When the Earth is Too Hot to Handle: Experimental Evidence on Protecting Manual Outdoor Workers from Extreme Heat *

Aditi Kharb†, Manpreet Singh‡, Akansh§

October 20, 2026

Latest Version

Abstract

Extreme heat threatens the labour supply and earnings of daily-wage workers in developing countries, where physically demanding outdoor work is common. We examine these effects among workers employed under India's National Rural Employment Guarantee Scheme (NREGS), one of the largest public employment programs globally. We combine administrative data on attendance and assessed wages with a discrete choice experiment to elicit worker's preferences over alternative heat adaptation measures. A one-degree increase in mean wet-bulb temperature over the preceding ten days raises absence by 3.0 percent of its mean; absence is 6.7 percent higher in the hottest than in the coolest quartile of accumulated heat. Effects on assessed wages are weaker, although the probability of receiving the scheme's full daily wage declines with increase in temperature during the hottest months. In the choice experiment, workers require compensation equivalent to 2.8 percent of the daily wage to accept a one-hour later start but are willing to forgo 5.8 percent for a 10 percent reduction in the work target, preferring reduced work requirements to changes in work hours.

JEL classification: J22, J24, J48, Q54
Keywords: Heat Stress, Climate Change, Heat adaptation, Outdoor workers, NREGS

*This research was supported by International Growth Centre -India Sustainable Growth Hub Grant Number IND-25292. The study was pre-registered at the Open Science Framework (OSF; y5b2h).
Author contributions: The first author led the research project, including idea conception, grant application, survey questionnaire development, fieldwork, data collection and processing, coding, conducted the statistical analysis, and led the manuscript writing. The second author contributed to idea conception, survey questionnaire development, manuscript writing, and data analysis. The third author provided administrative support and contributed to the development of the survey questionnaire.

†School of Economics, University College Dublin. Email: aditi.kharb@ucdconnect.ie. *Corresponding Author

‡Department of Economic Sciences, IIT Kanpur. Email: manpreets@iitk.ac.in.

§School of Economics, O.P. Jindal Global University

1 Introduction

As climate change pushes temperatures higher, extreme heat is becoming a major constraint on work, with especially severe consequences in low- and middle-income countries. According to International Labour Organization projections, heat stress could reduce global working hours by 2.2% by 2030, equivalent to 80 million full-time jobs, with Southern Asia and Western Africa among the regions expected to experience the largest losses (Kjellstrom et al. 2019). Risks are particularly acute for workers in physically demanding, outdoor occupations, where exposure to extreme temperatures is often unavoidable. Heat can make work more demanding and reduce productivity on the day of exposure, while also impairing recovery between workdays. It can further increase demands at home, particularly for workers with greater household production and caregiving responsibilities. Understanding these effects, and how workers can adapt to rising heat exposure, has therefore become increasingly important as extreme heat events grow more frequent and severe.

This paper address these questions by examining the effects of extreme heat on the labour supply and productivity of manual outdoor workers in India and eliciting their preferences for alternative work arrangements that reduce their exposure. Using detailed administrative data from India's National Rural Employment Guarantee Scheme (NREGS), the country's largest rural public employment programme, we focus on rural workforce engaged in physically demanding outdoor work. Launched in 2005 and implemented from February 2006, NREGS acts as an employer of last resort for unskilled rural workers, guaranteeing each rural household up to 100 days of unskilled manual work per financial year, on demand. It provides a particularly suitable setting for studying these effects because work is predominantly manual and outdoors, contractors are prohibited and labour-displacing machinery is restricted. Furthermore, work timing and conditions are largely determined administratively, which standardises working arrangements across worksites. Workers are paid according to measured work output [1], while attendance and payments are recorded through the programme's

[1] In NREGS the wages are paid based on assessed output therefore we will use assessed wages and assessed

administrative system regularly. The programme is also targeted toward households with limited employment opportunities, making the trade-off between avoiding heat exposure and foregoing earnings particularly salient.

We combine detailed administrative records from the scheme with high-resolution ERA5-Land climate data to examine workers' responses to heat and find that absence responds to accumulated heat. For same-day temperature we find no detectable effect on whether a worker attends, but a one-degree increase in the mean wet-bulb temperature over the preceding ten days raises absence by 3.0 percent of its mean, and absence is 6.7 percent higher in the hottest quartile of accumulated heat than in the coolest. Conditional on work being undertaken, the evidence that heat lowers measured output points in the same direction but is weaker. Within a worksite over March to June, a one-degree increase in mean wet-bulb temperature lowers the probability that a muster roll is assessed at the notified rate by 0.008 (SE 0.004), about one percent of its mean, and the dry-bulb and wet-bulb maximum readings agree. Following the same worker across spells at the same worksite attenuates this by about a fifth without reversing it.

Identifying the effects of heat on labour supply and productivity is challenging because temperature varies systematically across locations and seasons, while workers and worksites differ in their tasks and working conditions. We address these sources of heterogeneity by restricting the analysis to a single state, where workers across the districts in our sample face the same notified wage schedule, administrative system, and implementation rules, and by exploiting the within-location and within-worker variation in heat exposure. The two margins are therefore identified from different variation, absence from day-to-day deviations within a worker and panchayat, and assessed output from differences between spells at the same worksite.

For absence, we estimate a linear probability model with worker, gram panchayat [2],

output interchangeably

[2]A gram panchayat is the lowest tier of elected local government in rural India and is responsible for administering public programmes at the village level, including NREGS.

and calendar-month fixed effects, clustering standard errors at the gram panchayat level. Identification comes from day-to-day deviations in temperature from the usual conditions for a given panchayat and time of year. The fixed effects absorb time-invariant differences across workers and panchayats as well as common seasonal patterns, while we additionally control for rainfall over the same window and show that adding solar radiation and wind speed leaves the estimate unchanged.

For assessed output, identification differs. NREGS uses a group piece-rate system in which worksite output is measured and payments divided among workers present, subject to the statutory daily-wage ceiling. We interpret the recorded rate as an administratively assessed measure of group output per person-day, not individual physical productivity. The unit of analysis is the muster roll, the sheet on which attendance and the assessed rate are recorded for a group of workers at a worksite over a fortnight. This is the finest unit at which the rate varies while permitting worksite fixed effects. Identification compares rolls at the same worksite, holding time-invariant characteristics such as task and terrain fixed, and accounting for calendar year. We focus on March to June, Haryana's four hottest months, and report full-year estimates in the appendix. Within this shorter window, month and temperature move closely together, so we omit month-of-year effects. Over the full year, temperature levels recur across months, allowing these effects to be included, but comparisons then span months that also differ in agricultural activity and the work being done.

The effects of heat on workers have also prompted growing interest in workplace adaptation. Existing evidence considers adaptations such as cooling technologies, air conditioning, ventilation, and other investments that alter the physical work environment (Somanathan et al. 2021). Such measures can be effective, but they may require substantial capital investment and are often poorly suited to outdoor work, where workers cannot easily control ambient temperatures. We instead study organisational adaptations that change when and how work is performed rather than the physical environment in which it takes place. Using a field survey and discrete choice experiment (DCE), we examine workers' experiences with

heat and elicit the value they place on changes in work timing and required output that can reduce exposure to the hottest hours without requiring investment in cooling technology. This allows us to quantify workers' preferences for lower-cost organisational adaptations in a setting where more capital-intensive forms of protection are difficult to implement.

Our results suggest that workers' responses to heat depend in part on how NREGS determines pay. Although the scheme pays according to measured output, the notified daily rate caps earnings, and four-fifths of paid person-days are assessed at that rate. Additional output above the cap does not raise assessed wages, allowing groups to slow down on hot days without necessarily losing pay. Consistent with this, a degree increase in wet-bulb temperature reduces the probability of reaching the notified rate by about one percent of its mean. Following the same worker at the same worksite attenuates this effect by about a fifth but does not eliminate it, suggesting that changes in attendance account for part, but not all, of the association. Attendance responds instead to heat accumulated over preceding days. While the data cannot establish how workers adjust effort or recovery, the contrast suggests that heat costs emerge more through participation than through assessed pay. The attendance response is also larger among Scheduled Caste workers, with absence rising by 0.007 per degree compared with 0.005 for others. This difference is consistent with unequal capacity to sustain attendance under heat.

Our DCE results indicate that workers value reductions in heat exposure, particularly through lower work targets. Workers require about INR 11 a day (2.8 percent of the daily wage) to accept a one-hour later start, but would forgo INR 23 (5.8 percent) for a ten percent reduction in the output target. This distinction matters because the two margins are governed differently within the programme. The value is concentrated in the first reduction, a second reduction of the same size adds nothing, and workers would require payment to accept it, suggesting that larger cuts may have implications for earnings or work duration. Women value a later start more than men (INR 14 versus INR 7), though the difference is not significant after correcting for multiple comparisons. Among women, those spending at least

four hours a day on household work would forgo INR 33 for a lower target, compared with INR 17 among other women; this difference remains significant after correction. Overall, workers place the greatest value on a modest reduction in required output, which does not require a change to the programme's wage schedule.

Our study contributes to the literature studying the effects of heat stress on workers. Existing work documents increases in absenteeism and reductions in on-the-job productivity among manufacturing workers in India (Somanathan et al. 2021), earnings losses among informal workers (Das and Somanathan 2024; Bhattacharya, Kulshreshtha, and Soest 2026), and reductions in workers' reported ability to work (Heyes and Saberian 2022). A related literature shows that nighttime temperatures affect subsequent work capacity through disruptions to sleep and recovery (Garg, Gibson, and Sun 2020; Neidell et al. 2021; Bessone et al. 2021). Our results extend this evidence to an outdoor manual labour setting with extreme economic vulnerability and show that the timing of heat exposure matters for different margins of labour supply and productivity.

Second, we contribute to the literature on how workers and employers adapt to heat across labour-demand and labour-supply margins (Nandwani, Roychowdhury, and Shankar 2026; Graff Zivin and Neidell 2014; Cosaert, Nieto Castro, and Tatsiramos 2023). We focus on workers with limited scope for adaptation and use their preferences over alternative work arrangements to quantify the value they place on reducing heat exposure. This connects the heat-stress literature to a broader literature on workers' valuation of job characteristics, which shows that workers trade wages against hours, schedules, and other working conditions (Mas and Pallais 2017a; Wiswall and Zafar 2018; Maestas, Kathleen J. Mullen, et al. 2023b), including in low-income settings (Feld, Nagy, and Osman 2022). In the context of outdoor manual work, where work timing and physical requirements are typically determined administratively, our estimates capture workers' willingness to trade compensation for changes in these conditions and, conversely, the compensation required to accept greater heat exposure.

Lastly, our results add to the literature on NREGS and rural labour markets in India.

Previous work examines the program's effects on labour-market participation, private-sector wages, migration, and rural labour markets (Imbert and Papp 2015; Imbert and Papp 2020; Muralidharan, Niehaus, and Sukhtankar 2023; Zimmermann 2024), as well as its design, self-targeting, employment guarantee, delivery, leakage, and state capacity (Besley and Coate 1992; Basu, Chau, and Kanbur 2009; Muralidharan, Niehaus, and Sukhtankar 2016; Banerjee et al. 2020; Dutta et al. 2014). We shift attention from what NREGS does to rural labour markets to how working conditions within the program affect the workers employed in it, linking heat exposure to both observed labour-market outcomes and workers' preferred adaptations.

2 Context and Background

Our analysis combines administrative data with original field surveys of workers employed under India's rural employment guarantee programme in the state of Haryana. This section describes the key institutional features of the programme and provides background on Haryana, highlighting why it offers a suitable setting for studying the effects of heat stress on outdoor manual workers.

2.1 Employment guarantee schemes in rural India

2.1.1 National Rural Employment Guarantee Scheme (NREGS)

India's National Rural Employment Guarantee Act (NREGA), enacted in 2005, created a legal entitlement to up to 100 days of paid employment per rural household each year. The programme was designed primarily as an employment and income-support programme for rural households with limited employment opportunities. An adult household member can register with the local administration and obtain a job card, after which workers can demand employment on approved public works. If employment is not provided within the statutory period following a demand for work, the programme provides for an unemployment

allowance. In practice, however, households typically receive substantially fewer than the statutory 100 days of employment. Dutta et al. (2014) document this gap between the legal entitlement and realized employment, with the median household receiving only around ten days of work in their data. NREGS is therefore better understood as a source of intermittent employment rather than a full-year job. This feature is important in our setting because a day not worked may represent income forgone rather than employment that can readily be recovered later.

The programme is also an important institution in rural labour markets. It provides paid work on publicly funded projects while setting a statutory wage, and therefore affects workers' outside options in the private labour market. The existing literature shows that these effects extend beyond programme participants. Imbert and Papp (2015) find that NREGS employment crowded out private-sector employment while increasing rural casual wages, with effects concentrated during the agricultural off-season. Imbert and Papp (2020) show that the programme also reduced seasonal rural-to-urban migration, while Muralidharan, Niehaus, and Sukhtankar (2023) find that improvements in programme implementation increased beneficiary earnings and reduced poverty, with substantial effects operating through non-program earnings and local labour-market equilibrium. Thus, NREGS is not simply a transfer programme but a large-scale employer whose presence changes the set of opportunities available to rural workers.

The programme is explicitly targeted toward households for whom this employment option is valuable. It is self-targeting through the relatively low statutory wage and the physical nature of the work, so participation is disproportionately concentrated among poorer rural households and workers with limited alternatives. Women also constitute a large share of participants. These features make NREGS particularly relevant for studying heat exposure among workers who may have limited access to occupations offering air conditioning, flexible schedules, or other forms of workplace protection.

NREGS employment is organized around labour-intensive public works intended to create

or maintain rural public assets. Typical activities include excavation and pond-related work, land development, afforestation, water-conservation projects, and rural infrastructure. The programme is designed around labour-intensive production, mechanization is restricted and requires administrative approval where manual work is not feasible or where there is a binding time constraint. Workers therefore perform physically demanding tasks outdoors, often for several hours during the hottest part of the day. This feature distinguishes NREGS from much of the existing evidence on heat and labour, which often comes from manufacturing or other settings in which firms can partially insulate workers from ambient temperature through indoor production, capital, or cooling technologies.

Employment is organized at the worksite level. Once a worker has a job card, the local implementing agency allocates workers to an approved worksite and places them on a muster roll for the relevant period. A locally appointed semi-skilled worker, the *Mate*, records attendance, coordinates workers at the site, and allocates tasks. Attendance is now recorded digitally through the National Mobile Monitoring System (NMMS), and the resulting e-muster rolls contain worker-level information on attendance, worksite, workcode, and wages. Projects can continue across multiple fortnights, with a common workcode allowing the same project to be followed across successive rolls. These administrative records therefore provide unusually granular observations of who is present at a worksite on a given day.

The production technology also makes the programme particularly useful for measuring productivity. A technical functionary measures the quantity of work completed using a publicly specified Schedule of Rates. The relevant physical unit depends on the activity—for example, cubic metres of earth excavated, metres of road levelled, or saplings planted. The measured quantity is converted into a total payment using the applicable piece rate. Tasks are designed so that completing the prescribed quantity corresponds to the statutory daily wage. Consequently, variation in the amount of work completed, holding the composition of workers at the site fixed, provides a direct measure of assessed output.

An important feature of the wage-setting system is that output is assessed for the group

of workers present at a worksite rather than separately for each worker. The total quantity completed is divided by the number of workers present, and each worker receives the same assessed daily rate based on the group average. For example, if a group completes 80 percent of the prescribed quantity, each worker is assessed as having completed 80 percent of the individual target and receives 80 percent of the statutory daily rate. We therefore interpret the resulting measure as group-level assessed productivity, rather than individual physical productivity. This institutional feature is central to our analysis because it allows us to distinguish two margins of adjustment, whether a worker is present at the worksite, and how much output the group of workers who are present produces.

2.1.2 From NREGS to VB-G RAM G

In July 2026, the Viksit Bharat—Guarantee for Rozgar and Ajeevika Mission (Gramin) (VB-G RAM G) replaced NREGS. The transition largely preserves the institutional setting relevant to our study, rural workers continue to undertake labour-intensive public works through locally administered worksites, with attendance and work output recorded through the programme's administrative system. The new Act nevertheless changes several features of the programme, most notably increasing the statutory employment guarantee from 100 to 125 days and revising the categories, priorities, and planning of permissible works.

This transition is important for interpreting our data. Our administrative sample through June 2026 is drawn entirely from NREGS, while our survey fieldwork after July 1, 2026 takes place under VB-G RAM G. We therefore use the administrative data to characterize heat-related attendance and assessed productivity under NREGS, and the post-transition survey to examine workers' heat exposure, health responses, and preferences over adaptations to the timing and conditions of work. The continuity in the underlying public-employment system allows these sources to speak to complementary dimensions of workers' responses to heat despite the change in programme rules.

2.2 Background on Haryana

Haryana provides a useful setting for studying heat stress among rural public works workers because it combines substantial and recurrent heat exposure with a large and established NREGS programme. The state lies in the hot northwestern plains of India, and high temperatures are a regular feature of the pre-monsoon months during which public works are undertaken. For example, the India Meteorological Department recorded maximum temperatures above 40°C at Haryana stations during May 2026 (India Meteorological Department 2026a). More importantly, exposure in this setting is not confined to exceptional heat-wave days, workers perform physically demanding outdoor work during a period when high daytime temperatures are a regular feature of the working environment. Haryana therefore provides an opportunity to study heat as a recurring labour-market constraint rather than solely as an infrequent extreme event.

The programme also operates at substantial scale in the state. Haryana has implemented NREGS since 2008, and in 2024–25 generated 11.79 million person-days of NREGS employment, with approximately 0.479 million individuals working under the programme (Rural Development Department, Government of Haryana 2026; Ministry of Rural Development, Government of India 2026b). Employment nevertheless remains far below the statutory maximum for most households. A study of Haryana finds that only a small fraction of households completed the full 100 days of employment between 2016–17 and 2020–21 (Press Information Bureau, Government of India 2022). NREGS in Haryana is therefore an important but intermittent source of rural employment, with workers participating in the programme alongside other employment opportunities in the rural economy. This is relevant for interpreting workers' responses to heat because NREGS operates alongside the broader labour-market opportunities described above, rather than in isolation.

The composition of the workforce is also relevant for our study. Women represented 59.5% of NREGS person-days in Haryana in 2022–23, well above the statutory minimum of one-third (Press Information Bureau, Government of India 2023). NREGS workers in

Haryana also have relatively low levels of formal education, and Scheduled Caste workers constitute a large share of the workforce. The population therefore includes many workers who may have limited access to occupations with air conditioning, flexible schedules, or other forms of workplace protection. Women may additionally face heat-related constraints outside the worksite, including domestic and care responsibilities. These characteristics make the Haryana workforce particularly relevant for examining how heat affects both labour supply and workers' choices over workplace conditions. Our survey data allow us to examine these additional margins directly.

Haryana also has an established administrative system through which NREGS employment is implemented and recorded. The programme operates through Gram Panchayats, worksites, muster rolls, and digitally recorded payments and attendance, providing the administrative infrastructure required to construct detailed measures of workers' attendance and assessed output (Rural Development Department, Government of Haryana 2026; Ministry of Rural Development, Government of India 2026a). In addition, Haryana prohibits contractors and labour-displacing machinery in programme works, reinforcing the labour-intensive nature of the programme (Rural Development Department, Government of Haryana 2026). These state-specific features strengthen the link between ambient heat exposure and the physically demanding outdoor work performed by NREGS workers.

Taken together, Haryana combines substantial exposure to recurrent heat with a large population of workers engaged in physically demanding outdoor public employment and administrative records that allow us to observe both attendance and assessed output. This combination provides a setting in which the labour-supply and productivity margins of heat exposure can be studied jointly.

3 Absenteeism and Assessed Wages

In this section we examine how heat affects workers' attendance and assessed output through wages under the programme's existing work arrangements. Using administrative data, we assess whether heat-related costs emerge through workers' decisions to attend work, the output their groups are assessed to produce, or both. This provides evidence on how heat affects workers under administratively determined work arrangements.

3.1 Data

3.1.1 Weather Data and Measurement of Heat Stress

We use weather data from Copernicus project's ERA-5 database. The data includes hourly air temperature, surface temperature, dew point, evaporation, surface pressure, wind speed, and leaf area index. The use of hourly reanalysis data is important for our application because heat exposure at worksites is likely to vary substantially within a day. Daily mean temperature, for example, can obscure the high-temperature and high-humidity conditions experienced by workers during peak working hours, and hourly data enables a thorough analysis to compute various daily aggregates of temperature. While the climate data goes back to 2005, the year of NREGS launch, we restrict our analysis to data after 2021, due to the non-availability of digitised NREGS data prior to that. The data is provided at 0.1°x 0.1°grids, which covers an area of roughly 108 sq km. We use all the data enclosed between 27°and 31°N and 74°and 78°E. These latitudes and longitudes enclose the state of Haryana, and a little bit of the neighbouring states.

