

The Impact of Climate Change on Risk and Return in Indian Agriculture

Francisco Costa* Fabien Forge[†] Jason Garred[‡] João Paulo Pessoa[§]

November 4, 2022

Abstract

We investigate the extent to which climate change will result in insurable and uninsurable losses for farmers in India. Changes in temperature and precipitation patterns may increase the volatility of farmers' yields, leading to rising but insurable risk. These changes may also reduce the expected yield in an 'average year', causing uninsurable reductions in the returns to farming. We use a multi-run climate model to predict the future distribution of potential yields at the district level for sixteen major crops. For the average district, we project a sharp decline in expected agricultural revenue, but small shifts in volatility. This is because weather draws resulting in extremely low agricultural returns – what had once been 1-in-100-year events – are predicted to become the norm by the end of the century. Our projections therefore imply substantial uninsurable losses from the changing climate for Indian farmers.

Keywords: agriculture; climate change; crop yield; extreme events; India; volatility.

We thank Carolyn Fischer, Anthony Heyes and Maya Papineau, seminar participants at the World Bank, Insper, Calgary, UFDP, the Federal University of ABC, and the Climate Change Economics Lunch and the Development Lunch at UC Berkeley, as well as participants at the Environment for Development Annual Meeting, the Institutions Trade and Economic Development (InsTED) Workshop, the AERE Summer Conference, the EAERE 25th Annual Conference, the Economics of Low-Carbon Markets Workshop, the Occasional Workshop in Environmental and Resource Economics, the CIREQ Environmental and Natural Resource Economics Workshop, the Brazilian Econometric Society Meeting, the Canadian Resource and Environmental Economics Annual Meeting, the Moscow International Economics Workshop, and the Annual Meeting of the Canadian Economics Association. We also thank Magaly Somarriba for excellent research assistance. This study was financed in part by CAPES/Brasil, Grant 001, and the Canada 150 Research Chair in Climate Economics, Innovation and Policy.

*Department of Economics, University of Delaware. Email: fcosta@udel.edu.

[†]Department of Economics, University of Ottawa. Email: fforge097@uottawa.ca.

[‡]Department of Economics, University of Ottawa. Email: Jason.Garred@uottawa.ca.

[§]Sao Paulo School of Economics - FGV & Centre for Economic Performance. Email: joao.pessoa@fgv.br.

1 Introduction

Climate change may have significant consequences for agriculture around the globe. This might not only involve the crop yield that farmers can expect in an ‘average year’, but also the volatility of agricultural production and the frequency with which harvests are decimated by extreme weather. If future temperature and precipitation patterns result in greater yield variability, policies improving access to agricultural insurance might play a key role in mitigating the effects of climate change on farmers. On the other hand, if climate change leads to a sharp reduction in mean yields without an accompanying change in volatility, the main issue facing farmers will be the prospect of permanent, uninsurable losses.

In this paper, we investigate the extent to which climate change will result in insurable and uninsurable losses for farmers. We contrast the potential effect of climate change on farmers’ average agricultural returns with its impact on the volatility of those returns. We also examine the changing probability of extreme events with especially low returns. We focus on India, a setting with a large agricultural sector but relatively limited availability of insurance as compared to wealthier economies.

In order to jointly examine changes in average yields, variability of yields and extreme events, we require information on *distributions* of potential yields. We therefore project a yield distribution (i.e. a set of possible yield realizations) for each location, crop and time period in our sample. The Community Earth System Model Large Ensemble (CESM-LENS) dataset is the key data facilitating this exercise. This is a multi-run global simulation created by climatologists to study the potential variability of temperature and precipitation patterns within a given climate (Kay et al., 2015). It contains forty possible weather realizations at a given location for each day between 1920 and 2099.¹ We use this data to predict the evolution of the (potential) distribution of agricultural production for each of 310 Indian districts across eighteen decades.

¹For instance, rather than specifying a single projected value for precipitation for the southern tip of India on August 1 2050, the CESM-LENS data instead includes forty potential local precipitation realizations consistent with the prevailing climate as of 2050.

To convert weather projections into predicted crop yields, we employ a standard statistical approach using weather variation and fixed effects (Deschênes and Greenstone, 2007; Schlenker and Roberts, 2009). For each of the sixteen different crops in our data, we estimate models relating historical yields across Indian districts to realized temperature and precipitation. We allow for various possible ways of modelling the yield-weather relationship by running multiple alternative specifications. For example, we model temperature in terms of the number of days in various temperature bins, or in terms of degree days.²

We then generate a set of yield projections based on the results of each of these specifications. Specifically, we use the potential weather realizations from the CESM-LENS data to calculate 400 possible yield realizations per crop for each district and decade.³ We employ these projections to explore the consequences of climate change for total agricultural revenue in each district.

We begin by calculating expected agricultural revenue (i.e. the average across all possible realizations) for each district and decade. The results are consistent across the various projections based on our different regression specifications. For the average district, we forecast a gradual decline in expected agricultural revenue between the 2000s and the 2090s, ranging from 15.5% to 33.1%. These effects are driven by future rises in average temperature rather than changes in precipitation patterns.

Next, we assess potential changes in revenue variability. Estimates of the impact of climate change on variability in agricultural revenue are sensitive to the regression model underlying the projections. Specifications using temperature bins suggest a decline in the weather-induced standard deviation of agricultural revenue (on average across districts) of up to 68.4% between the 2000s and the 2090s. On the other hand, models based on degree days forecast a rise in this measure of up to 10.5%. These figures represent only the variation in agricultural revenue due to weather, rather than the total within-district variation in agricultural outcomes due to all factors. This is because we hold factors other than the climate constant. So even the

²For precipitation, we use either a quadratic in total precipitation or a set of precipitation bins. We also explore an augmented specification that includes variables capturing the within-season distribution of rain.

³We hold all factors other than the climate constant at their levels in 2000 in all projections.

largest positive projection does not suggest that climate change will lead to a substantial rise in the total variability of agricultural revenue.

Our third exercise using our revenue projections assesses the potential downside risk to agriculture from extreme weather. We define ‘1-in-100’ bad years for farmers based on the worst 1% of revenue realizations in each district from the first five decades of the projected data (the 1920s to the 1960s). We find that the frequency of such bad years will rise by 54 to 89 percentage points in the average district during the 21st century. In other words, what was once extremely poor weather for agricultural production is forecast to become the norm in India. This result is robust across specifications.

Overall, we project that climate change will have substantial negative effects on Indian farmers’ mean returns from agriculture, driven by the normalization of formerly extreme weather patterns. On the other hand, our projections do not provide strong evidence of a large rise in risk via increases in yield volatility. These findings have potentially important policy implications. They suggest that agricultural insurance may not be of first-order importance in mitigating farmers’ losses. Instead, to combat the threat of large but uninsurable losses from climate change in developing countries, a more promising avenue of adaptation could be the development and adoption of crop varieties resistant to extreme heat.

An important caveat of this analysis is that it holds factors such as crop choice and prices constant. We address this concern by taking a structural approach using the general equilibrium portfolio model of Allen and Atkin (2016). First, we allow farmers to adjust their crop choice according to the evolving distribution of potential yields, trading off the relative changes in risk and return across crops. Second, we let crop prices adjust endogenously with agricultural production. Third, we allow for changing trade flows across Indian districts. We find that allowing for endogenous crop choice, prices and trade has only a small moderating effect on the forecast losses from climate change.

Our paper contributes to the literature on the impact of climate change on agriculture in developing countries. A large body of work examines projected effects on average agricultural

outcomes (e.g. Auffhammer et al., 2006; Schlenker and Lobell, 2010; Lobell et al., 2011; Knox et al., 2012; Challinor et al., 2014; Chen et al., 2016). Some of this work focuses specifically on India (Guiteras, 2009; Burgess et al., 2017; Hari et al., 2018; Taraz, 2018). Some previous studies have also considered the future variability of agricultural yields, using projections of year-to-year weather variation.⁴ Urban et al. (2012) use the predicted time-series variation in weather patterns from one run of a given climate model to project future year-to-year variability in US maize yields. They then repeat this exercise separately for fifteen different climate models, deriving a distinct estimate of future variability for each model. Tigchelaar et al. (2018) project changes in the interannual volatility of global maize production under the scenario that global mean temperatures rise by two or four degrees Celsius, but year-to-year temperature variation remains as observed in a baseline historical period.

Our main contribution is to jointly consider future changes in farmers’ average return and portfolio risk in a unified framework. By studying the *distribution* of yields, we examine the evolution of both mean yields and their variability, along with the probability of extreme events. To do this, we deviate from standard approaches by exploiting many potential outcome realizations from a multi-run climate model rather than relying on year-to-year weather variation.⁵ Our paper also adds to a growing literature that assesses the effects of climate change using general equilibrium models.⁶ We contribute to this literature by considering a model in which climate change may affect agents’ welfare through rises or falls in portfolio risk.

⁴A recent set of papers considers the implications of agents’ uncertainty about the future climate, including Kala (2017), Shrader (2017) and Lemoine (2021). The main distinction between these studies of uncertainty and our exercise is that we examine the effect of changes in the distribution of potential economic outcomes when this distribution is known to agents, i.e. changes in known risk.

⁵Burke et al. (2015) use multiple climate models to refine projections of the average impacts of climate change. Here, we examine future changes in the distribution of potential yields (including their volatility) by exploiting a multi-run climate model designed for the study of climate variability. In this context, a key argument for focusing on a single multi-run model rather than a group of models is given in Ault et al. (2016). They state that “[a]lthough structural uncertainty and internal variability are conflated in the CMIP5 ensemble [of models], diagnosis of a single-model ensemble allows us to delineate the role of internal variability alone”.

⁶Along with the Ricardian formulation used by Costinot et al. (2016) and our work based on Allen and Atkin (2016), a study by Dingel et al. (2019) employs the general class of models explored by Arkolakis et al. (2012) to consider the impact of climate change on global inequality. Deryugina and Hsiang (2017) estimate the marginal product of climate in terms of total economic output using an Arrow-Debreu framework.

2 Data

In this section, we first discuss the agricultural and weather data that we use to estimate historical relationships between weather variables and crop yields. We then explain the nature of the climate change projection data and provide summary statistics for the key weather variables.

2.1 Historical weather and yield

Annual Indian agricultural data at the district level for 1979 to 2015 is provided by the Village Dynamics in South Asia (VDSA) dataset from the International Crops Research Institute for the Semi-Arid Tropics (ICRISAT).⁷ This gives us information by district-crop-year on quantity produced and area planted; we calculate yield as the quotient of these two variables. We keep the sixteen crops for which VDSA also provides similar data on farm-gate prices. These include the two main staple crops – rice and wheat – along with barley, castor, chickpea, cotton, finger millet, groundnut, linseed, maize, pearl millet, pigeon pea, rapeseed and mustard seed, sesame, sorghum and sugarcane. The data covers 310 districts across 20 of the 29 Indian states that existed at the end of our sample period in 2015.⁸

The crops in our dataset are commonly found across India. Nine of the sixteen crops were cultivated in at least 75% of the 310 sample districts in our baseline year of 2000, and all sixteen were grown in at least 30% of districts (see Table A1). The median district produced twelve of these crops in 2000. However, when we only include crops that occupy 5% or more of the land planted in a district, the median number of crops drops to three. Also, in almost half of the sample districts, one crop occupies at least 50% of land. Rice and wheat are the

⁷The VDSA data defines a year as beginning in June and ending in May, in line with India’s agricultural calendar. A year is therefore defined using this calendar throughout the paper. For example, 1980 corresponds to June 1980 to May 1981, and the 1980s encompass June 1980 to May 1990. We use 1979 to 2015 as our sample period because the VDSA data ends in 2015, while the ERA5 weather data we introduce below is available from 1979 onwards.

⁸The nine states excluded from the data consist of eight mostly Northeastern states with relatively small populations (Arunachal Pradesh, Goa, Manipur, Meghalaya, Mizoram, Nagaland, Sikkim and Tripura), and Jammu and Kashmir. Although there have been many boundary changes in Indian districts, the VDSA dataset defines districts consistently across time using 1966 district borders.

likeliest to be a district’s most-planted crop by land area (Figure A1).⁹

We construct district-level historical information on temperature and precipitation from 1979 to 2015 using the ERA5 dataset produced by the European Centre for Medium-Range Weather Forecasts. This is a simulation of past worldwide weather conditions, based on actual satellite, rain gauge and other observational data. It is produced at a resolution of approximately 7,000 square kilometers (0.75 degrees latitude by 0.75 degrees longitude).¹⁰

We calculate the mean daily temperature by averaging the minimum and maximum temperatures for each day from the ERA5 data. Total daily precipitation is computed by cumulating precipitation throughout the day. We allocate this information across 1966 districts by assigning to each district the weather information of the ERA5 pixel closest to the centroid of that district.

We limit these weather variables to the relevant growing season for each crop, based on information about crop-specific growing seasons from India’s Ministry of Agriculture and Farmers Welfare (2017). Nine of the sixteen crops we study are primarily grown in the *kharif* (monsoon) season from June to October, five crops are mainly grown in the *rabi* season of October to April, and one crop (sugarcane) matures throughout the year.¹¹

2.2 Weather projection

For projected weather patterns under climate change, we rely on the Community Earth System Model Large Ensemble (CESM-LENS) dataset. The stated purpose of the CESM-LENS project is to facilitate the study of ‘internal climate variability’; i.e. variation in realized weather within a given climate state (Kay et al., 2015). The CESM-LENS data provides us with daily information on average temperature and total precipitation from 1920 to 2099 for forty different simulations of global weather. Each of the forty runs in the ensemble projects

⁹Note that we base these calculations only on land planted with the sixteen crops in our dataset.

¹⁰In our dataset, the size of the average Indian district is 9,000 square kilometers and the median district is 7,900 square kilometers.

¹¹We classify barley, chickpea, linseed, rapeseed and mustard seed, and wheat as *rabi* crops. Note that rapeseed and mustard seed are treated as a single crop in the VDSA data.

weather patterns according to the same path of climate change. For the years 1920 to 2005, the simulations are based on historical climate conditions. From 2006 to 2099, climate change is modelled based on the representative concentration pathway 8.5 (RCP8.5) scenario.¹²

The scientific literature has used this data to study the potential distribution of future weather realizations.¹³ Similarly, for our analysis of the distribution of agricultural outcomes under climate change, the availability of this multi-run model is crucial. This is because it provides a large number of realizations of weather variables within a period in which we can reasonably treat the climate as fixed.¹⁴ This allows us to project a distribution of potential outcomes for each crop and district within a given period (decade). To do this, we follow the literature by modelling the realized weather variables in each run of the CESM-LENS simulations as occurring with equal probability (see Appendix A for details).

A second advantage of using the CESM-LENS data is its extended time coverage. This means that we do not need to make assumptions about the timing of the onset of climate change to compare ‘baseline’ and ‘post-climate change’ periods. Instead, we generate projections covering the entire 1920-2099 span.

The CESM-LENS dataset is available at a slightly lower resolution than the ERA5 data (1 degree latitude by 1 degree longitude). We thus adjust the projections to ERA5’s grid resolution using linear interpolation. We then link Indian districts to the weather projection data in the same way as with the ERA5 data (see above). We also account for systematically different prediction errors in the CESM-LENS and ERA5 datasets. To do this, we calculate the daily district-level average difference between ERA5 and CESM-LENS temperature (or precipitation) over the years 1979 to 2015. We then add this difference to the CESM-LENS data for the corresponding district-day combinations in all years (analogously to Deschênes

¹²RCP8.5 is also known as the “business as usual” scenario. It assumes that society will fall short of implementing significant reductions in greenhouse gas emissions. Given actual recent trends in emissions, this scenario seems increasingly likely (UNEP, 2019).

¹³For instance, Diffenbaugh et al. (2015) leverage the multiple realizations in the dataset to assess the changing future probability of the joint occurrence of warm and dry weather conditions in California. Yoon et al. (2015), Ault et al. (2016), Diffenbaugh et al. (2017), Thirumalai et al. (2017) and Swain et al. (2018) similarly use CESM-LENS to assess the future likelihood of historically extreme climate events.

¹⁴For example, while a standard dataset (from either a climate model or actual weather) includes ten observations for annual precipitation in each decade, CESM-LENS has 400.

and Greenstone, 2011). We set any negative precipitation values resulting from the adjustment to be equal to zero.

2.3 Summary statistics

Figure 1 displays the projected evolution of the temperature (top panel) and precipitation (bottom panel) distributions in the adjusted CESM-LENS dataset. We begin by calculating the mean of daily temperature across district-days, as well as the mean of total annual precipitation across districts, within a given run-year. We then plot the predicted changes in these two India-wide measures over time, representing each run with a different line. The thick line in each figure depicts the mean value across runs in each year. The figure shows that during the 21st century, the temperature distribution is projected to shift sharply upwards around a rising mean. In the bottom panel, we see that the mean of the precipitation distribution is also forecast to rise. However, substantial changes in variability are not readily apparent in either panel.¹⁵

Table 1 shows predicted changes in temperature and precipitation distributions for the average Indian district, by decade from the 1920s to the 2090s. On average across districts, mean temperature is projected to rise more than four degrees Celsius between the 2000s and the 2090s. The data also predicts a 9.8% rise in standard deviation over this period. This is within the range of variation in standard deviation observed in the CESM-LENS data between the 1920s and the 1990s. Meanwhile, there is a forecast rise of approximately 14% in the mean of annual precipitation for the average district during the 21st century, and its standard deviation is projected to increase by more than a third.

The extent of overlap in the distributions of temperature and precipitation in the ERA5 and CESM-LENS data will be important in determining the difficulty of translating our regression estimates into yield projections. Figure 2(a) displays box plots representing the distribution

¹⁵Figure A2 displays the heterogeneity in these projections across Indian regions. The largest temperature rises are forecast to occur in northern India. At the same time, precipitation is predicted to increase most steeply in India's northeast.




Figure 1: Distributions of mean daily temperature and total annual precipitation

The thick line represents mean temperature across district-days (top panel) and mean annual precipitation across districts (bottom panel), averaged across runs, for each year in the CESM-LENS dataset. The thin lines represent the path of each run considered individually.

of temperature across district-days, comparing ERA5 and CESM-LENS data for 1979 to 2015 with CESM-LENS data for the decade starting in 2090.¹⁶ This confirms that the ERA5 and CESM-LENS data project a similar distribution of temperature and precipitation between 1979 and 2015. This is partly by construction, because we adjust the CESM-LENS data as discussed above. However, due to the large upward shift in the CESM-LENS temperature distribution over time, the median for the 2090s exceeds the 75th percentile of the 1979-2015 distribution. Moreover, many temperatures above the 75th percentile in the 2090s are in the

¹⁶Figure A3 shows decade-by-decade box plots of the distributions of temperature and precipitation in the CESM-LENS data.

Table 1: Average changes in temperature and precipitation distributions across districts

	(1)	(2)	(3)	(4)
	Average temperature		Total precipitation	
	Mean	S.D.	Mean	S.D.
1920s	23.84	0.56	1.29	0.20
1930s	23.83	0.57	1.29	0.21
1940s	23.89	0.58	1.29	0.21
1950s	23.84	0.59	1.29	0.22
1960s	23.67	0.61	1.29	0.21
1970s	23.77	0.63	1.30	0.21
1980s	23.75	0.64	1.30	0.21
1990s	23.76	0.64	1.28	0.21
2000s	24.02	0.61	1.29	0.21
2010s	24.27	0.61	1.29	0.21
2020s	24.63	0.61	1.30	0.21
2030s	25.12	0.66	1.32	0.22
2040s	25.70	0.65	1.32	0.23
2050s	26.25	0.67	1.35	0.25
2060s	26.93	0.66	1.37	0.26
2070s	27.44	0.64	1.41	0.27
2080s	27.96	0.64	1.42	0.27
2090s	28.46	0.67	1.47	0.29

This table displays the mean values of four variables across the 310 Indian districts in the sample, using the CESM-LENS climate projection data. In the first two columns, the relevant variables are constructed by first calculating the average of daily temperatures (in degrees Celsius) in each district-run-year, and then taking the mean (column (1)) or standard deviation (column (2)) of this across the 400 runs in each decade. For the last two columns, we begin by finding total annual precipitation (in meters) in each district-run-year, and then calculate the mean (column (3)) or standard deviation (column (4)) across the 400 runs per decade.

upper tail of the distribution for 1979 to 2015. Figure 2(b) shows that an issue of distributional overlap is also present in the case of precipitation, though it is much less pronounced.

3 Constructing projected yields

In this section, we present our regression specifications relating crop yields to weather variables, briefly discuss the regression results, and describe how we generate yield projections for the years 1920 to 2099.




Figure 2: Distributions of average temperature and total precipitation

This figure displays distributions (5^{th} , 25^{th} , 50^{th} , 75^{th} and 95^{th} percentiles) of average temperature across district-days and total precipitation across district-years, for the datasets and periods indicated.

3.1 Yield-weather specifications

To estimate the relationship between weather and yield for each crop, we run regressions using the ERA5 and VDSA data from 1979 to 2015. This produces crop-level estimates of the effects of temperature and precipitation on yields in this baseline period. We then use these estimates to project the evolution of the distribution of potential yields from 1920 to 2099 based on the CESM-LENS data.

A key challenge in projecting yields is that many of the weather realizations in the CESM-LENS data for the later part of the 21st century are rarely observed in the ERA5 data. As displayed earlier (in Figure 2(a)), this is especially true for temperature. We therefore attempt several different regression specifications, which handle the problem of estimating the responses of yields to high temperatures in different ways.

