CAHIER DE RECHERCHE #2003E
Département de science économique
Faculté des sciences sociales
Université d'Ottawa

WORKING PAPER #2003E
Department of Economics
Faculty of Social Sciences
University of Ottawa

Climate Change and the Distribution of Agricultural Output[*]

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

April 2020

[*] We thank Carolyn Fischer, Anthony Heyes and Maya Papineau, as well as participants at the AERE Summer Conference, the Economics of Low-Carbon Markets Workshop, the Occasional Workshop in Environmental and Resource Economics, the Institutions Trade and Economic Development (InsTED) Workshop, the CIREQ Environmental and Natural Resource Economics Workshop, the Brazilian Econometric Society Meeting, the Environment for Development Annual Meeting, the Canadian Resource and Environmental Economics Annual Meeting, the Moscow International Economics Workshop, the Annual Meeting of the Canadian Economics Association, and 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. This study was financed in part by CAPES/Brasil, Grant 001.
[†] FGV EPGE, Rio de Janeiro, Brazil. Email: francisco.costa@fgv.br.
[‡] Department of Economics, University of Ottawa, Ottawa, Canada. Email: fforg097@uottawa.ca.
[§] Department of Economics, University of Ottawa, Ottawa, Canada. Email: Jason.Garred@uottawa.ca.
[**] Sao Paulo School of Economics - FGV, São Paulo, Brazil; Centre for Economic Performance, London, UK. Email: joao.pessoa@fgv.br.

Abstract

This paper uses a multi-run climate projection model to examine the potential impact of climate change on the distribution of agricultural outcomes in India. Weather draws resulting in extremely low agricultural revenues (1-in-100-year events) are projected to become the norm, increasing by 53 to 88 percentage points by the end of the 21st century. As a result, Indian farmers will face a 16% to 33% decline in mean revenue over the course of the century, presenting a more urgent problem than changes in yield variability. Analysis using a structural general equilibrium model suggests consequences of a similar magnitude for welfare.

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

JEL Classification: Q15, Q54, Q56, O13.

1 Introduction

Climate change may have important consequences for agricultural output around the globe. This might include not only the average crop yields a farmer can expect in the future, but also the year-to-year volatility of agricultural production and the extent of downside risk. Changes in the variability of yields may be of particular concern for farmers in low-income countries, where insurance markets are often poorly developed.

A large body of work examines the projected effects of climate change on average agricultural outcomes.[1] However, analyses looking beyond the first moment of the potential outcome distribution are less common. Such studies face an important constraint: data for a given spatial location and moment in time constitutes only a single draw from the distribution of interest. Yet an investigation of higher-order moments or extreme events depends on information about a much larger number of possible realizations.

In this paper, we consider the potential impact of climate change on the distribution of agricultural outcomes faced by farmers. Previous studies have investigated the potential future variability of agricultural yields by relying on projections of year-to-year weather variation.[2] Here, we instead generate a large number of different projected yields from which farmers' outcomes may be drawn, for each location, crop and time period in our sample. We can therefore leverage many potential outcome realizations to analyze average agricultural outcomes and their variability within a given period, as well as the frequency of extreme events. We perform this exercise for India, a context with limited availability of insurance as compared to more developed economies.

The key data input facilitating this analysis is the Community Earth System Model Large Ensemble (CESM-LENS) dataset. This is a multi-run global simulation created by climatologists in order to study the potential variability of temperature and precipitation patterns within a given climate, and how this variability might evolve over time with climate change (Kay et al., 2015).[3] CESM-LENS provides us with multiple potential realizations of weather at a given location for each day between 1920 and 2099. For instance, rather than specifying

[1] This includes many studies of the United States (e.g. Adams et al., 1990; Mendelsohn et al., 1994; Schlenker et al., 2005; Schlenker and Roberts, 2009; Fisher et al., 2012; Burke and Emerick, 2016) and also of developing countries (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).

[2] For instance, Urban et al. (2012) uses the predicted time-series variation in temperature and precipitation patterns from one run of a given climate model to project future year-to-year variability in US maize yields, and then repeats this exercise separately for fifteen different climate models. Tigchelaar et al. (2018) projects changes in the interannual volatility of global maize production under the scenario that global mean temperatures rise by either two or four degrees Celsius, but assuming that year-to-year temperature variation remains as observed in the period 1979 to 2008.

[3] The CESM-LENS data has also been used to predict changes in the incidence of extreme weather events (Swain et al., 2018).

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 that are forecast to be consistent with the prevailing climate as of 2050.[4] We use this data to predict the evolution of the distribution of agricultural production for 310 Indian districts across eighteen decades, based on 400 possible weather realizations per decade.

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). We estimate several different models relating historical yields across Indian districts to realized temperature and precipitation from the ERA5 weather dataset. We define annual temperature in terms of the number of days in various temperature ranges (bins), or alternatively in terms of degree days. For precipitation, we use either a quadratic in total precipitation or an augmented specification also including variables capturing the within-season distribution of rain. All in all, we estimate ten yield-weather specifications for each of the sixteen different crops for which we have district-level agricultural data.

We then generate a set of yield projections based on the results of each of these specifications. In particular, we use the 400 realizations of our temperature and precipitation variables from the CESM-LENS data to calculate 400 potential yield realizations per crop for a given district and decade. We employ these projections to explore the consequences of climate change for the distribution of total agricultural revenue in each district, decade by decade, holding all factors other than the climate constant at their levels in 2000.

We begin with the first moment of this distribution, by calculating predicted average revenue for each district and decade. The results are consistent across the various projections based on our different regression specifications: on average across districts, we forecast a gradual decline in mean agricultural revenue between the first and the last decade of the 21st century, ranging from 15.5% to 33.1%. These effects are almost entirely driven by future rises in average temperature rather than changes in precipitation patterns.[5]

Next, we consider the second moment, assessing potential changes in weather-induced variation in revenue. Estimates of the impact of climate change on this dimension are very 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 59.9% between the 2000s and the 2090s. On the other hand, models based on degree days forecast a rise in this measure of up to 18.8%. These

[4]CESM-LENS is based on the following climate trajectory: for the first half of the sample period (1920 to 2005), the set of simulated weather draws for each location is based on the observed evolution of the global climate, and for 2006 to 2099, the simulation generates multiple realizations of local weather consistent with the representative concentration pathway 8.5 (RCP8.5) climate change scenario.

[5]The magnitudes of these estimates are within the range found in other studies that consider the likely impact of climate change on average yields in India (e.g. Guiteras, 2009; Burgess et al., 2017; Hari et al., 2018).

figures represent only the variation in agricultural revenue that is due to weather, rather than total within-district variation in agricultural outcomes due to all factors, since we hold factors other than the climate constant. So even the largest positive projection does not suggest that climate change will lead to a substantial rise in the total variability of agricultural revenue.

Our final 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, i.e. between 1920 and 1969. We find that the frequency of such bad years will rise by 53 to 88 percentage points in the average district over the course of 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.

We also take a structural approach to the analysis of our yield projections, using the general equilibrium portfolio choice model of Allen and Atkin (2016). This allows us to assess the impact of relaxing some of the assumptions underlying the above exercises. First, we let crop prices adjust endogenously with changes in agricultural production. Second, we allow farmers to adapt to the evolving distribution of yields by adjusting crop choices.[6] Third, we allow for changes in trade flows between Indian districts.[7]

By comparing the welfare measure implied by the model with a naive version using nominal revenue, we find that allowing for endogenous prices, crop choice and trade has a small moderating effect on the forecast losses from climate change. Also, results from our model-based counterfactuals suggest that these welfare losses will almost entirely be due to declines in mean outcomes rather than changes in portfolio risk. The model therefore forecasts a 21st-century fall in welfare on average across districts that is only slightly smaller than the predicted decline in mean agricultural revenue from our model-free exercise, in a range from 13.5% to 29.0%.

Overall, our findings predict that climate change will have important effects on average agricultural outcomes through the normalization of formerly extreme weather. On the other hand, our projections do not provide strong evidence of a large rise in future variability. This suggests that the relative weakness of insurance markets in less developed countries

[6]The study of Costinot et al. (2016) suggests that this type of adaptation could partially offset the potential welfare losses from climate change of farmers worldwide. Aragón et al. (2019) show that subsistence farmers in Peru adjust the crops they plant to mitigate the impacts of extreme heat events.

[7]Of course, the model cannot endogenize all potential margins of adaptation, though some of these have previously been investigated in the context of Indian agriculture. For example, Fishman (2018) finds that while irrigation can help in mitigating the negative effect of precipitation variability in India, successful adaptation to climate change is unlikely given current groundwater depletion rates. 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 are able to adapt to moderately hot temperatures, but not to episodes of extreme heat. Finally, the results of a number of studies suggest that climate change may affect labour supply in developing countries, due to migration (Jessoe et al., 2018), human capital accumulation (Garg et al., 2018), sectoral mobility (Colmer, 2018) and conflict (Crost et al., 2018).

such as India may not be of first-order importance in their farmers' welfare losses due to climate change. Instead, a more promising avenue of adaptation could be the development and adoption of crop varieties resistant to extreme heat.

Our paper builds on the existing literature in important ways. First, we use the large number of potential weather realizations from a multi-run climate model to predict changes in yield distributions across multiple short time periods (decades). Second, we project the yield distributions of all major crops in India, allowing us to make statements about changing returns from agriculture rather than individual crops. Third, we consider several different relevant dimensions of the distribution of agricultural outcomes, including downside risk.[8]

Another contribution of the paper is that we investigate the possible effects of relaxing three assumptions often made in the above literature – exogeneity of prices, crop choice and trade flows – using a general equilibrium model. Our paper thus also adds to a growing set of assessments of the potential effects of climate change that take advantage of the structure of general equilibrium models at the frontier of current research.[9] Differently from other papers in this literature, we consider a model in which climate change may affect agents' welfare through rises or falls in portfolio risk.[10]

2 Data

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

[8]Another recent paper considering downside risk is Emerick et al. (2016). This study uses a randomized experiment to show that the adoption of a new rice variety that reduces downside risk (through improved flood resistance) was particularly effective in improving productivity by crowding in other investments.

[9]Along 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, and the implications of this for the effects of climate change, using an Arrow-Debreu framework.

[10]Our work is also related to a recent set of papers considering impacts of climate change due to agents' uncertainty about the changing climate, such as Kala (2017), Lemoine (2017), and Shrader (2017). It is important to note the distinction between these studies of uncertainty and our exercise, which instead examines the effect of changes in the distribution of economic outcomes when this distribution is known to agents.

2.1 Historical weather and yield

Annual agricultural data at the level of the Indian district for 1979 to 2015 is provided by the ICRISAT Village Dynamics in South Asia (VDSA) dataset.[11] 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 data on farm-gate prices (also at the district-crop-year level). 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 as of the end of our sample period in 2015.[12]

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 Appendix Table A.1). 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, then the median number of crops drops to three. Also, in almost half of districts, one crop occupies at least 50% of land. Rice and wheat are the likeliest to be a district's most-planted crop by land area, with rice more important in the east and south of India and wheat more prominent in the north (Appendix Figure A.1).[13]

We construct district-level historical information on temperature and precipitation for 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, at a resolution of approximately 7000 square kilometers (0.75 degrees latitude by 0.75 degrees longitude).[14] We calculate mean daily temperature by averaging the minimum and maximum temperature for each day. Similarly, total daily precipitation is computed by cumulating precipitation throughout the day. We allocate this information across Indian districts by assigning to each district the

[11]The VDSA data defines a year as beginning in June and ending in May, in line with India's agricultural calendar. Throughout the paper, a year is therefore defined using this calendar; for example, 1980 actually 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.

[12]The 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 during the period covered by this dataset, it defines districts consistently across time using 1966 district borders by apportioning data from later years across 1966 districts.

[13]Note that these calculations are based only on land planted with the sixteen crops in our dataset, and do not account for agricultural land planted with other crops.

[14]This is approximately the size of the average Indian district in our dataset.

weather information of the ERA5 pixel closest to the centroid of that district.[15]

In order to relate these high-frequency weather variables to agricultural outcomes at the district-crop-year level, we use information on crop-specific growing seasons from India's Ministry of Agriculture and Farmers Welfare (2017).[16] 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.[17] We use this information to limit our weather data to the relevant growing season for each crop, as we will discuss further in Section 3.1.

2.2 Weather projection

For projected weather patterns under climate change, we rely on the Community Earth System Model Large Ensemble (CESM-LENS) dataset. This provides us with daily information on average temperature and total precipitation from 1920 to 2099 for forty simulations of global weather. The stated purpose of the CESM-LENS project is to provide climate change researchers with data allowing for the study of 'internal climate variability'; i.e. variation in realized weather within a given climate state (Kay et al., 2015). Each of the forty runs in the 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.[18] However, all runs are subject to the same progression over time in the climate. For the years 1920 to 2005, the simulations are based on historical climate conditions, while for 2006 to 2099, climate change is modelled based on the representative concentration pathway 8.5 (RCP8.5) scenario.[19]

The availability of a multi-run model is a crucial advantage for our analysis of the dis-

[15]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).

[16]This publication 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 the publication's appendix.

[17]We 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.

[18]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 (see footnote 8).

[19]RCP8.5 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).

tribution of agricultural outcomes under climate change. This is because it provides a large number of realizations of weather variables within a period of time in which we can reasonably treat the climate as fixed, such as a decade.[20] When analyzing the projected distribution of outcomes for each crop and district within a given time period, we model the realized weather variables in each run of the CESM-LENS simulations as occurring with equal probability. Such an assumption has previously been employed in the scientific literature using CESM-LENS to project future probabilities of extreme weather occurrences (e.g. Swain et al., 2018).

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 in order to compare 'baseline' and 'post-climate change' periods. Instead, we simply generate projections covering the entire 1920-2099 span, as we further explain in Section 3.3.

The CESM-LENS dataset is available at a slightly lower resolution (1 degree latitude by 1 degree longitude) than the ERA5 data. We therefore first regrid the projections to 0.75 by 0.75 degrees using linear interpolation so as to match the size and boundary coordinates of the ERA5 grid, and then link Indian districts to temperature projection data in the same way as above (i.e. based on the nearest grid cell to the district centroid).[21] To account for prediction error in the CESM-LENS projections, as well as possible systematic differences between the temperature and precipitation measures used in the two datasets, we take advantage of the fact that data is available for both ERA5 and CESM-LENS in some years. We calculate the daily district-level average difference between ERA5 and CESM-LENS temperature (or precipitation) over the years 1979 to 2015, and then add this to the CESM-LENS data for the corresponding district-day combinations in all years (analogously to Deschênes and Greenstone, 2011).[22]

In Figure 1, we display information about the projected evolution of the temperature and precipitation distributions in the adjusted CESM-LENS dataset, in the top and bottom panels respectively. 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.

According to the figure, at this India-wide level, the temperature distribution is projected

[20]For instance, 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.

[21]Other possible methods, such as taking within-district weighted averages across cells, offer only a limited improvement over this 'nearest neighbour' approach but are more computationally intensive.

[22]For example, if a given district has an average temperature on January 1 in ERA5 that is one degree lower than the mean predicted temperature on January 1 in CESM-LENS during 1979 to 2015, we reduce its projected January 1 temperatures by one degree throughout the 1920 to 2099 period. We set any negative precipitation values resulting from the adjustment to be equal to zero.

to shift sharply upwards around a rising mean over the course of the 21st century. However, substantial changes in temperature variability are not readily apparent, as the vertical range across runs in the top panel appears quite similar over time. In the bottom panel, we can see that the mean of the precipitation distribution is also forecast to rise, and that this does seem to be accompanied by a predicted increase in variability.

We look more closely at these projections in Table 1, now instead investigating the nature of the predicted changes for the average district. Specifically, we first find average annual temperature and total annual precipitation for each district-run-year. For each of the 310 districts in our dataset, we take the mean and standard deviation of each of these variables across the 400 runs available for each decade from the 1920s to the 2090s (forty runs per year over ten years). Table 1 then charts the average of each of these two moments across districts.

By this measure, the temperature distribution is projected to see a rise of more than four degrees Celsius in its first moment between the 2000s and the 2090s. While an accompanying rise in standard deviation is also predicted, the proportional change between the 2000s and 2090s is less than 10% (from 0.61 to 0.67). This is within the range of variation in standard deviation observed in the CESM-LENS data across the decades of the 20th century. Meanwhile, there is a forecast rise of approximately 14% in the mean of the precipitation distribution for the average district over the course of the 21st century, while its standard deviation is projected to increase by more than a third (from 0.21 in the 2000s to 0.29 in the 2090s).

The CESM-LENS projection also suggests that changes in both temperature and precipitation will manifest in a heterogeneous way across different parts of India. Appendix Figure A.2(a) shows that the largest temperature rises are forecast to occur in northern India. Precipitation is predicted to increase most steeply in India's northeast, as can be seen in Appendix Figure A.2(b).

As we will discuss in detail in the next section, we use the ERA5 data to estimate the relationships between weather variables and crop yields, and then apply our estimates to the CESM-LENS weather data. The extent of overlap in the distributions of temperature and precipitation in these two datasets will thus be important in determining the difficulty of translating our regression estimates into yield projections. Figure 2(a) displays box plots representing the distribution of temperature across district-days, comparing the ERA5 data from 1979 to 2015 with CESM-LENS data for the decade starting in 2090, when the climate is forecast to be least similar to that observed in ERA5.[23]

We first confirm that the ERA5 and CESM-LENS data are similar in the 1979-2015 period; this is by construction, because of our adjustment of the CESM-LENS data discussed above. However, the upward shift in the CESM-LENS temperature distribution over time is such that

[23]In Appendix Figures A.3(a) and A.3(b), we also provide decade-by-decade box plots of the distributions of temperature and precipitation in the CESM-LENS data.

the median for the 2090s exceeds the 75th percentile of the 1979-2015 distribution. Moreover, many temperatures above the 75th percentile of the 2090s data are in the upper tail of the distribution for 1979 to 2015. As shown in Figure 2(b), which plots the equivalent distributions for total precipitation across district-years, 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.

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 responses of yields to changes in temperature and precipitation patterns for this baseline period. As described later in this section, we use these estimates to project the evolution of the distribution of agricultural production over time from 1920 to 2099, given the changing weather distribution predicted by the CESM-LENS data.

As we have noted above, a key challenge in projecting yields for future years 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; this is especially true for temperature. We therefore attempt several different regression specifications that handle the problem of estimating the responses of yields to high temperatures in different ways. We also use two alternative models of the relationship between precipitation and yield.

3.1.1 Temperature bins

We begin with the following baseline specification, estimating it separately for each crop:

$$\ln A_{isgt} = \sum_{k=1}^{K} \beta_g^k T_{isgt}^k + \phi_g P_{isgt} + \psi_g P_{isgt}^2 + \alpha_{isg} + \gamma_{sg} t + \delta_{sg} t^2 + \epsilon_{isgt} \tag{1}$$

On the left-hand side, the logarithm of yield (quantity produced per unit area) for district i (in state s), crop g and year t is represented by $\ln A_{isgt}$. As regressors, we include a set of temperature and precipitation variables similar to those in Schlenker and Roberts (2009).

We divide daily average temperatures into a set of K bins, to be described in more detail below. The variable T^k_{isgt} represents the share of days in the crop's growing season falling into temperature bin k in a given district and year. Our specification also includes a quadratic function of total monsoon precipitation P_{isgt}, on which all crops (whether grown during or after the monsoon season) may depend. This variable aggregates rain across the five months of the monsoon season. 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.[24] For sugarcane, given its year-round growing season, we include only one quadratic function of precipitation, and instead define P_{isgt} as total precipitation for the entire year.

In all regressions, 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-years in which area and/or quantity values are missing, or either zero area planted or zero quantity is reported. We also winsorize observations that lie more than one standard deviation above the 99th percentile or more than one standard deviation below the 1st percentile of the distribution of log yield for a given crop.

In our baseline version of this specification, we define the temperature bins using the deciles of the 1979-2015 temperature distribution in each growing season. That is, pooling all district-days over the whole period, we calculate deciles of the distribution of average daily temperature observed in the ERA5 data, and define the upper and lower boundaries of our bins 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 the limits of each temperature bin differ for crops grown in different seasons.[25]

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 also mechanically depress variance in projected yields for each district-crop as temperatures rise and are spread across fewer bins, especially as a much larger share of temperature realizations enter the highest temperature bin, where the relationship between temperature and log yield is flat by assumption.

We therefore also check the robustness of our results to two alternative specifications that define the temperature bins differently, allowing for more convexity in the relationship between temperature and yield at the highest temperatures. The first variant instead uses

[24]For *rabi* crops, we use only the four months prior to this growing season (June to September) to calculate the monsoon precipitation variable.

[25]The temperatures covered by each bin for each growing season are listed in Appendix Table A.2.

the bins defined by Schlenker and Roberts (2009); i.e. a set of three-degree bins from 0 to 39 degrees Celsius, along with two additional bins covering all lower and higher temperatures respectively. A second version augments our initial decile bins specification by adding a 'high-degree days' (HDD) variable, calculated from the size of the gap between very high observed daily temperatures and the lower limit of the top decile bin. Further details of these specifications may be found in Appendix A.1.

3.1.2 Degree days

In our second approach to modelling the temperature-yield relationship, we instead remove the temperature bins and replace these with a 'degree-days' (DD) variable. In this case, we calculate the difference between the observed temperature and a baseline value of 24 degrees Celsius for each district-day during a crop's growing season, setting this equal to zero in cases where temperature is below 24 degrees. We then add these values together for each district-year.[26] In our baseline version of this specification, the DD variable enters linearly as follows:

$$\ln A_{isgt} = \beta DD_{isgt} + \phi_g P_{isgt} + \psi_g P_{isgt}^2 + \alpha_{isg} + \gamma_{sg} t + \delta_{sg} t^2 + \epsilon_{isgt} \tag{2}$$

The degree-days approach estimates the convexity in the temperature-yield relationship at high temperatures in part by exploiting observed convexity at relatively lower temperatures (above 24 degrees) that occur more frequently in the ERA5 data. This regression thus relies on a somewhat different source of identifying variation than the bins specifications discussed above. At the same time, it has the disadvantage that it constitutes a more restrictive parameterization, with only a single coefficient governing the effect of temperature on yield. In a robustness check, we thus instead use a quadratic function of the DD variable in our regression, as explained further in Appendix A.1.

3.1.3 Additional precipitation variables

We also specify the precipitation-yield relationship in an alternative way in an additional set of regressions. In particular, 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 a 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 relates to the

[26] The resulting specification is thus closer to that of Schlenker et al. (2006) rather than Schlenker and Roberts (2009).

growing season of the relevant crop; for example, they are defined using days between June and October for *kharif* (monsoon) crops. Because heavy rain is rare during the *rabi* season, we drop this variable for the five *rabi* crops.

3.2 Yield-weather regression results

We now briefly summarize the results of the regressions outlined above.[27] Figure 3 displays estimates of the coefficients β_g^k on the decile bins in equation (1), 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.[28] As found in the previous literature, a greater share of days at high temperatures tends to be relatively worse for crops as compared to more frequent exposure to temperatures in this lower range. We also show these results in a table format, along with the estimated coefficients on the precipitation variables, in Appendix Table A.3. For almost all crops, the relationship between yield and total monsoon precipitation displays the expected concavity, as the estimated coefficients tend to be positive for the linear term but negative for the quadratic term.[29]

For the baseline specification with degree days instead of bins (equation (2)), we find the estimated coefficient on the DD variable to be negative and statistically significant at the 1% level for all sixteen crops. Moreover, we again observe a concave relationship between yield and monsoon precipitation for almost all crops. These results are displayed in Appendix Table A.6.[30]

Finally, we re-estimate each of the above specifications, adding the additional precipitation variables from Fishman (2016). In both sets of regressions (see Appendix Tables A.8 and A.9), 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 become worse when rain occurs more frequently during their growing season. Some *rabi* crops also 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, with positive, negative and statistically insignificant coefficients all observed.[31]

[27] Note that we cluster standard errors at the level of the district in all cases.

[28] See Appendix Table A.2 for the temperatures covered by each bin.

[29] See also Appendix Tables A.4 and A.5 for the coefficient estimates from our other two yield-weather regressions using temperature bins.

[30] Estimates for the quadratic degree-days specification may be found in Appendix Table A.7.

[31] In Appendix Tables A.10 to A.12, we display the results of including the additional precipitation variables in our three other specifications.

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 as recorded in the VDSA data 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.

For example, in the case of our baseline temperature bins regression, we define:

$$\ln \tilde{A}_{isgrt} = \ln A_{isg,2000} + \sum_{k=1}^{10} \hat{\beta}_g^k \Delta T_{isgrt}^k + \hat{\phi}_g \Delta P_{isgrt} + \hat{\psi}_g \Delta P_{isgrt}^2 \tag{3}$$

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 a full set of projections for each of our yield-weather regression specifications.

The fact that we hold factors other than our weather variables fixed at their 2000 levels has two key implications for the interpretation of our results. First, our study does not take account of the possibility of technological change in response to evolving climatic conditions.[32] Instead, we set aside this dimension of adaptation in order to focus more closely on the direct impact of changes in the distribution of temperature and precipitation on the distribution of agricultural outcomes. Note that we will allow for another possible dimension of adaptation, crop choice, using a general equilibrium model in Section 5.

Also, in this method we have eliminated all variation in yields across run-years for a given district, other than the variation due to differences in weather conditions. This means that when we discuss the projected future evolution of the distribution of agricultural outcomes, we will only be considering changes in weather-induced variability rather than changes in the total variation of these outcomes.[33]

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

[32]Holding technology fixed is a standard approach in this literature (e.g. Deschênes and Greenstone, 2007; Schlenker and Roberts, 2009; Hsiang, 2016).

[33]We take this approach because we have far more information on potential weather variation, thanks to the multi-run CESM-LENS dataset, than we do on other dimensions of yield variation.

the precipitation variables for a given run at their 2000 levels, according to the CESM-LENS data for that run. In these projections, the 'fixed' variables thus continue to vary across runs, as well as by district and crop (based on the growing season), but not over time.

4 Distribution of agricultural revenue

In this section, we use our yield projections to investigate the potential implications of climate change for agricultural outcomes in each Indian district. So as 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.[34] 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). Below, we discuss our projections of changes across decades in the first and second moments of the distribution of agricultural revenue, and then consider a measure of downside risk.

4.1 Mean revenue

We first take the simple average of the 400 revenue realizations for each district-decade, giving us a projection for mean agricultural revenue. To calculate the weather-induced changes in mean revenue by decade for the average district, we regress the logarithm of this variable on district fixed effects and decade fixed effects. In Figure 4(a), we display the decade fixed effects from this regression, along with 95% confidence intervals.

The left side of the figure shows results of projections using our baseline yield-weather specification, with ten bins based on deciles of the temperature distribution as well as a quadratic in total precipitation (equation (1)). We see a substantial projected decline in mean revenue due to climate change over the course of the 21st century, in contrast to the relative stability in this measure from 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%. This result remains very similar when we instead use either of our other two specifications incorporating temperature bins, as we show in Appendix Figures A.4 and A.5.[35]

[34] 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.

[35] Based on the specification using the temperature bins of Schlenker and Roberts (2009), mean revenue is projected to fall by 31.2% in the average district. This estimate rises only slightly in magnitude to 33.1% when we instead return to the decile bins but add an HDD variable based on the top 10% of daily temperatures.

We next consider the projections based on our specification modelling the temperature-yield relationship as a linear function of degree days (equation (2)). The right side of Figure 4(a) shows that the predicted evolution of mean agricultural revenue for the average district is similar to the projections using temperature bins, with an estimated 23.9% decline over the course of the 21st century.[36] As may be seen in Appendix Figure A.7, adding a quadratic degree-days term makes little difference to these results.[37]

Figure 4(b) displays analogous estimates for projections in which only temperature or precipitation are allowed to change from their values as of 2000, as discussed in the previous section. These show that the entire projected effect of climate change on mean agricultural revenue is driven by predicted changes in temperature. 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 (0.4% using the temperature bins specification, and 0.9% in the case of degree days).[38]

Incorporating the additional precipitation variables from Fishman (2016) into these specifications does not substantially change our projections of the impact of evolving precipitation patterns on mean agricultural revenue. The predicted effect of precipitation alone remains small, with a 0.4% decline (using the temperature bins model) or a 0.1% rise (using the degree-days model) in mean revenue now forecast for the 2090s relative to the 2000s for the average district. However, because the augmented specification tends to reduce the magnitudes of the estimated coefficients on the temperature variables, it suggests a somewhat smaller projected total fall in mean revenue due to climate change of 23.5% (temperature bins) or 15.5% (degree days). All of these results may be seen in Appendix Figures A.8 and A.9.[39]

Both of our baseline empirical models, along with our various robustness checks, thus suggest similar implications for mean agricultural revenue. On average across districts, we project a gradual but substantial decline in mean revenue over the course of the 21st century, in a range from 15.5% to 33.1%. Moreover, we find that this is almost entirely the result of predicted changes in temperature.

[36]The projected heterogeneity in this effect across districts varies somewhat between these two specifications, due to differences in the estimated yield-weather elasticities by crop. We depict district-level results for each of the outcomes studied in this section in the maps in Appendix Figure A.6.

[37]Specifically, this modestly increases the size of the projected decrease in mean revenue to 27.9%.

[38]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.

[39]See also Appendix Figures A.10 to A.12 for results when the additional precipitation variables are included in our three alternative temperature specifications. In these three cases, the projected impact of 21st-century climate change on the average district varies from 18.9% to 22.9%.

4.2 Standard deviation of revenue

We next turn to a discussion of the second moment of the agricultural revenue distribution. Specifically, for each decade of CESM-LENS data, we consider the standard deviation of our 400 revenue realizations within each district. As discussed earlier, our ability to consider projected variation in outcomes across so many potential weather realizations is a major advantage of our use of the CESM-LENS dataset.