Figures 1 and A.1a-A.1d summarize spatial and temporal variation in summer temperatures across Haryana since 2015. Average temperatures show a clear spatial gradient, with higher temperatures in the western and southern districts, and an upward trend over time. The particularly high temperatures observed in 2022 and 2024 indicate that this trend is accompanied by substantial year-to-year variation.

Temperature variability shows no clear time trend, although it tends to be higher in the warmer western districts. The years with higher average temperatures also tend to have more days above 35°, while the duration of hot periods varies considerably across years. Overall, the figures point to both a gradual increase in summer temperatures and substantial spatial and temporal heterogeneity in heat exposure.

Together these figures show us that there is a substantial spatial and temporal variation in summer temperatures in the state of Haryana. This facilitates our study, as the state is otherwise homogenous, relative to other parts of India. We can also notice that a higher than expected heat in 2022 and 2024, which provides us enough episodes of unexpectedly hot weather. In addition to air temperature, we construct wet-bulb temperature following (Somanathan et al. 2021), using air temperature and vapour pressure to account for humidity. We calculate wet-bulb temperature at the hourly level and aggregate it to daily means and maxima. We also use solar radiation and wind speed as separate controls in robustness analyses.

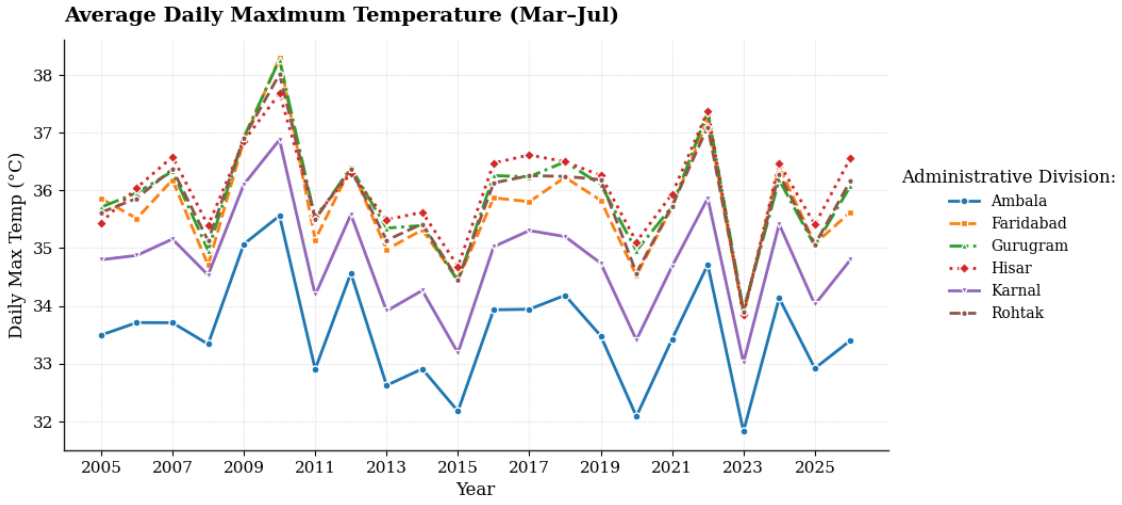


Figure 1: Average Daily Maximum Temperature

Note: Average daily maximum temperatures across the five administrative divisions of Haryana from March to July, 2005–2026. Data source: ERA-5 database.

3.1.2 Administrative data

We use the complete e-muster roll (worker's lists) and wage data for financial years 2022-23 through 2026-27 from state of Haryana. This public dataset records each worker's attendance at the worksite over the fortnight in which they were employed under NREGS. It also reports the worker's job card number, their caste category, their relation to the head of household, the wages paid to them, and the wages owed to them given their attendance and the work completion percentage. The roll is filled automatically at the time attendance is marked, through the National Mobile Monitoring System (NMMS) app. Each roll is mapped to a worksite and carries the fortnight for which it was used along with the workcode. Same worker can appear on different muster-rolls if employed over different fortnights and projects. Projects that run beyond a single fortnight appear as separate rolls, with workers on the new roll employed at the same site, and the shared workcode allows us to link these rolls to one another.

We construct our administrative sample from NREGS muster rolls covering two randomly selected blocks in each district of Haryana. Blocks were selected using an automated random sampling procedure applied during the data-scraping process. Overall, in our sample we have 93,398 roll observations, 8,358 works and 1,936,710 person-days. Workers on muster-rolls with the same work code are classified in the same group.

Each muster-roll records an average of 6 workers and 56.9 person-days. The average number of work days per person per work is 9.5, of which they are present on 6.8 days. There are roughly 4.4 rolls per group with 26.3 members per group. The raw data has an average absence rate of 0.283 over all the person-days, with a standard deviation of 0.45. While workers are assured 100 days of employment, we find only 16.4 days of work per financial year on average.

The wages are assessed at the group level in Haryana. If the group completes the project, everyone is paid the scheduled rate. If there is a shortfall, the reduction in total wages is shared equally by all the workers. In the admin data, we notice that the average wages

paid are INR 347.9, as opposed to the scheduled rate of INR 366.1, averaged over all the years. While the average shortfall seems small, roughly 20% of the groups have received lower wages. Moreover, paid wages have substantially larger variation than the notified wage rates, resulting in fatter distribution tails as evidenced by the 10^{th} percentile.

We combine ERA 5 data to the administrative data to broadly understand the weather patterns on actual work days, as opposed to all the days.

3.2 Methodology

In this section we describe the empirical strategies used to examine the effects of extreme heat on workers' labour supply and wages.

3.2.1 Absenteeism

Attendance is recorded at the person-day level through the National Mobile Monitoring Service Application (NMMS), with the Mate (site-manager) marking each listed worker present or absent. We define absence at that level. Most workers are employed under the programme for only a small number of days in a year, projects do not generally run consecutively, and some workers appear at more than one worksite within a financial year, so we do not observe participation decisions on most days of the calendar. Figures A.2 and A.3 show the distribution of days each worker appears on a muster roll[3] and is present at a worksite. The panel is therefore sparse, with workers as the cross-sectional unit and days as the time dimension.

The sparsity of the panel does not, in itself, reflect workers' endogenous participation decisions. It arises from the institutional design of NREGS, projects are approved and scheduled at the block level, and workers are observed in the administrative records only while relevant work is ongoing. A worker's absence from the records on days when no relevant work is available therefore cannot be interpreted as a decision not to work, it is

[3]A muster roll is an official attendance and employment record maintained under the employment programme, documenting workers registered or recorded for work over a given period.

simply that no NREGS work is available. Our estimates thus capture workers' attendance decisions conditional on NREGS work being available and the worker being listed on an active muster-roll. We consequently define the outcome only for worker-days covered by an active muster-roll.

Following Somanathan et al. (2021), we measure heat exposure with the wet bulb temperature

$$\mathrm{WBT} = 0.567\,TA + 0.216\,\rho + 3.38, \tag{1}$$

in the constants they state, where TA is air temperature in degrees Celsius and ρ the vapour pressure in hectopascals, read directly off the ERA5-Land dew point. For worker-level outcomes Somanathan et al. (2021) evaluate (1) at the daily mean rather than the daily maximum, and we do the same. The rationale for a measure that combines temperature with humidity is that the constraint on manual work in this setting is evaporative, a worker sheds heat by sweating, and the rate at which sweat can evaporate depends on how far the air is from saturation rather than on its temperature alone. This is a measure of the environment and not of the worker.

Air temperature and dew point are aggregated from hourly ERA5-Land values over local days in Indian Standard Time and assigned to each gram panchayat by its centroid's nearest grid cell, with every centroid checked to lie within half a grid box of the cell it is assigned to so that no panchayat is silently given a neighbouring cell's weather.

Somanathan et al. (2021) further report that absence, unlike output, responds to accumulated rather than contemporaneous exposure, and their headline regressor is a ten-day average; a related literature on sleep and overnight recovery points the same way obradovich2017nighttime, GargGibsonSun2020, Bessone2021. We follow this literature in measuring exposure as an average over a window of preceding days. For gram panchayat g

on day t and a window of length W,

$$\overline{\mathrm{WBT}}_{g,t}^{(W)} = \frac{1}{|\mathcal{D}_{g,t}^{(W)}|} \sum_{s \in \mathcal{D}_{g,t}^{(W)}} \mathrm{WBT}_{g,s}, \qquad \mathcal{D}_{g,t}^{(W)} = \{t-W, \ldots, t-1\}, \tag{2}$$

and we set $W = 10$ throughout the main results.

Three features of (2) are choices rather than conventions. First, the window excludes day t itself, so that a contemporaneous and an accumulated term can enter the same regression without overlapping and what they share is only the persistence of the weather.

Second, a trailing mean is only defined if the preceding days exist in the data, and the days preceding a work-day are mostly days on which the panchayat had no roll open. Exposure is therefore computed for every panchayat on every calendar day of the season and merged onto the attendance records afterwards, rather than being built from the work-days themselves: a window assembled from work-days would run over the days a panchayat chose to work, and that choice is partly a response to the very weather the window is meant to measure.

Third, a window is treated as missing unless the reanalysis supplies at least seventy per cent of its days — three of three, four of five, seven of ten, eleven of fourteen, fourteen of twenty. The requirement is proportional rather than a fixed number of days, because a fixed count strict enough to be meaningful at three days cannot be met at all, while one loose enough to be met at three days would accept barely a third of a twenty-day window.

Let A_{igt} equal one if the register marks worker i absent on date t, and zero if it marks her present. We measure heat exposure in gram panchayat g on date t using the 10-day moving average of wet-bulb temperature, $X_{g,t} = \overline{\mathrm{WBT}}_{g,t}^{(10)}$. We estimate the following linear probability model by ordinary least squares using person-day observations from March through June:

$$A_{igt} = \sum_{q=2}^{4} \beta_q D_{qgt} + \theta R_{g,t} + \alpha_i + \gamma_{g,m} + \delta_y + \lambda_d + \varepsilon_{igt}, \tag{3}$$

where $D_{qgt} = \mathbf{1}\{X_{g,t} \in Q_q(g)\}$ indicates whether heat exposure falls in quartile q of the gram

panchayat's own March–June distribution. The coolest quartile is omitted as the reference category. Thus, each coefficient β_q captures the difference in absence probability between quartile q and the coolest quartile, conditional on the included controls and fixed effects. Alongside the bins we report the compact form of (3) in which they are replaced by the level of $X_{g,t}$, giving a per-degree gradient. $R_{g,t}$ is rainfall standardised within gram panchayat and month of the year and averaged over the same window as the heat term. Each specification carries exactly one rainfall term, at the horizon of its own heat term. Section 3.3.4 examines robustness to alternative measures of the same weather variable and alternative exposure-window lengths W, leaving the remaining specification unchanged.

The fixed effects are the worker α_i, a gram panchayat by month-of-year interaction $\gamma_{g,m}$, a financial year δ_y and a day of week λ_d. The interaction compares a day with the usual conditions at that panchayat in that month rather than at that time of year across the state, so identification comes from deviations of accumulated exposure from the local seasonal norm. Since a job card is registered in a single panchayat, α_i already absorbs $\gamma_{g,m}$ for every worker who does not move between panchayats. Table 1 reports a ladder running from the fixed-effect to the full specification in (3), so that the contribution of each layer is visible. Standard errors are clustered on the gram panchayat level. Alternative clustering on the worker and two-way on gram panchayat and date is tested in the robustness section.

The estimation sample is person-days on an open muster roll, inside the roll's active span, excluding Sundays, and restricted to March through June. Two restrictions matter. A muster roll is issued for a fixed administrative period, but the underlying work may begin after the roll is printed or finish before its nominal end date; outside the active span every listed worker is marked absent, and those marks record the inactivity of the roll rather than individual decisions. We therefore restrict the sample to the interval between the first and last day on which at least one worker is recorded present. Second, we drop work-days on which nobody at the work was marked present. On such a day every name on every sheet at that work carries an absence mark, which records a stoppage rather than an individual

decision, and keeping those days would put an extensive margin into an outcome meant to measure the intensive one. We test the robustness of our results by relaxing these restrictions in the Robustness section.

3.2.2 Assessed Wage

Work under the employment guarantee is paid by measured output. A group is assigned a task, its output is measured after an interval, and the daily wage recorded on the roll equals that output valued at the state schedule of rates, subject to the notified daily rate as a ceiling. We therefore express each wage as a ratio to the notified rate in its financial year, removing the discrete changes in that rate between years; financial-year effects account for less than one per cent of the remaining variation. We examine three outcomes: an indicator for whether the roll reaches the notified rate, the ratio itself, and the shortfall in rupees per person-day. Because four-fifths of paid person-days sit exactly at the ceiling, the indicator is the margin on which movement is most likely to register, and the two continuous outcomes are reported alongside it rather than in place of it.

The unit of observation is the muster roll. Everyone named on a roll receives the same assessed rate, and 98.5 per cent of rolls record a single wage, so there is no within-roll variation to exploit and the person-day is not a usable unit. At the other extreme a work spans an average of 1.18 calendar years, so work-year observations are nearly indistinguishable from worksites and a worksite fixed effect cannot be estimated at that level. Each work carries approximately thirteen rolls, and 98.9 per cent of rolls belong to works with at least one other roll. The roll is thus the finest unit at which the outcome varies and the coarsest at which worksite fixed effects can be included. A variance decomposition confirms where the variation lies: the total standard deviation of the assessed rate is 0.104, of which 0.083 is across works, 0.055 across days within a work, and 0.029 across people on the same day at the same work.

Worksite fixed effects are central to the identification strategy. Comparing rolls from

the same work compares different spells on the same soil, terrain and task, under the same supervisor and largely among the same workers, absorbing time-invariant site characteristics that may affect measured output, including ground conditions, distance to water, supervisory practices and the local application of the schedule of rates. For roll r at worksite j, in calendar year y and month m, we estimate

$$Y_{jr} = \beta S_{jr} + \gamma R_{jr} + \delta' X_{jr} + \mu_j + \tau_y + \lambda_m + \varepsilon_{jr}, \tag{4}$$

where Y_{jr} denotes the outcome and S_{jr} the heat measure, averaged over the roll's duration. In the linear specification, S_{jr} enters continuously; in the temperature-band specification, it is replaced by indicators for temperature bands, with the lowest band omitted. The term R_{jr} denotes rainfall averaged over the roll's duration and standardised within gram panchayat and month of year. X_{jr} contains the logarithms of the number of days worked and the number of workers, which hold constant the opportunity to meet a hot day. All specifications include worksite and calendar-year fixed effects; half also include month-of-year fixed effects. Standard errors are clustered at the gram panchayat level.

We measure heat in three ways, averaging each over the days covered by the roll: the daily maximum dry-bulb temperature, and the wet-bulb temperature of Somanathan et al. (2021) at both its daily mean and its daily maximum. The three are reported together because they capture different aspects of thermal exposure, and because agreement among them is informative in a way that any one of them alone is not. Estimates that differ across the three should not be read as interchangeable measurements of a single heat effect.

Our estimates are drawn from March to June, the four hottest months in Haryana, and this choice governs what can be held fixed alongside the temperature. The window retains rolls with at least one day inside it and does not truncate them, so a roll straddling a boundary keeps all of its days and its own weather. Within four months, calendar time and temperature move together, and a month-of-year effect and the heat measure compete for the

same variation; we therefore do not impose month effects in the main specification and return below to what that costs. We report the full calendar year in the appendix. Over twelve months each temperature level is reached twice, once while the year warms and once while it cools, which makes a month effect affordable, but it also means that a comparison within a worksite is drawn across months that differ in the work being done, the tasks available and the agricultural calendar, and not only in their weather. The two windows therefore answer different questions, and we take the hot-season estimate as the paper's result.

Because a worksite fixed effect compares rolls at the same site rather than the same workers, it leaves open the possibility that hot spells draw a different set of people onto a site. We therefore also estimate the specification with a worker-by-worksite fixed effect, so that the comparison is the same person at the same site across a hotter and a cooler spell. This requires a row for each worker named on a roll, but the roll remains the unit of analysis: the outcome, the temperature and the controls are the roll's own values, and each roll counts once. A cell in which a worker appears at a site on only one roll is absorbed by its own effect, so the estimate rests on workers who returned to the same site across more than one spell.

Three further specifications test the shape and the sources of the response. The first replaces the linear heat term with the shares of a roll's days falling in each of a set of temperature bands, allowing the response to bend, and reports the joint test that all band coefficients are zero rather than relying on individual significance across many coefficients. The second adds solar radiation and wind speed, averaged over the roll's duration and standardised as rainfall is, since radiation adds radiant load and wind removes it by convection; because these are unavailable for part of the sample, the baseline is first re-estimated on the restricted sample so that the two rows are read against each other. The third examines heterogeneity by type of work. Work type does not vary within a worksite, so its main effect is absorbed; differences are therefore assessed from a pooled regression in which the heat measure is interacted with work-type indicators, and not by comparing separate subgroup regressions that share gram panchayats.

3.3 Results

Using the administrative data, we estimate the effects of heat on two margins of worker behaviour, whether a worker attends work at a NREGS worksite and the productivity represented by assessed wages conditional on attendance.

3.3.1 Absenteeism

Table 1 presents our main results on absenteeism. We estimate equation (3) on the March–June sample, which corresponds to the period of late spring and extreme summer heat and is the relevant window for our productivity analysis.

On 1,936,710 person-days contributed by 114,218 workers across 1,168 gram panchayats, we find that accumulated heat is associated with a substantial increase in absence. In our preferred specification, a worker whose preceding ten days fall in the hottest quartile of the wet bulb distribution for her own gram panchayat is 1.27 percentage points more likely to be absent than a worker in the coolest quartile (standard error 0.0055), against a mean absence rate of 0.189. This is a movement of approximately 6.7 percent relative to the mean. Entered in levels, a one-degree increase in the ten-day mean wet bulb temperature raises the probability of absence by 0.0057 (standard error 0.0016), or about three percent of the mean rate per degree.

The response is graded rather than concentrated at the top of the distribution. The second quartile is indistinguishable from the coolest, at 0.0003 (standard error 0.0043), the third raises absence by 0.68 percentage points, and the hottest by 1.27. A dose-response pattern of this shape is what we would expect of a constraint that binds progressively as conditions worsen, and it is not a pattern that a single outlying bin could produce.

The estimate also strengthens as the seasonal controls tighten. With worker, month-of-year and day-of-week effects alone the top-quartile coefficient is 0.0089; adding the gram panchayat moves it to 0.0091, adding the matched ten-day rainfall control to 0.0108, and replacing the common month effect with a gram panchayat by month-of-year interaction

and a financial-year effect to 0.0127. The per-degree estimate roughly doubles across the same four columns, from 0.0029 to 0.0057. This direction matters for interpretation. If the relationship were an artefact of seasonal or spatial confounding, then absorbing more of that variation would attenuate it; instead the estimate grows, which is what we would expect if the coarser specifications were comparing hot places and hot months with cool ones and thereby averaging the within-panchayat response toward zero. The rainfall term enters positively throughout, so the heat estimate is not standing in for the disruption that rain causes to earthwork.

These results are consistent in direction with the evidence for time-rate workers. Somanathan et al. (2021) find that a one-degree increase in ten-day average temperature raises absence among Indian garment workers by approximately ten percent, while finding no attendance response among cloth weavers paid by piece rate. Our estimate of about three percent of the mean per degree is toward the lower end of that range. The comparison with their piece-rate workers is the more instructive one, because NREGS compensates workers according to measured output and so gives a worker a direct reason to attend even when conditions are hard. We nevertheless find a clear attendance response. Whatever reduces attendance in this setting is therefore strong enough to operate against the incentive that the payment system creates, rather than in the absence of one.

These effects can be expressed in monetary terms. An absence represents a forgone day's assessed wage, and workers in our sample are listed on an active roll for about seventeen days between March and June. At the 2026–27 notified rate of INR 409, a one-degree increase in ten-day mean wet-bulb temperature costs about a tenth of a day and INR 39 in program earnings over the window. Moving from the coolest to the hottest quartile of accumulated heat costs about a fifth of a day and INR 88. These losses are modest because participation is brief, not because the attendance response is weak. They measure forgone program earnings, not total income, the data do not show whether absent days are spent on other paid work, recovery, or household work.

Table 1: Absence and accumulated heat: the 10-day mean wet bulb, March to June

	(1) Worker, month, day of week	(2) + gram panchayat	(3) + ten-day rainfall	(4) Panchayat × month, year
A. The 10-day mean wet bulb, per degree				
Wet bulb, 10-day mean (C)	0.0029**	0.0034***	0.0042***	0.0057***
	[0.0012]	[0.0013]	[0.0012]	[0.0016]
B. Quartiles, cut within the gram panchayat				
Q2	-0.0041	-0.0035	-0.0017	0.0003
	[0.0043]	[0.0044]	[0.0043]	[0.0043]
Q3	0.0022	0.0025	0.0050	0.0068
	[0.0053]	[0.0054]	[0.0053]	[0.0053]
Q4 hottest	0.0089*	0.0091*	0.0108**	0.0127**
	[0.0051]	[0.0054]	[0.0053]	[0.0055]
Worker	Yes	Yes	Yes	Yes
Gram panchayat	No	Yes	Yes	No
Month of year	Yes	Yes	Yes	No
Gram panchayat × month	No	No	No	Yes
Financial year	No	No	No	Yes
Day of week	Yes	Yes	Yes	Yes
Rainfall at the same horizon	No	No	Yes	Yes
Gram panchayat clusters	1,168	1,168	1,168	1,168
Workers	114,218	114,218	114,218	114,218
Observations (person-days)	1,936,710	1,936,710	1,936,710	1,936,710
Mean of the dependent variable	0.1891	0.1891	0.1891	0.1891

Linear probability model of daily absence, estimated on person-days in March–June. The sample is restricted to open muster rolls, excludes Sundays and days with no workers present, and is held constant across all columns and absence tables.

