workers who learn less during hot periods carry a productivity deficit forward, adding to the direct costs of heat exposure.

6. Conclusion

This study provides evidence that high temperatures not only impede productivity in the moment, but also reduce the rate of productivity growth. Using detailed data from 452 data-entry workers in Chennai, India, we find that a one degree Celsius increase in temperature reduces same-day average hourly productivity by 1.1%, resulting in a reduction in hourly earnings of 1.3% for the worker. The fine-grained data allow us to isolate the mechanisms for this change, highlighting the role of increased breaks rather than changes in days or hours worked. Further, we document that higher temperatures significantly slow learning on the job, with our preferred estimate indicating that a one degree increase in average temperature over the past week reduces the rate of productivity growth by 3.0 percentage points. Critically, these losses persist throughout the study: workers exposed to higher temperatures during the learning period remain less productive at the end of the study, with no evidence of catch-up. This deficit has direct costs for workers, who experience slower wage growth, and for firms, which face higher training costs and lower productivity.

These challenges are likely to be more acute in low- and middle-income countries, where faster warming and limited climate control coincide with two forces that make learning-driven productivity gains unusually frequent. First, high turnover means firms are constantly training new workers: Blattman and Dercon (2018) document turnover rates of 33% in the first month and 77% in the first year in Ethiopian garment factories. Second, even long-term employees with seemingly routine tasks may face recurring learning curves. Adhvaryu et al. (2023) show garment line efficiency rises over 40% in the first month of a new style, with new runs starting every 2–3 weeks. Similar dynamics are likely to arise in jobs like phone-based customer support or data entry as workers encounter new software or procedures. Temperature's impact on skill acquisition is thus a recurring cost imposed each time a worker — new or established — has to learn something again. As low- and middle-income economies shift toward more knowledge-intensive work, this interaction of heat, limited climate control, and repeated skill acquisition may constrain both worker advancement and broader economic development.

References

Adhvaryu, Achyuta, Anant Nyshadham, and Jorge Tamayo, "Managerial quality and productivity dynamics," *The Review of Economic Studies*, 2023, *90* (4), 1569–1607.

_ **, Namrata Kala, and Anant Nyshadham**, "The light and the heat: Productivity co-benefits of energy-saving technology," *Review of Economics and Statistics*, 2020, *102* (4), 779–792.

Anderson, Michael L., "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects," *Journal of the American Statistical Association*, 2008, *103* (484), 1481–1495.

Arrow, Kenneth J, "The economic implications of learning by doing," *The review of economic studies*, 1962, *29* (3), 155–173.

Becker, Gary S, "Investment in human capital: A theoretical analysis," *Journal of political economy*, 1962, *70* (5, Part 2), 9–49.

Ben-Porath, Yoram, "The production of human capital and the life cycle of earnings," *Journal of political economy*, 1967, *75* (4, Part 1), 352–365.

Bessen, James E, "Productivity adjustments and learning-by-doing as human capital," *CES Bureau of the Census WP*, 1997, pp. 97–17.

Bessone, Pedro, Gautam Rao, Frank Schilbach, Heather Schofield, and Mattie Toma, "The economic consequences of increasing sleep among the urban poor," *The Quarterly Journal of Economics*, 2021, *136* (3), 1887–1941.

Blattman, Christopher and Stefan Dercon, "The impacts of industrial and entrepreneurial work on income and health: Experimental evidence from Ethiopia," *American Economic Journal: Applied Economics*, 2018, *10* (3), 1–38.

Cai, Xiqian, Yi Lu, and Jin Wang, "The impact of temperature on manufacturing worker productivity: Evidence from personnel data," *Journal of Comparative Economics*, 2018, *46* (4), 889–905.

Copernicus Atmosphere Monitoring Service, "CAMS global reanalysis (EAC4)," 2020. Accessed June 1, 2026.

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken, "Temperature shocks and economic growth: Evidence from the last half century," *American Economic Journal: Macroeconomics*, 2012, *4* (3), 66–95.

Deschenes, Olivier, "Temperature, productivity, and income," *IZA World of Labor*, 2023.