When interpreting the results in this subsection, it is important to recall that the variation between projected yields across runs is only due to weather, because we hold other determinants of yield (conditional on weather) fixed across realizations. This means that we will interpret projected rises or declines in the standard deviation of revenue as proportional changes in weather-induced variation, rather than putting these in terms of total within-district variation.[40] A key question we explore here is whether or not our various scenarios predict a substantial rise in this weather-induced revenue variability.

Using our baseline temperature bins specification, we instead project that the average district will experience a 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 (according to our projections), 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). Again, this appears to be driven by the effects of temperature changes, as our temperature-only projection yields a similarly large fall in the standard deviation of revenue (65.4%), while under the precipitation-only projection we see a small rise of 3.4% (see the two panels on the left side of Figure 5(b)).

We noted in Section 3.1 that this specification might mechanically decrease projected variation in yield over time, as temperatures rise and most observations are thus spread across fewer bins. We show evidence that this is indeed the case in Appendix Figure A.13, which plots the India-wide shift in the share of days (year-round) in each of our decile bins, 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.[41]

As discussed in Section 2.2, the CESM-LENS data does not project large changes in the standard deviation of temperature for the average Indian district over the coming decades. Instead, our temperature bins specifications predict temperature-yield relationships to become relatively flatter in the temperature range we expect to observe in the late 21st century. Under

[40] As noted in footnote 29 in Section 3.3, we do this because we have far fewer available realizations with which to model residual yield variation.

[41] When we use the fifteen bins of Schlenker and Roberts (2009), the effect is dampened by only one-third, with a decline of 38.4% in weather-induced variability in the average district from the 2000s to the 2090s. Adding an HDD variable to the baseline regression has an even smaller impact on our projections, producing a predicted fall in the standard deviation of revenue of 56.1% on average. See Appendix Figures A.4 and A.5.

such a scenario, variance in agricultural yield due to temperature should be expected to decline rather than rise in future.

In contrast, the parametric assumption underlying our specification linear in degree days 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 this specification yields an upward-sloping projection for 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.

Moreover, the positive sign of the predicted effect of climate change in this case is not driven by temperature, as the temperature-only projection generates a small decline in the standard deviation of revenue by the 2090s (top right panel of Figure 5(b)). Instead, this scenario projects a sharper rise of 9.3% in the standard deviation of revenue due to precipitation, as displayed in the bottom right panel of Figure 5(b).[42] We do observe a small positive impact of temperature changes (of 9.8%) when the curvature of the temperature-yield relationship is increased by adding a quadratic DD term to the regression.[43]

The wide range of estimates implied by our various specifications, from a decline of 59.9% to a rise of 18.8% 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, it is notable that we do not find strong evidence that weather-related output volatility is likely to increase substantially for Indian farmers. In the next subsection, we further explore this observation.

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 in 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

[42]The addition of the precipitation variables from Fishman (2016) to our regressions results in a modest reduction in the projected effects for the precipitation-only scenario, to 1.3% for the temperature bins specification (Appendix Figure A.8) and to 6.4% in the model based on degree days (Appendix Figure A.9).

[43]This becomes larger (16.8%) when the additional precipitation variables are added, leading to a total rise of 18.8% when we account for the full effect of climate change; see Appendix Figures A.7 and A.12.

1960s) as a baseline, defining a 1-in-100 bad year for weather in each district based on the projected agricultural outcomes of the 2000 run-years we observe during this period (forty per year for fifty years). Specifically, after calculating agricultural revenue in a given district for each of these realizations, we set the threshold for a 1-in-100 year using the first percentile of this revenue distribution. For each decade up to the 2090s, we then calculate the share of the district's revenue realizations (out of the 400 run-years in each decade) that fall below this level. We perform these calculations separately for each of the 310 districts in our sample.

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 generated from our baseline yield-weather specification with temperature bins. We find that in the average district, bad years remain as rare as in the baseline period as the climate remains relatively unchanged through the 20th century. However, the number of realizations that would have been classified as 'extreme weather' in the first fifty years of the CESM data is projected to become progressively larger throughout 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.

This very large projected increase in the occurrence of bad years is driven by changes in the pattern of temperature rather than precipitation. In the two panels on the left of Figure 6(b), we show the results of making similar calculations under our temperature-only and precipitation-only climate change scenarios respectively. The projected rise in the share of bad years due to temperature changes in the average district is nearly identical to the scenario pictured in the left panel of Figure 6(a). While changes in the pattern of precipitation are also predicted to lead to a higher probability of a 1-in-100 year at the end of the 21st century, the estimated effect is very small by comparison, and does not exceed the magnitude of oscillations in the share of bad years over the previous decades.

These predictions are not unique to our baseline temperature bins specification. The right side of Figure 6(a) displays analogous projections using the specification linear in degree days, which predict an 84 percentage point rise in the share of 1-in-100 weather years between the 2000s and the 2090s for the average district. Projected impacts under the temperature-only and precipitation-only scenarios are also similar to the previous case (see the right side of Figure 6(b)).

When we add the rainy days, heavy rain and dry spell variables to our regressions, we find a larger impact of changes in precipitation patterns, via a sharper rise in the predicted share of bad years from the mid-21st century onwards. As shown in the bottom right panels of Appendix Figures A.8 and A.9, the effect of precipitation changes is now projected to constitute a rise of 0.7 (using decile temperature bins) or 0.9 (using degree days) percentage

points in this share for the average district between the 2000s and 2090s. This represents a near-doubling of 1-in-100 bad years, a notable result but one that remains small compared to the impact of temperature.[44]

Overall, we find an effect of 21st-century climate change on the share of 1-in-100 years in the average district ranging from 53 to 88 percentage points.[45] 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: 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. Moreover, price adjustments might offset part of the effect arising from changes in temperature and precipitation patterns. To check whether our main results are sensitive to these issues, we carry out a structural analysis of the changes in Indian agricultural outcomes due to the evolving climate, allowing for endogenous prices and crop choices as well as trade between districts.[46]

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 a finite number of goods, with portfolio choice by agents across risky assets. In Appendix A.2, we provide a brief description of the setup of the model; further details may be found in their paper.

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. This assumption, in conjunction

[44]The total predicted effect of climate change is again dampened when we add the additional precipitation variables (due to the smaller estimated impact of temperature on yield), but is still very large. The projections combining temperature and precipitation changes suggest a rise of 54 (temperature bins) or 72 (degree days) percentage points in the share of bad years during the 21st century for the average district (see the top right panels of Appendix Figures A.8 and A.9).

[45]See Appendix Figures A.4, A.5, A.7, and A.10 to A.12 for projections using our other yield-weather specifications. The smallest estimated rise of 53 percentage points is based on the quadratic degree-days regression with additional precipitation variables.

[46]Note that our approach effectively 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.

with constant relative risk aversion, leads to a simple expression for the log expected utility of the (identical) farmers within a district i:

$$\ln E(U_i) = \left(\mu_i^z + \frac{1}{2}\sigma_i^{2,z}\right) - \frac{1}{2}\rho\sigma_i^{2,z} \tag{4}$$

The two terms in the expression above 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 $\sum_g \theta_{ig} = 1$, leading to first-order conditions for each district and good in terms of district-specific Lagrange multipliers λ_i as follows:

$$\mu_{ig}^z - \rho \sum_h \theta_{ih}\Sigma_{igh}^z = \lambda_i \tag{5}$$

The first term in this expression 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. 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-weather regressions) in each decade. Several parameters are needed for these calculations; we provide details on how we calibrate or estimate each of these parameters in Appendix A.2.

5.2 Results

In order to assess the importance of allowing for endogenous prices, planting decisions and trade patterns, we begin by comparing the welfare impacts implied by the model to a naive benchmark using the nominal revenue measure from the previous section. Specifically, we calculate our measure of 'naive welfare' by using the model-derived expression for the change

in expected utility over time, but substituting nominal revenue whenever this formula calls for real returns.

Figure 7(a) displays the result of this comparison for the baseline temperature bins specification. The naive approach using nominal revenue predicts a gradual but large fall in welfare on average across districts between the 2000s and 2090s due to climate change, with a total projected decline over this period of 34.5% (left panel). These losses are approximately six percentage points smaller (28.6%) when we use the model-based counterfactual results, with a similarly shaped 21st-century trend (right panel). Results using the degree-days specification also predict a continuous decline in welfare from the 2000s to the 2090s, with a cumulative drop of 21.8% (see Figure 7(b), right panel). Again, the naive and actual welfare measures yield similar trends, with the full 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 the first and second moments to changes in welfare (or 'naive welfare'), based on the two terms found in equation (4). 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 does not have a major impact on this result, modestly increasing predicted losses to 34.6%. Moreover, the trend in estimated welfare due to the mean is also similar to our simple exercise in Section 4.1, where we projected a 32.3% fall in mean agricultural revenue (Figure 4(a)). Analogous conclusions may be reached with the specification using degree days, again by comparing Figure 8(a) to Figure 4(a).

As discussed above, the role of yield variability in farmers' welfare within the model is captured by the variance of log real returns, a simple summary measure that relies on the assumption of lognormality of yields across states. This measure is different from the elements of the distribution of agricultural outcomes discussed 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 small. As shown in Figure 8(b), on average the model predicts a 21st-century rise in welfare of 0.25% in the average district when we use temperature bins, and a welfare decrease of 0.1% in the case of degree days. Again, neither trend is substantially affected by the naive use of nominal revenue instead of real returns in the welfare formula.[47]

We draw two main conclusions from our exercises using the Allen-Atkin model. First,

[47]Results using projections based on our other yield-weather specifications may be seen in Appendix Figures A.14 to A.21. The main results are unchanged: naive and actual 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 13.5% to 29.0% across these eight alternative scenarios.

based on our comparisons of the model's welfare measures to similar naive benchmarks, 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 bring increases in insurable risks for farmers. In such a case, increased availability of insurance would be an important policy prescription in response to the changing distribution of returns in the agricultural sector. However, across the several different ways in which we have modelled projected agricultural outcomes, one robust finding has been a sharp predicted decline in the first moment of the relevant distribution. This implies potentially substantial uninsurable losses from the changing climate, whether or not insurance markets are well-developed.

We have also considered several relevant dimensions of the distribution of agricultural outcomes other than the mean. These have included the standard deviation of agricultural revenue, the frequency of '1-in-100' revenue years, and the risk-driven component of welfare implied by a model of portfolio choice. None of these exercises has provided us with strong evidence of a large projected rise in variability. Instead, we have forecast that upward shifts in the temperature distribution will 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 variation in agricultural outcomes are very sensitive to modelling choices. Third, it is possible that our conclusions could differ in countries that are at the technological frontier, such as the US and members of the EU. Finally, we have held technology constant in our exercises, while in practice, technological changes could be a crucial mode of adaptation.[48]

We nonetheless argue that our results are important to the discussion of the potential 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 with data from multiple runs, 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

[48] As noted in footnote 7 in the introduction, there are also other potential margins of adaptation that we have not modelled here.

hope that the challenges in the forecasting of the future distribution of agricultural outcomes, as presented in this paper, help to 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, these results suggest that adaptation to changes in average weather patterns, and especially the rising incidence of extreme temperatures, should be a priority. We leave to future research the question of which specific technologies might have the largest potential to address the daunting changes brought on by the evolving climate.

References

Adams, R. M., Rosenzweig, C., Peart, R. M., Ritchie, J. T., McCarl, B. A., Glyer, J. D., Curry, R. B., Jones, J. W., Boote, K. J., and Allen Jr, L. H. (1990). Global Climate Change and US Agriculture. *Nature*, 345(6272):219–224.

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. (2019). Climate Change and Agriculture: Subsistence Farmers' Response to Extreme Heat. *arXiv preprint arXiv:1902.09204*.

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.

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. and Emerick, K. (2016). Adaptation to Climate Change: Evidence from US Agriculture. *American Economic Journal: Economic Policy*, 8(3):106–140.

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.

Colmer, J. (2018). Weather, Labor Reallocation and Industrial Production: Evidence from India. *Mimeo*.

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.

Crost, B., Duquennois, C., Felter, J. H., and Rees, D. I. (2018). Climate Change, Agricultural Production and Civil Conflict: Evidence from the Philippines. *Journal of Environmental Economics and Management*, 88:379–395.

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.

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.

Fisher, A. C., Hanemann, W. M., Roberts, M. J., and Schlenker, W. (2012). The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather: Comment. *American Economic Review*, 102(7):3749–3760.

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.

Garg, T., Jagnani, M., and Taraz, V. (2018). Temperature and Human Capital in India. *Mimeo*.

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.

Jessoe, K., Manning, D. T., and Taylor, J. E. (2018). Climate Change and Labour Allocation in Rural Mexico: Evidence from Annual Fluctuations in Weather. *The Economic Journal*, 128(608):230–261.

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. (2017). The Climate Risk Premium: How Uncertainty Affects the Social Cost of Carbon. *University of Arizona Department of Economics Working Paper*, 15–01.

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.

Mendelsohn, R., Nordhaus, W. D., and Shaw, D. (1994). The Impact of Global Warming on Agriculture: A Ricardian Analysis. *American Economic Review*, 84(4):753–771.

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., Hanemann, W. M., and Fisher, A. C. (2005). Will US Agriculture Really Benefit from Global Warming? Accounting for Irrigation in the Hedonic Approach. *American Economic Review*, 95(1):395–406.

Schlenker, W., Hanemann, W. M., and Fisher, A. C. (2006). The Impact of Global Warming on US Agriculture: An Econometric Analysis of Optimal Growing Conditions. *Review of Economics and Statistics*, 88(1):113–125.

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.

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.

7 Figures and tables

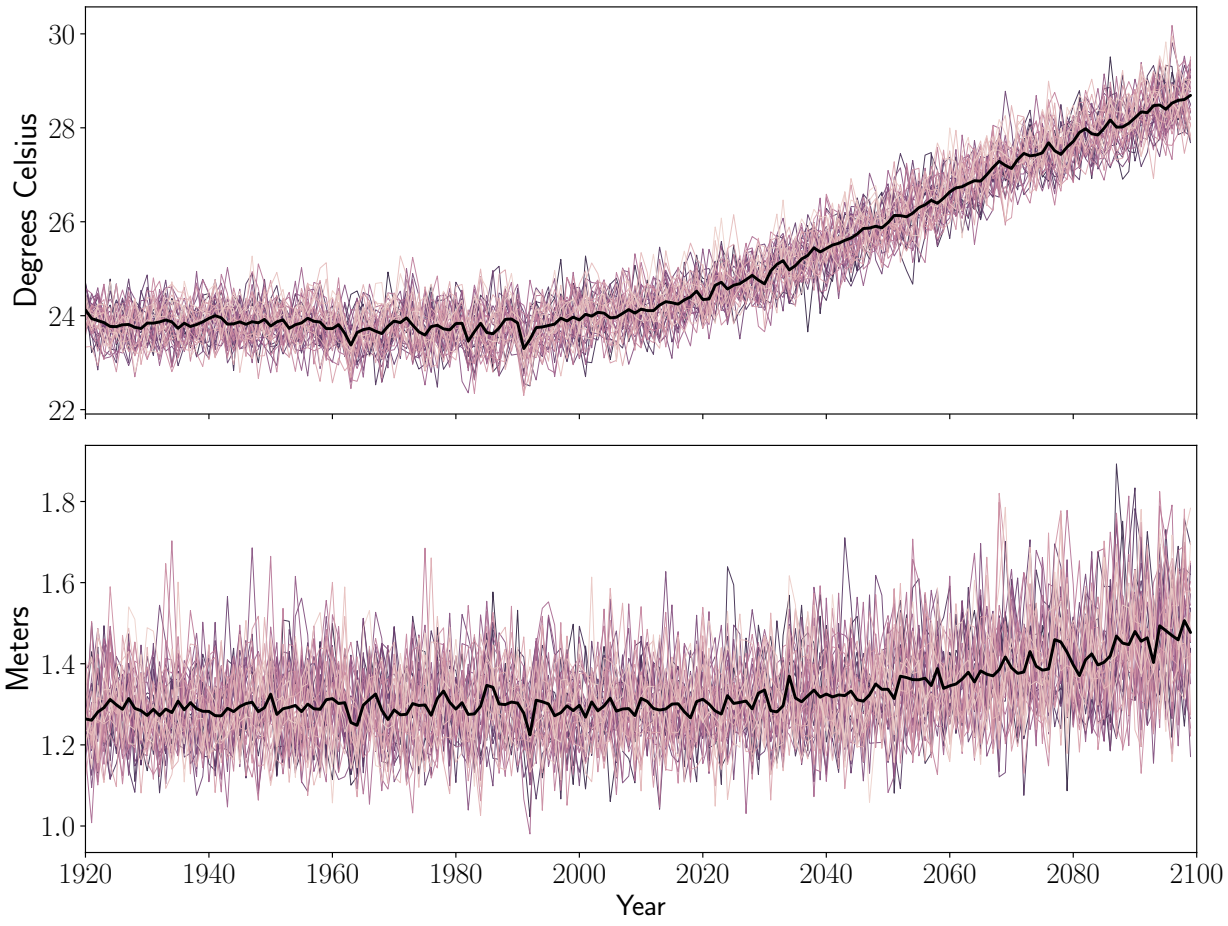

Figure 1: Distributions of mean daily temperature and total annual precipitation

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.

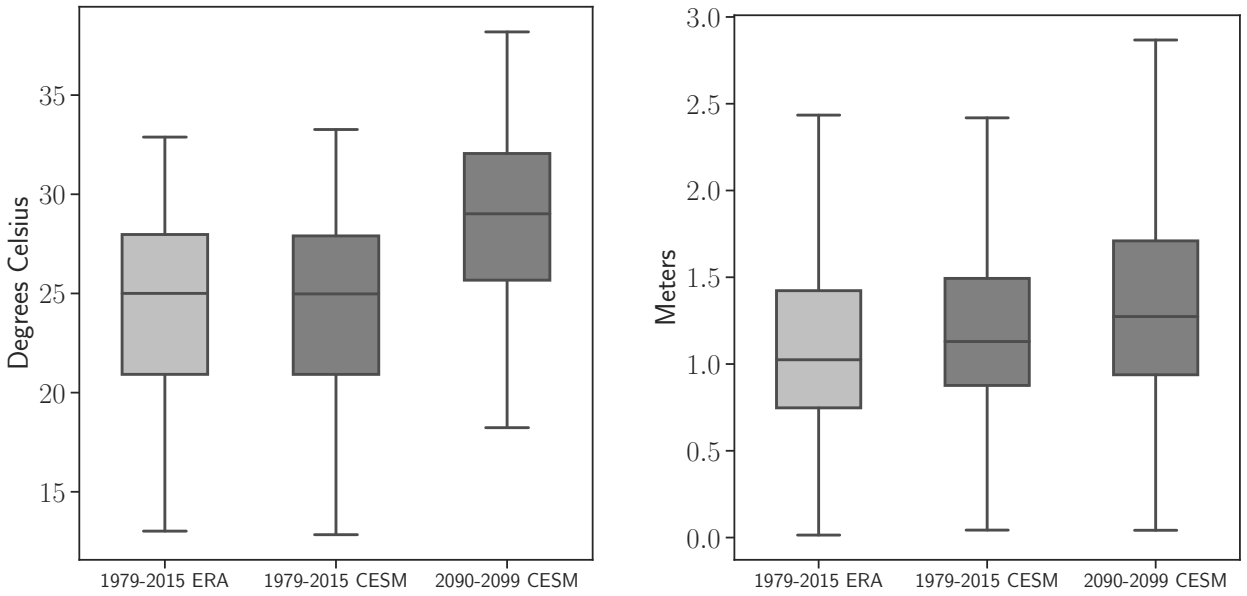

(a) Temperature

(b) Precipitation

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.

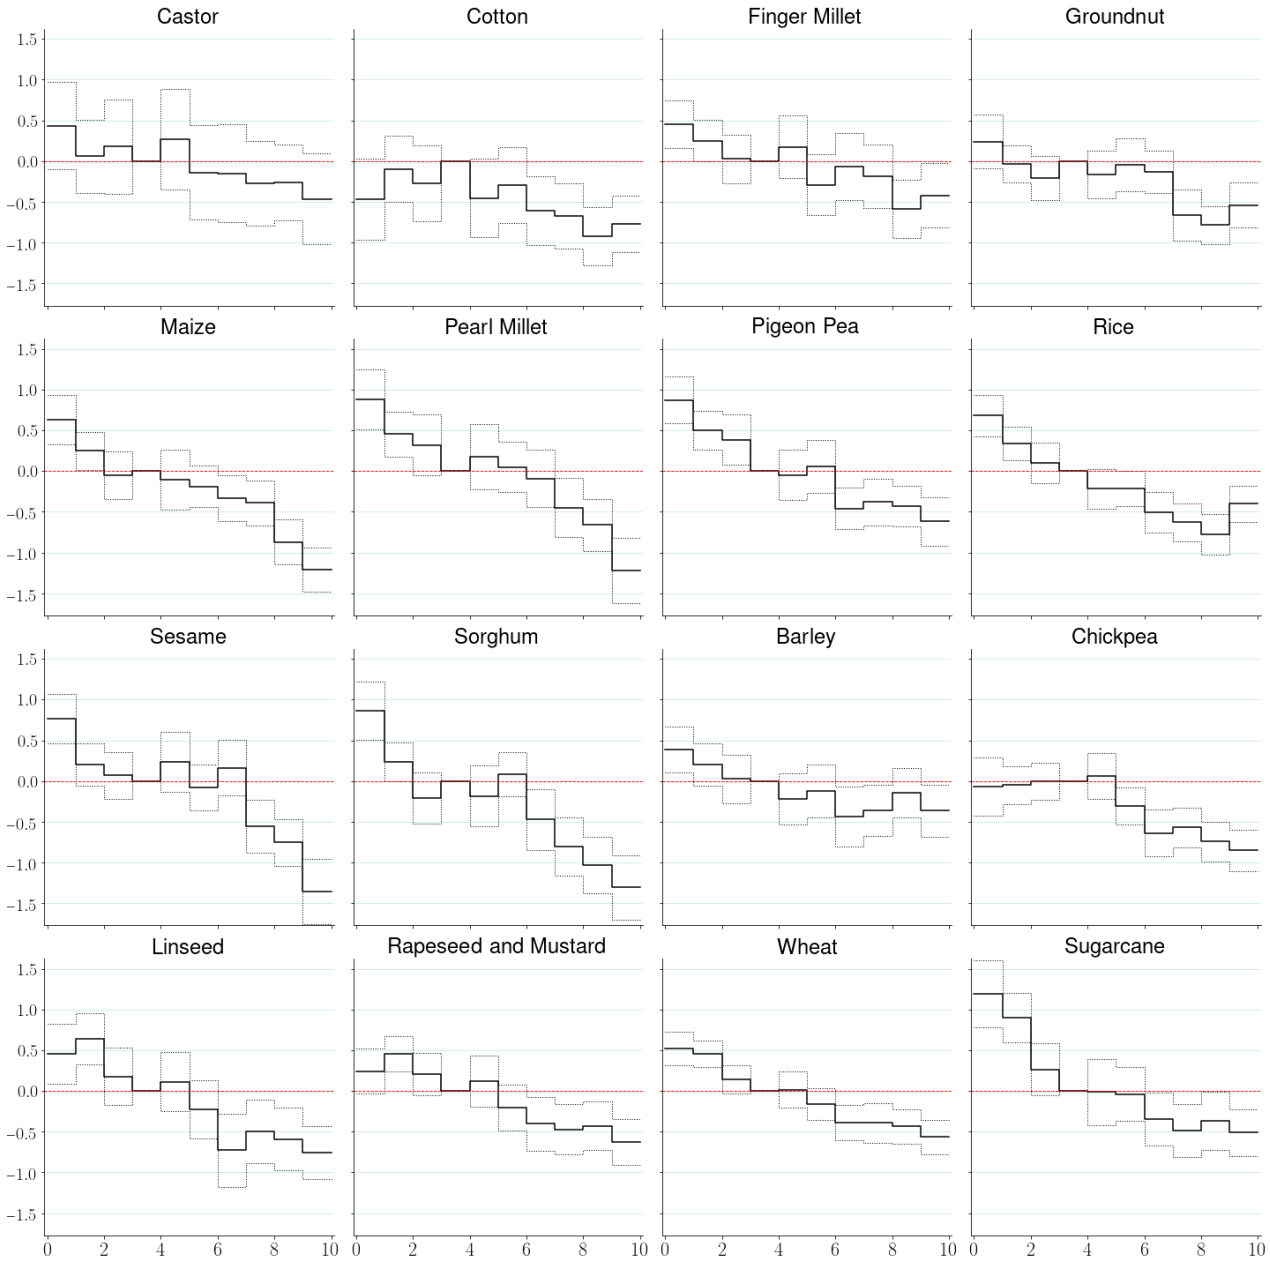

Figure 3: Estimates of coefficients on temperature bins using decile temperature bins

This figure shows the estimated coefficients on temperature bins using the decile temperature bins specification without additional precipitation variables. The solid lines represent the point estimates while the dotted lines correspond to 95% confidence intervals. 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 Appendix Table A.2 for the limits of each bin in degrees Celsius for each growing season.

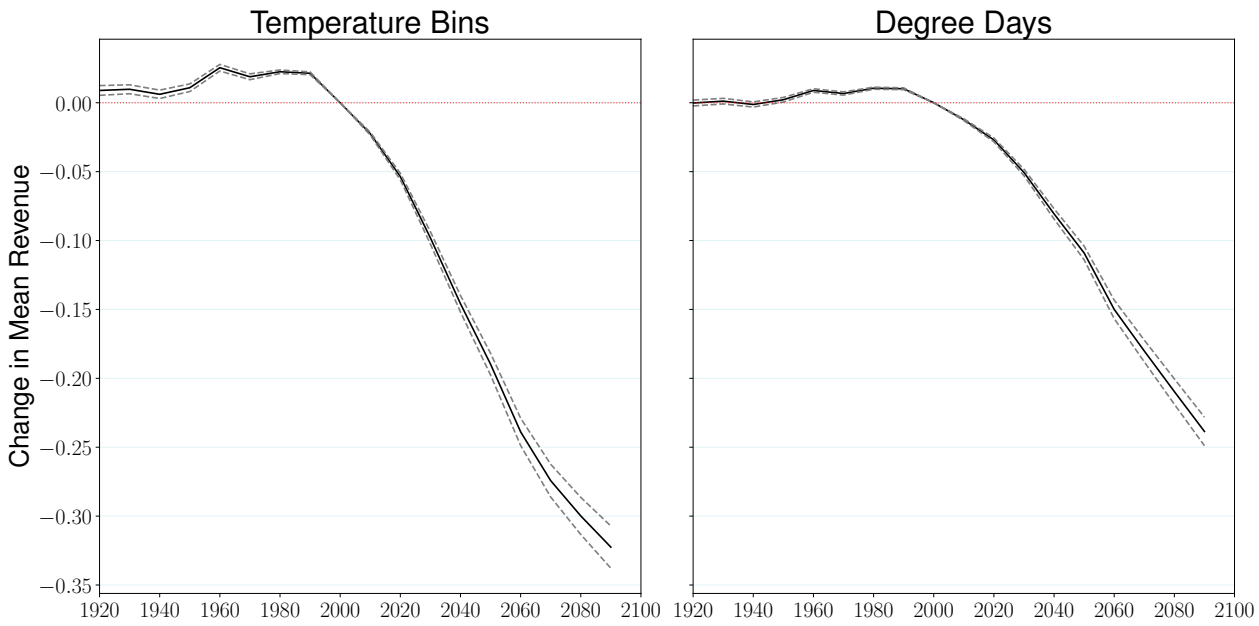

(a) Total

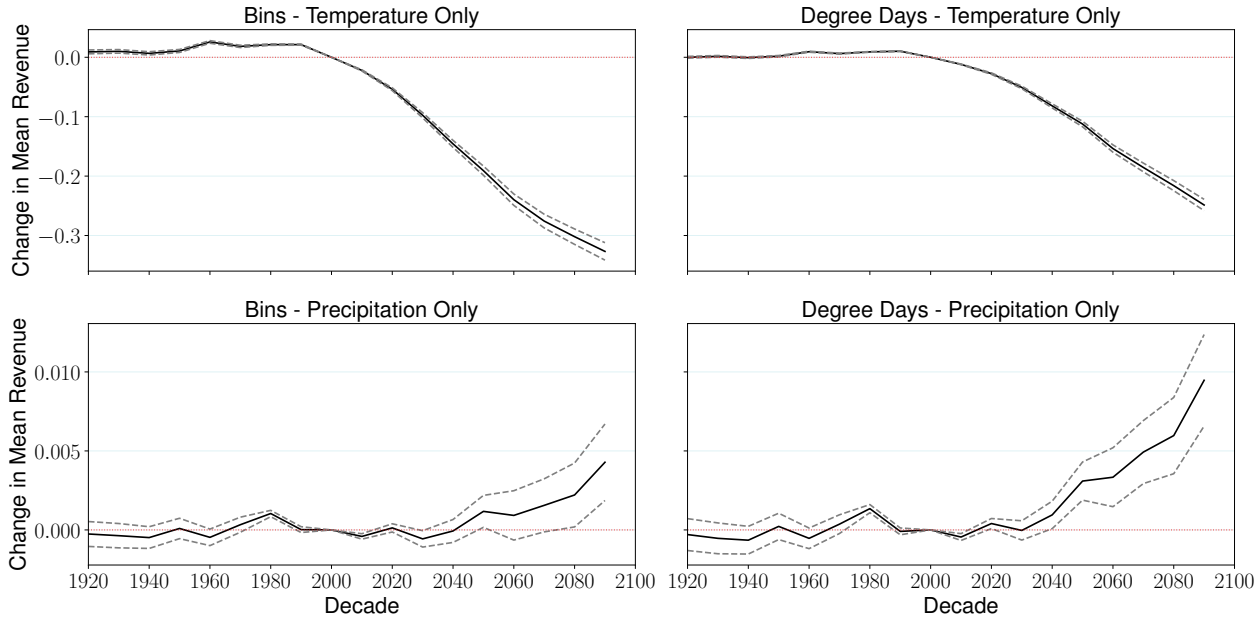

(b) Decomposition

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), without additional precipitation variables. Dashed lines display 95% confidence intervals, 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.

