NBER WORKING PAPER SERIES

HOT WEATHER, UNDERNUTRITION, AND ADAPTATION IN RURAL INDIA

Paul Stainier
Manisha Shah
Alan Barreca

Working Paper 34047
http://www.nber.org/papers/w34047

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2025

We would like to thank conference participants at UCLA, UCSB, and the WEAI 2023 for their valuable comments and feedback. We thank Michael Prelip and members of the Wang-Prelip research group for helpful feedback. We thank Teevrat Garg, Vis Taraz, Maulik Jagnani, and Supreet Kaur for sharing data and code. For questions or comments please contact Paul Stainier at paul.stainier@sp2.upenn.edu. The authors have nothing to disclose. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Hot Weather, Undernutrition, and Adaptation in Rural India
Paul Stainier, Manisha Shah, and Alan Barreca
NBER Working Paper No. 34047
July 2025
JEL No. O13, Q12, Q54

ABSTRACT

We examine the impact of temperature during the growing season on household diets in the subsequent year in rural India, a setting with a high prevalence of small family farms. High growing season temperatures reduce crop yields, which would presumably reduce incomes and home-grown food for consumption. However, household adaptation could mitigate how the reductions in yields affect diets. We find that heat increases the number of strongly undernourished households in the subsequent year, as measured by the consumption of calories, iron, zinc, thiamine, and niacin. We also find suggestive evidence that households adapt to heat-induced losses of home-grown crops by purchasing more food.

Paul Stainier
University of Pennsylvania
paul.stainier@sp2.upenn.edu

Manisha Shah
University of California, Berkeley
Goldman School of Public Policy
Department of Economics
and NBER
ManishaShah@berkeley.edu

Alan Barreca
University of California, Los Angeles
and NBER
abarreca@ioes.ucla.edu

A data appendix is available at http://www.nber.org/data-appendix/w34047

1 Introduction

While existing studies show that high temperatures harm crop yields in a variety of real-world settings (Schlenker and Roberts, 2009; Taraz, 2018), the ways in which these temperature-related crop losses subsequently affect household nutrition are not as well documented. These impacts may be especially important in low-income countries where a substantial number of people rely on home-grown crops for both sustenance and income. Crop losses could translate into higher rates of undernutrition, which has important downstream welfare costs to long-term health (Sharma et al., 2020), human capital accumulation (Georgieff, 2020), psychological well being (Pailler and Tsaneva, 2018), and labor productivity (Thomas et al., 2006). The magnitude of the relationship between extreme heat and undernutrition, even with a clear link between temperature and crop yields, is unclear given the scope for households to adapt. Households could decrease the amount of calories they eat or consume cheaper, less nutritious calories. Households might also seek work in non-agricultural sectors (Colmer, 2021), spend down savings, or decrease non-food purchases to mitigate impacts. Our study estimates the relationship between temperature and household undernutrition, and adds new insights into adaptation strategies, using data from India.

We use survey responses from over 300,000 households in rural India to estimate the impact of extreme heat on food consumption choices, diet quality, and labor allocation. We measure food consumption using India's National Sample Survey (NSS), a nationally representative repeated cross-section of Indian households, from 2003 to 2012 (the last year that includes detailed food consumption over a 30 day period). The NSS includes questions on the consumption of roughly 150 food items, which we convert into household-level consumption of calories and six additional nutrients. In past studies using the NSS, calories and some macronutrients (usually protein and fat) feature as the standard measures of nutrition (Deaton and Drèze, 2009; Carpena, 2019). Such a limited set of nutritive measures omits many nutrients crucial to dietary health. For example, iron deficiency is endemic to rural India: India's National Family Health Survey 4 (NFHS-4), finds that 59 percent of children below 5, 58 percent of breastfeeding women, and 53 percent of all women were anemic, caused in part by iron deficiency (International Institute for Population Sciences (IIPS) and ICF, 2017).[1] During pregnancy, iron deficiency is linked with increased maternal mortality, preterm births, low birthweight, and decreased human capital accumulation in children.[2] In adults, it lowers energy levels, labor productivity, and mental health (Marcus et al., 2021; Thomas et al., 2006). In order to more comprehensively assess diet quality, we consider calories along with the following set of nutrients: protein, iron, zinc, thiamine, niacin,

[1]Anemia, when a person lacks healthy red blood cells for the transport of oxygen, is a common clinical proxy measure for iron deficiency. Anemia has multiple causes, including iron deficiency, genetic factors, gastrointestinal diseases such as hookworm, postpartum hemorrhage, and folate or B12 deficiencies. This multitude of factors, and their frequent overlap, makes it difficult to assess exactly how much anemia is caused by iron deficiency (Chaparro and Suchdev, 2019).

[2]Breastmilk is low in iron, so newborn infants rely on iron stored during the fetal stage for the proper development of organs, including the brain (Georgieff, 2020).

and riboflavin.

Our identification strategy leverages the facts that high temperatures increase crop losses over a long growing season and that these losses could harm household welfare subsequent to harvest. Our treatment variable is the frequency of days in a given temperature bin, e.g. days between 90 and 100°F, for the growing season (June-December). We focus on how this year's growing season temperature impacts consumption in the subsequent year, which is likely when households either consume or sell their production. By focusing on nutritional outcomes in the subsequent year, our estimates abstract from heat's more contemporaneous impacts on health or labor productivity. The unit of observation for this variable is the district and year. To address potential confounders, our key control variables include district and year fixed effects, along with state specific time trends. The core estimates are identified by unpredictable temperature swings across growing seasons for a given district. We control for temperatures in this year's non-growing season and the subsequent year's growing and non-growing seasons. We also control for rainfall shocks, which can have a meaningful effect on crop yields and food consumption (Carpena, 2019), for both this year and the subsequent year. We show that the results are robust to including alternative time trends (e.g., district specific or state quadratic) as well.

There are three main findings from our research. First, we find an increase in the number of households that qualify as strongly undernourished in terms of most of the nutritive measures we consider. For example, one day above 110°F increases the fraction of households consuming less than 80 percent of the recommended caloric intake by 0.36 percentage points, an increase equivalent to 3.1 million people. On average, households in our sample experienced 0.6 days above 110°F per growing season, with 10 percent of households experiencing over 2 such days. Interestingly, we find no statistically significant impact of extreme heat on aggregate consumption. Instead, the impacts are at the 20^{th} percentile of the caloric intake distribution, suggesting important distributional implications. The effects of extreme heat are similarly concentrated at the lower end of the intake distribution for iron, zinc, thiamine, and niacin. Note that given important intra-household inequalities in food allocation (D'souza and Tandon, 2019), these findings may miss the effects of extreme heat on undernutrition for certain at risk groups. For example, female children tend to receive a low share of resources (Kaul, 2018; Azam and Kingdon, 2013), and this share may change (to be higher or lower) during periods of financial distress.

Second, we find suggestive evidence that households adapt to crop losses by purchasing more food. We estimate our model for calories and each nutrient, separately for home-grown sources and purchased sources. We find that, in most cases, heat during the growing season causes a decline in home-grown consumption but a commensurate increase in purchased consumption, helping to explain the lack of aggregate effects. These purchases constitute an *ex-post* adaptation, or a response to the effects of a weather shock after it has occurred (Carleton et al., 2024), which is similar to consumption smoothing across periods.

Third, we find evidence that adults re-allocate their labor away from agricultural activities in the household and towards non-agricultural work, both at home and outside of the home, after a very hot growing season. One day above 110°F in a growing season results in a 0.26 percentage point decrease in the percentage of adults working in agriculture at home in the next year, and a 0.12 percentage point increase in non-agricultural work outside of the household. This finding suggests that some food purchases might be due to new income sources. However, it remains possible that households may also be cutting into savings or sacrificing non-food purchases or investments (e.g. human capital) to buy food.

This paper makes three main contributions. First, we add to a growing literature on heat's impact on welfare in rural developing contexts. While heat's impacts on yields might be well documented, crop losses are difficult to translate into societal welfare losses without understanding how households respond. We consider the impacts of weather shocks on home-grown and purchased foods separately, something that is largely absent from the literature but important in rural contexts (Jayasinghe et al., 2017). We show that households purchase more foods in response to heat shocks. In addition, we show that in the year after a hot growing season, household members dedicate less time to agriculture in the household and more to non-agricultural work. In doing so, we complement the work of Colmer (2021), who finds contemporaneous effects of heat on the allocation of labor between agricultural and non-agricultural sectors, but without specifying the setting (in home or not) of the work.

Second, we capture important distributional effects. Past papers on nutrition tend to focus on aggregate measures, such as total calorie consumption or diet composition (e.g., percentage of calories from fruits vs. grains). By using low-consumption thresholds (e.g., consumption below 80 percent of recommended levels of calories or nutrients), we capture important distributional impacts that aggregate measures miss.

Third, we make methodological contributions to measuring diet quality. Our work emphasizes the need to include micronutrients and minerals when estimating climate change's impact on diet quality. Although it is particularly important in rural India, iron rarely features in past economics studies on diet quality in the area.[3] The NSS provides tables to convert food quantities into calories and macronutrients such as protein and fat. In this project, we use the raw nutritional data that the NSS tables reference, allowing us to directly measure the consumption of important micronutrients and minerals such as iron, zinc, and thiamine (Gopalan et al., 1989). Our focus on micronutrients and minerals adds to previous work that shows that calorie consumption and diet composition are affected by droughts in Mexico (Hou, 2010) and India (Carpena, 2019) and rainfall and temperature shocks in Nigeria (Dillon et al., 2015) and Tanzania (Randell et al.,

[3] See for example Deaton and Drèze (2009); Carpena (2019) who focus mostly on caloric intake and macronutrients such as protein or fat.

2022).

2 Data

Nutritional outcomes

The nutritional outcomes and socioeconomic household characteristics come from the National Sample Survey (NSS), a nationally representative repeated cross-section of Indian households. We use Schedule 1, which includes detailed household-level consumption data, and Schedule 10, which includes information on household members' primary activities, such as employment or school attendance. We construct nutritional outcomes using 8 rounds of the data (rounds 59-64, 66, and 68), spanning 2003 to 2012 (National Sample Survey Office, 2003-2012).[4] All rounds included in the analysis dataset ask households about consumption for all food items over a 30 day period. No rounds between 2013 and 2021 include household food consumption questions. We omit the most recent round of data, covering 2022 and 2023, from the analysis because it only asks households about food consumption over the last 7 days for many food items (including animal products, fruits and vegetables). Food consumption measured over a shorter recall period (e.g., 7 days) is often inflated relative to a longer recall period (e.g., 30 days) (Scott and Amenuvegbe, 1990; Mukherjee and Chaudhury, 2020).

Each household response records the consumption of roughly 150 food items, separated into home-grown and purchased foods, over the prior 30 days. In addition, it includes the number of meals that 1) household members eat outside of the home and that 2) non-household members eat in the home. Following past work, we use these numbers to adjust our food consumption figures (NSS, 2014). Note that the survey does not measure exactly how much food was eaten, as some could be lost during the cooking process or wasted after preparation. It also does not indicate who, within the household, consumed specific foods. These consumption measures therefore indicate household-level food availability for intake, rather than intake itself. The survey also contains all household members' basic demographic information. The dataset includes 314,414 households, with a total of 1,599,495 people, across 575 district-areas that we construct to be geographically consistent over the entire time frame.