Heat is the 10-day lagged mean wet-bulb temperature, calculated using ERA5-Land data at the gram panchayat centroid following Somanathan et al. (2021). The window excludes the current day.

Panel A enters temperature linearly; Panel B uses within-panchayat seasonal quartiles, with the lowest quartile omitted.

Rainfall controls are measured at the corresponding heat horizon and standardised within gram panchayat and month of year.

Column (1) follows the fixed-effects specification of Somanathan et al. (2021); column (2) adds gram panchayat fixed effects; column (4) includes gram panchayat-by-month-of-year fixed effects.

Standard errors clustered by gram panchayat are in brackets. Alternative clustering by worker and by gram panchayat and date leaves point estimates unchanged. Signs are reversed relative to Somanathan et al., who model presence. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

3.3.2 Assessed Wage

Table 2 reports the estimates for March to June. Within a worksite, conditional on calendar year and controlling for rainfall over the duration of each muster roll, a one-degree increase in mean wet-bulb temperature reduces the probability that a roll is assessed at the notified rate by 0.0078 (SE 0.0038). Relative to the sample mean of 0.749 this is a decline of about one percent per degree. The daily maximum readings agree in sign and magnitude, at −0.0106 (SE 0.0043) for the wet bulb and −0.0075 (SE 0.0027) for the dry bulb, both significant at the five percent level. The negative association is present in all three measures of heat exposure.

Column (2) holds the worker fixed alongside the worksite. The coefficient on the mean wet bulb is −0.0061 (SE 0.0044), on the maximum wet bulb −0.0088 (SE 0.0052) and on the maximum dry bulb −0.0061 (SE 0.0033), the latter two significant at ten per cent. All three are about a fifth smaller than the corresponding roll-level estimates, with standard errors wider in similar proportion, which is what holding the worker fixed costs: it leaves 23.7 per cent of the raw wet-bulb standard deviation against 35.6 per cent for the worksite alone. Changes in the composition of workers attending therefore account for part of the association but do not generate it [4].

The worker specification rests on a narrower set of comparisons than the roll specification. The estimate is identified from the 30,993 worker-worksite cells holding two or more rolls, and within Panel A from the 4,874 of those in which the same worker reached the notified rate on one roll and not on another. These cells are not a random slice of the sample: they sit at a mean wet-bulb temperature of 22.0 against 22.6 elsewhere, and the workers in them are present for 19.1 days against 7.0. Section 3.3.1 shows that attendance declines as heat accumulates. Consequently, workers observed at the same site across hot and cool periods

[4] We do not interpret the specification with month-of-year fixed effects within the March–June window as a decisive test. Across the full calendar year, month effects absorb all but 13.8 per cent of the wet-bulb variation remaining within worksites; within a four-month window temperature and month of year are more closely aligned still, leaving little independent variation with which to estimate their separate effects. That specification is reported in Appendix Table A.3 and is not treated as a test of the March–June result.

are disproportionately those who continue to attend despite the heat. Column (2) therefore identifies the relationship between heat and assessed output among continuing attendees, but does not capture responses operating through absence.

Panels B and C show no comparably clear relationship for the continuous outcomes. At the roll the coefficient is 0.0004 (SE 0.0014) for the assessed rate relative to the notified rate and −0.07 (SE 0.51) for the rupee shortfall per person-day, against means of 0.961 and 14.75; the worker column is likewise indistinguishable from zero on both. Since three-quarters of March–June rolls are assessed at the notified rate, the continuous outcomes have limited scope to capture changes in assessed wages, and the probability of reaching the notified rate reported in Panel A is the clearest margin in this sample.

Over the full calendar year the corresponding coefficient on mean wet-bulb temperature is 0.0015 (SE 0.0014; Appendix Table A.3), close to zero and statistically insignificant. The March–June and full-year point estimates differ by 0.0093. The contrast admits at least two interpretations. One is that heat affects assessed wages primarily during the hottest months, with cooler-season observations diluting the relationship in the full-year sample. The other is seasonal confounding, within March–June, temperature trends may coincide with changes in work progress or task composition that are not fully captured by the controls. Extending the sample across the full year provides more independent variation between temperature and season, but does not establish which interpretation is correct. We therefore report the March–June association while recognising that the available data cannot distinguish a genuinely seasonal heat effect from residual seasonal confounding.

Table 2: Heat and the Assessed Wage: March–June Estimates

	(1) Roll-level	(2) Worker-roll
A. Reaching the notified rate		
Wet bulb, daily mean	-0.0078**	-0.0061
	[0.0038]	[0.0044]
Wet bulb, daily maximum	-0.0106**	-0.0088*
	[0.0043]	[0.0052]
Dry bulb, daily maximum	-0.0075***	-0.0061*
	[0.0027]	[0.0033]
Mean of the outcome	0.7493	0.7493
B. Assessed rate over the notified rate		
Wet bulb, daily mean	0.0004	0.0015
	[0.0014]	[0.0015]
Wet bulb, daily maximum	0.0001	0.0018
	[0.0015]	[0.0018]
Dry bulb, daily maximum	-0.0002	0.0010
	[0.0009]	[0.0012]
Mean of the outcome	0.9609	0.9609
C. Shortfall in rupees per person-day		
Wet bulb, daily mean	-0.07	-0.52
	[0.51]	[0.56]
Wet bulb, daily maximum	0.03	-0.58
	[0.56]	[0.67]
Dry bulb, daily maximum	0.16	-0.30
	[0.33]	[0.45]
Mean of the outcome	14.75	14.75
Worksite effect	Yes	No
Worker × worksite effect	No	Yes
Month-of-year effect	No	No
Unit	Roll	Worker-roll
Observations	39,734	74,369
Worker-worksite cells with two or more rolls	–	30,993
of which identify the notified rate	–	4,874
Gram panchayats	1,141	1,141
Rainfall and size controls	Yes	Yes
Wet-bulb variation surviving (%)	35.6	23.7

Notes: Both columns use the March–June sample and include calendar-year fixed effects, rainfall and work-size controls, but no month-of-year fixed effects. Column (1) is at the roll level with worksite fixed effects; column (2) is at the worker-roll level with worker–worksite fixed effects. Each roll receives unit weight, divided equally across its worker-roll observations. Outcomes and heat exposure are measured at the roll level, with temperature averaged over all days the roll ran. Rolls must have at least one day in March–June. Worker–worksite cells count cells with two or more rolls; the next row reports those identifying the notified-rate outcome. The final row reports the share of raw wet-bulb variation remaining after fixed effects. Standard errors clustered at the gram panchayat level are in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

3.3.3 Heterogeneity

Absenteeism: Table A.4 estimates the specification separately by the caste category recorded on the muster roll. The register distinguishes Scheduled Caste workers from all others and makes no finer distinction, so the comparison group pools Other Backward Class and general-category workers and the table cannot speak to them separately. The two groups divide the sample almost evenly, 951,316 person-days against 985,394, and they differ sharply in the outcome: Scheduled Caste workers are absent on 21.9 per cent of person-days against 16.0 per cent for others, a gap of more than a third. Their response to accumulated heat differs much less. A degree of the ten-day mean wet bulb raises absence by 0.0070 (0.0021) among Scheduled Caste workers and 0.0046 (0.0015) among others, both significant at one per cent; against their respective means these are 3.2 and 2.9 per cent, and the pooled interaction that tests the difference formally returns 0.0005 (0.0006). The quartile estimates tell the same story with less precision, 0.0163 (0.0068) in the hottest quartile for Scheduled Caste workers against 0.0092 (0.0056) for others. What differs by caste is therefore the level of absence, not its sensitivity to heat: accumulated exposure costs the two groups a similar proportion of their attendance, but it does so on top of a baseline that is already much higher for Scheduled Caste workers, so the same proportional response translates into more days lost.

Assessed Wage: Table A.5 examines whether the relationship between heat and assessed wages varies by type of work. We distinguish three groups with at least 500 March–June rolls and 30 works: canals, drains and ponds (30,405 rolls); roads and land levelling (8,548); and plantation (669). Together these account for 99.5 per cent of the sample. The distinction is motivated by differences in exposure, roadwork and land levelling typically involve digging in open ground, with limited shade or access to water, whereas work on canals, drains and ponds is carried out in or beside a water body.

Estimated separately, the two larger groups differ in magnitude in the direction exposure

would predict. For the wet-bulb mean the coefficient is −0.0224 (SE 0.0068) for roads and land levelling against −0.0036 (SE 0.0048) for canals, drains and ponds, and the maximum readings show the same ordering. The difference between the two is not, however, estimable with any precision. Read off a pooled regression in which the heat term is interacted with work type, so that the groups share one set of worksite effects and one standard error, roads and land levelling differ from canals, drains and ponds by −0.0108 (SE 0.0104) on the wet-bulb mean, −0.0081 (SE 0.0112) on the wet-bulb maximum and −0.0016 (SE 0.0062) on the dry-bulb maximum. None approaches conventional significance, and the joint test that all three groups respond alike does not reject in any specification, with p-values from 0.232 to 0.921.

On the dry-bulb maximum both larger groups show a negative and statistically significant response, at −0.0124 (SE 0.0046) for roads and land levelling and −0.0068 (SE 0.0034) for canals, drains and ponds, and the two are statistically indistinguishable. Estimates for plantation rest on 669 rolls and are uninformative. We therefore find no evidence that the heat response differs by type of work, including in the group expected to face greater thermal exposure, and we do not interpret the gap between the separately estimated coefficients as a difference. Each is estimated against its own worksite effects, and the comparison that carries a valid standard error does not reject.

3.3.4 Robustness

Absenteeism: Equation (3) requires two decisions that the literature does not settle, which reading of the weather enters as $X_{g,t}$, and over how many days it is accumulated. We vary each in turn, holding the sample, the controls and the fixed effects of column (4) of Table 1 fixed throughout, so that every estimate reported here differs from the baseline in the exposure measure alone.

Table A.6 varies the temperature reading. Three results follow. First, the horizon rather than the thermometer determines the sign. The wet bulb of the day itself, entered alone, is

indistinguishable from zero, and when the contemporaneous and ten-day terms are entered together the contemporaneous coefficient remains zero while the accumulated term retains its full magnitude. Absence responds to exposure that has built up, not to the heat of the morning on which the decision is taken. Second, the choice of thermometer governs magnitude rather than direction. Replacing the wet bulb with the ten-day mean of the dry-bulb daily mean preserves the sign and significance at roughly two-fifths of the magnitude, which is the attenuation expected of a measure that discards humidity and is consistent with the evaporative rationale for (1). Third, the previous day's dry-bulb maximum, which is the regressor much of this literature uses, carries the opposite sign. It differs from our baseline in both the statistic and the horizon, and the two intermediate columns show that the horizon is responsible: both thermometers accumulated over ten days give a positive coefficient, and only the contemporaneous reading gives a negative one.

The same table reports the overnight window from 21:00 to 05:00, contemporaneously and accumulated. The accumulated night term behaves like the accumulated day term. We do not, however, give it a recovery interpretation: the night and day measures correlate too closely over this sample to be treated as separate channels, and a night measure that correlates with its daytime counterpart at close to one is one measurement read twice rather than evidence of a nocturnal mechanism.

Table A.7 varies the window length from three to twenty days. The estimate rises monotonically with the window and then flattens, so ten days is not a length at which a result appears and vanishes on either side of it. Because the windows overlap heavily by construction, the columns are not independent tests and the pattern across them should be read as the shape of an accumulation profile rather than as a set of separate results.

Table A.8 adds solar radiation and wind speed at the same ten-day horizon as the heat term, both standardised within gram panchayat and month of the year. Neither disturbs the result. The coefficient on the ten-day mean wet bulb moves from 0.0057 (0.0016) to 0.0057 with radiation, 0.0056 with wind, and 0.0055 (0.0016) with both, significant at one per cent

throughout, and the hottest quartile holds at 0.0127, 0.0123, 0.0121 and 0.0115, significant at five per cent in every column. Both controls are available on 99.9 per cent of the sample, so nothing here turns on a change of sample. Radiation enters on its own account at -0.0058 (0.0018) and absorbs most of what the rainfall term had been carrying, which falls from 0.0064 (0.0017) to 0.0017 (0.0022) once radiation is included; rainfall was evidently proxying for cloud cover. Wind speed is indistinguishable from zero. Since temperature, radiation and wind move together, the fact that the accumulated heat coefficient is unchanged by conditioning on the other two is the relevant point, the absence response is not a radiation or a stillness effect wearing a thermometer's label.

We also assess the sensitivity of the estimates to alternative clustering schemes that allow for dependence in outcomes within workers, worksites, or survey dates. Clustering at the worker level and two-way clustering by gram panchayat and date leave the point estimates unchanged.

Assessed Wages: Table A.9 lets the response bend, and the conclusion survives. In the preferred column the joint test that all bands are zero returns p-values of 0.158 for the dry-bulb maximum, 0.421 for the wet-bulb mean and 0.732 for the wet-bulb maximum: the ladders are flat. One coefficient recurs across specifications, the dry-bulb band from 35 to 40 degrees, at -0.0630 (0.0280) in the preferred column and -0.0761 (0.0232) without the month effect. We do not build on it, for two reasons. It is one of sixteen band coefficients, which is about what testing sixteen hypotheses at conventional size yields unaided. More tellingly, the band above 40 degrees is null in every column of the table, at -0.0053 (0.0299) in the preferred column. A response that shows up between 35 and 40 degrees and vanishes above 40 is not a dose–response. One possible explanation for this is, that above 40 degrees work stops or shifts to the early morning so the hottest worked days are selected, but it is constructed after the fact and we do not rest anything on it.

Table A.10 adds solar radiation and wind speed, and they change nothing. On the

81,389 rolls that carry both, the wet-bulb mean is 0.0087 (0.0050) before the controls and 0.0094 (0.0050) after; the dry-bulb maximum moves from -0.0004 (0.0026) to 0.0010 (0.0030). Radiation's own coefficient is indistinguishable from zero throughout. The fact that the heat coefficient is unmoved by conditioning on the other two is informative, the result on heat is not a temperature term standing in for the weather that accompanies it.

4 Discrete Choice Experiment

The preceding analysis shows that heat affects both workers' attendance and their earnings on public works. These outcomes, however, do not reveal how workers would trade off heat exposure against alternative working conditions. Wages on public works are administratively set, and access to employment is rationed, leaving limited scope for workers to negotiate for cooler working conditions. To examine worker's preferences directly, we conduct a discrete choice experiment in which workers choose between alternative configurations of a workday at their own worksite. This approach draws on established applications of hypothetical choice experiments in labour economics to value job attributes that are not separately priced in the labour market (Mas and Pallais 2017a; Wiswall and Zafar 2018).

4.1 Data

Our team of enumerators collected data from 1626 workers employed at 72 VBGRAMG worksites spanning 17 districts in Haryana in the months of July, August and early September 2026,[5], shortly after a period of intense heat in Haryana. In June, temperatures reached 40-46°C across many parts of Haryana and neighbouring north India, with the India Meteorological Department issuing heat-wave warnings for Haryana as late as 8-11 June (India Meteorological Department 2026b). The southwest monsoon reached parts of Haryana in

[5]The survey was piloted during April–May 2026. However, the timing of the main survey was affected by the introduction of the Viksit Bharat-Guarantee for Rozgar and Ajeevika Mission (Grammen) (VBGRAMG) in July 2026. While NERGS got replaced by VBGRAMG, the workers were not replaced.

early July, but workers entered the survey period with recent experience of extreme heat and continued to face hot and humid conditions during the dry spells of monsoon. Moreover, the exact amount of rain, and hence relief from heat, was adversely affected by the super El-Nino effect of 2026 (Greenstone et al. 2026). This timing allows us to measure workers' reported symptoms, coping responses, and preferences when heat exposure was recent and salient, rather than asking them to recall experiences from a more distant period.

We selected worksites based on their size and spatial representativeness. We prioritized worksites with at least 30 workersand we randomly selected worksites across blocks, with the number selected from each block proportional to the number of large worksites there. In districts without sufficiently large worksites, we applied the same procedure using a lower size threshold of 15 workers. We replaced selected worksites that were closed on the survey day due to rain or other issues with the nearest active worksite in the same or an adjoining block. As we excluded worksites with fewer than 14 workers, this excluded the predominantly urban districts of Gurugram, Faridabad, and Panchkula. Sonipat and Jhajjar were also excluded because they did not have sufficiently large worksites during the survey period.

At each of the selected worksites, we surveyed the site supervisor (Mate) as well as workers. Workers were chosen at random if more than the desired number of them (30 or 15) were present on the worksite on the day of survey. The mate questionnaire collected information on worksite characteristics and working conditions, including the number of workers registered at the worksite, work schedules, and any adjustments to working hours. It also documented the existing institutional level heat mitigation methods like mate receiving formal communication from the administration regarding heat-related risks.

The worker questionnaire comprised of several modules to understand multiple implications and aspects of working in heat and included a discrete choice experiment. The household and demographic module collected information on workers' individual and household characteristics. The heat-perception module elicited workers' perceptions of heat exposure and heat-related conditions at the worksite. The health module collected information on

health symptoms and other health impacts associated with heat exposure. The work and adaptation module measured workers' heat-related adaptation and coping strategies, including adjustments made during work. A women-specific module was administered only to female workers and collected additional information on gender-specific experiences and responses to heat exposure. Finally, the survey included a discrete choice experiment (DCE) designed to elicit workers' preferences over alternative job characteristics, including wages, working hours, and work quantity. The combination of worksite- and worker-level data also allows us to relate individual responses to the institutional and environmental conditions of the worksites in which workers are employed.

Table 3 provides the demographic characteristics of surveyed workers. Women make up 70% of the surveyed workers. This is in line with their state-wide proportion which has gone up from 59.61% in 2022 to 69.62% in 2026. Our survey has slightly oversampled scheduled caste (SC)[6] workers, which make up 50.29% of all the NREGA workers in Haryana. Roughly half of the workers are illiterate and a third of them have only completed basic schooling. This reflects the high levels of deprivation and poverty among the workers who usually form the unskilled workforce of NREGS. These are the primary intended beneficiaries of the scheme. In our DCE sample, as two comprehension questions gate the discrete choice module, each recording whether the respondent understood at the first explanation, at the second, or not after two. Failing either ends the module, and the second question is not asked of those who fail the first. Of the 1,626 workers interviewed, 612 did not pass the first question and 85 did not pass the second, leaving 929 who were given the tasks. Every one of them completed all eight, so selection into the experiment occurs at the screen and not through attrition within it.

NREGS is self-targeting, so the participation is concentrated among workers with relatively weaker alternatives in the private labour market. The close correspondence between the shares in our sample and the corresponding administrative shares suggests that our sam-

[6]Scheduled Castes (SCs) are historically marginalized caste groups that are officially recognized under India's Constitution and are eligible for affirmative action and other government protections.

ple broadly reflects the NREGS workforce. Because earnings are tied to amount of daily work and there is little cushion behind them, staying home or working less on a hot day carries a direct income cost. This could also explain the larger proportion of SC workers, who are among the poorest of the poor and are more likely to be present at the site on an unusually hot survey day compared to other workers. Household conditions can affect decision to skip work, since dependents, other earners, and the quality of the dwelling determine how much recovery is possible after a day of exposure. Our survey samples the part of the population where these constraints bind hardest.

Excavation work accounts for two-thirds of sampled workers and desilting or cleaning water resources for a further 18%; about half the excavation work is at ponds. The median worker walks 1.8 km to the site, and 77% travel on foot. Tables A.11 and A.12 describe what heat protection exists at these worksites and what workers do in its absence. Twenty-nine of the 65 mates reported receiving formal communication from the administration regarding heat-related risks. Worksites whose mate reported such communication begin at a median of nine o'clock, against a quarter past nine elsewhere, so the advisories are not associated with a shift in the working day.Shade is present at 86% of sites and drinking water at 89%, but only half of workers judge the site's heat-wave measures adequate. Adaptation is therefore mostly individual and nearly universal, 97% drink more water, 96% take breaks in shade, and 77% cover their heads. Where it becomes costly, it divides. Two in five workers report working more slowly in the heat, and one in nine goes home early, in both cases men far more often than women (49% against 39%, and 20% against 8%), which is what one expects when the same decision is cheaper for the worker with fewer claims on the rest of the day. Discretion over the working day is limited, 29% decide their own rest breaks on hot days and 10% decide when work stops, though 69% say they are comfortable asking for a break when unwell. Women constitute around 70% of the workforce. Table A.13 shows that nearly 80% have primary responsibility for household chores, while two-thirds spend 2–4 hours per day on these activities. Almost all report greater domestic responsibilities on hot days,

potentially limiting time for rest and recovery from heat.