We estimate the following specification separately for each crop:

$$\ln A_{isgt} = F(\beta, T_{isgt}) + \phi_g P_{isgt} + \psi_g P_{isgt}^2 + \alpha_{isg} + \gamma_{sg}t + \delta_{sg}t^2 + \epsilon_{isgt} \quad (1)$$

On the left-hand side, $\ln A_{isgt}$ represents the logarithm of yield (quantity produced per unit area) for district i (in state s), crop g and year t . As regressors, we include a set of temperature and precipitation variables similar to those in Schlenker and Roberts (2009). $F(\beta, T_{isgt})$ is a function of a vector of parameters β (to be estimated) and the distribution of temperature facing crop g in a given district and year, T_{isgt} . This function will take various forms, to be described in more detail below.

Our specifications also include a quadratic function of total monsoon precipitation P_{isgt} , on which all crops (whether grown during or after the monsoon season) may depend. In the case of *rabi* crops, we also include an additional quadratic function of total precipitation from October to April, i.e. during the growing season itself.¹⁷ We also add crop-specific district fixed effects α_{isg} and crop-state-specific quadratic time trends $\gamma_{sg}t + \delta_{sg}t^2$, as in Schlenker and Roberts (2009). In other words, our model partials out district-specific factors and local trends before estimating the elasticity of agricultural productivity to weather.

We exclude district-crop-year observations for which zero area planted or zero quantity are reported, or area or quantity values are missing. We also winsorize values of log yield that are more than one standard deviation above the 99th percentile or more than one standard deviation below the 1st percentile of the distribution for a given crop. We cluster standard errors at the level of the district.

3.1.1 Temperature bins

In our baseline version of equation (1), we define $F(\beta, T_{isgt}) = \sum_{k=1}^K \beta_g^k T_{isgt}^k$. Here, T_{isgt}^k represents the share of days in the crop's growing season falling into temperature bin k in a

¹⁷For *rabi* crops, we use only the four months before this growing season (June to September) to calculate the monsoon precipitation variable. For sugarcane, given its year-round growing season, we include only a quadratic function of total precipitation for the entire year.

given district and year. We generate these bins using the deciles of the 1979-2015 temperature distribution for each growing season. More specifically, we first pool all district-days over this whole period, and calculate deciles of the distribution of average daily temperature observed in the ERA5 data. We then define bins' upper and lower boundaries according to these deciles, so that 10% of district-day observations fall in each bin. We do this separately for each growing season (*kharif*, *rabi* and year-round), so our temperature bins differ for crops grown in different seasons. Table A2 shows the temperatures covered by each bin for each growing season.

This approach has the advantage that we use a substantial share of our data to estimate the responses of yields to the highest observed temperatures. It also assures that bin widths are not set arbitrarily relative to the actual distributions of our weather variables. However, it could mechanically depress variance in projected yields for each district-crop. This is due to the fact that as temperatures rise relative to the 1979-2015 baseline period, they will tend to be spread across fewer bins. In particular, a larger share of temperature realizations will enter the highest temperature bin, where the relationship between temperature and log yield is flat by assumption.

3.1.2 Degree days

In our second approach to modelling the yield-temperature relationship, we instead remove the temperature bins and replace these with a 'degree-days' (DD) variable, so that $F(\beta, T_{isgt}) = \beta DD_{isgt}$. In this case, we calculate the difference between the observed temperature and a baseline temperature for each district-day during a crop's growing season, setting this equal to zero if the temperature in that district-day is below the baseline level. We then add all of these daily values together for each district-year.

The degree-days approach estimates the convexity in the yield-temperature relationship at high temperatures in part by exploiting observed convexity at relatively lower temperatures. At the same time, it has the disadvantage that it constitutes a more restrictive parameterization. This is because only one coefficient captures the effect of temperature on yield in this

specification.

In choosing an appropriate baseline temperature value, we face a tradeoff. With a lower baseline temperature, more data is used to estimate the convexity of the yield-temperature relationship. However, the average data point employed to do so becomes farther away from the high temperatures whose effects we want to capture. We use 24 degrees Celsius as the baseline temperature in our main degree-days specification. This is the approximate mean temperature observed in the ERA5 weather data during the 1979-2015 sample period. We also explore the sensitivity of our findings to using adjacent values (23 degrees or 25 degrees).

3.1.3 Robustness specifications

We check the robustness of our results to several alternative ways of allowing for convexity in the relationship between temperature and yield at the highest temperatures. The first robustness specification adopts the temperature bins used in Schlenker and Roberts (2009). Specifically, we define a set of three-degree bins from 0 to 39 degrees Celsius, along with two additional bins covering all lower and higher temperatures. A second version instead augments our baseline decile bins specification by adding a ‘high-degree days’ (HDD) variable. This is calculated from the size of the gap between very high observed daily temperatures and the lower limit of the top decile bin. Appendix B provides further details of these specifications. In addition, we run versions of our degree-days regressions with alternative baseline temperatures of 23 and 25 degrees Celsius, as mentioned above.

We also specify the yield-precipitation relationship in two alternative ways. Our first approach is based on the finding of Fishman (2016) that the within-season distribution of precipitation may also be important for yields in India. We add a vector of three additional variables used in that study. These are the count of the number of days with precipitation above 0.1 millimeters, an analogous count of days with especially heavy rain (above 100 millimeters), and the length of the longest dry spell (i.e. the maximum number of consecutive days with less than 0.1 millimeters of rain). Each of these variables is calculated for the

growing season of the relevant crop. Because heavy rain is rare during the *rabi* season, we drop this variable for the five *rabi* crops.

We also attempt specifications replacing the quadratic function of precipitation with a set of precipitation bins, as in our baseline approach to the yield-temperature relationship. To do so, we classify all district-years into deciles based on the total monsoon precipitation in each district-year. We then define bins corresponding to these deciles and include them in our regressions. For *rabi* crops, we include an additional set of bins for precipitation during the *rabi* growing season.

Finally, we take into account the potential concern that the relationships between yields and weather variables might have changed during the 1979-2015 period, due to factors such as the spread of irrigation and hybrid seeds. If these relationships became weaker over time, then we might want to exclude variation from the early part of the period from our estimates, because it might lead us to overestimate the projected severity of yield losses from future climate change. So we also attempt a robustness check in which we estimate our baseline specifications using data from the second half of the sample period only (i.e. from 1997 onwards).

3.2 Yield-weather regression results

We now briefly summarize the results of the regressions outlined above. Figure 3 displays estimates of the coefficients β_g^k on the decile temperature bins, along with 95% confidence intervals, for each of our sixteen crop-specific regressions. The omitted category is the fourth decile, which corresponds to temperatures in the low to mid-20s depending on the crop's growing season (see Table A2). As found in the previous literature, a greater share of days at high temperatures tends to be worse for crops than more frequent exposure to temperatures in this lower range. Table A3 shows these results in a table format, along with the estimated coefficients on the precipitation variables. For almost all crops, the relationship between yield and total monsoon precipitation displays the expected concavity; i.e. the estimated coefficients tend to be positive for the linear term but negative for the quadratic term.




Figure 3: Estimates of coefficients on temperature bins using decile temperature bins

This figure shows the estimated coefficients on temperature bins and their corresponding 95% confidence intervals, using the baseline decile temperature bins specification. Each panel represents a separate regression for one of the sixteen crops included in our data. The crops are ordered by growing season: the first ten crops (castor to sorghum) are *kharif* (monsoon) crops, the next five (barley to wheat) are *rabi* crops and sugarcane is grown year-round. The temperatures covered by each bin vary by growing season, as discussed in Section 3; see Table A2 for the limits of each bin in degrees Celsius for each growing season.

For the baseline specification with degree days, we find the estimated coefficient on the DD variable to be negative and statistically significant at the 1% level for all sixteen crops (Table A4). Also, we again observe a concave relationship between yield and monsoon precipitation for almost all crops.

Results of our robustness specifications may be found in Tables A5 to A14. Among these results, one notable finding is apparent from a comparison of our degree-days regressions with baseline temperatures of 23 degrees and 25 degrees (Tables A13 and A14). From these specifications, we see that the t-statistic for the estimated coefficient on the degree-days variable tends to decline in magnitude as the baseline level rises. This result is in line with the intuition suggested in the previous subsection: the higher the baseline, the less data we use to estimate this coefficient.

Also notable is that when we add the precipitation variables from Fishman (2016) to our main specifications, we find that crops grown in the monsoon season benefit from a larger number of rainy days during that time. On the other hand, yields of *rabi* crops tend to diminish when rain occurs more frequently during their growing season. Some *rabi* crops appear to benefit from longer dry spells, but these have no statistically significant effects for most crops. The estimated impacts of episodes of heavy rain are also inconsistent across crops.

3.3 Yield projections

We use the estimates in the previous subsection to construct projected yields based on the CESM-LENS data. Essentially, our projections insert the climate of each year from 1920 to 2099 into the India of the year 2000. To do this, we take the logarithm of observed yield for each district and crop in 2000. We then shift log yield based on the projected difference in weather between 2000 (based on ERA5) and each run-year in the CESM-LENS dataset.

Specifically, we project yields according to the following equation:

$$\ln \tilde{A}_{isgrt} = \ln A_{isg,2000} + F(\hat{\beta}, \Delta T_{isgrt}) + \hat{\phi}_g \Delta P_{isgrt} + \hat{\psi}_g \Delta P_{isgrt}^2 \quad (2)$$

where $\ln \tilde{A}_{isgrt}$ is the projected log yield for district i and crop g in run-year rt of the CESM-LENS data, while $\ln A_{isg,2000}$ is the log yield for that district-crop from the VDSA data for 2000. Here, ΔP_{isgrt} represents the difference between total precipitation in run-year rt (according to CESM-LENS) and actual precipitation in 2000 (according to ERA5) for crop g in district i . Changes in the other weather variables are calculated analogously. We create these projections separately for each of our yield-weather regression specifications, based on the estimated coefficients from each of these regressions.

The fact that we hold factors other than weather fixed at their 2000 levels has two key implications for the interpretation of our results. First, our study does not take into account the possibility of technological change in response to evolving climatic conditions. Instead, we set aside this dimension of adaptation to focus on the direct impact of changes in temperature and precipitation patterns on the distribution of potential agricultural outcomes. Holding technology fixed is a standard approach in this literature (e.g. Deschênes and Greenstone, 2007; Schlenker and Roberts, 2009; Hsiang, 2016). However, we will allow for another possible dimension of adaptation – changes in crop choice – using a general equilibrium model in Section 5.

Second, this projection method eliminates all variation in yields across run-years for a given district, other than the variation due to differences in weather conditions. So when we discuss the projected evolution of the distribution of agricultural outcomes, we will only be considering changes in weather-induced variability rather than changes in the total variability of these outcomes.

We also generate two other sets of projections in which we hold the distribution of either temperature or precipitation fixed, while allowing the distribution of the other variable to evolve with the changing climate. For example, in our ‘temperature-only’ projections, we fix the precipitation variables for each run at their 2000 levels, according to the CESM-LENS data for that run. This allows us to shed light on whether temperature or precipitation is mainly responsible for weather-induced changes in agricultural outcomes.

4 Distribution of potential agricultural revenue

In this section, we use our yield projections to consider the impact of climate change on the distribution of potential agricultural outcomes in each Indian district. To summarize these outcomes in a single measure, we calculate projected nominal agricultural revenue for each district in each CESM-LENS run-year, holding area planted and crop prices constant at their levels in 2000. In other words, we multiply the projected yield of a crop in a district-run-year by the area planted with that crop in the district in 2000 and its farm-gate price, and then take the sum across crops. We apply uniform crop prices countrywide, by taking the average across districts of the farm-gate prices observed in the VDSA data for each crop in 2000. For each decade from the 1920s to the 2090s, we perform this revenue calculation for the 400 available run-years for each district (forty runs per year over ten years).

4.1 Mean revenue

We first look at the evolution of mean agricultural revenue, i.e. the expected agricultural revenue of a district in an ‘average year’. We take the simple average of the 400 potential revenue realizations for each district-decade, giving us a projection for the mean of the distribution of potential agricultural revenue. To calculate the weather-induced changes in mean revenue by decade for the average district, we regress the logarithm of mean revenue on district fixed effects and decade fixed effects. Figure 4(a) displays the decade fixed effects from this regression.

We also present 95% confidence intervals for these decade fixed effects that incorporate the uncertainty in the estimated coefficients from our yield-weather regressions. To do so, we use a block bootstrap approach with resampling by district, which effectively replicates our district-level clustering of standard errors. In each of 500 bootstrap runs, we draw 310 districts from our sample with replacement, and then run each crop-specific regression (using a given specification) with that sample. Next, we calculate projected mean agricultural revenue by district-decade for that bootstrap run. Finally, we plot the 2.5th and 97.5th percentiles,

across all bootstrap runs, of each of the estimated decade fixed effects.

The left side of Figure 4(a) shows results of projections using our baseline yield-weather specification; i.e. the specification with ten bins based on deciles of the temperature distribution and a quadratic in total precipitation. We see a substantial decline in projected mean revenue as a result of climate change during the 21st century, in contrast to the relative stability of the 1920s to the 1990s. On average across districts, the predicted proportional fall in mean revenue from the 2000s to the 2090s is 32.3%.

The right side of Figure 4(a) also shows the predicted evolution of mean agricultural revenue for the average district, this time using our specification modelling the yield-temperature relationship in terms of degree days. Similar to the projections using temperature bins, we see an estimated 23.9% decline in mean agricultural revenue during the 21st century. As the maps in Figure A4 show, the projected heterogeneity in this effect across districts varies somewhat between these two specifications. This is due to differences in the estimated yield-weather elasticities by crop.

Figure 4(b) displays analogous estimates for projections in which either temperatures or precipitation levels are held constant at their values as of 2000. This shows that temperature changes drive the entire projected effect of climate change on mean agricultural revenue. In contrast, mean revenue is forecast to rise in the average district due to precipitation changes, though the estimated increase is an order of magnitude smaller than the temperature effects. This result is in line with the findings of Vogel et al. (2019), who conclude that extreme temperature has been more important than precipitation in driving yield anomalies worldwide.

These results remain very similar when we instead: i) use either of our other two specifications incorporating temperature bins; ii) shift the baseline temperature in the degree-days specification; iii) use precipitation bins; iv) add additional precipitation variables into our regressions; or v) use estimates based only on weather and yield data from 1997 onwards. See Figure A5 for a summary of all of these robustness checks. Overall, we project a gradual but substantial decline in mean revenue during the 21st century for the average district, ranging

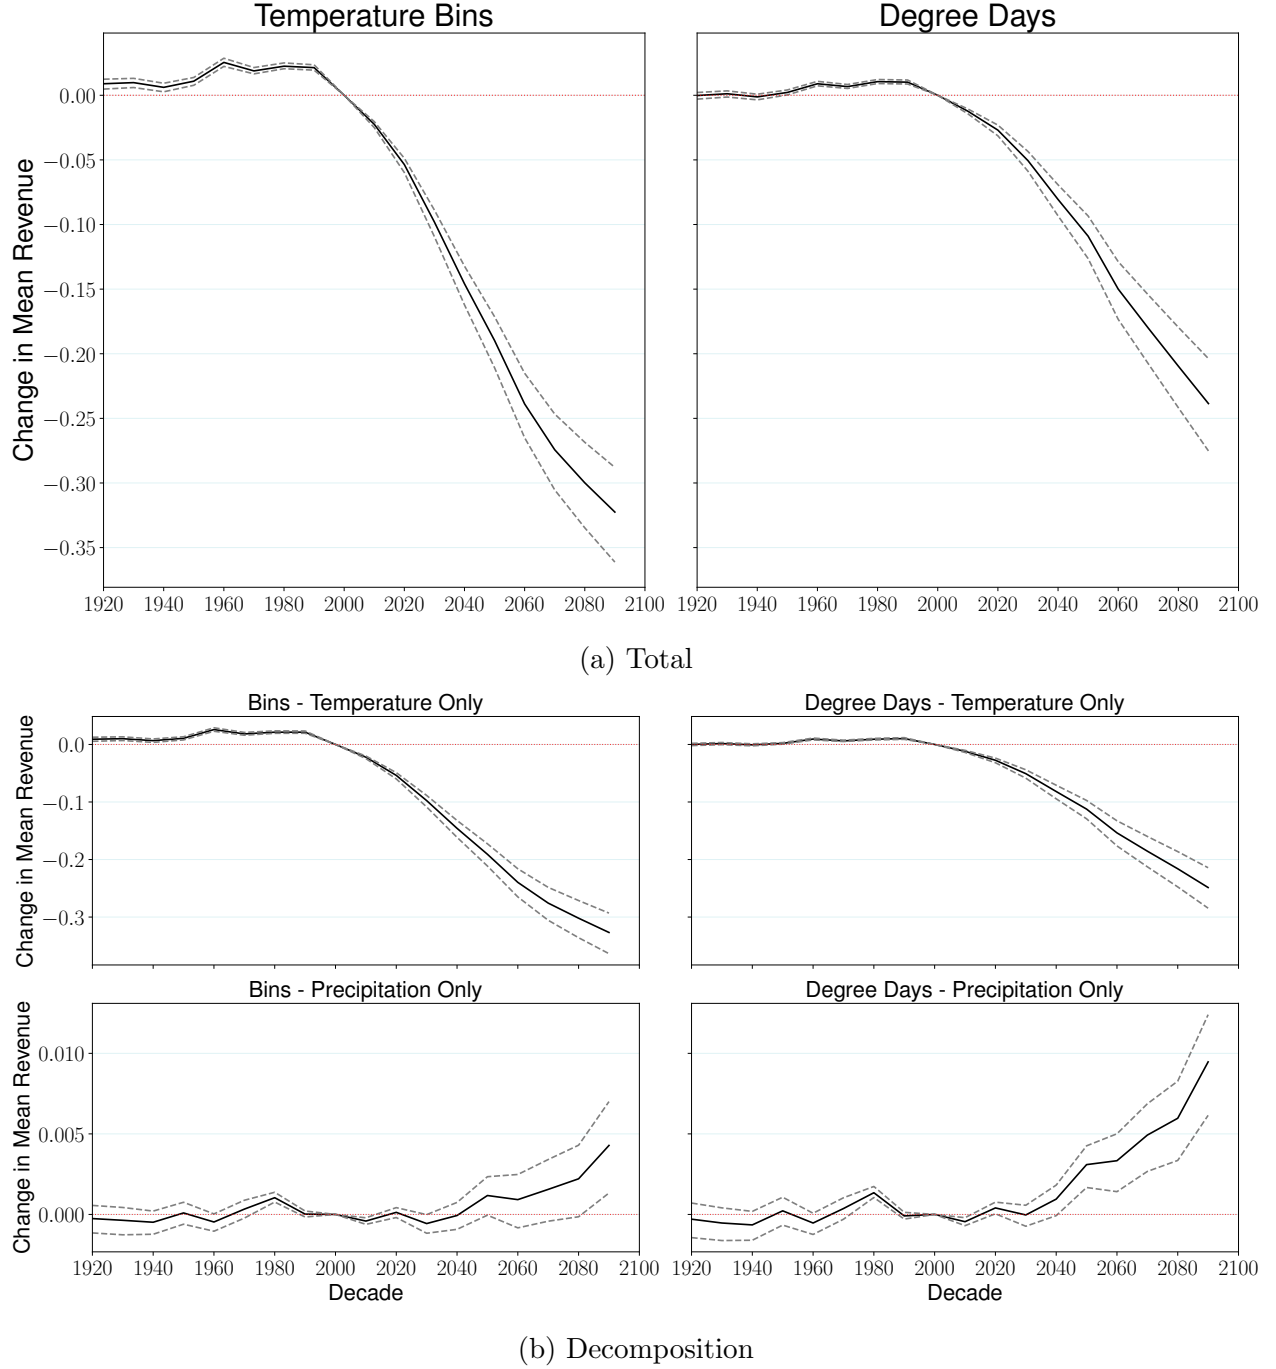


Figure 4: Projected changes in mean agricultural revenue (in percent relative to 2000s)

This figure shows the projected proportional changes in mean agricultural revenue (in percent) relative to the 2000s for the baseline models: decile temperature bins (left) and degree days (right). Dashed lines display 95% confidence intervals based on a block bootstrap, as explained in Section 4. Panel (a) presents the total projected effect of climate change, and Panel (b) displays analogous estimates for projections in which only temperature or precipitation are allowed to change from their values as of 2000.

from 15.5% to 33.1% depending on the yield-weather specification used.¹⁸ Again, we find that this is almost entirely due to predicted changes in temperature. These magnitudes are within the range found in other studies that consider the likely impact of climate change on expected yields in India (e.g. Guiteras, 2009; Burgess et al., 2017; Hari et al., 2018).

4.2 Standard deviation of revenue

We next investigate whether our projections indicate a substantial change in weather-induced revenue variability. In other words, we examine how climate change is predicted to affect the second moment of the distribution of potential agricultural revenue. For each decade of CESM-LENS data, we consider the standard deviation of the 400 potential revenue realizations for each district. As discussed earlier, the ability to consider variation in outcomes across so many potential realizations is a major advantage of the CESM-LENS dataset.

When interpreting the results in this subsection, it is important to recall that the variation in projected yields across runs is only due to weather. This is because we hold other determinants of yield (conditional on weather) fixed across realizations. So we will interpret projected rises or declines in the standard deviation of revenue as proportional changes in weather-induced variation rather than total within-district variation.