In order to generate the main outcomes of interest, we translate the detailed food consumption data from the NSS into household-level consumption of calories and nutrients. For example, for each food item (e.g., wheat or bananas), we multiply the quantity (weight or number) of the food consumed by its caloric content per unit. We then sum up across all food items for each household to calculate a total amount of calories consumed. We repeat this process for six nutrients: protein, iron, zinc, thiamine, niacin, and riboflavin. The NSS provides tables listing

[4] Rounds 65 and 67 do not include household food consumption questions.

each food item's caloric and protein values. The NSS-provided tables do not include other nutrients, resulting in them being excluded from previous analyses (Deaton and Drèze, 2009; Carpena, 2019).[5] These NSS-provided tables, however, are based on a dataset from Nutritive Value of Indian Foods (Gopalan et al., 1989), which does include information on these additional nutrients for most food items. In some cases (e.g., iron for 14 percent of food items and zinc for 35 percent of food items), nutrients are missing from Gopalan (1989). In these cases, we use information from the Food Composition Tables for Bangladesh (Shaheen et al., 2013) that build on Gopalan et al. (1989). We use tables from these original sources to create the daily per capita consumption of each nutrient, and take their log as they are right-skewed and roughly log normal. We winsorize all per-capita outcomes at the 1 and 99 percent levels.

We generate additional dependent variables per nutritional factor to estimate the distributional impacts of extreme heat on food consumption. We match the NSS data with age- and gender-specific dietary requirements from Nutrient Requirements and Recommended Dietary Allowances For Indians, and then aggregate requirements to the household-level (ICMR, 2010). We calculate adequacy percentages by dividing each household's consumption by its recommended intake. For example, the calorie adequacy percentage is the total household calorie consumption divided by the recommended calorie intake over a 30 day period. Using these percentages, we create dummy variables indicating whether household-level caloric or nutrient consumption is below 100 percent of the recommended levels. We create analogous dummy variables for thresholds ranging from 50 percent to 150 percent.

We use these threshold dummy variables to assess how extreme heat affects households at different points in the nutrient adequacy distribution. In our preferred definition, we label a household below the 100 percent threshold as experiencing "undernutrition," one below the 80 percent threshold as experiencing "strong undernutrition", and one below the 60 percent threshold as experiencing "extreme undernutrition." Our results are robust to using other thresholds for strong and extreme undernutrition (e.g., 70 percent or 50 percent, respectively). An increase in the percentage of households below the 100 percent threshold indicates households being pushed from adequate dietary availability to undernutrition, and an increase in the percentage of households below the 80 percent threshold indicates households being pushed from undernutrition to strong undernutrition. Because each household only answers the survey once, after treatment, we are unable to investigate distributional effects by separating households into groups by baseline food consumption levels. Note that, as we only have household-level food consumption we are unable to tell exactly who, within the household, is consuming sub-adequate levels of calories or nutrients. For example, the children in a given household may have been adequately nourished while the parents may have been the ones to forgo food after crop losses.

[5] One notable exception is Behrman and Deolalikar (1987).

Table 1 shows average levels of nutrient consumption across the households in the sample (for the full distribution of nutrient consumption adequacy, see Appendix Figure A1). On average, households consume adequate (or higher) levels of calories and nutrients, with the exception of iron and riboflavin. For example, the average per capita daily caloric consumption is around 2,181 kCal/person-day, roughly the recommended amount of caloric intake, and the average calorie adequacy percentage is 104 percent. However, many households still experience strong undernutrition, including in the case of calories. Twenty point four percent of households consume less than 80 percent of the recommended level of calories. Iron consumption, an important contributing factor to anemia, which is prevalent in India (International Institute for Population Sciences (IIPS) and ICF, 2017), is especially low: 54.7 percent of households consume less than 80 percent of their recommended level of iron. These statistics illustrate two main points. First, focusing on average levels of food consumption masks significant heterogeneity in diet quantity and quality across households. Second, focusing only on calories would lead one to miss how low the consumption of other nutrients, especially iron, is in rural India.

Table 1 also highlights variation in food consumption levels across different regions of India (see the table notes for details on the designation of these regions). For example, households in the North of India consume higher levels of calories than any other region, with only 10.8 percent experiencing strong caloric undernutrition. In contrast, 25.3 percent of households in the West, and 27.0 in the South, experience strong caloric undernutrition. These differences are even starker in the case of iron: 18.1 percent of households are strongly iron undernourished in the North, compared to 82.1 percent in the South and 91.2 percent in the Northeast.

The NSS separates foods into home-grown and purchased items. We use this distinction to construct 2 more dependent variables per nutrient: levels of consumption from home-grown foods versus from purchased foods. Around 65 percent of households consume some amount of home-grown foods, and about 27 percent of iron (4.16 mg/person-day) and 28 percent of calories (611 kCal/person-day) come from home-grown foods.

Employment Outcomes

We use data from 6 rounds[6] of the NSS Survey's Schedule 10, covering the period 2004 to 2012 (National Sample Survey Office, 2004-2012), to study the impact of a hot growing season on individual household members' principal activities (e.g., employed, unemployed, or in school). Following Colmer (2021), we limit the sample to 1,150,280 working-age individuals, from 14 to 65 years old. We create four primary outcome variables: dummy variables indicating whether an individual's principal activity is employment in an agricultural sector or in a non-agricultural sector, separately for employment in a household enterprise (e.g., self-employment or as an unpaid

[6] Of the rounds we use for the nutritional outcomes, 59 and 63 do not include Schedule 10.

family worker) or as a wage laborer outside of the household. We present summary statistics for these outcomes in Appendix Table A2. Of working-age individuals, 32.3 percent are employed in agriculture (22.5 percent in the household and 9.8 percent outside of it) and 21.8 percent are employed in non-agricultural sectors (9.8 percent in the household and 12.0 percent outside of it). The remainder of working-age individuals are either unemployed or out of the labor force.

Weather Data

We use weather data from the ERA5 reanalysis model (Copernicus Climate Change Service, 1979-2012), which includes hourly temperature and precipitation across a 31km by 31km grid covering the entire globe (Hersbach et al., 2020). For each day and gridpoint, we calculate the maximum temperature and total precipitation. Following Garg et al. (2020), we separate weather into growing season (June to December) and non-growing season (March to May) periods. We assign temperature to each district-year-season as follows. For each gridpoint-year-season, we count the number of days within distinct 10 degree temperature bins ranging from below 70°F to above 110°F. We then average these temperature bins across all the gridpoints within each district. For the 18 small districts that do not contain any ERA5 gridpoints, we assign them the nearest ERA5 grid-point to their respective centroid.

Table 1 summarizes the average number of growing season days within each temperature range and the rainfall shock variable. It shows that, on average, India is hot: households in our sample experience an average of 109 days between 80 and 90°F, 51 days between 90 and 100°F, 8 days between 100°F and 110°F, and 0.6 days above 110°F during the seven month growing season. There is also significant across-region variation: the North of India experiences an average of 18 days between 100 and 110°F, while the East only experiences 5, and the Northeast experiences 0.01. Figure 1 maps the number of days per growing season with maximum temperatures between 100°F and 110°F (Panel A), and above 110°F (Panel B). It further emphasizes the vast climatic diversity across India: the number of days per year with maximum temperature between 100°F and 110°F ranges from 0 to above 70. About 25 percent of households in the sample reside in districts with no growing season days above 100°F over our study period. We also construct a "rainfall shock" control variable, which we describe in the Appendix, as is done in Shah and Millett Steinberg (2017).

Agricultural Yields

We use agricultural data from the Village Dynamics in South Asia Meso dataset from the International Crops Research Institute for the Semi-Arid Tropics (ICRISAT, 2015). The dataset includes price-weighted yield in tons/hectare for the six major crops (rice, wheat, sugarcane, groundnut, sorghum, and maize), and the five major monsoon crops (excludes wheat), at the district-year level. Specifically, this is a weighted average of the yields of individual crops, with individual crop prices as weights. We focus on the years 2000 through 2011 to match up with our NSS analysis.

3 Empirical Strategy

Nutritional Outcomes

Our main empirical strategy for nutritional outcomes is as follows:

$$N_{hdsmy} = \sum_{g=0}^{1}\sum_{t=0}^{1}\sum_{j=1}^{J} \beta^{j}_{y-t,g} T^{j}_{d,y-t,g} + \sum_{g=0}^{1}\sum_{t=0}^{1} \alpha_{y-t,g} P_{d,y-t,g} + \delta X_h + \gamma_d + \tau_y + \eta_m + \sigma_{sy} + \epsilon_{hdsmy} \quad (1)$$

where N is a nutritional outcome for household h in district d in state s surveyed during month m of year y. $T^{j}_{d,y-t,g}$ is the number of days in 10 degree temperature bin j in district d during year $y-t$'s growing ($g=1$) and non-growing ($g=0$) seasons, running from 70 to 80°F to above 110°F. We exclude the bin below 70°F, so all coefficients can be interpreted relative to a day below 70°F. P is the rainfall shock for this year and last year, separated by growing season. X_h is a set of household-level controls, including religion, caste, head of household education, and household composition by age group and gender. We include these controls to account for potential changes in the composition of the sample (as it is a repeated cross-section) that are correlated with heat shocks. When the outcome variable is a nutrient adequacy threshold (e.g., below 100 percent of recommended levels of calories), we replace household composition with total household size, as the nutrient adequacy threshold already explicitly includes household composition.

District fixed effects (γ_d) account for district-level differences in nutrition and temperature. For example, the North is both warmer and has higher levels of nutrition (see Table 1), so without these fixed effects our estimates would pick up a spurious correlation that shows that heat is associated with better nutritional outcomes. Year fixed effects (τ_y) account for any yearly India-wide differences in temperature that may be correlated with differences in nutrition, and month fixed effects (η_m) improve precision by controlling for India-wide seasonality in nutrition. We include state-specific yearly trends (σ_{sy}) to control for any state-specific changes in nutrition that may be spuriously correlated with warming over time. Our results are robust to district-specific yearly trends and state-specific quadratic trends as well. ϵ_{hdsmy} is a household-specific error term. The standard errors are clustered at the district level to account for any correlation between the error terms of households in the same district. For example, all households in a district with a more iron-rich staple crop (such as wheat) might experience bigger swings in iron consumption due to heat compared to households in other districts.

In the Results section, we focus on the effect of last year's growing season temperature (the

$\beta^j_{y-1,1}$) on this year's nutritional outcomes. We do so because weather during the growing period (June-December) likely could only affect household outcomes via crop losses in the period after the growing season (Garg et al., 2020). Weather might have contemporaneous effects, so we include this year's weather (both growing and non-growing seasons) and last year's non-growing season weather as controls.

Employment Outcomes

When estimating the effects of heat on labor outcomes, we use the same specification as in equation 1, but at the individual level. We remove head of household education and household composition controls, and instead add individual age, education, and gender controls.

Agricultural Yields

We also test whether hot weather impacts crop yields, the first step in the causal link from heat to nutritional outcomes. We match district-year level yields with district-year-season level weather variables for 1981 to 2011 to estimate the following model:

$$log(Y_{dy}) = \sum_{g=0}^{1}\sum_{j=1}^{J}\beta^j_g T^j_{d,y,g} + \sum_{g=0}^{1}\alpha_g P_{y,g} + \gamma_d + \tau_y + \epsilon_{dy}, \tag{2}$$

where Y is the yield for district d in year y. $T^j_{d,y,g}$ and $T^j_{d,y,ng}$ are the number of days in temperature bin j in district d during year y's growing ($g = 1$) and non-growing ($g = 0$) seasons. We include 10 degree temperature bins running from 70 to 80°F to above 110°F, and exclude the bin below 70°F. We also include a rainfall shock (P), district fixed effects (γ_d), and year fixed effects (τ_y). We cluster standard errors at the district level to account for any within-district correlation of the error terms. We estimate this regression for the yields of two different groups of crops: 1) the top six crops (rice, wheat, sugarcane, groundnut, sorghum, and maize), and 2) the top five monsoon crops (which excludes wheat from 1)).

