HIGH TEMPERATURES IMPEDE PRODUCTIVITY GROWTH ON THE JOB*

Isadora Frankenthal† Heather Schofield‡ Pedro Bessone§

We study how temperature affects productivity growth among 452 Indian data-entry workers new to the job. Higher temperatures impede on-the-job learning, reducing the rate of productivity growth by 3.2 percentage points per degree Celsius over the past week. These effects are in addition to temperatures' contemporaneous impacts: each additional degree Celsius lowers same-day hourly productivity by 1.1% via increased breaks. The finding that higher temperatures not only reduce contemporaneous output, but also hinder productivity growth suggests a potential underpinning for temperature's impact on macroeconomic growth, with implications for developing countries facing rising temperatures, high turnover, and increasingly knowledge-based economies.

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, Jacob Moscona, Anant Nyshadham, Ben Olken, and Frank Schilbach 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. All remaining errors are our own.

†MIT (isadoraf@mit.edu)

‡Cornell University (hws44@cornell.edu). Corresponding author. Sage Hall 375, Ithaca, NY 14850.

§Uber (tepedino@mit.edu)

1. Introduction

Rising global temperatures have been shown to reduce productivity and slow economic growth (Adhvaryu et al., 2020; Somanathan et al., 2021; LoPalo, 2023; Huppertz, 2024; Dell et al., 2012; Burke et al., 2015). Although previous work has demonstrated that higher temperatures reduce macroeconomic growth, there is relatively little empirical evidence on the mechanisms through which these effects occur. One possibility is that, beyond reducing contemporaneous productivity, elevated temperatures may also hinder productivity growth by impeding workers' ability to acquire skills and knowledge on the job, a requirement that is surprisingly common even in seemingly routine tasks such as garment production (Adhvaryu et al., 2023). This channel is likely to be particularly important for jobs involving desk-based work and cognitive tasks. These types of non-routine cognitive jobs are increasingly prevalent in less developed countries, where the adverse effects of climate change are projected to be more severe (Lo Bello et al., 2019; 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 on a piece rate basis, 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. The study's location, Chennai, India, is representative of typical tropical environments. The daily *average* temperature ranges from 21.0°C to 36.5°C, and the yearly average maximum daily temperature is 32.5°C (Ross et al., 2006; Burke et al., 2015).

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 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

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

in hourly productivity over a month, with over 90% of the gain in the first two weeks — to test whether temperature 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 3.2 percentage points (p=0.001). The causal interpretation of this result is reinforced by two placebo tests. First, we show that there in no relationship between temperature and productivity growth in the second half of the study, when productivity growth is essentially flat. Second, 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 ommitted variables, reverse causality, or anticipatory effects.

Our study contributes to three growing literatures on the potential economic impacts of climate change and rising temperatures. 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 factories, agriculture, and field surveying as well as detrimental effects on teamwork (Adhvaryu et al., 2020; Somanathan et al., 2021; LoPalo, 2023; Masuda et al., 2021; Huppertz, 2024; Garg et al., 2020). We document similar results in a novel context of desk-based office work that is exclusively cognitive. Additionally, in contrast to previous literature with the exception of Masuda et al. (2021), we focus on a context in which workers face strong incentives for productivity — 73% of their total payment for the day was a function of the time spent actively typing and output — rather than a flat wage. 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 temperature and economic growth (Dell et al., 2012) and to the drivers of that relationship, including declines in total factor productivity (TFP) at higher temperatures (Zhang et al., 2018). We provide the first evidence that heat slows the growth rate of labor productivity. This mechanism is likely to matter across many settings and may, in turn, contribute to slower macroeconomic growth. For example, factories in low-income countries often face annual turnover rates of 50–100%, requiring continual training of new workers which is slower and more costly at higher temperatures. Moreover, this effect might matter even among incumbent workers. Learning on the job is frequent even in seemingly routine work: in garment lines, efficiency rises 40–60% within a production run for a particular garment (Adhvaryu et al., 2023). As economies shift toward more knowledge-based and less routine work this effect of heat on productivity growth will be even more consequential (Lo Bello et al., 2019).