Table 3: Characteristics of surveyed MGNREGA workers

Variable	Mean or %	SD
Demographics		
Age (years)	41.9	10.7
Married	86.6 (1408)	
Household size	5.3	2.1
Children in the household	0.8	1.2
Adults over 60 in the household	0.3	0.6
Household members in paid work	2.0	0.9
Education		
No schooling	44.1 (717)	
Primary (5th Class)	32.6 (530)	
Secondary (10th Class)	15.3 (249)	
Higher secondary (12th Class)	6.2 (101)	
College or above	1.8 (29)	
Caste		
Scheduled Caste	64.4 (1047)	
Scheduled Tribe	1.4 (22)	
OBC	27.7 (450)	
General	6.0 (97)	
Other	0.6 (10)	
Household economic position		
Durable goods owned (of 11)	3.95	1.69
Landless	88.4 (1437)	
Less than 1 acre	6.4 (104)	
1-2 acres	4.0 (65)	
More than 2 acres	1.2 (20)	
MGNREGA is the main household income source	38.0 (618)	
Pucca house	94.5 (1537)	
Electricity	98.2 (1597)	
Toilet	88.3 (1436)	
Heat and health		
Heat symptoms reported, last 7 days (of 15)	4.12	3.21
Three or more symptoms	65.9 (1071)	
No symptoms	14.2 (231)	
A household member fell ill from heat, last 30 days	28.1 (457)	
Has a chronic illness	8.0 (130)	
Believes extreme heat can be life-threatening	86.8 (1412)	

Continuous variables report the mean and the standard deviation across workers; binary variables the percentage with the count in parentheses.
$N = 1626$ (1166 women, 460 men), being workers reporting at least five days of MGNREGA work in the previous fifteen.

4.2 Methodology

Our aim is to recover what a later start or a lower output target is worth to a worker in rupees. We therefore model the choice experiment in the random utility framework of McFadden (1974), in which each arrangement delivers a level of utility, the worker chooses the arrangement she most prefers, and the money value of any attribute is the ratio of its coefficient to the coefficient on the wage. Let i index respondents, j alternatives, and c choice tasks. The utility that respondent i derives from alternative j in task c is

$$\begin{aligned} U_{ijc} = {} & \beta_{10}\mathbf{1}\{\text{Start}_{jc} = 10\} + \beta_{11}\mathbf{1}\{\text{Start}_{jc} = 11\} \\ & + \gamma_1\mathbf{1}\{\text{Cut}_{jc} = 10\} + \gamma_2\mathbf{1}\{\text{Cut}_{jc} = 20\} \\ & + \beta_3\text{ExtraWage}_{jc} + \varepsilon_{ijc}, \end{aligned} \tag{5}$$

where Start_{jc} is the start time, Cut_{jc} is the percentage reduction in the output target, and ExtraWage_{jc} is the additional daily wage in hundreds of rupees above the INR 400 base. The term ε_{ijc} collects everything about the arrangement that we do not observe. The omitted levels are a 9 a.m. start, the full output target, and the base wage, so each coefficient measures the value of moving away from the arrangement workers face now [7].

We do not assume that the second hour of delay is worth the same as the first, or that the second ten points off the target are worth the same as the first. Each level therefore enters as its own indicator. Linearity would require $\beta_{11} = 2\beta_{10}$ and $\gamma_2 = 2\gamma_1$, and we test both restrictions inside the saturated model rather than by fitting a separate model for each attribute. This matters because the experimental profiles correlate the attributes with one another, so imposing linearity on the start time would shift the estimated value of the target and vice versa. We also test whether the two attributes interact, and retain the additive specification when they do not.

[7] The experiment fixes the status quo at a nine o'clock start. This is the median start time among the 67 worksites for which the mate reported one, and the most common single time, but there is dispersion around it, starts run from six to half past ten, with seventeen sites beginning before nine and thirty-one after. Contrasts against the status quo therefore compare the experimental arrangements with a specified schedule rather than with each respondent's own, which is a further reason to read the ten-to-eleven step, which involves no such comparison, as the contrast the design identifies cleanly.

Each task presents three arrangements, but they are not three comparable points on a single scale, and this shapes how we estimate the model. Two of the three are experimentally varied: one starts at 10 a.m. and one at 11 a.m., and both carry randomly assigned targets and wages. The third is the status quo, which is the arrangement workers already have — a 9 a.m. start, the full target, and the base wage. We therefore estimate the model on two choice sets, each of which answers a different question.

The first choice set contains only the two experimentally varied alternatives. This is the cleanest comparison in the experiment, because both alternatives are new arrangements and differ only in their assigned attributes. Since every task pairs a 10 a.m. alternative with an 11 a.m. one, the two start-time indicators carry the same information here, so we take 10 a.m. as the reference and estimate a single coefficient for the 11 a.m. start. Respondents who chose the status quo in all eight tasks never express a preference between the two, and so drop out of this specification. Estimates from this choice set describe the preferences of workers who were willing to consider a change at least once.

The second choice set restores the status quo, which lets us model the decision workers actually face: whether to move away from the current arrangement at all. What we gain in realism we lose in interpretation, because the status quo is the only alternative that sits at all three omitted levels at once. In every task,

$$\mathbf{1}\{j = \text{SQ}\} = 1 - \mathbf{1}\{\text{Start}_{jc} = 10\} - \mathbf{1}\{\text{Start}_{jc} = 11\}, \tag{6}$$

so a constant for the status quo is exactly a combination of the two start-time indicators and cannot be estimated beside them. The deeper limitation is that a 9 a.m. start, the full target, and the base wage occur together and nowhere apart, so the value of those features cannot be told apart from an attachment to the familiar arrangement itself (Samuelson and Zeckhauser 1988). For this reason we treat the 10-to-11a.m. contrast as the start-time comparison the experiment identifies cleanly, and we report contrasts against the status quo

under the normalisation that no separate preference for the status quo exists. They are not valuations of the 9 a.m. start.

How much weight these trade-offs can bear depends on one feature of the design. The two experimental alternatives in a task differ in start time, target, and wage at the same time, so within any single task the 11 a.m. indicator and the wage move together exactly and their effects cannot be told apart. What separates them is that this relationship reverses across tasks, in some the later start is the better paid option and in others it is the worse paid. The variation identifying the start-time and wage trade-off is therefore the number of tasks in which that relationship reverses, which is far smaller than the number of choices we observe. We accordingly report how the estimates move when each task is dropped in turn. For the same reason we do not estimate the model on halves of the sequence, since within some halves the start time and the wage do not vary independently at all.

We estimate the model by conditional logit, treating each task as a stratum and clustering standard errors by respondent. Conditional logit assumes independence of irrelevant alternatives, and our design strains that assumption in a particular way: both experimental alternatives are departures from the current arrangement and differ only in how far they depart, while the status quo is qualitatively different from both. A worker who first decides whether to accept any change, and only then which change, would have unobserved preferences that are shared between the two experimental alternatives but not with the status quo (Hausman and McFadden 1984). Because those two alternatives are generic and unlabelled, nothing distinguishes them from each other in this respect, so the concern is confined to the status quo against the pair. We test the assumption with the Hausman and McFadden (1984) statistic, which compares estimates from the full choice set with those from the pair alone, and with the Small and Hsiao (1985) statistic, computed over repeated random splits of respondents because its value depends on how the sample is split.

We then relax the assumption in two ways, because each addresses a different concern. A nested logit groups the two experimental alternatives together and places the status quo

in a nest of its own, allowing unobserved preferences to be shared within the pair. Its dissimilarity parameter equals one when independence holds, at which point the model reduces to conditional logit, and we confirm this numerically. The nest and a status-quo constant describe the same distinction between the familiar arrangement and the new ones, so they draw on the same variation, and we do not read the 9 a.m. levels in any specification that includes the nest.

A mixed logit instead allows preferences themselves to differ across workers (Revelt and Train 1998; Train 2016). We let the coefficient on the start-time contrast vary across respondents and estimate the model on the experimental pair by simulated maximum likelihood with 500 Halton draws (Bhat 2001), so that it can be compared directly with the conditional logit on the same choice set. Random coefficients make the utilities of alternatives that share an attribute correlated through the taste they have in common, which is the correlation that independence of irrelevant alternatives rules out; and because a respondent's taste persists across her eight choices, we can estimate how widely it varies. With a suitable choice of variables and mixing distribution, a mixture of logits can approximate any random utility model arbitrarily closely (McFadden and Train 2000). In our specification, though, only the start-time coefficient varies, so the assumption is relaxed along that dimension alone. This is why we report the nested model alongside it rather than in place of it.

Two coefficients are held fixed in the mixed logit. The wage coefficient is fixed so that the compensation measures below do not divide by a random quantity, which can leave the resulting distribution without finite moments (Daly, Hess, and Train 2012). The output-target coefficients are fixed. The target itself varies freely across the alternatives, it differs between the two experimental options in every task, by one level or by two and in both directions, so this is a modelling choice rather than a limitation of the design, and Table A.14 reports the specifications in which the target coefficient is allowed to vary. For the start-time coefficient we report two mixing distributions. A normal distribution allows some workers to prefer a later start and others to prefer an earlier one. A log-normal holds the sign the

same for everyone, and we parameterise it as

$$\beta_{11,i} = -\exp(\mu + \sigma v_i), \qquad v_i \sim N(0,1),$$

which keeps the coefficient negative for every respondent. The median of this distribution is $-\exp(\mu)$ and its mean is $-\exp(\mu+\sigma^2/2)$. These separate once the distribution is dispersed, so we report the compensation implied at both, and the tables give the median coefficient in place of the location parameter μ.

Finally, we express the estimates in rupees. For an attribute contrast R, the compensating differential implied by the model is

$$\mathrm{WTA}(R) = -100\frac{R'\beta}{\beta_{\mathrm{ExtraWage}}}, \tag{7}$$

in rupees per day. A positive value is what a worker must be paid to accept the change; a negative value is what she would give up to obtain it.

We build confidence intervals by the Krinsky–Robb procedure (Krinsky and Robb 1986), which draws from the joint asymptotic distribution of the coefficients and takes percentiles of the resulting ratio. A ratio of two estimated coefficients has a skewed sampling distribution even when both are precisely estimated, so the symmetric interval the delta method produces can misstate it. Because the quantity is a ratio of two coefficients from the same model, it is unaffected by a common rescaling of utility, which is what makes it comparable across the specifications above. In the mixed logit the varying coefficient implies a distribution of compensating differentials rather than a single figure, and we summarise it by its mean and median. These are the payments implied by stated choices within the experiment. They are not wage differentials observed in the labour market, and because every task required the respondent to pick one of three arrangements, they describe preferences over how work is organised rather than whether to work at all.

We examine heterogeneity by estimating the model separately within respondent sub-

groups rather than by interacting attributes with characteristics. The reason is that the scale of the logit error cannot be separated from the coefficients themselves (Swait and Louviere 1993), so a larger coefficient in one subgroup may reflect less unobserved noise rather than a stronger preference. Compensating differentials do not have this problem, because a common scale factor cancels in the ratio, and we therefore compare subgroups onthat basis. Where we compare two groups we report uncertainty for the difference between their compensating differentials, since two intervals that overlap do not establish that the groups are alike.

4.3 Results

The choice experiment measures workers' willingness to trade wages against changes in the timing and intensity of work. Of the 7,432 choice tasks, workers selected the status quo in 918 (12.4 per cent) and chose one of the two experimental arrangements in the remaining 87.6 per cent. Workers therefore chose an alternative to the existing arrangement in roughly seven out of every eight tasks. The central question is consequently not only whether workers would accept a change, but how much compensation they require for different changes to the working day.[8]

Table 4 reports the estimated choice models.The clearest result is that workers must be compensated to start later, but the amount required for a one-hour delay is modest relative to the value they place on a lower output target. In column (1), which compares the two experimental alternatives and allows the effects of the attributes to vary across levels, the coefficient on an 11 a.m. start relative to a 10 a.m. start is -0.257 (standard error 0.039), while the wage coefficient is 2.247 (0.075). Their ratio implies that workers require approximately INR 11 per day to accept a one-hour delay in the start time. The Krinsky–Robb confidence interval is INR 8–15, relative to the INR 400 base wage.

[8]Two comprehension questions gate the module and 929 of 1,626 respondents reached the tasks. Appendix E.2 reports who was removed, those admitted are younger and better educated and differ on gender, caste, marital status and assets, but not in heat symptoms. We do not reweight, the experiment identifies preferences among workers who understood it.

This estimate is stable across alternative specifications. The full choice-set model in column (2) implies compensation of INR 11, with a confidence interval of INR 8–14. The nested logit in column (3) also imposes linearity in the attributes, so its INR 12, with an interval of INR 9–15, is a slope per hour rather than the 10-to-11 a.m. step. The mixed logit in column (4), which allows the start-time coefficient to vary across workers, implies INR 12 for the step itself. Three specifications therefore price the same step at INR 11 to 12. The estimated wage coefficients differ across models, ranging from 1.734 to 2.335, but the implied compensating differentials remain close. This stability reflects the fact that compensation is calculated as a ratio of the attribute and wage coefficients, so a common rescaling of utility does not change the ratio.

The average valuation, however, conceals substantial heterogeneity in preferences over start time. In the mixed logit, the estimated standard deviation of the start-time coefficient is 0.458 (standard error 0.074). A likelihood-ratio test against the fixed-coefficient model gives a statistic of 12.5 on one degree of freedom ($p = 0.0004$), providing evidence that preferences vary across respondents. Under the estimated normal distribution, 27.9 per cent of workers have a positive coefficient on the later start and would therefore accept it without additional pay.

Table A.15 shows how the implied compensation distribution changes when the coefficient is instead constrained to be negative using a log-normal specification. The median falls to INR 4, while the mean remains INR 13. The difference between these estimates indicates a dispersed, asymmetric distribution in which the mean is influenced by workers with higher valuations. Despite this heterogeneity, INR 12 per day remains a useful summary of the average valuation across specifications, though the log-normal fit is a reminder that the typical worker may value the hour at considerably less.

The value workers place on a lower output target is larger than the value of a later start, and it is concentrated in the first reduction. In column (1), a ten-percentage-point reduction in the target raises utility by 0.523 (standard error 0.051), an implied compensating

differential of about INR −23 per day. Workers would forgo that amount to obtain the lower target. This is roughly twice the compensation they require for a one-hour delay in the start time, and in the opposite direction. A twenty-percentage-point reduction raises utility by only 0.034 (standard error 0.049), a compensating differential of about INR −1 that cannot be told from zero. Because both coefficients are measured against the full output target, this does not mean that the second ten percentage points add little; it means that they take back what the first ten gave. In the full choice set of column (2), where the two levels are estimated more precisely, the value of moving from a ten-point to a twenty-point reduction is $\gamma_{20} - \gamma_{10} = -0.419$ (standard error 0.036), so a worker would require about INR 20 per day to accept it, with a Krinsky–Robb interval of INR 16 to 23 (Table A.14). The value of a lower target is therefore not merely concentrated in the first reduction: a larger reduction is worse than a smaller one.

This pattern is not confined to a single specification. The full choice-set model in column (2) implies compensating differentials of INR −25 for a ten-point reduction and INR −5 for a twenty-point reduction; the mixed logit in column (4) implies INR −23 and INR −1. The restriction that the two target levels follow a linear pattern is strongly rejected: in column (2), $\gamma_{20} - 2\gamma_{10} = -0.951$ (standard error 0.066). Imposing linearity, as the nested logit in column (3) does, averages the substantial value of the first reduction against the negative value of the second and yields a coefficient of 0.016 (standard error 0.016), indistinguishable from zero. The unrestricted estimates reveal a distinction that linearity conceals: workers value an initial reduction in required output, and would need paying to accept a further reduction of the same size.

The full choice set also permits comparisons with the status quo, which combines a 9 a.m. start, the full output target, and the INR 400 base wage. In column (2), the implied compensation for moving from 9 a.m. to 10 a.m. is INR 42 (confidence interval INR 34–50). This decomposes into a common component of about INR 31.5 for leaving the status quo at all plus the INR 11 estimated above for the hour.

The common component cannot, however, be interpreted as the value of changing the start time from 9 a.m. alone. The status quo is the only alternative that combines a 9 a.m. start with the full output target and the base wage. The design therefore cannot separate workers' preferences for these baseline attributes from a general preference for the existing arrangement. Panel D of Table A.16 illustrates the same identification problem in the nested logit. Fitting the nest on the saturated specification, in which the 9 a.m. levels are estimated freely, drives the dissimilarity parameter from 0.765 to 0.240 and reverses the 9-to-10 a.m. contrast altogether: where the unnested model implies that workers require INR 42 to accept a 10 a.m. start, the nested one implies they would give up INR 229 a day — more than half the daily wage — to obtain it. The nest and the 9 a.m. levels describe the same distinction between the existing arrangement and the alternatives, and the data cannot support both. The 10-to-11 a.m. step, which does not involve the status quo, is unaffected at INR 12.

The status-quo comparisons therefore suggest that workers' choices reflect more than the incremental cost of delaying the start by an hour. The estimated contrast between the existing arrangement and the experimental alternatives is larger than the value of the 10-to-11 a.m. shift. But because the existing arrangement bundles start time, output requirements, and wage, the experiment cannot identify how much of this difference is attributable to the start time itself and how much reflects the other features of the status quo.

These magnitudes are modest in rupees and less modest as shares of pay. Against the base wage of INR 400, an hour's later start is 2.8 per cent of a day's earnings, the first ten-point cut in the target is 5.8 per cent, and the fixed component of moving away from a nine o'clock start is 7.9 per cent. The closest comparisons are stated-preference estimates for American workers. Maestas, Kathleen J Mullen, et al. (2023a) find a willingness to give up 9.0 per cent of wages to set one's own schedule and 12 to 15 per cent to move from heavy physical work to moderate work or to mostly sitting, while Mas and Pallais (2017b) find in a hiring experiment that most workers will pay nothing for schedule flexibility itself.

Read against those, the two attributes separate. The hour is cheap, 2.8 per cent for a

one-hour shift, and the near-zero valuation most American workers place on flexible timing suggests this is not peculiar to our setting. What is not cheap is leaving the existing arrangement at all, at 7.9 per cent, which is close to the American figure for schedule control; and neither is the target, at 5.8 per cent for a ten per cent reduction in required output against 12 to 15 per cent for a far larger change in physical demands. Scaled to the size of the change offered, these workers value relief from how much they must produce at least as highly as better-paid workers elsewhere value relief from physical effort. That is the more striking of the two comparisons, since a worker earning INR 400 a day has a high marginal utility of income and should on that account be less willing, not more, to trade pay for conditions.

4.3.1 Robustness

We assess the robustness of the estimated compensating differentials to alternative inference procedures, changes in the choice-task sample, variation across enumerators, and departures from the conditional logit specification. The estimates are broadly stable across model specifications, although sensitivity to the composition of the experimental tasks highlights an important limitation of the design.

The main estimates are little changed by alternative inference procedures or by excluding the first choice task. Clustering standard errors at the worksite rather than the respondent leaves the estimated compensation for a one-hour delay at INR 12 and that for a ten-percentage-point reduction in the output target at INR -23. Excluding the first task, to allow for possible respondent unfamiliarity with the choice format, yields corresponding estimates of INR 10 and INR -22. These changes are small relative to the estimates' standard errors.

A more consequential source of uncertainty is the limited independent variation in the experimental design. In each task, the two experimental alternatives differ simultaneously in start time, output target, and wage. The effects of start time and wage can therefore be distinguished only through changes in their relationship across tasks. We assess the

sensitivity of the estimates by dropping each task in turn. The implied compensation for a one-hour delay ranges from INR 7 to INR 19, compared with the pooled estimate of INR 11; the estimate for a ten-percentage-point reduction in the output target ranges from INR −7 to INR −35, compared with the pooled estimate of INR −23. These ranges capture sensitivity to the realised task design, whereas the Krinsky–Robb confidence intervals capture sampling uncertainty conditional on that design. They therefore provide complementary information and should be considered together.

We also examine variation across enumerators. Among the twelve enumerators who administered the experiment, seven interviewed at least sixty respondents. Across these seven enumerators, the estimated compensation for a one-hour delay ranges from INR 3 to INR 24, with the lowest estimate not statistically distinguishable from zero. The corresponding estimates for a ten-percentage-point reduction in the output target range from INR −13 to INR −34 and are negative for all seven enumerators. These differences warrant caution in interpreting the pooled estimates, particularly for start time. Enumerators were not randomly assigned to worksites, however, so the estimates combine differences in enumerators with differences in the worksites and respondents they served. They should therefore be read as descriptive evidence of variation across enumerator-specific samples, not as estimates of causal enumerator effects.