4 Results

Yields

Figure 2 shows the effects of hot temperatures during the growing season on yields. Panel A shows results for the top six crops in India, and Panel B only includes monsoon crops. We multiply the log yield outcome variables by 100 so the coefficients can be interpreted as percentage changes. We find statistically significant decreases in overall yields for days above 80°F, with larger damages at higher temperatures. Panel A shows that a day above 110°F decreases yields for the top six crops by roughly 1.3 percent relative to a day below 70°F. This is 2.1 times larger than the effects of a day between 80 and 90°F. These results for the impact of growing season temperature are consistent with past literature (Carleton, 2017; Taraz, 2018; Garg et al., 2020). The top panel of Appendix Table A4, which corresponds to Figure 2, shows the full range coefficients for the effects of growing season weather on yields. The bottom panel shows the coefficients for the non-growing season. Heat above 70°F during the non-growing season negatively impacts crop yields, but most of these coefficients are statistically insignificant at conventional levels and they do not increase with higher temperatures as sharply as during the growing season.

Nutritional Outcomes

We estimate equation 1 to test how heat impacts caloric outcomes. Figure 3 Panel A plots the coefficient estimates from the previous year's growing season temperature, for the dependent variable log of daily calories per capita. The coefficients document the effect of one day in the corresponding temperature bin relative to a day below 70°F. We multiply the dependent variable by 100 so that the coefficients can be interpreted as percentage changes. Despite heat's impact on yields, heat in any range does not impact caloric consumption. In the Appendix, we consider the effect of heat on average consumption (Figure A2) across all nutrients. As in the case of calories, we do not find any evidence that extreme heat decreases average consumption for any nutrient or mineral.

Next, we consider whether extreme heat affects the number of households below key nutritional adequacy thresholds. In Panel B we graph coefficient estimates for the effect of the previous year's growing season weather on the percent of households below 100 percent calories adequacy. The graph shows there is no effect of heat on the percentage of households who consume 100 percent of their recommended level of calories. In contrast, we find that extreme heat does significantly increase the percentage of households in strong caloric undernutrition (Panel C) and extreme caloric undernutrition (Panel D). Figure 3 Panels C and D show that a day above 110°F pushes 0.36 percent of households into strong caloric undernutrition and 0.11 percent into extreme caloric undernutrition, respectively. These percentage changes correspond to an increase of around 3 million people and 1 million people per day above 110°F, respectively, out of rural India's 2012 population of 871 million (World Bank, 2024). These results suggest that extreme heat does exacerbate the diet quality of households already experiencing undernutrition, pushing many into strong or extreme undernutrition.

Next, we test whether households are pushed into strong undernutrition in terms of other nutrients after a hot growing season. In Figure 4, we plot the effect of last year's growing season temperatures on the percentage of households experiencing strong undernutrition for all of the nutrients we study. We find evidence that a day above 110°F in the prior growing season increases the percentage of households below the 80 percent adequacy threshold in the following nutrients: zinc, thiamine, and niacin. These results are statistically significant at the 0.05 level. We find suggestive evidence that households are pushed below this threshold in the case of iron, and no evidence that this is the case for protein or riboflavin. In Appendix Figure A3, we plot the effects of temperature on the percentage of households experiencing undernutrition (i.e., consumption below the 100 percent adequacy threshold). We show that for some outcomes (thiamine and niacin), heat between 70 and 80°F increases the percentage of households experiencing undernutrition. However, across all dietary outcomes, this effect does not increase with higher temperatures, and the coefficient of a day above 110°F is not statistically significant at the 0.05 level. These results suggest that the same distributional differences in the effects of a heat shock on calorie consumption hold true for several important nutrients.

In Appendix Figure A4, we consider results for the extreme undernutrition threshold. In most cases, the effect sizes are smaller than for the strong undernutrition results. This difference can be explained in part by the smaller amount of households near or below the 60 percent adequacy threshold in general (e.g., only 3.9 percent of households are below the 60 percent calorie adequacy threshold). However, in the cases of iron and thiamine, the effect sizes for extreme undernutrition are higher than those for strong undernutrition. This result emphasizes the importance of considering a range of measures for nutritional adequacy, as there can be heterogeneity across nutrients in baseline levels of adequacy and how a shock affects different degrees of undernutrition.

Adaptations

If households experience a shock to their home-grown sources of food, such as a local weather shock, they may respond by purchasing more food. These purchases would represent an *ex-post* adaptation, defined as responses to weather, and its effects, after it has occurred (Carleton et al., 2024). The net effect of heat shocks on undernutrition may also depend on *ex-ante* adaptations (e.g., choosing heat-resistant crops or moving out of agriculture based on the expectation of future heat shocks), though here we focus on *ex-post* adaptations. We test whether households purchase more food after a heat shock by estimating equation 1 for home-grown and purchased foods separately. We estimate these regressions as levels rather than logs since 35 percent of households consume no home-grown foods.

In Figure 5, we provide suggestive evidence for this adaptation mechanism. Each panel plots the coefficient estimates from the previous year's growing season from equation 1 for a different nutritive measure, separately by home-grown and purchased sources of food. The coefficients from Panel A suggest that households consume fewer home-grown calories as a result of hotter temperatures. They consume 2.5kCal per person per day fewer calories due to a single day between

100 and 110°F relative to one below 70°F, and 4.1kCal (slightly more than 1g of uncooked rice, where one serving is 45g) per person per day fewer for a day above 110°F. Neither coefficient is statistically significant at the 0.05 level. Panel A also shows that households purchase more calories after a hot growing season. A day between 100 and 110°F increases purchased calories by 1.8kCal per person per day and a day above 110°F increases them by 5.3kCal per person per day. The latter coefficient is statistically significant at the 0.05 level.

Panels B to G show suggestive evidence of a similar pattern for most nutrients we study. While the coefficients are not statistically significant at the 0.05 level, the point estimates for increases in purchases for a day above 110°F are generally close to those for decreases in home-grown consumption. Figure 5 points to an important yet understudied agricultural household adaptation to heat: households compensate for nutritional decreases from home-grown crop losses by purchasing more food.

These findings raise the question of how households pay for these additional food purchases after a hot growing season. One possibility is that after a hot growing season, households with lower home-grown crop yields allocate their labor away from agriculture (especially in the home) and towards non-agricultural sectors. We test this possibility using information on individuals' principal activity.

We estimate our individual regression on four outcomes for working-age individuals: working in agriculture in a household-run enterprise, working in agriculture as a wage laborer outside the household, working in non-agriculture in a household-run enterprise, and working in non-agriculture as a wage laborer outside the household. We plot the results of this analysis in Figure 6. Panel A shows that a day above 110°F leads to a 0.26 percentage point decrease in the percentage of working-age individuals working in agriculture at home. This corresponds to a 1.2 percent decrease in overall at-home agricultural employment. In Panel B, we show that there is no such decrease in agricultural employment outside of the home. Panels C and D plot the coefficients for households working in non-agricultural sectors in a household enterprise and outside the household, respectively. Both point estimates for the effect of a day above 110°F are near a 0.1 percentage point increase, but only the outside of the household coefficient is statistically significant at the 0.05 level. Together, these figures suggest that households that experience heat-induced losses in home-grown crop production adapt, at least in part, by spending less time working in home-grown agriculture and more time in non-agricultural work. This additional income may then be used to help purchase foods to make up for the losses in home-grown food sources.

Potential Mechanisms

Extreme heat could affect household diets through several different channels. The primary channel discussed in this paper thus far is that extreme heat's negative impact on yields (Taraz, 2018) decreases available food for households to consume or sell. However, other possible mechanisms include heat's effects on physiological well-being (Carleton et al., 2022), labor force participation and productivity (Somanathan et al., 2021), or fertility (Barreca and Schaller, 2020). By analyzing the difference in the effects of extreme heat based on its timing (i.e., during the growing season vs. the non-growing season), and considering additional outcomes, we test which mechanisms are most likely.

First, we test for the possibility that extreme heat pushes households into extreme undernutrition primarily due to its effect on crop yields. If crop yield reductions are a key mechanism, then extreme heat during the non-growing season, which we confirm does not affect yields (see Appendix Table A4), should not affect our nutritional outcomes. In Appendix Table A5, we show the coefficient estimates for last year's non-growing season weather for all of our nutritional outcome variables. None of the coefficients are statistically significant at the 0.05 level apart from 3 exceptions: 1) days above 90°F decrease the percentage of households below the 80 percent calories threshold, 2) days between 90°F and 100°F decrease the percentage of households below the 60 percent calories threshold, and 3) days between 80°F and 110°F decrease the percentage of households below the 80 percent thiamine threshold. The coefficients on the highest temperature bin (above 110°) are also generally closer to 0 than the coefficient estimates for last year's growing season. This suggests that extreme heat during the previous growing season is a more important determinant of household food consumption this year, consistent with the hypothesis that heat's impact on yields is driving the results.

We also present the coefficients for the current year's growing season weather from our main specification in Appendix Table A6. We find no effects of extreme heat in the current year, with one exception: a decrease in the percentage of households experiencing strong thiamine undernutrition. The lag in the effect, with heat this year impacting nutrition in the next year, is also consistent with the agricultural mechanism.

Second, we consider the possibility that heat in the prior period could affect household composition, in particular through increased mortality (Carleton et al., 2022). Heat-induced mortality could affect a household's nutritional adequacy in two opposing ways. On the one hand, it could decrease a household's productive capacity, decreasing the total amount of food for consumption. On the other hand, it could decrease the total amount of food necessary for the household. This latter effect may dominate the net effect on per capita consumption if older or less healthy household members, who are less likely to be able to contribute labor, are more susceptible to the mortality effects of heat.

We directly test the mortality channel with data from the National Family Health Survey (International Institute for Population Sciences (IIPS) and ICF, 2017, 2022). These data include responses to two household surveys conducted in 2015 to 2016 (round 4) and 2019 to 2021 (round 5). In each survey, the head of household is asked to list any deaths in the household over the last 4 years, and to note the year and age of death. We calculate mortality rates separately for 4 age groups: below 18, 18 to 39, above 40, and above 60. For each age group, district, and year, we divide the number of deaths by the total number of people in that age group across all households. Appendix Table A3 shows the average mortality rates for each age group over this period, which range from 2.3 deaths per thousand people (ages 18 to 39) to 35.9 deaths per thousand people (ages 60 and above). The analysis dataset covers the period from 2012 to 2019. As with the NSS data, we pair these mortality rates with district-level weather data from ERA5.

We estimate a regression with the dependent variables district-year mortality rates for different age groups. We include the temperature distribution from the entire year (as opposed to separately by growing and non-growing season), as extreme heat can impact current mortality at any time of the year. We plot coefficient estimates for the effects of temperatures on contemporaneous mortality rates in Figure 7. Panel A shows that we find no effect of extreme heat on mortality rates for children below 18. Panels B, C, and D, however, show suggestive evidence that extreme heat increases mortality for adults ages 18-39, above 40, and above 60. For example, a day above 110°F increases the mortality rate for adults at least 40 years old by around 0.5 percent. This coefficient is statistically significant at the 0.05 level. The coefficients for adults 18-39 and above 60 are similar in magnitude, but noisier. It is important to note that these coefficients represent percentage changes and therefore represent different numbers of increased deaths. For instance, the 0.8 percent increase for adults 18-39 represents 0.018 additional deaths per 1,000 people, but the 0.5 percent increase for adults above 60 represents 0.18 additional deaths per 1,000 people. The magnitudes we find are slightly smaller than the effect found in a previous study for the period 1957 to 2000 (Burgess et al., 2017). In Appendix Figure A7, we show the coefficients for the effects of extreme heat on mortality in the next year. The coefficients for the effect of a day above 110°F are smaller than in the case of this year's mortality, and none of them are statistically significant at conventional levels.