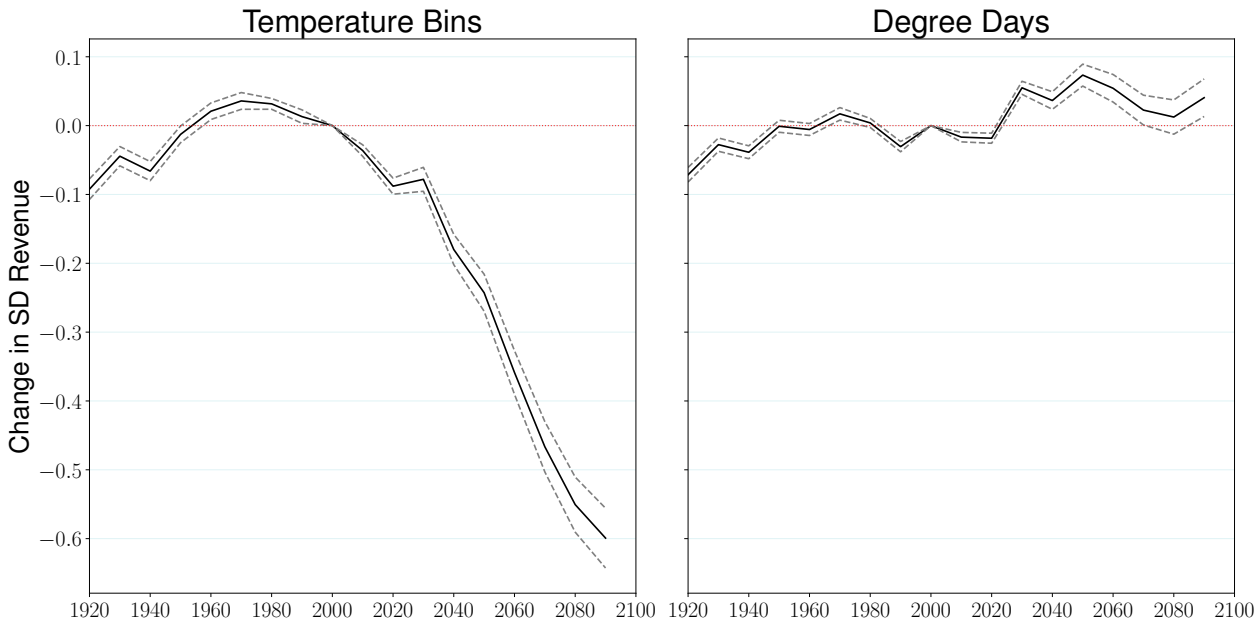

(a) Total

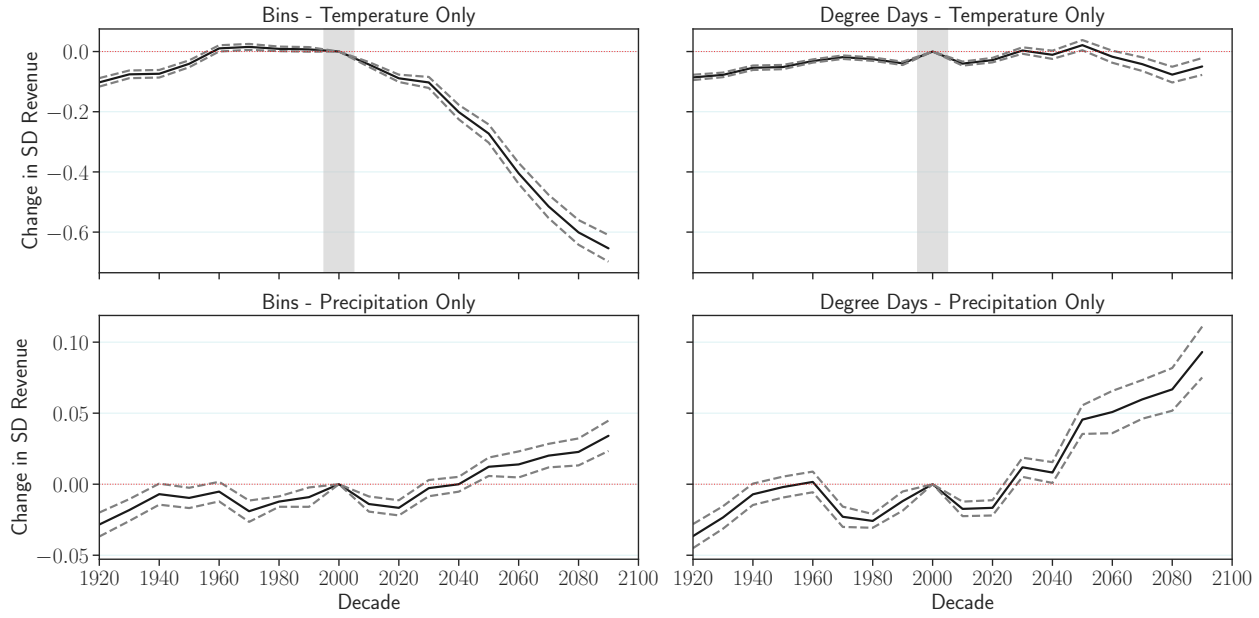

(b) Decomposition

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), without additional precipitation variables. Dashed lines display 95% confidence intervals, 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.

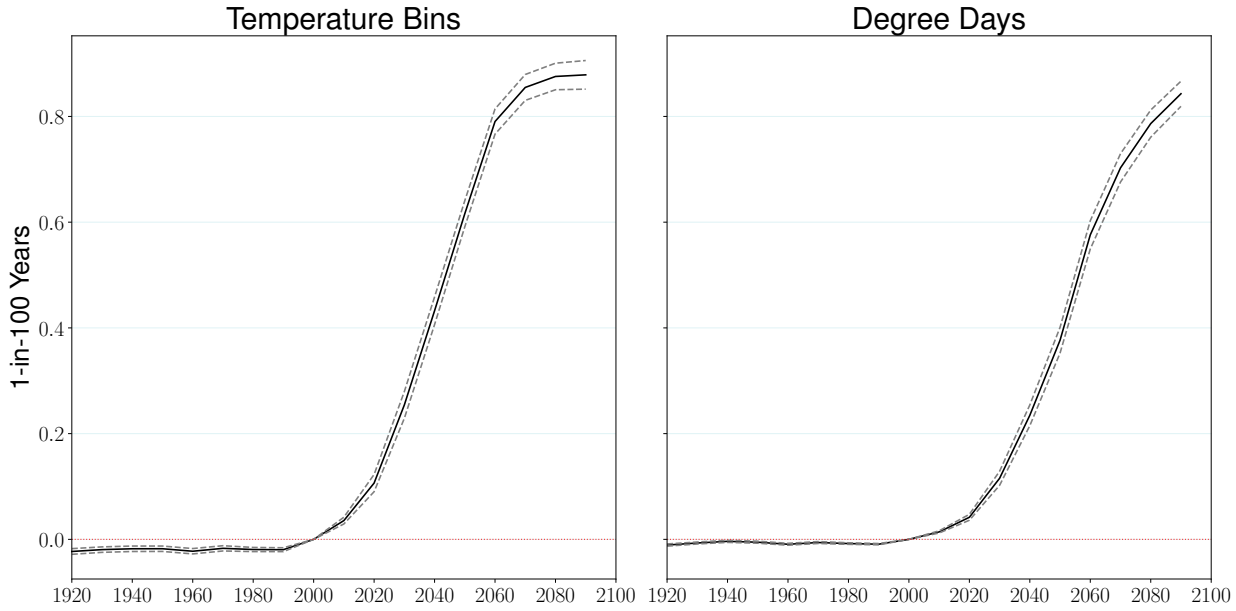

(a) Total

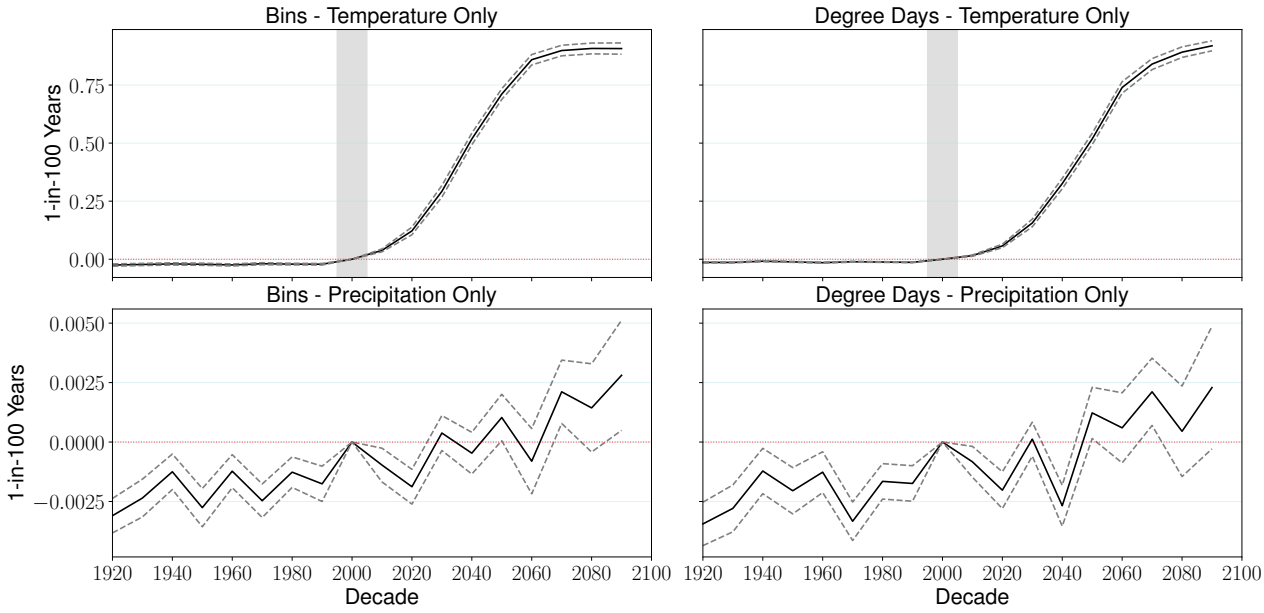

(b) Decomposition

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), without additional precipitation variables. Dashed lines display 95% confidence intervals, 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.

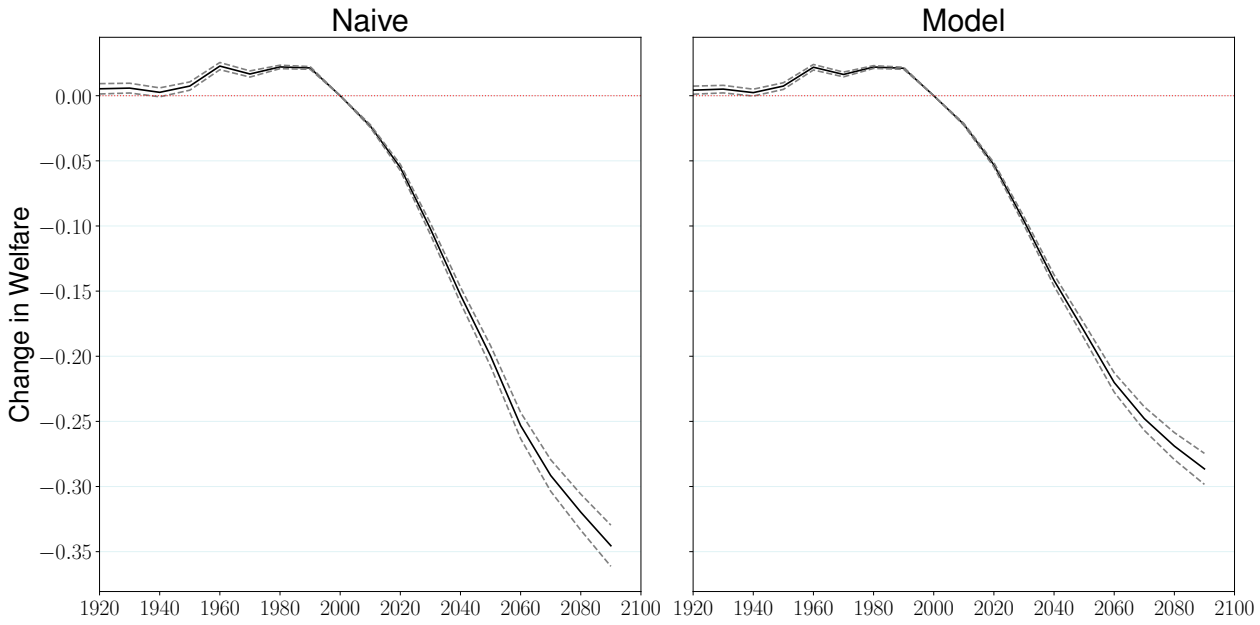

(a) Temperature Bins

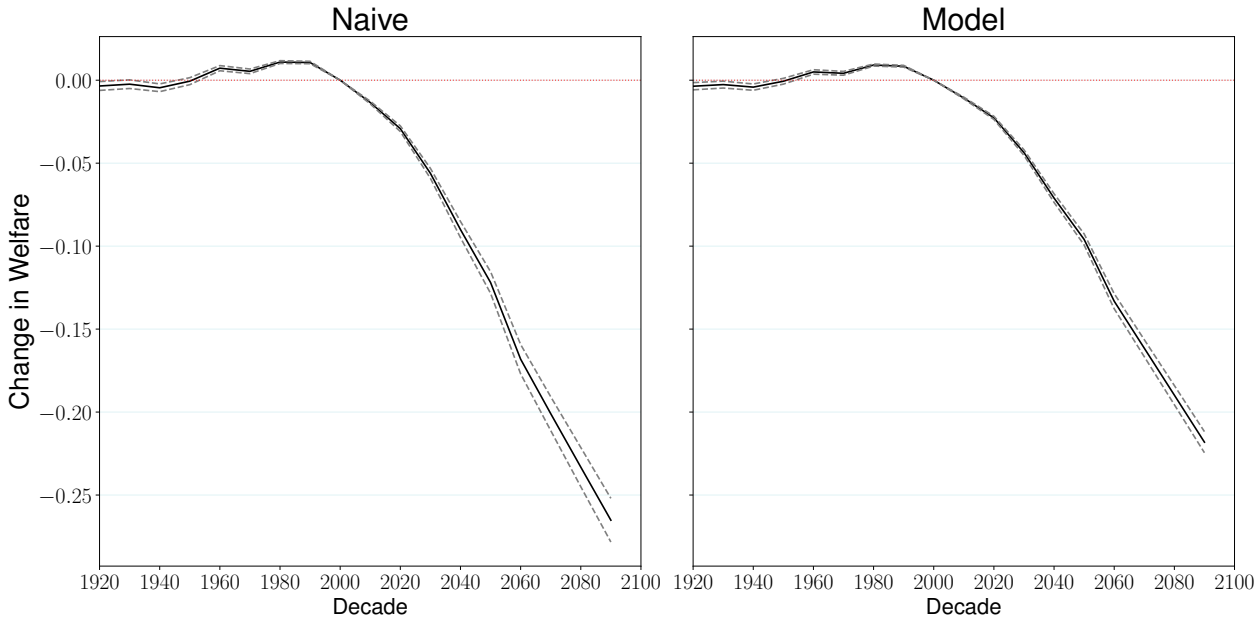

(b) Degree Days

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, both without additional precipitation variables. Dashed lines display 95% confidence intervals.

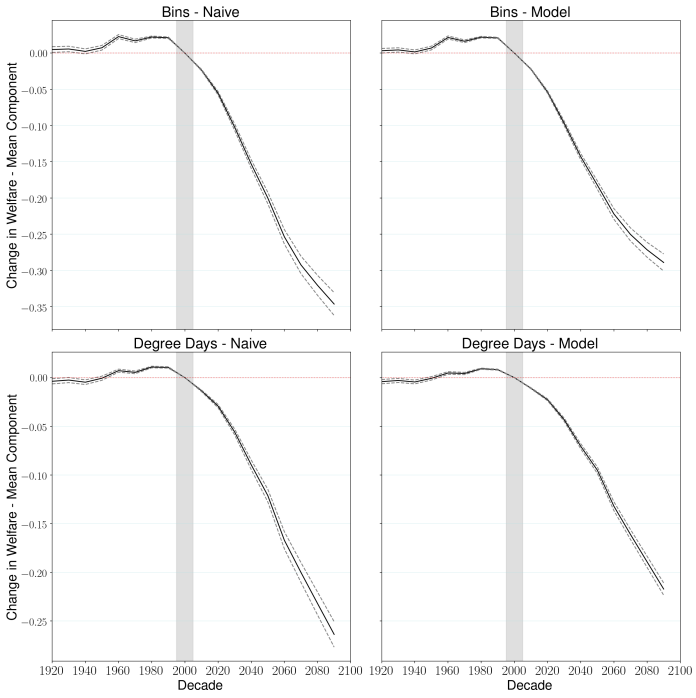

(a) Mean component

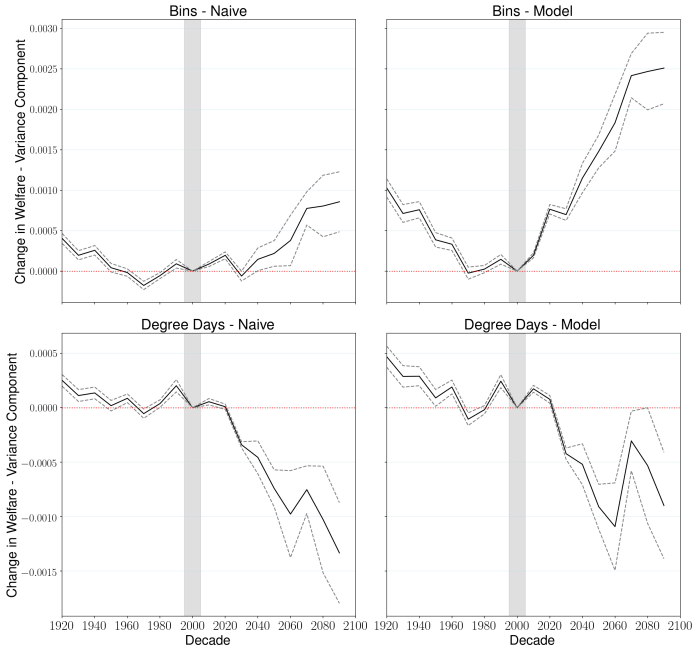

(b) Variance component

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, both without additional precipitation variables.

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 of data.

Appendix

A.1 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.

We also estimate a second variant of our degree-days specification. In this case, we calculate degree days using a 24-degree threshold as in the baseline version, but add both this variable and its square to our regression. Similarly to our baseline degree-days regression, this allows us to use the ample observed variation in yields at relatively lower temperatures to estimate the reactions of yields to higher temperatures, while also accommodating some curvature in this relationship.

A.2 Details of model and parameterization

A.2.1 Model setup

In this appendix, we discuss the basic setup of the Allen and Atkin (2016) model, and then explain how the model is parameterized in our empirical application. 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.

Marginal costs of production c_{ig} are specific to each district and crop, and are assumed to be invariant across states. Therefore, 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 p_{ig}:

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

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.

A.2.2 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 (5)):

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

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^z_{igd} gross of production costs c_{igd}. We can also calculate Σ^z_{ighd} 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 (5)). 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^A_{igd} 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^A_{ighd} equal to values from $\Sigma^A_{i'ghd}$, 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.

A.3 Additional figures and tables

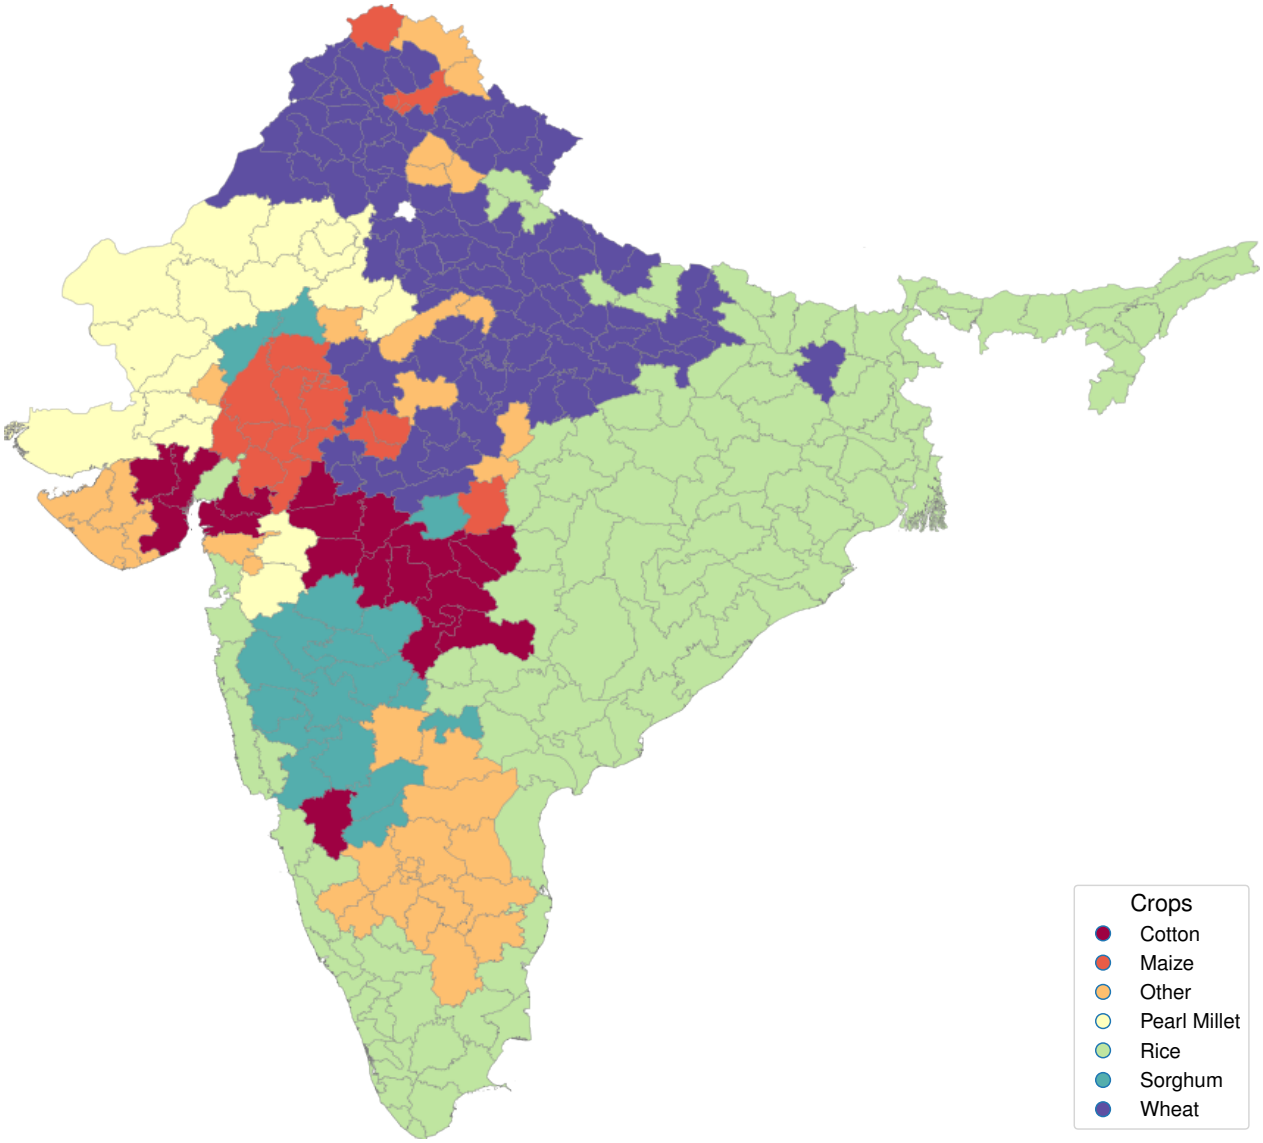

Figure A.1: Most-planted crop (by land area) by district in 2000

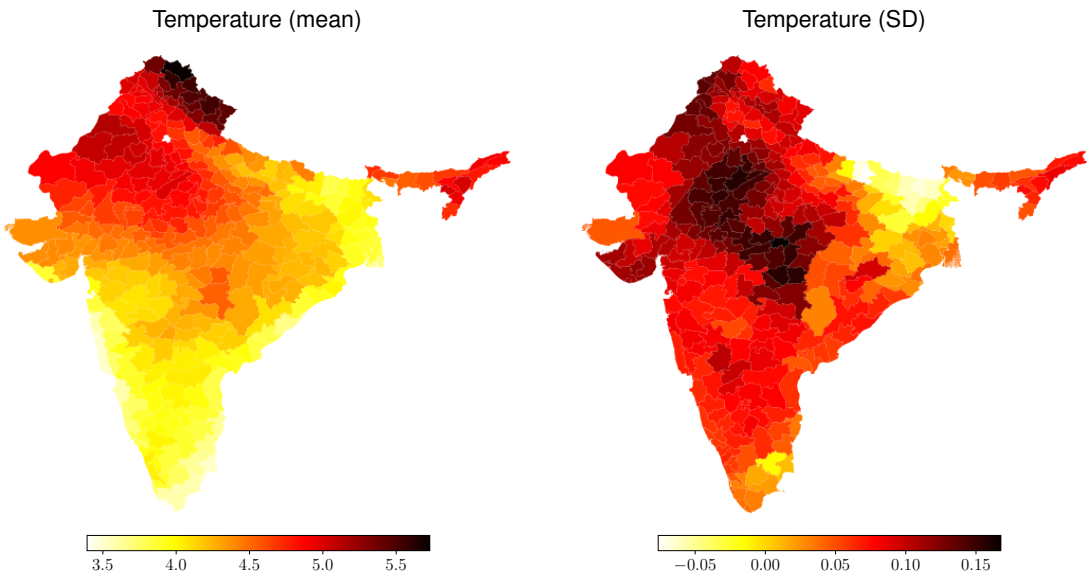

(a) Temperature

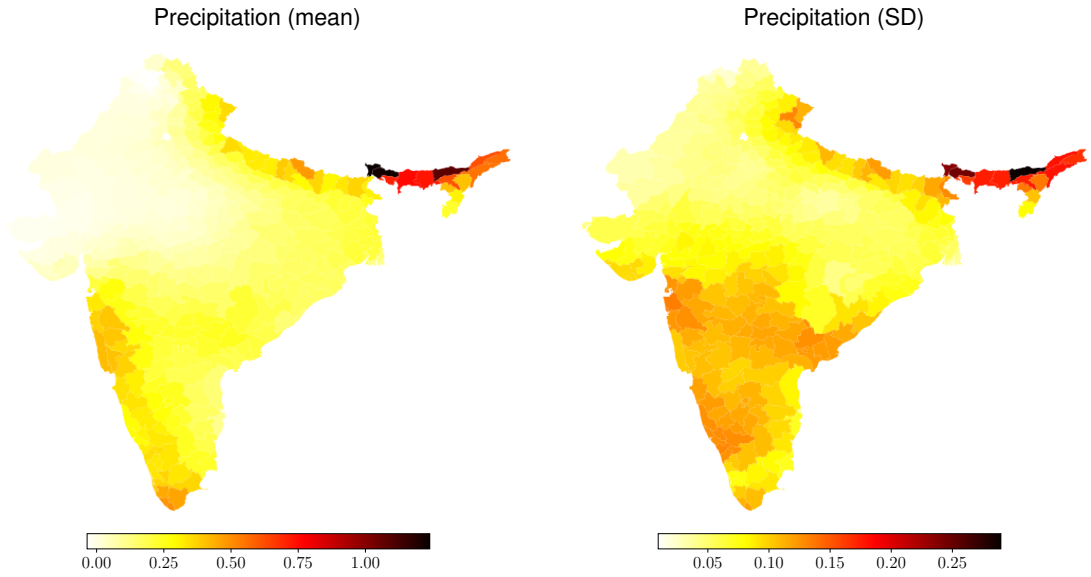

(b) Precipitation

Figure A.2: 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

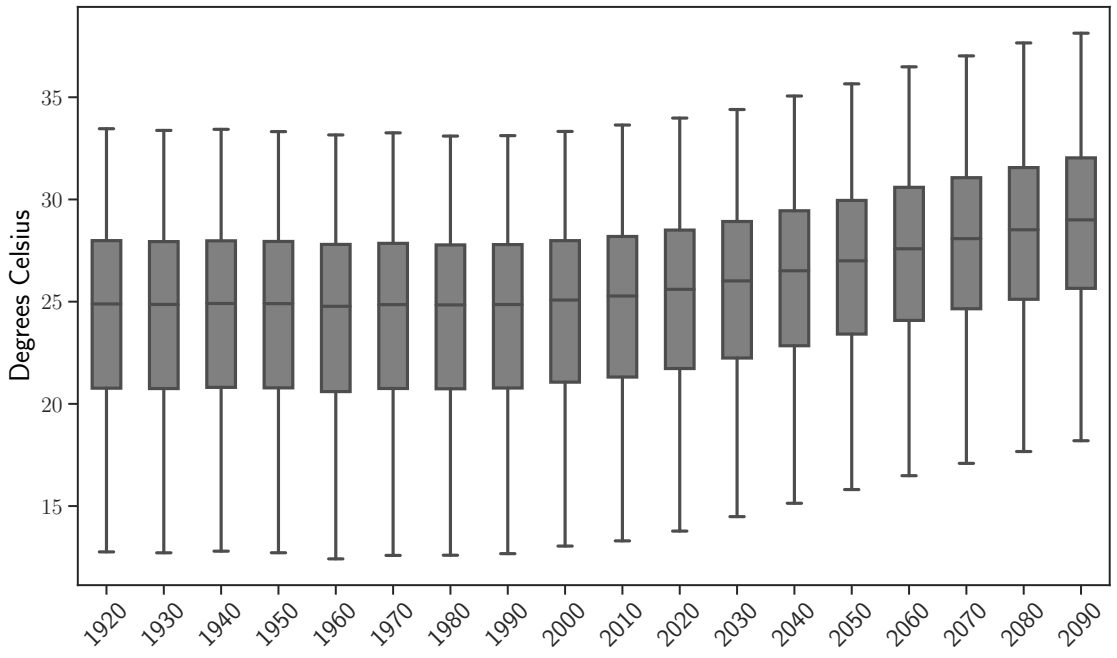

(a) Temperature

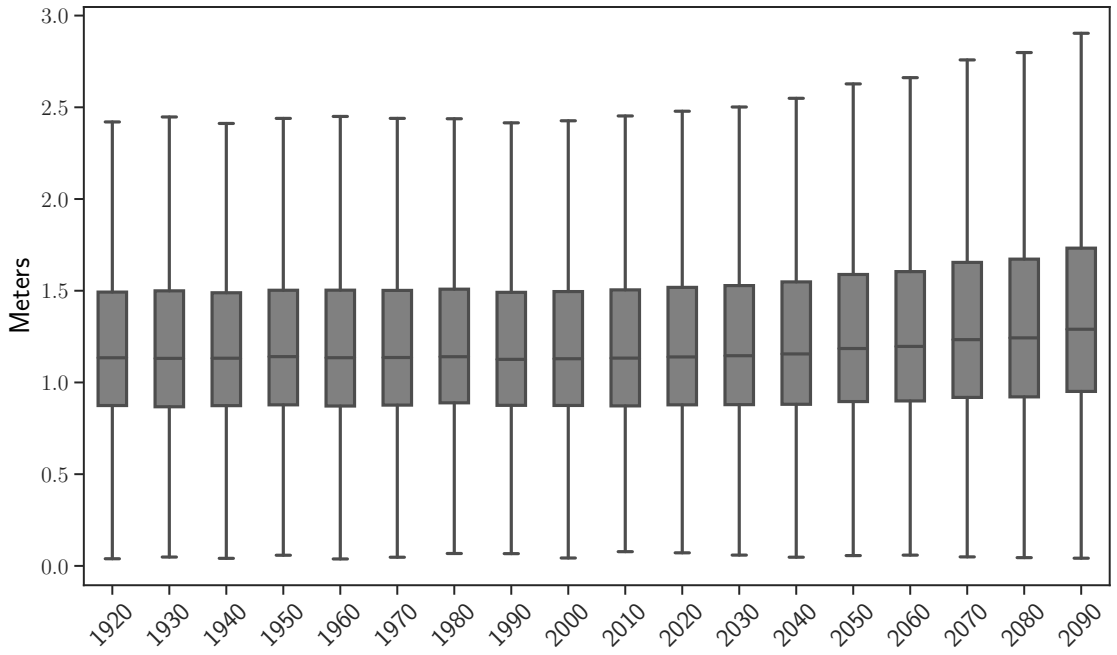

(b) Precipitation

Figure A.3: 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.