Deschênes, Olivier and Michael Greenstone, "The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather," *American economic review*, 2007, *97* (1), 354–385.

Garg, Teevrat, Maulik Jagnani, and Elizabeth Lyons, "Heat and team production: Experimental evidence from Bangladesh," 2025.

_ , _ , **and Vis Taraz**, "Temperature and human capital in India," *Journal of the Association of Environmental and Resource Economists*, 2020, *7* (6), 1113–1150.

Hancock, Peter A, Jennifer M Ross, and James L Szalma, "A meta-analysis of performance response under thermal stressors," *Human factors*, 2007, *49* (5), 851–877.

Huppertz, Maximilian, "Sacking the sales staff: Firm reactions to extreme weather and implications for policy design," Technical Report, mimeo 2024.

Ibarrarán, María Eugenia, Matthias Ruth, Sanjana Ahmad, and Marisa London, "Climate change and natural disasters: macroeconomic performance and distributional impacts," *Environment, development and sustainability*, 2009, *11* (3), 549–569.

Klosin, Sylvia, "Dynamic Biases of Static Panel Data Estimators," *arXiv preprint arXiv:2410.16112*, 2024.

Lagakos, David, Benjamin Moll, Tommaso Porzio, Nancy Qian, and Todd Schoellman, "Life cycle wage growth across countries," *Journal of Political Economy*, 2018, *126* (2), 797–849.

Lai, Wangyang, Yun Qiu, Qu Tang, Chen Xi, and Peng Zhang, "The effects of temperature on labor productivity," *Annual Review of Resource Economics*, 2023, *15* (1), 213–232.

Lazear, Edward P, Kathryn L Shaw, and Christopher T Stanton, "The value of bosses," *Journal of Labor Economics*, 2015, *33* (4), 823–861.

Levitt, Steven D, John A List, and Chad Syverson, "Toward an understanding of learning by doing: Evidence from an automobile assembly plant," *Journal of political Economy*, 2013, *121* (4), 643–681.

Liu, Zhao, Bruce Anderson, Kai Yan, Weihua Dong, Hua Liao, and Peijun Shi, "Global and regional changes in exposure to extreme heat and the relative contributions of climate and population change," *Scientific Reports*, 2017, *7* (1), 43909.

Lo Bello, Salvatore, Maria Laura Sanchez Puerta, and Hernan Winkler, "From Ghana to America: The Skill Content of Jobs and Economic Development," Technical Report, IZA Discussion Papers 2019.

LoPalo, Melissa, "Temperature, worker productivity, and adaptation: evidence from survey data production," *American Economic Journal: Applied Economics*, 2023, *15* (1), 192–229.

Ma, Xiao, Alejandro Nakab, and Daniela Vidart, "Human capital investment and development: The role of on-the-job training," *Journal of Political Economy Macroeconomics*, 2024, *2* (1), 107–148.

Mas, Alexandre and Enrico Moretti, "Peers at work," *American Economic Review*, 2009, *99* (1), 112–145.

Masuda, Yuta J, Teevrat Garg, Ike Anggraeni, Kristie Ebi, Jennifer Krenz, Edward T Game, Nicholas H Wolff, and June T Spector, "Warming from tropical deforestation reduces worker productivity in rural communities," *Nature Communications*, 2021, *12* (1), 1601.

Mendelsohn, Robert, Ariel Dinar, and Larry Williams, "The distributional impact of climate change on rich and poor countries," *Environment and development economics*, 2006, *11* (2), 159–178.

Mincer, Jacob, "On-the-job training: Costs, returns, and some implications," *Journal of political Economy*, 1962, *70* (5, Part 2), 50–79.

Nickell, Stephen, "Biases in dynamic models with fixed effects," *Econometrica: Journal of the econometric society*, 1981, pp. 1417–1426.

Park, R Jisung, "Hot temperature and high-stakes performance," *Journal of Human Resources*, 2022, *57* (2), 400–434.

_ , Joshua Goodman, Michael Hurwitz, and Jonathan Smith, "Heat and learning," *American Economic Journal: Economic Policy*, 2020, *12* (2), 306–339.