These findings may lend a partial explanation to our finding that, on average, per-capita consumption of calories and nutrients does not decrease. Specifically, if a heat wave that threatens crop yields also results in deaths, the net effect on per-capita food consumption could be closer to zero than a shock that only decreases yields without directly contributing to a decrease in the number of people in the household. However, it is also important to note that any increase in households experiencing strong or extreme undernutrition is also net of any heat-induced mortality in the prior year.

Third, we consider heat's effects on labor activity, which could impact income. Heat could lower the probability of employment due to decreased health or labor demand (Colmer, 2021), or it could decrease on-the-job productivity (Somanathan et al., 2021). The resulting reductions in labor income could then translate into decreased food consumption in the period after a heat shock. We directly test the labor force participation channel by estimating a regression, using the sample of working-age individuals, with the following dependent variable: a dummy variable indicating that someone is working.

In Appendix Figure A6, we show that a day above 110°F during the growing season results in a 0.16 percentage point decrease in the percentage of working-age individuals working in the same year (Panel B). This coefficient is statistically significant at the 0.05 level. The effect on work in the next year (Panel A) is a statistically insignificant decrease of 0.09 percentage points. Panels C-F show that the coefficients in Panels A and B are the net result of a decrease in workers in agriculture and an increase in workers in non-agricultural sectors. This reduction in agricultural labor could be driven by heat's effects on labor demand or on labor supply. On the one hand, heat-induced decreases in crop production could lead to decreases in agricultural labor demand. On the other hand, workers may decrease agricultural labor during the growing season to avoid on-the-job heat exposure, leading to lower home-grown crop production and an increase in food purchases in the next year. While we cannot rule out this possibility, air conditioning penetration in rural India during this period was effectively 0 (Davis et al., 2021), suggesting that the heat avoidance benefits of working in non-agriculture are likely limited. In addition, Colmer (2021) finds that extreme heat decreases agricultural wages in rural India, suggesting that the labor demand mechanism dominates the labor supply one. Regardless of the precise mechanism, the net impacts of heat on the labor market may also help explain some of the increase in households experiencing strong undernutrition.

Fourth, we consider the possibility of a change in a woman's current likelihood of being pregnant or raising an infant (Barreca and Schaller, 2020). We test the fertility channel by estimating regressions with the number of children in the household as the outcome. These regressions are identical to our main specification but do not include household composition controls. Appendix Figure A8 plots the coefficient estimates for the impact of last year's growing season temperature on the number of children present in a household. It shows no effect of heat on the number of children in a household, which helps rule out the fertility channel.

Robustness tests

We test the validity of our model through a series of robustness checks that we present in the Appendix. First, in Table A7 we show that our results do not change much when including district-month and year-month fixed effects. These fixed effects account for the possibility that each district, and each year, has a unique seasonality in nutritional outcomes. For example, wealthier districts may have a better ability to smooth consumption across the year, whereas poorer districts may see food consumption decrease as time from the most recent harvest increases. The coefficient for a day above 110°F increases for most outcomes, but the other temperature coefficients are largely the same. This consistency is unsurprising, because our identifying variation is driven by yearly, not monthly, variations in weather at the district level.

Second, we show robustness to specifications that include alternative trends. Replacing state-year linear trends with district-year linear trends (Table A8) or state-year quadratic trends (Table A9) results in qualitatively similar coefficients. Table A10 reports results from a first-difference estimator with the data collapsed to the district-year, which also generates similar results. Third, we test the effect of removing this year's weather as a control. Table A11 shows that we find similar results qualitatively to our main specification, though many of the coefficients are smaller. One explanation for this difference is the correlation between temperatures across years, which could lead to bias in the coefficients without controlling for this year's weather. Fourth, we show that our results are robust to the inclusion of different rainfall variables (Tables A12, A13, and A14).[7]

Fifth, we show that we find qualitatively similar results when considering different thresholds for strong or extreme undernutrition (e.g., 70 percent or 50 percent adequacy). In Appendix Figure A5, we plot the effect of a growing season day above 110°F on the percentage of households below different adequacy thresholds for each nutrient. The coefficients for 70 and 80 percent thresholds are generally similar (with the exception of zinc), as are the 50 and 60 percent thresholds. In general, the results in this figure point to larger negative impacts of heat at the lower end of the nutrient adequacy distribution, and the results are largely robust to using different thresholds to indicate strong undernutrition.

We are unable to analyze two connections between household food consumption and individual health outcomes due to the nature of our data. First, we cannot take waste (during cooking or after), different preparation styles, or combinations of foods into account. It is possible that levels of waste differ after a heat shock, as households may waste less when less food is available. The survey does not indicate what foods were eaten when, which matters due to connections between different nutrients and their absorption. For example, iron absorption increases when vitamin C is eaten in the same meal (Lane and Richardson, 2014). Second, we are unable to say anything about intra-household allocation of resources, which is biased against female children in India and other South Asian countries (Kaul, 2018; Azam and Kingdon, 2013). As such, our household-level estimates of undernourishment may misclassify it for certain particularly vulnerable groups.

[7] We also check that any per capita outcomes (e.g., calories per capita, calories from home per capita) are robust to setting children as equivalent to one half or one third of a person. We do not present these in the Appendix, but they are available upon request.

It is also possible that this allocation is sensitive to weather. Past work suggests that intra-household allocation is more unequal in households that are less economically well-off (D'souza and Tandon, 2019), and that this allocation can become more unequal during unexpected food shortages (Behrman and Deolalikar, 1990; Harris-Fry et al., 2017). Future work should, where possible, consider the effects of climate shocks on intra-household allocation.

The sample of households in this study was surveyed between 2003 and 2012, and the effects of extreme heat on nutrition may have changed between the study period and the present day. On the one hand, changes in technology, wealth, or other *ex-ante* adaptations may have reduced the impact of extreme heat on yields. On the other hand, warm days have become more frequent: relative to our sample period, the period 2015 to 2024 saw 11 fewer days below 90°F and 11 more days above 90°F (though the number of days above 110°F is roughly the same). A higher frequency of hot growing season days may increase the difficulty or decrease the effectiveness of certain types of adaptations, thereby increasing the harm from hot days.

5 Discussion

Understanding household responses to rising temperatures is a key aspect of designing climate adaptation policy. Our findings suggest that rural Indian households have been able to partially mitigate the loss of calorie and nutrient consumption. They do so by purchasing more food than normal in order to compensate for losses in home-grown foods. In addition, after a hot growing season, there is a partial reallocation of labor among rural households from agriculture to non-agricultural sectors, which may help explain how households pay for these additional purchases. Importantly, these households may still face significant longer-term consequences that we cannot observe. Prior research documents that poor rural households often lack access to credit or other means of smoothing consumption across periods in response to income shocks (Castells-Quintana et al., 2018). Along with our findings, this suggests households may be sacrificing non-food purchases, or reducing long-term investments in human capital.

In addition, we find evidence of distributional impacts, with a hot growing season leading to an increase in households experiencing strong malnutrition. For some crucial nutrients, such as iron, the largest decreases in consumption are among households at even more severe levels of undernutrition (around the 60 percent adequacy threshold). This finding suggests that the costs of adaptation to a hot growing season may be prohibitively high for some households. This constraint may be especially binding for the poorest households, who are likely to already have low levels of food consumption. Future work should consider the impacts of policies which might increase resilience to climate shocks. For instance, policies that increase market access to both foods and non-agricultural labor employment opportunities may be effective in enhancing rural households' abilities to adapt to heat shocks. Another set of policies may include direct cash assistance to households in areas that have recently experienced (or are forecast to experience) extreme heat, as has been done in the case of floods (Pople et al., 2021), droughts (Premand and Stoeffler, 2022), and, more recently in India, heat (Madhok, 2024).

References

Azam, Mehtabul, and Geeta Gandhi Kingdon. 2013. Are Girls the Fairer Sex in India? Revisiting Intra-Household Allocation of Education Expenditure. World Development 42: 143–164.

Barreca, Alan, and Jessamyn Schaller. 2020. The impact of high ambient temperatures on delivery timing and gestational lengths. Nature Climate Change 10 (January).

Behrman, Jere R, and Anil B Deolalikar. 1987. Will developing country nutrition improve with income? a case study for rural south india. Journal of political Economy 95 (3): 492–507.

———. 1990. The intrahousehold demand for nutrients in rural south india: Individual estimates, fixed effects, and permanent income. Journal of human resources : 665–696.

Burgess, Robin, Olivier Deschenes, Dave Donaldson, and Michael Greenstone. 2017. Weather, climate change and death in india. University of Chicago : 577–617.

Carleton, Tamma, Esther Duflo, B Kelsey Jack, and Guglielmo Zappalà. 2024. Adaptation to climate change. In Handbook of the Economics of Climate Change, volume 1. Elsevier, 143–248.

Carleton, Tamma, Amir Jina, Michael Delgado, Michael Greenstone, Trevor Houser, Solomon Hsiang, Andrew Hultgren, et al. 2022. Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits. The Quarterly Journal of Economics : 2037–2105.

Carleton, Tamma A. 2017. Crop-damaging temperatures increase suicide rates in india. Proceedings of the National Academy of Sciences 114 (33): 8746–8751.

Carpena, Fenella. 2019. How Do Droughts Impact Household Food Consumption and Nutritional Intake? A Study of Rural India. World Development 122 (226057): 349–369.

Castells-Quintana, David, Lopez-U, Pilar Ribe, and Thomas K J Mcdermott. 2018. Adaptation to climate change : A review through a development economics lens. World Development 104: 183–196.

Chaparro, Camila M, and Parminder S Suchdev. 2019. Anemia epidemiology, pathophysiology, and etiology in low- and middle-income countries. Annals of the New York Academy of Sciences 1450 (1): 15–31.

Colmer, Jonathan. 2021. Temperature, labor reallocation, and industrial production: Evidence from india. American Economic Journal: Applied Economics 13 (4): 101–124.

Copernicus Climate Change Service. 1979-2012. ECMWF Reanalysis v5 (ERA5). Accessed April 26, 2021.

Davis, Lucas, Paul Gertler, Stephen Jarvis, and Catherine Wolfram. 2021. Air conditioning and global inequality. Global Environmental Change 69: 102299.

Deaton, Angus, and Jean Drèze. 2009. Food and nutrition in India: Facts and interpretations. Economic and Political Weekly 44 (7): 42–65.

Dillon, Andrew, Kevin McGee, and Gbemisola Oseni. 2015. Agricultural Production, Dietary Diversity and Climate Variability. Journal of Development Studies 51 (8): 976–995.

D'souza, Anna, and Sharad Tandon. 2019. Intrahousehold nutritional inequities in rural bangladesh. Economic Development and Cultural Change 67 (3): 625–657.

Garg, Teevrat, Maulik Jagnani, and Vis Taraz. 2020. Temperature and human capital in india. Journal of the Association of Environmental and Resource Economists 7 (6): 1113–1150.

Georgieff, Michael K. 2020. Iron deficiency in pregnancy. American journal of obstetrics and gynecology 223 (4): 516–524.