Using our baseline temperature bins specification, we project that the average district will experience a large decline in yield variation due to climate change, as measured by the standard deviation of our revenue measure. After remaining relatively stable from the 1920s to the 1990s, the standard deviation in agricultural revenue due to weather is predicted to fall by 59.9% between the 2000s and the 2090s, as shown in the left panel of Figure 5(a). This result is again driven by changes in temperature rather than precipitation patterns (see Figure 5(b)).

We noted in Section 3.1 that this specification might mechanically decrease projected

¹⁸In almost all cases, confidence intervals are also similar in width to those in Figure 4(a). The only exception is the specification with three-degree temperature bins, where some of our estimates for very high temperatures are imprecise.

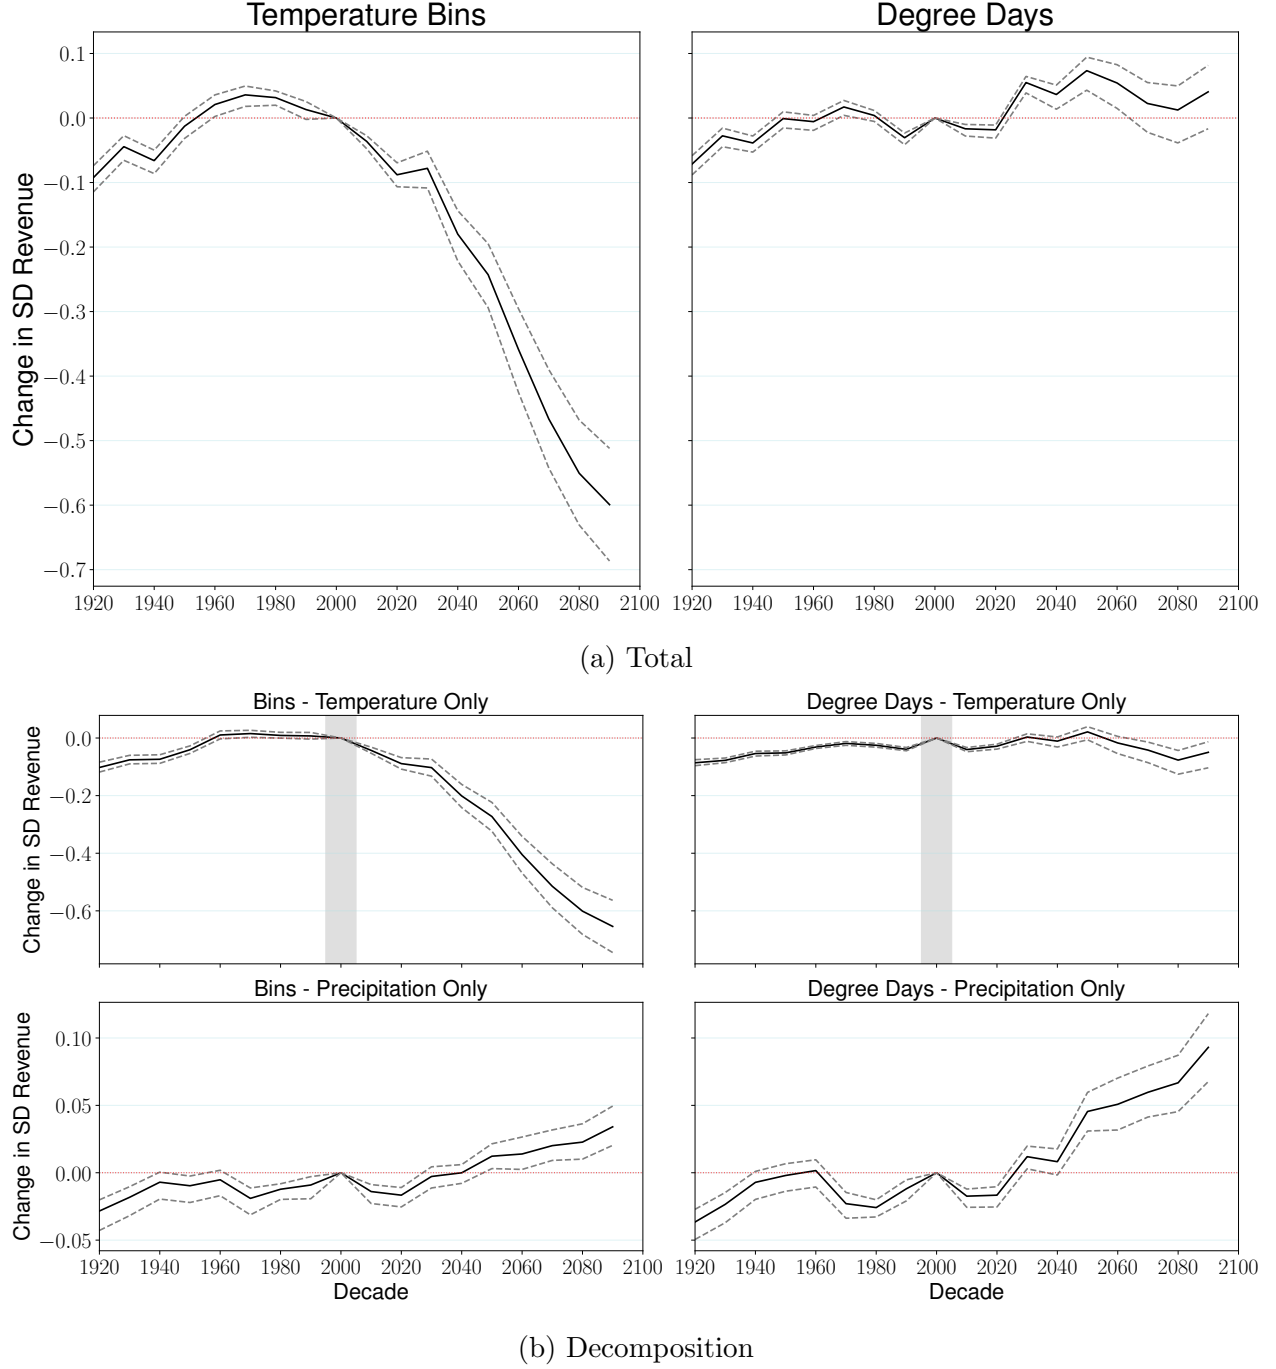


Figure 5: Projected changes in the weather-induced standard deviation of agricultural revenue (in percent relative to 2000s)

This figure shows the projected proportional changes in the weather-induced standard deviation of agricultural revenue (in percent) relative to the 2000s for the baseline models: decile temperature bins (left) and degree days (right). Dashed lines display 95% confidence intervals based on a block bootstrap, as explained in Section 4. Panel (a) presents the total projected effect of climate change, and Panel (b) displays analogous estimates for projections in which only temperature or precipitation are allowed to change from their values as of 2000.

variation in yield over time, as temperatures rise and most observations are spread across fewer bins. We show evidence that this is indeed the case in Figure A6, which plots the India-wide shift in the share of days (year-round) in each decile bin for each decade as compared to the 2000s. However, our other two temperature bins specifications, which aim to address this concern, do not yield substantially different results (see Figure A7). Overall, we find that when specifications with temperature bins are used, volatility in yield due to temperature is projected to decline in the future for the average Indian district.

In contrast, the parametric assumption underlying our degree-days specification is that each additional degree Celsius above 24 degrees on a given day has the same impact on the log yield of a given crop. However, although projections using this specification indicate a future rise in the weather-induced standard deviation of agricultural revenue, the magnitude of the predicted impact is small. As shown in the right panel of Figure 5(a), the average district is projected to see an increase in this measure of just 4.1% between the 2000s and the 2090s. The model with additional precipitation variables yields a marginally larger rise in variability, as does the model with a baseline temperature of 25 degrees (Figure A7).

The wide range of estimates from our various specifications – ranging from a decline of 68.4% to a rise of 10.5% in the weather-induced standard deviation of revenue for the average district in the 21st century – indicates the challenges of projecting future variability in yields under climate change. Specifically, our findings suggest that such second-moment projections are very sensitive to how the relationship between weather and yield is modelled, especially for high temperatures. However, across all of our exercises, we do not find strong evidence that weather-related output volatility is likely to increase substantially for Indian farmers.

4.3 Downside risk from extreme weather

Poorly insured farmers in a developing country may be most concerned about increases in downside risk – the frequency of ruinously bad years – rather than a widening range of both positive and negative outcomes due to rising variance. We therefore also use our projections to

forecast the evolution of extreme ‘1-in-100-years’ harmful weather realizations as the climate changes.

To do so, we take the first five decades of the CESM-LENS data (the 1920s through the 1960s) as a baseline. For each district, we define a 1-in-100 bad year for weather based on the projected agricultural outcomes of all 2000 run-years we observe during this baseline period (forty per year for fifty years). Specifically, we calculate agricultural revenue in a given district for each of these 2000 weather realizations, and then set the threshold for a 1-in-100 year to be the first percentile of this revenue distribution. For each decade up to the 2090s, we then calculate the share of the district’s projected revenue realizations (out of the 400 run-years in each decade) that fall below this threshold. We perform these calculations separately for each of the 310 districts in our sample, so that the threshold for a 1-in-100 bad year is specific to each district.

The left panel of Figure 6(a) displays the estimated decade fixed effects from a regression of the share of 1-in-100 bad years on decade and district fixed effects, using the projections from our baseline specification with temperature bins. We find that in the late 20th century, bad years remain as rare as in the baseline period for the average district. However, the share of realizations that would have been classified as ‘extreme weather’ between 1920 and 1969 is projected to become progressively larger during the 21st century. By the 2090s, what had been 1-in-100 bad years are predicted to occur at a rate that is 88 percentage points higher than in the 2000s.

These predictions are not unique to our baseline temperature bins specification. The right side of Figure 6(a) displays an analogous projection using the degree-days specification, which predicts an 84 percentage point rise in the share of 1-in-100 weather years during the 21st century for the average district. Projections using our robustness specifications similarly imply a substantial increase in the incidence of extremely low-revenue years (see Figure A8).

In all of our specifications, this very large projected increase in the occurrence of bad years is driven by changes in temperature rather than precipitation patterns. As Figure 6(b) shows,

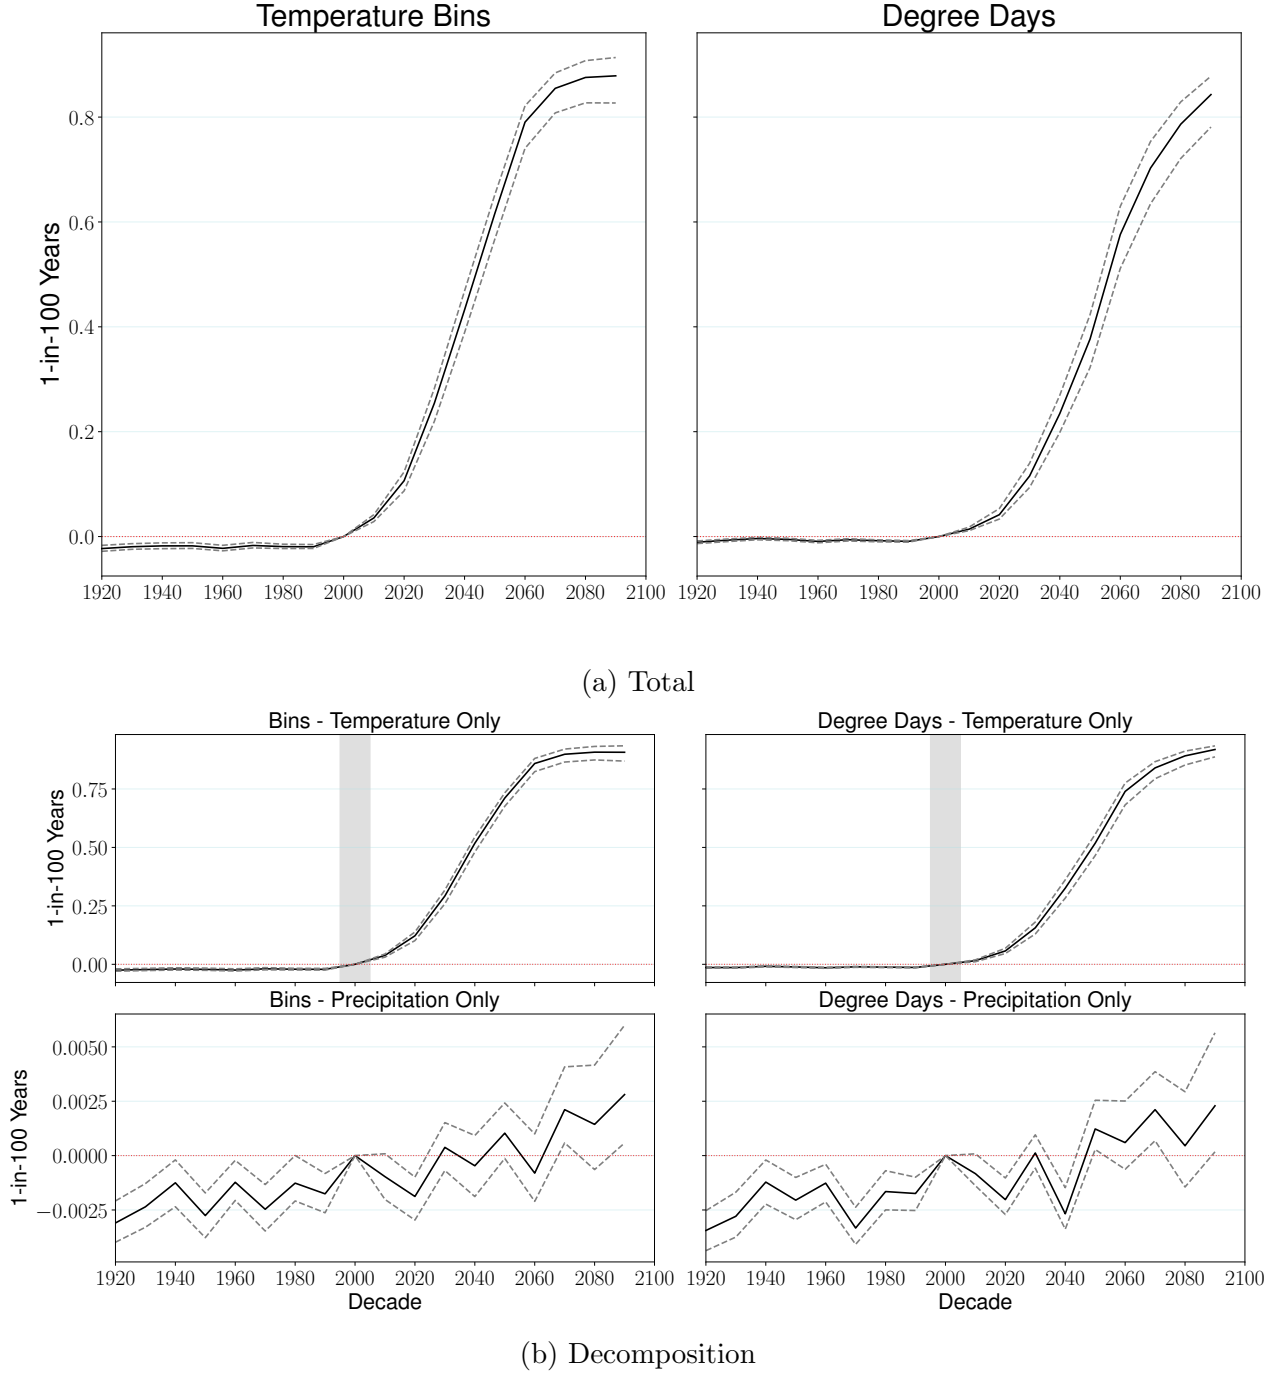


Figure 6: Projected changes in downside risk – share of 1-in-100 bad weather years for agricultural revenue (relative to 2000s)

This figure shows the projected changes in the share of 1-in-100 bad weather years for agricultural revenue (in percentage points) relative to the 2000s for the baseline models: decile temperature bins (left) and degree days (right). Dashed lines display 95% confidence intervals based on a block bootstrap, as explained in Section 4. Panel (a) presents the total projected effect of climate change, and Panel (b) displays analogous estimates for projections in which only temperature or precipitation are allowed to change from their values as of 2000.

the predicted rise in the share of 1-in-100 years due only to temperature changes is nearly identical to the scenario pictured in Figure 6(a).

Overall, we find that during the 21st century, climate change will increase the share of 1-in-100 bad agricultural years in the average Indian district by 54 to 89 percentage points. These results are consistent with the fact that we observe a substantial negative impact on mean agricultural revenue in all of our projections but do not see a steep rise in volatility. Essentially, formerly extreme bad weather is predicted to become the ‘new normal’.

5 Endogenizing prices, crop choice and trade

In our analysis so far, we have not accounted for the facts that crop prices and farmers’ planting choices are endogenous, and that agricultural goods are tradable. In practice, farmers may be able to mitigate the effects of climate change by altering their crop choice decisions and through trade.¹⁹ Moreover, price adjustments might offset part of the effect of changing temperature and precipitation patterns. To check whether our main results are sensitive to these issues, we conduct a structural analysis of the impact of climate change on Indian agriculture, allowing for endogenous prices and crop choices as well as trade between districts.²⁰ Despite the additional structure required for this exercise, we find that the key insights from the previous section are preserved.

5.1 Summary of model

We use the model and methodology of Allen and Atkin (2016). This model combines a Ricardian framework, incorporating many locations and goods, with portfolio choice by agents across risky assets. Here, we explain the setup and main implications of the model in intuitive

¹⁹The study of Costinot et al. (2016) suggests that adaptation through changing crop choice could partly offset the potential welfare losses from climate change of farmers worldwide. Aragón et al. (2021) show that subsistence farmers in Peru adjust the crops they plant to mitigate the impacts of extreme heat events.

²⁰Note that our approach only allows farmers in a given district to substitute across crops that we observe to have been grown in that district in the baseline year of 2000, according to the VDSA data. We thus abstract from extensive-margin changes in the set of crops grown in each district.

terms. An exposition of the model's key equations is provided in Appendix C.

In the model, farmers in a given district face risk when growing crops, in the sense that each crop's yield in the farmer's district is stochastic. Specifically, the yield of a particular crop in a given location is lognormally distributed with a mean and variance known to the farmer. Each farmer is also aware of the covariance of these potential yields across crops within their local district. So farmers face risk but not uncertainty: they are aware of the distributions of potential yields they face, but not the yields that will actually be realized. Moreover, any given crop may have a different productivity distribution in different districts.

The number of farmers may differ across districts. All farmers are geographically immobile, so the population of farmers in each district is fixed. Moreover, farmers all have identical preferences, which implies that all farmers within each district make identical decisions. All key outcomes of the model can therefore be described at the district level.

Given the local distribution of potential yields, farmers choose the share of land planted with each crop to maximize their expected utility. Their preferences are assumed to be characterized by constant relative risk aversion (CRRA). At the same time, because the model assumes that the yield distribution is lognormal, the real returns of the portfolio of crops held by each farmer are also approximately lognormal. Together, these two key assumptions make the model tractable, because they imply that farmers' expected utility depends only on the log of mean real returns and the variance of log real returns. The relative importance of the latter is governed by farmers' degree of risk aversion. This optimization problem thus involves the usual tradeoff between mean and variance in models of risky portfolio choice: a crop bringing greater risk to the portfolio should have a higher mean return.

The model goes beyond a single-agent portfolio choice framework by allowing farmers in different districts to interact through trade in crop output, and allowing prices to adjust endogenously to clear markets. As usual in Ricardian models, a district tends to produce and export the crops in which it has a comparative advantage in productivity (depending on the level of demand for each crop). However, this finding is now more subtle than in a

standard model, because a district’s comparative advantage is affected by the variance in the local yield of each crop – its riskiness – as well as the mean yield. Therefore, farmers focus on the production of goods in which they have a ‘risk-adjusted’ comparative advantage.

5.2 Quantitative exercise

We conduct a quantitative analysis based on this model for each decade from the 1920s to the 2090s. Specifically, we separately generate counterfactuals for each decade by inserting the distribution of potential yields in that decade into the model. As in our revenue exercises in Section 4, we define these decade-by-decade distributions based on our projections from Section 3.3 (i.e. the yield projections constructed from the CESM-LENS data and our yield-weather regressions). We hold all other parameters of the model fixed at values corresponding to our baseline decade, the 2000s. We discuss the calibration or estimation of each of these parameters in detail in Appendix C.

We generate counterfactual outcomes for each decade by allowing crop choices in each district to be based on the distribution of potential yields in that decade, and allowing for trade across Indian districts and prices to adjust accordingly. We then compare farmers’ projected welfare in each decade (based on the expected utility function discussed in Section 5.1) to projected welfare in the 2000s. Analogously to our revenue exercises in Section 4, we do this by running a regression of our welfare measure on district and decade fixed effects and plotting the decade fixed effects.²¹

Also, in order to assess the quantitative importance of allowing for endogenous planting decisions, trade patterns and prices, we compare the welfare impacts implied by the model to a naive benchmark using the nominal revenue measure from Section 4. Specifically, we define a ‘naive welfare’ measure by using the model-derived formula for the change in expected utility over time, but substituting nominal revenue whenever this formula calls for real returns. Here, we calculate nominal revenue as in Section 4 rather than using the counterfactuals produced

²¹Note that we do not calculate confidence intervals based on a block bootstrap for these estimates, due to the significant computing time required to calculate each counterfactual.


by the quantitative model. As above, we regress this quasi-welfare measure on district and decade fixed effects. We then plot the estimated decade fixed effects from this naive exercise alongside our results using the fully model-based welfare measure.