Pittock, A Barrie, *Climate change: turning up the heat*, Routledge, 2017.

Saha, Suranjana, Shrinivas Moorthi, Hua-Lu Pan, Xiao Wu, Jing Wang, Suranjana Nadiga, Patrick Tripp, Richard Kistler, Eugene Woollen, David Behringer, Hong Liu, Dale Stokes, Robert Grumbine, Gregory Gayno, Jennifer Wang, Yingyong T. Hou, Hung-Yi Chuang, Hong-I Juang, Joseph Sela, Mark Iredell, Ralph Treadon, David Kleist, Paul Van Delst, Daniel Keyser, John Derber, Michael Ek, Jin Meng, Heping Wei, Rui Yang, Steve Lord, Huug van den Dool, Arun Kumar, Wanqiu Wang, Curtis Long, Muthuvel Chelliah, Yan Xue, Boyin Huang, Jae-Kyung Schemm, Wesley Ebisuzaki, Rui Lin, Pingping Xie, Mingyue Chen, Shuguang Zhou, Wayne Higgins, Cheng-Zhi Zou, Qian Liu, Yun Chen, Yong Han, Lidia Cucurull, Richard W. Reynolds, George Rutledge, and Mitchell Goldberg, "The NCEP Climate Forecast System Reanalysis," *Bulletin of the American Meteorological Society*, 2010, *91* (8), 1015–1057.

Sharma, Disha and Denise Mauzerall, "CPCB continuous monitoring air quality d ata for India, from 2015-2019," 2021.

Somanathan, Eswaran, Rohini Somanathan, Anant Sudarshan, and Meenu Tewari, "The impact of temperature on productivity and labor supply: Evidence from Indian manufacturing," *Journal of Political Economy*, 2021, *129* (6), 1797–1827.

Walter, Torsten Figueiredo and Niclas Moneke, "Endogenous Sample Selection," 2024.

Weather and Climate, "Chennai Weather & Climate | Year-Round Guide with Graphs," 2024. Average monthly rainfall, temperature, and sunshine data for Chennai, India.

Wilkinson, Allie, "Expanding tropics will play greater global role, report predicts," *ScienceInsider*, 2014.

Zhang, Peng, Olivier Deschenes, Kyle Meng, and Junjie Zhang, "Temperature effects on productivity and factor reallocation: Evidence from a half million Chinese manufacturing plants," *Journal of Environmental Economics and Management*, 2018, *88*, 1–17.

Zivin, Joshua Graff and Matthew Neidell, "Temperature and the allocation of time: Implications for climate change," *Journal of Labor Economics*, 2014, *32* (1), 1–26.

— , **Solomon M Hsiang, and Matthew Neidell**, "Temperature and human capital in the short and long run," *Journal of the Association of Environmental and Resource Economists*, 2018, *5* (1), 77–105.

— , **Yingquan Song, Qu Tang, and Peng Zhang**, "Temperature and high-stakes cognitive performance: Evidence from the national college entrance examination in China," *Journal of Environmental Economics and Management*, 2020, *104*, 102365.

A. Appendix

FIGURE A.1: Data-Entry Software

Left side

Right side

Notes: This figure shows screenshots of the data-entry software used by workers in the study. Panel A shows the left side of the screen, which contains the fictional data to be transcribed by workers. Panel B shows the right side of the screen, where workers entered the data. The text was divided into boxes, and participants had to progress linearly, without the ability to skip fields or return to earlier fields. They also could not move on to the next box if less than 20% of characters were correct.

