Essays in Environmental and Applied Economics

by

Fabien Forge

A thesis submitted in partial fulfillment of the requirements for the
Doctorate in Philosophy degree in Economics

Department of Economics
Faculty of Social Sciences
University of Ottawa

Abstract

This dissertation includes three distinct chapters looking at different challenges faced by developing countries. The first chapter examines the situation of farmers under climate change by mapping future climatic conditions onto the distribution of agricultural revenue in India. The second chapter uses the Mexican conditional cash transfer (CCT) program Progresa to investigate the relationship between income, education and fertility. Finally, the third chapter studies the extent to which preferential tariffs extended by OECD countries have helped least developing countries (LDC) diversify their exports. Given the issues explored by these essays, I contribute to several distinct strands of the economic literature, yet each paper is motivated by its policy relevance and is embedded in the issues faced by developing economies.

Declaration of Authorship

I, Fabien Forge, declare that this thesis titled, “Essays in Environmental and Applied Economics” and the work presented in it are my own.

Chapter one of this thesis was done jointly with Dr. Francisco Costa, Dr. Jason Garred and Dr. João Paulo Pessoa. My contribution is equal to theirs.

The third chapter of this thesis was done jointly with Dr. Jason Garred and Kyae Lim Kwon. My contribution is equal to theirs.

I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed: ______________________________

Date: 2020 October 14

General Introduction

This dissertation includes three distinct chapters looking at different challenges faced by developing countries. The first chapter examines the situation of farmers under climate change by mapping future climatic conditions onto the distribution of agricultural revenue in India. The second chapter uses the Mexican conditional cash transfer (CCT) program *Progresa* to investigate the relationship between income, education and fertility. Finally, the third chapter studies the extent to which preferential tariffs extended by OECD countries have helped least developing countries (LDC) diversify their exports. Given the issues explored by these essays, I contribute to several distinct strands of the economic literature, yet each paper is motivated by its policy relevance.

In the first chapter, my coauthors and I investigate the extent to which countries with incomplete insurance markets, such as India, will be hurt by a potential increase in the variability of weather as a consequence of climate change. To do this, we first estimate two sets of models to capture the historical weather-yield relationship using the ERA5 weather and ICRISAT Village Dynamics in South Asia (VDSA) datasets. We then use these estimates to predict past and future distributions of agricultural output in India. Key to our strategy is the use of the multi-run CESM-LENS dataset that provides us with 400 weather realizations per decade from 1920 to 2099. Our results suggest that although model specifications matter when predicting the second moment of the distribution of agricultural output, our various models yield the same overall conclusion: the impact of climate change on insurable risk is not of first order importance in India. We show instead that uninsurable downside-risk is the main source of concern. Rare extreme weather events resulting in extremely poor revenues prior to climate change are projected to become the new normal by the end of the 21^{st} century. We also check whether our main conclusions about revenue are the result of factors such as crop choice and and prices being held constant in our analysis. To do so, we embed our projections in the general equilibrium portfolio model developed by Allen and Atkin (2016), which incorporates optimal crop choice adaptation, trade and endogenous prices. Our conterfactual results confirm our main conclusions and suggest welfare losses will mainly be driven by changes in mean returns. Our paper builds on the literature on the consequences of climate change for agriculture in developing countries by being the first to jointly consider both the uninsurable losses and the changing risks facing farmers in a unified framework.

My second chapter explores the relationship between poverty and fertility, which has long been an important issue in development economics since Malthus' seminal work. Specifically, I study how Mexico's CCT program *Progresa* has impacted fertility decisions. To do so, I use a difference-in-difference strategy based on age and the timing and intensity of the treatment. I use the Integrated Public Use Microdata Series (IPUMS) data for Mexico at the municipality level in order to identify women old enough to be enrolled in the program as a parent but not so old that they no longer are able to make fertility decisions. These women, forming the *older cohort*, were potentially treated by receiving cash transfer on the condition that their children would attend school. The second cohort, labelled the *younger cohort*, is composed of women younger than the *older cohort* that first benefited from the program as students and as a result received, on average, extra schooling. They then made fertility decisions facing the same type of treatment as the *older cohort*. Fertility decisions are observed in the 2010 census, 13 years after the program began. My results suggest that even though the cash transfer increased with the number of children, fertility declined for both cohorts in response to the treatment as would be suggested by a quality-quantity trade-off model. My results also suggest that the effects on fertility are mainly driven by the cash-conditionality treatment as I do not observe an additional impact of education on fertility. Because I look at the effect of the program 13 years after it started, I contribute to the literature looking at the long term effects of such programs. My second chapter also speaks to a large literature that links fertility decisions with poverty and explores how this link can evolve over time. My key contribution is to define two cohorts of women that are treated differentially by the program and to show the relative contributions of income and education to fertility decisions.

In the third chapter, my coauthors and I study whether the reduction of import tariffs helped spur new export activities in the least developped countries. To do so, we jointly consider thirteen trade policy reforms by nine OECD members and their associated product-specific tariff cuts for LDCs using trade data for 1996 to 2011. We identify the effect of these reforms on the extensive margin of trade using variation by importer, exporter, product and year in a quadruple-difference framework suggested by a model based on Helpman, Melitz and Rubenstein (2008). We document a positive average effect of these reforms but show that this effect is entirely driven by importer-exporter-product categories with pre-existing export experience as predicted by our framework. Our paper contributes to a surprisingly small literature on the impact of trade preferences for developing countries. Our key contribution is to uncover systematic heterogeneity in the effects of trade preference regimes, in line with our hypothesis that their benefits flow disproportionately to exporters with particularly favourable initial conditions.

Acknowledgment

I was blessed by truly amazing supervisors. It has been a privilege to be supervised and mentored by Jason Garred. His support, guidance and generosity with his time throughout the best part of my PhD helped bring my research and academic profile to levels I could only hope to reach. His incredible kindness was never at the expense of a rigorous truth and I am incredibly thankful for both.

Un grand merci to my other supervisor Abel Brodeur. His sharp insights helped me in more ways than one. He too helped decisively shape not only my research but also the many strategic decisions, small and large, a PhD students needs to make.

I grateful to Carolyn Fischer for opening so many doors for me and for showing such an interest for my research even when it took her outside of the environmental issues I lured her with.

Muito obrigado to Francisco Costa and João Paulo Pessoa for being so welcoming. It was a privilege to work among such amazing researchers.

I would also like to acknowledge thesis committee members for their useful comments and suggestions. I would also like to thank Reed Walker for his advice during my research visit at UC Berkeley.

I want to thank the entire Economic Department of the University of Ottawa and all its professors and staff for their support. A special thanks to the head of the department Catherine Deri-Armstrong for all her help.

This experience would not have been nearly as good without the friendships I made with my fellow graduate students. To Farzaneh Davarzani, Joanne Haddad and Philippe Kabore: it has been a pleasure agonizing over problem sets together, talking about research together, witnessing new life together and I look forward to the next chapter in which we are no longer graduate students. I am also thankful for Julia Mink's comments on my work and look forward to our future collaborations.