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 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). Direct

comparisons are complicated since labor productivity growth can be driven by factors beyond learning. Our study complements this work by examining temperature's effects on productivity growth — which is likely driven at least in part by learning on the job, 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. Given the screening procedures and recruitment method — which consisted of advertising the job in low-income neighborhoods — the workers were literate but typically from lower socio-economic strata. 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. Recruited participants worked in an office space in Chennai equipped with computers and the data-entry software. Each worker had their (fixed) desk. Using the software, workers were asked to transcribe what appeared to be handwritten text from the left portion of the screen to the right portion of the screen (Figure A.1). 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 by the software: after two minutes without active typing the software would note that the work was paused. Workers could pause and resume at any time.

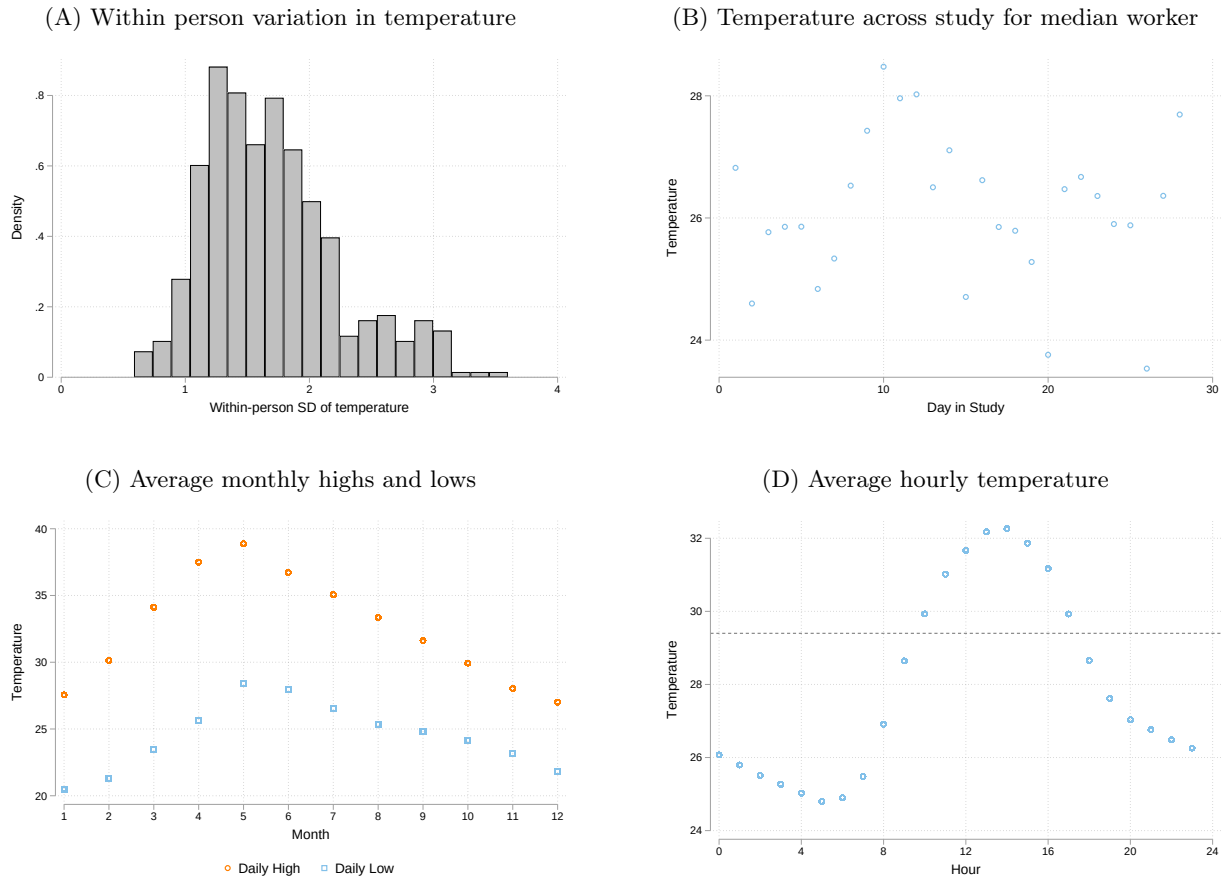
The job was flexible but workers were incentivised for production. They could choose whether or not to come to the office, their arrival and departure time between 9:30 a.m. and 8 p.m., and how much work to do when they were present.[2] A surveyor recorded the time they entered and left the office each 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. There were no penalties for breaks or absenteeism beyond the foregone earnings associated with that time.