FIGURE A.2: The Effect of Temperature Leads on Productivity Growth

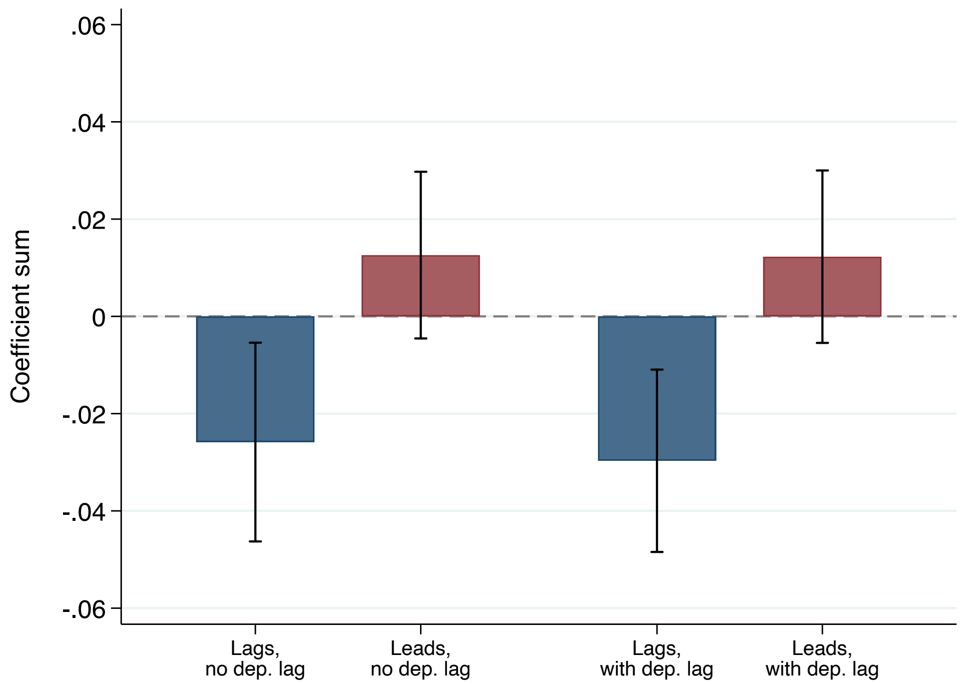

Notes: This figure presents a falsification test for the productivity-growth results in Table II. The unit of observation, estimation sample, dependent variable, and controls are the same across all four specifications and match those in Table II. The two maroon bars plot the sum of the coefficients on three lead temperature variables. For comparison, the two navy bars reproduce the sum of the coefficients on contemporaneous and three lagged temperature variables from the last two columns of Table II. Specifications labeled "with dep. lag" control for the lag of the dependent variable; those labeled "no dep. lag" exclude it. Standard errors are clustered at the worker level and 95% confidence intervals are reported.

Table A.1: Balance in Characteristics Between Workers Enrolled at Different Times

	April-September	**October-March**	**p-value, 1 = 2**
	(1)	(2)	(3)
Literate in English (=1)	0.799	0.781	0.645
	(0.030)	(0.025)	
Prior Computer Experience (=1)	0.247	0.306	0.179
	(0.033)	(0.028)	
Years of Education	9.983	10.309	0.241
	(0.210)	(0.177)	
Math Ability (=1)	0.632	0.615	0.716
	(0.037)	(0.029)	

Notes: This table displays average worker characteristics for participants enrolled in different seasons of the year. Column 1 reports the average and standard error for characteristics of workers enrolled in summer, from April to September. Column 2 reports average characteristics for workers enrolled in winter, from October to March. Column 3 displays the p-value for a t-test of equality between columns 1 and 2.

Table A.2: The Contemporaneous Effect of Temperature on Cognition

	Dependent Variable is			
	Cognition Index	**PVT**	**Corsi**	**Hearts and Flowers**
	(1)	(2)	(3)	(4)
Temperature ($^{\circ}C$)	0.003	0.001	-0.011	0.005
	(0.007)	(0.006)	(0.010)	(0.006)
Observations	6846	6701	3383	3412
R-squared	0.523	0.486	0.526	0.739

Notes: The unit of observation is the worker-day. The dependent variable in column 1 is an Anderson index composed of all three measures of cognition in columns 2 - 4. PVT (attention) was administered daily. Corsi (memory) and Hearts and Flowers (inhibitory control) were administered on alternate days. Score coding on each task is as pre-registered in the original experiment, where higher scores indicate better performance. All columns control for worker, day-in-study, and month-by-year fixed effects, and binary variables for the six underlying experimental arms. Standard errors are clustered at the worker level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels respectively.