Finally, we examine whether the results are sensitive to the conditional logit's independence of irrelevant alternatives (IIA) assumption. Table A.16 rejects it three ways: the Hausman–McFadden test gives $\chi^2 = 56.9$ on three degrees of freedom, the Small–Hsiao test rejects in 84 per cent of 25 random splits, and the nested logit's dissimilarity parameter is 0.765 with a standard error of 0.035, significantly below one. All three say that the two experimental alternatives are closer substitutes for each other than either is for the status quo, which is how the experiment was built. The rejection does not move the quantities we report. A one-hour delay is worth INR 11 in the conditional logit on the experimental pair, INR 11 in the full choice set, INR 12 in the nested logit and INR 12 in the mixed logit.

Within a nest the choice probabilities follow a binary logit whose coefficients are scaled by the inverse of the dissimilarity parameter, and a compensating differential is a ratio of two coefficients from the same model, so that common factor cancels.

That agreement establishes less than it may appear to. It shows the estimates do not depend on the logit error structure, which is what the rejection would otherwise call into question, but not that they are precisely identified. All four specifications are fitted to the same sixteen experimental profiles, so any feature of the realised design propagates to every one of them, and only three of the eight tasks pair the later-starting alternative with the worse pay. The leave-one-task-out range in Table 4 measures that: INR 7 to INR 19 for the hour, against a Krinsky–Robb interval of INR 8 to 15. Dropping a single task moves the estimate further than sampling error alone would allow, so we read the design range rather than the interval as the width of what this experiment can say. What survives both is the sign and the ordering rather than the exact figure. Workers require payment to start later and would forgo pay for a lower target in every specification and under every task exclusion, and the target is worth roughly twice the hour throughout. The point estimates are not pinned to the rupee, and we do not treat them as such.

4.3.2 Heterogeneity

We examine whether workers' preferences for later start times and lower output targets vary across respondent characteristics. Table A.17 reports estimates from separate conditional logit models for each subgroup. Because the logit scale is not separately identified from the coefficients, we compare groups using compensating differentials rather than raw coefficients. Panel B reports the difference in compensating differentials between the two groups for every binary split considered, covering eighteen comparisons across the two attributes. We account for this multiple testing when assessing which differences provide robust evidence of heterogeneity.

The clearest difference concerns women's household workload. Women who spend at least

four hours a day on household work are willing to accept INR 33 less for a ten-percentage-point reduction in the output target, compared with INR 17 among other women. The difference of INR 16 has a standard error of INR 5 and is the only subgroup difference that remains statistically significant after correcting for multiple comparisons. By contrast, the two groups do not differ significantly in the compensation they require for a one-hour delay in the start time. This pattern suggests that, among women with heavier domestic workloads, reducing the amount of work required may matter more than changing when the workday begins.

Several other differences are suggestive but less conclusive. Workers under forty require approximately INR 7 more for a one-hour delay than workers aged forty and above. Women likewise require approximately INR 7 more than men, while workers who rely on MGNREGA as their main household income require approximately INR 7 less. These differences are individually close to conventional significance thresholds but do not remain significant after accounting for the eighteen comparisons. They should therefore be treated as suggestive patterns rather than firm evidence of subgroup differences. The estimates also provide no clear evidence of heterogeneity by schooling, caste, assets, or landholding.

Finally, preferences differ descriptively by reported heat symptoms, although the differences are imprecisely estimated. Workers reporting three or more symptoms require INR 5 more for a one-hour delay and would give up INR 8 more for a ten-percentage-point reduction in the output target than workers reporting fewer symptoms. Neither difference is statistically distinguishable from zero. The direction of these estimates is consistent with the possibility that workers already experiencing greater heat-related discomfort place more value on reducing the amount of work required than on shifting its timing. Given the uncertainty, however, we treat this pattern as suggestive rather than as evidence that heat symptoms systematically alter workers' preferences.

Table 4: Discrete choice experiment: the price of a later start and of a lower target

	(1) Conditional logit	(2) Conditional logit	(3) Nested logit	(4) Mixed logit
Attributes	Saturated experimental pair	Saturated full, with the status quo	Linear full, with the status quo	Saturated experimental pair
Choice set				
A. Start time				
Eleven rather than ten o'clock	-0.257***	-0.226***		-0.269***
	(0.039)	(0.033)		(0.041)
Ten rather than nine o'clock		-0.895***		
		(0.101)		
Eleven rather than nine o'clock		-1.120***		
		(0.104)		
Each hour later than nine			-0.209***	
			(0.029)	
B. Output target				
Cut of ten points	0.523***	0.533***		0.544***
	(0.051)	(0.044)		(0.055)
Cut of twenty points	0.034	0.114**		0.029
	(0.049)	(0.047)		(0.055)
Each ten-point cut			0.016	
			(0.016)	
C. Wage, nesting and heterogeneity				
Wage, hundreds of rupees	2.247***	2.118***	1.734***	2.335***
	(0.075)	(0.067)	(0.076)	(0.062)
Dissimilarity parameter λ			0.765***	
			(0.035)	
SD of the start coefficient				0.458***
				(0.074)
D. Compensation required, rupees per day				
Eleven rather than ten o'clock	11	11		12
Krinsky–Robb 95%	[8, 15]	[8, 14]		
Leave-one-task-out range	7 to 19			
Ten rather than nine o'clock		42		
Krinsky–Robb 95%		[34, 50]		
Each hour later, linearity imposed			12	
Krinsky–Robb 95%			[9, 15]	
A ten-point cut in the target	-23	-25		-23
Krinsky–Robb 95%	[-27, -19]	[-29, -22]		
Leave-one-task-out range	-35 to -7			
A twenty-point cut	-1	-5		-1
Each ten-point cut, linearity imposed			-1	
Respondents	903	929	929	903
Choice sets	6,514	7,432	7,432	6,514
Log-likelihood	-2791.2	-5615.0	-5674.5	-2784.9

The status quo is a 9 a.m. start, the full output target, and INR 400 per day. Columns (1) and (4) exclude the status quo; column (2) includes it.

Column (3) is a nested logit; column (4) is a mixed logit with a normally distributed start-time coefficient, estimated using 500 Halton draws.

Panel D reports compensating differentials in INR per day. Krinsky–Robb 95% confidence intervals use 5,000 draws. Leave-one-task-out ranges report estimates after excluding each task in turn.

Columns (1) and (4) exclude 26 respondents who always chose the status quo. Standard errors clustered by respondent are in parentheses. Log-likelihoods are comparable only within choice sets. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

5 Conclusion

In this paper we study how heat affects the work of economically vulnerable outdoor workers when opportunities to adjust labour supply are limited, and how workers value feasible adaptations to those conditions. Using administrative records from India's National Rural Employment Guarantee Scheme and a discrete choice experiment, we examine both sides of this question, how workers respond to heat through attendance and productivity, and which changes to the organisation of work they would prefer.

We find that heat affects workers on both margins, but unevenly. Within a worksite over March to June, a degree of mean wet-bulb temperature lowers the probability that a muster roll is assessed at the notified rate by 0.008 (0.004), about one per cent of its mean, and the dry-bulb and maximum readings agree; holding the same worker fixed at the same site attenuates this by about a fifth without reversing it. Attendance responds more clearly. Same-day heat has no detectable effect, but a one-degree increase in the ten-day mean wet bulb is followed by a 3.0 per cent increase in absence. The consequences of heat therefore cannot be read from daytime exposure alone, measured output falls modestly if at all, and the clearer cost appears in the days workers are subsequently unable to work.

The choice experiment provides a complementary perspective on how workers would like these constraints to be addressed. Workers require approximately INR 11 a day to accept a one-hour later start, while they would forgo approximately INR 23 a day for a ten per cent reduction in the output target. The relief they value most is therefore in how much work is required rather than in when it takes place, and it is concentrated in the first reduction. This preference is not uniform. The mixed-logit estimates indicate substantial heterogeneity in the valuation of the hour, with 28 per cent of workers willing to accept a later start without additional pay, and the value of a lower target is twice as high among women carrying four or more hours of household work a day as among other women. The implication is therefore not that there is a universally preferred arrangement, but that the costs of adapting differ across workers according to their household, demographic, and economic circumstances.

Taken together, the administrative and stated-preference evidence points to an important distinction between the capacity to adapt and the cost of adapting. Workers in NREGS have limited ability to avoid a hot workday because employment is rationed and a missed day represents foregone income. They nevertheless adjust their behaviour at work, including by slowing down and taking breaks, and while their measured output is sustained, the accumulated exposure shows up in later absence. The DCE indicates that, within the programme's institutional constraints, workers value a reduction in how much work they are required to complete more than a change in when work takes place, though both are valued. Heat adaptation in this setting can therefore involve changing the organisation of work rather than simply relying on workers to withdraw from employment.

These results have several implications for program design. First, interventions that shift work away from the hottest hours may be feasible without changing the underlying employment guarantee or requiring capital-intensive cooling technologies. Second, existing protections such as shade, drinking water, and rest periods are relevant not only as health measures but also as mechanisms for preserving workers' capacity to work under extreme conditions. Third, a single schedule is unlikely to accommodate all workers equally well. In particular, differences in household responsibilities, travel, age, and dependence on program earnings may alter the costs of changing work hours. Policy therefore needs to allow for heterogeneity rather than treating an earlier or later start as a universally appropriate response. Finally, because workers cannot necessarily substitute away from NREGS employment when conditions are dangerous, information or heat advisories alone may have limited effectiveness when they are not accompanied by feasible changes in working conditions.

Several limitations qualify these conclusions. First, the productivity measure is an administratively assessed group piece rate rather than individual output. The estimates therefore capture changes in assessed output per person-day and may understate the full productivity effect because a large share of observations are recorded at the statutory wage ceiling. Second, the register does not record when work was measured, so if the measurement period

differs from the muster-roll period the roll-mean temperature is a mismeasured proxy for the exposure the assessed output reflects, which would bias the productivity estimate towards zero. Third, the DCE measures stated preferences rather than realized choices under an actual policy change. Workers' willingness to accept a change in work timing therefore need not translate one-for-one into behavioral responses or program costs in practice. The experiment also does not identify the causal effect of changing work hours on attendance, productivity, health, or earnings.

Future work could address these limitations by evaluating actual changes in work schedules or other heat-adaptation measures at worksites. In particular, a field experiment or policy rollout that varies start times across comparable worksites could test whether the schedule changes workers prefer translate into higher productivity, lower heat exposure, or different attendance patterns, while also identifying distributional effects across workers with different household constraints.

More broadly, the findings show that the economic consequences of climate change depend not only on how much workers are exposed to heat, but also on the institutional constraints governing their ability to respond. In a program where workers have limited scope to withdraw from work or substitute across jobs, heat appears to reduce the number of days they can work rather than the amount they produce on the days they do. At the same time, workers express a preference for adaptations that reduce the amount of work required more than alter its timing. Understanding these margins is important both for measuring the labor-market costs of rising temperatures and for designing employment programs that remain workable as extreme heat becomes more frequent.

References

Banerjee, Abhijit et al. (2020). "E-governance, accountability, and leakage in public programs: Experimental evidence from a financial management reform in India". In: *American Economic Journal: Applied Economics* 12.4, pp. 39–72.

Basu, Arnab K., Nancy H. Chau, and Ravi Kanbur (2009). "A Theory of Employment Guarantees: Contestability, Credibility and Distributional Concerns". In: *Journal of Public Economics* 93.3–4, pp. 482–497.

Besley, Timothy and Stephen Coate (1992). "Workfare versus Welfare: Incentive Arguments for Work Requirements in Poverty-Alleviation Programs". In: *American Economic Review* 82.1, pp. 249–261.

Bessone, Pedro et al. (2021). "The Economic Consequences of Increasing Temperatures on Sleep". In: *Quarterly Journal of Economics*.

Bhat, Chandra R (2001). "Quasi-random maximum simulated likelihood estimation of the mixed multinomial logit model". In: *Transportation Research Part B: Methodological* 35.7, pp. 677–693.

Bhattacharya, Leena, Shobhit Kulshreshtha, and Arthur van Soest (2026). *Surviving the Daily Heat: Climate Change and Labor Productivity*. Tech. rep. GLO Discussion Paper.

Cosaert, Sam, M. Nieto Castro, and Konstantinos Tatsiramos (2023). "Temperature and the Timing of Work". In: *Journal of Labor Economics*.

Daly, Andrew, Stephane Hess, and Kenneth Train (2012). "Assuring finite moments for willingness to pay in random coefficient models". In: *Transportation* 39.1, pp. 19–31.

Das, K. and E. Somanathan (2024). "Temperature and Earnings in India's Informal Economy". In: *Environmental Research Letters*.

Dutta, Puja et al. (2014). "Right to Work? Assessing India's Employment Guarantee Scheme in Bihar". In: *World Bank*.

Feld, Jan, Alexandra Nagy, and Adam Osman (2022). "Disentangling the Effects of Job Attributes on Workers' Job Preferences". In: *Journal of Development Economics*.

Garg, Teevrat, Matthew Gibson, and Fang Sun (2020). "Extreme Temperatures and Labor Supply". In: *Working Paper*.

Graff Zivin, Joshua and Matthew Neidell (2014). "Temperature and the Allocation of Time: Implications for Climate Change". In: *Journal of Labor Economics* 32.1, pp. 1–26.

Greenstone, Michael et al. (2026). *451,000 at Risk: Where Action Can Save Lives During the Super El Niño*. Climate Impact Lab. URL: https://impactlab.org/research/451000-at-risk/ (visited on 09/24/2026).

Hausman, Jerry and Daniel McFadden (1984). "Specification tests for the multinomial logit model". In: *Econometrica: Journal of the econometric society*, pp. 1219–1240.

Heyes, Anthony and Soodeh Saberian (2022). "Temperature and Human Capital: Evidence from the Indian Labour Market". In: *Journal of Development Economics*.

Imbert, Clément and John Papp (2015). "Labor Market Effects of Social Programs: Evidence from India's Employment Guarantee". In: *American Economic Journal: Applied Economics* 7.2, pp. 233–263.

— (2020). "Short-Term Migration and Rural Labor Markets: Evidence from India". In: *Journal of the European Economic Association* 18.1, pp. 327–355.

India Meteorological Department (May 2026a). *Current Weather Status and Extended Range Forecast for the Next Two Weeks (28 May to 10 June 2026)*. Press Release. Ministry of Earth Sciences, Government of India. URL: https://internal.imd.gov.in/press_release/20260528_pr_5026.pdf.

— (June 2026b). *Press Release: Weather Forecast and Heat Wave Warning*. URL: https://mausam.imd.gov.in/Forecast/marquee_data/Press%20Release%2010-06-2026.pdf (visited on 09/20/2026).

Kjellstrom, Tord et al. (2019). *Working on a Warmer Planet: The Impact of Heat Stress on Labour Productivity and Decent Work*. Geneva: International Labour Organization.

Krinsky, Itzhak and A Leslie Robb (1986). "On approximating the statistical properties of elasticities". In: *The review of economics and statistics*, pp. 715–719.

Maestas, Nicole, Kathleen J Mullen, et al. (2023a). "The value of working conditions in the United States and implications for the structure of wages". In: *American Economic Review* 113.7, pp. 2007–2047.

Maestas, Nicole, Kathleen J. Mullen, et al. (2023b). "The Value of Working Conditions in the United States and Implications for the Structure of Wages". In: *American Economic Review* 113.7, pp. 2007–2047.

Mas, Alexandre and Amanda Pallais (2017a). "Valuing Alternative Work Arrangements". In: *American Economic Review* 107.12, pp. 3722–3759.

— (2017b). "Valuing alternative work arrangements". In: *American Economic Review* 107.12, pp. 3722–3759.

McFadden, Daniel (1974). "The measurement of urban travel demand". In: *Journal of public economics* 3.4, pp. 303–328.

McFadden, Daniel and Kenneth Train (2000). "Mixed MNL models for discrete response". In: *Journal of applied Econometrics* 15.5, pp. 447–470.

Ministry of Rural Development, Government of India (2026a). *Mahatma Gandhi National Rural Employment Guarantee Scheme: National Mobile Monitoring System (NMMS).* URL: `https://nrega.dord.gov.in/MGNREGA_new/Nrega_home.aspx` (visited on 09/20/2026).

— (July 2026b). *State/Union Territory-wise Details of Persondays Generated and Persons Who Availed of Employment under Mahatma Gandhi NREGS.* URL: `https://www.pib.gov.in/PressReleasePage.aspx?PRID=2289110` (visited on 09/20/2026).

Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar (2016). "Building State Capacity: Evidence from Biometric Smartcards in India". In: *American Economic Review* 106.10, pp. 2895–2929.

— (2023). "General Equilibrium Effects of (Improving) Public Employment Programs: Experimental Evidence from India". In: *Econometrica* 91.4, pp. 1221–1259.

Nandwani, S., P. Roychowdhury, and S. Shankar (2026). "Heat Stress and Labour Demand in India's Informal Sector". In: *Global Labour Organization Discussion Paper* 1785.

Neidell, Matthew et al. (2021). "Temperature and Labor Supply". In: *Journal of Economic Perspectives*.

Press Information Bureau, Government of India (Feb. 2022). *Beneficiaries under MGNREGA*. URL: https://www.pib.gov.in/Pressreleaseshare.aspx?PRID=1796871 (visited on 09/20/2026).

— (Aug. 2023). *Women Participation under Mahatma Gandhi National Rural Employment Guarantee Scheme*. URL: https://www.pib.gov.in/PressReleasePage.aspx?PRID=1942377 (visited on 09/20/2026).

Revelt, David and Kenneth Train (1998). "Mixed logit with repeated choices: households' choices of appliance efficiency level". In: *Review of economics and statistics* 80.4, pp. 647–657.

Rural Development Department, Government of Haryana (2026). *Mahatma Gandhi National Rural Employment Guarantee Scheme (MGNREGS)*. URL: https://haryanarural.gov.in/scheme/mahatma-gandhi-national-rural-employment-guarantee-scheme-mgnregs/ (visited on 09/20/2026).

Samuelson, William and Richard Zeckhauser (1988). "Status quo bias in decision making". In: *Journal of risk and uncertainty* 1.1, pp. 7–59.

Small, Kenneth A and Cheng Hsiao (1985). "Multinomial logit specification tests". In: *International economic review* 26.3, pp. 619–627.

Somanathan, E. et al. (2021). "The Impact of Temperature on Productivity and Labor Supply: Evidence from Indian Manufacturing". In: *Journal of Political Economy* 129.6, pp. 1797–1847.

Swait, Joffre and Jordan Louviere (1993). "The role of the scale parameter in the estimation and comparison of multinomial logit models". In: *Journal of marketing research* 30.3, pp. 305–314.

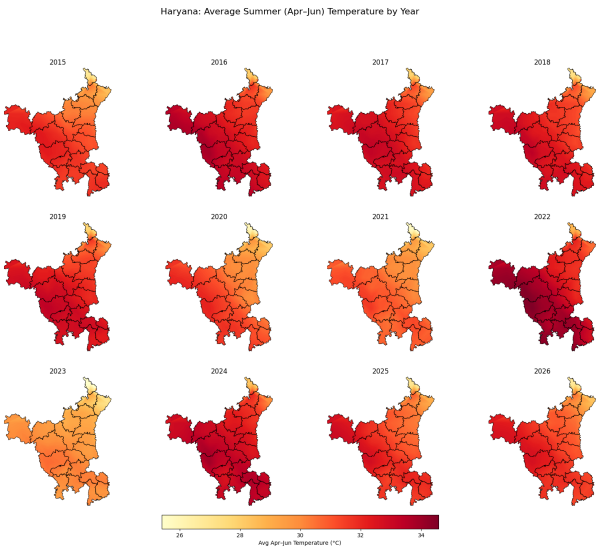
Train, Kenneth (2016). "Mixed logit with a flexible mixing distribution". In: *Journal of choice modelling* 19, pp. 40–53.

Wiswall, Matthew and Basit Zafar (2018). "Preference for the Workplace, Investment in Human Capital, and Gender". In: *Quarterly Journal of Economics* 133.1, pp. 457–507.

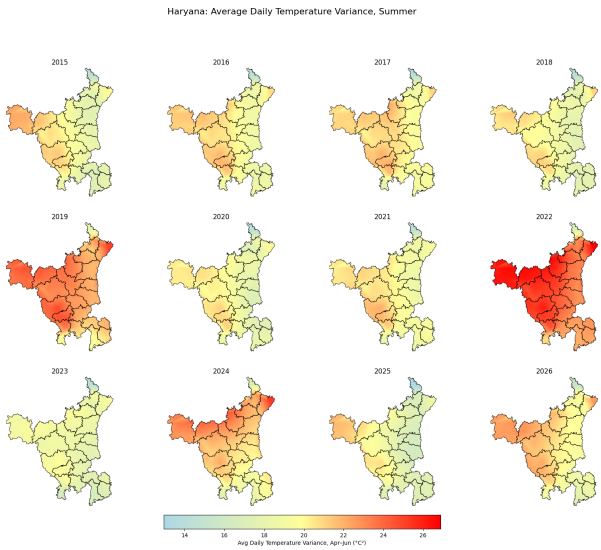
Zimmermann, Laura (2024). "Why Do Workers Participate in Public Employment Programs? Evidence from India". In: *Economic Development and Cultural Change*.