5.3 Results


Figure 7(a) displays our findings when we employ the baseline temperature bins specification. The naive approach using nominal revenue predicts a gradual but large (34.5%) fall in welfare on average across districts between the 2000s and 2090s due to climate change (left panel). These losses are approximately six percentage points smaller (28.6%) when we instead calculate welfare from the model-based counterfactuals (right panel). Results using the degree-days specification also predict a continuous decline in welfare from the 2000s to the 2090s, with a drop of 21.8% in the model-based measure (see Figure 7(b), right panel). Again, the naive and model-based welfare measures follow similar trends, with the model leading to a moderation of 4.6 percentage points in the predicted extent of the decrease.

Next, as in Allen and Atkin (2016), we separately assess the contributions of mean returns and the variability of returns to changes in welfare or ‘naive welfare’. Using the baseline temperature bins specification, we find a predicted 21st-century welfare decline of 28.9% via changes in the logarithm of mean real returns (Figure 8(a), top right panel) – i.e. the first-moment effect. Again, the insertion of nominal revenue rather than real returns in the authors’ welfare formula modestly increases predicted losses (to 34.6%). Analogous conclusions may be reached with the specification using degree days.

As discussed above, the role of yield variability in farmers’ welfare in the model is captured by the variance of log real returns, a simple summary measure that relies on the assumption of the lognormality of yields across states. This measure is different from the elements of the distribution of agricultural outcomes explored in Section 4, because it is a function of the variance-covariance matrix of the *logarithm* of yields by crop. The resulting welfare effects of climate change via the second moment are very small, as shown in Figure 8(b).



(a) Temperature Bins



(b) Degree Days

Figure 7: Comparison of naive and model-based welfare measures

This figure shows the projected proportional changes (in percent) relative to the 2000s in a naive welfare measure based on nominal returns and welfare from real returns as specified in Section 5. Panel (a) shows the results based on the decile temperature bins specification, and Panel (b) from the degree-days specification.

The main results are unchanged when we use projections based on our other yield-weather specifications (see Figure A9). The naive and model-based welfare measures produce similar trends, and welfare losses are dominated by first-moment impacts. The estimated fall in welfare for the average district between the 2000s and the 2090s ranges from 13.2% to 29.3% across these ten alternative projections.

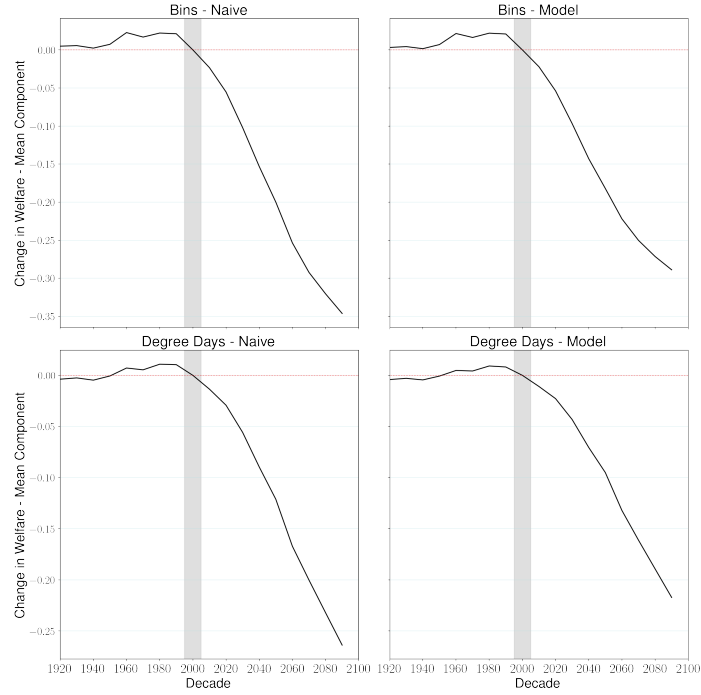
We draw two main conclusions from our exercises using the Allen-Atkin model. First, based on our comparison of the model’s welfare measure to a naive benchmark, we find that endogenizing prices, crop choice and trade flows leads to modest loss mitigation, but not to substantially different implications. Second, the model suggests that farmers’ welfare losses from climate change will be overwhelmingly driven by declines in mean outcomes rather than changes in portfolio risk.

6 Discussion


A key objective of this study is to provide insight into the possibility that climate change will increase insurable risk for farmers. In such a case, increased availability of insurance would be an important policy prescription in response to the changing distribution of potential returns in the agricultural sector. However, across our several different modelling approaches, one robust finding has been a sharp predicted decline in the first moment of this distribution.²² This implies potentially substantial uninsurable losses from the changing climate, whether or not insurance markets are well-developed.

We have also considered several dimensions of the distribution of potential agricultural outcomes other than the mean. These include the standard deviation of agricultural revenue, the frequency of ‘1-in-100’ revenue years, and the risk-driven component of welfare from a portfolio choice model. None of these exercises has provided strong evidence of a large projected rise in variability. Instead, we have forecast that upward shifts in the temperature distribution will

²²As noted earlier, this finding is in line with other studies projecting the effects of climate change on expected yields in India (e.g. Guiteras, 2009; Burgess et al., 2017; Hari et al., 2018).



(a) Mean component



(b) Variance component

Figure 8: Decomposition of welfare measures into mean and variance components

This figure shows the projected proportional changes (in percent) relative to the 2000s in the mean and variance components of a naive welfare measure based on nominal returns and welfare from real returns as specified in Section 5. Panels (a) and (b) show changes driven by the mean and variance components of these measures, respectively. The top part of each panel displays the results based on the decile temperature bins specification, and the bottom part shows results from the degree-days specification.

cause future ‘normal’ agricultural outcomes to resemble 20th-century extremes.

Of course, these conclusions come with some caveats. First, our study relies on projections from a climate model which, like all such forecasts, may or may not be successful in predicting future climate change. Second, as we have shown, predictions about future variability in agricultural outcomes are very sensitive to modelling choices. Third, our conclusions could be different in countries at the technological frontier, such as the US and members of the EU. Finally, we have held technology constant in our exercises. In practice, however, technological changes such as better irrigation or the use of heat-resistant seeds could be crucial modes of adaptation.²³

We nonetheless argue that our results are important to the discussion about the impact of climate change on agriculture. We believe that the key ingredients of our study represent the present state of the art: a climate projection designed for the study of internal variability, a set of yield-weather regressions in line with the current standard in the literature on climate change, and a general equilibrium framework at the frontier of current research. Indeed, we hope that the challenges involved in forecasting the future distribution of potential agricultural outcomes, as presented in this paper, help spur refinements of these various elements.

Moreover, although we have not modelled changes in agricultural technology, our findings still have implications for the desired direction of technological innovation. Specifically, our results suggest that adaptation to changes in the weather expected in an ‘average year’, and especially the rising incidence of formerly extreme temperatures, should be a priority.

²³Various studies have previously investigated possible margins of adaptation in the context of Indian agriculture. For example, Emerick et al. (2016) study the adoption of a new rice variety with improved resistance to extreme weather events (floods). They find this to be successful in avoiding yield losses in years with floods, and in encouraging farmers to make investments that improve yields in years without floods. Fishman (2018) finds that while irrigation can mitigate the negative effect of precipitation variability in India, successful adaptation to climate change is unlikely given current groundwater depletion rates. Similarly, Blakeslee et al. (2020) do not find evidence of successful past adaptation by Indian agriculture to problems of water scarcity. Taraz (2018) concludes that Indian farmers can adapt to moderately hot temperatures, but not to episodes of extreme heat.

References

- Allen, T. and Atkin, D. (2016). Volatility and the Gains from Trade. *NBER working paper*, 22276.
- Aragón, F. M., Oteiza, F., and Rud, J. P. (2021). Climate Change and Agriculture: Subsistence Farmers’ Response to Extreme Heat. *American Economic Journal: Economic Policy*, 13(1):1–35.
- Arkolakis, C., Costinot, A., and Rodríguez-Clare, A. (2012). New Trade Models, Same Old Gains? *American Economic Review*, 102(1):94–130.
- Auffhammer, M., Ramanathan, V., and Vincent, J. R. (2006). Integrated Model Shows that Atmospheric Brown Clouds and Greenhouse Gases Have Reduced Rice Harvests in India. *Proceedings of the National Academy of Sciences*, 103(52):19668–19672.
- Ault, T. R., Mankin, J. S., Cook, B. I., and Smerdon, J. E. (2016). Relative Impacts of Mitigation, Temperature, and Precipitation on 21st-Century Megadrought Risk in the American Southwest. *Science Advances*, 2(10):e1600873.
- Blakeslee, D., Fishman, R., and Srinivasan, V. (2020). Way Down in the Hole: Adaptation to Long-Term Water Loss in Rural India. *American Economic Review*, 110(1):200–224.
- Burgess, R., Deschenes, O., Donaldson, D., and Greenstone, M. (2017). Weather, Climate Change and Death in India. *Mimeo*.
- Burke, M., Dykema, J., Lobell, D. B., Miguel, E., and Satyanath, S. (2015). Incorporating Climate Uncertainty into Estimates of Climate Change Impacts. *Review of Economics and Statistics*, 97(2):461–471.
- Challinor, A. J., Watson, J., Lobell, D. B., Howden, S. M., Smith, D. R., and Chhetri, N. (2014). A Meta-Analysis of Crop Yield Under Climate Change and Adaptation. *Nature Climate Change*, 4(4):287–291.
- Chen, S., Chen, X., and Xu, J. (2016). Impacts of Climate Change on Agriculture: Evidence from China. *Journal of Environmental Economics and Management*, 76:105–124.
- Costinot, A., Donaldson, D., and Smith, C. (2016). Evolving Comparative Advantage and the Impact of Climate Change in Agricultural Markets: Evidence from 1.7 Million Fields Around the World. *Journal of Political Economy*, 124(1):205–248.
- Deryugina, T. and Hsiang, S. (2017). The Marginal Product of Climate. *NBER working paper*, 24072.
- Deschênes, O. and Greenstone, M. (2007). The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather. *American Economic Review*, 97(1):354–385.

- Deschênes, O. and Greenstone, M. (2011). Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US. *American Economic Journal: Applied Economics*, 3(4):152–185.
- Diffenbaugh, N. S., Singh, D., Mankin, J. S., Horton, D. E., Swain, D. L., Touma, D., Charland, A., Liu, Y., Haugen, M., Tsiang, M., et al. (2017). Quantifying the Influence of Global Warming on Unprecedented Extreme Climate Events. *Proceedings of the National Academy of Sciences*, 114(19):4881–4886.
- Diffenbaugh, N. S., Swain, D. L., and Touma, D. (2015). Anthropogenic Warming has Increased Drought Risk in California. *Proceedings of the National Academy of Sciences*, 112(13):3931–3936.
- Dingel, J. I., Meng, K. C., and Hsiang, S. M. (2019). Spatial Correlation, Trade, and Inequality: Evidence from the Global Climate. *NBER working paper*, 25447.
- Emerick, K., de Janvry, A., Sadoulet, E., and Dar, M. H. (2016). Technological Innovations, Downside Risk, and the Modernization of Agriculture. *American Economic Review*, 106(6):1537–1561.
- Fishman, R. (2016). More Uneven Distributions Overturn Benefits of Higher Precipitation for Crop Yields. *Environmental Research Letters*, 11(2):024004.
- Fishman, R. (2018). Groundwater Depletion Limits the Scope for Adaptation to Increased Rainfall Variability in India. *Climatic Change*, 147(1-2):195–209.
- Guiteras, R. (2009). The Impact of Climate Change on Indian Agriculture. *Mimeo*.
- Hari, S., Khare, P., and Subramanian, A. (2018). *Climate, Climate Change, and Agriculture*, pages 82–101. In: Economic Survey 2017–18. Ministry of Finance, Government of India.
- Hsiang, S. (2016). Climate Econometrics. *Annual Review of Resource Economics*, 8:43–75.
- Kala, N. (2017). Learning, Adaptation, and Climate Uncertainty: Evidence from Indian Agriculture. *MIT Center for Energy and Environmental Policy Research Working Paper*, 23.
- Kay, J. E., Deser, C., Phillips, A., Mai, A., Hannay, C., Strand, G., Arblaster, J. M., Bates, S. C., Danabasoglu, G., Edwards, J., et al. (2015). The Community Earth System Model (CESM) Large Ensemble Project: A Community Resource for Studying Climate Change in the Presence of Internal Climate Variability. *Bulletin of the American Meteorological Society*, 96(8):1333–1349.
- Knox, J., Hess, T., Daccache, A., and Wheeler, T. (2012). Climate Change Impacts on Crop Productivity in Africa and South Asia. *Environmental Research Letters*, 7(3):034032.
- Lemoine, D. (2021). The Climate Risk Premium: How Uncertainty Affects the Social Cost of Carbon. *Journal of the Association of Environmental and Resource Economists*, 8(1):27–57.

- Lobell, D. B., Bänziger, M., Magorokosho, C., and Vivek, B. (2011). Nonlinear Heat Effects on African Maize as Evidenced by Historical Yield Trials. *Nature Climate Change*, 1(1):42–45.
- Ministry of Agriculture and Farmers Welfare (2017). *Agricultural Statistics at a Glance 2016*. New Delhi: Government of India.
- Office of the Registrar General and Census Commissioner (2004). *India Administrative Atlas, 1872-2001: A Historical Perspective of Evolution of Districts and States in India*. New Delhi: Controller of Publications.
- Schlenker, W. and Lobell, D. B. (2010). Robust Negative Impacts of Climate Change on African Agriculture. *Environmental Research Letters*, 5(1):014010.
- Schlenker, W. and Roberts, M. J. (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.
- Shrader, J. (2017). Expectations and Adaptation to Environmental Risks. *Mimeo*.
- Swain, D. L., Langenbrunner, B., Neelin, J. D., and Hall, A. (2018). Increasing Precipitation Volatility in Twenty-First-Century California. *Nature Climate Change*, 8(5):427–433.
- Taraz, V. (2018). Can Farmers Adapt to Higher Temperatures? Evidence from India. *World Development*, 112:205–219.
- Thirumalai, K., DiNezio, P. N., Okumura, Y., and Deser, C. (2017). Extreme Temperatures in Southeast Asia Caused by El Niño and Worsened by Global Warming. *Nature Communications*, 8:15531.
- Tigchelaar, M., Battisti, D. S., Naylor, R. L., and Ray, D. K. (2018). Future Warming Increases Probability of Globally Synchronized Maize Production Shocks. *Proceedings of the National Academy of Sciences*, 115(26):6644–6649.
- UNEP (2019). *Emissions Gap Report 2019*. Nairobi: United Nations Environment Programme.
- Urban, D., Roberts, M. J., Schlenker, W., and Lobell, D. B. (2012). Projected Temperature Changes Indicate Significant Increase in Interannual Variability of US Maize Yields. *Climatic Change*, 112(2):525–533.
- Vogel, E., Donat, M. G., Alexander, L. V., Meinshausen, M., Ray, D. K., Karoly, D., Meinshausen, N., and Frieler, K. (2019). The Effects of Climate Extremes on Global Agricultural Yields. *Environmental Research Letters*, 14(5):054010.
- Yoon, J.-H., Wang, S. S., Gillies, R. R., Kravitz, B., Hipps, L., and Rasch, P. J. (2015). Increasing Water Cycle Extremes in California and in Relation to ENSO Cycle Under Global Warming. *Nature Communications*, 6:8657.

Appendix

A Additional data details

Generating 1966 district shapes. To generate information on Indian district centroids, we begin with a GIS shapefile of 2018 Indian district boundaries. We then edit this to reconstruct 1966 districts by hand, using a map of 1971 districts and information on district changes between 1961 and 1971, both from the Indian Administrative Atlas (Office of the Registrar General and Census Commissioner, 2004).

Crop calendar. Ministry of Agriculture and Farmers Welfare (2017) labels each crop as a ‘*kharif* crop’ or ‘*rabi* crop’ in a table displaying crop-specific minimum support prices. Castor, chickpea and linseed are not included in this table, and so we infer the primary growing season of each of these crops using crop calendars found in this publication’s appendix.

CESM-LENS data: additional details. Each of the forty runs in the CESM-LENS ensemble is initiated with slightly different initial conditions (at the level of rounding errors), and due to the chaotic nature of the system, realized weather swiftly evolves independently across runs, conditional on the prevailing climate. However, all runs are subject to the same progression over time in the climate. As the dataset’s creators note, “[a]fter initial condition memory is lost, which occurs within weeks in the atmosphere, each ensemble member evolves chaotically, affected by atmospheric circulation fluctuations characteristic of a random, stochastic process” (Kay et al., 2015). Because the first day of all but one of the forty runs is January 1, 1920 (one run begins in 1850), similar initial conditions imply that projected weather patterns near the beginning of 1920 remain similar across runs. However, we do not use the first five months of the 1920 data because we define each year to begin in June in line with the VDSA data.

B Additional temperature specifications

Along with our baseline specification with ten bins based on deciles of the temperature distribution, we estimate two other regressions that include temperature bins. The first uses a set of bins of three degrees Celsius in width, spanning from 0 to 3 degrees to 36 to 39 degrees Celsius, along with two additional bins covering all lower and higher temperatures respectively. This gives us a total of fifteen bins, which we use for crops in all three growing seasons. Importantly, this allows for more variation in yields in response to different temperatures at the higher end of this scale. However, some bins are sparsely populated with nonzero observations in our baseline period, thus limiting our statistical power to estimate these relationships. In some cases, such as for especially hot temperatures in the *rabi* season, we do not have sufficient

data to estimate β_g^k at all, in which case we apply the estimate from the nearest adjacent bin when calculating projected yields.

A second version augments our initial decile bins specification by adding a ‘high-degree days’ (HDD) variable. For each district-day whose temperature falls into the top decile bin, we define a variable that is calculated as the difference between the observed temperature and the lower limit of that bin. When the temperature falls below the top decile, the variable is equal to zero. We then sum this variable across days for each district-year to produce the HDD variable used in our regressions. The addition of this HDD variable thus allows for some convexity in the relationship between temperature and yield at the highest temperatures, while maintaining the other advantages of our specification with decile bins.

C Details of model and parameterization

In this appendix, we discuss the setup and key implications of the Allen and Atkin (2016) model, and then explain how the model is parameterized in our empirical application. Note that this model is sufficiently complex that a full exposition would require us to replicate a substantial part of the paper by Allen and Atkin. Here, we therefore aim to provide enough information to make clear the model’s main components and their intuitive interpretations. A reader interested in understanding the model in depth should refer to the original paper.

Farmers are distributed across N districts, where there are a fixed number of identical farmers in each district i , L_i , and each farmer has a choice of G crops they may grow. The quantity produced Q_{ig} of a given crop g in district i depends on the share θ_{ig} of each local farmer’s land planted with that crop. Quantity produced is also a function of the crop’s local productivity (yield) $A_{ig}(s) > 0$, which in turn depends on the state of the world s . Therefore, quantity produced is proportional to θ_{ig} and the number of farmers in a district, as well as to yield productivity in the realized state, which is drawn from a continuum of states and is not known before planting. However, farmers have full knowledge of the distribution of possible states; i.e. there is no uncertainty in the model.

Farmers have a constant relative risk aversion (CRRA) utility function with constant elasticity of substitution (CES) preferences across crops, which depend on the quantity consumed $C_{ig}(s)$ of each crop in the realized state:

$$U_i(s) = \frac{1}{1-\rho} \left[\left(\sum_{g=1}^G \alpha_g^{1/\sigma} (C_{ig}(s)/L_i)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \right]^{1-\rho}.$$

Here, $\sigma > 0$ is the elasticity of substitution between goods, $\rho > 0$ represents the extent of risk aversion, and $\alpha_g > 0$ is a set of crop-specific preference parameters that sum to one.

Given their crop planting choices, each farmer receives nominal revenue $\sum_{g=1}^G \theta_{ig} A_{ig}(s) p_{ig}(s)$ in state s , where $p_{ig}(s)$ is the price of crop g in district i in that state. Farmers have marginal costs of production c_{ig} , which are specific to each district and crop, and are assumed to be invariant across states.

Trade between districts is governed by price differentials and trade costs across districts, implying that goods will flow from low-cost producers to high-cost ones, with the intensity of the effects mediated by the magnitude of trade costs. Specifically, trade between districts

is governed by the following log-linear arbitrage condition relating consumption, quantity produced and goods prices:²⁴

$$\frac{C_{ig}(s)}{Q_{ig}(s)} \propto \prod_{j=1}^N \left(\frac{p_{ig}(s)}{p_{jg}(s)} \right)^{\varepsilon_{ij}}.$$

The parameters ε_{ij} represent the costs of trade between districts i and j ; the larger the value of ε_{ij} , the less costly is trade between i and j .

To make the model tractable, the authors assume that the yield of each crop in a given district is lognormally distributed across states, which implies that the real returns z_i from a farmer's portfolio of crops are also approximately lognormal: $\ln z_i \sim N(\mu_i^z, \sigma_i^{2,z})$. This assumption, in conjunction with constant relative risk aversion, leads to a simple expression for the log expected utility of the (identical) farmers within a district i :

$$\ln E(U_i) = \left(\mu_i^z + \frac{1}{2} \sigma_i^{2,z} \right) - \frac{1}{2} \rho \sigma_i^{2,z}. \quad (3)$$

The two terms in this expression represent the log of mean real returns (which, due to the lognormality assumption, is a function of both the mean and variance of log real returns) and the variance of log real returns, whose relative importance is governed by the risk-aversion parameter ρ .