Gopalan, C., V. Rama Sastri, S.C. Balasubramanian, B.S. Narasinga Rao, Y.G. Deosthale, and K.C. Pant. 1989. Nutritive Value of Indian Foods.

Harris-Fry, Helen, Niva Shrestha, Anthony Costello, and Naomi M Saville. 2017. Determinants of intra-household food allocation between adults in south asia–a systematic review. International journal for equity in health 16: 1–21.

Hersbach, Hans, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater, Julien Nicolas, et al. 2020. The era5 global reanalysis. Quarterly Journal of the Royal Meteorological Society 146 (730): 1999–2049.

Hou, Xiaohui. 2010. Can Drought Increase Total Calorie Availability? The Impact of Drought on Food Consumption and the Mitigating Effects of a Conditional Cash Transfer Program. Economic Development and Cultural Change .

ICMR, Indian Council of Medical Research. 2010. Nutrient Requirements and Recommended Dietary Allowances for Indians.

ICRISAT. 2015. Meso level data for india: 1966–2011, collected and compiled under the project on village dynamics in south asia.

International Institute for Population Sciences (IIPS), and ICF. 2017. National family health survey (NFHS-4), 2015-16: India. Survey report, International Institute for Population Sciences, Mumbai.

———. 2017, 2022. National family health survey (NFHS-4 and NFHS-5), 2015-16 and 2019-21 [dataset]. Dataset.

Jayasinghe, Maneka, Andreas Chai, Shyama Ratnasiri, and Christine Smith. 2017. The power of the vegetable patch : How home-grown food helps large rural households achieve economies of scale & escape poverty. Food Policy 73 (April): 62–74.

Kaul, Tara. 2018. Intra-household allocation of educational expenses : Gender discrimination and investing in the future. World Development 104: 336–343.

Lane, Darius J R, and Des R Richardson. 2014. The active role of vitamin C in mammalian iron metabolism: much more than just enhanced iron absorption! Free Radical Biology and Medicine 75: 69–83.

Madhok, Diksha. 2024. It's so hot in india, an insurer is helping thousands of women buy food | cnn business.

Marcus, Hannah, Claudia Schauer, and Stanley Zlotkin. 2021. Effect of Anemia on Work Productivity in Both Labor- and Occupations : A Systematic Narrative Synthesis. Food and Nutrition Bulletin 42 (2): 289–308.

Mukherjee, Diganta, and Prabir Chaudhury. 2020. Why does "last week" reporting give higher estimates than "last month"? Communications in Statistics-Theory and Methods 49 (8): 1873–1893.

National Sample Survey Office. 2003-2012. National Sample Survey, Schedule 1, Rounds 59-64, 66, 68. Accessed October 12, 2022.

———. 2004-2012. National Sample Survey, Schedule 10, Rounds 60, 61, 62, 64, 66, 68. Accessed October 12, 2022.

NSS, National Sample Survey Office. 2014. Nutritional intake in India, 2011–2012 (Number 560).

Pailler, Sharon, and Magda Tsaneva. 2018. The effects of climate variability on psychological well-being in India. World Development 106: 15–26.

Pople, Ashley, Ruth Hill, Stefan Dercon, and Ben Brunckhorst. 2021. Anticipatory cash transfers in climate disaster response .

Premand, Patrick, and Quentin Stoeffler. 2022. Cash transfers, climatic shocks and resilience in the sahel. Journal of Environmental Economics and Management 116: 102744.

Randell, Heather, Clark Gray, and Elizabeth H Shayo. 2022. Climatic conditions and household food security : Evidence from Tanzania. Food Policy 112 (September): 102362.

Schlenker, Wolfram, and Michael J Roberts. 2009. Nonlinear temperature effects indicate severe damages to us crop yields under climate change. Proceedings of the National Academy of sciences 106 (37): 15594–15598.

Scott, Chris, and Ben Amenuvegbe. 1990. Effect of recall duration on reporting of household expenditures. World Bank, Washington DC .

Shah, Manisha, and Bryce Millett Steinberg. 2017. Drought of Opportunities : Contemporaneous and Long-Term Impacts of Rainfall Shocks on Human Capital Manisha Shah. Journal of Political Economy 125 (2).

Shaheen, Nazma, Abu Torab MA Rahim, Md. Mohiduzzaman, Cadi Parvin Banu, Md. Latiful Bari, Avonti Basak Tukun, MA Mannan, et al. 2013. Food composition table for bangladesh.

Sharma, Himani, S K Singh, and Shobhit Srivastava. 2020. Socio-Economic Inequality and Spatial Heterogeneity in Anaemia Among Children in India: Evidence from NFHS-4 (2015-16). Clinical Epidemiology and Global Health 8 (4): 1158–1171.

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

Taraz, Vis. 2018. Can farmers adapt to higher temperatures? Evidence from India. World Development 112 (December): 205–219.

Thomas, Duncan, Elizabeth Frankenberg, Jed Friedman, Jean-Pierre Habicht, Mohammed Hakimi, Nicholas Ingwersen, Jaswadi, et al. 2006. Causal Effect of Health on Labor Market Outcomes : Experimental Evidence. Working Paper (September).

World Bank. 2024. Rural population - India.

6 Tables

Table 1: Summary Statistics: Nutrition and Weather

	Country	North	Central	West	South	East	Northeast
Calories per capita (kCal)	2,181	2,397	2,222	2,101	2,109	2,160	2,104
Calories AP (CAP)	103.62	113.57	106.36	98.87	99.77	103.00	99.95
CAP <80%	20.41	10.80	18.11	25.28	26.99	20.26	20.27
Protein per capita (g)	58.69	70.82	63.48	56.84	53.72	55.55	53.51
Protein AP (PAP)	119.82	144.30	133.20	114.34	106.56	114.53	109.13
PAP <80%	11.62	1.74	5.72	11.20	20.48	13.36	13.85
Iron per capita	15.27	21.47	19.00	18.24	11.29	14.21	9.72
Iron AP (IAP)	81.80	114.51	102.47	97.64	60.38	76.30	51.75
IAP <80%	54.69	18.13	28.22	31.81	82.11	60.66	91.17
Zinc per capita	10.30	12.69	11.14	10.05	9.37	9.79	9.11
Zinc AP (ZAP)	107.69	132.94	118.83	103.84	96.08	102.76	95.07
ZAP <80%	21.52	3.98	11.89	20.20	34.54	23.44	30.27
Thiamine per capita	1.34	2.11	1.83	1.61	0.90	1.11	0.74
Thiamine AP (TAP)	125.75	197.21	173.07	149.39	84.06	104.87	69.64
TAP <80%	35.11	5.03	10.29	9.20	55.23	41.22	75.48
Niacin per capita	15.03	17.16	17.33	15.04	12.64	15.15	13.31
Niacin AP (NAP)	108.72	123.87	126.08	107.93	91.73	109.63	95.87
NAP <80%	22.05	7.76	9.70	20.01	42.66	17.99	28.83
Riboflavin per capita	1.13	1.90	1.26	1.18	0.97	0.87	0.79
Riboflavin AP (RAP)	91.62	153.00	102.44	94.35	78.13	71.18	63.66
RAP <80%	51.97	9.60	40.56	41.97	60.95	68.49	78.69
GS days ≥110°F	0.59	1.61	1.20	0.18	0.09	0.47	0.00
GS days 100-110°F	7.65	17.60	11.91	6.64	5.64	4.88	0.01
GS days 90-100°F	50.61	59.82	60.45	51.76	49.79	58.39	18.89
GS days 80-90°F	109.26	57.80	98.01	137.34	136.15	121.90	99.05
GS days 70-80°F	33.23	37.38	32.95	18.06	22.24	27.00	64.01
GS days <70°F	12.66	39.79	9.49	0.02	0.08	1.36	32.05
GS rainfall (m)	1.24	0.66	1.07	1.15	1.18	1.36	1.99
GS rainfall shock	0.11	-0.06	0.12	0.25	0.17	0.19	-0.04
Number of Districts	575	98	144	63	94	101	75
Number of Observations	314,414	46,012	58,361	31,105	67,289	64,872	46,775

This table shows the means of our primary dependent variables: daily per capita consumption (in mg unless otherwise noted, and winsorized at the 1 and 99 percent levels), adequacy percentage (AP), and the percentage of households below 80% AP. All measures are based on a 30 day recall period for each household from the National Sample Survey (NSS), covering the period 2003 to 2012. It also shows the core weather growing season (GS, June-December) variables from ERA5 for 2002 to 2011. The non-growing season (March to May) summary statistics are in Appendix Table A1. We do not include January and February in either season because they contain very few hot days. The rainfall shock takes on a value of 1 if rainfall in a given district-season-year exceeds the 80th percentile of its historical distribution (defined as 1979-2012) of rainfall for that district-season, and a value of -1 if rainfall is below the 20th percentile of the historical distribution (0 otherwise). We present statistics for the full sample and for six regions: North (Delhi, Haryana, Himachal Pradesh, Jammu & Kashmir, Punjab, and Rajasthan), Central (Chhattisgarh, Madhya Pradesh, Uttar Pradesh, and Uttarakhand), West (Dadra & Nagar Haveli, Daman & Diu, Goa, Gujarat, and Maharashtra), South (Andhra Pradesh, Karnataka, Kerala, Pondicherry, and Tamil Nadu), East (Bihar, Jharkhand, Odisha, and West Bengal), and Northeast (Arunachal Pradesh, Assam, Manipur, Meghalaya, Mizoram, Nagaland, Sikkim, and Tripura).

Table 2: Core Specification: Last Year's Growing Season Coefficients

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	100x Log of Daily per Capita Calories	Percent of Households Below 100% Calories Adequacy	Percent of Households Below 80% Calories Adequacy	Percent of Households Below 60% Calories Adequacy	Percent of Households Below 80% Protein Adequacy	Percent of Households Below 80% Iron Adequacy	Percent of Households Below 80% Zinc Adequacy	Percent of Households Below 80% Thiamine Adequacy	Percent of Households Below 80% Niacin Adequacy	Percent of Households Below 80% Riboflavin Adequacy
Day ≥110°F	0.015	0.193	0.356***	0.105**	0.022	0.262*	0.246**	0.317***	0.408***	-0.089
	(0.108)	(0.187)	(0.133)	(0.052)	(0.086)	(0.139)	(0.124)	(0.114)	(0.127)	(0.120)
Day 100-110°F	-0.040	0.112	0.249***	0.092***	0.061	0.129*	0.182**	0.198***	0.251***	0.037
	(0.054)	(0.093)	(0.077)	(0.027)	(0.046)	(0.068)	(0.077)	(0.064)	(0.077)	(0.060)
Day 90-100°F	-0.051	0.108	0.226***	0.077***	0.064	0.157**	0.159**	0.223***	0.238***	0.104*
	(0.049)	(0.085)	(0.070)	(0.024)	(0.041)	(0.064)	(0.069)	(0.061)	(0.070)	(0.054)
Day 80-90°F	-0.054	0.129	0.224***	0.065***	0.051	0.137**	0.164**	0.220***	0.228***	0.101*
	(0.049)	(0.082)	(0.069)	(0.024)	(0.041)	(0.063)	(0.068)	(0.060)	(0.069)	(0.052)
Day 70-80°F	-0.024	0.084	0.131**	0.040*	-0.011	0.064	0.064	0.131**	0.133**	0.036
	(0.046)	(0.075)	(0.064)	(0.023)	(0.036)	(0.059)	(0.061)	(0.057)	(0.061)	(0.049)
Rainfall Shock	-0.073	-0.126	-0.113	0.142	0.398	-0.068	0.381	-0.462*	-0.008	-0.277
	(0.228)	(0.390)	(0.310)	(0.125)	(0.243)	(0.314)	(0.322)	(0.280)	(0.303)	(0.275)
Observations	314,413	314,414	314,414	314,414	314,414	314,414	314,414	314,414	314,414	314,414
Mean Outcome	765.37	51.00	20.41	3.94	11.62	54.69	21.52	35.11	22.05	51.97
District FE	X	X	X	X	X	X	X	X	X	X
Year FE	X	X	X	X	X	X	X	X	X	X
Month FE	X	X	X	X	X	X	X	X	X	X
Demographic controls	X	X	X	X	X	X	X	X	X	X
State time trends	X	X	X	X	X	X	X	X	X	X
Rainfall shocks	X	X	X	X	X	X	X	X	X	X
Last year's weather	X	X	X	X	X	X	X	X	X	X
This year's weather	X	X	X	X	X	X	X	X	X	X