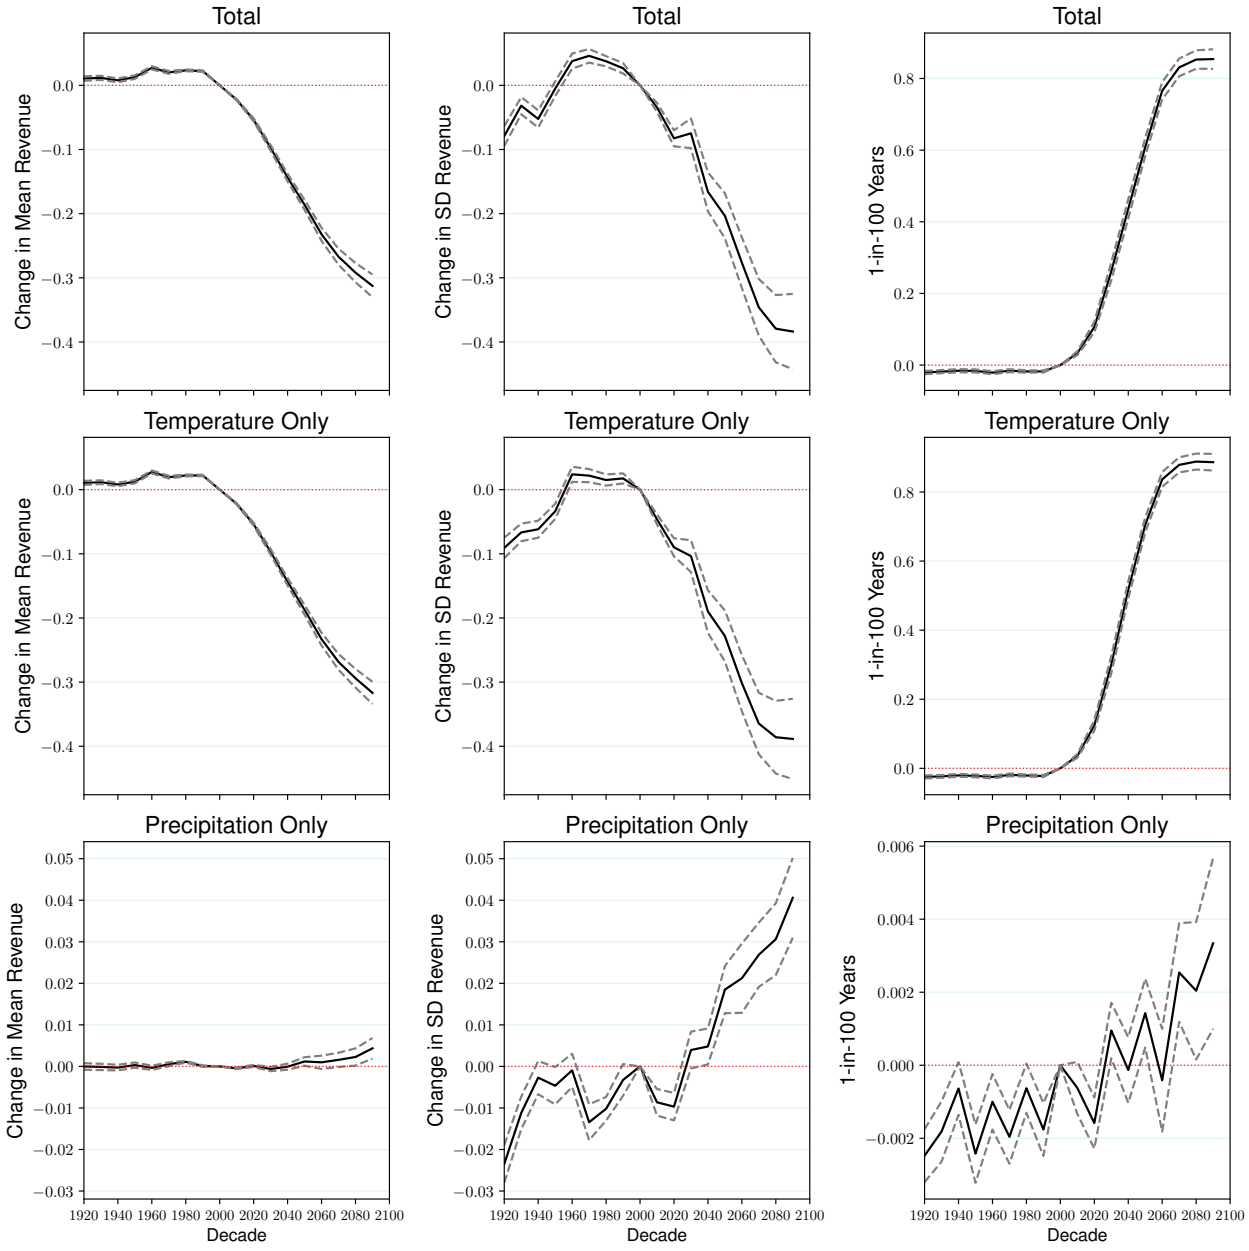

Figure A.4: Projected changes in agricultural revenue – three-degree temperature bins specification without additional precipitation variables

This figure shows the projection of the change in the mean (left), weather-induced standard deviation (center), and share of 1-in-100 bad years (right) for agricultural revenue relative to the 2000s for the *model using the three-degree temperature bins specification without additional precipitation variables*. Dashed lines display 95% confidence intervals. The top row presents the total projected effect of climate change, and the other rows display analogous estimates for projections in which only temperature (middle row) or precipitation (bottom row) are allowed to change from their values as of 2000.

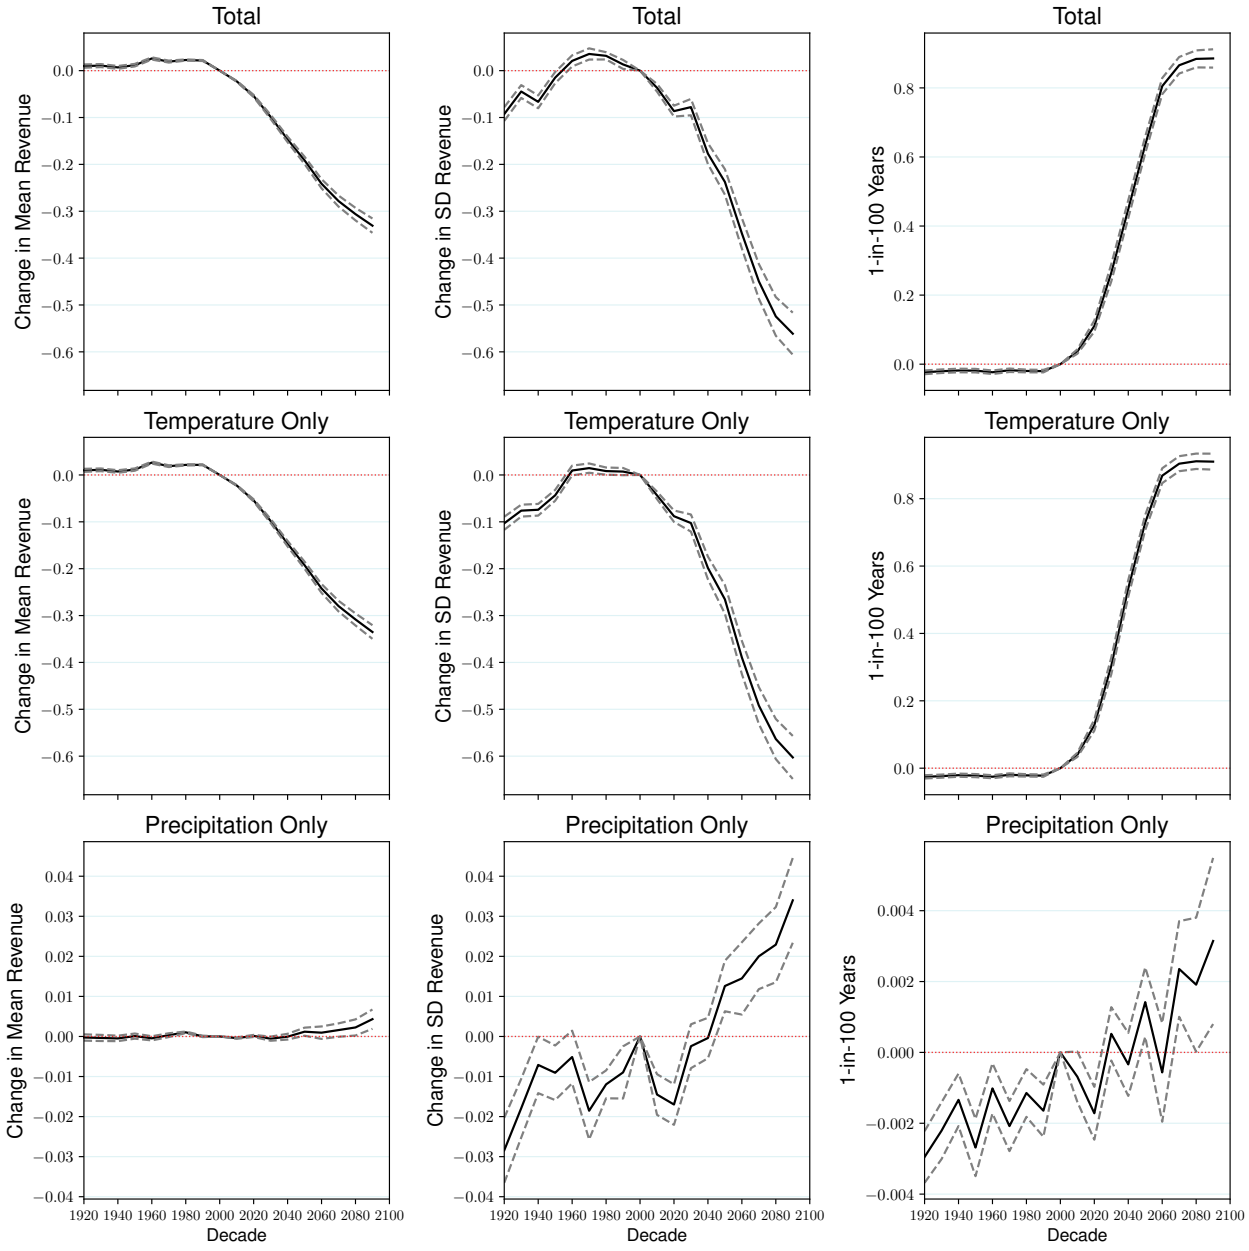

Figure A.5: Projected changes in agricultural revenue – decile temperature bins specification including HDD variable without additional precipitation variables

This figure shows the projection of the change in the mean (left), weather-induced standard deviation (center), and share of 1-in-100 bad years (right) for agricultural revenue relative to the 2000s for the *model using the decile temperature bins specification including an HDD variable without additional precipitation variables*. Dashed lines display 95% confidence intervals. The top row presents the total projected effect of climate change, and the other rows display analogous estimates for projections in which only temperature (middle row) or precipitation (bottom row) are allowed to change from their values as of 2000.

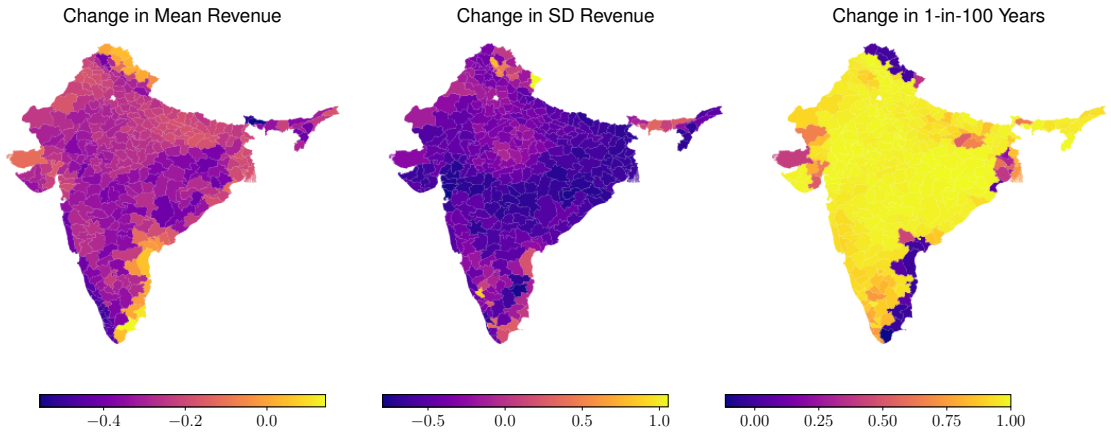

(a) Temperature Bins

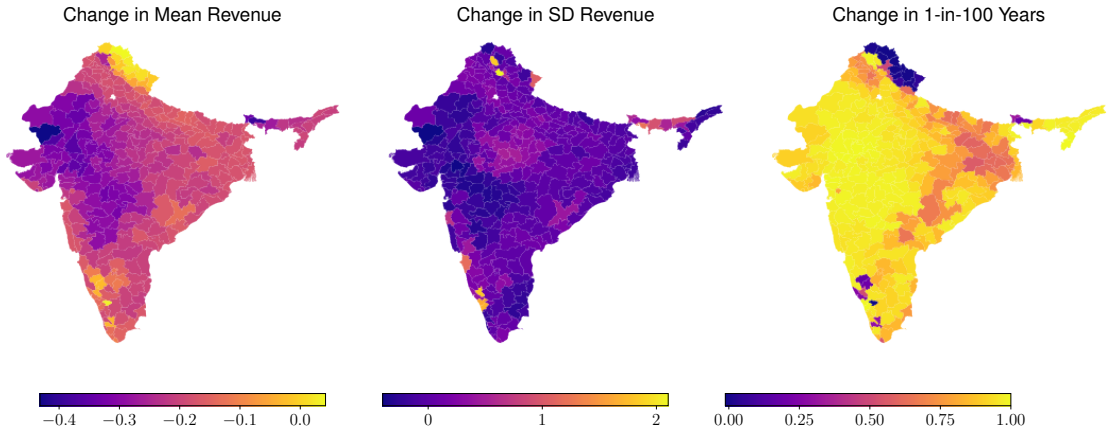

(b) Degree Days

Figure A.6: 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), without additional precipitation variables.

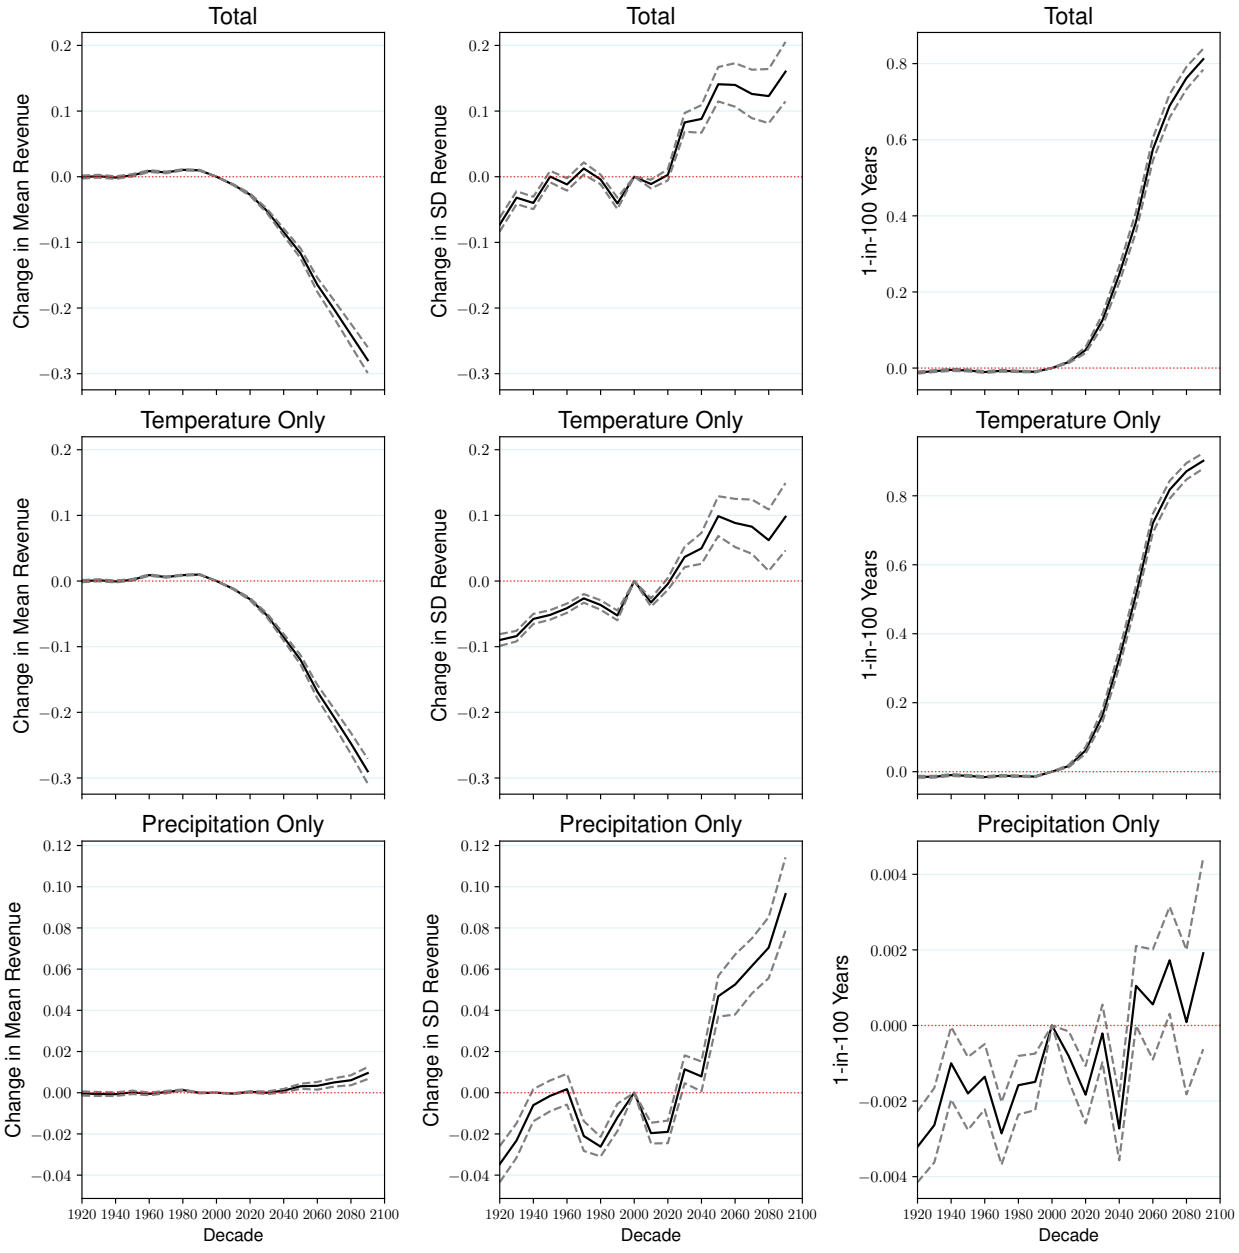

Figure A.7: Projected changes in agricultural revenue – quadratic degree-days specification without additional precipitation variables

This figure shows the projection of the change in the mean (left), weather-induced standard deviation (center), and share of 1-in-100 bad years (right) for agricultural revenue relative to the 2000s for the *model using the quadratic degree-days specification without additional precipitation variables*. Dashed lines display 95% confidence intervals. The top row presents the total projected effect of climate change, and the other rows display analogous estimates for projections in which only temperature (middle row) or precipitation (bottom row) are allowed to change from their values as of 2000.

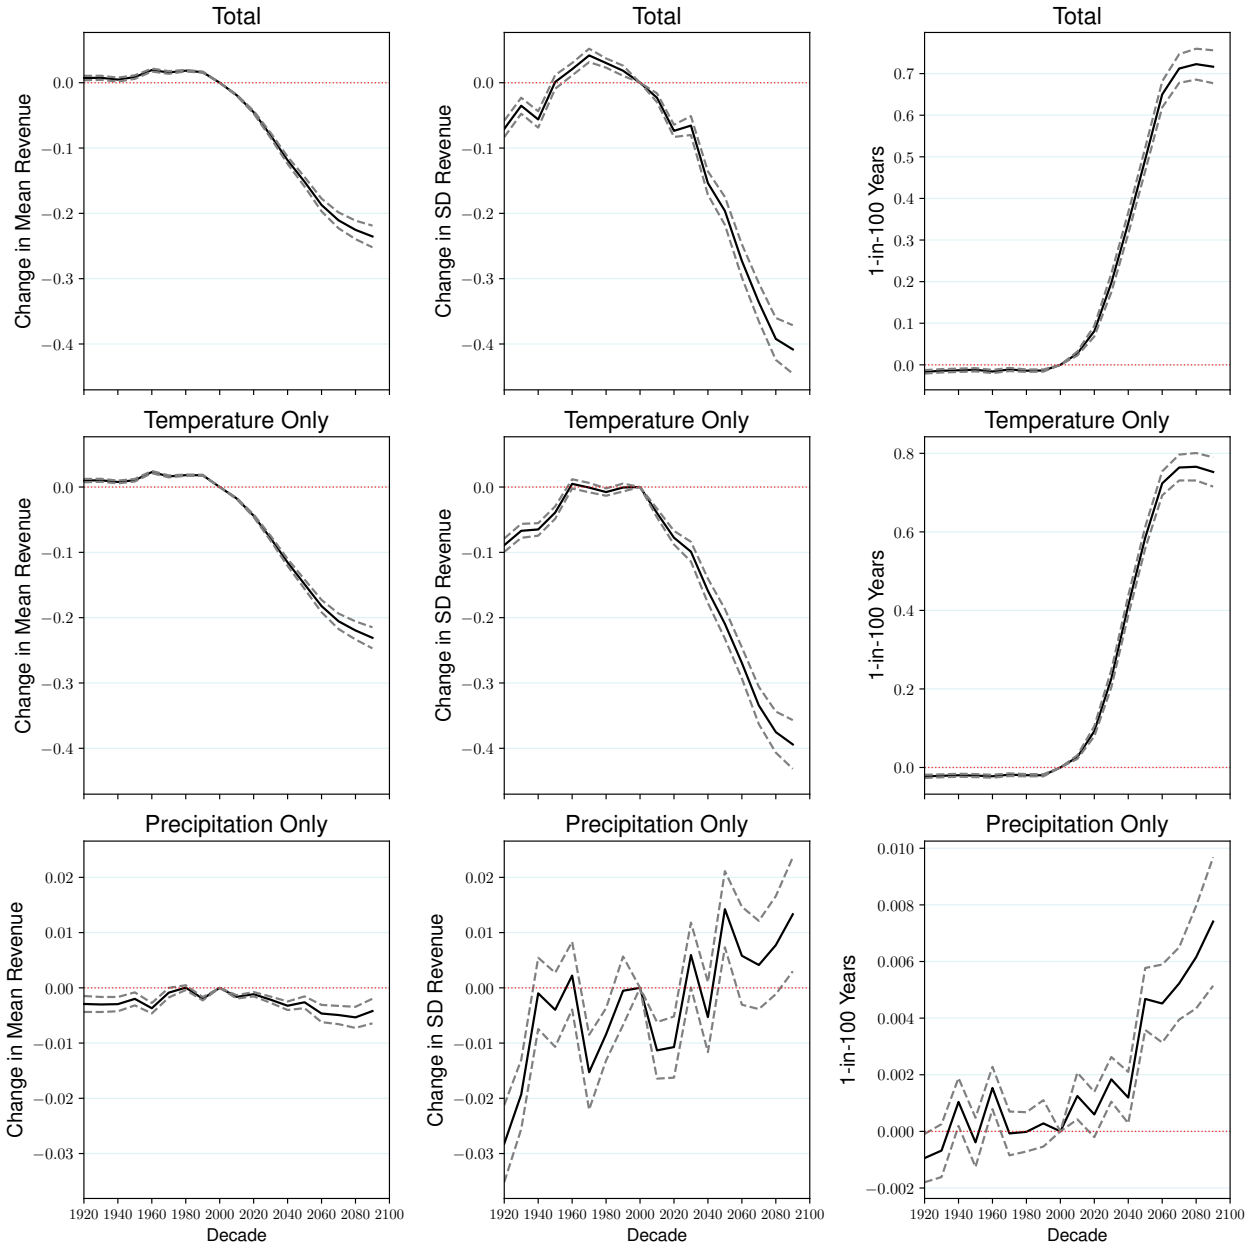

Figure A.8: Projected changes in agricultural revenue – decile temperature bins specification with additional precipitation variables

This figure shows the projection of the change in the mean (left), weather-induced standard deviation (center), and share of 1-in-100 bad years (right) for agricultural revenue relative to the 2000s for the *model using the decile temperature bins specification with additional precipitation variables*. Dashed lines display 95% confidence intervals. The top row presents the total projected effect of climate change, and the other rows display analogous estimates for projections in which only temperature (middle row) or precipitation (bottom row) are allowed to change from their values as of 2000.