I am grateful to the Canada 150 Chair in Climate Economics, Innovation, and Policy for its generous financial support.

As a first generation student to go to university, I want to thank my parents, Jean-Robert Forge and Martine Baulès, for supporting my choice to get a PhD even when it implied doing it so far from home.

Last but not least, I want to thank my wife Marlène Al-Kazzi. Only she could transform the hotel reception by night and semi-pro ballroom dancer by day I was when she met me into the PhD graduate I am about to become. A lot of what I achieved was because of her. Thank you so much for your sacrifices and everlasting support.

à Marlène et Maryam

Contents

1 The Impact of Climate Change on Risk and Return in Indian Agriculture 1
1.1 Introduction 2
1.2 Data 6
1.2.1 Historical weather and yield 6
1.2.2 Weather projection 7
1.3 Constructing projected yields 12
1.3.1 Yield-weather specifications 12
1.3.2 Yield-weather regression results 15
1.3.3 Yield projections 16
1.4 Distribution of agricultural revenue 18
1.4.1 Mean revenue 19
1.4.2 Standard deviation of revenue 21
1.4.3 Downside risk from extreme weather 24
1.5 Endogenizing prices, crop choice and trade 27
1.5.1 Summary of model 27
1.5.2 Results 28
1.6 Discussion 31
1-A Additional temperature specifications 33
1-B Details of model and parameterization 34
1-B.1 Parameterization 36
1-C Additional figures and tables 38

2 CCTs and Fertility: Long-Term Impacts Across Two Generations 70
2.1 Introduction 71
2.2 Context and Theoretical Motivation 74
2.2.1 Context 74
2.2.2 Theoretical Motivation 78
2.3 Empirical Strategy and Data 83
2.3.1 Treatment Exposure 83
2.3.2 Cross-Census Variation 85
2.4 Results 87
2.4.1 Main Results 87
2.4.2 Additional Effect of Schooling 89
2.4.3 Mechanisms 91
2.5 Conclusion 93

2-A Appendix 93
  2-A.1 Support for the Identification Strategy 93
  2-A.2 Supplemental Figures 96
  2-A.3 Supplemental Tables 99

3 When are Tariff Cuts Not Enough? Heterogeneous Effects of Trade Preferences for the Least Developed Countries 102
3.1 Introduction 103
3.2 Theoretical framework 105
3.3 Empirical strategy 108
3.4 Data and summary statistics 110
  3.4.1 Importers and program reforms 110
  3.4.2 Exporters 112
  3.4.3 Affected products and tariffs 113
  3.4.4 Nonzero trade flows 116
  3.4.5 Export experience 118
3.5 Results 119
3.6 Conclusion 128
3-A Appendix 128

Bibliography 132

List of Tables

1-1 Average changes in temperature and precipitation distributions across districts 11
1-C.1 VDSA summary statistics in 2000 39
1-C.2 Upper limits of decile temperature bins 51
1-C.3 Yield-weather estimation – decile temperature bins 60
1-C.4 Yield-weather estimation – three-degree temperature bins 61
1-C.5 Yield-weather estimation – decile temperature bins including HDD variable 62
1-C.6 Yield-weather estimation – linear degree days 63
1-C.7 Yield-weather estimation – quadratic degree days 64
1-C.8 Yield-weather estimation – decile temperature bins with additional precipitation variables 65
1-C.9 Yield-weather estimation – linear degree days with additional precipitation variables 66
1-C.10 Yield-weather estimation – three-degree temperature bins with additional precipitation variables 67
1-C.11 Yield-weather estimation – decile temperature bins including HDD variable with additional precipitation variables 68
1-C.12 Yield-weather estimation – quadratic degree days with additional precipitation variables 69

2-1 Exposure Measures 85
2-2 Effects on Fertility 88
2-3 Extra Effect of Schooling 90
2-4 Potential Mechanisms 92
2-A.1 Amount of Transfers 99
2-A.2 Marginality Index and Roll-Out 100
2-A.3 Effects on Fertility 2010 vs 2000 101

3-1 Reforms of OECD programs for LDCs, 1996-2011 111
3-2 Distribution of products affected by reforms 115
3-3 Summary statistics 117
3-4 Baseline results and export experience decomposition 120
3-5 Robustness checks - alternative country samples 122
3-6 Robustness checks - varying extensive margin cutoff 124
3-7 Additional results 125
3-8 Intensity of treatment 127
3-A.1 Sample exporters and treatment status by importer 129

3-A.1 Sample exporters and treatment status by importer (continued) 130
3-A.1 Sample exporters and treatment status by importer (continued) 131

List of Figures

1-1 Distributions of mean daily temperature and total annual precipitation . . . 10
1-3 Distributions of average temperature and total precipitation 12
1-4 Estimates of coefficients on temperature bins using decile temperature bins 17
1-6 Projected changes in mean agricultural revenue (in percent relative to 2000s) 20
1-8 Projected changes in the weather-induced standard deviation of agricultural revenue (in percent relative to 2000s) 23
1-10 Projected changes in downside risk – share of 1-in-100 bad weather years for agricultural revenue (relative to 2000s) 26
1-12 Comparison of naive and model-based welfare measures 30
1-14 Decomposition of welfare measures into mean and variance components . . 32
1-C.1 Most-planted crop (by land area) by district in 2000 38
1-C.3 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 40
1-C.5 Distributions of average temperature and total precipitation for each decade 41
1-C.6 Projected changes in agricultural revenue – three-degree temperature bins specification without additional precipitation variables 42
1-C.7 Projected changes in agricultural revenue – decile temperature bins specification including HDD variable without additional precipitation variables . . 43
1-C.9 Projected district-level proportional changes in agricultural revenue between the 2000s and 2090s 44
1-C.10 Projected changes in agricultural revenue – quadratic degree-days specification without additional precipitation variables 45
1-C.11 Projected changes in agricultural revenue – decile temperature bins specification with additional precipitation variables 46
1-C.12 Projected changes in agricultural revenue – linear degree-days specification with additional precipitation variables 47
1-C.13 Projected changes in agricultural revenue – three-degree temperature bins specification with additional precipitation variables 48
1-C.14 Projected changes in agricultural revenue – decile temperature bins specification including HDD variable with additional precipitation variables 49
1-C.15 Projected changes in agricultural revenue – quadratic degree-days specification with additional precipitation variables 50
1-C.16 Share of days in each decile temperature bin relative to the 2000s 51

1-C.17 Projected changes in welfare measures – three-degree temperature bins specification without additional precipitation variables 52
1-C.18 Projected changes in welfare measures – decile temperature bins specification including HDD variable without additional precipitation variables 53
1-C.19 Projected changes in welfare measures – quadratic degree-days specification without additional precipitation variables 54
1-C.20 Projected changes in welfare measures – decile temperature bins specification with additional precipitation variables 55
1-C.21 Projected changes in welfare measures – linear degree-days specification with additional precipitation variables 56
1-C.22 Projected changes in welfare measures – three-degree temperature bins specification with additional precipitation variables 57
1-C.23 Projected changes in welfare measures – decile temperature bins specification including HDD variable with additional precipitation variables 58
1-C.24 Projected changes in welfare measures – quadratic degree-days specification with additional precipitation variables 59