2.2. Climate in Chennai

Chennai, India has a tropical climate. Using data from nearest weather station to the study office, mean temperature during working hours ("daytime") over the 18 month period in which workers

[2]A small subset of days were 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 the nearest weather station to the 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 (1.2 SD). 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.

were enrolled was 30.2°C (86.4°F), ranging from 22.8°–38.7°C (73°–101.7°F). 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 A). Temperatures typically peak around 2 p.m. (Figure I, panel B), though stay high until late afternoon or early evening. Maximal temperatures were over 40°C, however these days were rare, with only 4.1% of total days above this cutoff at any point. Nighttime temperatures are noticeably cooler, though often feel warm given the high humidity. 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 two datasets: 1) productivity of 452 workers over 4 weeks of work, and 2) hourly weather data from the nearest weather station to the office where the worker's were located in Chennai, India. 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 they spent actively typing, and their typing earnings, for the hours when they were present in the office. We use the hourly data to compute average hourly output, which is our main measure of productivity. 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 — data which is used to investigate the effect of temperature on the extensive margin of labor supply — and, if so, exactly what time they arrived and left. In addition, three tests were administered to measure cognition. First, a psychomotor vigilance test — designed to capture alertness — was administered daily, second, a hearts and flowers test — designed to measure inhibitory control — was administered on alternate days; and third, a Corsi task — designed to test 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, which we use to understand whether the effect of temperature is operating through cognition. The three metrics are also combined into an Anderson (2008) index.

Temperature We use hourly temperature data for the closest weather station to the office in Chennai, India, collected from the NCEP Climate Forecast System Reanalysis (CFSR) provided by the NOAA National Centers for Environmental Prediction (NCEP).[3] The data are described in Saha et al. (2010). These data record the temperature at the beginning of each hour in degrees Celsius, accurate to the nearest 0.1 degree.[4] As Chennai is located on the ocean and humidity is high throughout the year, we use temperature as reported rather than wet-bulb temperature for ease of interpretation.[5] The hourly temperature data is matched to the hourly productivity data at the hour-day level. From the hourly temperature data we also create a daily-level dataset which

[3]NOAA is an American government agency whose objective is to monitor global weather and climate.

[4]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 leading office temperatures to track external temperatures, though with a lag and at lower levels.

[5]Relative humidity ranges from 63% to 76%, peaking during the cooler rainy season in November (Weather and Climate, 2024).

records the average temperature for each workday. Given that temperature does not impact arrival or departure times from the office, this variable is created at the worker level to improve accuracy.[6]

4. Empirical Strategy

Effect on Productivity There are two ways in which 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:

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

Where i indexes workers, d indexes 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. $\mathrm{Temperature}_{t(i,d)}$ 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-year fixed effects, respectively. Worker fixed effects hold constant any differences across people in inherent ability. 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 to account for any potential differences in productivity across months due to different enrollment waves.[8] Following previous work that estimates the effect of temperature on output, our main specification also controls for lagged values of average 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.