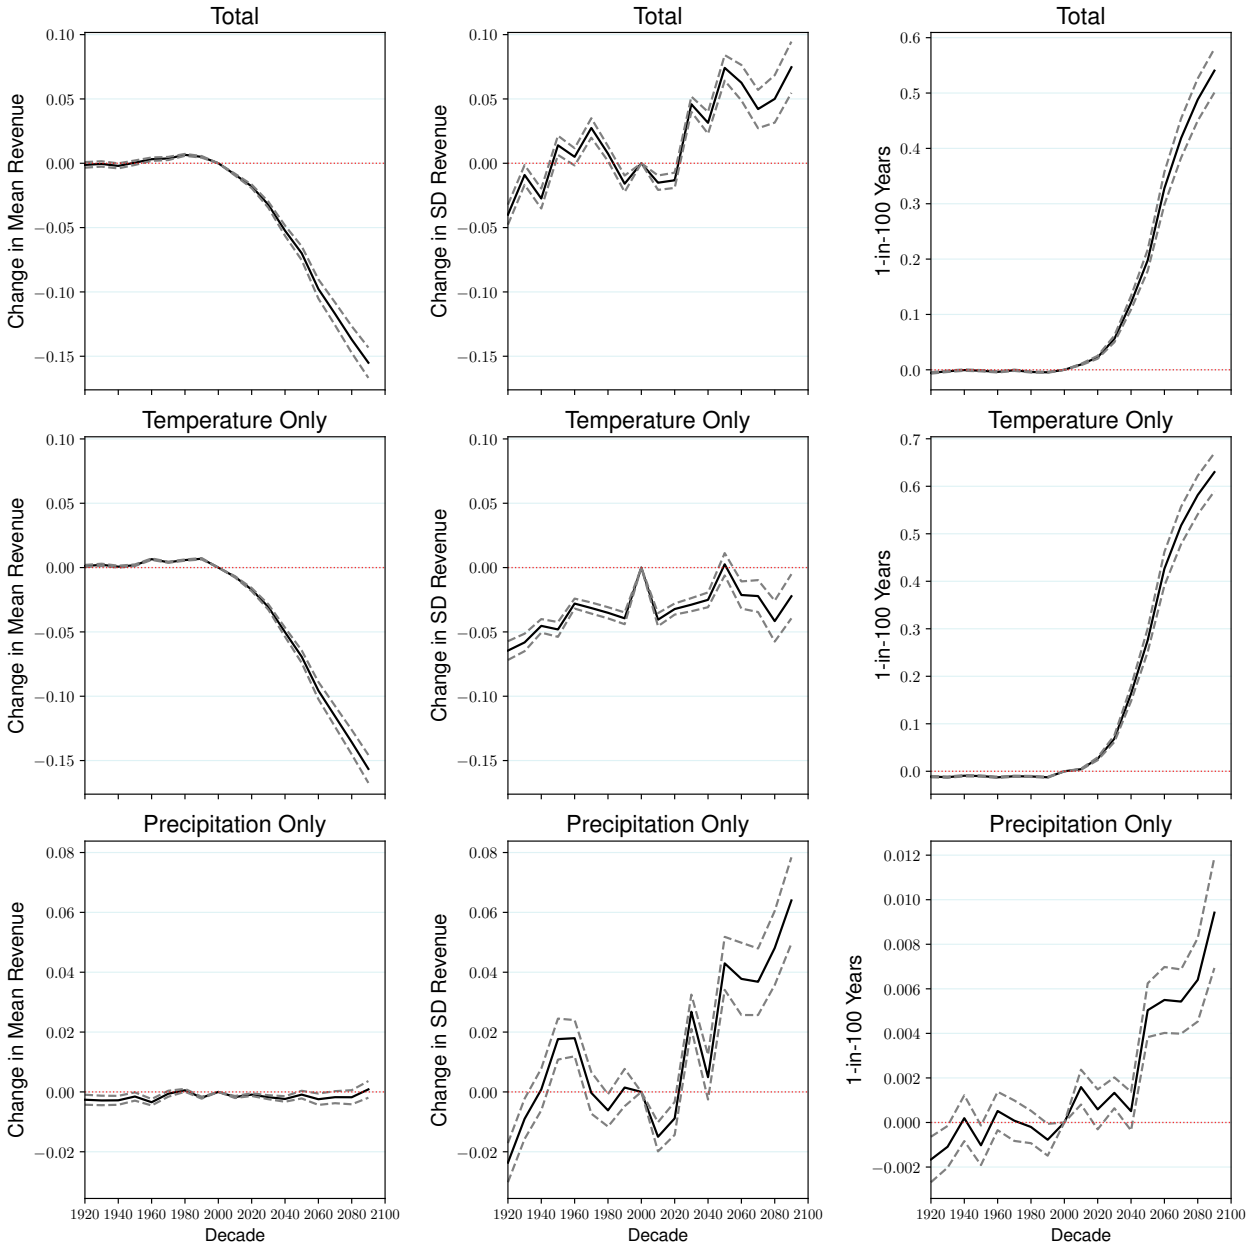

Figure A.9: Projected changes in agricultural revenue – linear degree-days specification with additional precipitation variables

This figure shows the projection of the change in the mean (left), weather-induced standard deviation (center), and share of 1-in-100 bad years (right) for agricultural revenue relative to the 2000s for the *model using the linear degree-days specification with additional precipitation variables*. Dashed lines display 95% confidence intervals. The top row presents the total projected effect of climate change, and the other rows display analogous estimates for projections in which only temperature (middle row) or precipitation (bottom row) are allowed to change from their values as of 2000.

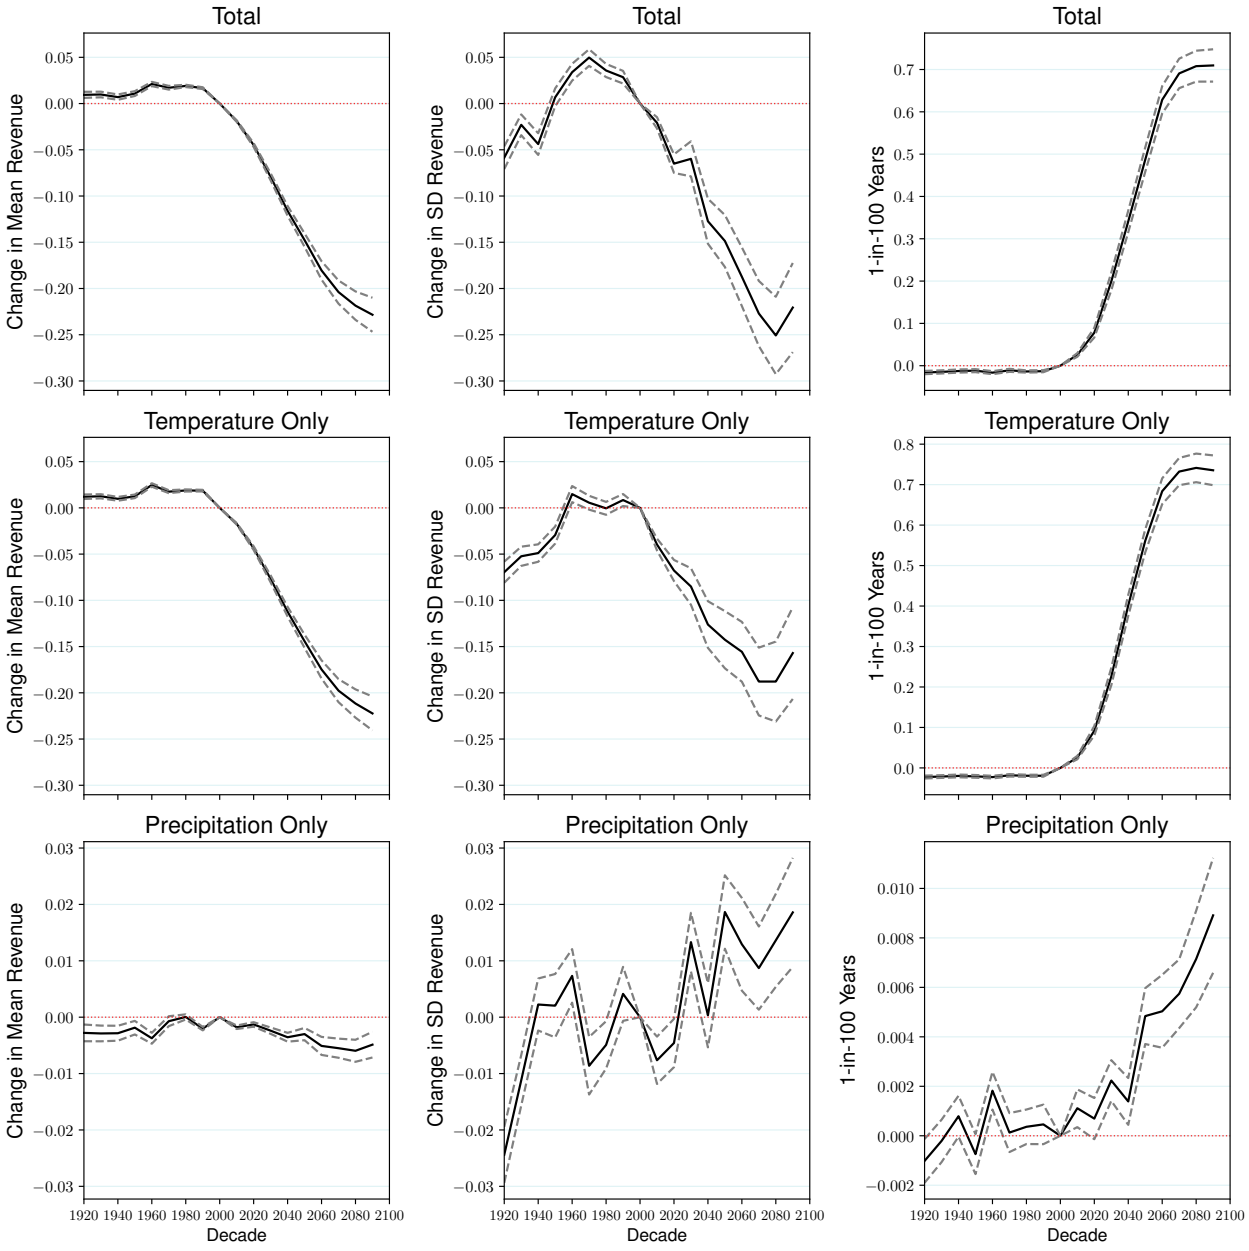

Figure A.10: Projected changes in agricultural revenue – three-degree temperature bins specification with additional precipitation variables

This figure shows the projection of the change in the mean (left), weather-induced standard deviation (center), and share of 1-in-100 bad years (right) for agricultural revenue relative to the 2000s for the *model using the three-degree temperature bins specification with additional precipitation variables*. Dashed lines display 95% confidence intervals. The top row presents the total projected effect of climate change, and the other rows display analogous estimates for projections in which only temperature (middle row) or precipitation (bottom row) are allowed to change from their values as of 2000.

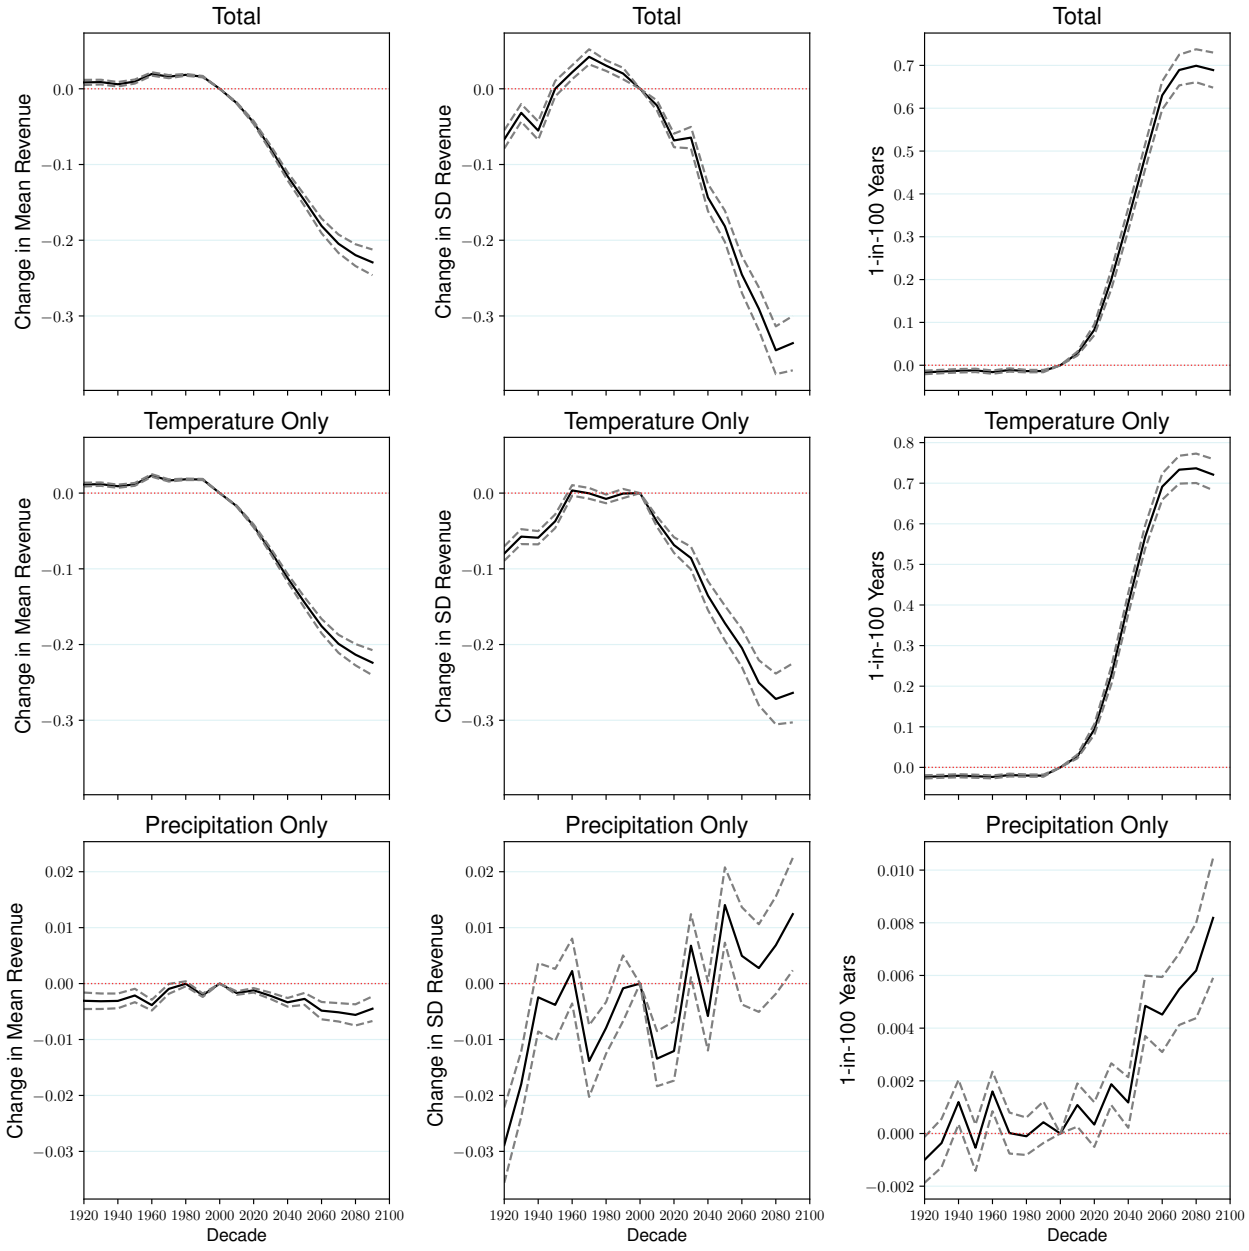

Figure A.11: Projected changes in agricultural revenue – decile temperature bins specification including HDD variable with additional precipitation variables

This figure shows the projection of the change in the mean (left), weather-induced standard deviation (center), and share of 1-in-100 bad years (right) for agricultural revenue relative to the 2000s for the *model using the decile temperature bins specification including an HDD variable with additional precipitation variables*. Dashed lines display 95% confidence intervals. The top row presents the total projected effect of climate change, and the other rows display analogous estimates for projections in which only temperature (middle row) or precipitation (bottom row) are allowed to change from their values as of 2000.

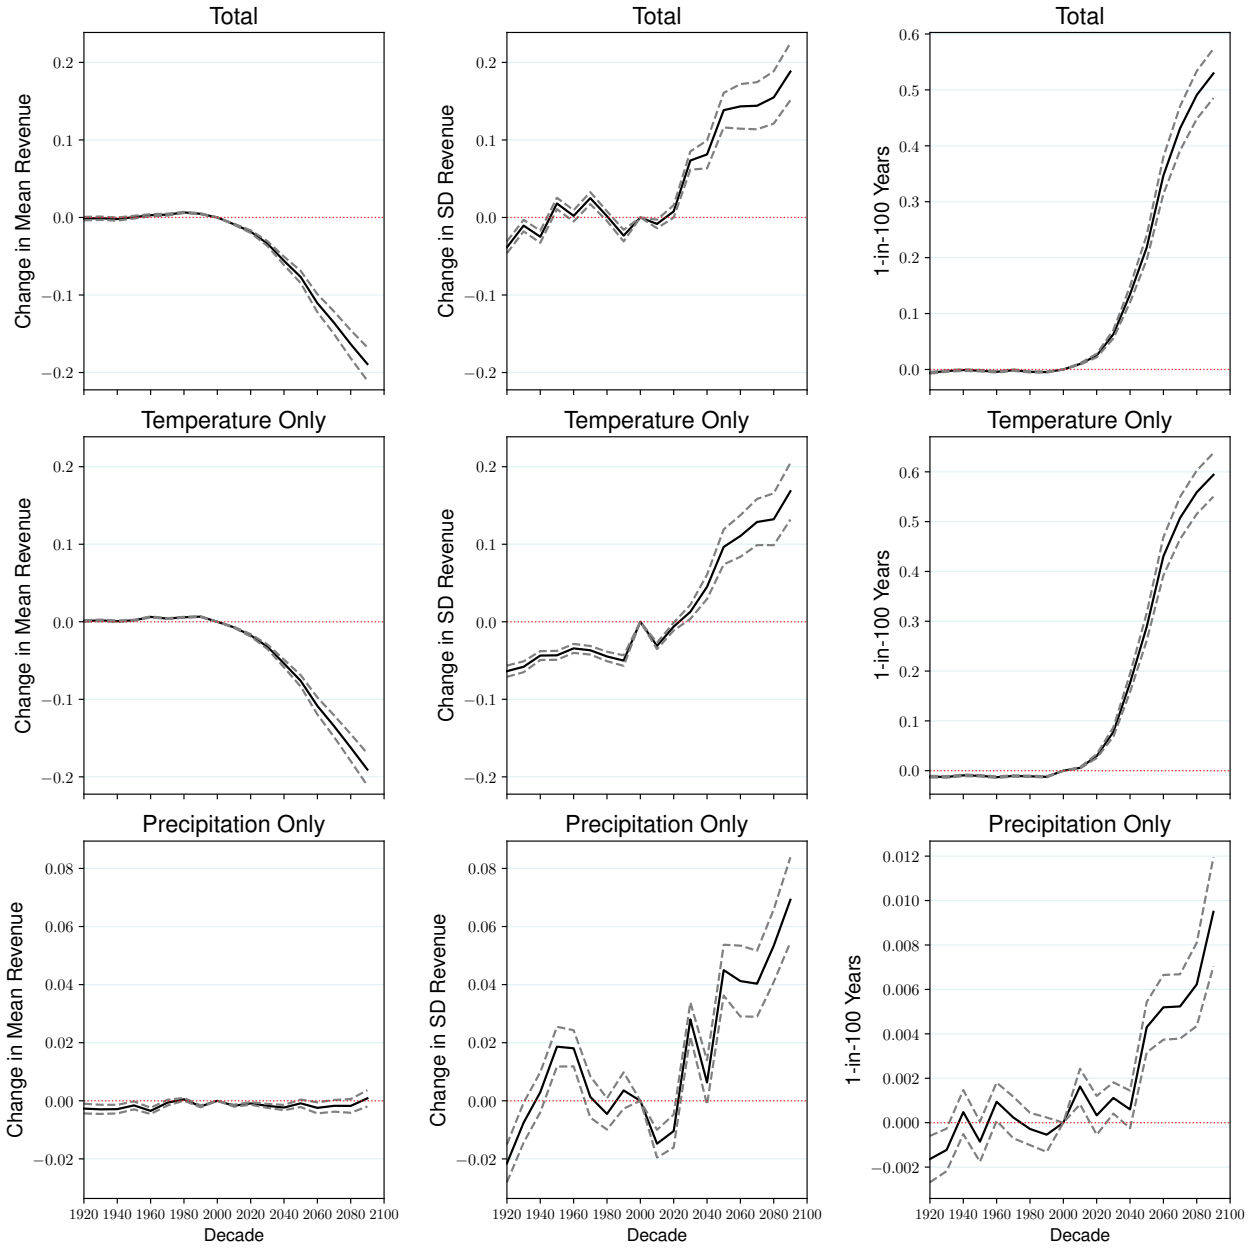

Figure A.12: Projected changes in agricultural revenue – quadratic degree-days specification with additional precipitation variables

This figure shows the projection of the change in the mean (left), weather-induced standard deviation (center), and share of 1-in-100 bad years (right) for agricultural revenue relative to the 2000s for the *model using the quadratic degree-days specification with additional precipitation variables*. Dashed lines display 95% confidence intervals. The top row presents the total projected effect of climate change, and the other rows display analogous estimates for projections in which only temperature (middle row) or precipitation (bottom row) are allowed to change from their values as of 2000.

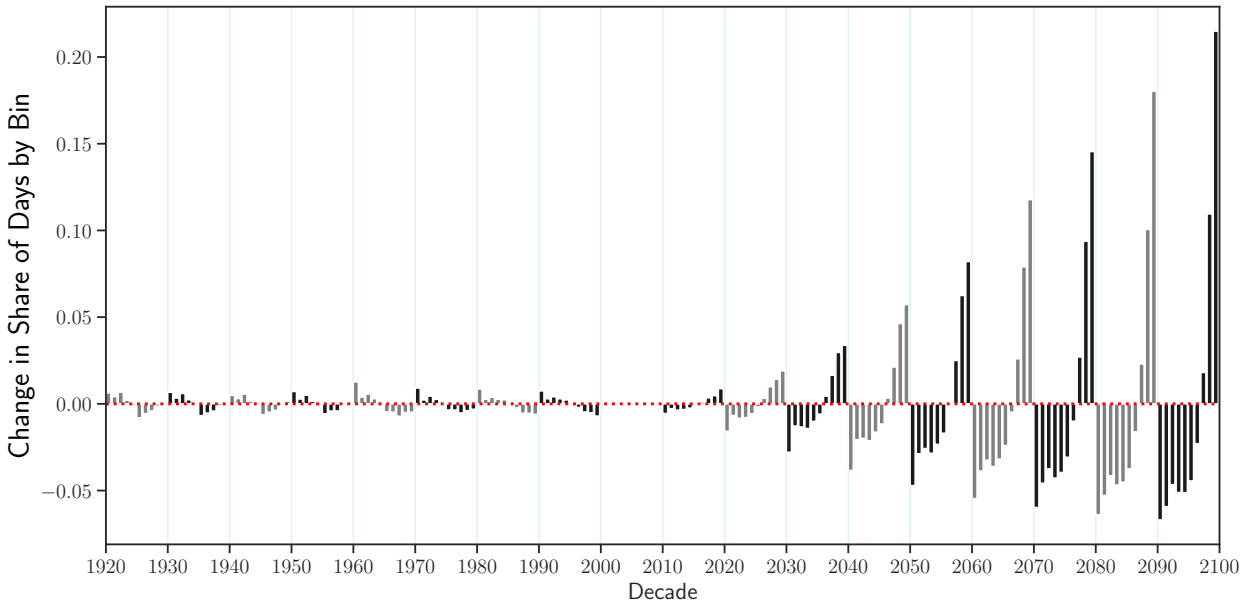

Figure A.13: Share of days in each decile temperature bin relative to the 2000s

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 Appendix Table A.2 for the limits of each bin in degrees Celsius.

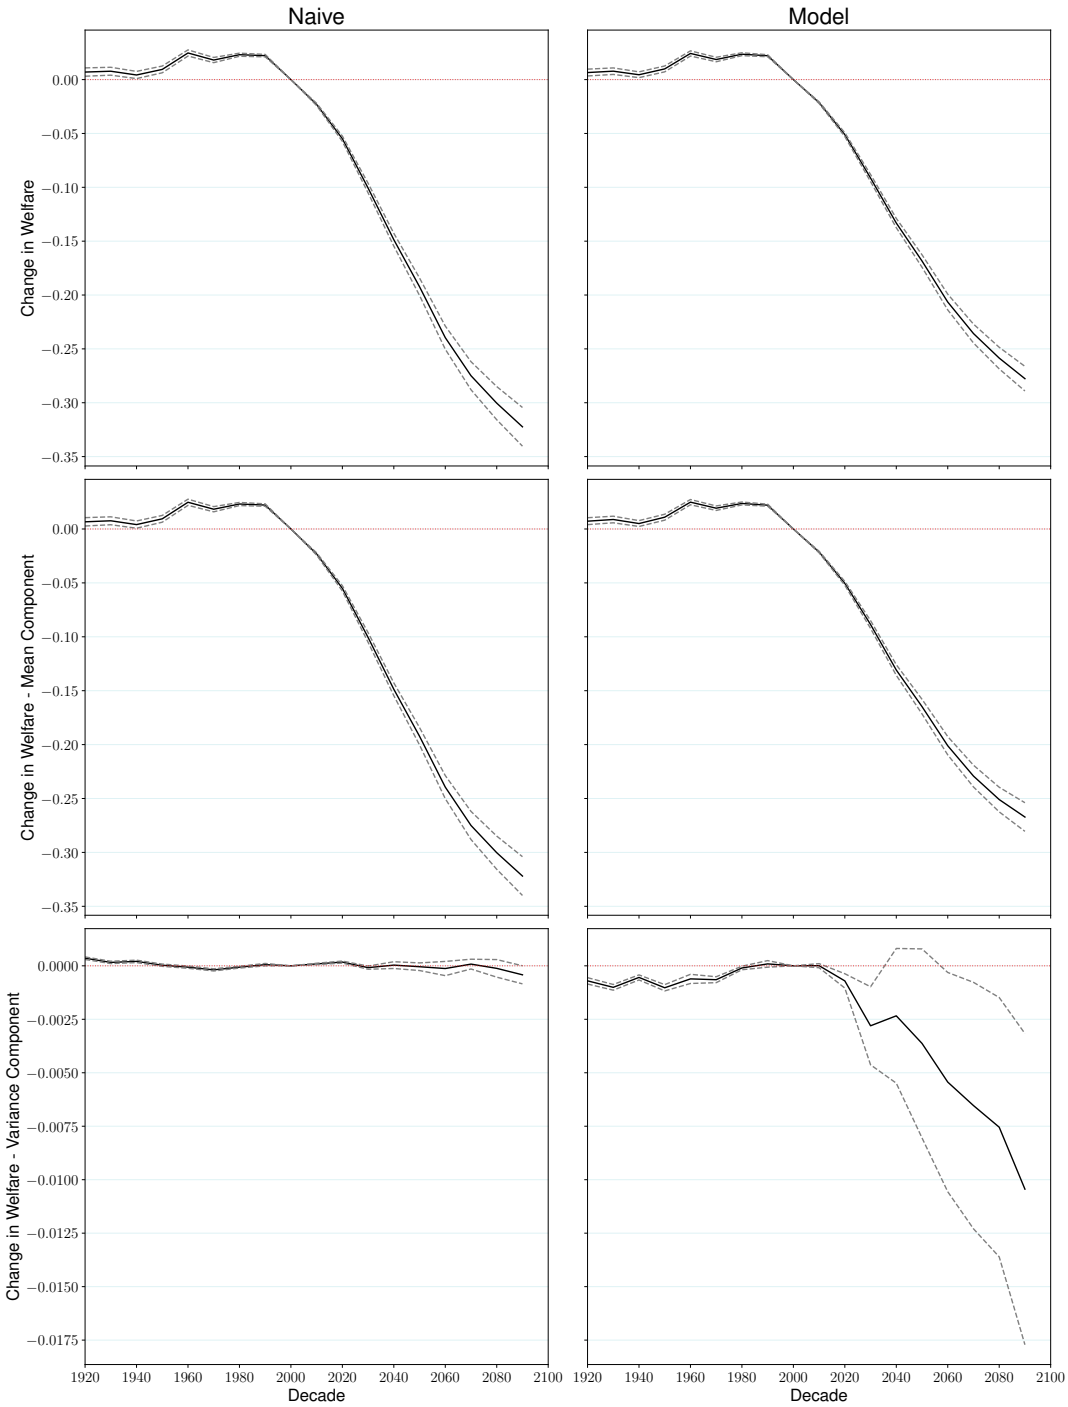

Figure A.14: Projected changes in welfare measures – three-degree temperature bins specification without additional precipitation variables

This figure shows the projected proportional changes (in percent) 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. The top row displays the total change in each welfare measure. The middle and bottom rows show changes driven by the mean and variance components of these measures, respectively. Results are based on the *three-degree temperature bins specification without additional precipitation variables*. Dashed lines represent 95% confidence intervals.