APPENDIX

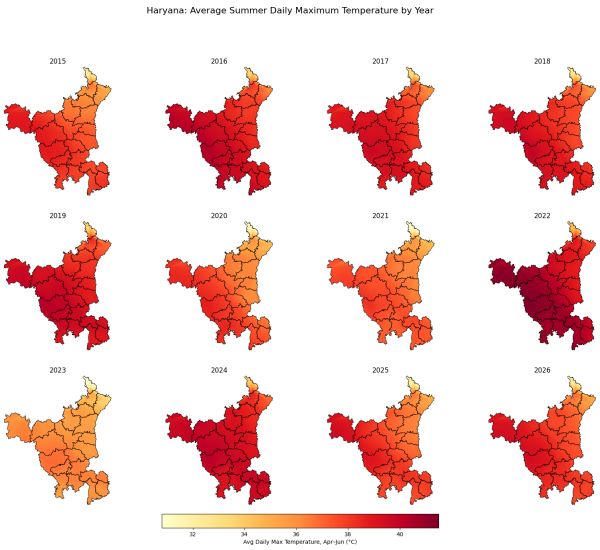
A Setting and policy



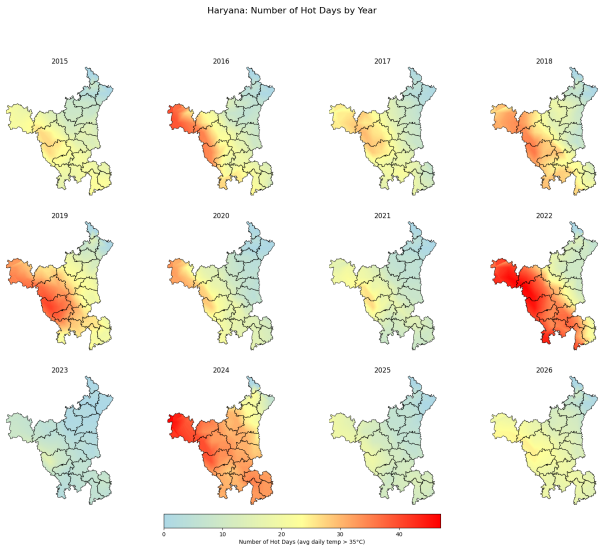
(a) Average summer temperature



(b) Average daily variation in summer temperatures



(c) Average daily maximum temperature



(d) Number of hot days

Figure A.1: Temperature patterns in Haryana, 2015–2026.

Note: The figure reports temperature measures based on ERA-5 hourly weather data for Haryana in the years 2015-2026. Summer refers to the months of April-June.

Table A.1 and Table A.2 give the statutory daily wage under NREGS in Haryana over the years and the changes to the scheme on the introduction of VB–G RAM G, 2025.

Table A.1: Statutory daily wage for unskilled rural work in Haryana

						In the muster rolls	
Financial year	Rate (Rs.)	Change	In force from	Under	Rank	Modal rate	At that rate (%)
2021-22	315		1 Apr 2021	MGNREGA s. 6(1)	Highest		
2022-23	331	+16 (5.1%)	1 Apr 2022	MGNREGA s. 6(1)	Highest	331	86.1
2023-24	357	+26 (7.9%)	1 Apr 2023	MGNREGA s. 6(1)	Highest	357	73.5
2024-25	374	+17 (4.8%)	1 Apr 2024	MGNREGA s. 6(1)	Highest	374	73.1
2025-26	400	+26 (7.0%)	1 Apr 2025	MGNREGA s. 6(1)	Highest	400	69.3
2026-27	409	+9 (2.2%)	1 Jul 2026	VB–G RAM G	2nd	400	59.8

The rate is set centrally for each state and financial year, indexed to the Consumer Price Index for Agricultural Labourers. Until 2025–26 it was notified under section 6(1) of the Mahatma Gandhi National Rural Employment Guarantee Act, 2005 and took effect on 1 April. The 2026–27 rate is the first under the Viksit Bharat—Guarantee for Rozgar and Ajeevika Mission (Gramin) Act, 2025, which commenced on 1 July 2026 and introduced a national floor of Rs. 300; Haryana was the highest-paying state in every year to 2025–26 and is second to the high-altitude gram panchayats of Sikkim (Rs. 450) in 2026–27.

The last two columns are what the register records, not what was notified. Haryana pays by measured output against a task, so the notified rate is a ceiling: it is the maximum daily rate observed in every year, and the share reaching it is the share of person-days on which the full task was completed. That the two agree in every overlapping year is a check on both. The exception is 2026–27, where the register stops on 30 June 2026, the day before the Rs. 409 rate took effect, and so still shows the old ceiling.

This study's fieldwork ran from 10 July to 1 September 2026, which is entirely under the new Act and its Rs. 409 rate. The choice experiment was designed earlier and used Rs. 400, the 2025–26 rate, as the status quo, so its reference wage is Rs. 9 below the statutory rate in force during fieldwork.

Sources, all retrieved 4 September 2026. (a) Outlook India, "Centre Notifies Revised MGNREGA Wage Rates", 24 March 2023. (b) ThePrint, "Centre increases MGNREGA wages by 3–10%", 28 March 2024. (c) GKToday, "Revised MGNREGA Wages", 31 March 2022. (d) FactChecker.in, "All Poll-Bound States See Hike In MGNREGA Wages, Except Kerala", 2021. (e) Adda247 Current Affairs, "MGNREGS Wage Hike for 2025–26", March 2025. (f) The Tribune, "Govt notifies enhanced wage rates under VB-G RAM G Act, fixes base floor wage at Rs 300", July 2026.

Table A.2: The rural employment guarantee before and after July 2026

	MGNREGA, 2005	VB–G RAM G, 2025
Assent	23 August 2005	21 December 2025
In force for this study	To 30 June 2026	From 1 July 2026
Entitlement		
Guaranteed days per rural household	100 per financial year	125 per financial year
Employment	On demand, within 15 days	On demand, within 15 days
Unemployment allowance	If work is not provided in 15 days	Retained
Seasonal suspension of works	None	Up to 60 days a year in peak agricultural seasons
Wages		
Rate set by	Centre, s. 6(1), indexed to CPI–AL	Centre, annual indexation on an interim base rate
National floor	None	Rs. 300 per day
Haryana rate	Rs. 400 (2025–26)	Rs. 409 (2026–27)
Range across states	Rs. 241 to Rs. 400 (2025–26)	Rs. 300 to Rs. 450 (2026–27)
Money		
Wage cost	Centre pays in full	Shared 60:40 centre to state
north-eastern and Himalayan states	Centre pays in full	Shared 90:10
Material cost	Centre 75, state 25	Within the same shared allocation
Budget basis	Demand-driven, funded as demand arises	Normative allocation set centrally; states meet any excess
Allocation, 2026–27	—	Rs. 95,692 crore
Works and oversight		
Permissible works	A broad statutory list	Four themes: water security, rural infrastructure, livelihood infrastructure, extreme-weather mitigation
Planning alignment	Gram panchayat plans	Also the PM Gati Shakti National Master Plan
Council	Central Employment Guarantee Council	Central Gramin Rozgar Guarantee Council; national steering and state committees
Monitoring	Muster rolls, social audit	Biometric attendance, geospatial monitoring, algorithmic fraud detection, weekly disclosure

The Viksit Bharat—Guarantee for Rozgar and Ajeevika Mission (Gramin) Act, 2025 was introduced in the Lok Sabha on 16 December 2025, passed by both Houses on 18 December, received assent on 21 December, and came into force on 1 July 2026, replacing the Mahatma Gandhi National Rural Employment Guarantee Act, 2005.

Fieldwork ran from 10 July to 1 September 2026, so all interviewed workers were employed under the new Act. Administrative muster rolls end on 30 June 2026 and fall entirely under the old Act. The institutional shift is also relevant: MGNREGA provided demand-driven employment, whereas the new Act operates within a capped allocation that states must supplement.

Sources (retrieved 4 September 2026): PRS Legislative Research, bill track for the VB–G RAM G Bill, 2025; Outlook Business, "VB-G RAM G Bill vs MGNREGA"; The Tribune, "Govt notifies enhanced wage rates under VB-G RAM G Act"; and News on AIR. MGNREGA figures are from the Act and wage notifications cited in Table A.1.

B Administrative data

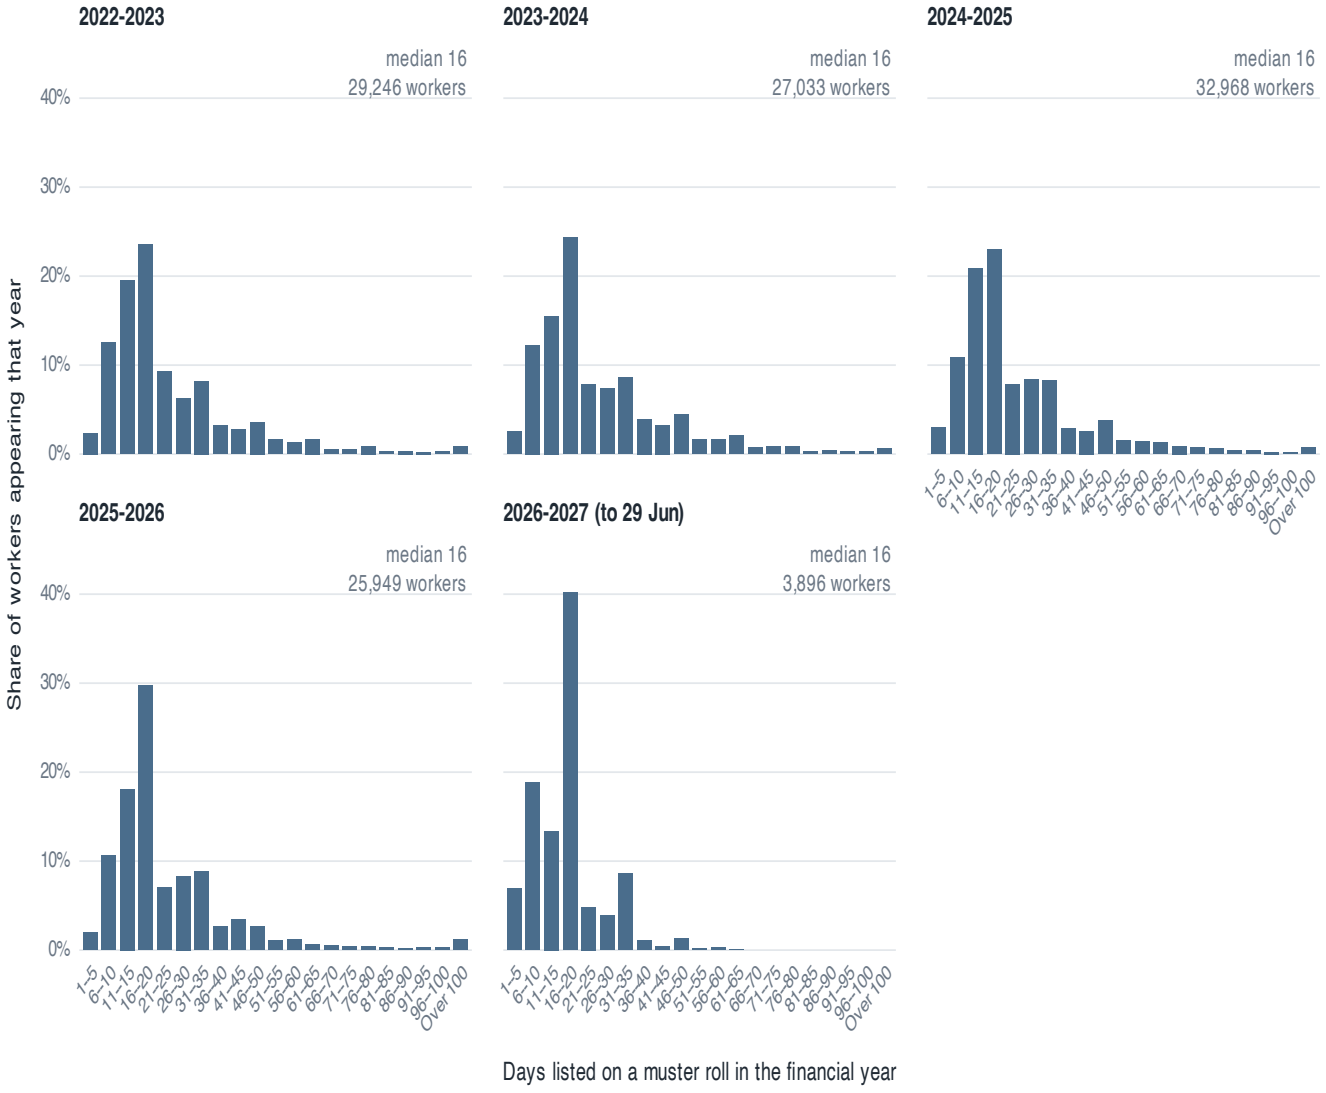


Figure A.2: Distribution of worker listed on NREGA muster rolls 2022-2026

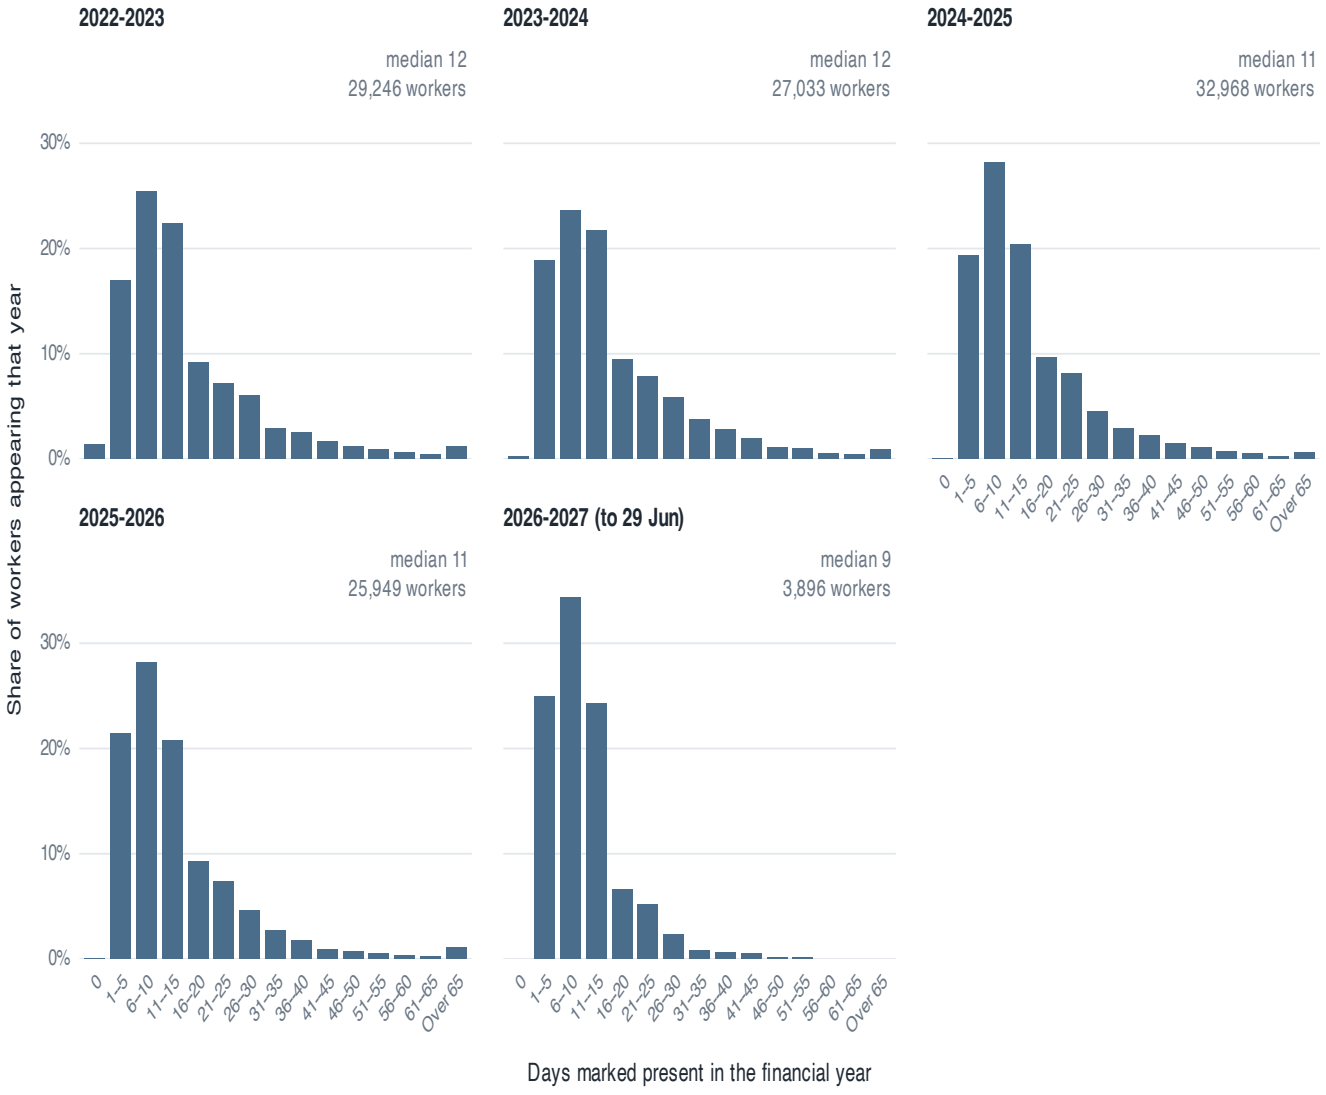


Figure A.3: Distribution of worker present on NREGA muster rolls 2022-2026

C Heat and absence and Assessed Wages

Table A.3: Heat and the Assessed Wage: Roll-Level and Worker–Worksite Estimates

	(1)	(2)	(3)	(4)
		All twelve months		March to June
Worksite effect	Yes	No	No	No
Worker × worksite effect	No	Yes	Yes	Yes
Month-of-year effect	No	No	Yes	Yes
A. Reaching the notified rate				
Wet bulb, daily mean	0.0015	0.0029	0.0057	0.0092
	[0.0014]	[0.0019]	[0.0049]	[0.0080]
Wet bulb, daily maximum	0.0010	0.0022	0.0017	0.0088
	[0.0015]	[0.0020]	[0.0048]	[0.0078]
Dry bulb, daily maximum	-0.0002	0.0003	-0.0010	0.0038
	[0.0011]	[0.0014]	[0.0025]	[0.0034]
Mean of the outcome	0.7850	0.7850	0.7850	0.7493
B. Assessed rate over the notified rate				
Wet bulb, daily mean	0.0006	0.0009**	0.0006	0.0017
	[0.0004]	[0.0004]	[0.0018]	[0.0031]
Wet bulb, daily maximum	0.0005	0.0008*	0.0002	0.0021
	[0.0004]	[0.0004]	[0.0018]	[0.0028]
Dry bulb, daily maximum	0.0001	0.0003	-0.0002	0.0006
	[0.0003]	[0.0003]	[0.0009]	[0.0012]
Mean of the outcome	0.9660	0.9660	0.9660	0.9609
C. Shortfall in rupees per person-day				
Wet bulb, daily mean	-0.19	-0.29*	-0.27	-0.58
	[0.13]	[0.15]	[0.67]	[1.18]
Wet bulb, daily maximum	-0.15	-0.26	-0.12	-0.70
	[0.15]	[0.16]	[0.66]	[1.06]
Dry bulb, daily maximum	-0.02	-0.09	0.03	-0.20
	[0.11]	[0.13]	[0.34]	[0.46]
Mean of the outcome	12.64	12.64	12.64	14.75
Unit	Roll	Worker-roll	Worker-roll	Worker-roll
Observations	93,398	249,350	249,350	74,369
Worker-worksite cells with two or more rolls	–	98,322	98,322	30,993
of which identify the notified rate	–	18,184	18,184	4,874
Gram panchayats	1,432	1,432	1,432	1,141
Rainfall and size controls	Yes	Yes	Yes	Yes
Wet-bulb variation surviving (%)	50.9	32.8	10.1	12.7

Columns (1)–(3) use the full-year sample; column (4) uses the March–June sample. Column (1) is estimated at the roll level and includes worksite fixed effects. Columns (2)–(4) use worker-roll observations and include worker-by-worksite fixed effects. Each roll receives a total weight of one, divided equally across its worker-roll observations. Outcomes and heat exposure are measured at the roll level; temperature is averaged over all days the roll ran, not only days the worker attended. The March–June sample retains rolls with at least one day in the window.

All specifications include calendar-year fixed effects, rainfall controls, and controls for work size. Month-of-year fixed effects are included in columns (3) and (4). Worker–worksite cells count cells with two or more rolls; the next row reports the subset that identifies the notified-rate outcome. The final row reports the share of raw wet-bulb variation remaining after the specified fixed effects.

In the worker–worksite specifications, comparisons are within the same worker and worksite across rolls. The median interval between rolls within a cell is 38 days; 1.9 percent of cells have all their rolls within a week. Cells with multiple rolls have mean wet-bulb temperature of 22.01, compared with 22.58 elsewhere, and workers present for 19.1 days, compared with 7.0 days; these cells account for 44.8 percent of paid person-days.