Table A.3: The Contemporaneous Effect of Temperature on Daily Work Outcomes

	Dependent Variable is **Daily**				
	Quality Adjusted Output	**Total Number of Entries**	**Active Typing Time**	**Mistakes** (per 100 entries)	**Typing Earnings**
	(1)	(2)	(3)	(4)	(5)
Temperature ($^{\circ}C$)	-405.519***	-422.395***	-5.836***	0.007**	-8.260***
	(37.358)	(39.005)	(0.513)	(0.003)	(0.901)
Dep. Var. Mean	14402.8	15412.6	230.7	0.827	279.5
Observations	7268	7268	7268	7267	7268

Notes: This table is identical to Table I columns 1-5, with the exception that all dependent variables except column 4 are daily totals rather than average hourly values. Column 4, mistakes, remains a rate to avoid confounding the estimate with changes in overall output. All columns control for worker, day-in-study, and month-by-year fixed effects, lagged temperature, and binary variables for the six underlying experimental arms. Standard errors are clustered at the worker level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels respectively.

Table A.4: The Contemporaneous Effect of Temperature Accounting for Absent Days

	Average Hourly Outcomes					Daily Outcomes			
	Quality Adjusted Output	**Total Number of Entries**	**Active Typing Time**	**Mistakes** (per 100 entries)	**Typing Earnings**	**Present** (=1)	**Check-in Time**	**Check-out Time**	**Hours of Work**
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Temperature ($^{\circ}C$)	-17.824***	-18.287***	-0.249***	0.004	-0.393***	-0.000	0.004	-0.008	-0.011
	(2.517)	(2.570)	(0.032)	(0.003)	(0.058)	(0.002)	(0.004)	(0.007)	(0.008)
Dep. Var. Mean	1736.7	1862.0	28.72	0.868	33.71	0.893	10.59	18.33	7.740
Observations	10881	10881	10881	10878	10881	12655	10884	10884	10884

Notes: This table is identical to Table I with the exception that temperature lags are backfilled using the average temperature during standard working hours (9 a.m. to 8 p.m.) for days a worker was absent in columns (1-5). This approach maximizes the data available for estimation. All columns control for worker, day-in-study, and month-by-year fixed effects, and binary variables for the six underlying experimental arms. Standard errors are clustered at the worker level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels respectively.

Table A.5: The Contemporaneous Effect of Temperature without Dependent Variable Lag as Control

	Average Hourly Outcomes					Daily Outcomes			
	Quality Adjusted Output	**Total Number of Entries**	**Active Typing Time**	**Mistakes** (per 100 entries)	**Typing Earnings**	**Present** (=1)	**Check-in Time**	**Check-out Time**	**Hours of Work**
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Temperature ($^{\circ}C$)	-12.446***	-12.844***	-0.159***	0.003	-0.284***	0.002	-0.003	-0.005	-0.003
	(2.347)	(2.399)	(0.028)	(0.003)	(0.055)	(0.002)	(0.003)	(0.006)	(0.007)
Dep. Var. Mean	1736.7	1862.0	28.72	0.868	33.71	0.893	10.59	18.33	7.740
Observations	10881	10881	10881	10878	10881	12655	10884	10884	10884

Notes: This table is identical to Table I with the exception that lags of temperature for previous days are excluded from the regression. All columns control for worker, day-in-study, and month-by-year fixed effects, and binary variables for the six underlying experimental arms. Standard errors are clustered at the worker level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels respectively.