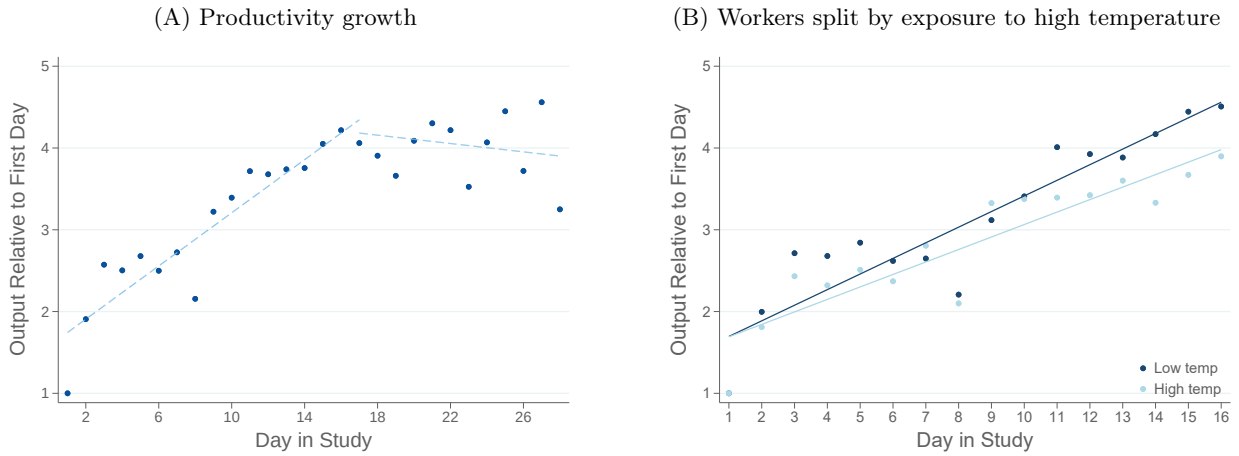
The advantage of this worker-by-day level specification is that it takes into account an important

[6] Temperature during the standard working day (9 a.m. to 8 p.m.) is used when estimating impacts on attendance as well arrival and departure time times from the office.

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

[8] Workers were enrolled in waves recruited from areas around the study office throughout the study period. One potential concern with this enrollment strategy is that workers enrolled in different waves could be of different ability while also facing different temperatures, confounding the effects we find. To mitigate this concern we show that workers recruited in different seasons are comparable in terms of their literacy in English, prior computer experience, years of education, and math ability, in Appendix Table A.1.

FIGURE II: Increase in Average Hourly Productivity Over Study Period



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 displays a binscatter of the learning period (the first 16 days) with a best-fit line, separating participants who are exposed to the top (light blue) and bottom (dark blue) quartile of average temperature over the study.

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.

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. We investigate whether higher temperatures decrease productivity growth by following the estimation strategy in Dell et al. (2012), who estimate the effect of temperature on the growth rate of GDP. 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, indexed by p, is therefore a two-day block. Temperature is calculated as the simple average of the temperature during work hours across each period. We estimate the following equation:

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

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). Productivity Growth Rate$_{ip}$ is calculated as average hourly productivity in period p minus average

hourly 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.[9] α_i, δ_p, and γ_m are worker, time period, and month-by-year fixed effects, respectively. Standard errors are clustered at the worker level.

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. 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.

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).[10] The sample consists of the 89.2% of workers who were present at the office more than 75% of the days 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.

5. Results

We document two sets of results: 1) higher temperatures meaningfully reduce the contemporaneous productivity of workers in a desk-based job, 2) higher temperatures have a negative, statistically significant, and economically large 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 20.84 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

[9] 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.

[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. Panel A of Figure II plots average productivity growth over time and the piece-wise linear function that fits the data.

Table I: The Effect of Temperature on Main 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)	-20.835***	-20.999***	-0.286***	0.008**	-0.454***	-0.000	0.004	-0.008	-0.011
	(2.829)	(2.922)	(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. Standard errors are clustered at the worker level. *, **, and *** indicate significance at the 1%, 5%, and 10% levels respectively.

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 an 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 drop 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 (Table A.2).