Standard errors clustered at the gram panchayat level are in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.4: Absence and accumulated heat by caste category, March to June

	(1) All	(2) Scheduled Caste	(3) Others
A. The 10-day mean wet bulb, per degree			
Wet bulb, 10-day mean (C)	0.0057***	0.0070***	0.0046***
	[0.0016]	[0.0021]	[0.0015]
Difference from Others, pooled		0.0005	—
		[0.0006]	
B. Quartiles, cut within the gram panchayat			
Q2	0.0003	0.0013	-0.0012
	[0.0043]	[0.0052]	[0.0043]
Q3	0.0068	0.0060	0.0074
	[0.0053]	[0.0063]	[0.0054]
Q4 hottest	0.0127**	0.0163**	0.0092
	[0.0055]	[0.0068]	[0.0056]
Panchayat × month, year, worker, day of week	Yes	Yes	Yes
Rainfall at the same horizon	Yes	Yes	Yes
Gram panchayat clusters	1,168	1,052	1,148
Workers	114,218	58,265	55,953
Observations (person-days)	1,936,710	951,316	985,394
Mean of the dependent variable	0.1891	0.2194	0.1599

Column (4) of Table 1 estimated separately on each caste category. The category is the one printed on the muster roll; a worker whose rolls disagree is assigned the value appearing on most of them, which is 0.00 per cent of workers. Categories with fewer than 10,000 person-days or 30 gram panchayats are not estimated.

Caste does not vary within a worker, so it cannot be estimated alongside the worker fixed effect and the table is separate regressions rather than interactions. Column (1) is the pooled sample of the groups shown.

The difference row is the interaction of the heat term with the category in one pooled regression, Others omitted, and is the test to read. The groups share gram panchayats, so a difference formed from the separate standard errors above would treat correlated estimates as independent.

Panel B keeps the quartiles cut on each panchayat's full March to June distribution rather than re-cutting them within each group, so the bins mean the same thing in every column.

Standard errors clustered on the gram panchayat in brackets. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.5: Heat and the assessed wage within the worksite, by type of work

	All twelve months		March to June	
	(1)	(2)	(3)	(4)
Wet bulb, daily mean				
Canals, drains and ponds	0.0027	0.0097*	-0.0036	0.0172*
	[0.0018]	[0.0058]	[0.0048]	[0.0098]
Roads and land levelling	-0.0011	-0.0006	-0.0224***	-0.0069
	[0.0022]	[0.0087]	[0.0068]	[0.0118]
Plantation	-0.0002	0.0022	0.0026	0.0162
	[0.0015]	[0.0049]	[0.0049]	[0.0104]
All work types alike, p	0.485	0.895	0.401	0.731
Wet bulb, daily maximum				
Canals, drains and ponds	0.0021	0.0042	-0.0070	0.0154
	[0.0020]	[0.0057]	[0.0056]	[0.0098]
Roads and land levelling	-0.0017	-0.0029	-0.0237***	-0.0029
	[0.0024]	[0.0090]	[0.0074]	[0.0108]
Plantation	-0.0002	0.0050	0.0033	0.0169
	[0.0019]	[0.0058]	[0.0060]	[0.0114]
All work types alike, p	0.606	0.921	0.499	0.607
Dry bulb, daily maximum				
Canals, drains and ponds	0.0003	-0.0006	-0.0068**	0.0045
	[0.0013]	[0.0029]	[0.0034]	[0.0042]
Roads and land levelling	-0.0017	-0.0024	-0.0124***	0.0046
	[0.0019]	[0.0050]	[0.0046]	[0.0052]
Plantation	0.0001	0.0065	0.0026	0.0071
	[0.0014]	[0.0041]	[0.0038]	[0.0054]
All work types alike, p	0.785	0.400	0.546	0.232
Worksite	Yes	Yes	Yes	Yes
Calendar year	Yes	Yes	Yes	Yes
Month of year	No	Yes	No	Yes
Rainfall over the days the roll ran	Yes	Yes	Yes	Yes
Works	6,960	6,960	3,708	3,708
Gram panchayats	1,432	1,432	1,141	1,141
Mean of the dependent variable	0.7847	0.7847	0.7485	0.7485
Canals, drains, ponds (N)	66,058	66,058	30,405	30,404
Roads & land levelling (N)	25,171	25,171	8,548	8,547
Plantation (N)	1,719	1,719	669	669

Notes: Table 2 is estimated separately by work type, as recorded in the work code. Groups require at least 500 March–June rolls and 30 works and are held constant across sample windows. Because work type is constant within a worksite, it is absorbed by worksite fixed effects; the table reports separate regressions. Differences across groups are tested using pooled regressions with heat-by-group interactions.

The outcome is an indicator for reaching the notified rate. The unit is the muster roll; each work contains about 13 rolls and spans 1.18 calendar years on average. All columns include worksite and calendar-year fixed effects; columns (2) and (4) also include month-of-year fixed effects. March–June samples retain rolls with at least one day in the window without truncating their days or weather measures.

Controls include rainfall averaged over roll-days and standardised within gram panchayat and calendar month, and log days worked and log workers. Standard errors clustered at the gram panchayat level are in brackets. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.6: Absence and heat: which reading of the weather, March to June

	(1) Wet bulb, 10-day	(2) Wet bulb, same day	(3) Night wet bulb	(4) Night, 10-day	(5) Dry bulb mean, 10-day
A. The measure in levels, per degree					
Per degree	0.0057***	0.0005	0.0008	0.0056***	0.0022**
	[0.0016]	[0.0009]	[0.0008]	[0.0015]	[0.0010]
B. Quartiles, cut within the gram panchayat					
Q2	0.0003	-0.0039	-0.0043	0.0010	0.0061
	[0.0043]	[0.0035]	[0.0039]	[0.0045]	[0.0040]
Q3	0.0068	-0.0012	0.0002	0.0090*	0.0060
	[0.0053]	[0.0041]	[0.0041]	[0.0054]	[0.0046]
Q4 hottest	0.0127**	-0.0004	0.0004	0.0109*	0.0063
	[0.0055]	[0.0044]	[0.0048]	[0.0056]	[0.0052]
C. Both horizons in one regression, per degree					
Same day	-0.0008				
	[0.0009]				
10-day mean	0.0061***				
	[0.0016]				
Rainfall at the same horizon	Yes	Yes	Yes	Yes	Yes
Panchayat × month, year, worker, day of week	Yes	Yes	Yes	Yes	Yes
Gram panchayat clusters	1,168	1,168	1,168	1,168	1,168
Observations (person-days)	1,936,710	1,936,710	1,936,710	1,936,710	1,936,710
Mean of the dependent variable	0.1891	0.1891	0.1891	0.1891	0.1891

Every column is the specification of column (4) of Table 1 with the heat measure replaced, on the one sample that table uses. The night measure is the 21:00 to 05:00 window, which crosses midnight and is indexed by the morning it ends on, so it is the night immediately before the working day and carries no further lag.

Panel A enters each measure in levels. Panel B cuts its quartiles within the gram panchayat on that measure's own March to June distribution, with the coolest quartile omitted.

Each column carries one rainfall term at its own horizon: same-day rainfall with a same-day or overnight heat term, the 10-day mean with a 10-day heat term.

Panel C puts the same-day and the 10-day wet bulb in one regression. The window excludes the day itself, so the two do not overlap by construction; what they share is the persistence of the weather, and the log reports their correlation. A night measure that correlates with its daytime counterpart at close to one is one measurement read twice and cannot be given a night interpretation, whatever its coefficient.

Standard errors clustered on the gram panchayat in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.7: Absence and accumulated heat: the length of the window, March to June

	(1) 3-day	(2) 5-day	(3) 10-day	(4) 15-day	(5) 20-day
A. The window mean, per degree					
Wet bulb, window mean (C)	0.0033***	0.0044***	0.0057***	0.0064***	0.0058***
	[0.0012]	[0.0013]	[0.0016]	[0.0019]	[0.0022]
B. Quartiles, cut within the gram panchayat					
Q2	0.0022	0.0066*	0.0003	0.0015	0.0053
	[0.0040]	[0.0037]	[0.0043]	[0.0041]	[0.0045]
Q3	0.0037	0.0064	0.0068	0.0084	0.0112**
	[0.0050]	[0.0051]	[0.0053]	[0.0051]	[0.0056]
Q4 hottest	0.0102*	0.0150**	0.0127**	0.0154***	0.0154***
	[0.0058]	[0.0059]	[0.0055]	[0.0058]	[0.0059]
Days required of the window	3	4	7	11	14
Rainfall at the same horizon	Yes	Yes	Yes	Yes	Yes
Panchayat × month, year, worker, day of week	Yes	Yes	Yes	Yes	Yes
Gram panchayat clusters	1,168	1,168	1,168	1,168	1,168
Observations (person-days)	1,936,710	1,936,710	1,936,710	1,936,710	1,936,710
Mean of the dependent variable	0.1891	0.1891	0.1891	0.1891	0.1891

The specification of column (4) of Table 1 with the accumulation window varied. Ten days is the horizon Somanathan et al. report and is the one this paper adopts; the others are shown so that the choice can be seen to be a choice. A result that appears at ten days and at no neighbouring length is a result about the window rather than about heat.

Each window requires values on at least seventy per cent of its days, so the requirement is proportional rather than fixed; the ten-day column keeps the seven-of-ten rule used elsewhere. The day itself is outside every window. Rainfall enters at the same horizon as the heat term in each column.

The windows overlap heavily with one another by construction, so the columns are not independent readings and the pattern across them should be read as a shape rather than as a set of separate tests.

Standard errors clustered on the gram panchayat in brackets. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.8: Absence and accumulated heat: wind and radiation at the 10-day horizon

	(1) Table 1, column (4)	(2) + net radiation	(3) + wind speed	(4) + radiation and wind
A. The 10-day mean wet bulb, per degree				
Wet bulb, 10-day mean (C)	0.0057***	0.0057***	0.0055***	0.0055***
	[0.0016]	[0.0016]	[0.0016]	[0.0016]
Net radiation, ten-day (SD)		-0.0058***		-0.0059***
		[0.0018]		[0.0018]
Wind speed, ten-day (SD)			0.0010	0.0012
			[0.0014]	[0.0014]
Rainfall, ten-day (SD)	0.0064***	0.0017	0.0065***	0.0019
	[0.0017]	[0.0022]	[0.0017]	[0.0022]
B. Quartiles, cut within the gram panchayat				
Q2	0.0003	0.0008	0.0003	0.0008
	[0.0043]	[0.0043]	[0.0043]	[0.0043]
Q3	0.0068	0.0073	0.0065	0.0069
	[0.0053]	[0.0053]	[0.0053]	[0.0053]
Q4 hottest	0.0127**	0.0123**	0.0121**	0.0115**
	[0.0055]	[0.0055]	[0.0056]	[0.0056]
Net radiation at the same horizon	No	Yes	No	Yes
Wind speed at the same horizon	No	No	Yes	Yes
Rainfall at the same horizon	Yes	Yes	Yes	Yes
Panchayat × month, year, worker, day of week	Yes	Yes	Yes	Yes
Gram panchayat clusters	1,153	1,153	1,153	1,153
Workers	113,908	113,908	113,908	113,908
Observations (person-days)	1,934,179	1,934,179	1,934,179	1,934,179
Mean of the dependent variable	0.1891	0.1891	0.1891	0.1891

Column (4) of Table 1 with net downward shortwave radiation and wind speed added as controls. Both enter at the horizon of the heat term, as the rainfall control does, and both are standardised within gram panchayat and month of the year, so their coefficients are per standard deviation.

The two controls are available on 99.9 per cent of the sample of Table 1, so every column here is estimated on that subsample rather than on the full one. Column (1) is the main specification re-estimated on the subsample, which separates what the controls do from what the restriction does; on the full sample the same coefficient is 0.0057 with a standard error of 0.0016.

Panel B keeps the quartiles cut on each panchayat's full March to June distribution. They are not of equal size on this subsample, which is deliberate: re-cutting them would change the bins as well as the controls.

Radiation and the wet bulb are both driven by the clear-sky day, and wind is what makes a given wet bulb bearable, so these are the two controls under which an accumulated heat coefficient is most likely to be something else. The coefficient of interest is read across the columns, not within one.

Standard errors clustered on the gram panchayat in brackets. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.9: Heat and the assessed wage within the worksite, in temperature bands

	All twelve months		March to June	
	(1)	(2)	(3)	(4)
Dry bulb, daily maximum — omitted: 30 to 35 C				
Below 20 C	-0.0157	0.0242	0.0124	-0.3232**
	[0.0243]	[0.0480]	[0.0926]	[0.1302]
20 to 25 C	-0.0326	0.0024	0.1112*	-0.0794
	[0.0206]	[0.0384]	[0.0613]	[0.0787]
25 to 30 C	-0.0265	-0.0325	0.0686	-0.0839
	[0.0248]	[0.0297]	[0.0513]	[0.0683]
35 to 40 C	-0.0761***	-0.0630**	-0.0876**	-0.0175
	[0.0232]	[0.0280]	[0.0360]	[0.0351]
40 C and above	-0.0031	-0.0053	-0.0305	0.0202
	[0.0234]	[0.0299]	[0.0331]	[0.0332]
All bands jointly zero, p	0.037	0.158	0.004	0.149
Wet bulb, daily mean — omitted: 25 to 27 C				
Below 15 C	-0.0103	-0.1060*	-0.0109	-0.2848***
	[0.0200]	[0.0603]	[0.0694]	[0.0986]
15 to 20 C	0.0108	-0.0846	0.1153**	-0.1497*
	[0.0232]	[0.0541]	[0.0467]	[0.0856]
20 to 25 C	-0.0253	-0.0401	-0.0373	-0.0397
	[0.0210]	[0.0300]	[0.0365]	[0.0407]
27 C and above	0.0404**	0.0129	0.0056	-0.0089
	[0.0201]	[0.0223]	[0.0346]	[0.0360]
All bands jointly zero, p	0.029	0.421	0.012	0.077
Wet bulb, daily maximum — omitted: 28.5 to 30.5 C				
Below 18.5 C	-0.0144	-0.0599	-0.0301	-0.2769***
	[0.0232]	[0.0516]	[0.0759]	[0.1036]
18.5 to 23.5 C	0.0009	-0.0418	0.1368***	-0.1138
	[0.0254]	[0.0450]	[0.0456]	[0.0749]
23.5 to 28.5 C	-0.0124	-0.0028	-0.0255	-0.0381
	[0.0220]	[0.0256]	[0.0319]	[0.0385]
30.5 C and above	0.0297	0.0226	0.0010	-0.0154
	[0.0305]	[0.0328]	[0.0411]	[0.0401]
All bands jointly zero, p	0.534	0.732	0.006	0.127
Worksite	Yes	Yes	Yes	Yes
Calendar year	Yes	Yes	Yes	Yes
Month of year	No	Yes	No	Yes
Rainfall and size controls	Yes	Yes	Yes	Yes
Works	6,960	6,960	3,708	3,708
Gram panchayats	1,432	1,432	1,141	1,141
Mean of the dependent variable	0.7847	0.7847	0.7485	0.7485
Observations	93,398	93,398	39,734	39,732

Notes: Temperature enters as the share of roll-days in each band; coefficients are relative to the omitted band in each panel heading. The outcome is an indicator for reaching the notified rate. Joint tests assess the bands collectively. The unit is the muster roll. Each work contains about 13 rolls and spans 1.18 calendar years on average. All columns include worksite and calendar-year fixed effects; columns (2) and (4) also include month-of-year fixed effects. March–June samples retain rolls with at least one day in the window without truncating their days or weather measures. Controls include rainfall averaged over roll-days and standardised within gram panchayat and calendar month, and log days worked and log workers. Standard errors clustered at the gram panchayat level are in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.10: Heat and the assessed wage within the worksite, with radiation and wind

	(1)	(2)	(3)	(4)
	All twelve months		March to June	
Month of year held	No	Yes	No	Yes
A. On the rolls that carry both controls				
Wet bulb, daily mean	0.0020	0.0087*	-0.0069*	0.0116
	[0.0016]	[0.0050]	[0.0040]	[0.0080]
Wet bulb, daily maximum	0.0016	0.0040	-0.0095**	0.0113
	[0.0017]	[0.0050]	[0.0045]	[0.0080]
Dry bulb, daily maximum	0.0001	-0.0004	-0.0062**	0.0051
	[0.0012]	[0.0026]	[0.0028]	[0.0035]
B. With solar radiation and wind speed				
Wet bulb, daily mean	0.0021	0.0094*	-0.0069*	0.0116
	[0.0016]	[0.0050]	[0.0040]	[0.0081]
Wet bulb, daily maximum	0.0017	0.0056	-0.0096**	0.0112
	[0.0017]	[0.0053]	[0.0046]	[0.0082]
Dry bulb, daily maximum	0.0003	0.0010	-0.0068**	0.0058
	[0.0013]	[0.0030]	[0.0030]	[0.0040]
Solar radiation (SD)	-0.0034	-0.0041	-0.0009	0.0013
	[0.0031]	[0.0032]	[0.0047]	[0.0046]
Wind speed (SD)	-0.0062**	-0.0056*	0.0014	0.0014
	[0.0030]	[0.0029]	[0.0041]	[0.0041]
Worksite	Yes	Yes	Yes	Yes
Calendar year	Yes	Yes	Yes	Yes
Month of year	No	Yes	No	Yes
Rainfall over the days the roll ran	Yes	Yes	Yes	Yes
Works	5,726	5,726	3,125	3,125
Gram panchayats	1,199	1,199	925	925
Mean of the dependent variable	0.7855	0.7855	0.7487	0.7487
Observations	81,389	81,389	35,492	35,491

Notes: The table adds solar radiation and wind speed to Table 2. Both are averaged over roll-days and standardised within gram panchayat and calendar month. Because these controls are missing for some rolls, Panels A and B use the same restricted sample; compare them within columns to distinguish sample changes from the effect of adding controls. Radiation and wind coefficients are reported but not interpreted.

The unit is the muster roll. Each work contains about 13 rolls and spans 1.18 calendar years on average. All columns include worksite and calendar-year fixed effects; columns (2) and (4) also include month-of-year fixed effects. March–June samples retain rolls with at least one day in the window without truncating their days or weather measures.

All specifications control for rainfall averaged over roll-days and standardised within gram panchayat and calendar month, and for log days worked and log workers. Standard errors clustered at the gram panchayat level are in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D Survey data

Table A.11: Characteristics of the surveyed worksites

Variable	N	Mean	SD	Median	Range
Size and attendance					
Workers interviewed at the site	72	22.4	9.5	25.0	1.0–43.0
Days the site was visited	72	1.15	0.40	1.00	1.00–3.00
Workers registered on the muster roll	65	65.6	45.5	53.0	10.0–245.0
Workers present on the day of the visit	65	38.0	23.3	32.0	10.0–116.0
Workers absent on the day of the visit	65	27.5	36.5	15.0	0.0–189.0
Attendance rate	65	0.660	0.230	0.680	0.055–1.000
Women present	65	28.2	19.0	22.0	0.0–93.0
Women as a share of those present	65	0.746	0.178	0.794	0.000–0.962
Heat amenities					
Any shade at the site	65	86.2			
Drinking water available	65	89.2			
A shelter to rest in	65	30.8			
A first aid kit	65	61.5			
Amenities present (of 4)	65	2.68	1.05	3.00	0.00–4.00
Hours and output target					
Work starts (hour of the day)	65	9.16	0.87	9.00	6.00–10.50
Hours worked	63	7.10	1.56	7.50	4.00–9.50
Today's target equals the usual target	65	96.9			

$N = 72$ worksites in 17 districts. A worksite is a cluster of worker GPS fixes linked at 600 m. Rows from the mate interview have $N = 65$; 7 worksites never had a mate interviewed.