2-1 Treatment Expansion Over Time 76
2-2 Program Expansion Towards More Urban Areas 77
2-3 Cohorts Definitions 78
2-A.1 Probability of Giving Birth by Age in Each Census 96
2-A.2 Cumulative Fertility by Age and Census 97
2-A.3 Probability of First Child 97
2-A.4 Within Census Comparisons, Schooling and Fertility 98

Chapter 1

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

By Francisco Costa, Fabien Forge, Jason Garred and João Paulo Pessoa

Abstract

Climate change may have significant implications for agricultural productivity, as well as the risks associated with crop production. Shifts in the distributions of temperature and precipitation may increase the volatility of farmers' yields, leading to rising but insurable risk, and/or reduce mean yields and thus cause permanent reductions in the returns to farming. We investigate the extent to which climate change will result in such insurable and uninsurable losses for farmers in India. To do so, we use the Community Earth System Model Large Ensemble dataset to predict the future distribution of yields at the district level for sixteen major crops. For the average district, we project a sharp decline in mean agricultural revenue, but relatively small shifts in volatility. This is because weather events resulting in extremely low agricultural revenue – what had once been 1-in-100-year events – are predicted to become the norm by the end of the century, implying substantial uninsurable losses from the changing climate. These results provide new insight into the relative importance of widely discussed channels of adaptation to climate change for farmers. Specifically, our findings suggest that strengthening agricultural insurance markets may not be of first-order importance to mitigate farmers' losses due to climate change. Instead, to combat the threat of large but uninsurable losses, a more promising avenue is to bolster adaptation measures aimed at improving crop yields in conditions of extreme heat.

1.1 Introduction

Climate change may have important consequences for agriculture around the globe. This might include not only the average crop yields a farmer can expect in the future, but also the volatility of agricultural production and the extent of downside risk. If changing patterns of temperature and precipitation result in increases in yield variability, this could be especially harmful for farmers without access to well-developed insurance markets, suggesting a key role for policies mitigating risk through increases in access to insurance. On the other hand, if climate change leads to a sharp reduction in mean yields without an accompanying change in volatility, the main issue facing farmers will be the prospect of permanent, uninsurable losses.

In this paper, we investigate the extent to which climate change will result in insurable and uninsurable losses for farmers in India, a setting with a large agricultural sector but relatively limited availability of insurance as compared to more developed economies. Specifically, we contrast the potential effect of climate change on the average agricultural outcomes faced by farmers with its impact on the volatility of these outcomes, as well as the probability of extreme events.

To do this, we generate information on the projected future distribution of yields for each location, crop and time period in our sample. The key data input facilitating our analysis is the Community Earth System Model Large Ensemble (CESM-LENS) dataset. This is a multi-run global simulation created by climatologists as a tool to study the potential variability of temperature and precipitation patterns within a given climate (66). It provides us with multiple potential realizations of weather at a given location for each day between 1920 and 2099. For instance, rather than specifying a single projected value for precipitation for the southern tip of India on August 1 2050, the CESM-LENS data instead includes forty potential local precipitation realizations that are forecast to be consistent with the prevailing climate as of 2050.[1,2] We use this data to predict the evolution of the distribution

---

[1] The CESM-LENS data has frequently been used in the scientific literature to study the potential distribution of future weather realizations. For instance, (39) leverage the multiple realizations in the dataset to assess the changing future probability of the joint occurrence of warm and dry weather conditions in California. (108), (9), (38), (100) and (97) similarly use CESM-LENS to assess the future likelihood of historically extreme climate events.

[2] 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.

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 (36; 91). 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.[3]

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 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.

---

[3] 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. 56; 20; 58).

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.

Overall, our findings predict that climate change will have substantial negative effects on Indian farmers' mean returns from agriculture, driven by the normalization of formerly extreme weather patterns. On the other hand, our projections do not provide strong evidence of a large rise in risk via increases in yield volatility. This suggests that in order to mitigate losses due to climate change faced by farmers in developing countries, strengthening agricultural insurance markets may not be of first-order importance. Instead, to combat the threat of large but uninsurable losses, a more promising avenue of adaptation could be the development and adoption of crop varieties resistant to extreme heat.[4]

An important caveat of our analysis is that it holds factors such as crop choice and prices constant. We address this concern by also taking a structural approach using the general equilibrium portfolio model of (4). First, we allow farmers to adjust their crop choice according to the evolving distribution of yields, trading off the relative changes in risk and return across crops.[5] Second, we let crop prices adjust endogenously with agricultural production. Third, we allow for changes in trade flows between Indian districts.[6]

By comparing the welfare measure implied by the model with a naive version using nominal revenue, we find that allowing for endogenous crop choice, prices and trade has a small

---

[4] Using a randomized experiment conducted in India, (43) find that adoption of a new rice variety with improved resistance to extreme weather events (floods) was successful in avoiding yield losses in years with floods, and also encouraged farmers to make investments that improved yields in years without floods. The results of (33) suggest that adoption might be encouraged through targeted financing; they find that access to subsidized credit induced investment in climate-resilient livestock in Brazil. Alternatively, (73) show that interventions encouraging diversification into nonagricultural activities were effective in enabling Nicaraguan households to manage drought shocks.

[5] The study of (32) suggests that this type of adaptation could partially offset the potential welfare losses from climate change of farmers worldwide. (5) show that subsistence farmers in Peru adjust the crops they plant to mitigate the impacts of extreme heat events.

[6] 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, (48) 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. (18) do not find evidence of successful past adaptation by Indian agriculture to problems of water scarcity. (98) concludes that Indian farmers can 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 labor supply in developing countries, due to migration (24; 60; 63; 67), human capital accumulation (52), and sectoral mobility (31; 25).

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 returns 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%.

Our paper builds on the literature on the consequences of climate change for agriculture in developing countries. A large body of work examines the projected effects of climate change on average agricultural outcomes (e.g. 8; 90; 72; 68; 26; 28), some focusing specifically on India (56; 20; 58; 98).[7] Some previous studies have also considered the future variability of agricultural yields, using projections of year-to-year weather variation.[8] (104) use the predicted time-series variation in weather patterns from one run of a given climate model to project future year-to-year variability in US maize yields, and then repeats this exercise separately for fifteen different climate models, deriving a distinct estimate of future variability for each model. (101) project 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 a baseline historical period.

Our key contribution is to jointly consider both the uninsurable losses and the changing risks facing farmers in a unified framework. By studying the *distribution* of yields, we can examine the evolution of both mean returns and their variability, along with the probability of extreme outcomes. To accomplish this, we deviate from standard approaches by harnessing numerous potential outcome realizations from a multi-run climate model rather than relying on year-to-year weather variation.[9] Our paper also adds to a growing set of assessments of the potential effects of climate change that take advantage of the structure of general equilibrium

---

[7] There are also many groundbreaking studies looking at the United States (e.g. 2; 74; 88; 91; 46; 22).