The results described above show that higher temperatures affect productivity by reducing the amount of time spent actively working conditional on coming to work. In addition, 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 in this context, 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. In contrast to the absenteeism effects found in (Somanathan et al., 2021) but consistent with (Cai et al., 2018), the effect of temperature is very small (less than 30 seconds) and statistically indistinguishable from zero across columns. In this context, workers do not adjust labor supply in response to higher temperatures

either by altering the number of days or by altering when they choose to work.[11]

Robustness We conduct four 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. Again, we find significant reductions in productivity driven by reduction in time spent typing. 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 pollution as an additional covariate. We find that the effects of temperature are robust to the inclusion of pollution controls (Table A.6).

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.[12] 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 higher temperatures versus those exposed to lower temperatures (Panel B of Figure II). We normalize the productivity measure by day one output to account for any initial differences across individuals. The raw difference in productivity growth between above median average daytime temperatures and below median average daytime temperatures is roughly 15.4% 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 production 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

[11] This result also ensures that the sample we use to estimate the effect of temperature on productivity is not positively selected.

[12] The break point at day 16 is data-driven using a piece-wise linear function in the raw productivity data.

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 Coefficents	-0.007*	-0.007*	-0.005	-0.008*	-0.016***	-0.018***	-0.028***	-0.032***
	[0.090]	[0.086]	[0.306]	[0.093]	[0.007]	[0.002]	[0.008]	[0.001]
Temperature ($°C$)	-0.007*	-0.007*	-0.010*	-0.009*	-0.013**	-0.012**	-0.019***	-0.017***
	(0.004)	(0.004)	(0.005)	(0.005)	(0.006)	(0.005)	(0.007)	(0.006)
Lag 1 of Temperature			0.006	0.001	0.005	0.000	0.002	-0.002
			(0.005)	(0.005)	(0.006)	(0.006)	(0.008)	(0.008)
Lag 2 of Temperature					-0.007	-0.007	-0.006	-0.006
					(0.006)	(0.005)	(0.006)	(0.005)
Lag 3 of Temperature							-0.004	-0.007
							(0.006)	(0.005)
Control for Lag of Dependent Variable	No	Yes	No	Yes	No	Yes	No	Yes
Dependent Variable Mean	0.0962	0.0962	0.0980	0.0980	0.0944	0.0944	0.107	0.107
Observations	2381	2380	2050	2049	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. 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.

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. The second column 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.[13] 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. 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 3.2 percentage points.

Falsification tests Two placebo tests validate the causal interpretation of this result. First, the effect of temperature on productivity growth should be significantly more pronounced during the first half of the study, when productivity grows rapidly, compared to the second half, which sees very

[13]Note that there is a significant tradeoff between the number of lags and the sample size. We include a maximum of 3 lags – or 6 days – given that the period of productivity growth is 16 days in total.

Table III: Falsification Tests

	Panel A: Heterogenous Effect of Temperature			
	Learning Half	Post-Learning	Learning Half	Post-Learning
	(1)	(2)	(3)	(4)
Sum of Temperature Coefficents	-0.028***	0.005	-0.032***	-0.001
	[0.008]	[0.467]	[0.001]	[0.931]
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: Effect of Temperature Leads on Growth			
	N = 1 Leads		N = 3 Leads	
	(1)	(2)	(3)	(4)
Sum of Lead Temperature Coefficents	0.007	0.011**	0.015	0.015
	[0.140]	[0.027]	[0.121]	[0.164]
Lag of Dependent Variable	No	Yes	No	Yes
Dependent Variable Mean	0.0969	0.0968	0.100	0.100
Observations	1971	1970	1448	1447

Notes: 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. Panel A splits the sample: 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. Panel B uses leads of temperature: columns (1)–(2) use one lead; columns (3)–(4) use three leads. Even-numbered columns also include 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.

limited, if any, 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 earlier 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.