This table shows the change in our main outcome variables for an additional day at a given temperature during last year's growing season relative to a day below 70°F (Equation 1). The outcomes in column 1 are multiplied by 100 so the coefficients can be interpreted as percentage changes. The models all include 10°F temperature bins from 70-80°F to above 110°F for this and last year's growing season (June-December) and non-growing season (March-May). We omit the <70°F bins so the coefficients can be interpreted relative to this temperature range. Each model includes district, year, and month fixed effects, demographic controls, rainfall shocks, and state time trends. Standard errors (shown in parentheses) are clustered at the district level to account for any inter-district correlation of the error terms. All regressions include household-level outcomes from 2003-2012. *p<0.1 **p<0.05 ***p<0.01

7 Figures

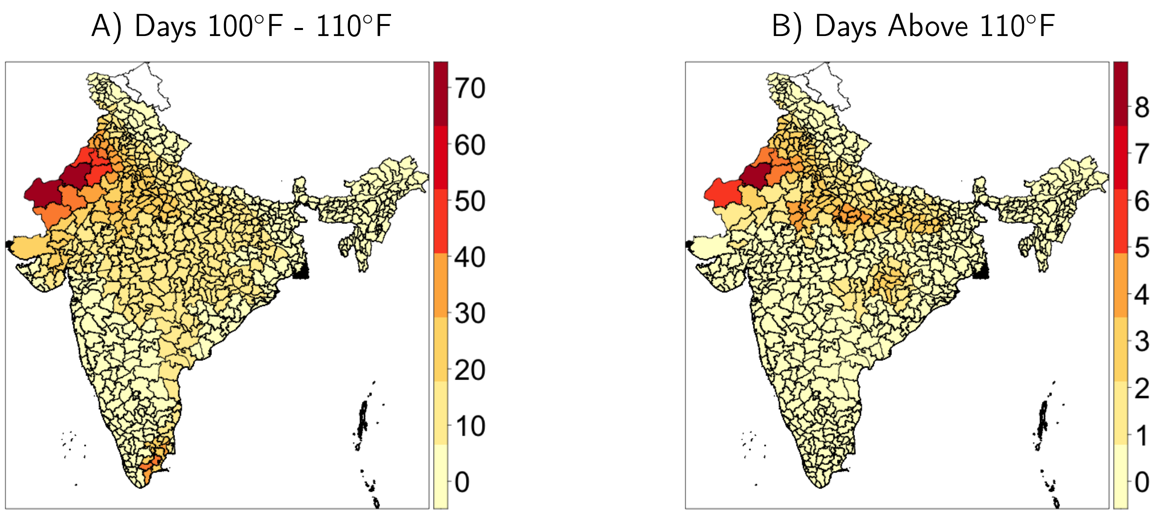

Figure 1: Growing Season Weather by District (2000-2011)

Panel A shows the average number of growing season (June-December) days with maximum temperature between 100°F and 110°F in each district from 2002 to 2011. Panel B shows the number of growing season days greater than or equal to 110°F.

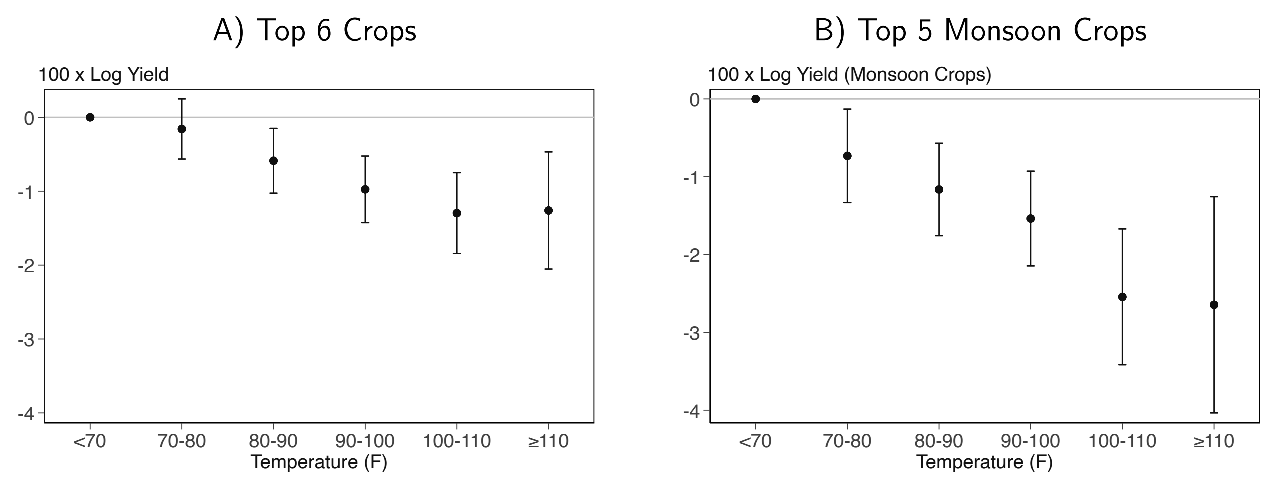

Figure 2: This Year's Growing Season Weather and Yields

These graphs show the change in 100 times log yields for each additional day at a given temperature during this year's growing season relative to a day below 70°F. Panel A is for the top 6 crops in India (rice, wheat, sugarcane, groundnut, sorghum, and maize). Panel B is for the top 5 monsoon-dependent crops (excludes wheat). The coefficients can be interpreted as percentage changes. We match district-year yields with weather data at the district-year-season level. The model includes 10°F temperature bins from 70-80°F to above 110°F for this year's growing season (June - December) and non-growing season (March - May). We omit the <70°F bins so the coefficients can be interpreted relative to this temperature range. The models include district and year fixed effects and rainfall shocks. Standard errors (errors bars represent the 95 percent confidence interval) are clustered at the district level to account for any inter-district correlation of the error terms. The models include data from 2000-2011. The model in Panel A includes 3,365 observations, and Panel B includes 3,365. The mean outcome variable in panel A is 507.88, and the mean outcome variable in panel B is 496.74.

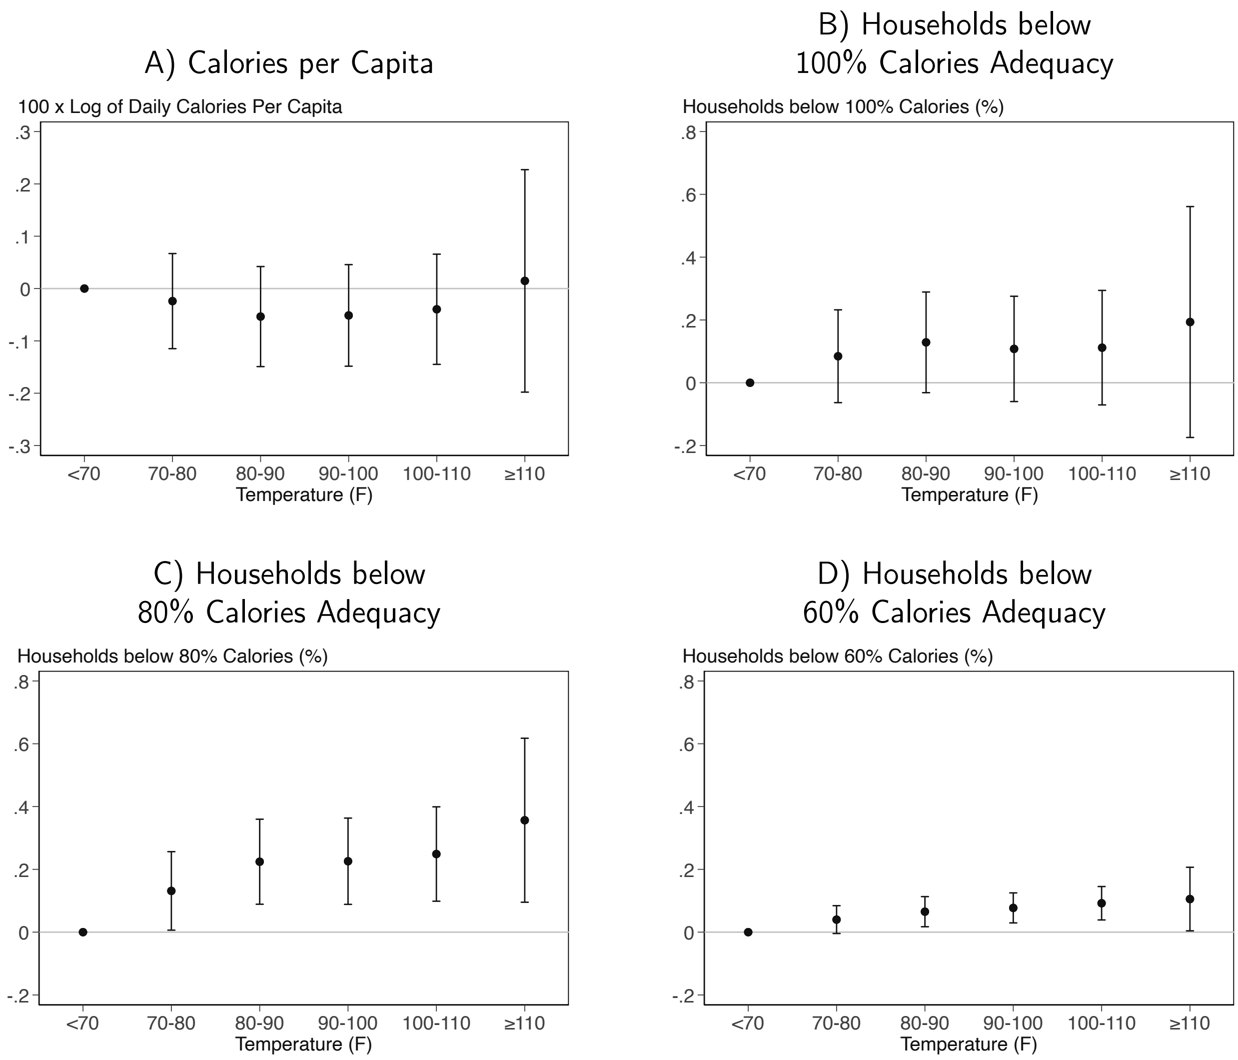

Figure 3: Last Year's Growing Season Weather and Calorie Availability

Panel A shows the change in 100 times the log of per capita daily consumption of calories (winsorized at the 1 and 99 percent levels) for each additional day at a given temperature during last year's growing season relative to a day below 70°F. The coefficients can be interpreted as percentage changes. The dependent variable in panel B is the percentage of households consuming below 100 percent of the recommended level of calories, based on household composition. Panels C and D are the same, but for the 80 percent threshold and 60 percent threshold, respectively. All models include 10°F temperature bins from 70-80°F to above 110°F for this year and last year's growing season (June-December) and non-growing season (March-May). They omit the <70°F bins so the coefficients can be interpreted relative to this temperature range. The models include district, month, and year fixed effects, state time trends, rainfall shocks, demographic (religion, social group, and head of household education) and household composition (Panel A: number of members of different ages and genders, Panels B-D: household size) controls. Standard errors (error bars represent the 95% confidence interval) are clustered at the district level to account for any inter-district correlation of the error terms. All regressions include household-level outcomes from 2003-2012. The model in panel A includes 314,413 observations, and the models in Panels B-D include 314,414 observations. These coefficients correspond to those in Table 2, columns 1-4.