[8] Our work is also related to a recent set of papers considering the implications of agents' uncertainty about the future climate, such as (65), (71) and (94). 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; i.e. changes in known risk.

[9] (21) use multiple climate models to refine projections of the average impacts of climate change. Here, we exploit a multi-run climate model especially designed for the study of internal climate variability, in order to examine future changes in the distribution of agricultural yields, including their volatility. In this context, a key argument for focusing on a single multi-run model rather than a group of models is given in (9): "[a]lthough structural uncertainty and internal variability are conflated in the CMIP5 ensemble [of models], diagnosis of a single-model ensemble allows us to delineate the role of internal variability alone".

models at the frontier of current research.[10]

1.2 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.

1.2.1 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 1-C.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

---

[10] Along with the Ricardian formulation used by (32) and our work based on (4), a study by (40) employs the general class of models explored by (7) to consider the impact of climate change on global inequality. (35) estimate the marginal product of climate in terms of total economic output using an Arrow-Debreu framework. 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.

[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.

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 1-C.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 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.

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

---

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

[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 (78).

[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.

realized weather within a given climate state (66). 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 distribution 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. 97).

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

---

[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" (66). 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 (103).

[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, are more computationally intensive and offer limited improvement over this 'nearest neighbor' method given the relatively low resolution of our data.

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 37).[22]

In Figure 1-1, 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 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 1-1, 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 1-1 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

---

[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.

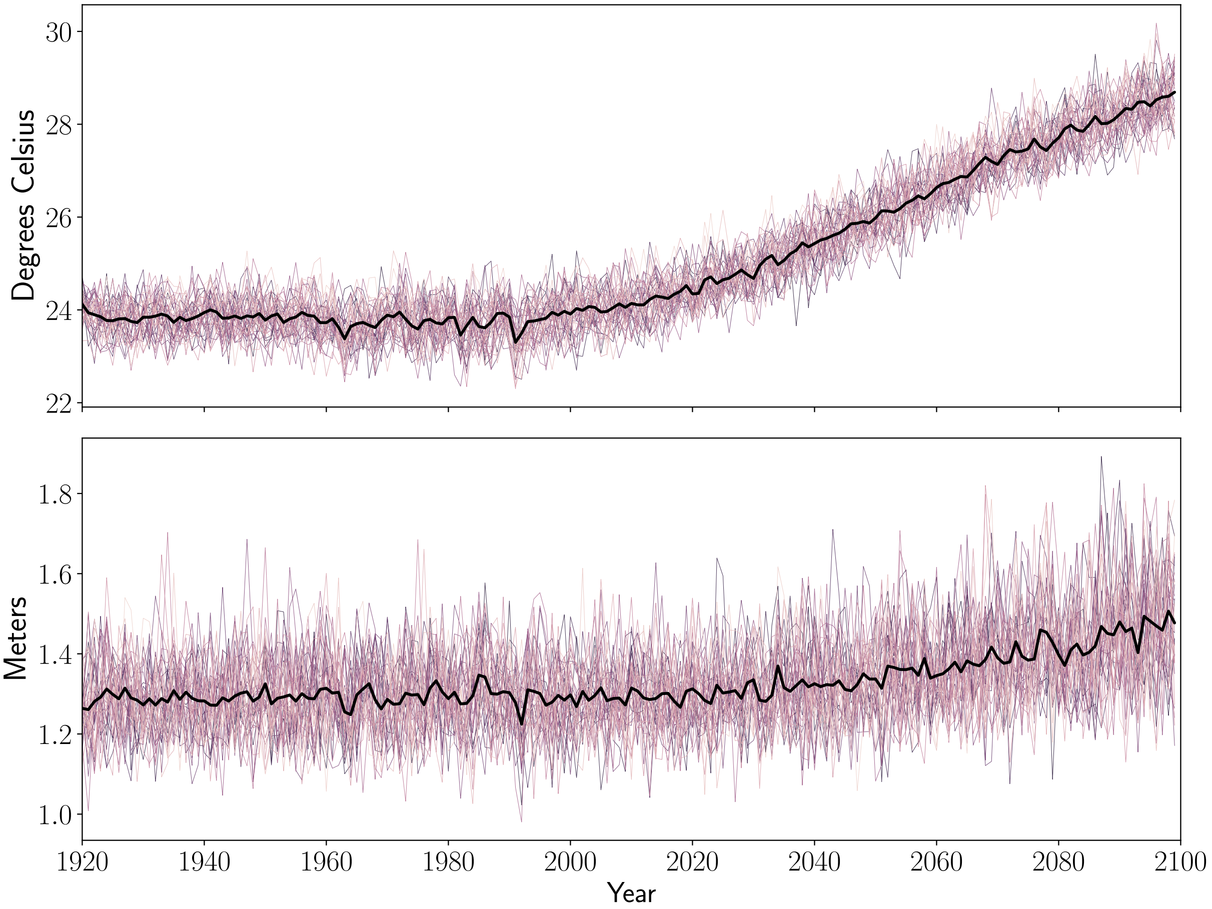

Figure 1-1: Distributions of mean daily temperature and total annual precipitation

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.

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 1-C.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 1-C.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

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

	(1) Average temperature Mean	(2) S.D.	(3) Total precipitation Mean	(4) 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.

translating our regression estimates into yield projections. Figure 1-2(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 the median for the 2090s exceeds the 75th percentile of the 1979-2015 distribution.

---

[23] In Appendix Figures 1-C.4(a) and 1-C.4(b), we also provide decade-by-decade box plots of the distributions of temperature and precipitation in the CESM-LENS data.

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 1-2(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.

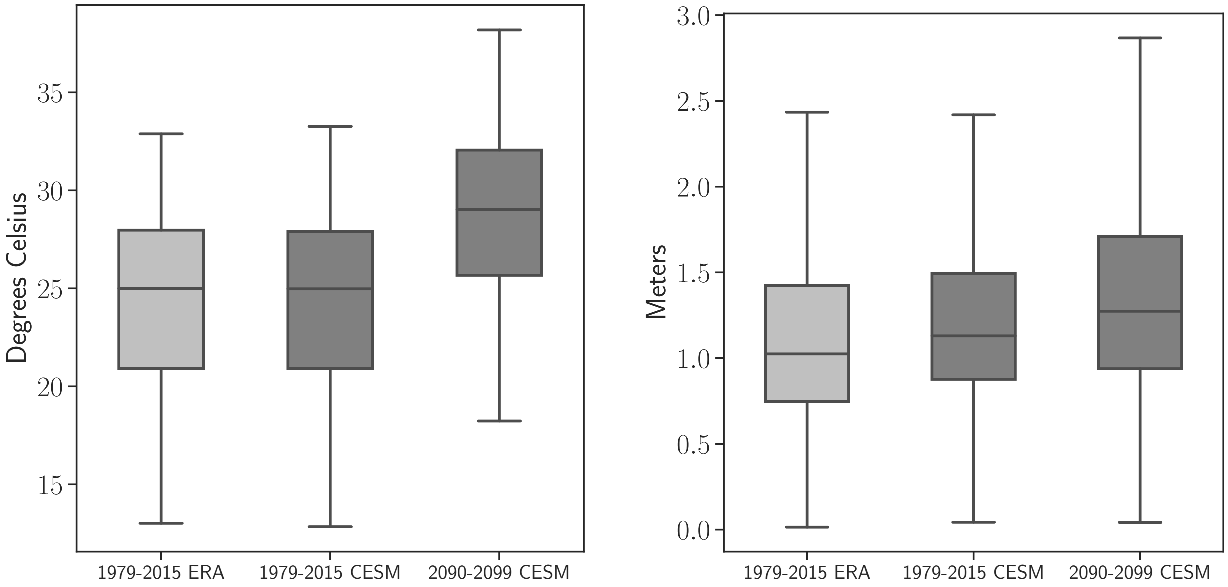

(a) Temperature

(b) Precipitation

Figure 1-3: 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.