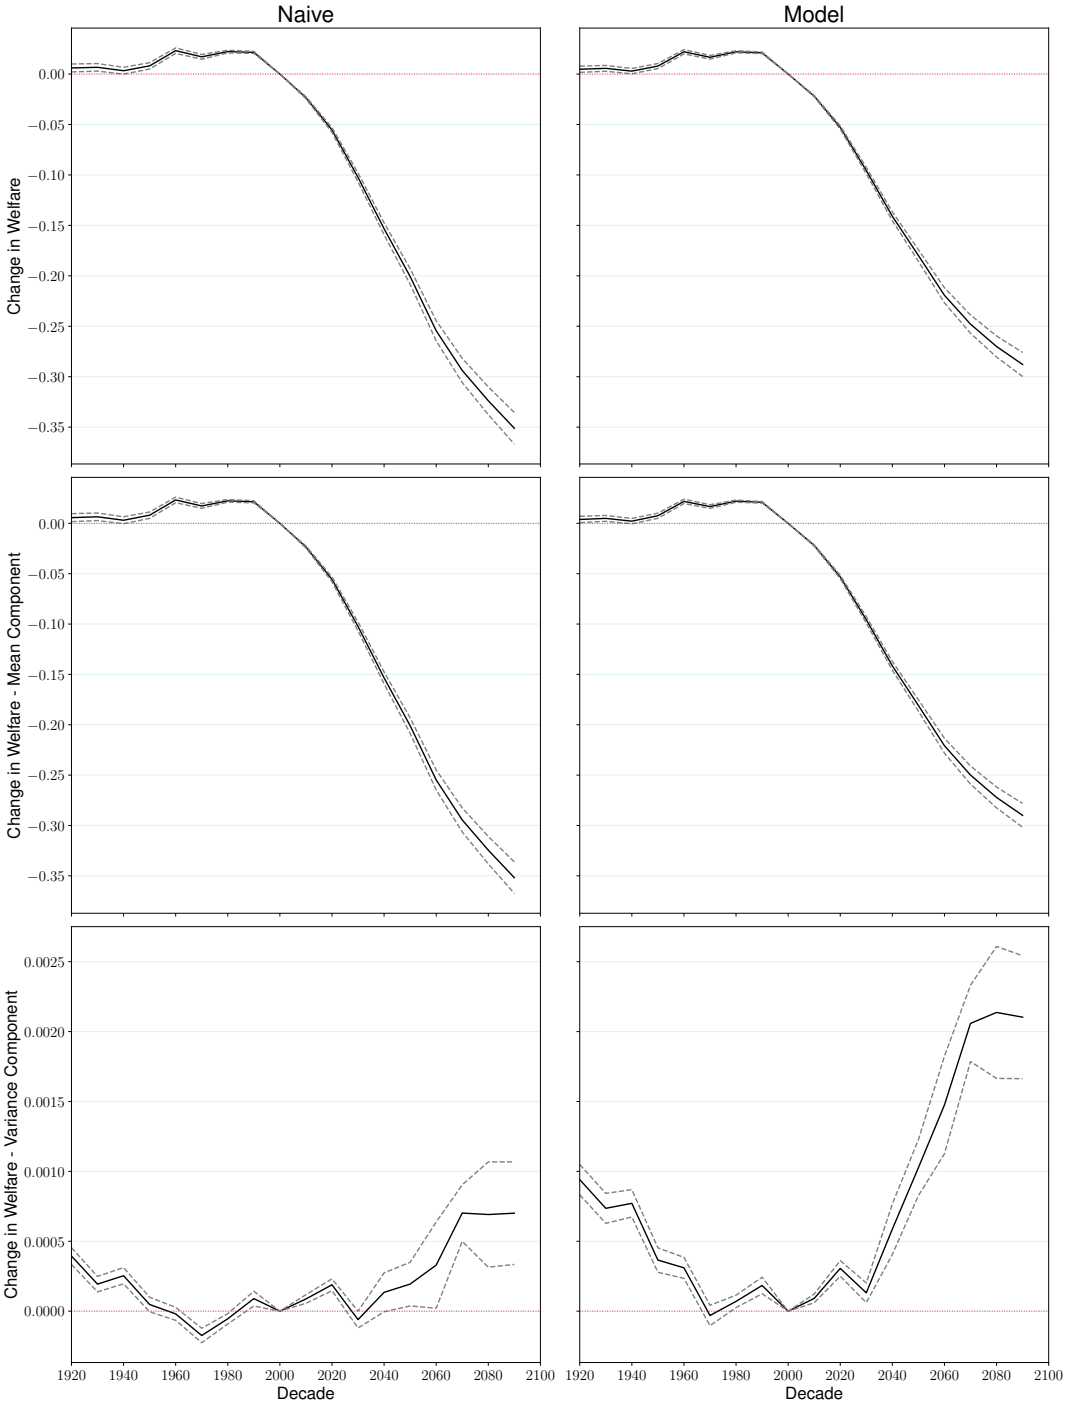

Figure A.15: Projected changes in welfare measures – decile temperature bins specification including HDD variable without additional precipitation variables

This figure shows the projected proportional changes (in percent) 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. The top row displays the total change in each welfare measure. The middle and bottom rows show changes driven by the mean and variance components of these measures, respectively. Results are based on the *decile temperature bins specification including an HDD variable without additional precipitation variables*. Dashed lines represent 95% confidence intervals.

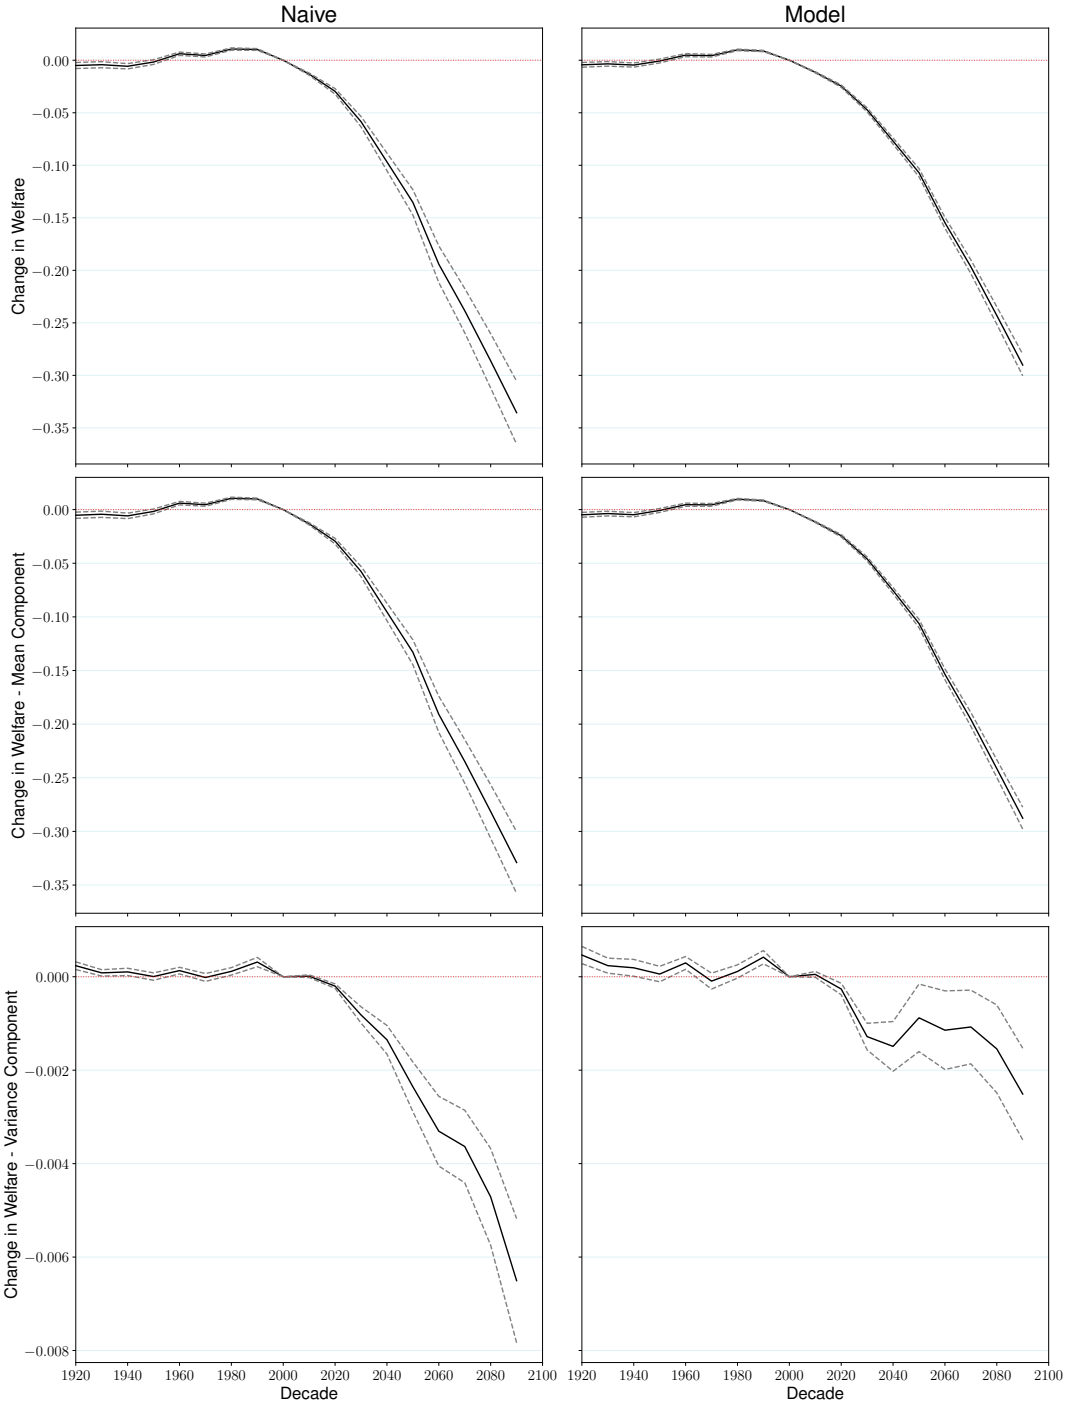

Figure A.16: Projected changes in welfare measures – quadratic degree-days specification without additional precipitation variables

This figure shows the projected proportional changes (in percent) 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. The top row displays the total change in each welfare measure. The middle and bottom rows show changes driven by the mean and variance components of these measures, respectively. Results are based on the *quadratic degree-days specification without additional precipitation variables*. Dashed lines represent 95% confidence intervals.

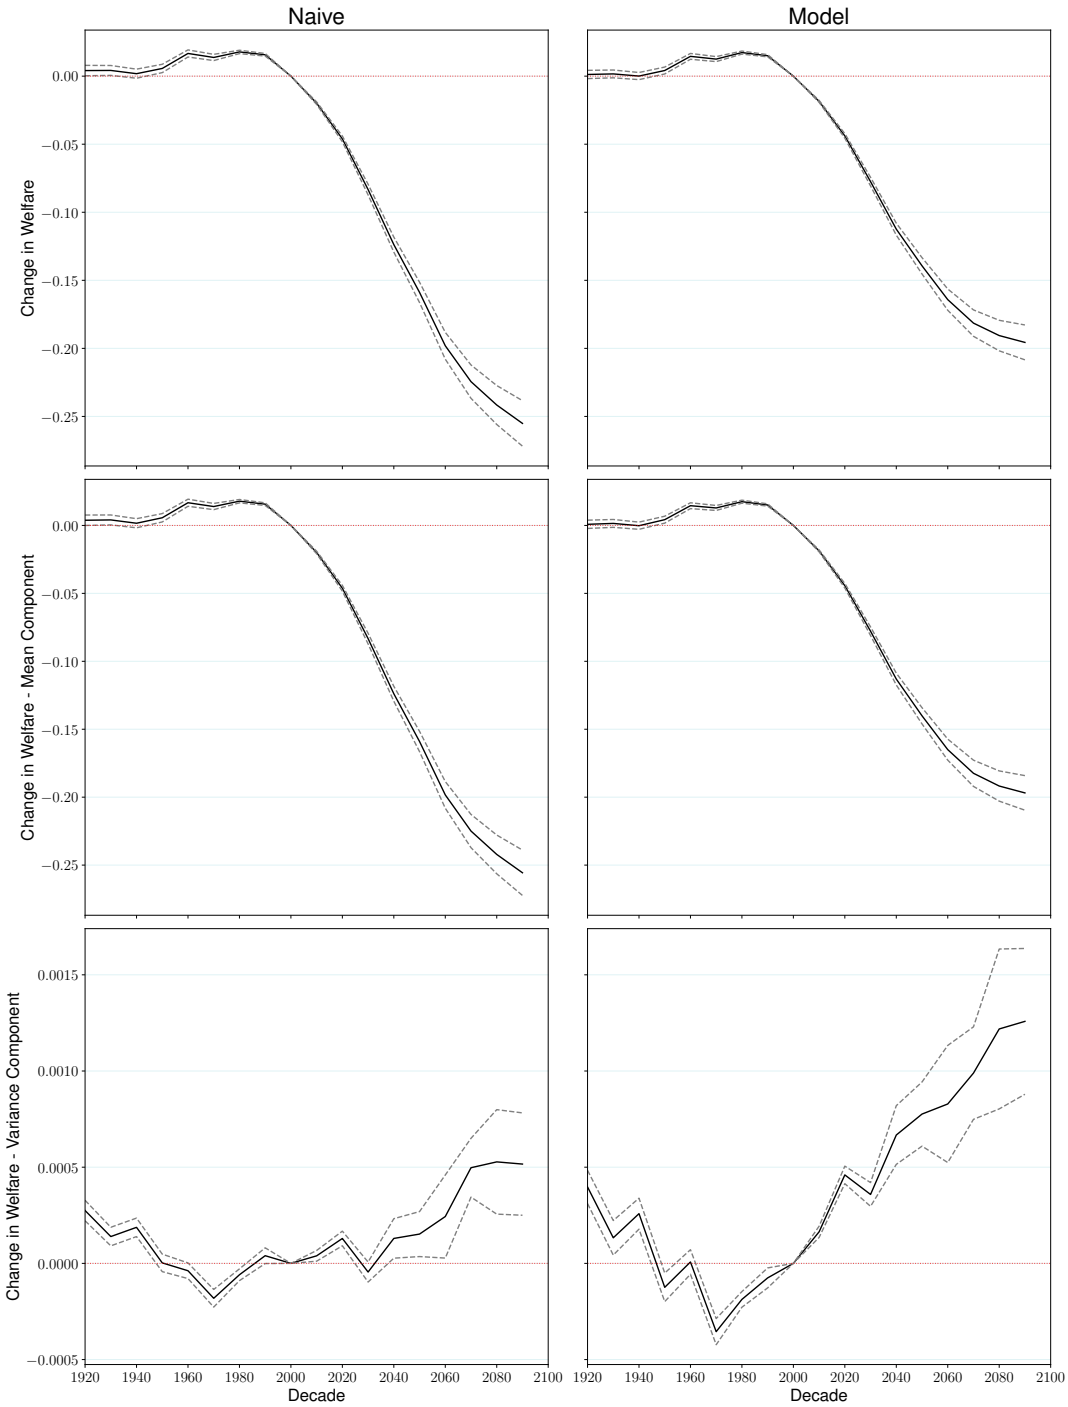

Figure A.17: Projected changes in welfare measures – decile temperature bins specification with additional precipitation variables

This figure shows the projected proportional changes (in percent) 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. The top row displays the total change in each welfare measure. The middle and bottom rows show changes driven by the mean and variance components of these measures, respectively. Results are based on the *decile temperature bins specification with additional precipitation variables*. Dashed lines represent 95% confidence intervals.

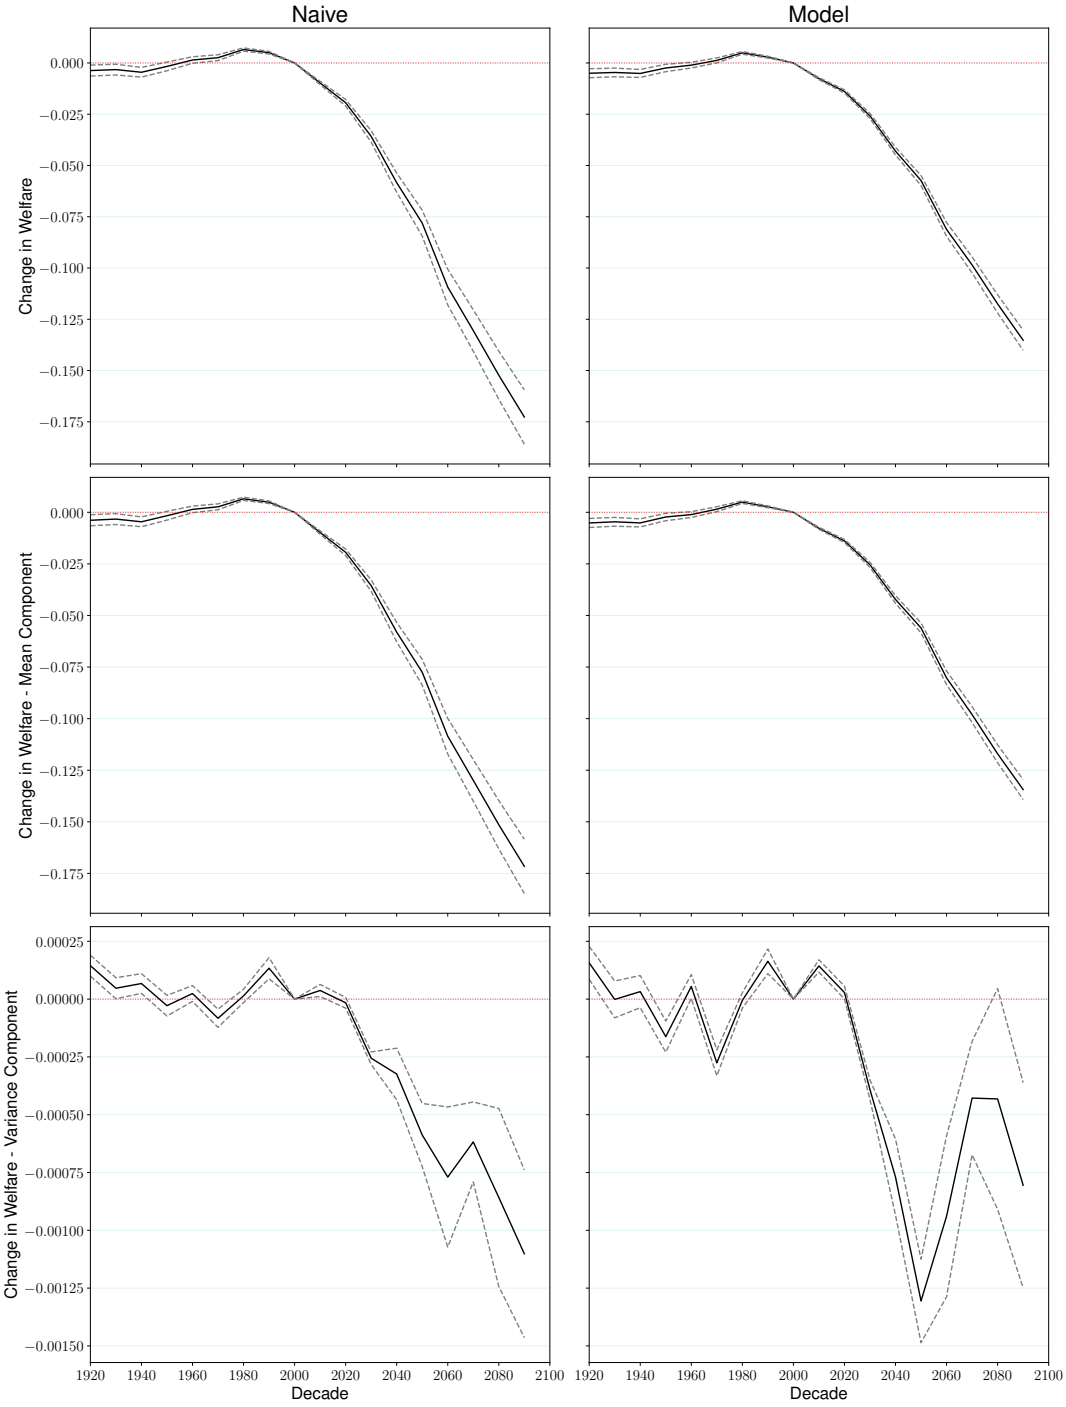

Figure A.18: Projected changes in welfare measures – linear degree-days specification with additional precipitation variables

This figure shows the projected proportional changes (in percent) 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. The top row displays the total change in each welfare measure. The middle and bottom rows show changes driven by the mean and variance components of these measures, respectively. Results are based on the *linear degree-days specification with additional precipitation variables*. Dashed lines represent 95% confidence intervals.

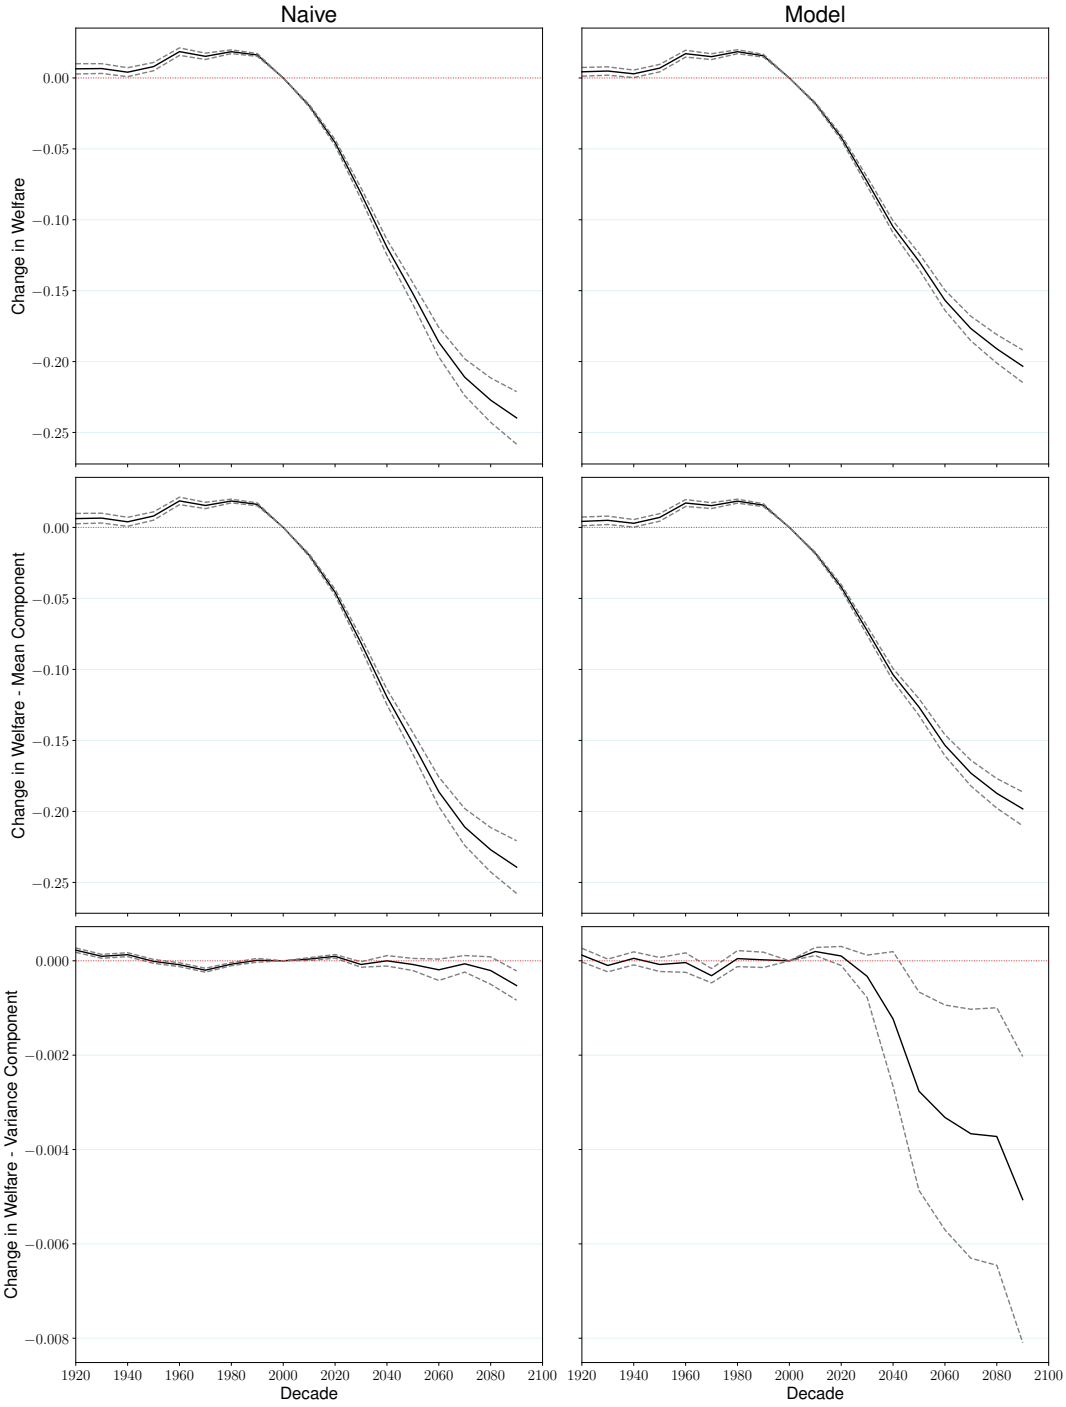

Figure A.19: Projected changes in welfare measures – three-degree temperature bins specification with additional precipitation variables

This figure shows the projected proportional changes (in percent) 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. The top row displays the total change in each welfare measure. The middle and bottom rows show changes driven by the mean and variance components of these measures, respectively. Results are based on the *three-degree temperature bins specification with additional precipitation variables*. Dashed lines represent 95% confidence intervals.

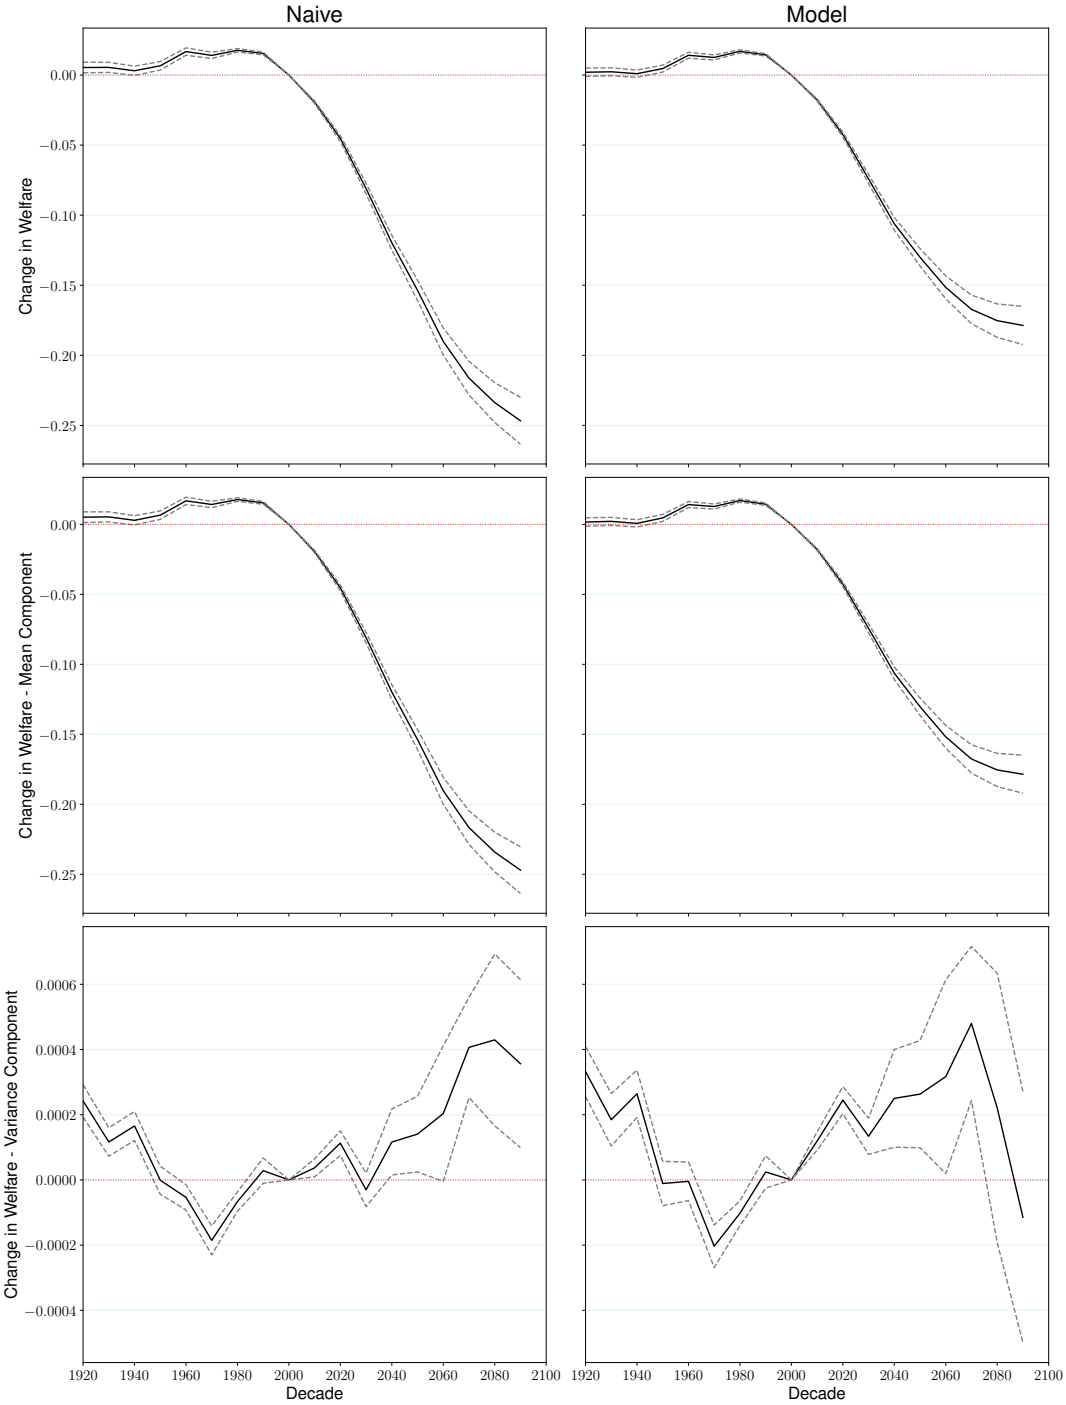

Figure A.20: Projected changes in welfare measures – decile temperature bins specification including HDD variable with additional precipitation variables

This figure shows the projected proportional changes (in percent) 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. The top row displays the total change in each welfare measure. The middle and bottom rows show changes driven by the mean and variance components of these measures, respectively. Results are based on the *decile temperature bins specification including an HDD variable with additional precipitation variables*. Dashed lines represent 95% confidence intervals.