Table A.12: What MGNREGA workers already do about heat

	All	Women	Men	p
A. Measures taken at the worksite				
Drinks more water	96.5 (1569)	96.1 (1121)	97.4 (448)	0.181
Covers the head	76.7 (1247)	75.7 (883)	79.1 (364)	0.135
Takes breaks in shade	96.4 (1567)	95.8 (1117)	97.8 (450)	0.024
Works more slowly	42.1 (685)	39.3 (458)	49.3 (227)	<0.001
Goes home early	11.1 (180)	7.6 (89)	19.8 (91)	<0.001
Takes no measure	0.4 (7)	0.5 (6)	0.2 (1)	0.325
B. What is done when heat makes the worker unwell				
Rests in shade	90.2 (1466)	89.6 (1045)	91.5 (421)	0.229
Drinks water	90.1 (1465)	89.3 (1041)	92.2 (424)	0.062
Oral rehydration salts	12.3 (200)	11.7 (137)	13.7 (63)	0.296
Lemon water	49.8 (809)	46.1 (538)	58.9 (271)	<0.001
Wet cloth on the head	17.4 (283)	11.2 (131)	33.0 (152)	<0.001
Stops working	18.3 (298)	14.8 (172)	27.4 (126)	<0.001
Seeks medical help	32.8 (534)	27.1 (316)	47.4 (218)	<0.001
Does nothing	2.6 (43)	3.3 (39)	0.9 (4)	<0.001
C. Adjustments made at home				
Changed water or liquid intake	72.0 (1170)	68.7 (801)	80.2 (369)	<0.001
Changed food or drink	48.2 (781)	43.7 (507)	59.6 (274)	<0.001
Changed the sleeping schedule	37.3 (606)	32.4 (378)	49.6 (228)	<0.001
Changed where the household sleeps	33.0 (536)	32.6 (380)	33.9 (156)	0.611
Changed when household chores are done	41.8 (680)	38.5 (449)	50.2 (231)	<0.001
Takes medicine to prevent heat stress	25.8 (419)	24.2 (282)	29.8 (137)	0.024
D. Whether the worker is free to adapt				
Comfortable asking for a rest break when unwell	68.8 (1119)	69.0 (804)	68.5 (315)	0.853
sometimes comfortable	1.8 (29)	2.0 (23)	1.3 (6)	0.318
Decides own rest breaks on hot days	29.2 (474)	27.4 (319)	33.7 (155)	0.014
Decides when work stops for the day	9.5 (155)	10.5 (122)	7.2 (33)	0.029
Told the mate of feeling unwell in the last 7 days	33.3 (542)	34.6 (403)	30.2 (139)	0.090
E. What the worksite and the household provide				
Site has adequate measures for heat waves	49.8 (810)	50.9 (594)	47.0 (216)	0.148
Drinking water tap within the home	83.9 (1365)	85.1 (992)	81.1 (373)	0.058
F. Asked of women only				
Adequate sanitation and privacy at or near the site		22.5 (262)		
Would prefer a women-only rest area		59.0 (688)		
Household responsibilities increase on hot days		88.9 (1036)		
Heat affects the ability to work during menstruation		52.1 (607)		
Would accept earlier hours if safe transport were provided		96.7 (1127)		
would perhaps accept		0.4 (5)		
Feels unsafe travelling to the site early in the morning		5.4 (63)		
Respondents	1,626	1,166	460	

Each cell is the percentage answering in the stated way, with the count in parentheses, over those asked the question. The final column is the p-value on the difference between women and men.

Panels A and B are multiple-response questions and need not sum to one hundred. Panel F was asked of women only; the columns it cannot speak to are blank.

Table A.13: Household work reported by women workers

Variable	N	Per cent	Count
Time spent on household work			
Less than 2 hours	1166	10.4	121
2-4 hours	1166	65.4	762
4-6 hours	1166	18.0	210
More than 6 hours	1166	6.3	73
Four or more hours a day	1166	24.3	283
Household work and heat			
More household responsibilities on hot days	1166	88.9	1036
Has changed when household work is done	1166	38.5	449
Someone fills in while she is at the worksite			
Cooking	894	22.5	201
Cleaning	806	19.5	157
Laundry	836	16.3	136
Fetching water	430	23.7	102
Collecting fuel	928	12.4	115
Childcare	85	49.4	42
Livestock care	281	39.1	110
Farm work	35	14.3	5
Elder care	237	30.8	73
Tasks reported			
Tasks she reported doing (of 9)	1166	3.91 [1.97]	

Women only; Section D of the questionnaire is routed on gender. $N = 1166$ women, of whom 1166 answered the time-use question.

The fill-in rows are put only to a woman who reported doing the task, so N is each row's own denominator and the per cent is the share with someone to cover it. The final row reports a mean with its standard deviation in brackets.

E The discrete choice experiment

E.1 Design

The experimental attribute space comprises two start times, three output-target levels and five wage levels, giving thirty candidate profiles and, with the status quo held fixed, 435 possible choice tasks. The fielded design selects eight of these by D-optimal search conducted directly over pairs of alternatives rather than over individual profiles: row-level optimality concentrates on corner profiles that dominate one another, leaving no dominance-free pairing among the selected rows. Candidate pairs are restricted to those in which neither alternative is weakly better on every attribute, and the criterion is evaluated for the full three-alternative choice set under a utility-neutral prior, with a penalty on attribute-level imbalance. The number of tasks was set at eight after Monte Carlo simulation of the estimator under plausible preference parameters, which indicated that additional tasks bought little precision relative to the respondent burden they imposed. The realised design contains sixteen experimental alternatives across the eight tasks, comprising thirteen distinct profiles; start times are balanced eight-eight and the target and wage levels occur approximately equally often. No experimental alternative is dominated within its task. Evaluated under a null prior against the best of 20,000 randomly generated eight-task designs over the same attribute levels, its relative D-efficiency is 87.5 percent.

Because the status quo is the only alternative with a nine o'clock start and also the only alternative paying INR 400, the two features are perfectly associated. Three consequences follow. First, a status-quo constant is not separately identified from the start-time coefficient, since the constant equals $\mathbf{1}\{\text{start} = 9\}$ and the two together saturate a three-level attribute; the baseline therefore omits it. Second, the association inflates the measured collinearity in the estimating sample. Across the sixteen experimental alternatives the correlation between start time and wage is 0.278, with variance inflation factors of 1.16, 1.15 and 1.21 across the three attributes, but across all twenty-four alternatives it is 0.731. This is a property

of the status quo rather than of how the tasks were assembled, and it is why we report the experimental-alternatives-only specification alongside the baseline. Third, option A carries the ten o'clock start in six of the eight tasks, though it is the higher-paying option in only three and the lower-target option in five, so any position effect to guard against is specific to the start time.

The experimental wage range lies entirely above the status-quo wage, so the design speaks to what workers must be paid to accept a later start and is silent on preferences over wages below the current arrangement. It does not, however, require extrapolation. Conditional on being the higher-paying alternative in its task, an experimental option is chosen by 59 percent of respondents at a premium of INR 75, 77 percent at INR 100 and 76 percent at INR 150, and the status quo is chosen on only 12.4 percent of choice occasions. The acceptance margin therefore lies inside the bid range and the implicit compensation value is an interpolation. We note that acceptance rates by raw premium level are not informative on their own, because the design assigns the two lowest premiums to the lower-paying alternative in every task and the two highest to the higher-paying one.

E.2 The comprehension screen

Respondents answered two comprehension questions before reaching the choice tasks. The first presented three options differing only in start time and asked which meant working later into the afternoon heat; the second presented three options differing only in wage and asked which paid most. An incorrect answer triggered one re-reading and one repeat, after which respondents who still answered incorrectly did not proceed. The screen removed a substantial and non-random share of the interviewed sample:

	Interviewed	Reached the tasks	Share
All respondents	1,626	929	57.1%
Men	460	292	63.5%
Women	1,166	637	54.6%
No schooling	717	317	44.2%
Primary (5th class)	530	332	62.6%
Secondary (10th class)	249	174	69.9%
Higher secondary or above	130	106	81.5%

Mean age is 40.0 among those who reached the tasks and 44.3 among those screened out. Those admitted also differ in caste, marital status, assets and dependence on MGNREGA, but not in heat symptoms or in whether a household member has fallen ill from heat, which is the dimension the experiment is about. Every respondent who reached the tasks completed all eight, so selection operates entirely at the screen rather than through attrition within the experiment. Table A.19 reports the full comparison. We do not reweight: the experiment identifies preferences among workers who understood it, and we make no claim about the preferences of those who did not.

E.3 Worksites and districts

The worker questionnaire records no location other than a GPS fix, so worksites are defined as clusters of worker fixes linked at 600 metres using fixes accurate to 100 metres or better. Interviews were placed at 72 worksites, 65 of which were matched to a mate interview; the median worksite contributed sixteen respondents. District is recorded in the mate survey and assigned to workers through the worksite, supplemented where the same-day match fails by a mate interviewed at the same worksite on another day, by the nearest worksite already carrying a district, and by days on which every located interview fell in one district. Holding each worksite out in turn and predicting its district from the others returns the correct district for 98.3 percent of worksites, and for every worksite within five kilometres of another.

District is missing for 4.2 percent of respondents. Of the seventeen districts covered, eight also appear in the administrative data, accounting for 53.8 percent of respondents with an assigned district; the survey and administrative samples overlap geographically but describe different workers.

E.4 Enumerators

Estimated preferences differ substantially across the twelve enumerators. Allowing the start-time coefficient to vary by enumerator improves the fit sharply ($\chi^2 = 194.1$, 22 degrees of freedom), and the improvement survives allowing the wage coefficient, and hence the scale of the logit error, to vary as well ($\chi^2 = 119.4$, 11 degrees of freedom), so this is not scale heterogeneity. Among the seven enumerators with at least sixty respondents, the implied compensation for an hour ranges from INR 3 to INR 24, a spread wider than any respondent-level contrast we report, and the lowest is not distinguishable from zero. An equivalent test by district is of similar magnitude ($\chi^2 = 181.8$, 34 degrees of freedom). Enumerators were assigned to districts rather than randomised within them, so interviewer and location effects cannot be separated here, and the enumerator rows of Table A.18 describe who administered the experiment rather than identifying an enumerator effect.

E.5 Tables

Table A.14: Mixed logit with a mixing distribution on the output target

	(1) Start step only	(2) Start and 10 pt	(3) Start and both levels	(4) 10 pt only
Eleven o'clock	-0.269	-0.269	-0.269	-0.257
	(0.041)	(0.041)	(0.041)	(0.038)
Ten point cut	0.544	0.544	0.544	0.524
	(0.055)	(0.055)	(0.055)	(0.054)
Twenty point cut	0.029	0.029	0.029	0.034
	(0.055)	(0.055)	(0.055)	(0.054)
Extra wage, INR 100	2.335	2.335	2.335	2.247
	(0.062)	(0.062)	(0.062)	(0.055)
s.d. eleven o'clock	0.458	0.458	0.459	
	(0.074)	(0.074)	(0.074)	
s.d. ten point cut		0.000	0.000	0.002
		(0.001)	(0.001)	(0.061)
s.d. twenty point cut			0.006	
			(0.201)	
Random coefficients	1	2	3	1
Log-likelihood	-2784.9	-2784.9	-2784.9	-2791.2
χ^2 against fixed coefficients	12.5	12.5	12.5	-0.0
nominal p	0.000	0.002	0.006	1.000
Standard error exists for every s.d.	Yes	Yes	Yes	Yes
Optimiser exit code	0	0	0	0

Mixed logits on the two experimental alternatives, fitted with 500 Halton draws and normal mixing distributions, with the wage coefficient fixed throughout. Column (1) is the specification reported in Table 4; the remaining columns add a mixing distribution on the output-target levels.

A standard deviation at the boundary of the parameter space leaves the information matrix singular, so no standard error exists; the table records where this occurs.

The chi-squared statistic compares each column with the fixed-coefficient conditional logit on the same choice set, with degrees of freedom equal to the number of random coefficients. The null places a standard deviation at zero, on the boundary of the parameter space, so the p-value is indicative.

The compensation row is $-\gamma_{10}/\beta_{\text{wage}}$ in rupees per day. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.15: Discrete choice experiment: the distribution of the start-time coefficient

	Normal	Log-normal
Eleven rather than ten o'clock	-0.269***	-0.089***
	(0.041)	(0.032)
Output target cut by ten points	0.544***	0.564***
	(0.055)	(0.056)
Output target cut by twenty points	0.029	0.051
	(0.055)	(0.057)
Wage, hundreds of rupees	2.335***	2.358***
	(0.062)	(0.060)
Dispersion of the start coefficient	0.458***	1.593***
	(0.074)	(0.210)
Share with a positive start coefficient	0.279	0 by construction
Compensation for an hour later, at the median	12	4
at the mean	12	13
Log-likelihood	-2784.9	-2762.2
Respondents	903	903

Both columns are mixed logits on the experimental pair with the start step as the only random coefficient and fixed target and wage coefficients, fitted with 500 Halton draws. The normal column is column (4) of Table 4.

A normal places mass on both signs and so implies a share of respondents who would start later for no extra pay; that share is reported. A log-normal holds the coefficient negative for everyone.

Under the log-normal the entry reported is the median of the fitted distribution, $-\exp(\mu)$, with its delta-method standard error, and the dispersion row is the standard deviation on the log scale. Median and mean coincide under a normal and diverge under a log-normal, so both are given.

The two columns are not nested, so log-likelihoods are compared descriptively. Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.16: Discrete choice experiment: the assumptions behind the conditional logit

Test	Statistic	df	p	Note
A. Independence of irrelevant alternatives				
Status quo chosen, share of tasks	0.124			
Hausman–McFadden, full set vs pair	56.89	3	<0.001	
Small–Hsiao, median over 25 splits	16.07	3	0.001	rejects in 84% of splits
Dissimilarity parameter λ	0.765***		<0.001	(0.035), $H_0: \lambda = 1$
likelihood ratio against $\lambda = 1$	68.22	1	<0.001	
compensation value, conditional logit	17			rupees per day
compensation value, nested logit	12			rupees per day
B. Restrictions tested against the saturated model				
start time linear, $\beta_{11} = 2\beta_{10}$	0.669	1	<0.001	(0.108)
target linear, $\gamma_{20} = 2\gamma_{10}$	-0.951	1	<0.001	(0.066)
start by target interaction	0.055	1	0.313	(0.055)
C. What identifies the start-and-wage tradeoff				
Tasks whose later start pays more	5 of 8			within a task, ±1
Pooled correlation with the wage	0.275			over all tasks
Ten-to-eleven step, task left out	7 to 19			11 rupees pooled
Ten-point cut, task left out	-35 to -7			-23 rupees pooled
D. Linearity and the nine o'clock contrasts				
Linear start, full set	-0.347***			17 rupees per hour
Linear start, pair only	-0.322***			14 rupees per hour
λ, attributes saturated	0.240			(0.059); 0.765 when linear
Ten-to-eleven step, that nest	12			11 without the nest
Nine-to-ten step, that nest	-229			42 without the nest
Linear start, task left out	-0.561 to -0.212			-0.347 pooled
Respondents	929			
Choice sets	7,432			

Panel A tests independence of irrelevant alternatives. The Hausman–McFadden comparison restricts the estimator to the two experimental options, which is column (1) of Table 4; where the difference of the two variance matrices is not positive definite a generalised inverse is used and the degrees of freedom are the rank employed. The Small–Hsiao statistic is the median over 25 random halves of the respondents, with the share of splits rejecting at five per cent.

The nested logit places the two experimental options in one nest and the status quo in a degenerate nest of its own. λ equals one under IIA and below one if the two experimental options are closer substitutes for each other than for the status quo.

Panel B tests each restriction against the saturated model, whose estimates are column (2) of Table 4.

Panel C concerns what identifies the tradeoff between the start time and the wage. Each task offers two experimental options that differ on the start, the target and the wage at once, so within a task the eleven o'clock indicator and the wage are perfectly correlated; the tradeoff is identified only by the sign of that correlation changing across tasks. The eight tasks came in one fixed order, so position and profile cannot be separated and no test of order effects is attempted.

Panel D is why the estimates table does not lead with the linear specification on the full choice set. Linearity averages a nine-to-ten step, which carries the value of the hour together with whatever makes the status quo the status quo, and a ten-to-eleven step, which does not; the restriction is rejected in Panel B. Saturating the attributes inside the nest shows the two cannot both be estimated: λ falls and the nine o'clock level stops being interpretable, while the ten-to-eleven step and the target levels are unchanged. Compensating differentials are ratios, so they are invariant to the scale the nest changes.

On the pair there is only one start contrast, so linearity in the start time binds on nothing there; the per-hour figure still moves because linearity in the target is imposed at the same time and the two attributes are correlated across the profiles. Column (1) of Table 4 leaves both saturated.

Standard errors clustered on the respondent. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.17: The price of an hour and of a lower target, by respondent group

		An hour later		A ten-point cut	
Group	Respondents	Rupees	(SE)	Rupees	(SE)
All respondents	903	11	(2)	-23	(2)
Women	623	14	(2)	-21	(2)
Men	280	7	(3)	-28	(4)
Secondary schooling or more	271	11	(3)	-29	(4)
Less than secondary	632	12	(2)	-21	(3)
Scheduled Caste or Tribe	572	10	(2)	-24	(3)
Other castes	331	14	(3)	-22	(3)
Durables above the median	360	15	(3)	-23	(3)
Durables at or below the median	543	9	(2)	-24	(3)
Landless	790	11	(2)	-24	(2)
Owns land	113	12	(5)	-16	(7)
Three or more heat symptoms	593	10	(2)	-26	(2)
Fewer than three heat symptoms	310	15	(3)	-18	(4)
Women: four or more hours of household work	167	10	(4)	-33	(4)
Women: fewer than four hours	456	15	(2)	-17	(3)
MGNREGA is the main household income	373	7	(3)	-25	(3)
MGNREGA is not the main income	530	14	(2)	-22	(3)
Aged under 40	460	15	(2)	-24	(3)
Aged 40 and over	443	8	(2)	-22	(3)
Aged under 30	129	12	(4)	-29	(5)
Aged 30 to 39	331	16	(3)	-23	(4)
Aged 40 to 49	262	8	(3)	-23	(4)
Aged 50 and over	181	8	(4)	-20	(5)
Aged 30-39, three or more symptoms	218	12	(3)	-27	(4)
Aged 30-39, fewer than three	113	24	(5)	-16	(7)
Aged 50+, three or more symptoms	118	2	(5)	-18	(5)
Aged 50+, fewer than three	63	18	(7)	-23	(8)

Each row is a separate conditional logit on that subgroup in the specification of column (1) of Table 4, not an interaction. Raw coefficients are not comparable across subgroups because the logit scale is not identified separately from them; the rupee figures are, because the scale cancels.

The first pair of columns is the compensation required to start at eleven rather than ten o'clock. The second is the amount the worker would give up for the first ten-point cut in the output target, so a negative entry is a willingness to pay.

The household-work rows are women only, since that section of the questionnaire is routed on gender, and compare women with women.

Subgroups with fewer than 60 respondents are not estimated and are listed in the run log.

Standard errors clustered on the respondent.

Table A.18: Discrete choice experiment: robustness of column (1)

Specification	$\beta_{11 \text{ vs } 10}$	An hour later	A ten-point cut	Respondents	Sets
Baseline (column 1 of Table 4)	-0.257***	11	-23	903	6,514
	(0.039)	(2)	(2)		
Clustered on the worksite	-0.266***	12	-23	893	6,445
	(0.040)	(2)	(2)		
First choice task dropped	-0.235***	10	-22	898	5,685
	(0.042)	(2)	(2)		
By enumerator					
Enumerator 3	-0.461***	24	-13	158	1,122
	(0.102)	(5)	(6)		
Enumerator 6	-0.241***	8	-34	122	929
	(0.093)	(3)	(5)		
Enumerator 7	-0.258**	11	-13	98	744
	(0.119)	(5)	(5)		
Enumerator 2	-0.219**	12	-33	94	660
	(0.110)	(6)	(7)		
Enumerator 1	-0.079	3	-31	92	700
	(0.122)	(5)	(7)		
Enumerator 4	-0.286**	15	-23	89	603
	(0.123)	(6)	(7)		
Enumerator 5	-0.326**	10	-26	88	625
	(0.144)	(4)	(6)		

Each row re-estimates column (1) of Table 4. $\beta_{11 \text{ vs } 10}$ is the coefficient on an eleven rather than a ten o'clock start; the rupee columns are the compensation required for that hour and the amount the worker would forgo for a ten-point cut.

Worksites are clusters of worker GPS fixes; respondents with no identified worksite are dropped from that row.

Enumerators were not randomly assigned to worksites, so these rows describe who administered the experiment and do not identify an enumerator effect. Of the 12 who administered it, the 7 with at least 60 respondents are shown.

Standard errors clustered on the respondent except where stated. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.19: Characteristics of those who reached the choice tasks

Variable	Reached the tasks	Did not	p
Age (years)	40.0	44.3	<0.001
Female	68.6 (637)	75.9 (529)	0.001
Married	88.4 (821)	84.2 (587)	0.017
Household size	5.3	5.2	0.422
Any schooling	65.9 (612)	42.6 (297)	<0.001
Secondary schooling or more	30.1 (280)	14.2 (99)	<0.001
Scheduled Caste or Tribe	63.2 (587)	69.2 (482)	0.012
Durable goods owned (of 11)	4.10	3.75	<0.001
Landless	87.5 (813)	89.5 (624)	0.205
MGNREGA is the main household income source	41.0 (381)	34.0 (237)	0.004
Heat symptoms, last 7 days	4.05	4.22	0.306
A household member fell ill from heat	30.2 (281)	25.3 (176)	0.025
Has a chronic illness	7.1 (66)	9.2 (64)	0.133
Needed a second explanation of the tasks	52.9 (491)		
Interviewed at a worksite with a mate survey	94.1 (874)	95.7 (667)	0.139
Hour of the interview (IST)	12.0	12.1	0.033

929 of 1,626 respondents reached the choice tasks, 57.1 per cent, and every one of them completed all eight.

Two comprehension questions gate the module. Of the 697 who did not reach the tasks, 612 stopped at the first question and 85 at the second. The second-explanation row is defined only for those admitted.

Continuous variables report means, binary variables percentages with the count in parentheses. The final column is a two-sided Welch test.