1.3 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.

1.3.1 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.

Temperature bins

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

$$\ln A_{isgt} = \sum_{k=1}^{K} \beta_g^k T_{isgt}^k + \phi_g P_{isgt} + \psi_g P_{isgt}^2 + \alpha_{isg} + \gamma_{sg} t + \delta_{sg} t^2 + \epsilon_{isgt} \tag{1.1}$$

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 (91).

We divide daily average temperatures into a set of K bins, to be described in more detail below. The variable T_{isgt}^k 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 (91). 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.

---

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

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 the bins defined by (91); 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 1-A.

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:

$$\ln A_{isgt} = \beta DD_{isgt} + \phi_g P_{isgt} + \psi_g P_{isgt}^2 + \alpha_{isg} + \gamma_{sg} t + \delta_{sg} t^2 + \epsilon_{isgt} \qquad (1.2)$$

---

[25] The temperatures covered by each bin for each growing season are listed in Appendix Table 1-C.2.

[26] The resulting specification is thus closer to that of (89) rather than (91).

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 1-A.

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 (47) 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 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.

1.3.2 Yield-weather regression results

We now briefly summarize the results of the regressions outlined above.[27] Figure 1-4 displays estimates of the coefficients β_g^k on the decile bins in equation (1.1), 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 1-C.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

---

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

[28] See Appendix Table 1-C.2 for the temperatures covered by each bin.

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 1-C.6.[30]

Finally, we re-estimate each of the above specifications, adding the additional precipitation variables from (47). In both sets of regressions (see Appendix Tables 1-C.8 and 1-C.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]

1.3.3 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:

$$\ln \tilde{A}_{isgrt} = \ln A_{isg,2000} + \sum_{k=1}^{10} \hat{\beta}_g^k \Delta T_{isgrt}^k + \hat{\phi}_g \Delta P_{isgrt} + \hat{\psi}_g \Delta P_{isgrt}^2 \tag{1.3}$$

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

---

[29] See also Appendix Tables 1-C.4 and 1-C.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 1-C.7.

[31] In Appendix Tables 1-C.10 to 1-C.12, we display the results of including the additional precipitation variables in our three other specifications.