Farmers in a given district choose the share θ_{ig} of land to be planted with crop g to maximize their (log) expected utility under the restriction that $\sum_g \theta_{ig} = 1$. After some manipulation, Allen and Atkin (2016) show that the first-order conditions for each district and good can be written in the following way:²⁵

$$\mu_{ig}^z - \rho \sum_h \theta_{ih} \Sigma_{igh}^z = \lambda_i. \quad (4)$$

The first term in the expression above is the contribution of crop g to the log of mean real returns; in other words, the sum of μ_{ig}^z across goods is $\mu_i^z + \frac{1}{2} \sigma_i^{2,z}$, the first term in the previous equation. Similarly, the second term is the contribution of crop g to the variance of log real returns, multiplied by ρ . Here, Σ_{igh}^z is a function of the variance-covariance matrix of real returns across crops within a given district; note that this depends in part on covariances of returns between crops g and h . Finally, λ_i is the district-specific Lagrange multiplier from the farmer's maximization problem. These first-order conditions capture the tradeoff facing farmers in choosing their portfolio of crops: a crop bringing greater risk to the portfolio (the second term) should have a higher mean return (the first term).

For counterfactual analysis using this model, we require sufficient information to calculate μ_{ig}^z and Σ_{igh}^z for each district and crop, given a particular distribution of planting choices, as well as an estimate of ρ . These are then used to identify a set of crop choices that satisfy the first-order conditions above. In our counterfactuals of interest, μ_{ig}^z and Σ_{igh}^z are functions of the distribution of our projected yields (based on the CESM-LENS data and our yield-

²⁴The authors suggest two ways of microfounding this relationship, including a scenario in which transport costs between villages are ad valorem and trade is facilitated by a large number of heterogeneous traders.

²⁵See Proposition 1 in Allen and Atkin (2016).

weather regressions) in each decade. Several parameters are needed for these calculations; we next provide details on how we calibrate or estimate each of these parameters.

C.1 Parameterization

We assign α_g to be equal to the share of the value of each crop in total India-wide agricultural revenue according to the VDSA data for the year 2000. We use an estimate of 2.38 for σ from Allen and Atkin (2016), who recover this using information on household consumption from India's 1987-88 National Sample Survey. We set L_i equal to the rural population of district i as observed in India's 2001 census.

The authors assume a relationship $\varepsilon_{ij} = \beta D_{ij}^{-1.5}$ between trade costs and travel times D_{ij} . The model delivers equations relating prices to yields and travel times, from which the authors estimate $\beta = 6.42$. We use this value, along with their data on travel times between districts, to calculate bilateral trade costs. The authors provide data on the travel times prevailing in several different years and under various possible assumed off-highway speeds; we use the average of 1996 and 2004 travel times, given an off-highway speed of 20 miles per hour.

In order to estimate risk aversion ρ and crop-district-specific production costs c_{ig} , we run a regression derived by the authors from the model's first-order conditions (equation (4)):

$$(\mu_{igd}^z + c_{igd}) = \rho \sum_h \theta_{ihd} \Sigma_{ighd}^z + \delta_{id} + \delta_{gd} + \delta_{ig} + \xi_{igd}$$

This specification depends on the assumption that farmers make utility-maximizing crop choice decisions based on known weather distributions that vary by decade d . The parameters discussed above, along with decade-level averages of yield A_{igd} and crop planting patterns θ_{igd} from the VDSA data, allow us to calculate real returns μ_{igd}^z gross of production costs c_{igd} . We can also calculate Σ_{ighd}^z in a similar way, using the annual VDSA data to determine the required decade-level yield variances and covariances.

The first-order conditions imply the presence of district-decade fixed effects corresponding to the Lagrange multipliers. It thus only remains to add production costs to the right-hand side of the regression (since these are also present on the left-hand side, but are not in equation (4)). As in Allen and Atkin (2016), we model these as the sum of crop-decade and district-crop fixed effects and a residual at the crop-district-decade level. We use the two decades around our usual baseline year of 2000 to estimate this equation: i.e. the 1990s and the 2000s. We only include district-crops in the regression if θ_{igd} exceeds 0.001.

Our estimate of ρ is equal to 0.813, which is close to the OLS estimate of 0.964 from Allen and Atkin (2016), who use a longer panel of the VDSA data. Based on this result, we calibrate costs c_{igd} so that the decade-specific first-order conditions hold exactly, given that δ_{id} represents the decade-specific Lagrange multiplier. More precisely, we use $c_{igd} \approx \delta_{gd} + \delta_{ig} + \xi_{igd}$ as our initial guess and then adjust the estimated costs until the first-order conditions hold with equality in our baseline decade (the 2000s).

We include all possible district-crop combinations in the counterfactuals. For district-crop pairs with no observed planted area, we assume that mean yield μ_{igd}^A is equal to the minimum value observed across districts for that same crop in the 2000s, and set the missing rows and columns of the variance-covariance matrix of yields Σ_{ighd}^A equal to values from $\Sigma_{i'ghd}^A$, where i' is another (randomly chosen) district with a full variance-covariance matrix in the 2000s. Our

initial guess for c_{igd} for these district-crops (and all others omitted from our regression) is the maximum estimated cost across districts for the same crop.

D Additional figures and tables


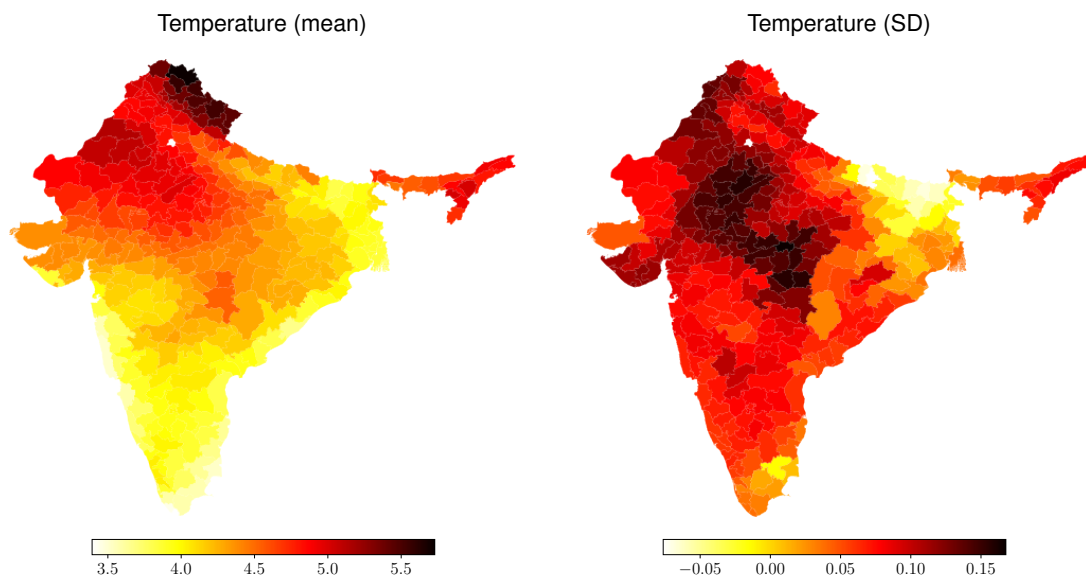
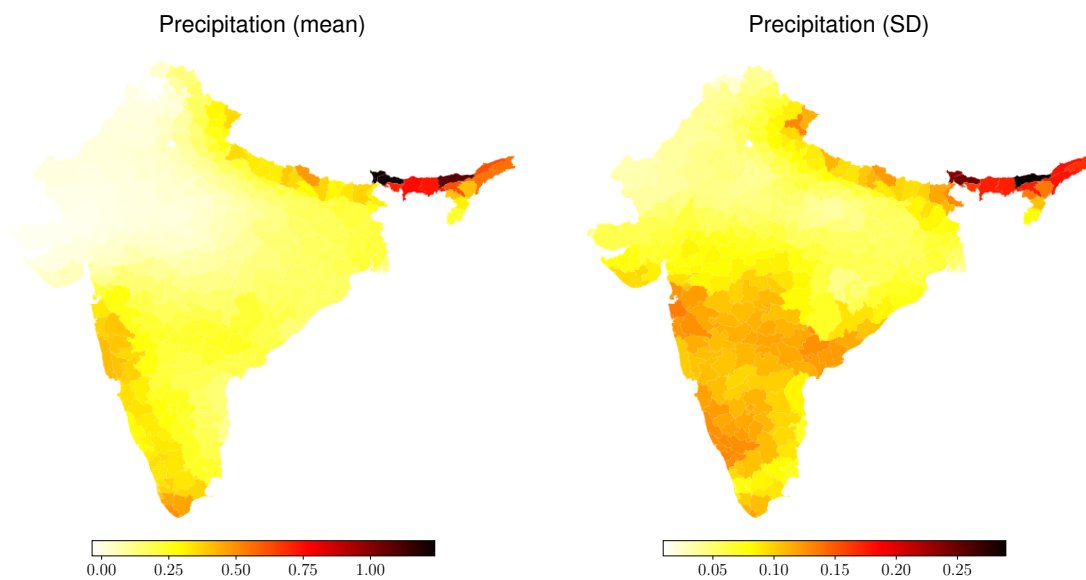


Figure A1: Most-planted crop (by land area) by district in 2000




(a) Temperature




(b) Precipitation

Figure A2: Projected change in mean and standard deviation of average temperature (*a*) and total precipitation (*b*) between the 2000s and 2090s by district according to the CESM-LENS data



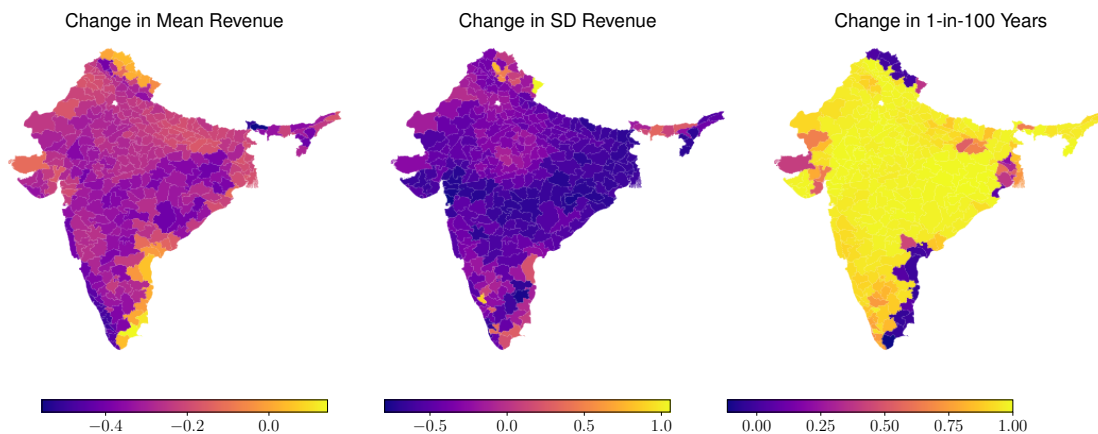
(a) Temperature



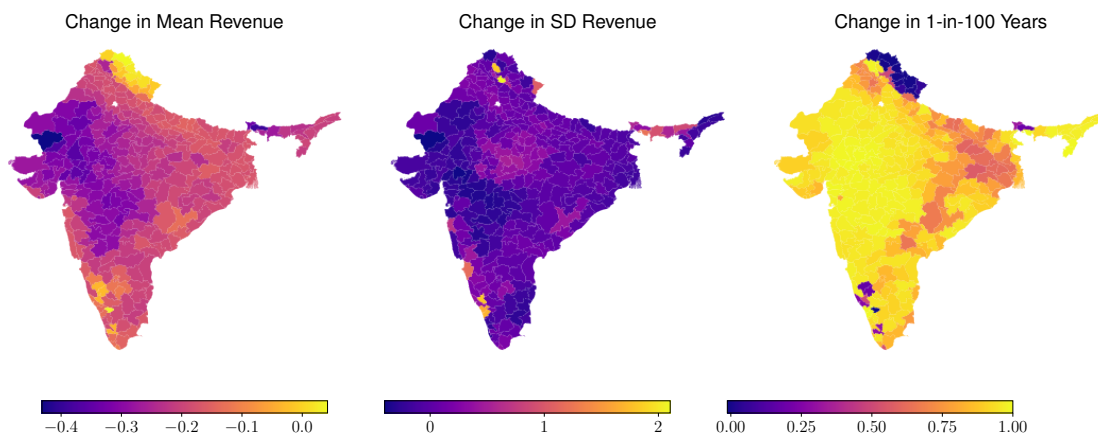
(b) Precipitation

Figure A3: Distributions of average temperature and total precipitation for each decade

This figure displays projected distributions (5^{th} , 25^{th} , 50^{th} , 75^{th} and 95^{th} percentiles) of (a) average temperature across district-days and (b) total precipitation across district-years, for each decade in the CESM-LENS data for India.



(a) Temperature Bins



(b) Degree Days

Figure A4: Projected district-level proportional changes in agricultural revenue between the 2000s and 2090s

This figure shows maps with the projected district-level proportional changes in mean agricultural revenue (left) and weather-induced standard deviation of agricultural revenue (center), and changes in the share of 1-in-100 bad weather years for agricultural revenue (right), between the 2000s and 2090s, for the baseline models: decile temperature bins (top) and degree days (bottom).




Figure A5: Differences across specifications in projected changes in mean agricultural revenue

This figure shows the projected proportional changes in mean agricultural revenue (in percent) in the 2030s, 2060s and 2090s relative to the 2000s for twelve different models. *Decile Bins* refers to the model with decile temperature bins; *Degree Days* refers to the model using degree days, with baseline temperature specified; *Add. Precipitation* refers to the inclusion of additional precipitation variables; *Decile Precipitation Bins* refers to the use of decile bins instead of a quadratic in precipitation; *1997-2015* refers to the use of this shorter sample period; *3°C Bins* refers to the model with three-degree temperature bins; and *Decile Bins + HDD* refers to the model with decile temperature bins and an HDD variable. See Section 3.1 for details of all of these specifications. Results highlighted in red (decile bins model) or blue (degree days model) are also pictured in Figure 4(a). Vertical bars display 95% confidence intervals based on a block bootstrap, as explained in Section 4 and further discussed in Footnote 16.

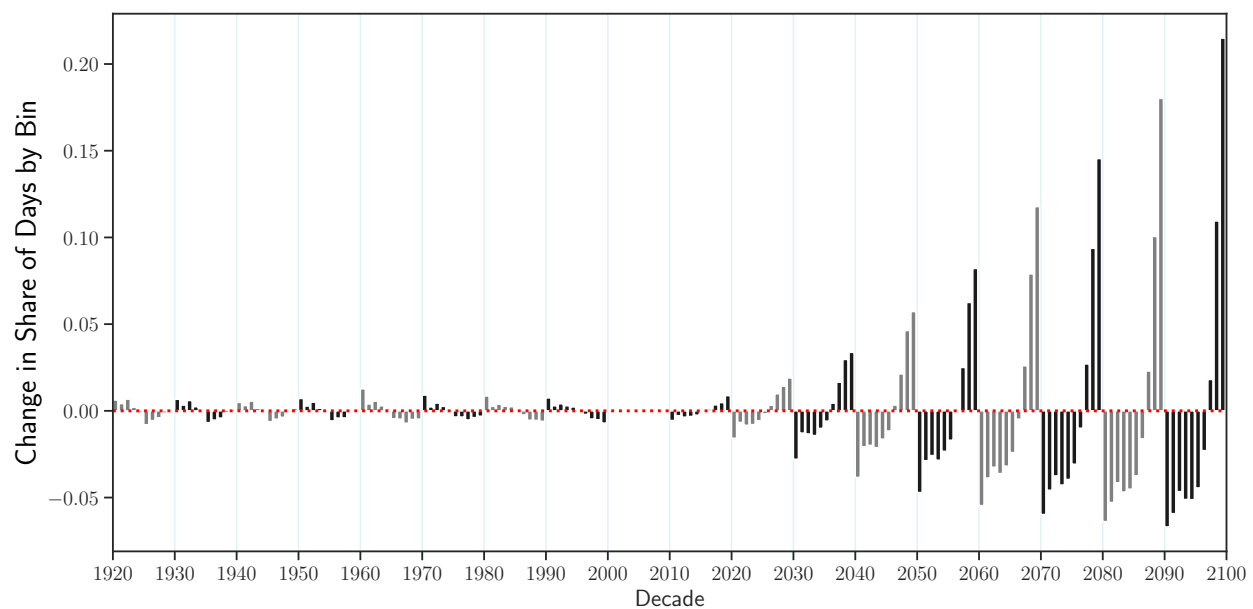


Figure A6: Share of days in each decile temperature bin relative to the 2000s

This figure displays the change (in percentage points) in the share of days in each decile temperature bin (year-round) relative to the 2000s in the CESM-LENS weather projection data. See Table A2 for the limits of each bin in degrees Celsius.




Figure A7: Differences across specifications in projected changes in the weather-induced standard deviation of agricultural revenue

This figure shows the projected proportional changes in the weather-induced standard deviation of agricultural revenue (in percent) in the 2030s, 2060s and 2090s relative to the 2000s for twelve different models. *Decile Bins* refers to the model with decile temperature bins; *Degree Days* refers to the model using degree days, with baseline temperature specified; *Add. Precipitation* refers to the inclusion of additional precipitation variables; *Decile Precipitation Bins* refers to the use of decile bins instead of a quadratic in precipitation; *1997-2015* refers to the use of this shorter sample period; *3°C Bins* refers to the model with three-degree temperature bins; and *Decile Bins + HDD* refers to the model with decile temperature bins and an HDD variable. See Section 3.1 for details of all of these specifications. Results highlighted in red (decile bins model) or blue (degree days model) are also pictured in Figure 5(a). Vertical bars display 95% confidence intervals based on a block bootstrap, as explained in Section 4 and further discussed in Footnote 16.



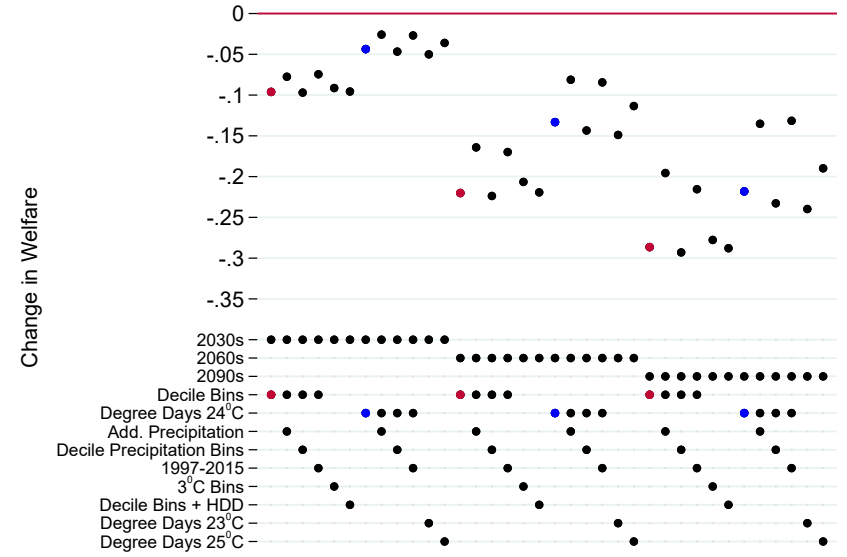


Figure A8: Differences across specifications in projected changes in the share of 1-in-100 bad weather years for agricultural revenue

This figure shows the projected changes in the share of 1-in-100 bad years for agricultural revenue (in percentage points) in the 2030s, 2060s and 2090s relative to the 2000s for twelve different models. *Decile Bins* refers to the model with decile temperature bins; *Degree Days* refers to the model using degree days, with baseline temperature specified; *Add. Precipitation* refers to the inclusion of additional precipitation variables; *Decile Precipitation Bins* refers to the use of decile bins instead of a quadratic in precipitation; *1997-2015* refers to the use of this shorter sample period; *3°C Bins* refers to the model with three-degree temperature bins; and *Decile Bins + HDD* refers to the model with decile temperature bins and an HDD variable. See Section 3.1 for details of all of these specifications. Results highlighted in red (decile bins model) or blue (degree days model) are also pictured in Figure 6(a). Vertical bars display 95% confidence intervals based on a block bootstrap, as explained in Section 4 and further discussed in Footnote 16.



(a) Naive Welfare



(b) Model-Based Welfare

Figure A9: Differences across specifications in naive and model-based welfare measures

This figure shows the projected proportional changes (in percent) in the 2030s, 2060s and 2090s relative to the 2000s in a naive welfare measure based on nominal returns (left column) and welfare from real returns (right column) as specified in Section 5, for twelve different models. *Decile Bins* refers to the model with decile temperature bins; *Degree Days* refers to the model using degree days, with baseline temperature specified; *Add. Precipitation* refers to the inclusion of additional precipitation variables; *Decile Precipitation Bins* refers to the use of decile bins instead of a quadratic in precipitation; *1997-2015* refers to the use of this shorter sample period; *3°C Bins* refers to the model with three-degree temperature bins; and *Decile Bins + HDD* refers to the model with decile temperature bins and an HDD variable. See Section 3.1 for details of all of these specifications. Results highlighted in red (decile bins model) or blue (degree days model) are also pictured in Figures 7(a) and 7(b) respectively.