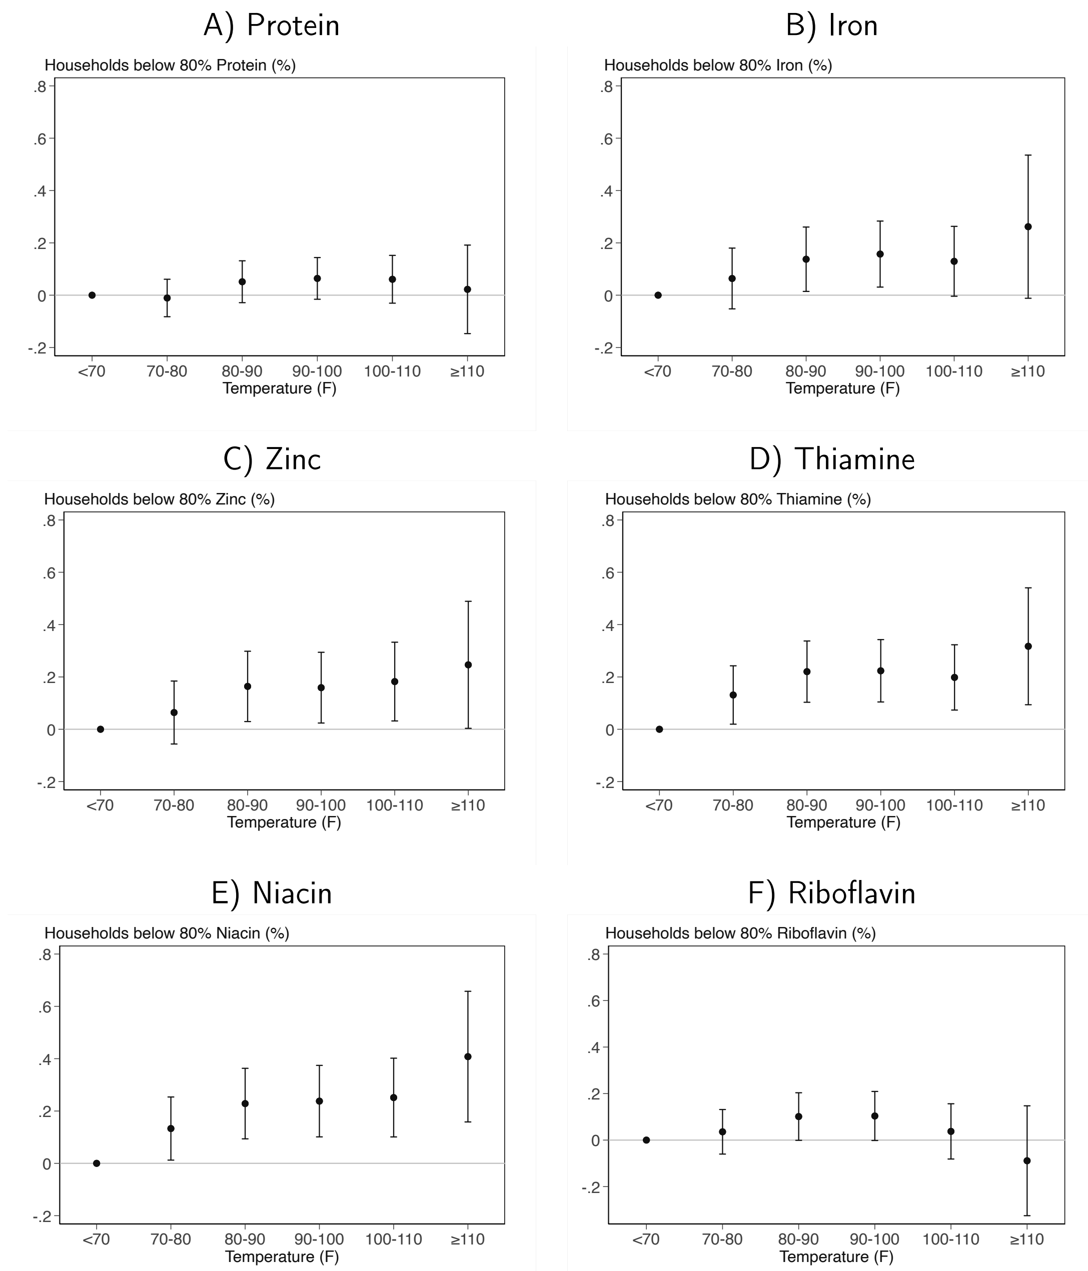

Figure 4: Last Year's Growing Season Weather and 80% Nutritional Adequacy

Panel A shows the change in the percentage of households consuming below 80 percent of the recommended level of protein, based on household composition, for each additional day at a given temperature during last year's growing season relative to a day below 70°F. Panels B-F show the same relationship for different nutrients. We match household-year-month nutritional outcomes with weather data at the district-year-season level. All models include 10°F temperature bins from 70-80°F to above 110°F for this year and last year's growing season (June - December) and non-growing season (March - May). They omit the <70°F bins so the coefficients can be interpreted relative to this temperature range. The models include district, month, and year fixed effects, state time trends, rainfall shocks, demographic (religion, social group, and head of household education) and household size controls. Standard errors (error bars represent the 95% confidence interval) are clustered at the district level to account for any inter-district correlation of the error terms. All regressions include household-level outcomes from 2003-2012. The models all include 314,414 observations. These coefficients correspond to those in Table 2, columns 5-10.

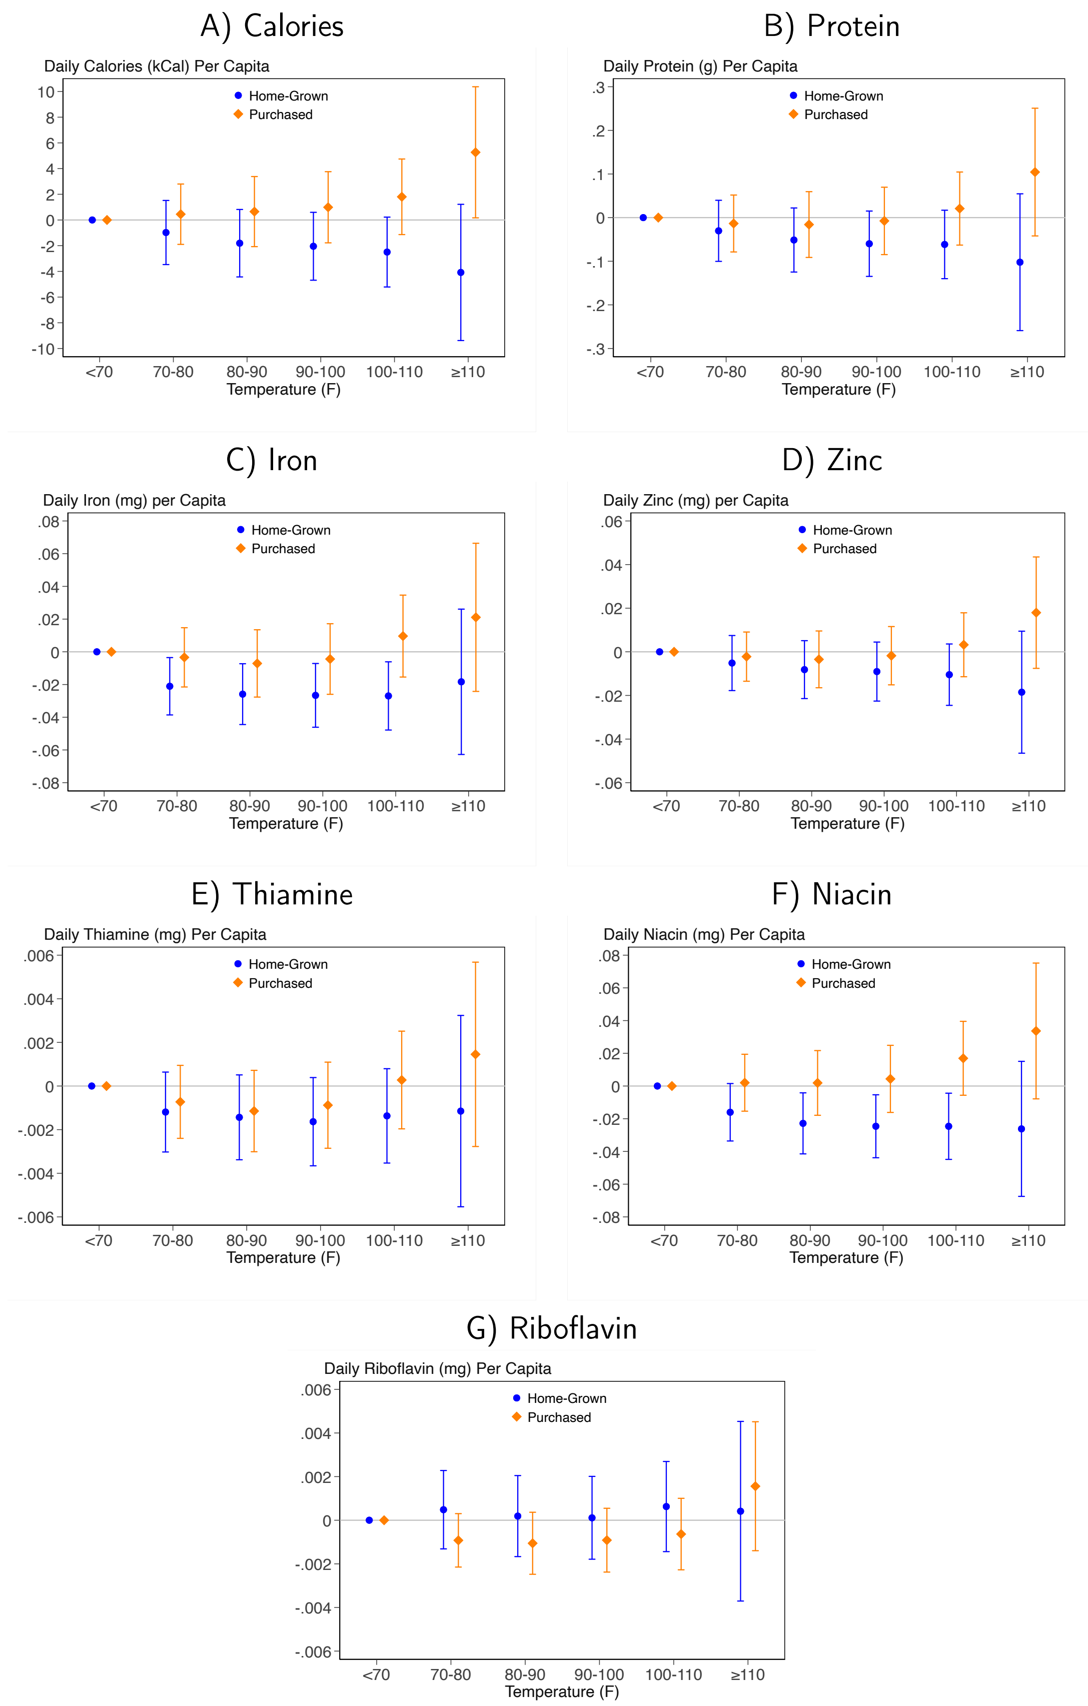

Figure 5: Last Year's Growing Season Weather and Home-Grown vs. Purchased Nutrients

Panel A shows the change in calories per capita from home grown foods (blue circles) and from purchased foods (orange diamonds) for each additional day at a given temperature during last year's growing season relative to a day below 70°F. Panels B-G show the same relationship for different nutrients. Each outcome is winsorized at the 1 and 99 percent levels. We match household-year-month nutritional outcomes with weather data at the district-year-season level. All models include 10°F temperature bins from 70-80°F to above 110°F for this year and last year's growing season (June - December) and non-growing season (March - May). They omit the <70°F bins so the coefficients can be interpreted relative to this temperature range. The models include district, month, and year fixed effects, state time trends, rainfall shocks, demographic (religion, social group, and head of household education) and household composition controls. Standard errors (error bars represent the 95% confidence interval) are clustered at the district level to account for any inter-district correlation of the error terms. All regressions include household-level outcomes from 2003-2012. The models all include 314,414 observations.