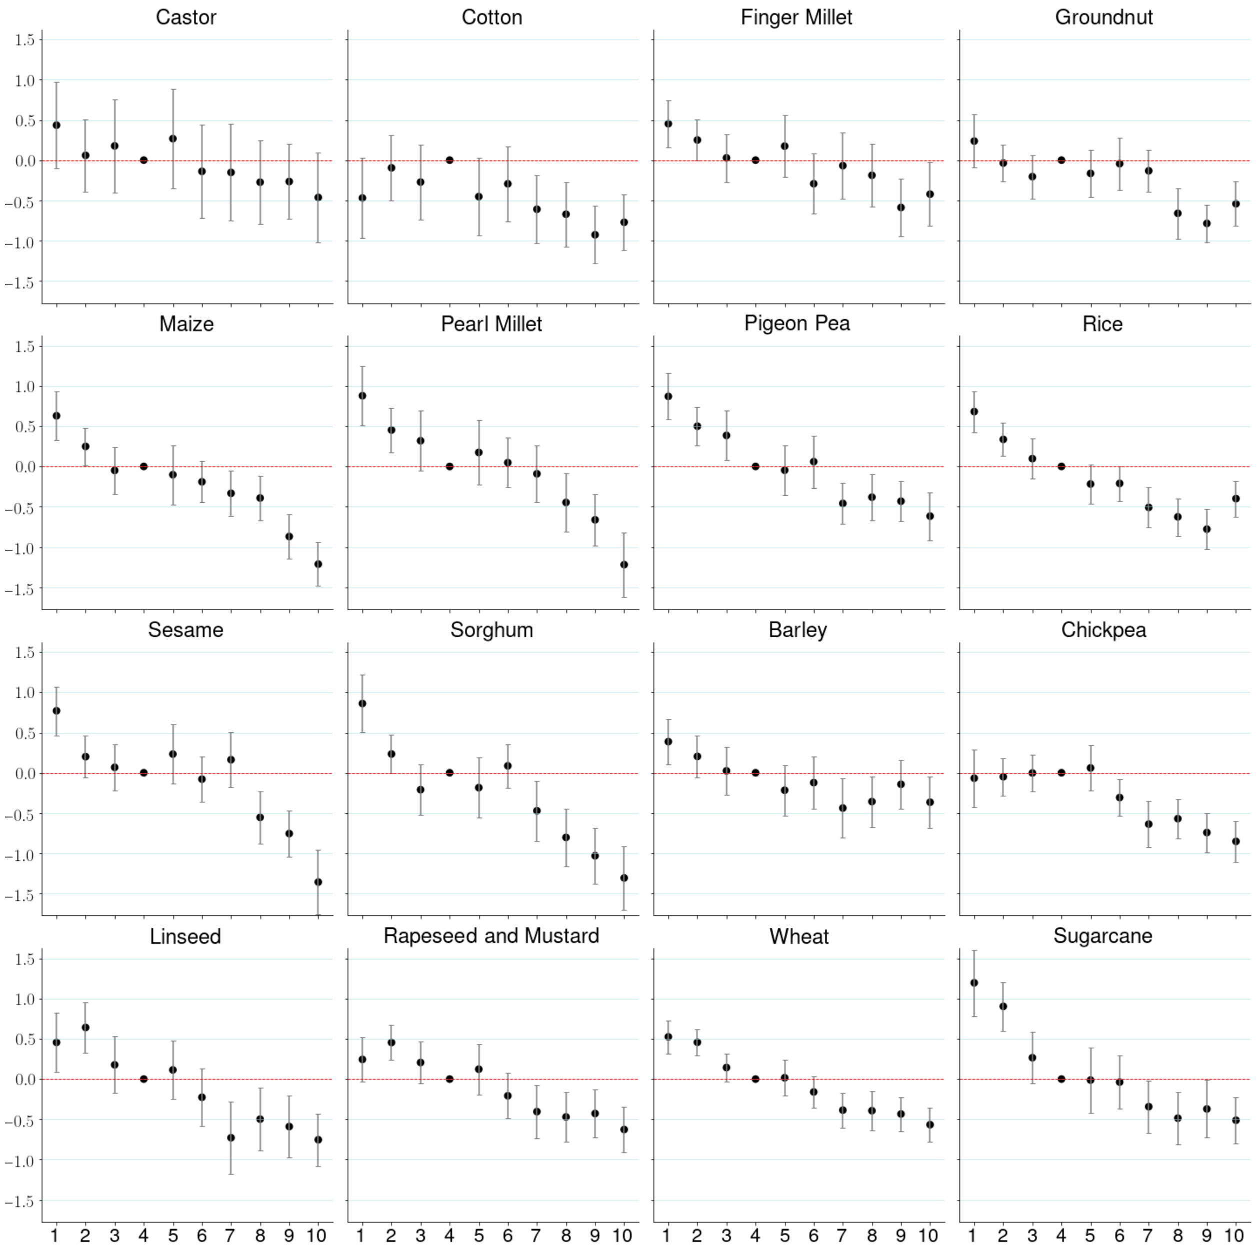

Figure 1-4: Estimates of coefficients on temperature bins using decile temperature bins

This figure shows the estimated coefficients on temperature bins and their corresponding 95% confidence intervals, using the decile temperature bins specification without additional precipitation variables. 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 1-C.2 for the limits of each bin in degrees Celsius for each growing season.

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 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.

1.4 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.

---

[32] Holding technology fixed is a standard approach in this literature (e.g. 36; 91; 61).

[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.

[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.

1.4.1 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 1-5(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 (1.1)). 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 1-C.6 and 1-C.7.[35]

We next consider the projections based on our specification modelling the temperature-yield relationship as a linear function of degree days (equation (1.2)). The right side of Figure 1-5(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 1-C.10, adding a quadratic degree-days term makes little difference to these results.[37]

Figure 1-5(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 (47) into these specifications does

---

[35] Based on the specification using the temperature bins of (91), 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.

[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 1-C.9.

[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 (106), who conclude that extreme temperature has been more important than precipitation in driving yield anomalies worldwide.

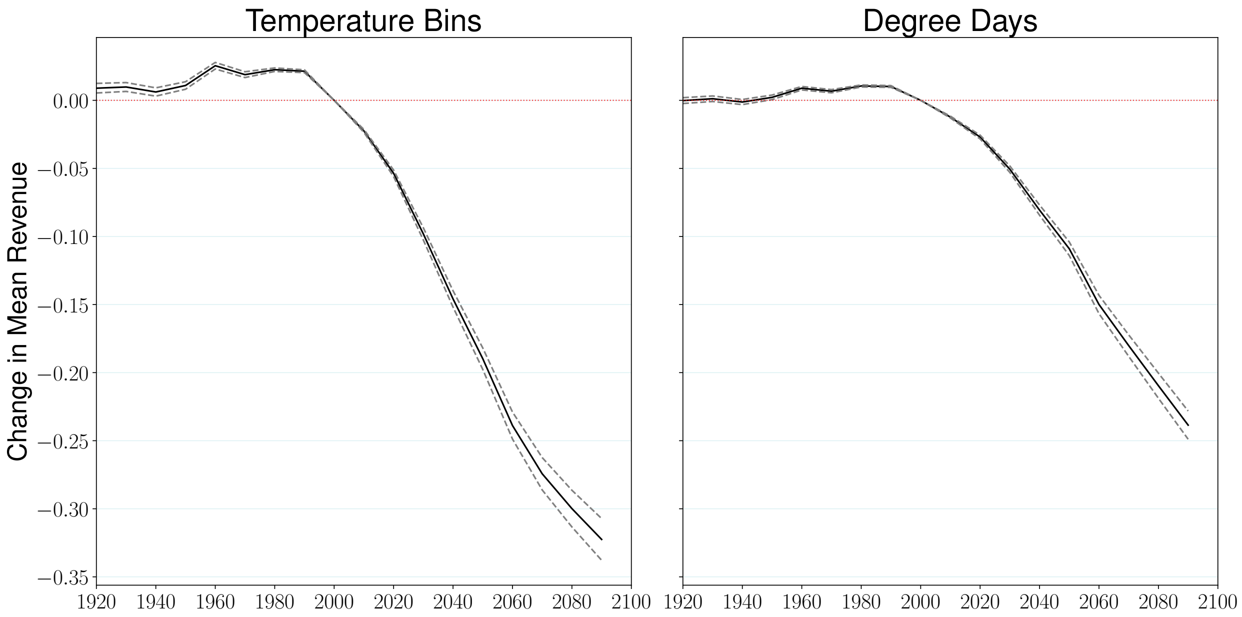

(a) Total

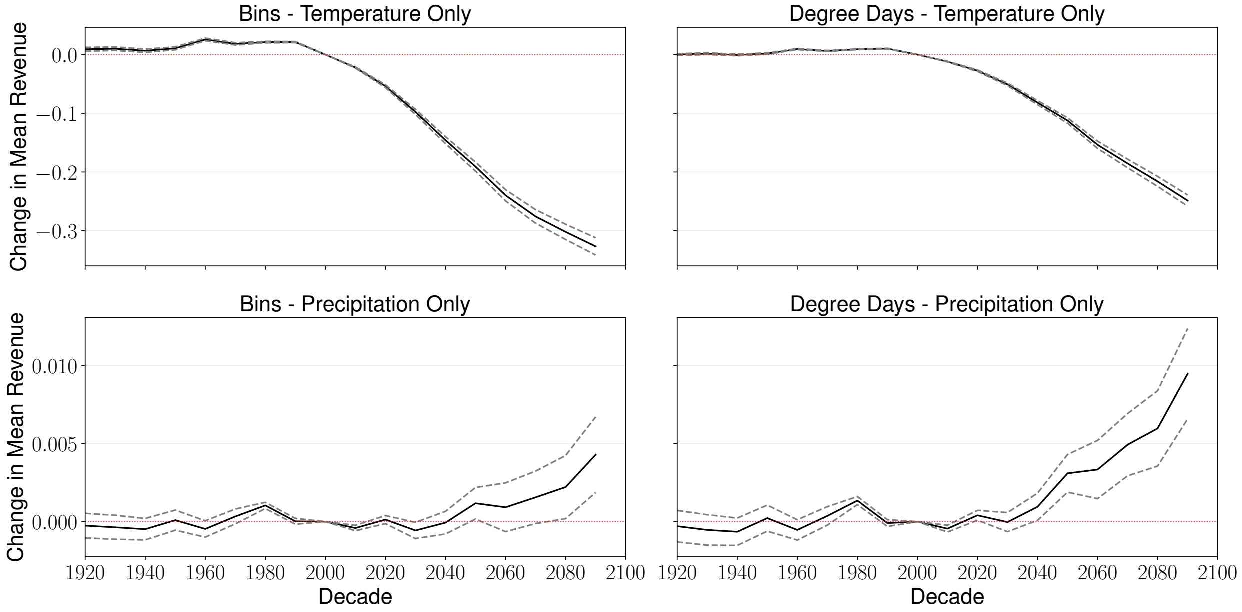

(b) Decomposition

Figure 1-6: 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 1.4. 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.