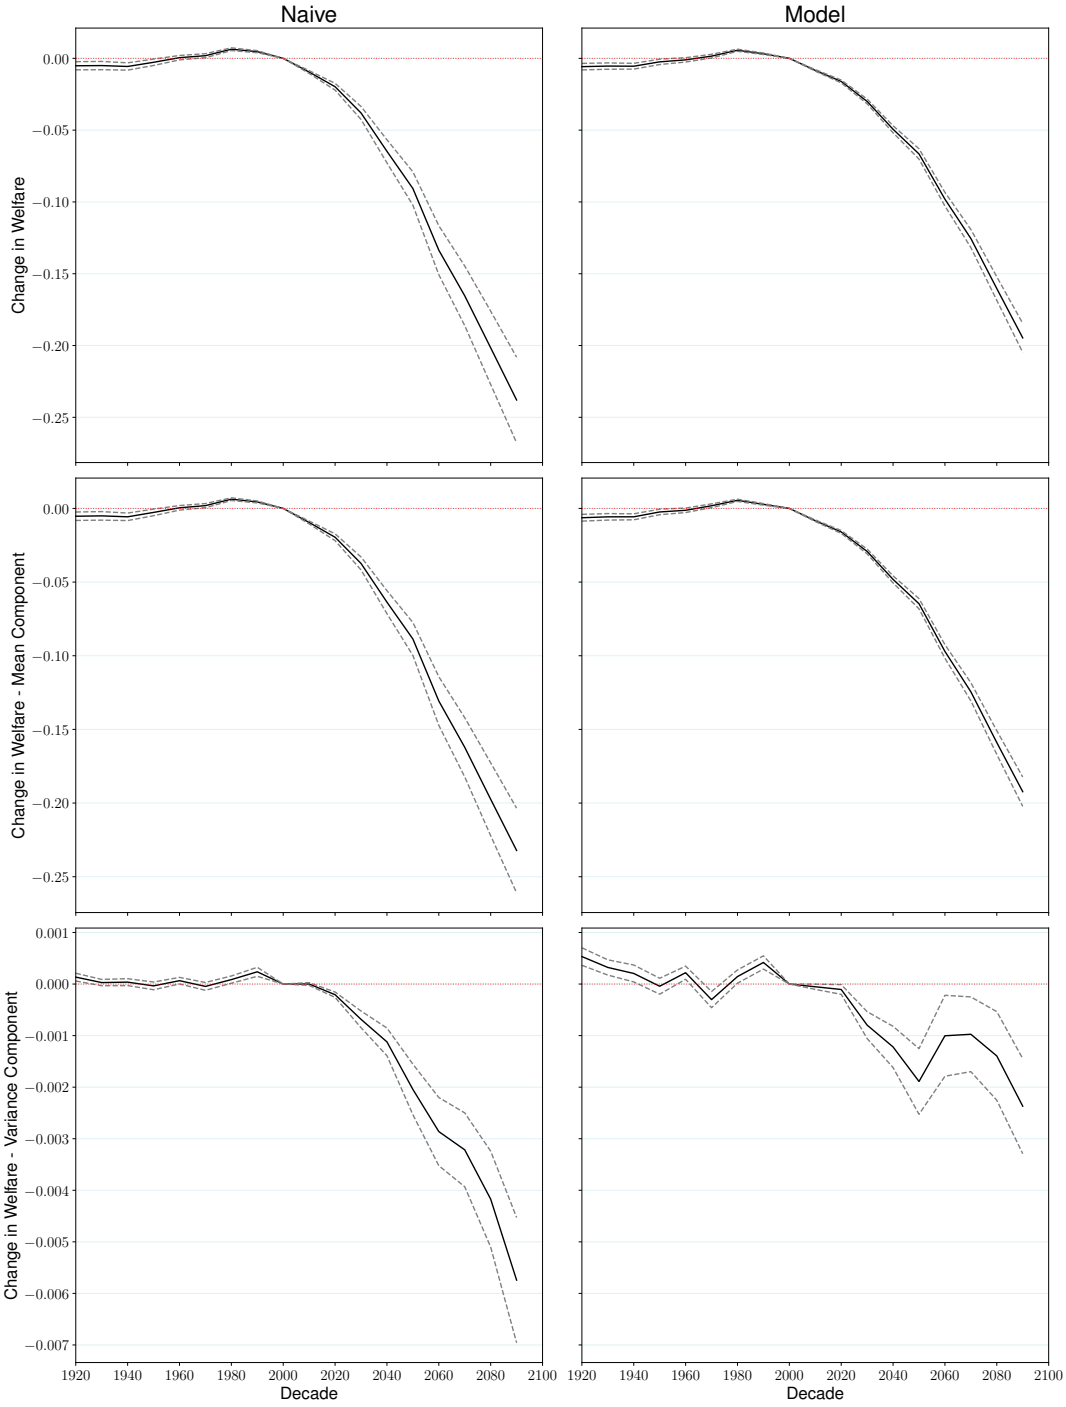

Figure A.21: Projected changes in welfare measures – quadratic degree-days specification with additional precipitation variables

This figure shows the projected proportional changes (in percent) 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. The top row displays the total change in each welfare measure. The middle and bottom rows show changes driven by the mean and variance components of these measures, respectively. Results are based on the *quadratic degree-days specification with additional precipitation variables*. Dashed lines represent 95% confidence intervals.

Table A.1: VDSA summary statistics in 2000

	(1)	(2)	(3)	(4)
	Share of	Share of land (among producers)		
Panel A. Crop-level summary statistics	districts	25th pctile	Median	75th pctile
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 pctile	Median	75th pctile
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 A.2: Upper limits of decile temperature bins

	(1) *Kharif*	(2) *Rabi*	(3) 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 A.3: Yield-weather estimation – decile 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) Linseed	(14) Rapeseed/Mustard	(15) Wheat	(16) Sugarcane
Bin [min,p10)	0.435	-0.470*	0.453***	0.237	0.630***	0.879***	0.871***	0.682***	0.770***	0.861***	0.387***	-0.0672	0.456**	0.244*	0.525***	1.196***
	(0.273)	(0.253)	(0.147)	(0.168)	(0.155)	(0.186)	(0.147)	(0.129)	(0.154)	(0.181)	(0.144)	(0.182)	(0.184)	(0.141)	(0.105)	(0.209)
Bin [p10,p20)	0.0597	-0.0953	0.250*	-0.0363	0.249**	0.453***	0.500***	0.336***	0.200	0.233*	0.203	-0.0490	0.641***	0.454***	0.457***	0.903***
	(0.227)	(0.207)	(0.127)	(0.116)	(0.118)	(0.142)	(0.120)	(0.104)	(0.132)	(0.123)	(0.132)	(0.117)	(0.159)	(0.111)	(0.0805)	(0.154)
Bin [p20,p30)	0.177	-0.272	0.0282	-0.205	-0.0493	0.319*	0.386**	0.0973	0.0674	-0.210	0.0238	-0.00236	0.177	0.206	0.143*	0.266
	(0.296)	(0.234)	(0.152)	(0.137)	(0.148)	(0.189)	(0.155)	(0.128)	(0.147)	(0.160)	(0.149)	(0.116)	(0.179)	(0.130)	(0.0865)	(0.162)
Bin [p40,p50)	0.267	-0.453*	0.174	-0.165	-0.103	0.175	-0.0464	-0.218*	0.232	-0.185	-0.217	0.0600	0.114	0.123	0.0165	-0.0125
	(0.314)	(0.243)	(0.193)	(0.147)	(0.187)	(0.204)	(0.159)	(0.122)	(0.188)	(0.189)	(0.159)	(0.145)	(0.186)	(0.160)	(0.114)	(0.205)
Bin [p50,p60)	-0.139	-0.294	-0.293	-0.0439	-0.191	0.0477	0.0596	-0.211*	-0.0792	0.0861	-0.123	-0.308***	-0.226	-0.207	-0.160	-0.0373
	(0.293)	(0.236)	(0.189)	(0.166)	(0.129)	(0.156)	(0.166)	(0.110)	(0.141)	(0.138)	(0.166)	(0.117)	(0.180)	(0.143)	(0.101)	(0.166)
Bin [p60,p70)	-0.152	-0.610***	-0.0682	-0.132	-0.333**	-0.0912	-0.459***	-0.509***	0.162	-0.472**	-0.438**	-0.638***	-0.729***	-0.404**	-0.386***	-0.343**
	(0.305)	(0.214)	(0.210)	(0.132)	(0.143)	(0.177)	(0.130)	(0.126)	(0.175)	(0.189)	(0.186)	(0.146)	(0.228)	(0.170)	(0.112)	(0.165)
Bin [p70,p80)	-0.274	-0.672***	-0.188	-0.661***	-0.392***	-0.447**	-0.382**	-0.626***	-0.554***	-0.803***	-0.358**	-0.570***	-0.498**	-0.471***	-0.393***	-0.486***
	(0.263)	(0.203)	(0.196)	(0.158)	(0.140)	(0.182)	(0.148)	(0.118)	(0.166)	(0.180)	(0.159)	(0.123)	(0.198)	(0.155)	(0.122)	(0.162)
Bin [p80,p90)	-0.265	-0.927***	-0.588***	-0.786***	-0.870***	-0.663***	-0.432***	-0.779***	-0.756***	-1.031***	-0.143	-0.743***	-0.590***	-0.428***	-0.434***	-0.371**
	(0.237)	(0.182)	(0.181)	(0.119)	(0.138)	(0.162)	(0.127)	(0.127)	(0.148)	(0.176)	(0.154)	(0.124)	(0.193)	(0.149)	(0.107)	(0.182)
Bin [p90,max]	-0.463	-0.772***	-0.422**	-0.542***	-1.213***	-1.220***	-0.617***	-0.399***	-1.358***	-1.307***	-0.366**	-0.853***	-0.754***	-0.627***	-0.567***	-0.511***
	(0.284)	(0.178)	(0.202)	(0.140)	(0.138)	(0.203)	(0.151)	(0.113)	(0.202)	(0.200)	(0.159)	(0.131)	(0.165)	(0.145)	(0.109)	(0.147)
Monsoon prec	0.446***	0.451***	0.186***	0.389***	-0.206***	0.853***	0.252***	0.366***	-0.0208	0.438***	0.207***	0.326***	0.199***	0.175***	0.218***	0.0612
	(0.116)	(0.110)	(0.0667)	(0.0873)	(0.0621)	(0.133)	(0.0596)	(0.0404)	(0.0636)	(0.107)	(0.0353)	(0.0363)	(0.0445)	(0.0314)	(0.0239)	(0.0380)
Monsoon prec^2	-0.136***	-0.160***	-0.0401**	-0.121***	0.0305	-0.427***	-0.0798***	-0.0805***	-0.0342*	-0.214***	-0.0151	-0.0549***	-0.0359***	-0.0168*	-0.0323***	-0.0268**
	(0.0458)	(0.0467)	(0.0190)	(0.0319)	(0.0220)	(0.0633)	(0.0221)	(0.0116)	(0.0200)	(0.0485)	(0.0113)	(0.0113)	(0.0118)	(0.00965)	(0.00770)	(0.0118)
Rabi prec											-0.199*	-0.520***	-0.240***	-0.656***	-0.321***	
											(0.106)	(0.0877)	(0.0879)	(0.0969)	(0.0646)	
Rabi prec^2											0.423***	0.220***	0.124	0.463***	0.190***	
											(0.133)	(0.0732)	(0.0774)	(0.104)	(0.0657)	
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 without additional precipitation variables*. Definitions of the variables may be found in Section 3. See also Appendix Table A.2 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 A.4: 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) Linseed	(14) Rapeseed/Mustard	(15) Wheat	(16) Sugarcane
Bin (−∞,0)			1.575	3.256	-0.0517	-2.669*	-1.195	0.290	-0.0386		2.600***	3.376**	6.355***	4.432***	3.406***	-5.072***
			(1.104)	(4.049)	(0.730)	(1.489)	(0.896)	(0.542)	(0.975)		(0.440)	(1.458)	(1.279)	(0.646)	(0.456)	(1.889)
Bin [0,3)			2.516**	3.580	0.691	-12.44	-0.103	1.001**	-0.292		1.593***	2.812***	3.941***	3.401***	2.186***	-0.717
			(1.056)	(3.797)	(0.741)	(7.934)	(0.901)	(0.489)	(1.044)		(0.297)	(1.023)	(0.718)	(0.856)	(0.272)	(1.384)
Bin [3,6)			2.223**	1.147	0.886		-0.767	0.484	1.395**		2.074***	2.668**	2.399**	3.051***	2.031***	-4.245
			(1.011)	(3.895)	(0.685)		(1.251)	(0.441)	(0.703)		(0.290)	(1.054)	(0.935)	(0.848)	(0.287)	(2.581)
Bin [6,9)			1.800*	1.186	1.062	1.223	-0.139	0.563	0.454		1.367***	2.213***	0.757	2.346***	1.553***	-2.086**
			(0.992)	(3.994)	(0.720)	(2.525)	(0.885)	(0.351)	(0.802)		(0.202)	(0.592)	(0.582)	(0.593)	(0.240)	(1.054)
Bin [9,12)		0.199	2.496***	0.742	0.835	-2.192***	0.123	0.522	0.795	-164.5***	0.532***	0.159	0.242	-0.299	0.505***	1.332***
		(1.136)	(0.854)	(2.415)	(0.650)	(0.695)	(1.033)	(0.340)	(0.629)	(22.84)	(0.171)	(0.298)	(0.365)	(0.247)	(0.140)	(0.319)
Bin [12,15)		3.637*	2.230***	1.960*	1.196	-1.213	0.148	0.545**	-0.0126	10.42***	0.453***	-0.00188	0.613***	0.402***	0.575***	1.177***
		(2.019)	(0.706)	(1.029)	(0.754)	(1.555)	(0.717)	(0.226)	(0.741)	(0.850)	(0.116)	(0.160)	(0.170)	(0.130)	(0.0817)	(0.192)
Bin [15,18)	-33.59	2.937**	1.329***	0.789	0.481	0.622	0.0183	0.387	-0.199	6.958***	0.276**	0.0497	0.616***	0.376***	0.464***	1.088***
	(26.27)	(1.384)	(0.452)	(0.771)	(0.400)	(0.612)	(0.587)	(0.238)	(0.485)	(2.536)	(0.113)	(0.105)	(0.140)	(0.102)	(0.0731)	(0.158)
Bin [18,21)	0.212	-1.193***	0.408**	0.188	0.586***	0.839***	0.566***	0.521***	0.484**	0.702***	0.143	0.212**	0.119	0.142	0.131**	0.340**
	(0.292)	(0.306)	(0.194)	(0.130)	(0.160)	(0.315)	(0.178)	(0.113)	(0.227)	(0.214)	(0.0961)	(0.0882)	(0.111)	(0.0960)	(0.0664)	(0.133)
Bin [24,27)	-0.0644	-0.0748	-0.265***	-0.0911	-0.377***	-0.357***	-0.525***	-0.445***	-0.209**	-0.374***	-0.0844	-0.461***	-0.463***	-0.373***	-0.281***	-0.231**
	(0.129)	(0.127)	(0.0659)	(0.0838)	(0.0840)	(0.0957)	(0.0828)	(0.0763)	(0.0875)	(0.0847)	(0.113)	(0.0944)	(0.138)	(0.0989)	(0.0778)	(0.0962)
Bin [27,30)	-0.390**	-0.565***	-0.541***	-0.488***	-0.774***	-0.860***	-0.985***	-1.015***	-0.607***	-0.992***	0.0840	-0.616***	-0.411**	-0.478***	-0.374***	-0.583***
	(0.177)	(0.148)	(0.0989)	(0.0900)	(0.0987)	(0.117)	(0.0941)	(0.0929)	(0.106)	(0.117)	(0.131)	(0.106)	(0.161)	(0.118)	(0.0840)	(0.108)
Bin [30,33)	-0.541**	-0.722***	-0.858***	-0.897***	-1.584***	-1.638***	-1.181***	-0.965***	-1.692***	-1.713***	-0.200	-0.694***	-0.586**	-0.605***	-0.385***	-0.496***
	(0.263)	(0.172)	(0.235)	(0.146)	(0.140)	(0.196)	(0.132)	(0.113)	(0.210)	(0.193)	(0.143)	(0.151)	(0.226)	(0.147)	(0.103)	(0.141)
Bin [33,36)	-0.502	-0.492**	-0.564**	-0.277*	-1.376***	-1.593***	-1.305***	-0.637***	-1.478***	-1.456***	-1.237***	-1.083***	-1.863***	-1.305***	-1.100***	-0.875***
	(0.405)	(0.213)	(0.233)	(0.161)	(0.145)	(0.222)	(0.183)	(0.124)	(0.241)	(0.218)	(0.416)	(0.287)	(0.394)	(0.293)	(0.224)	(0.158)
Bin [36,39)	-1.829***	-0.938**	-0.762	-0.387	-1.564***	-2.095***	-0.874***	-0.736***	-1.540***	-1.377***	2.598	-2.091	4.453	-3.685	-2.085	0.257
	(0.667)	(0.391)	(0.772)	(0.242)	(0.250)	(0.327)	(0.261)	(0.182)	(0.343)	(0.302)	(3.665)	(1.987)	(3.972)	(3.970)	(2.265)	(0.311)
Bin [39,∞)	-10.07	5.872	21.37***	-3.731**	3.754	2.496	1.610	-0.549	5.206**	-2.032						0.0227
	(8.745)	(4.589)	(5.388)	(1.827)	(2.486)	(4.190)	(5.998)	(2.203)	(2.342)	(1.923)						(6.028)
Monsoon prec	0.470***	0.475***	0.185***	0.411***	-0.202***	0.896***	0.248***	0.381***	-0.0000934	0.472***	0.202***	0.317***	0.185***	0.160***	0.211***	0.0609*
	(0.118)	(0.110)	(0.0649)	(0.0896)	(0.0616)	(0.134)	(0.0583)	(0.0412)	(0.0657)	(0.112)	(0.0349)	(0.0380)	(0.0452)	(0.0303)	(0.0243)	(0.0366)
Monsoon prec^2	-0.143***	-0.165***	-0.0390**	-0.129***	0.0294	-0.444***	-0.0776***	-0.0829***	-0.0392*	-0.228***	-0.0191*	-0.0559***	-0.0355***	-0.0173*	-0.0329***	-0.0249**
	(0.0461)	(0.0469)	(0.0185)	(0.0330)	(0.0220)	(0.0641)	(0.0214)	(0.0120)	(0.0208)	(0.0508)	(0.0112)	(0.0123)	(0.0117)	(0.00921)	(0.00797)	(0.0114)
Rabi prec											-0.118	-0.524***	-0.214**	-0.662***	-0.319***	
											(0.105)	(0.0845)	(0.0906)	(0.0847)	(0.0596)	
Rabi prec^2											0.214*	0.200***	0.0845	0.432***	0.162***	
											(0.119)	(0.0707)	(0.0849)	(0.0868)	(0.0539)	
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 without 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 A.5: Yield-weather estimation – decile temperature bins including HDD variable

	(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) Linseed	(14) Rapeseed/Mustard	(15) Wheat	(16) Sugarcane
Bin [min,p10)	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***
	(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***
	(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 [p20,p30)	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
	(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 [p40,p50)	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
	(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 [p50,p60)	-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
	(0.293)	(0.235)	(0.189)	(0.165)	(0.129)	(0.156)	(0.166)	(0.110)	(0.140)	(0.138)	(0.162)	(0.117)	(0.180)	(0.142)	(0.100)	(0.166)
Bin [p60,p70)	-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**
	(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)
Bin [p70,p80)	-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***
	(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)
Bin [p80,p90)	-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**
	(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)
Bin [p90,max]	-0.309	-0.805***	-0.443	-0.736***	-1.336***	-1.167***	-0.668***	-0.466***	-1.553***	-1.458***	0.326	-0.696***	-0.0551	-0.259	-0.290**	-0.554***
	(0.291)	(0.234)	(0.330)	(0.196)	(0.173)	(0.281)	(0.185)	(0.147)	(0.281)	(0.241)	(0.240)	(0.183)	(0.265)	(0.213)	(0.132)	(0.179)
High degree days	-0.000516	0.0000968	0.0000688	0.000519*	0.000317	-0.000136	0.000128	0.000169	0.000517	0.000389	-0.00119***	-0.000293	-0.00122***	-0.000685**	-0.000510***	0.0000442
	(0.000397)	(0.000384)	(0.000623)	(0.000314)	(0.000270)	(0.000445)	(0.000252)	(0.000196)	(0.000401)	(0.000303)	(0.000344)	(0.000225)	(0.000374)	(0.000271)	(0.000174)	(0.000114)
Monsoon prec	0.451***	0.450***	0.185***	0.386***	-0.208***	0.854***	0.252***	0.366***	-0.0232	0.435***	0.199***	0.325***	0.189***	0.172***	0.216***	0.0611
	(0.116)	(0.110)	(0.0663)	(0.0873)	(0.0621)	(0.133)	(0.0596)	(0.0404)	(0.0637)	(0.107)	(0.0346)	(0.0363)	(0.0448)	(0.0314)	(0.0240)	(0.0379)
Monsoon prec^2	-0.137***	-0.160***	-0.0400**	-0.120***	0.0310	-0.428***	-0.0796***	-0.0804***	-0.0335*	-0.213***	-0.0138	-0.0549***	-0.0346***	-0.0164*	-0.0321***	-0.0268**
	(0.0459)	(0.0467)	(0.0189)	(0.0318)	(0.0220)	(0.0637)	(0.0220)	(0.0116)	(0.0200)	(0.0485)	(0.0112)	(0.0114)	(0.0117)	(0.00968)	(0.00771)	(0.0118)
Rabi prec											-0.161	-0.516***	-0.216**	-0.644***	-0.312***	
											(0.105)	(0.0879)	(0.0882)	(0.0971)	(0.0649)	
Rabi prec^2											0.395***	0.219***	0.113	0.458***	0.186***	
											(0.132)	(0.0733)	(0.0777)	(0.104)	(0.0657)	
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 without additional precipitation variables*. Definitions of the variables may be found in Section 3. See also Appendix Table A.2 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 A.6: Yield-weather estimation – linear degree days

	(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) Linseed	(14) Rapeseed/Mustard	(15) Wheat	(16) Sugarcane
Degree days	-0.000484***	-0.000477***	-0.000495***	-0.000392***	-0.000915***	-0.00106***	-0.000613***	-0.000386***	-0.00108***	-0.000923***	-0.000295***	-0.000510***	-0.000637***	-0.000477***	-0.000456***	-0.000219***
	(0.000174)	(0.0000901)	(0.000121)	(0.0000772)	(0.0000756)	(0.000130)	(0.0000773)	(0.0000567)	(0.000133)	(0.000121)	(0.0000716)	(0.0000648)	(0.0000933)	(0.0000674)	(0.0000517)	(0.0000291)
Monsoon prec	0.469***	0.510***	0.244***	0.506***	-0.125**	1.000***	0.326***	0.470***	0.0729	0.629***	0.222***	0.340***	0.221***	0.189***	0.242***	0.0950**
	(0.113)	(0.112)	(0.0650)	(0.0947)	(0.0619)	(0.130)	(0.0596)	(0.0439)	(0.0664)	(0.112)	(0.0351)	(0.0369)	(0.0453)	(0.0311)	(0.0244)	(0.0374)
Monsoon prec^2	-0.140***	-0.172***	-0.0525***	-0.151***	0.00940	-0.481***	-0.0925***	-0.0981***	-0.0606***	-0.273***	-0.0195*	-0.0584***	-0.0451***	-0.0213**	-0.0389***	-0.0328***
	(0.0443)	(0.0478)	(0.0190)	(0.0351)	(0.0233)	(0.0642)	(0.0224)	(0.0128)	(0.0223)	(0.0530)	(0.0112)	(0.0117)	(0.0120)	(0.00955)	(0.00795)	(0.0117)
Rabi prec											-0.0478	-0.410***	-0.0300	-0.508***	-0.167***	
											(0.0941)	(0.0853)	(0.0817)	(0.0846)	(0.0591)	
Rabi prec^2											0.310**	0.160**	-0.00646	0.384***	0.114**	
											(0.125)	(0.0721)	(0.0833)	(0.0904)	(0.0576)	
Observations	3683	5537	3905	8690	10186	6985	9485	10685	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 (2) using *degree days without 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 A.7: Yield-weather estimation – quadratic degree days

	(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) Linseed	(14) Rapeseed/Mustard	(15) Wheat	(16) Sugarcane
Degree days	-0.000370	-0.000812***	-0.000488**	0.000192	0.000118	0.000571*	-0.000694***	-0.000748***	0.00108***	0.000640***	-0.000509**	0.000151	-0.0000351	-0.0000478	-0.0000816	-0.000215*
	(0.000314)	(0.000211)	(0.000234)	(0.000195)	(0.000184)	(0.000295)	(0.000210)	(0.000171)	(0.000272)	(0.000218)	(0.000216)	(0.000188)	(0.000197)	(0.000173)	(0.000121)	(0.000114)
Degree days^2	-9.33e-08	0.000000256	-8.11e-09	-0.000000468***	-0.000000845***	-0.00000121***	6.75e-08	0.000000298**	-0.00000173***	-0.00000123***	0.000000399	-0.000000954***	-0.000000914***	-0.000000697**	-0.000000560***	-2.05e-09
	(0.000000295)	(0.000000164)	(0.000000297)	(0.000000171)	(0.000000152)	(0.000000260)	(0.000000189)	(0.000000120)	(0.000000265)	(0.000000201)	(0.000000429)	(0.000000266)	(0.000000288)	(0.000000273)	(0.000000194)	(5.10e-08)
Monsoon prec	0.465***	0.527***	0.244***	0.490***	-0.141**	0.841***	0.328***	0.475***	0.0378	0.513***	0.223***	0.337***	0.218***	0.187***	0.240***	0.0949**
	(0.113)	(0.110)	(0.0647)	(0.0922)	(0.0581)	(0.131)	(0.0615)	(0.0447)	(0.0616)	(0.103)	(0.0351)	(0.0370)	(0.0451)	(0.0311)	(0.0244)	(0.0374)
Monsoon prec^2	-0.138***	-0.183***	-0.0525***	-0.140***	0.0214	-0.383***	-0.0940***	-0.101***	-0.0348*	-0.204***	-0.0197*	-0.0577***	-0.0446***	-0.0209**	-0.0387***	-0.0328***
	(0.0439)	(0.0470)	(0.0188)	(0.0336)	(0.0209)	(0.0622)	(0.0235)	(0.0133)	(0.0189)	(0.0472)	(0.0112)	(0.0116)	(0.0118)	(0.00958)	(0.00795)	(0.0117)
Rabi prec											-0.0604	-0.383***	-0.00268	-0.485***	-0.150**	
											(0.0943)	(0.0849)	(0.0819)	(0.0854)	(0.0593)	
Rabi prec^2											0.318**	0.146**	-0.0102	0.376***	0.110*	
											(0.125)	(0.0730)	(0.0831)	(0.0905)	(0.0575)	
Observations	3683	5537	3905	8690	10186	6985	9485	10685	9633	8044	5589	9857	5010	8655	9806	9521
R^2	0.697	0.611	0.782	0.526	0.676	0.709	0.596	0.806	0.632	0.605	0.752	0.561	0.683	0.726	0.834	0.819