Table A1: VDSA summary statistics in 2000

	(1)	(2)	(3)	(4)
	Share of	Share of land	(among producers)	
Panel A. Crop-level summary statistics	districts	25th pctl	Median	75th pctl
Rice	94.2	7.7	31.7	65.6
Maize	88.7	0.3	2.0	9.3
Sesame	86.8	0.2	0.7	1.6
Chickpea	86.1	0.1	1.3	4.7
Wheat	85.5	2.9	18.4	40.9
Pigeon pea	83.9	0.3	1.4	3.4
Sugarcane	82.6	0.2	0.9	3.6
Groundnut	76.1	0.1	1.1	4.2
Rapeseed/mustard	75.8	0.4	1.5	3.9
Sorghum	69.4	0.3	2.0	9.2
Pearl millet	60.6	0.1	1.7	9.2
Cotton	47.4	0.3	2.5	13.2
Linseed	47.4	0.1	0.2	0.9
Barley	46.4	0.2	0.8	2.2
Castor	38.4	0.0	0.2	0.9
Finger millet	31.0	0.1	0.8	7.2
Panel B. District-level summary statistics		25th pctl	Median	75th pctl
Number of crops grown		10	12	13
Number of crops (at least 5% of land)		2	3	4
Share of land of largest crop		40.1	49.7	70.7

This table displays summary statistics of the VDSA agricultural data for the 310 districts in the sample, for the year 2000. For each crop, we show the share of districts for which area planted and quantity produced are both recorded in the VDSA data as nonzero and nonmissing. We then take the total land area dedicated to these crops in each district, and calculate the share of this land that is planted with each crop. We display the value of this variable at the 25th, 50th and 75th percentile of the districts that are recorded as producing that crop. We also show information on the distribution by district of the number of crops grown (with and without the restriction that the crop should be grown on at least 5% of land area). Finally, we show information on the distribution across districts of the share of land allocated to the most important crop (by land area).

Table A2: Upper limits of decile temperature bins

	(1)	(2)	(3)
	<i>Kharif</i>	<i>Rabi</i>	Year-round
Decile 1	22.26	14.25	15.98
Decile 2	23.86	17.09	19.62
Decile 3	24.88	19.25	22.00
Decile 4	25.70	21.00	23.67
Decile 5	26.42	22.43	25.00
Decile 6	27.11	23.72	26.18
Decile 7	27.86	25.00	27.34
Decile 8	28.87	26.48	28.74
Decile 9	30.89	28.65	31.09

This table displays the limits of the ten decile bins used for temperature in our baseline yield-weather specification. Each of the numbers in the table represents the upper limit of one bin; for example, the second bin for the *kharif* (monsoon) season covers 22.26 to 23.86 degrees Celsius.

Table A3: Yield-weather estimation – decile temperature bins

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	Castor	Cotton	Finger Millet	Groundnut	Maize	Pearl Millet	Pigeon Pea	Rice	Sesame	Sorghum	Barley	Chickpea	Linsseed	Rapeseed/Mustard	Wheat	Sugarcane
Bin [min,p10]	0.435 (0.273)	-0.470* (0.253)	0.453*** (0.147)	0.237 (0.168)	0.630*** (0.155)	0.879*** (0.186)	0.871*** (0.147)	0.682*** (0.129)	0.770*** (0.154)	0.861*** (0.181)	0.387*** (0.144)	-0.0672 (0.182)	0.456** (0.184)	0.244* (0.141)	0.525*** (0.105)	1.196*** (0.209)
Bin [p10,p20]	0.0597 (0.227)	-0.0953 (0.207)	0.250* (0.127)	-0.0363 (0.116)	0.249** (0.118)	0.453*** (0.142)	0.500*** (0.120)	0.336*** (0.104)	0.200 (0.132)	0.233* (0.123)	0.203 (0.132)	-0.0490 (0.117)	0.641*** (0.159)	0.454*** (0.111)	0.457*** (0.0805)	0.903*** (0.154)
Bin [p20,p30]	0.177 (0.296)	-0.272 (0.234)	0.0282 (0.152)	-0.205 (0.137)	-0.0493 (0.148)	0.319* (0.189)	0.386** (0.155)	0.0973 (0.128)	0.0674 (0.147)	-0.210 (0.160)	0.0238 (0.149)	-0.00236 (0.116)	0.177 (0.179)	0.206 (0.130)	0.143* (0.0865)	0.266 (0.162)
Bin [p40,p50]	0.267 (0.314)	-0.453* (0.243)	0.174 (0.193)	-0.165 (0.147)	-0.103 (0.187)	0.175 (0.204)	-0.0464 (0.159)	-0.218* (0.122)	0.232 (0.188)	-0.185 (0.189)	-0.217 (0.159)	0.0600 (0.145)	0.114 (0.186)	0.123 (0.160)	0.0165 (0.114)	-0.0125 (0.205)
Bin [p50,p60]	-0.139 (0.293)	-0.294 (0.236)	-0.293 (0.189)	-0.0439 (0.166)	-0.191 (0.129)	0.0477 (0.156)	0.0596 (0.166)	-0.211* (0.110)	-0.0792 (0.141)	0.0861 (0.138)	-0.123 (0.166)	-0.308*** (0.117)	-0.226 (0.180)	-0.207 (0.143)	-0.160 (0.101)	-0.0373 (0.166)
Bin [p60,p70]	-0.152 (0.305)	-0.610*** (0.214)	-0.0682 (0.210)	-0.132 (0.132)	-0.333*** (0.143)	-0.0912 (0.177)	-0.459*** (0.130)	-0.509*** (0.126)	0.162 (0.175)	-0.472** (0.189)	-0.438* (0.186)	-0.638*** (0.146)	-0.729*** (0.228)	-0.404** (0.170)	-0.386*** (0.112)	-0.343** (0.165)
Bin [p70,p80]	-0.274 (0.263)	-0.672*** (0.203)	-0.188 (0.196)	-0.661*** (0.158)	-0.392*** (0.140)	-0.447* (0.182)	-0.382** (0.148)	-0.626*** (0.118)	-0.554*** (0.166)	-0.803*** (0.180)	-0.358* (0.159)	-0.570*** (0.123)	-0.498*** (0.198)	-0.471*** (0.155)	-0.393*** (0.122)	-0.486*** (0.162)
Bin [p80,p90]	-0.265 (0.237)	-0.927*** (0.182)	-0.588*** (0.181)	-0.786*** (0.119)	-0.870*** (0.138)	-0.663*** (0.162)	-0.432*** (0.127)	-0.779*** (0.127)	-0.756*** (0.148)	-1.031*** (0.176)	-0.143 (0.154)	-0.743*** (0.124)	-0.590*** (0.193)	-0.428*** (0.149)	-0.434*** (0.107)	-0.371** (0.182)
Bin [p90,max]	-0.463 (0.284)	-0.772*** (0.178)	-0.422** (0.202)	-0.542*** (0.140)	-1.213*** (0.138)	-1.220*** (0.203)	-0.617*** (0.151)	-0.399*** (0.113)	-1.358*** (0.202)	-1.307*** (0.200)	-0.366** (0.159)	-0.853*** (0.131)	-0.754*** (0.165)	-0.627*** (0.145)	-0.567*** (0.109)	-0.511*** (0.147)
Monsoon prec	0.446*** (0.116)	0.451*** (0.110)	0.186*** (0.0667)	0.389*** (0.0873)	-0.206*** (0.0621)	0.853*** (0.133)	0.252*** (0.0596)	0.366*** (0.0404)	-0.0208 (0.0636)	0.438*** (0.107)	0.207*** (0.0353)	0.326*** (0.0363)	0.199*** (0.0445)	0.175*** (0.0314)	0.218*** (0.0239)	0.0612 (0.0380)
Monsoon prec ²	-0.136*** (0.0458)	-0.160*** (0.0467)	-0.0401** (0.0190)	-0.121*** (0.0319)	0.0305 (0.0220)	-0.427*** (0.0633)	-0.0798*** (0.0221)	-0.0805*** (0.0116)	-0.0342* (0.0200)	-0.214*** (0.0485)	-0.0151 (0.0113)	-0.0549*** (0.0113)	-0.0359*** (0.0118)	-0.0168* (0.00965)	-0.0323*** (0.00770)	-0.0268** (0.0118)
Rabi prec																
Rabi prec ²																
Observations	3683	5537	3905	8690	10186	6985	9485	10685	9633	8044	5589	9857	5010	8655	9806	9521
R ²	0.698	0.613	0.783	0.531	0.676	0.706	0.600	0.811	0.629	0.606	0.753	0.562	0.685	0.727	0.835	0.821

This table shows the coefficients estimated from equation (1) using *decile temperature bins*. Definitions of the variables may be found in Section 3. See also Table A2 for the limits of each temperature bin in degrees Celsius. Standard errors clustered by district in parentheses: * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Table A4: Yield-weather estimation – degree days

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Degree days	Castor -0.000484*** (0.000174)	Cotton -0.000477*** (0.0000901)	Finger Millet -0.000495*** (0.000121)	Groundnut -0.000392*** (0.0000772)	Maize -0.000915*** (0.0000756)	Pearl Millet -0.00106*** (0.000130)	Pigeon Pea -0.000613*** (0.0000773)	Rice -0.000386*** (0.0000567)	Sesame -0.00108*** (0.000133)	Sorghum -0.000923*** (0.000121)	Barley -0.000295*** (0.0000716)	Chickpea -0.000510*** (0.0000648)	Linsed -0.000637*** (0.000033)	Rapeseed/Mustard -0.000477*** (0.0000674)	Wheat -0.000456*** (0.0000517)	Sugarcane -0.000219*** (0.0000291)
Monsoon prec	0.469*** (0.113)	0.510*** (0.112)	0.244*** (0.0650)	0.506*** (0.0947)	-0.125*** (0.0619)	1.00*** (0.130)	0.326*** (0.0596)	0.470*** (0.0439)	0.0729 (0.0664)	0.629*** (0.112)	0.222*** (0.0351)	0.340*** (0.0369)	0.221*** (0.0453)	0.189*** (0.0311)	0.242*** (0.0244)	0.0950** (0.0374)
Monsoon prec ²	-0.140*** (0.0443)	-0.172*** (0.0478)	-0.0525*** (0.0190)	-0.151*** (0.0351)	0.00940 (0.0233)	-0.481*** (0.0642)	-0.0925*** (0.0224)	-0.0981*** (0.0128)	-0.0606*** (0.0223)	-0.273*** (0.0530)	-0.0195* (0.0112)	-0.0584*** (0.0117)	-0.0451*** (0.0120)	-0.0213** (0.00955)	-0.0389*** (0.00795)	-0.0328*** (0.0117)
Rabi prec																
Rabi prec ²																
Observations	3083	5537	3905	8690	10186	6985	9485	10085	9633	8044	5589	9857	5010	8655	9806	9521
R ²	0.697	0.611	0.782	0.525	0.674	0.705	0.596	0.805	0.625	0.600	0.752	0.559	0.683	0.726	0.834	0.819

This table shows the coefficients estimated from equation (1) using *degree days*. Definitions of the variables may be found in Section 3. Standard errors clustered by district in parentheses: * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Table A5: Yield-weather estimation – decile temperature bins with additional precipitation variables

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Castor		Cotton	Finger Millet	Groundnut	Maize	Pearl Millet	Pigeon Pea	Rice	Sesame	Sorghum	Barley	Chickpea	Linseed	Rapeseed/Mustard	Wheat	Sugarcane
Bin [min,p10]	0.348 (0.271)	-0.533** (0.256)	0.363** (0.148)	0.157 (0.164)	0.465** (0.152)	0.682** (0.186)	0.782** (0.186)	0.518** (0.126)	0.056** (0.152)	0.763** (0.182)	0.386** (0.146)	-0.0694 (0.185)	0.434** (0.184)	0.237* (0.142)	0.533** (0.106)	1.071** (0.211)
Bin [p10,p20]	0.0225 (0.226)	-0.131 (0.210)	0.205 (0.129)	-0.0575 (0.115)	0.191 (0.117)	0.394** (0.139)	0.466** (0.119)	0.271** (0.105)	0.159 (0.133)	0.195 (0.124)	0.209 (0.134)	-0.0626 (0.118)	0.378** (0.159)	0.434** (0.111)	0.454** (0.0803)	0.830** (0.155)
Bin [p20,p30]	0.155 (0.297)	-0.291 (0.234)	0.00606 (0.152)	-0.210 (0.137)	-0.0679 (0.147)	0.282 (0.185)	0.371** (0.155)	0.0761 (0.127)	0.0614 (0.146)	-0.229 (0.161)	0.0296 (0.150)	-0.0178 (0.116)	0.151 (0.178)	0.180 (0.131)	0.136 (0.0859)	0.234 (0.162)
Bin [p40,p50]	0.292 (0.313)	-0.444* (0.242)	0.189 (0.195)	-0.170 (0.147)	-0.112 (0.186)	0.170 (0.170)	-0.0460 (0.160)	-0.219* (0.123)	0.232 (0.188)	-0.183 (0.189)	-0.203 (0.160)	0.0623 (0.145)	0.129 (0.186)	0.119 (0.160)	0.00784 (0.114)	-0.0177 (0.205)
Bin [p50,p60]	-0.140 (0.289)	-0.269 (0.235)	-0.288 (0.189)	-0.0190 (0.166)	-0.154 (0.128)	0.0759 (0.157)	0.0716 (0.165)	-0.174 (0.108)	-0.0297 (0.140)	0.101 (0.138)	-0.109 (0.167)	-0.298** (0.118)	-0.210 (0.143)	-0.201 (0.143)	-0.167 (0.102)	-0.0353 (0.163)
Bin [p60,p70]	-0.0975 (0.304)	-0.573** (0.213)	-0.0398 (0.210)	-0.101 (0.133)	-0.265* (0.143)	-0.0396 (0.177)	-0.419** (0.130)	-0.426** (0.126)	0.224 (0.176)	-0.428** (0.190)	-0.421** (0.191)	-0.634** (0.147)	-0.748** (0.226)	-0.431** (0.169)	-0.399** (0.112)	-0.325** (0.163)
Bin [p70,p80]	-0.276 (0.264)	-0.623** (0.209)	-0.137 (0.195)	-0.639** (0.164)	-0.328** (0.139)	-0.415** (0.183)	-0.333** (0.149)	-0.529** (0.118)	-0.492** (0.169)	-0.746** (0.181)	-0.345** (0.161)	-0.573** (0.125)	-0.516** (0.199)	-0.491** (0.155)	-0.406** (0.123)	-0.475** (0.162)
Bin [p80,p90]	-0.203 (0.237)	-0.855** (0.189)	-0.475** (0.188)	-0.727** (0.120)	-0.728** (0.136)	-0.573** (0.165)	-0.349** (0.126)	-0.612** (0.146)	-0.632** (0.172)	-0.943** (0.172)	-0.128 (0.159)	-0.746** (0.125)	-0.596** (0.192)	-0.455** (0.150)	-0.448** (0.107)	-0.355** (0.179)
Bin [p90,max]	-0.293 (0.288)	-0.591** (0.205)	-0.201 (0.208)	-0.347** (0.150)	-0.779** (0.144)	-0.871** (0.214)	-0.395** (0.158)	0.0623 (0.116)	-0.969** (0.206)	-1.048** (0.197)	-0.337** (0.167)	-0.899** (0.135)	-0.825** (0.167)	-0.696** (0.149)	-0.613** (0.110)	-0.327** (0.148)
Monsoon prec	0.376** (0.122)	0.427** (0.115)	0.135** (0.0658)	0.320** (0.0882)	-0.298** (0.0583)	0.744** (0.134)	0.233** (0.0636)	0.313** (0.0393)	-0.114* (0.0655)	0.395** (0.105)	0.208** (0.0352)	0.326** (0.0361)	0.199** (0.0441)	0.172** (0.0316)	0.218** (0.0239)	0.0221 (0.0390)
Monsoon prec ²	-0.118** (0.0438)	-0.151** (0.0471)	-0.0241 (0.0190)	-0.112** (0.0314)	0.0422** (0.0211)	-0.405** (0.0627)	-0.0713** (0.0215)	-0.0641** (0.0111)	-0.0228 (0.0195)	-0.294** (0.0479)	-0.0156 (0.0113)	-0.0536** (0.0112)	-0.0341** (0.0117)	-0.0146 (0.00974)	-0.0314** (0.00771)	-0.0252** (0.0114)
Rainy days	0.00209 (0.00129)	0.00204** (0.000981)	0.00340** (0.000939)	0.00260** (0.000671)	0.00528** (0.000564)	0.00451** (0.000789)	0.00235** (0.000687)	0.00509** (0.000470)	0.00480** (0.000754)	0.00295** (0.000739)	0.000383 (0.000459)	-0.00156** (0.000412)	-0.00175** (0.000625)	-0.00165** (0.000479)	-0.00104** (0.000302)	0.00128** (0.000266)
Heavy rain	-0.00243 (0.0120)	-0.00395 (0.0107)	-0.0121 (0.00773)	0.0134** (0.00622)	0.0133* (0.00675)	0.00746 (0.00900)	-0.00896 (0.00629)	-0.0120** (0.00450)	0.0185* (0.00835)	-0.00423 (0.00917)						
Dry spell	-0.00467** (0.00161)	0.000436 (0.00123)	-0.000837 (0.00191)	-0.000597 (0.000608)	0.000163 (0.000709)	-0.00122 (0.000784)	0.000538 (0.000658)	0.00109** (0.000538)	0.00129 (0.00104)	0.00128 (0.000816)	0.000147 (0.000201)	0.000315* (0.000173)	0.000490* (0.000260)	0.000335** (0.000138)	-0.0000914 (0.000111)	0.000174 (0.000126)
Rabi prec																
Rabi prec ²																
Observations	3683	5537	3905	8690	10186	6985	9485	10685	9633	8044	5589	9857	5010	8655	9806	9521
R ²	0.700	0.613	0.785	0.533	0.680	0.709	0.601	0.814	0.631	0.607	0.753	0.563	0.687	0.728	0.836	0.821

This table shows the coefficients estimated from equation (1) using *decile temperature bins with additional precipitation variables*. Definitions of the variables may be found in Section 3. See also Table A2 for the limits of each temperature bin in degrees Celsius. Standard errors clustered by district in parentheses: * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Table A6: Yield-weather estimation – decile temperature and precipitation bins