Table A.6: The Contemporaneous Effect of a Heat Index on Labor Outcomes

	Average Hourly Outcomes					Daily Outcomes			
	Quality Adjusted Output	**Total Number of Entries**	**Active Typing Time**	**Mistakes** (per 100 entries)	**Typing Earnings**	**Present** (=1)	**Check-in Time**	**Check-out Time**	**Hours of Work**
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Heat Index	-24.716***	-25.615***	-0.352***	0.003*	-0.553***	-0.000	-0.000	-0.006	-0.006
	(1.994)	(2.080)	(0.025)	(0.002)	(0.051)	(0.002)	(0.004)	(0.006)	(0.008)
Dep. Var. Mean	1820.3	1948.5	29.23	0.827	35.27	0.893	10.59	18.33	7.740
Observations	7266	7266	7266	7265	7266	12655	10884	10884	10884

Notes: This table is identical to Table I with the exception that temperature is defined as the heat index. All columns control for worker, day-in-study, and month-by-year fixed effects, lagged temperature, and binary variables for the six underlying experimental arms. Standard errors are clustered at the worker level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels respectively.

Table A.7: The Contemporaneous Effect of Temperature Controlling for Pollution

	Average Hourly Outcomes					Daily Outcomes			
	Quality Adjusted Output	**Total Number of Entries**	**Active Typing Time**	**Mistakes** (per 100 entries)	**Typing Earnings**	**Present** (=1)	**Check-in Time**	**Check-out Time**	**Hours of Work**
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Temperature ($^{\circ}C$)	-21.509***	-21.600***	-0.297***	0.009***	-0.464***	-0.000	0.005	-0.005	-0.009
	(2.968)	(3.069)	(0.039)	(0.004)	(0.080)	(0.002)	(0.004)	(0.007)	(0.009)
Dep. Var. Mean	1820.3	1948.4	29.23	0.827	35.27	0.893	10.59	18.33	7.740
Observations	7268	7268	7268	7267	7268	12655	10884	10884	10884

Notes: This table is identical to Table I with the exception that all columns control for average daily pollution, proxied by $PM_{2.5}$, PM_{10}, $N0_2$, SO_2, and O_3. For each pollution proxy, we source data from the measurement station in Chennai that recorded a measurement and aggregate to the daily level. In addition to pollution, all columns control for worker, day-in-study, and month-by-year fixed effects, lagged temperature, and binary variables for the six underlying experimental arms. Standard errors are clustered at the worker level. ***, **, and * indicate significance at the 1%, 5%, and 10% levels respectively.

Table A.8: The Effect of Temperature on Productivity Growth Using Change in Log of Productivity

	Dependent Variable: Difference in **Log Productivity**							
	N = 0 Lags		N = 1 Lag		N = 2 Lags		N = 3 Lags	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sum of Temperature Coefficents	-0.008**	-0.007**	-0.005	-0.007*	-0.013**	-0.014***	-0.023**	-0.027***
	[0.017]	[0.017]	[0.209]	[0.055]	[0.010]	[0.005]	[0.011]	[0.001]
Temperature ($°C$)	-0.008**	-0.007**	-0.011**	-0.009**	-0.013**	-0.010**	-0.017***	-0.013**
	(0.003)	(0.003)	(0.004)	(0.004)	(0.005)	(0.004)	(0.006)	(0.005)
Lag 1 of Temperature			0.006	0.002	0.005	0.001	0.002	-0.002
			(0.005)	(0.004)	(0.006)	(0.005)	(0.007)	(0.006)
Lag 2 of Temperature					-0.005	-0.005	-0.007	-0.007
					(0.004)	(0.004)	(0.005)	(0.004)
Lag 3 of Temperature							-0.002	-0.005
							(0.005)	(0.004)
Control for Lag of Dependent Variable	No	Yes	No	Yes	No	Yes	No	Yes
Dependent Variable Mean	0.0656	0.0656	0.0677	0.0677	0.0647	0.0647	0.0793	0.0793
Observations	2382	2381	2050	2050	1919	1919	1455	1455

Notes: This table is identical to Table II with the exception that the dependent variable, productivity growth, is calculated as the difference in the log of hourly quality-adjusted output between block p and $p-1$. All columns control for worker, two-day block, and month-by-year fixed effects, and binary variables for the six underlying experimental arms. The number of lags of temperature in each specification is indicated in the table at the top of the column. Even-numbered columns additionally control for the lag of the dependent variable. p-values are shown in square brackets. Standard errors are clustered at the worker level and are presented in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels respectively.