This table shows the coefficients estimated from equation (2) using *a quadratic in degree days without 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 A.8: Yield-weather estimation – decile temperature bins 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) Linseed	(14) Rapeseed/Mustard	(15) Wheat	(16) Sugarcane
Bin [min,p10)	0.348	-0.533**	0.363**	0.157	0.465***	0.682***	0.782***	0.518***	0.656***	0.763***	0.386***	-0.0694	0.434**	0.237*	0.533***	1.071***
	(0.271)	(0.256)	(0.148)	(0.164)	(0.152)	(0.186)	(0.150)	(0.126)	(0.152)	(0.182)	(0.146)	(0.184)	(0.185)	(0.142)	(0.106)	(0.211)
Bin [p10,p20)	0.0225	-0.131	0.205	-0.0575	0.191	0.394***	0.466***	0.271**	0.159	0.195	0.209	-0.0626	0.578***	0.434***	0.454***	0.830***
	(0.226)	(0.210)	(0.129)	(0.115)	(0.117)	(0.139)	(0.119)	(0.105)	(0.133)	(0.124)	(0.134)	(0.118)	(0.159)	(0.111)	(0.0803)	(0.155)
Bin [p20,p30)	0.155	-0.291	0.00606	-0.210	-0.0679	0.282	0.371**	0.0761	0.0614	-0.229	0.0296	-0.0178	0.151	0.180	0.136	0.234
	(0.297)	(0.234)	(0.152)	(0.137)	(0.147)	(0.185)	(0.155)	(0.127)	(0.146)	(0.161)	(0.150)	(0.116)	(0.178)	(0.131)	(0.0859)	(0.162)
Bin [p40,p50)	0.292	-0.444*	0.189	-0.170	-0.112	0.170	-0.0460	-0.219*	0.232	-0.183	-0.203	0.0623	0.129	0.119	0.00784	-0.0177
	(0.313)	(0.242)	(0.195)	(0.147)	(0.186)	(0.202)	(0.160)	(0.123)	(0.188)	(0.189)	(0.160)	(0.145)	(0.186)	(0.160)	(0.114)	(0.205)
Bin [p50,p60)	-0.140	-0.269	-0.288	-0.0190	-0.154	0.0759	0.0716	-0.174	-0.0297	0.101	-0.109	-0.298**	-0.210	-0.201	-0.167	-0.0353
	(0.289)	(0.235)	(0.189)	(0.166)	(0.128)	(0.157)	(0.165)	(0.108)	(0.140)	(0.138)	(0.167)	(0.118)	(0.180)	(0.143)	(0.102)	(0.163)
Bin [p60,p70)	-0.0975	-0.573***	-0.0398	-0.101	-0.265*	-0.0396	-0.419***	-0.426***	0.224	-0.428**	-0.421**	-0.634***	-0.748***	-0.431**	-0.399***	-0.325**
	(0.304)	(0.213)	(0.210)	(0.133)	(0.143)	(0.177)	(0.130)	(0.126)	(0.176)	(0.190)	(0.191)	(0.147)	(0.226)	(0.169)	(0.112)	(0.163)
Bin [p70,p80)	-0.276	-0.623***	-0.137	-0.639***	-0.328**	-0.415**	-0.333**	-0.529***	-0.492***	-0.746***	-0.345**	-0.573***	-0.516**	-0.491***	-0.406***	-0.475***
	(0.264)	(0.209)	(0.195)	(0.164)	(0.139)	(0.183)	(0.149)	(0.118)	(0.169)	(0.181)	(0.161)	(0.125)	(0.199)	(0.155)	(0.123)	(0.162)
Bin [p80,p90)	-0.203	-0.855***	-0.475**	-0.727***	-0.728***	-0.573***	-0.349***	-0.612***	-0.632***	-0.943***	-0.128	-0.746***	-0.596***	-0.455***	-0.448***	-0.355**
	(0.237)	(0.189)	(0.188)	(0.120)	(0.136)	(0.165)	(0.126)	(0.124)	(0.146)	(0.172)	(0.159)	(0.125)	(0.192)	(0.150)	(0.107)	(0.179)
Bin [p90,max]	-0.293	-0.591***	-0.201	-0.347**	-0.779***	-0.871***	-0.395**	0.0623	-0.969***	-1.048***	-0.337**	-0.899***	-0.825***	-0.696***	-0.613***	-0.327**
	(0.288)	(0.205)	(0.208)	(0.150)	(0.144)	(0.214)	(0.158)	(0.116)	(0.206)	(0.197)	(0.167)	(0.135)	(0.167)	(0.149)	(0.110)	(0.148)
Monsoon prec	0.376***	0.427***	0.135**	0.320***	-0.298***	0.744***	0.233***	0.313***	-0.114*	0.395***	0.208***	0.326***	0.199***	0.172***	0.218***	0.0221
	(0.122)	(0.115)	(0.0658)	(0.0882)	(0.0583)	(0.134)	(0.0636)	(0.0393)	(0.0655)	(0.105)	(0.0352)	(0.0361)	(0.0441)	(0.0316)	(0.0239)	(0.0390)
Monsoon prec^2	-0.118***	-0.151***	-0.0241	-0.112***	0.0422**	-0.405***	-0.0713***	-0.0641***	-0.0228	-0.204***	-0.0156	-0.0536***	-0.0341***	-0.0146	-0.0314***	-0.0252**
	(0.0438)	(0.0471)	(0.0190)	(0.0314)	(0.0211)	(0.0627)	(0.0215)	(0.0111)	(0.0195)	(0.0479)	(0.0113)	(0.0112)	(0.0117)	(0.00974)	(0.00771)	(0.0114)
Rainy days	0.00209	0.00204**	0.00340***	0.00260***	0.00528***	0.00451***	0.00235***	0.00509***	0.00480***	0.00295***	0.000383	-0.00156***	-0.00175***	-0.00165***	-0.00104***	0.00128***
	(0.00129)	(0.000981)	(0.000939)	(0.000671)	(0.000564)	(0.000789)	(0.000687)	(0.000470)	(0.000754)	(0.000739)	(0.000459)	(0.000412)	(0.000625)	(0.000479)	(0.000302)	(0.000266)
Heavy rain	-0.00243	-0.00395	-0.0121	0.0134**	0.0133*	0.00746	-0.00896	-0.0120***	0.0185**	-0.000423						0.00940
	(0.0120)	(0.0107)	(0.00773)	(0.00622)	(0.00675)	(0.00900)	(0.00629)	(0.00450)	(0.00835)	(0.00917)						(0.00652)
Dry spell	-0.00467***	0.000436	-0.000837	-0.000597	0.000163	-0.00122	0.000538	0.00109**	0.00129	0.00128	0.000147	0.000315*	0.000490*	0.000335**	-0.0000914	0.000174
	(0.00161)	(0.00123)	(0.00191)	(0.000608)	(0.000709)	(0.000784)	(0.000858)	(0.000538)	(0.00104)	(0.000816)	(0.000201)	(0.000173)	(0.000260)	(0.000138)	(0.000111)	(0.000126)
Rabi prec											-0.230**	-0.291***	0.0651	-0.427***	-0.203***	
											(0.107)	(0.0926)	(0.124)	(0.100)	(0.0648)	
Rabi prec^2											0.434***	0.101	-0.0701	0.342***	0.133**	
											(0.135)	(0.0764)	(0.0951)	(0.0993)	(0.0643)	
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 Appendix Table A.2 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 A.9: Yield-weather estimation – linear 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) Linseed	(14) Rapeseed/Mustard	(15) Wheat	(16) Sugarcane
Degree days	-0.000361**	-0.000348***	-0.000309**	-0.000240***	-0.000558***	-0.000762***	-0.000462***	-0.0000652	-0.000729***	-0.000691***	-0.000279***	-0.000545***	-0.000683***	-0.000521***	-0.000483***	-0.000150***
	(0.000172)	(0.000123)	(0.000136)	(0.0000804)	(0.0000779)	(0.000127)	(0.0000823)	(0.0000591)	(0.000126)	(0.000114)	(0.0000733)	(0.0000672)	(0.0000930)	(0.0000696)	(0.0000528)	(0.0000329)
Monsoon prec	0.406***	0.505***	0.188***	0.454***	-0.211***	0.907***	0.317***	0.436***	-0.0117	0.597***	0.221***	0.337***	0.217***	0.185***	0.241***	0.0572
	(0.120)	(0.118)	(0.0629)	(0.0971)	(0.0577)	(0.134)	(0.0631)	(0.0430)	(0.0679)	(0.108)	(0.0349)	(0.0367)	(0.0445)	(0.0313)	(0.0243)	(0.0385)
Monsoon prec^2	-0.122***	-0.165***	-0.0342*	-0.143***	0.0232	-0.457***	-0.0830***	-0.0821***	-0.0482**	-0.258***	-0.0198*	-0.0565***	-0.0420***	-0.0188*	-0.0380***	-0.0306***
	(0.0425)	(0.0488)	(0.0189)	(0.0350)	(0.0222)	(0.0643)	(0.0218)	(0.0122)	(0.0215)	(0.0520)	(0.0112)	(0.0115)	(0.0117)	(0.00965)	(0.00795)	(0.0113)
Rainy days	0.00185	0.00162	0.00353***	0.00242***	0.00548***	0.00443***	0.00190***	0.00437***	0.00535***	0.00314***	0.000453	-0.00166***	-0.00190***	-0.00167***	-0.00100***	0.00144***
	(0.00130)	(0.00100)	(0.000946)	(0.000709)	(0.000542)	(0.000728)	(0.000668)	(0.000473)	(0.000802)	(0.000731)	(0.000453)	(0.000417)	(0.000636)	(0.000484)	(0.000315)	(0.000263)
Heavy rain	-0.00420	-0.00840	-0.0143*	0.00794	0.00986	0.00129	-0.0142**	-0.0185***	0.0149*	-0.00924						0.00703
	(0.0122)	(0.0106)	(0.00773)	(0.00622)	(0.00667)	(0.00897)	(0.00633)	(0.00473)	(0.00836)	(0.00917)						(0.00664)
Dry spell	-0.00470***	0.000659	-0.000607	-0.000515	0.000297	-0.00100	0.000269	0.000712	0.00179*	0.00151*	0.000236	0.000347**	0.000602**	0.000397***	0.00000431	0.000208
	(0.00162)	(0.00122)	(0.00192)	(0.000582)	(0.000688)	(0.000757)	(0.000859)	(0.000573)	(0.00104)	(0.000828)	(0.000199)	(0.000167)	(0.000250)	(0.000138)	(0.000109)	(0.000129)
Rabi prec											-0.0823	-0.173*	0.294**	-0.275***	-0.0435	
											(0.0960)	(0.0923)	(0.121)	(0.0894)	(0.0596)	
Rabi prec^2											0.321**	0.0361	-0.213**	0.261***	0.0535	
											(0.127)	(0.0771)	(0.103)	(0.0861)	(0.0563)	
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 (2) 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 A.10: Yield-weather estimation – three-degree temperature bins 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) Linseed	(14) Rapeseed/Mustard	(15) Wheat	(16) Sugarcane
Bin $(-\infty,0)$			1.502	3.579	-0.622	-2.919**	-1.489*	-0.244	-0.638		2.624***	3.692**	6.568***	4.671***	3.554***	-5.875***
			(1.066)	(3.961)	(0.751)	(1.256)	(0.870)	(0.428)	(0.941)		(0.438)	(1.491)	(1.314)	(0.643)	(0.462)	(1.891)
Bin [0,3)			2.347**	3.899	0.00624	-11.57	-0.287	0.405	-0.794		1.598***	2.944***	4.092***	3.525***	2.275***	-1.073
			(1.028)	(3.741)	(0.733)	(7.463)	(0.841)	(0.458)	(1.019)		(0.294)	(1.030)	(0.732)	(0.862)	(0.280)	(1.360)
Bin [3,6)			2.030**	1.504	0.246		-1.035	-0.0729	0.943		2.088***	2.858***	2.530***	3.257***	2.158***	-4.624*
			(0.975)	(3.816)	(0.697)		(1.277)	(0.424)	(0.657)		(0.287)	(1.037)	(0.929)	(0.846)	(0.289)	(2.560)
Bin [6,9)			1.569	1.538	0.456	1.516	-0.392	0.00338	-0.0238		1.364***	2.273***	0.828	2.395***	1.597***	-2.298**
			(0.968)	(3.916)	(0.697)	(2.911)	(0.843)	(0.303)	(0.749)		(0.202)	(0.593)	(0.570)	(0.586)	(0.244)	(1.077)
Bin [9,12)		-0.469	2.195**	0.517	0.214	-2.665***	-0.167	-0.0825	0.323	-152.8***	0.528***	0.200	0.293	-0.255	0.546***	1.147***
		(1.127)	(0.845)	(2.272)	(0.645)	(0.753)	(0.999)	(0.290)	(0.598)	(23.29)	(0.171)	(0.296)	(0.364)	(0.246)	(0.140)	(0.315)
Bin [12,15)		3.243*	2.041***	1.741*	0.815	-1.327	-0.0587	0.149	-0.255	10.56***	0.439***	-0.0306	0.543***	0.376***	0.582***	1.067***
		(1.925)	(0.727)	(1.012)	(0.769)	(1.362)	(0.681)	(0.211)	(0.717)	(0.837)	(0.116)	(0.162)	(0.170)	(0.130)	(0.0824)	(0.191)
Bin [15,18)	-33.82	2.704*	1.194***	0.702	0.249	0.402	-0.0887	0.126	-0.325	6.234**	0.268**	0.0246	0.538***	0.350***	0.464***	0.983***
	(27.76)	(1.407)	(0.457)	(0.746)	(0.413)	(0.664)	(0.567)	(0.225)	(0.483)	(2.548)	(0.113)	(0.106)	(0.140)	(0.101)	(0.0724)	(0.158)
Bin [18,21)	0.127	-1.247***	0.330*	0.0860	0.414**	0.462	0.467**	0.368***	0.350	0.595***	0.142	0.198**	0.0988	0.133	0.135**	0.292**
	(0.284)	(0.316)	(0.180)	(0.129)	(0.166)	(0.317)	(0.183)	(0.102)	(0.214)	(0.193)	(0.0972)	(0.0890)	(0.111)	(0.0955)	(0.0666)	(0.133)
Bin [24,27)	-0.0101	-0.0309	-0.219***	-0.0534	-0.300***	-0.281***	-0.486***	-0.362***	-0.144*	-0.321***	-0.0943	-0.466***	-0.488***	-0.399***	-0.287***	-0.218**
	(0.129)	(0.131)	(0.0668)	(0.0811)	(0.0824)	(0.0944)	(0.0819)	(0.0767)	(0.0865)	(0.0859)	(0.114)	(0.0942)	(0.137)	(0.0969)	(0.0780)	(0.0949)
Bin [27,30)	-0.306*	-0.469***	-0.445***	-0.409***	-0.607***	-0.715***	-0.890***	-0.814***	-0.461***	-0.866***	0.0769	-0.626***	-0.426***	-0.511***	-0.388***	-0.565***
	(0.181)	(0.161)	(0.106)	(0.0907)	(0.0959)	(0.118)	(0.0915)	(0.0899)	(0.109)	(0.119)	(0.131)	(0.105)	(0.158)	(0.117)	(0.0834)	(0.107)
Bin [30,33)	-0.364	-0.543***	-0.618**	-0.722***	-1.214***	-1.332***	-0.990***	-0.557***	-1.349***	-1.460***	-0.210	-0.769***	-0.674***	-0.699***	-0.437***	-0.394***
	(0.268)	(0.203)	(0.252)	(0.144)	(0.141)	(0.206)	(0.134)	(0.107)	(0.206)	(0.191)	(0.143)	(0.151)	(0.220)	(0.145)	(0.103)	(0.142)
Bin [33,36)	-0.227	-0.221	-0.309	0.0404	-0.744***	-1.055***	-0.997***	0.0352	-0.893***	-1.038***	-1.265***	-1.255***	-2.095***	-1.508***	-1.204***	-0.621***
	(0.416)	(0.240)	(0.246)	(0.156)	(0.148)	(0.224)	(0.197)	(0.133)	(0.246)	(0.221)	(0.422)	(0.286)	(0.403)	(0.295)	(0.227)	(0.165)
Bin [36,39)	-1.598**	-0.612	-0.455	0.0160	-0.814***	-1.384***	-0.530**	0.0500	-0.816**	-0.928***	2.567	-2.241	4.227	-3.879	-2.145	0.590*
	(0.643)	(0.420)	(0.801)	(0.257)	(0.247)	(0.314)	(0.261)	(0.189)	(0.354)	(0.289)	(3.640)	(1.957)	(3.958)	(3.975)	(2.263)	(0.314)
Bin [39,∞)	-7.666	5.877	14.38**	-3.582*	3.775	2.787	1.399	-0.778	5.550**	-1.821						-0.740
	(7.953)	(4.451)	(5.920)	(1.903)	(2.548)	(4.097)	(5.904)	(2.127)	(2.311)	(1.891)						(6.114)
Monsoon prec	0.401***	0.447***	0.135**	0.333***	-0.304***	0.777***	0.225***	0.320***	-0.102	0.418***	0.201***	0.316***	0.183***	0.155***	0.211***	0.0196
	(0.124)	(0.116)	(0.0638)	(0.0917)	(0.0579)	(0.136)	(0.0621)	(0.0396)	(0.0689)	(0.109)	(0.0347)	(0.0378)	(0.0447)	(0.0304)	(0.0243)	(0.0379)
Monsoon prec^2	-0.123***	-0.155***	-0.0238	-0.118***	0.0431**	-0.417***	-0.0684***	-0.0649***	-0.0255	-0.214***	-0.0189*	-0.0543***	-0.0332***	-0.0144	-0.0317***	-0.0231**
	(0.0437)	(0.0474)	(0.0186)	(0.0326)	(0.0210)	(0.0637)	(0.0209)	(0.0113)	(0.0203)	(0.0500)	(0.0112)	(0.0122)	(0.0116)	(0.00925)	(0.00797)	(0.0110)
Rainy days	0.00216*	0.00218**	0.00328***	0.00297***	0.00558***	0.00482***	0.00238***	0.00545***	0.00533***	0.00342***	-0.000115	-0.00191***	-0.00196***	-0.00209***	-0.00129***	0.00140***
	(0.00131)	(0.000970)	(0.000969)	(0.000696)	(0.000561)	(0.000709)	(0.000692)	(0.000507)	(0.000779)	(0.000746)	(0.000441)	(0.000387)	(0.000646)	(0.000453)	(0.000301)	(0.000257)
Heavy rain	-0.00408	-0.00405	-0.0109	0.0139**	0.0139**	0.00621	-0.00867	-0.0121***	0.0179**	-0.000655						0.00944
	(0.0121)	(0.0108)	(0.00777)	(0.00633)	(0.00679)	(0.00880)	(0.00623)	(0.00448)	(0.00845)	(0.00924)						(0.00690)
Dry spell	-0.00454***	0.000376	-0.000891	-0.000500	0.000231	-0.00102	0.000513	0.00125**	0.00152	0.00146*	0.000131	0.000293*	0.000476*	0.000300**	-0.000110	0.000188
	(0.00156)	(0.00120)	(0.00193)	(0.000593)	(0.000696)	(0.000765)	(0.000861)	(0.000547)	(0.00101)	(0.000820)	(0.000202)	(0.000172)	(0.000255)	(0.000137)	(0.000111)	(0.000128)
Rabi prec											-0.0871	-0.253***	0.124	-0.384***	-0.171***	
											(0.105)	(0.0934)	(0.134)	(0.0959)	(0.0623)	
Rabi prec^2											0.189	0.0583	-0.130	0.284***	0.0903*	
											(0.116)	(0.0746)	(0.108)	(0.0818)	(0.0519)	
Observations	3683	5537	3905	8690	10186	6985	9485	10685	9633	8044	5589	9857	5010	8655	9806	9521
R^2	0.700	0.614	0.786	0.533	0.680	0.709	0.601	0.814	0.630	0.606	0.756	0.566	0.689	0.733	0.837	0.822

This table shows the coefficients estimated from equation (1) using *three-degree temperature bins 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 A.11: Yield-weather estimation – decile temperature bins including HDD variable 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) Linseed	(14) Rapeseed/Mustard	(15) Wheat	(16) Sugarcane
Bin [min,p10)	0.345	-0.531**	0.363**	0.148	0.458***	0.675***	0.778***	0.512***	0.647***	0.755***	0.383***	-0.0733	0.430**	0.230	0.530***	1.060***
	(0.271)	(0.256)	(0.148)	(0.163)	(0.151)	(0.185)	(0.150)	(0.126)	(0.152)	(0.180)	(0.146)	(0.184)	(0.187)	(0.141)	(0.106)	(0.209)
Bin [p10,p20)	0.0224	-0.133	0.206	-0.0548	0.194*	0.395***	0.466***	0.272**	0.162	0.200	0.194	-0.0689	0.554***	0.425***	0.447***	0.824***
	(0.226)	(0.210)	(0.129)	(0.116)	(0.118)	(0.140)	(0.119)	(0.105)	(0.133)	(0.124)	(0.133)	(0.119)	(0.160)	(0.110)	(0.0805)	(0.154)
Bin [p20,p30)	0.154	-0.292	0.00626	-0.206	-0.0624	0.284	0.372**	0.0795	0.0667	-0.224	0.0365	-0.0186	0.148	0.181	0.134	0.226
	(0.296)	(0.234)	(0.152)	(0.138)	(0.147)	(0.186)	(0.155)	(0.128)	(0.146)	(0.161)	(0.150)	(0.116)	(0.178)	(0.130)	(0.0858)	(0.162)
Bin [p40,p50)	0.292	-0.446*	0.189	-0.174	-0.115	0.168	-0.0477	-0.222*	0.228	-0.181	-0.163	0.0708	0.158	0.139	0.0216	-0.0187
	(0.313)	(0.242)	(0.195)	(0.147)	(0.185)	(0.203)	(0.160)	(0.122)	(0.188)	(0.189)	(0.161)	(0.145)	(0.187)	(0.161)	(0.114)	(0.205)
Bin [p50,p60)	-0.142	-0.271	-0.288	-0.0269	-0.162	0.0723	0.0677	-0.180*	-0.0380	0.0927	-0.0470	-0.287**	-0.177	-0.175	-0.151	-0.0390
	(0.288)	(0.235)	(0.189)	(0.166)	(0.129)	(0.157)	(0.165)	(0.108)	(0.139)	(0.138)	(0.164)	(0.118)	(0.181)	(0.142)	(0.101)	(0.163)
Bin [p60,p70)	-0.0965	-0.573***	-0.0402	-0.0961	-0.260*	-0.0366	-0.416***	-0.422***	0.229	-0.424**	-0.333*	-0.619***	-0.693***	-0.394**	-0.377***	-0.332**
	(0.304)	(0.213)	(0.210)	(0.132)	(0.144)	(0.177)	(0.130)	(0.127)	(0.176)	(0.190)	(0.189)	(0.148)	(0.227)	(0.172)	(0.113)	(0.163)
Bin [p70,p80)	-0.268	-0.626***	-0.140	-0.650***	-0.336**	-0.418**	-0.337**	-0.535***	-0.504***	-0.750***	-0.233	-0.540***	-0.406*	-0.428***	-0.362***	-0.487***
	(0.264)	(0.208)	(0.196)	(0.164)	(0.139)	(0.183)	(0.148)	(0.118)	(0.168)	(0.181)	(0.158)	(0.125)	(0.206)	(0.157)	(0.122)	(0.163)
Bin [p80,p90)	-0.223	-0.835***	-0.475**	-0.662***	-0.667***	-0.545***	-0.323**	-0.560***	-0.550***	-0.888***	-0.138	-0.742***	-0.599***	-0.463***	-0.447***	-0.349*
	(0.239)	(0.195)	(0.187)	(0.120)	(0.138)	(0.160)	(0.128)	(0.124)	(0.143)	(0.173)	(0.159)	(0.125)	(0.192)	(0.150)	(0.107)	(0.179)
Bin [p90,max]	-0.183	-0.668***	-0.222	-0.615***	-1.054***	-0.980***	-0.513***	-0.158	-1.290***	-1.267***	0.341	-0.675***	-0.0205	-0.265	-0.304**	-0.452**
	(0.298)	(0.252)	(0.330)	(0.198)	(0.179)	(0.292)	(0.188)	(0.144)	(0.278)	(0.244)	(0.245)	(0.186)	(0.268)	(0.215)	(0.133)	(0.178)
High degree days	-0.000385	0.000245	0.0000694	0.000784**	0.000766***	0.000313	0.000321	0.000607***	0.000942**	0.000619**	-0.00119***	-0.000423*	-0.00142***	-0.000811***	-0.000575***	0.000133
	(0.000396)	(0.000388)	(0.000621)	(0.000323)	(0.000270)	(0.000434)	(0.000259)	(0.000202)	(0.000411)	(0.000304)	(0.000346)	(0.000229)	(0.000386)	(0.000274)	(0.000176)	(0.000114)
Monsoon prec	0.381***	0.424***	0.135**	0.310***	-0.307***	0.738***	0.229***	0.307***	-0.126*	0.384***	0.199***	0.325***	0.187***	0.168***	0.216***	0.0210
	(0.123)	(0.116)	(0.0653)	(0.0884)	(0.0583)	(0.135)	(0.0637)	(0.0392)	(0.0663)	(0.106)	(0.0345)	(0.0361)	(0.0444)	(0.0317)	(0.0240)	(0.0389)
Monsoon prec^2	-0.119***	-0.151***	-0.0240	-0.110***	0.0440**	-0.403***	-0.0706***	-0.0630***	-0.0202	-0.201***	-0.0141	-0.0535***	-0.0324***	-0.0141	-0.0311***	-0.0251**
	(0.0439)	(0.0471)	(0.0189)	(0.0313)	(0.0210)	(0.0633)	(0.0215)	(0.0110)	(0.0195)	(0.0479)	(0.0112)	(0.0112)	(0.0117)	(0.00978)	(0.00771)	(0.0114)
Rainy days	0.00202	0.00213**	0.00340***	0.00291***	0.00553***	0.00467***	0.00247***	0.00534***	0.00519***	0.00322***	0.000264	-0.00162***	-0.00199***	-0.00173***	-0.00111***	0.00131***
	(0.00129)	(0.000984)	(0.000941)	(0.000683)	(0.000560)	(0.000711)	(0.000708)	(0.000495)	(0.000771)	(0.000725)	(0.000459)	(0.000417)	(0.000640)	(0.000485)	(0.000306)	(0.000267)
Heavy rain	-0.00277	-0.00365	-0.0121	0.0142**	0.0139**	0.00786	-0.00859	-0.0115**	0.0193**	0.000421						0.00954
	(0.0120)	(0.0108)	(0.00772)	(0.00622)	(0.00675)	(0.00893)	(0.00627)	(0.00451)	(0.00836)	(0.00923)						(0.00652)
Dry spell	-0.00469***	0.000439	-0.000830	-0.000574	0.000192	-0.00122	0.000561	0.00115**	0.00135	0.00131	0.000152	0.000314*	0.000499*	0.000342**	-0.0000917	0.000178
	(0.00160)	(0.00123)	(0.00190)	(0.000605)	(0.000702)	(0.000785)	(0.000859)	(0.000540)	(0.00103)	(0.000815)	(0.000202)	(0.000173)	(0.000260)	(0.000138)	(0.000112)	(0.000126)
Rabi prec											-0.177*	-0.277***	0.127	-0.402***	-0.184***	
											(0.105)	(0.0937)	(0.128)	(0.101)	(0.0652)	
Rabi prec^2											0.397***	0.0956	-0.104	0.330***	0.125*	
											(0.132)	(0.0766)	(0.0969)	(0.0988)	(0.0641)	
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	0.534	0.680	0.709	0.601	0.814	0.631	0.607	0.754	0.563	0.688	0.728	0.836	0.821

This table shows the coefficients estimated from equation (1) using *decile temperature bins and an HDD variable with additional precipitation variables.* Definitions of the variables may be found in Section 3. See also Appendix Table A.2 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 A.12: Yield-weather estimation – quadratic 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) Linseed	(14) Rapeseed/Mustard	(15) Wheat	(16) Sugarcane
Degree days	-0.000310	-0.000678***	-0.000313	0.000314	0.000434**	0.000875***	-0.000540***	-0.000443***	0.00135***	0.000893***	-0.000486**	0.0000650	-0.000145	-0.000160	-0.000136	-0.000121
	(0.000332)	(0.000231)	(0.000228)	(0.000198)	(0.000183)	(0.000313)	(0.000207)	(0.000166)	(0.000285)	(0.000230)	(0.000218)	(0.000193)	(0.000194)	(0.000180)	(0.000122)	(0.000115)
Degree days^2	-4.21e-08	0.000000251	5.09e-09	-0.000000445**	-0.000000813***	-0.00000120***	6.56e-08	0.000000312***	-0.00000168***	-0.00000125***	0.000000384	-0.000000875***	-0.000000814***	-0.000000582**	-0.000000515***	-1.33e-08
	(0.000000301)	(0.000000165)	(0.000000290)	(0.000000172)	(0.000000153)	(0.000000258)	(0.000000189)	(0.000000119)	(0.000000263)	(0.000000202)	(0.000000429)	(0.000000270)	(0.000000283)	(0.000000278)	(0.000000196)	(5.07e-08)
Monsoon prec	0.405***	0.517***	0.188***	0.442***	-0.223***	0.759***	0.319***	0.439***	-0.0323	0.490***	0.222***	0.335***	0.215***	0.183***	0.240***	0.0568
	(0.120)	(0.116)	(0.0626)	(0.0948)	(0.0547)	(0.136)	(0.0648)	(0.0438)	(0.0640)	(0.101)	(0.0349)	(0.0368)	(0.0444)	(0.0313)	(0.0244)	(0.0386)
Monsoon prec^2	-0.121***	-0.176***	-0.0342*	-0.133***	0.0352*	-0.357***	-0.0844***	-0.0855***	-0.0233	-0.189***	-0.0200*	-0.0560***	-0.0417***	-0.0185*	-0.0379***	-0.0303***
	(0.0423)	(0.0480)	(0.0186)	(0.0336)	(0.0201)	(0.0633)	(0.0228)	(0.0126)	(0.0187)	(0.0460)	(0.0112)	(0.0115)	(0.0116)	(0.00967)	(0.00795)	(0.0112)
Rainy days	0.00184	0.00163	0.00353***	0.00238***	0.00539***	0.00447***	0.00191***	0.00441***	0.00497***	0.00309***	0.000430	-0.00151***	-0.00183***	-0.00158***	-0.000927***	0.00144***
	(0.00130)	(0.00100)	(0.000946)	(0.000715)	(0.000549)	(0.000739)	(0.000670)	(0.000473)	(0.000789)	(0.000742)	(0.000455)	(0.000419)	(0.000636)	(0.000491)	(0.000314)	(0.000263)
Heavy rain	-0.00445	-0.00699	-0.0143*	0.00637	0.00713	-0.00536	-0.0140**	-0.0177***	0.00894	-0.0140						0.00697
	(0.0121)	(0.0107)	(0.00774)	(0.00615)	(0.00662)	(0.00868)	(0.00622)	(0.00471)	(0.00823)	(0.00878)						(0.00659)
Dry spell	-0.00469***	0.000574	-0.000608	-0.000434	0.000321	-0.000662	0.000270	0.000698	0.00177*	0.00181**	0.000233	0.000355**	0.000592**	0.000404***	0.00000824	0.000209
	(0.00166)	(0.00122)	(0.00190)	(0.000586)	(0.000678)	(0.000747)	(0.000859)	(0.000571)	(0.00102)	(0.000845)	(0.000199)	(0.000169)	(0.000249)	(0.000138)	(0.000109)	(0.000129)
Rabi prec											-0.0917	-0.165*	0.307**	-0.266***	-0.0367	
											(0.0963)	(0.0918)	(0.121)	(0.0895)	(0.0595)	
Rabi prec^2											0.327**	0.0318	-0.210**	0.259***	0.0543	
											(0.127)	(0.0783)	(0.103)	(0.0864)	(0.0562)	
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	0.528	0.680	0.712	0.597	0.809	0.634	0.607	0.752	0.562	0.685	0.727	0.834	0.820

This table shows the coefficients estimated from equation (2) using *a quadratic in 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 $^{***}\,p < 0.01$.