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	Castor	Cotton	Finger Millet	Groundnut	Maize	Pearl Millet	Pigeon Pea	Rice	Sesame	Sorghum	Barley	Chickpea	Linsed	Rapeseed/Mustard	Wheat	Sugarcane
Temp bin [min,p10]	0.467 (0.284)	-0.376 (0.259)	0.477*** (0.156)	0.190 (0.170)	0.496*** (0.152)	1.038*** (0.194)	0.814*** (0.146)	0.672*** (0.131)	0.743*** (0.161)	0.813*** (0.178)	0.416*** (0.138)	-0.110 (0.182)	0.452*** (0.184)	0.233* (0.140)	0.519*** (0.104)	1.202*** (0.210)
Temp bin [p10,p20]	0.0544 (0.234)	-0.0673 (0.214)	0.263** (0.131)	-0.0583 (0.118)	0.211* (0.119)	0.491*** (0.145)	0.486*** (0.120)	0.343*** (0.105)	0.203 (0.131)	0.215* (0.123)	0.227* (0.132)	-0.0645 (0.118)	0.604*** (0.157)	0.473*** (0.113)	0.466*** (0.0793)	0.897*** (0.156)
Temp bin [p20,p30]	0.177 (0.297)	-0.248 (0.236)	0.0356 (0.155)	-0.193 (0.139)	-0.0336 (0.149)	0.331* (0.192)	0.401*** (0.144)	0.0942 (0.129)	0.0975 (0.144)	0.123 (0.157)	0.0516 (0.149)	-0.00494 (0.116)	0.179 (0.175)	0.204 (0.131)	0.158* (0.0858)	0.254 (0.164)
Temp bin [p40,p50]	0.218 (0.316)	-0.466* (0.247)	0.150 (0.192)	-0.170 (0.149)	-0.0815 (0.186)	0.143 (0.208)	-0.0322 (0.160)	-0.221* (0.123)	0.242 (0.187)	-0.198 (0.188)	-0.208 (0.160)	0.0637 (0.146)	0.102 (0.189)	0.106 (0.160)	0.069 (0.113)	-0.00491 (0.207)
Temp bin [p50,p60]	-0.159 (0.294)	-0.292 (0.236)	-0.290 (0.185)	-0.0372 (0.129)	-0.180 (0.140)	0.0348 (0.188)	0.0618 (0.164)	-0.218*** (0.110)	-0.0821 (0.138)	0.0865 (0.159)	-0.119 (0.138)	-0.244** (0.119)	-0.186 (0.183)	-0.251* (0.145)	-0.134 (0.102)	-0.0180 (0.167)
Temp bin [p60,p70]	-0.212 (0.307)	-0.645*** (0.217)	-0.0966 (0.213)	-0.161 (0.131)	-0.322** (0.140)	-0.324*** (0.179)	-0.460*** (0.188)	-0.524*** (0.127)	0.140 (0.171)	-0.501*** (0.185)	-0.431** (0.185)	-0.570*** (0.148)	-0.670*** (0.221)	-0.427*** (0.170)	-0.347*** (0.112)	-0.315* (0.167)
Temp bin [p70,p80]	-0.332 (0.264)	-0.709*** (0.203)	-0.199 (0.193)	-0.675*** (0.160)	-0.358*** (0.139)	-0.572*** (0.178)	-0.376*** (0.148)	-0.635*** (0.118)	-0.594*** (0.166)	-0.832*** (0.178)	-0.405*** (0.159)	-0.534*** (0.122)	-0.424*** (0.199)	-0.551*** (0.149)	-0.397*** (0.164)	-0.457*** (0.122)
Temp bin [p80,p90]	-0.337 (0.251)	-1.062*** (0.183)	-0.611*** (0.181)	-0.769*** (0.119)	-0.731*** (0.138)	-0.877*** (0.171)	-0.388*** (0.127)	-0.783*** (0.125)	-0.741*** (0.148)	-1.054*** (0.178)	-0.154 (0.149)	-0.686*** (0.120)	-0.520*** (0.192)	-0.467*** (0.145)	-0.409*** (0.104)	-0.326* (0.186)
Temp bin [p90,max]	-0.548* (0.289)	-0.957*** (0.179)	-0.452*** (0.200)	-0.561*** (0.144)	-1.054*** (0.133)	-1.320*** (0.234)	-0.595*** (0.188)	-0.444*** (0.112)	-1.322*** (0.211)	-1.348*** (0.202)	-0.446*** (0.157)	-0.780*** (0.128)	-0.704*** (0.160)	-0.663*** (0.138)	-0.569*** (0.106)	-0.478*** (0.150)
Monsoon prec bin [min,p10]	-0.147*** (0.0447)	-0.0611 (0.0419)	-0.00595 (0.0409)	-0.195*** (0.0437)	-0.0883*** (0.0274)	-0.111*** (0.0411)	-0.119*** (0.0325)	-0.0991*** (0.0282)	-0.0914** (0.0404)	-0.173*** (0.0386)	-0.0411** (0.0185)	-0.193*** (0.0222)	-0.0346 (0.0438)	-0.0762*** (0.0207)	-0.0916*** (0.0148)	-0.0218 (0.0179)
Monsoon prec bin [p10,p20]	-0.0110 (0.0287)	-0.0239 (0.0257)	-0.00857 (0.0305)	-0.0879*** (0.0226)	-0.0201 (0.0197)	-0.0340 (0.0173)	-0.0709*** (0.0243)	-0.0721*** (0.0208)	0.0114 (0.0248)	-0.0529* (0.0284)	-0.0587*** (0.0174)	-0.115*** (0.0155)	-0.102*** (0.0314)	-0.0573*** (0.0192)	-0.0790*** (0.0125)	-0.0194 (0.0143)
Monsoon prec bin [p20,p30]	-0.0182 (0.0294)	-0.0425* (0.0238)	0.0297 (0.0286)	-0.0405** (0.0167)	-0.000932 (0.0152)	-0.00332 (0.0167)	-0.0340* (0.0173)	-0.0387*** (0.0143)	0.0307 (0.0216)	-0.0166 (0.0194)	-0.0341** (0.0170)	-0.0499*** (0.0122)	-0.0199 (0.0200)	-0.00711 (0.0164)	-0.0236* (0.0105)	-0.000955 (0.0115)
Monsoon prec bin [p40,p50]	0.0728** (0.0307)	0.0178 (0.0250)	0.0451** (0.0199)	0.0122 (0.0159)	-0.00416 (0.0158)	-0.00429 (0.0173)	0.0285** (0.0141)	0.0283** (0.0125)	-0.0415* (0.0241)	-0.00345 (0.0175)	0.0448** (0.0143)	0.00416 (0.0125)	0.0351* (0.0186)	0.0209 (0.0145)	0.00878 (0.0102)	-0.00334 (0.0129)
Monsoon prec bin [p50,p60]	0.0624** (0.0264)	0.0231 (0.0259)	0.0629*** (0.0236)	0.0223 (0.0164)	-0.0433*** (0.0158)	-0.0106 (0.0185)	0.0220 (0.0168)	0.0581*** (0.0133)	-0.0626*** (0.0227)	-0.0181 (0.0168)	0.0532*** (0.0159)	0.0433*** (0.0114)	0.0422*** (0.0185)	0.0555*** (0.0163)	0.0442*** (0.0102)	0.0123 (0.0118)
Monsoon prec bin [p60,p70]	0.0360 (0.0363)	0.0673** (0.0310)	0.0827*** (0.0268)	0.0321*** (0.0163)	-0.0510*** (0.0157)	-0.0495*** (0.0199)	0.0422** (0.0174)	0.110*** (0.0135)	-0.0523*** (0.0241)	-0.0363*** (0.0269)	0.0914*** (0.0155)	0.0441*** (0.0188)	0.0641*** (0.0202)	0.0459*** (0.0161)	0.0751*** (0.0106)	0.00619 (0.0154)
Monsoon prec bin [p70,p80]	0.0428 (0.0370)	0.0129 (0.0331)	0.0764*** (0.0271)	0.0144 (0.0185)	-0.115*** (0.0201)	-0.0898*** (0.0256)	0.0398** (0.0170)	0.120*** (0.0143)	-0.125*** (0.0260)	-0.0636*** (0.0217)	0.112*** (0.0158)	0.0596*** (0.0156)	0.0659*** (0.0222)	0.0485*** (0.0171)	0.0661*** (0.0104)	-0.00390 (0.0154)
Monsoon prec bin [p80,p90]	0.0365 (0.0390)	0.0484 (0.0329)	0.108*** (0.0311)	0.0116 (0.0194)	-0.170*** (0.0232)	-0.129*** (0.0279)	0.0124 (0.0202)	0.123*** (0.0157)	-0.150*** (0.0262)	-0.106*** (0.0269)	0.135*** (0.0207)	0.0670*** (0.0188)	0.0513*** (0.0247)	0.0952*** (0.0180)	0.0839*** (0.0119)	-0.00276 (0.0211)
Monsoon prec bin [p90,max]	0.0637 (0.0482)	0.0651 (0.0451)	0.0899*** (0.0333)	0.00967 (0.0235)	-0.199*** (0.0293)	-0.120*** (0.0349)	-0.106 (0.0285)	0.139*** (0.0196)	-0.189*** (0.0329)	-0.150*** (0.0369)	0.253*** (0.0359)	0.146*** (0.0258)	0.117*** (0.0486)	0.183*** (0.0339)	0.144*** (0.0220)	-0.0670 (0.0473)
Rabi prec bin [min,p10]																
Rabi prec bin [p10,p20]																
Rabi prec bin [p20,p30]																
Rabi prec bin [p40,p50]																
Rabi prec bin [p50,p60]																
Rabi prec bin [p60,p70]																
Rabi prec bin [p70,p80]																
Rabi prec bin [p80,p90]																
Rabi prec bin [p90,max]																
Observations	3683	5537	3905	8690	10186	6985	9485	10685	9633	8044	5589	9857	5010	8655	9806	9521
R ²	0.698	0.612	0.784	0.533	0.679	0.704	0.601	0.811	0.631	0.607	0.755	0.563	0.688	0.728	0.837	0.821

This table shows the coefficients estimated from equation (1) using *decile temperature and precipitation bins*. Definitions of the variables may be found in Section 3. See also Table A2 for the limits of each temperature bin in degrees Celsius. Standard errors clustered by district in parentheses: * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Table A7: Yield-weather estimation – decile temperature bins (starting 1997)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	Castor	Cotton	Finger Millet	Groundnut	Maize	Pearl Millet	Pigeon Pea	Rice	Sesame	Sorghum	Barley	Chickpea	Linsed	Rapeseed/Mustard	Wheat	Sugarcane
Bin [min,p10]	0.942*** (0.359)	0.0542 (0.268)	0.415** (0.209)	0.450** (0.198)	0.654*** (0.182)	1.000*** (0.221)	1.008*** (0.214)	0.357*** (0.153)	0.685*** (0.232)	0.718*** (0.225)	0.504*** (0.179)	0.332* (0.188)	0.643*** (0.292)	0.320** (0.158)	0.602*** (0.133)	1.442*** (0.307)
Bin [p10,p20]	0.211 (0.315)	0.339 (0.227)	0.126 (0.174)	0.129 (0.154)	0.389*** (0.149)	0.571*** (0.182)	0.643*** (0.197)	0.343*** (0.126)	0.391** (0.207)	0.208 (0.186)	0.352*** (0.193)	-0.171 (0.144)	0.415* (0.225)	0.582*** (0.130)	0.571*** (0.111)	1.219*** (0.289)
Bin [p20,p30]	0.278 (0.423)	-0.352 (0.339)	-0.223 (0.220)	-0.121 (0.203)	0.0334 (0.196)	0.172 (0.263)	0.482* (0.248)	0.130 (0.181)	0.0910 (0.247)	-0.189 (0.240)	0.347 (0.214)	-0.0660 (0.157)	0.132 (0.239)	0.453*** (0.158)	0.270** (0.122)	0.374 (0.311)
Bin [p40,p50]	0.517 (0.380)	-0.531* (0.306)	-0.352 (0.267)	-0.299 (0.184)	-0.199 (0.221)	-0.0304 (0.261)	0.0149 (0.231)	0.0135 (0.180)	0.240 (0.244)	-0.318 (0.256)	-0.119 (0.240)	-0.0346 (0.180)	0.256 (0.299)	0.328* (0.177)	-0.0712 (0.157)	0.282 (0.376)
Bin [p50,p60]	-0.199 (0.410)	-0.249 (0.313)	-0.268 (0.232)	-0.0631 (0.186)	-0.375** (0.168)	-0.179 (0.217)	0.294 (0.205)	-0.0519 (0.150)	-0.0783 (0.222)	-0.164 (0.197)	-0.0506 (0.278)	-0.383*** (0.146)	-0.166 (0.293)	0.0195 (0.186)	-0.0364 (0.129)	-0.0373 (0.274)
Bin [p60,p70]	-0.0479 (0.343)	-0.867*** (0.268)	-0.386 (0.250)	-0.222 (0.169)	-0.288* (0.159)	-0.0159 (0.217)	-0.270 (0.197)	-0.187 (0.164)	0.154 (0.223)	-0.210 (0.226)	-0.240 (0.256)	-0.810*** (0.183)	-0.598* (0.339)	-0.0463 (0.209)	-0.338** (0.156)	-0.294 (0.258)
Bin [p70,p80]	-0.162 (0.328)	-0.491* (0.275)	-0.499* (0.258)	-0.772*** (0.200)	-0.346** (0.174)	-0.511* (0.243)	-0.230 (0.189)	-0.330** (0.142)	0.0184 (0.226)	-0.880*** (0.244)	-0.0330 (0.257)	-0.545*** (0.163)	-0.369 (0.324)	0.227 (0.189)	-0.312** (0.158)	-0.452* (0.274)
Bin [p80,p90]	-0.229 (0.313)	-0.984*** (0.221)	-0.641*** (0.203)	-0.904*** (0.173)	-0.915*** (0.164)	-0.640*** (0.241)	-0.301* (0.175)	-0.258* (0.138)	-0.603*** (0.224)	-1.299*** (0.231)	0.353 (0.233)	-0.945*** (0.157)	-0.675** (0.281)	-0.0448 (0.170)	-0.486*** (0.150)	-0.125 (0.283)
Bin [p90,max]	-0.446 (0.385)	-0.888*** (0.238)	-0.417** (0.206)	-0.747*** (0.176)	-1.322*** (0.167)	-1.347*** (0.287)	-0.460** (0.198)	-0.117 (0.151)	-1.159*** (0.268)	-1.359*** (0.275)	0.355 (0.223)	-0.655*** (0.163)	-0.680** (0.274)	0.109 (0.153)	-0.287** (0.141)	-0.210 (0.236)
Monsoon prec	0.517*** (0.148)	1.005*** (0.137)	0.0942 (0.0808)	0.345*** (0.104)	-0.196*** (0.0744)	0.948*** (0.156)	0.392*** (0.0818)	0.418*** (0.0504)	0.0263 (0.0857)	0.599*** (0.191)	0.289*** (0.0517)	0.419*** (0.0629)	0.253*** (0.0712)	0.134*** (0.0407)	0.229*** (0.0313)	0.142** (0.0604)
Monsoon prec ²	-0.145*** (0.0552)	-0.351*** (0.0534)	-0.00903 (0.0254)	-0.115*** (0.0373)	0.0117 (0.0266)	-0.453*** (0.0725)	-0.123*** (0.0318)	-0.0869*** (0.0134)	-0.0486* (0.0279)	-0.283*** (0.0881)	-0.0521*** (0.0180)	-0.0883*** (0.0224)	-0.0604*** (0.0205)	-0.0193 (0.0119)	-0.0418*** (0.0103)	-0.0494*** (0.0184)
Rabi prec																
Rabi prec ²																
Observations	2114	2771	1931	4543	5280	3501	4933	5424	5011	3989	2747	5050	2560	4567	4989	4786
R ²	0.734	0.660	0.866	0.648	0.735	0.734	0.592	0.835	0.658	0.624	0.768	0.599	0.704	0.817	0.855	0.846

This table shows the coefficients estimated from equation (1) using *decile temperature bins with a sample period starting in 1997*. Definitions of the variables may be found in Section 3. See also Table A2 for the limits of each temperature bin in degrees Celsius. Standard errors clustered by district in parentheses: * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Table A8: Yield-weather estimation – three-degree temperature bins

	(1) Castor	(2) Cotton	(3) Finger Millet	(4) Groundnut	(5) Maize	(6) Pearl Millet	(7) Pigeon Pea	(8) Rice	(9) Sesame	(10) Sorghum	(11) Barley	(12) Chickpea	(13) Linsseed	(14) Rapeseed/Mustard	(15) Wheat	(16) Sugarcane
Bin $(-\infty, 0)$			1.575 (1.104)	3.256 (4.049)	-0.0517 (0.730)	-2.669* (1.489)	-1.195 (0.896)	0.290 (0.730)	-0.0386 (0.975)		2.600*** (0.440)	3.376*** (1.458)	6.355*** (1.279)	4.432*** (1.458)	3.406*** (1.889)	-5.072*** (1.889)
Bin [0, 3]			2.516** (1.056)	3.580 (3.797)	0.691 (0.741)	-12.44 (7.934)	-0.103 (0.901)	1.001** (0.489)	-0.292 (1.044)		1.593*** (0.297)	2.812*** (1.023)	3.941*** (0.718)	3.401*** (0.856)	2.186*** (0.272)	-0.717 (1.384)
Bin [3, 6]			2.223** (1.011)	1.147 (3.895)	0.886 (0.685)		-0.767 (1.251)	1.395** (0.484)	1.395** (0.703)		2.074*** (0.290)	2.668** (1.054)	2.399** (0.935)	3.051*** (0.848)	2.031*** (0.287)	-4.245 (2.581)
Bin [6, 9]			1.800* (0.992)	1.186 (3.994)	1.062 (0.720)	1.223 (2.525)	-0.139 (0.885)	0.563 (0.351)	0.454 (0.802)		1.367*** (0.202)	2.213*** (0.592)	0.757 (0.582)	2.346*** (0.593)	1.553*** (0.240)	-2.086** (1.054)
Bin [9, 12]		0.199 (1.136)	2.496*** (0.854)	0.742 (2.415)	0.835 (0.650)	-2.192*** (0.695)	0.123 (1.033)	0.522 (0.340)	0.795 (0.629)	-164.5*** (22.84)	0.532*** (0.171)	0.159 (0.298)	0.242 (0.365)	-0.299 (0.247)	0.505*** (0.140)	1.332*** (0.319)
Bin [12, 15]		3.637* (2.019)	2.230*** (0.706)	1.960* (1.029)	1.196 (0.754)	-1.213 (1.555)	0.148 (0.717)	0.545** (0.226)	-0.0126 (0.741)	10.42*** (0.850)	0.453*** (0.116)	-0.00188 (0.160)	0.613*** (0.170)	0.402*** (0.130)	0.575*** (0.0817)	1.177*** (0.192)
Bin [15, 18]	-33.59 (26.27)	2.937** (1.384)	1.329** (0.452)	0.789 (0.771)	0.481 (0.400)	0.622 (0.612)	0.0183 (0.587)	0.387 (0.238)	-0.199 (0.485)	6.958*** (2.536)	0.276** (0.113)	0.0497 (0.105)	0.616*** (0.140)	0.376*** (0.102)	0.464*** (0.0731)	1.088*** (0.158)
Bin [18, 21]	0.212 (0.408)	-1.193*** (0.408)	0.408** (0.194)	0.188 (0.130)	0.586*** (0.160)	0.839*** (0.315)	0.566*** (0.178)	0.521*** (0.113)	0.484** (0.227)	0.702*** (0.214)	0.143 (0.0961)	0.212** (0.0882)	0.119 (0.111)	0.142 (0.0960)	0.131** (0.0664)	0.340** (0.133)
Bin [24, 27]	0.292) (0.129)	0.306) (0.127)	-0.265*** (0.0659)	-0.0911 (0.0838)	-0.377*** (0.0840)	-0.357*** (0.0957)	-0.525*** (0.0828)	-0.445*** (0.0763)	-0.209** (0.0875)	-0.374*** (0.113)	-0.0844 (0.0944)	-0.461*** (0.0944)	-0.463*** (0.138)	-0.373*** (0.0989)	-0.281*** (0.0778)	-0.231** (0.0962)
Bin [27, 30]	-0.390** (0.177)	-0.565*** (0.148)	-0.541*** (0.0989)	-0.488*** (0.0900)	-0.774*** (0.0987)	-0.860*** (0.117)	-0.985*** (0.0941)	-1.015*** (0.0929)	-0.607*** (0.106)	-0.992*** (0.117)	0.0840 (0.131)	-0.616*** (0.106)	-0.411** (0.161)	-0.478*** (0.118)	-0.374*** (0.0840)	-0.583*** (0.108)
Bin [30, 33]	-0.541** (0.263)	-0.722*** (0.172)	-0.858*** (0.235)	-0.897*** (0.146)	-1.584*** (0.140)	-1.638*** (0.196)	-1.181*** (0.132)	-0.965*** (0.113)	-1.692*** (0.210)	-1.713*** (0.193)	-0.200 (0.143)	-0.694*** (0.151)	-0.586*** (0.226)	-0.605*** (0.147)	-0.385*** (0.103)	-0.496*** (0.141)
Bin [33, 36]	-0.502 (0.405)	-0.492** (0.213)	-0.564** (0.233)	-0.277 (0.161)	-1.376*** (0.145)	-1.593*** (0.222)	-1.305*** (0.183)	-0.637*** (0.124)	-1.478*** (0.241)	-1.456*** (0.218)	-1.237*** (0.416)	-1.083*** (0.287)	-1.863*** (0.394)	-1.305*** (0.293)	-1.100*** (0.224)	-0.875*** (0.158)
Bin [36, 39]	-1.829*** (0.667)	-0.938*** (0.391)	-0.762 (0.772)	-0.387 (0.242)	-1.564*** (0.250)	-2.095*** (0.327)	-0.874*** (0.261)	-0.736*** (0.182)	-1.540*** (0.343)	-1.377*** (0.302)	2.598 (3.665)	-2.091 (1.987)	4.453 (3.972)	-3.685 (3.970)	-2.085 (2.265)	0.257 (0.311)
Bin [39, ∞]	-10.07 (8.745)	5.872 (4.589)	21.37*** (5.388)	-3.731** (1.827)	3.754 (2.486)	2.496 (4.190)	1.610 (5.998)	-0.549 (2.203)	5.206** (2.342)	-2.032 (1.923)						
Monsoon prec	0.470*** (0.118)	0.475*** (0.110)	0.185*** (0.0649)	0.411*** (0.0896)	-0.202*** (0.0616)	0.896*** (0.134)	0.248*** (0.0583)	0.381*** (0.0412)	-0.0000934 (0.0657)	0.472*** (0.112)	0.202*** (0.0349)	0.317*** (0.0380)	0.185*** (0.0452)	0.160*** (0.0303)	0.211*** (0.0243)	0.0609* (0.0366)
Monsoon prec ²	-0.143*** (0.0461)	-0.165*** (0.0469)	-0.0390** (0.0185)	-0.129*** (0.0330)	0.0294 (0.0220)	-0.444*** (0.0641)	-0.0776*** (0.0214)	-0.0829*** (0.0120)	-0.0392* (0.0208)	-0.228*** (0.0308)	-0.0191* (0.0112)	-0.0559*** (0.0123)	-0.0355*** (0.0117)	-0.0173* (0.00921)	-0.0329*** (0.00797)	-0.0249** (0.0114)
Rabi prec																
Rabi prec ²																
Observations	3683	5537	3905	8690	10186	6985	9485	10685	9633	8044	5589	9857	5010	8655	9806	9521
R ²	0.698	0.614	0.784	0.530	0.676	0.706	0.600	0.811	0.628	0.604	0.756	0.565	0.687	0.731	0.837	0.822

This table shows the coefficients estimated from equation (1) using *three-degree temperature bins*. Definitions of the variables may be found in Section 3. Standard errors clustered by district in parentheses: * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Table A9: Yield-weather estimation – decile temperature bins including HDD variable