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 1-C.11 and 1-C.12.[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.

1.4.2 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

---

[39] See also Appendix Figures 1-C.13 to 1-C.15 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%.

[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.

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 1-7(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 1-7(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 1-C.16, 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 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 1-7(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 1-7(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 1-7(b).[42] We do observe a small

---

[41] When we use the fifteen bins of (91), 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 1-C.6 and 1-C.7.

[42] The addition of the precipitation variables from (47) 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 1-C.11) and to 6.4% in the model based on degree days (Appendix Figure 1-C.12).

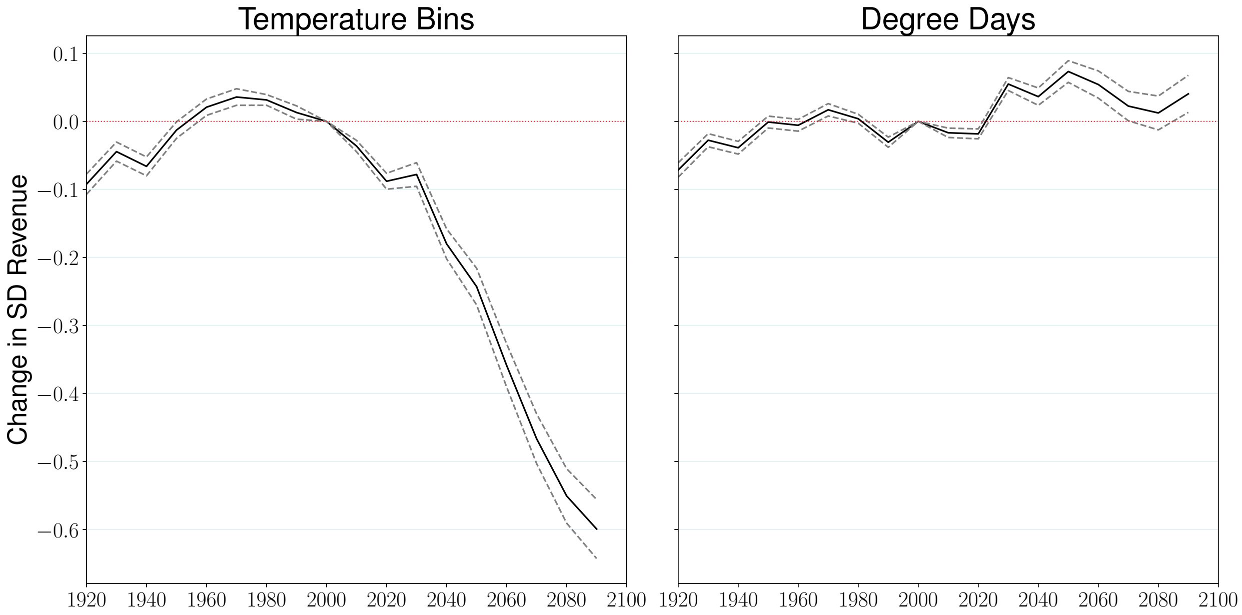

(a) Total

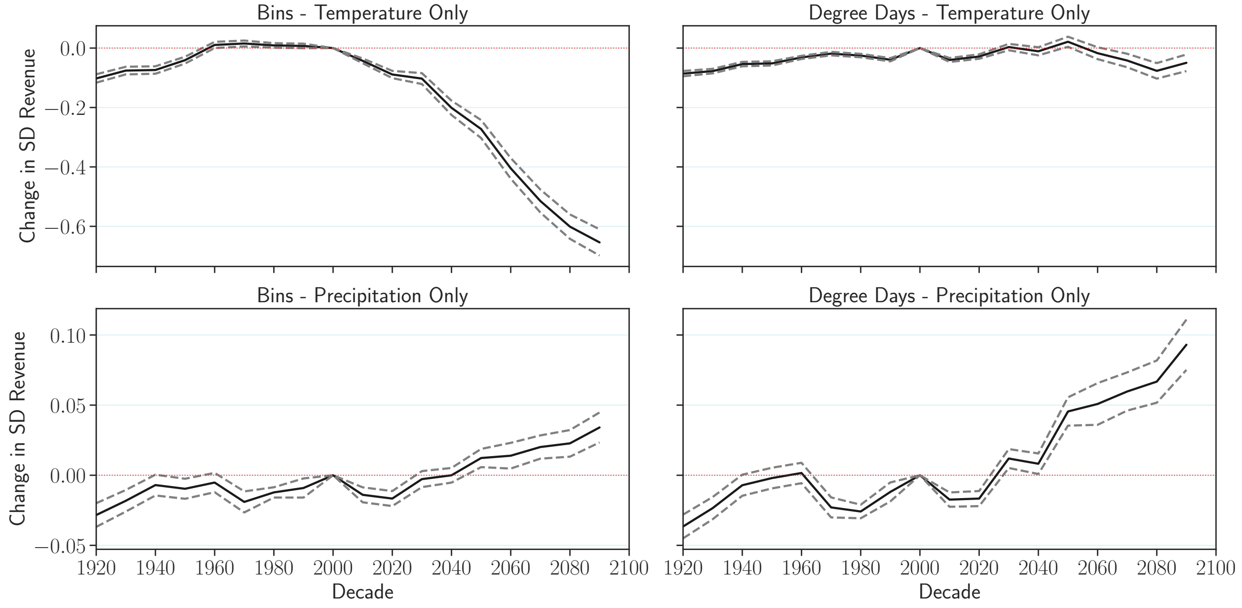

(b) Decomposition

Figure 1-8: 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 1.4. 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.

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.

1.4.3 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 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 1-9(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

---

[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 1-C.10 and 1-C.15.

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 1-9(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 1-9(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 1-9(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 1-9(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 1-C.11 and 1-C.12, 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

---

[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 1-C.11 and 1-C.12).

[45] See Appendix Figures 1-C.6, 1-C.7, 1-C.10, and 1-C.13 to 1-C.15 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.