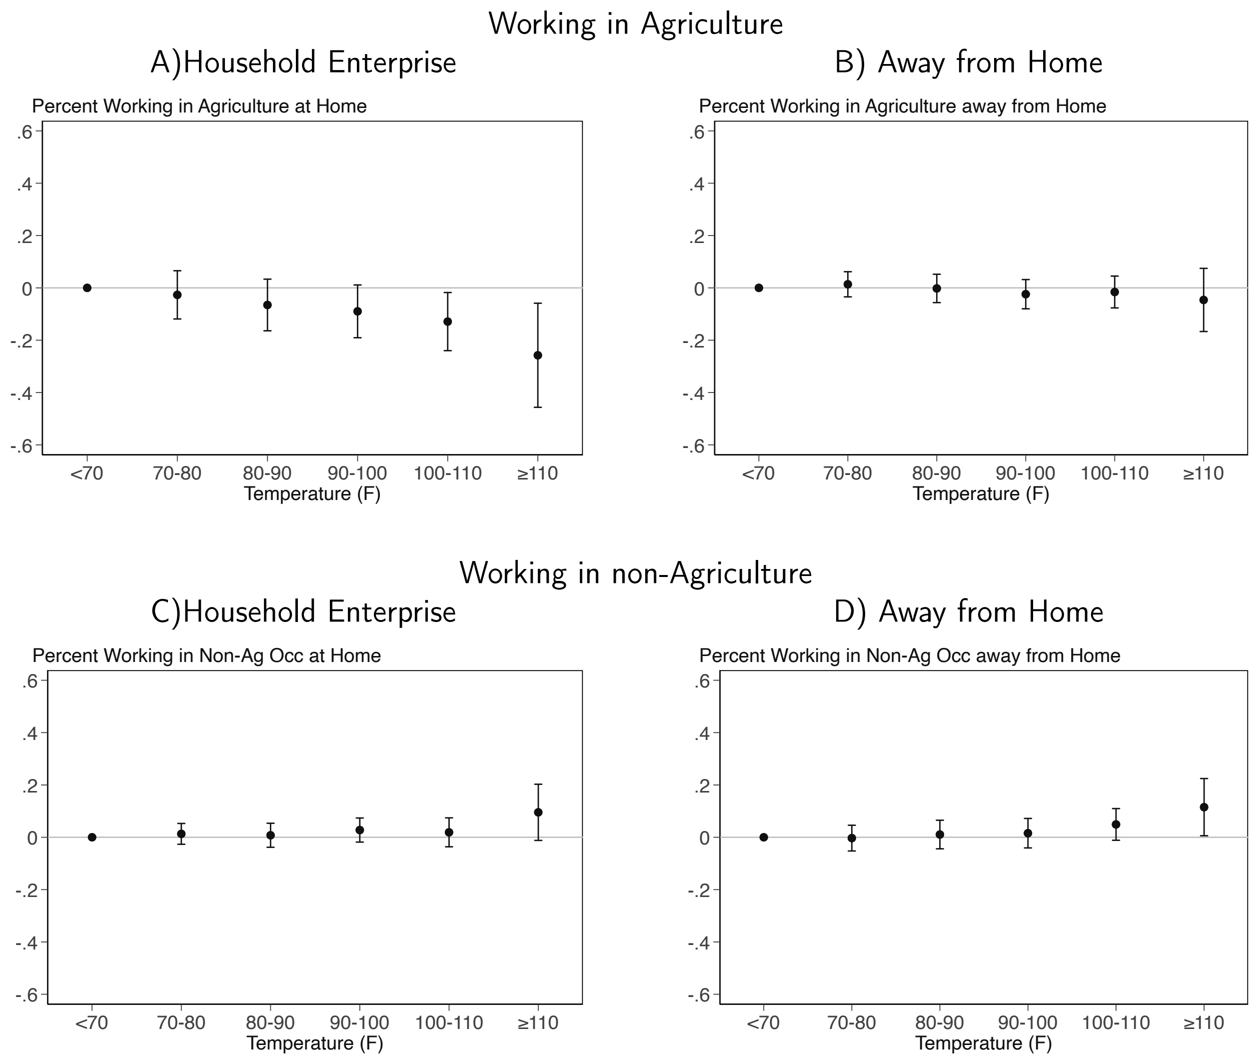

Figure 6: Last Year's Growing Season Weather and Working-Age (14-65) Work Status

Panel A shows the change in the percent of working-age individuals (14-65) working in agriculture at home for each additional day at a given temperature during last year's growing season relative to a day below 70°F. Panel B shows the same relationship for working-age individuals working in agriculture away from home. Panels C and D show the same relationships as A and B, respectively, but for non-agricultural work. All models include 10°F temperature bins from 70-80°F to above 110°F for this year and last year's growing season (June - December) and non-growing season (March - May). They omit the <70°F bins so the coefficients can be interpreted relative to this temperature range. The models include district, month, and year fixed effects, state time trends, rainfall shocks, individual demographic controls (religion, social group, age, sex and education level). Standard errors (error bars represent the 95% confidence interval) are clustered at the district level to account for any inter-district correlation of the error terms. All regressions include individual-level outcomes from 2004-2012. All models include 1,134,910 observations.

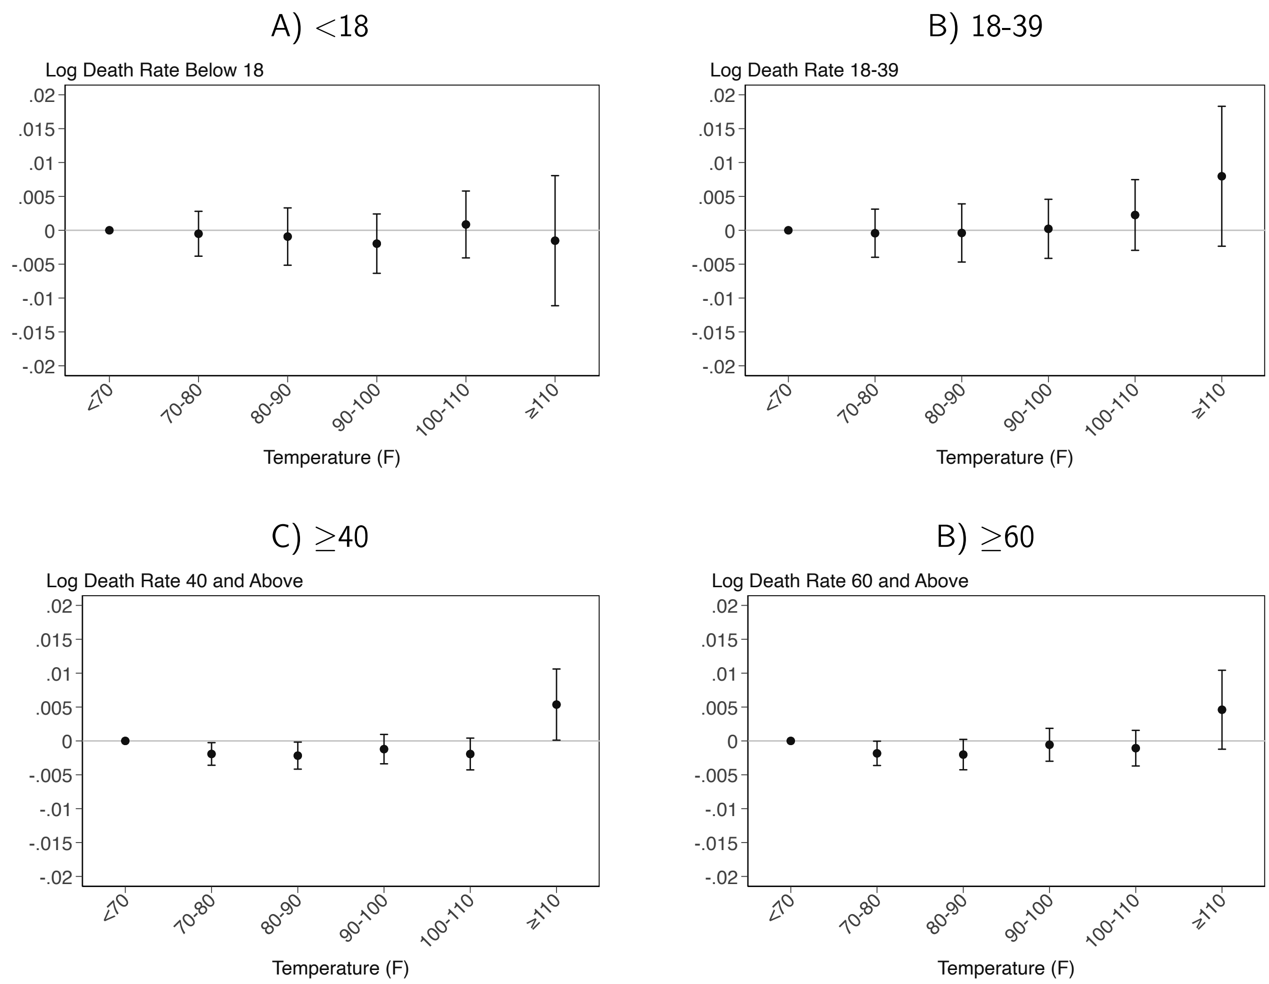

Figure 7: This Year's Full Year Weather and Mortality

Panel A shows the change in the log mortality rate (deaths per 1,000 people) for children (below 18) for each additional day at a given temperature this year relative to a day below 70°F. Panels B, C, and D show the same relationship as panel A, but for the 18-39, 40 and above, and 60 and above mortality rates, respectively. Each outcome is winsorized at the 1 and 99 percent levels. We pair district-year-level mortality rates from 2012 to 2019 with district-year level weather. All models include 10°F temperature bins from 70-80°F to above 110°F for this year and last year. They omit the <70°F bins so the coefficients can be interpreted relative to this temperature range. The models include year and district fixed effects and a state-year trend. Standard errors (error bars represent the 95% confidence interval) are clustered at the district level to account for any inter-district correlation of the error terms. Panel A includes 3,407 observations, panel B includes 3,187, panel C includes 3,695, and panel D includes 3,689.