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Castor	0.430	-0.469*	0.453***	0.238	0.630***	0.879***	0.871***	0.682***	0.771***	0.861***	0.386***	-0.0702	0.453**	0.238*	0.521***	1.193***
Bin [min,p10]	(0.273)	(0.253)	(0.147)	(0.168)	(0.155)	(0.186)	(0.147)	(0.129)	(0.154)	(0.181)	(0.144)	(0.183)	(0.186)	(0.141)	(0.105)	(0.208)
Bin [p10,p20]	0.0584	-0.0956	0.250*	-0.0324	0.251***	0.451***	0.501***	0.337***	0.204	0.238*	0.193	-0.0531	0.626***	0.447***	0.451***	0.902***
Bin [p20,p30]	0.227	(0.207)	(0.127)	(0.116)	(0.119)	(0.141)	(0.120)	(0.104)	(0.132)	(0.123)	(0.132)	(0.118)	(0.160)	(0.111)	(0.0807)	(0.153)
Bin [p30,p40]	0.175	-0.272	0.0284	-0.202	-0.0466	0.318*	0.386**	0.0983	0.0708	-0.206	0.0334	-0.00254	0.177	0.207	0.143	0.263
Bin [p40,p50]	(0.295)	(0.234)	(0.152)	(0.137)	(0.148)	(0.189)	(0.155)	(0.128)	(0.146)	(0.160)	(0.150)	(0.115)	(0.179)	(0.130)	(0.0865)	(0.162)
Bin [p50,p60]	0.268	-0.454*	0.174	-0.167	-0.104	0.175	-0.0470	-0.219*	0.230	-0.183	-0.174	0.0661	0.139	0.139	0.0291	-0.0128
Bin [p60,p70]	(0.313)	(0.243)	(0.193)	(0.147)	(0.187)	(0.204)	(0.159)	(0.122)	(0.188)	(0.189)	(0.160)	(0.145)	(0.186)	(0.161)	(0.114)	(0.205)
Bin [p70,p80]	-0.140	-0.295	-0.293	-0.0510	-0.195	0.0497	0.0578	-0.213*	-0.0859	0.0801	-0.0576	-0.300**	-0.197	-0.186	-0.146	-0.0386
Bin [p80,p90]	(0.293)	(0.235)	(0.189)	(0.165)	(0.129)	(0.166)	(0.147)	(0.110)	(0.140)	(0.138)	(0.162)	(0.117)	(0.180)	(0.142)	(0.100)	(0.166)
Bin [p90,max]	-0.150	-0.610**	-0.0686	-0.132	-0.333**	-0.0918	-0.459**	-0.508**	0.162	-0.472*	-0.343*	-0.628**	-0.679**	-0.372**	-0.366**	-0.345**
High degree days	(0.304)	(0.214)	(0.209)	(0.132)	(0.143)	(0.177)	(0.130)	(0.126)	(0.175)	(0.189)	(0.183)	(0.148)	(0.228)	(0.172)	(0.113)	(0.164)
Monsoon prec	-0.261	-0.674***	-0.191	-0.671***	-0.396***	-0.445**	-0.384***	-0.629***	-0.564***	-0.808***	-0.240	-0.547***	-0.401*	-0.416***	-0.353***	-0.490***
Monsoon prec ²	(0.263)	(0.203)	(0.196)	(0.158)	(0.140)	(0.182)	(0.147)	(0.118)	(0.166)	(0.180)	(0.155)	(0.124)	(0.205)	(0.157)	(0.122)	(0.163)
Rabi prec	-0.290	-0.920**	-0.588**	-0.749**	-0.848**	-0.674***	-0.423**	-0.766**	-0.717**	-1.001**	-0.147	-0.740**	-0.591**	-0.433**	-0.431**	-0.369**
Rabi prec ²	(0.240)	(0.187)	(0.180)	(0.118)	(0.141)	(0.158)	(0.129)	(0.126)	(0.146)	(0.177)	(0.154)	(0.124)	(0.193)	(0.149)	(0.107)	(0.182)
Observations	3683	5537	3905	8690	10186	6985	9485	10685	9633	8044	5589	9857	5010	8655	9806	9521
R ²	0.698	0.613	0.783	0.531	0.676	0.706	0.600	0.811	0.629	0.606	0.754	0.562	0.686	0.727	0.836	0.821

This table shows the coefficients estimated from equation (1) using *decile temperature bins* and an *HDD variable*. Definitions of the variables may be found in Section 3. See also Table A2 for the limits of each temperature bin in degrees Celsius. Standard errors clustered by district in parentheses: * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Table A10: Yield-weather estimation – degree days with additional precipitation variables

	(1) Castor	(2) Cotton	(3) Finger Millet	(4) Groundnut	(5) Maize	(6) Pearl Millet	(7) Pigeon Pea	(8) Rice	(9) Sesame	(10) Sorghum	(11) Barley	(12) Chickpea	(13) Linsed	(14) Rapeseed/Mustard	(15) Wheat	(16) Sugarcane
Degree days	-0.000361** (0.000172)	-0.000348*** (0.000123)	-0.000309** (0.000136)	-0.000240*** (0.0000804)	-0.000558*** (0.0000779)	-0.000762*** (0.000127)	-0.000462*** (0.0000823)	-0.000652 (0.0000591)	-0.000729*** (0.000126)	-0.000691*** (0.000114)	-0.000279*** (0.0000733)	-0.000545*** (0.0000672)	-0.000683*** (0.0000930)	-0.000521*** (0.0000696)	-0.000483*** (0.0000528)	-0.000150*** (0.0000329)
Monsoon prec	0.406*** (0.120)	0.505*** (0.118)	0.188*** (0.0629)	0.454*** (0.0971)	-0.211*** (0.0577)	0.907*** (0.134)	0.317*** (0.0631)	0.436*** (0.0430)	-0.0117 (0.0679)	0.597*** (0.108)	0.221*** (0.0349)	0.337*** (0.0367)	0.217*** (0.0445)	0.185*** (0.0313)	0.241*** (0.0243)	0.0572 (0.0385)
Monsoon prec ²	-0.122*** (0.0425)	-0.165*** (0.0488)	-0.0342* (0.0189)	-0.143*** (0.0350)	0.0232 (0.0222)	-0.457*** (0.0643)	-0.0830*** (0.0218)	-0.0821*** (0.0122)	-0.0482** (0.0215)	-0.258*** (0.0520)	-0.0198* (0.0112)	-0.0565*** (0.0115)	-0.0420*** (0.0117)	-0.0188* (0.00965)	-0.0380*** (0.00795)	-0.0306*** (0.0113)
Rainy days	0.00185 (0.00130)	0.00162 (0.00100)	0.00353*** (0.000946)	0.00242*** (0.000709)	0.00548*** (0.000542)	0.00443*** (0.000728)	0.00190*** (0.000668)	0.00437*** (0.000473)	0.00535*** (0.000802)	0.00314*** (0.000731)	0.000453 (0.000453)	-0.00166*** (0.000417)	-0.00190*** (0.000636)	-0.00167*** (0.000484)	-0.00100*** (0.000315)	0.00144*** (0.000263)
Heavy rain	-0.00420 (0.0122)	-0.00840 (0.0106)	-0.0143* (0.00773)	0.00794 (0.00622)	0.00986 (0.00667)	0.00129 (0.00897)	-0.0142* (0.00633)	-0.0185*** (0.00473)	0.0149* (0.00836)	-0.00924 (0.00917)						0.00703 (0.00664)
Dry spell	-0.00470*** (0.00162)	0.000659 (0.00122)	-0.000607 (0.00192)	-0.000515 (0.000582)	0.000297 (0.000688)	-0.00100 (0.000757)	0.000269 (0.00059)	0.000712 (0.000573)	0.00179* (0.00104)	0.00151* (0.000828)	0.000236 (0.000199)	0.000347** (0.000167)	0.000602** (0.000250)	0.000397*** (0.000138)	0.00000431 (0.000109)	0.000208 (0.000129)
Rabi prec																
Rabi prec ²																
Observations	3683	5537	3905	8690	10186	6985	9485	10685	9633	8044	5589	9857	5010	8655	9806	9521
R ²	0.700	0.611	0.784	0.527	0.678	0.707	0.597	0.808	0.628	0.602	0.752	0.561	0.685	0.727	0.834	0.820

This table shows the coefficients estimated from equation (1) using *degree days with additional precipitation variables*. Definitions of the variables may be found in Section 3. Standard errors clustered by district in parentheses: * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Table A11: Yield-weather estimation – degree days and precipitation bins

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	Castor	Cotton	Finger Millet	Groundnut	Maize	Pearl Millet	Pigeon Pea	Rice	Sesame	Sorghum	Barley	Chickpea	Linsed	Repesced/Mustard	Wheat	Sugarcane
Degree days	-0.00052*** (0.000184)	-0.000643*** (0.0000863)	-0.000516*** (0.000122)	-0.000409*** (0.0000783)	-0.000792*** (0.0000708)	-0.00131*** (0.000162)	-0.000399*** (0.0000768)	-0.000414*** (0.0000533)	-0.00106*** (0.000142)	-0.000998*** (0.000129)	-0.000338*** (0.0000709)	-0.000469*** (0.0000613)	-0.000593*** (0.0000860)	-0.000489*** (0.0000651)	-0.000456*** (0.0000492)	-0.000208*** (0.0000299)
Monsoon prec bin [min,p10]	-0.159*** (0.0445)	-0.0865** (0.0432)	-0.0354 (0.0399)	-0.241*** (0.0470)	-0.121*** (0.0269)	-0.161*** (0.0423)	-0.145*** (0.0320)	-0.141*** (0.0285)	-0.142*** (0.0399)	-0.241*** (0.0403)	-0.0514*** (0.0186)	-0.200*** (0.0217)	-0.0442 (0.0417)	-0.0855*** (0.0201)	-0.105*** (0.0150)	-0.0402** (0.0178)
Monsoon prec bin [p10,p20]	-0.0198 (0.0277)	-0.0426* (0.0244)	-0.0301 (0.0297)	-0.117*** (0.0226)	-0.0392** (0.0189)	-0.0628** (0.0260)	-0.0997*** (0.0244)	-0.116*** (0.0217)	-0.0943*** (0.0242)	-0.0943*** (0.0272)	-0.0623*** (0.0173)	-0.123*** (0.0156)	-0.113*** (0.0324)	-0.0640*** (0.0190)	-0.0871*** (0.0128)	-0.0345** (0.0141)
Monsoon prec bin [p20,p30]	-0.0209 (0.0299)	-0.0497** (0.0232)	0.0219 (0.0282)	-0.0520*** (0.0170)	-0.00737 (0.0152)	-0.0148 (0.0169)	-0.0480*** (0.0172)	-0.0588*** (0.0149)	0.0225 (0.0215)	-0.0325* (0.0194)	-0.0356** (0.0166)	-0.0521*** (0.0121)	-0.0224 (0.0200)	-0.00770 (0.0161)	-0.0260** (0.0106)	-0.00551 (0.0115)
Monsoon prec bin [p40,p50]	0.0753** (0.0308)	0.0222 (0.0254)	0.0488** (0.0200)	0.0191 (0.0160)	0.00101 (0.0158)	0.00516 (0.0170)	0.0339** (0.0141)	0.0364*** (0.0129)	-0.0346 (0.0223)	0.00722 (0.0172)	0.0461*** (0.0143)	0.00528 (0.0124)	0.0367* (0.0187)	0.0224 (0.0146)	0.104 (0.0103)	-0.00129 (0.0130)
Monsoon prec bin [p50,p60]	0.0644** (0.0257)	0.0335 (0.0264)	0.0684*** (0.0235)	0.0322* (0.0167)	-0.0374** (0.0157)	-0.00225 (0.0186)	0.0275* (0.0165)	0.0691*** (0.0137)	-0.0547** (0.0229)	-0.00702 (0.0166)	0.0564*** (0.0158)	0.0434*** (0.0115)	0.0427** (0.0189)	0.0570*** (0.0164)	0.0468*** (0.0103)	0.0167 (0.0119)
Monsoon prec bin [p60,p70]	0.0412 (0.0356)	0.0829*** (0.0309)	0.0919*** (0.0269)	0.0461*** (0.0164)	-0.0440*** (0.0199)	-0.0378* (0.0157)	0.0509*** (0.0175)	0.127*** (0.0139)	-0.0955*** (0.0240)	-0.0355* (0.0189)	0.0924*** (0.0156)	0.0433*** (0.0137)	0.0632*** (0.0207)	0.0484*** (0.0161)	0.0772*** (0.0107)	0.0110 (0.0156)
Monsoon prec bin [p70,p80]	0.0503 (0.0355)	0.0327 (0.0339)	0.0890*** (0.0268)	0.0336* (0.0184)	-0.106*** (0.0199)	-0.0738*** (0.0254)	0.0539*** (0.0169)	0.145*** (0.0147)	-0.113*** (0.0265)	-0.0405* (0.0209)	0.113*** (0.0163)	0.0595*** (0.0155)	0.0643*** (0.0228)	0.0501*** (0.0173)	0.0680*** (0.0106)	0.00182 (0.0157)
Monsoon prec bin [p80,p90]	0.0467 (0.0379)	0.0690** (0.0328)	0.124*** (0.0310)	0.0351* (0.0193)	-0.139*** (0.0228)	-0.112*** (0.0280)	0.0334* (0.0196)	0.157*** (0.0161)	-0.138*** (0.0270)	-0.0783*** (0.0258)	0.133*** (0.0208)	0.0686*** (0.0187)	0.0467* (0.0254)	0.0956*** (0.0181)	0.0836*** (0.0122)	0.00501 (0.0210)
Monsoon prec bin [p90,max]	0.0774* (0.0468)	0.0901** (0.0448)	0.111*** (0.0330)	0.0350 (0.0234)	-0.185*** (0.0287)	-0.107*** (0.0363)	0.0260 (0.0278)	-0.192*** (0.0203)	-0.179*** (0.0340)	-0.116*** (0.0376)	0.255*** (0.0364)	0.147*** (0.0256)	0.110*** (0.0489)	0.185*** (0.0337)	0.147*** (0.0224)	-0.0520 (0.0467)
Rabi prec bin [min,p10]																
Rabi prec bin [p10,p20]																
Rabi prec bin [p20,p30]																
Rabi prec bin [p40,p50]																
Rabi prec bin [p50,p60]																
Rabi prec bin [p60,p70]																
Rabi prec bin [p70,p80]																
Rabi prec bin [p80,p90]																
Rabi prec bin [p90,max]																
Observations	3683	5537	3905	8690	10186	6985	9485	10685	9633	8044	5589	9857	5010	8655	9806	9521
R ²	0.698	0.610	0.782	0.527	0.678	0.702	0.598	0.806	0.628	0.602	0.753	0.562	0.685	0.727	0.835	0.819

This table shows the coefficients estimated from equation (1) using *degree days with decile precipitation bins*. Definitions of the variables may be found in Section 3. Standard errors clustered by district in parentheses: * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Table A12: Yield-weather estimation – degree days (starting 1997)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	Castor	Cotton	Finger Millet	Groundnut	Maize	Pearl Millet	Pigeon Pea	Rice	Sesame	Sorghum	Barley	Chickpea	Linseed	Rapeseed/Mustard	Wheat	Sugarcane
Degree days	-0.000541** (0.000235)	-0.000638*** (0.000125)	-0.000373*** (0.000130)	-0.000577*** (0.0000978)	-0.000983*** (0.000103)	-0.00113*** (0.000189)	-0.000560*** (0.0000972)	-0.000135* (0.0000783)	-0.000986*** (0.000174)	-0.000969*** (0.000177)	0.0000436 (0.000106)	-0.000326*** (0.0000813)	-0.000505*** (0.000155)	-0.0000635 (0.0000787)	-0.000253*** (0.0000665)	-0.000159*** (0.0000425)
Monsoon prec	0.562*** (0.145)	1.075*** (0.139)	0.174** (0.0822)	0.408*** (0.112)	-0.120 (0.0740)	1.057*** (0.153)	0.465*** (0.0839)	0.480*** (0.0530)	0.0722 (0.0865)	0.823*** (0.200)	0.302*** (0.0502)	0.448*** (0.0631)	0.284*** (0.0742)	0.158*** (0.0409)	0.266*** (0.0326)	0.188*** (0.0606)
Monsoon prec ²	-0.156*** (0.0550)	-0.360*** (0.0546)	-0.0257 (0.0254)	-0.145*** (0.0407)	-0.00720 (0.0274)	-0.496*** (0.0738)	-0.136*** (0.0329)	-0.0973*** (0.0139)	-0.0642** (0.0295)	-0.353*** (0.0944)	-0.0563*** (0.0172)	-0.0938*** (0.0226)	-0.0687*** (0.0213)	-0.0251** (0.0122)	-0.0511*** (0.0109)	-0.0565*** (0.0183)
Rabi prec																
Rabi prec ²																
Observations	2114	2771	1931	4543	5280	3501	4933	5424	5011	3989	2747	5050	2560	4567	4989	4786
R ²	0.733	0.657	0.864	0.642	0.733	0.733	0.588	0.833	0.655	0.618	0.766	0.594	0.702	0.816	0.852	0.844

This table shows the coefficients estimated from equation (1) using *degree days with a sample period starting in 1997*. Definitions of the variables may be found in Section 3. Standard errors clustered by district in parentheses: * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Table A13: Yield-weather estimation – degree days (baseline 23°C)

	(1) Castor	(2) Cotton	(3) Finger Millet	(4) Groundnut	(5) Maize	(6) Pearl Millet	(7) Pigeon Pea	(8) Rice	(9) Sesame	(10) Sorghum	(11) Barley	(12) Chickpea	(13) Linsed	(14) Rapeseed/Mustard	(15) Wheat	(16) Sugarcane
Degree days	-0.000475*** (0.000169)	-0.000458*** (0.0000885)	-0.000517*** (0.000111)	-0.000388*** (0.0000746)	-0.000916*** (0.0000735)	-0.00106*** (0.000125)	-0.000640*** (0.0000746)	-0.000422*** (0.0000550)	-0.00106*** (0.000128)	-0.000933*** (0.000117)	-0.000273*** (0.0000649)	-0.000491*** (0.0000571)	-0.000602*** (0.0000828)	-0.000454*** (0.0000603)	-0.000427*** (0.0000471)	-0.000222*** (0.0000272)
Monsoon prec	0.465*** (0.114)	0.512*** (0.114)	0.226*** (0.0655)	0.501*** (0.0959)	-0.139*** (0.0620)	0.979*** (0.131)	0.305*** (0.0585)	0.452*** (0.0434)	0.0663 (0.0666)	0.609*** (0.112)	0.220*** (0.0350)	0.331*** (0.0364)	0.213*** (0.0446)	0.184*** (0.0314)	0.236*** (0.0242)	0.0859** (0.0373)
Monsoon prec ²	-0.140*** (0.0445)	-0.173*** (0.0484)	-0.0489*** (0.0191)	-0.150*** (0.0353)	0.0126 (0.0232)	-0.474*** (0.0645)	-0.0879*** (0.0220)	-0.0943*** (0.0126)	-0.0596*** (0.0223)	-0.268*** (0.0529)	-0.0190* (0.0112)	-0.0564*** (0.0114)	-0.0432*** (0.0117)	-0.0201** (0.00964)	-0.0377*** (0.00788)	-0.0314*** (0.0116)
Rabi prec																
Rabi prec ²																
Observations	3083	5537	3905	8690	10186	6985	9485	10085	9633	8044	5589	9857	5010	8655	9806	9521
R ²	0.697	0.611	0.782	0.525	0.674	0.705	0.597	0.806	0.625	0.601	0.752	0.560	0.683	0.726	0.834	0.819

This table shows the coefficients estimated from equation (1) using *degree days with a baseline temperature of 23°C*. Definitions of the variables may be found in Section 3. Standard errors clustered by district in parentheses: * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

Table A14: Yield-weather estimation – degree days (baseline 25°C)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	Castor	Cotton	Finger Millet	Groundnut	Maize	Pearl Millet	Pigeon Pea	Rice	Sesame	Sorghum	Barley	Chickpea	Linsed	Rapeseed/Mustard	Wheat	Sugarcane
Degree days	-0.000502*** (0.000184)	-0.000486*** (0.0000942)	-0.000490*** (0.000136)	-0.000402*** (0.0000824)	-0.000634*** (0.0000795)	-0.00108*** (0.000140)	-0.000569*** (0.0000825)	-0.000329*** (0.0000593)	-0.00113*** (0.000143)	-0.000937*** (0.000129)	-0.000333*** (0.0000813)	-0.000538*** (0.0000739)	-0.000698*** (0.000108)	-0.000509*** (0.0000753)	-0.000495*** (0.0000573)	-0.000214*** (0.0000313)
Monsoon prec	0.475*** (0.114)	0.519*** (0.110)	0.259*** (0.0641)	0.512*** (0.0933)	-0.108* (0.0617)	1.022*** (0.129)	0.357*** (0.0613)	0.497*** (0.0448)	0.0777 (0.0658)	0.646*** (0.111)	0.223*** (0.0352)	0.349*** (0.0373)	0.230*** (0.0458)	0.195*** (0.0308)	0.248*** (0.0245)	0.106*** (0.0374)
Monsoon prec ²	-0.141*** (0.0443)	-0.173*** (0.0473)	-0.0556*** (0.0188)	-0.152*** (0.0348)	0.00646 (0.0233)	-0.486*** (0.0638)	-0.0991*** (0.0231)	-0.104*** (0.0131)	-0.0605*** (0.0220)	-0.275*** (0.0528)	-0.0199* (0.0112)	-0.0602*** (0.0119)	-0.0467*** (0.0121)	-0.0224** (0.00947)	-0.0400*** (0.00803)	-0.0344*** (0.0117)
Rabi prec																
Rabi prec ²																
Observations	3083	5537	3905	8690	10186	6985	9485	10085	9633	8044	5589	9857	5010	8655	9806	9521
R ²	0.697	0.611	0.782	0.525	0.673	0.704	0.596	0.805	0.625	0.600	0.752	0.559	0.682	0.725	0.834	0.819

This table shows the coefficients estimated from equation (1) using *degree days with a baseline temperature of 25°C*. Definitions of the variables may be found in Section 3. Standard errors clustered by district in parentheses: * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.