Second, despite the serial correlation in temperature, a causal relationship requires that future temperature not predict current productivity growth. Leads thus serve as a placebo: tomorrow's weather cannot cause today's learning: if the leads mattered, we would suspect an ommitted variable, reverse causality, or anticipatory behavior. Panel B of Table III explores this by summing lead rather than lag coefficients. The sum of the lead coefficients is wrong-signed, much smaller in magnitude, and statistically insignificant. These results suggest that the effect of temperature on productivity growth is specific to past temperatures workers experience during the period of the study when on-the-job learning is most significant.

6. Conclusion

This study provides evidence that high temperatures not only impede productivity in the moment, but also reduce the rate of productivity growth among data-entry workers. 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 cognitive impairment or changes in days or hours worked. Further, we document that higher temperatures significantly slow productivity *growth*, with our preferred estimate indicating that a one degree increase in average temperature over the past six days reduces the rate of productivity growth by 3.2 percentage points. This finding suggests that temperature's impact on learning on the job could be one contributing factor to the broader impacts of heat on macroeconomic growth. These effects on productivity growth are economically meaningful and significantly increase the expected harmful effects of higher temperatures on labor productivity as temperatures continue to rise around the globe.

Low-income countries are likely to be particularly hard hit by these effects. They face faster warming and limited climate control, and they also experience high workforce turnover, making training investments both frequent and essential—but vulnerable to heat. For example, (Blattman and Dercon, 2018) find turnover rates of 33% within the first month and 77% per year in garment factories in Ethiopia. Further, many seemingly routine jobs still require significant on-the-job-learning: Adhvaryu et al. (2023) find that line efficiency increases over 40% (17 percentage points) in the first month of producing a new style of garment. As low- and middle-income countries shift toward knowledge-based work, temperature-induced slowdowns in productivity growth may constrain workers' skill acquisition, with effects that compound with increasingly difficult tasks, impeding long-run progress and access to higher-wage jobs.

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.

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.

Burke, Marshall, Solomon M Hsiang, and Edward Miguel, "Global non-linear effect of temperature on economic production," *Nature*, 2015, *527* (7577), 235–239.

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.

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.

Garg, Teevrat, Maulik Jagnani, 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.

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.

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.

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.

Ross, Jennifer M, James L Szalma, and Peter A Hancock, "A meta-analysis of performance under thermal stress," in "Proceedings of the Human Factors and Ergonomics Society Annual Meeting," Vol. 50 SAGE Publications Sage CA: Los Angeles, CA 2006, pp. 1736–1740.

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.

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.

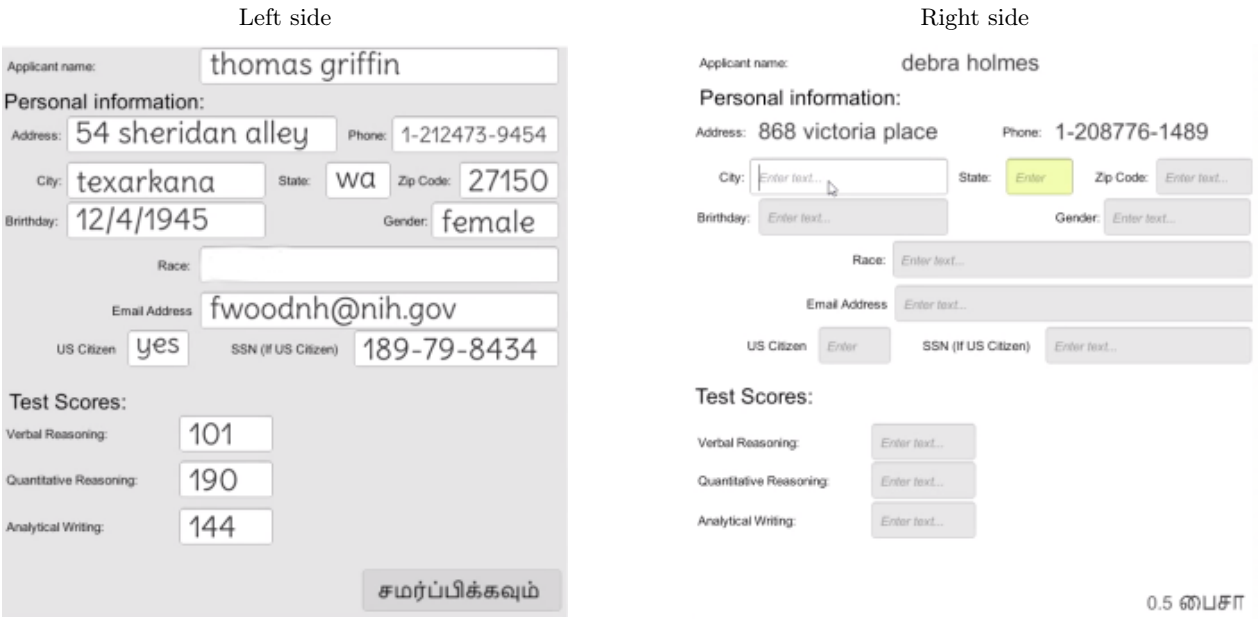
Zhang, Peng, Olivier Deschenes, Kyle Meng, and Junjie Zhang, "Temperature effects on productivity and factor reallocation: Evidence from a half million Chinese manufacturing plants," *Journal of Environmental Economics and Management*, 2018, *88*, 1–17.

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.

A. Appendix

FIGURE A.1: Data-Entry Software



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.

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: Effect of Temperature on Cognition

	Dependent Variable is			
	Cognition Index	**PVT**	**Corsi**	**Hearts and Flowers**
	(1)	(2)	(3)	(4)
Temperature ($°C$)	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. All columns control for worker, day in study, and month-by-year fixed effects. Standard errors are clustered at the worker level. *, **, and *** indicate significance at the 1%, 5%, and 10% levels respectively.

Table A.3: The 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 ($°C$)	-404.613***	-421.342***	-5.810***	0.008**	-8.232***
	(37.270)	(38.927)	(0.511)	(0.003)	(0.898)
Dep. Var. Mean	14402.8	15412.6	230.7	0.827	279.5
Observations	7268	7268	7268	7267	7268

Notes: This table mirrors Table I 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. Standard errors are clustered at the worker level. *, **, and *** indicate significance at the 1%, 5%, and 10% levels respectively.

Table A.4: The Effect of Temperature - Backfilled Temperature 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 ($°C$)	-17.771***	-18.217***	-0.247***	0.004	-0.392***	-0.000	0.004	-0.008	-0.011
	(2.521)	(2.575)	(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. Standard errors are clustered at the worker level. *, **, and *** indicate significance at the 1%, 5%, and 10% levels respectively.

Table A.5: The Effect of Temperature - No Temperature Lag

	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 ($°C$)	-12.403***	-12.778***	-0.157***	0.003	-0.283***	0.002	-0.003	-0.005	-0.003
	(2.353)	(2.406)	(0.028)	(0.003)	(0.056)	(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. Standard errors are clustered at the worker level. *, **, and *** indicate significance at the 1%, 5%, and 10% levels respectively.

Table A.6: The 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 ($°C$)	-20.784***	-20.903***	-0.283***	0.008**	-0.451***	-0.000	0.004	-0.008	-0.011
	(2.830)	(2.923)	(0.038)	(0.003)	(0.077)	(0.002)	(0.004)	(0.007)	(0.008)
PM2.5	-0.048	-0.091	-0.003	-0.000	-0.003	-0.000	-0.000	-0.000	-0.000
	(0.156)	(0.162)	(0.002)	(0.000)	(0.004)	(0.000)	(0.000)	(0.000)	(0.000)
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 the PM 2.5 level. Pollution data is retrieved from the Central Control Board of India. In addition to pollution, all columns control for worker, day in study, and month-by-year fixed effects. Standard errors are clustered at the worker level. *, **, and *** indicate significance at the 1%, 5%, and 10% levels respectively.