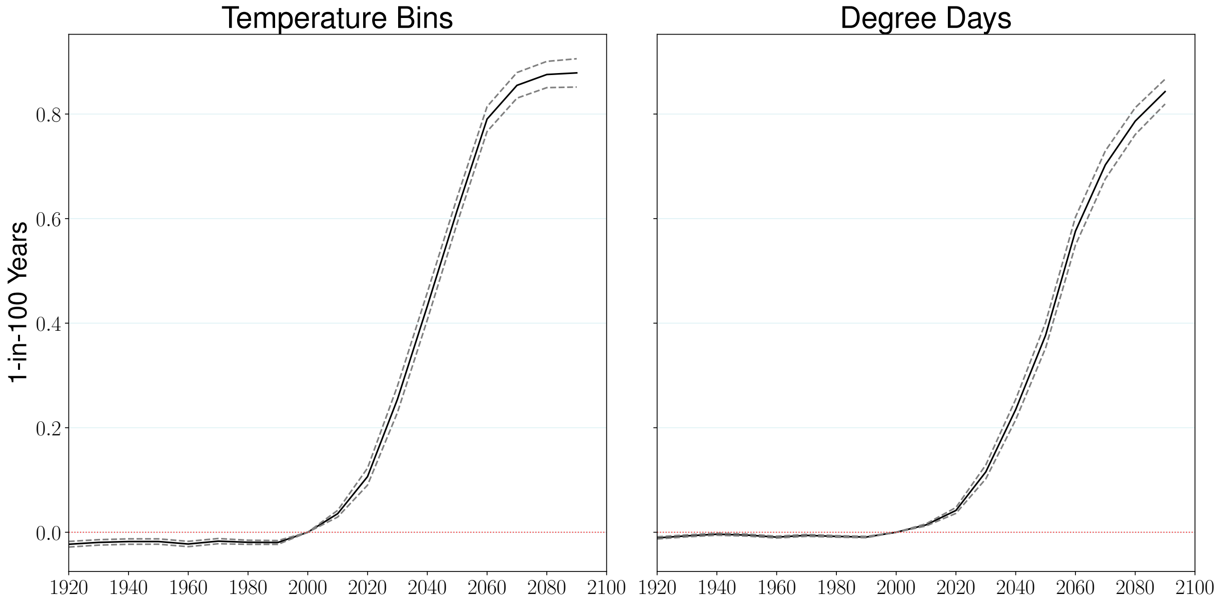

(a) Total

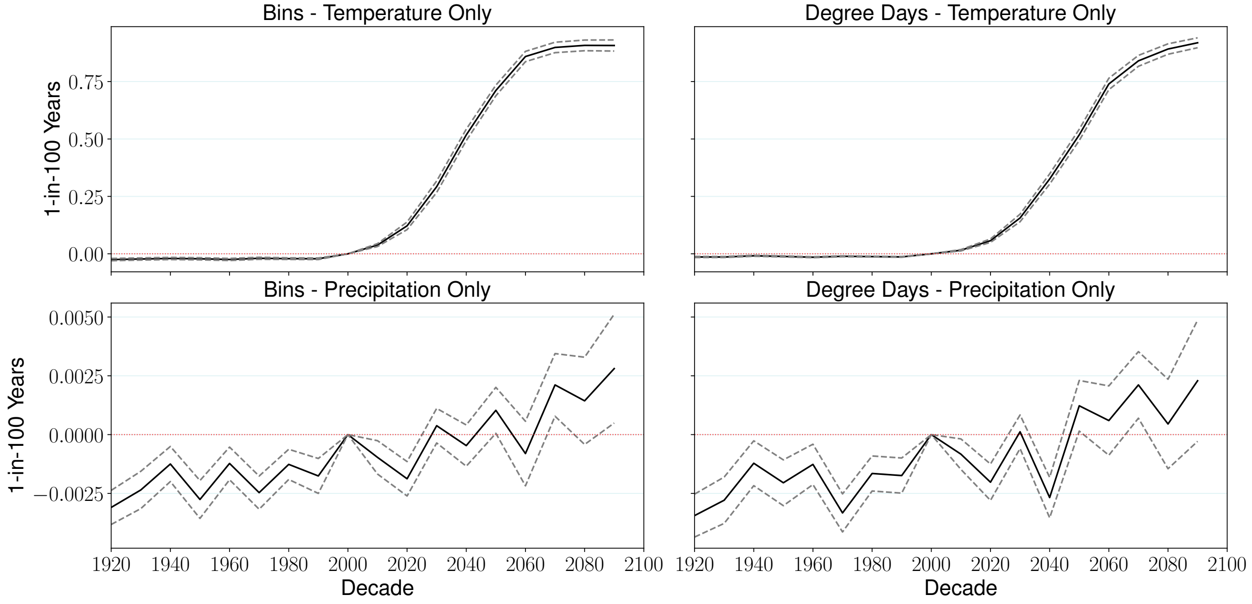

(b) Decomposition

Figure 1-10: 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 1.4. 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.

is predicted to become the 'new normal'.

1.5 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] Despite the additional structure required for this exercise, we find that the key insights from the previous section are preserved.

1.5.1 Summary of model

We use the model and methodology of (4). This model combines a Ricardian framework, incorporating many locations and goods, with portfolio choice by agents across risky assets. Here, we explain the setup and main implications of the model in intuitive terms, leaving an exposition of the model's key equations to Appendix 1-B.

The model allows for a finite number of crops, each of which may be grown in any district. Farmers located in a given district face risk when growing crops, in the sense that each crop's yield (quantity produced per unit land area) in the farmer's district is stochastic. Specifically, the yield of a particular crop in a given location is lognormally distributed with a mean and variance that is known to the farmer. Each farmer is also aware of the covariance of these potential yields across crops within their local district. So farmers face risk but not uncertainty: they are aware of the distributions of potential yields that they face, but not the yields that will actually be realized. Moreover, while farmers in different districts have the same menu of crops from which to choose, any given crop may have a different productivity distribution in different districts.

The number of farmers may differ across districts. All farmers are geographically immobile, so the population of farmers in each district is fixed. Moreover, farmers are all identical in their preferences, which implies that all farmers within each district make identical decisions.

---

[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.

All key outcomes of the model can therefore be described at the district level.

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

The model goes beyond a single-agent portfolio choice framework by allowing farmers in different districts to interact through trade in crop output, and allowing prices to adjust endogenously to clear markets. As usual in Ricardian models, a district will tend to produce and export the crops in which it has a comparative advantage in productivity (depending on the prevailing level of demand for each crop). However, this finding is now more subtle than in a standard model, because a district's comparative advantage is affected by the variance in the local yield of each crop – its riskiness – as well as the mean yield. Therefore, farmers focus on the production of goods in which they have a 'risk-adjusted' comparative advantage.

To conduct a quantitative analysis using this framework, we use our projected yield distributions from Section 3.3 (based on the CESM-LENS data and our yield-weather regressions), and also calibrate or estimate the other parameters of the model.[47] We generate counterfactual outcomes for each decade from the 1920s to the 2090s by allowing for crop choices in each district to be based on the distribution of potential yields prevailing in that decade, and allowing for trade across Indian districts and prices to adjust accordingly. We then compare the resulting outcomes, including farmers' welfare (based on the expected utility function discussed above), to outcomes in our baseline decade, the 2000s.

1.5.2 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 Section 4. Specifically, we calculate

---

[47] We discuss the calibration and estimation of the other parameters required for this analysis, such as farmers' level of risk aversion, in Appendix 1-B.

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 1-11(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 1-11(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 (4), we separately assess the contributions of the first and second moments to changes in welfare (or ‘naive welfare’). Using the baseline temperature bins specification, we find a predicted 21st-century welfare decline of 28.9% via changes in the logarithm of mean real returns (Figure 1-13(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 1.4.1, where we projected a 32.3% fall in mean agricultural revenue (Figure 1-5(a)). Analogous conclusions may be reached with the specification using degree days, again by comparing Figure 1-13(a) to Figure 1-5(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 1-13(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

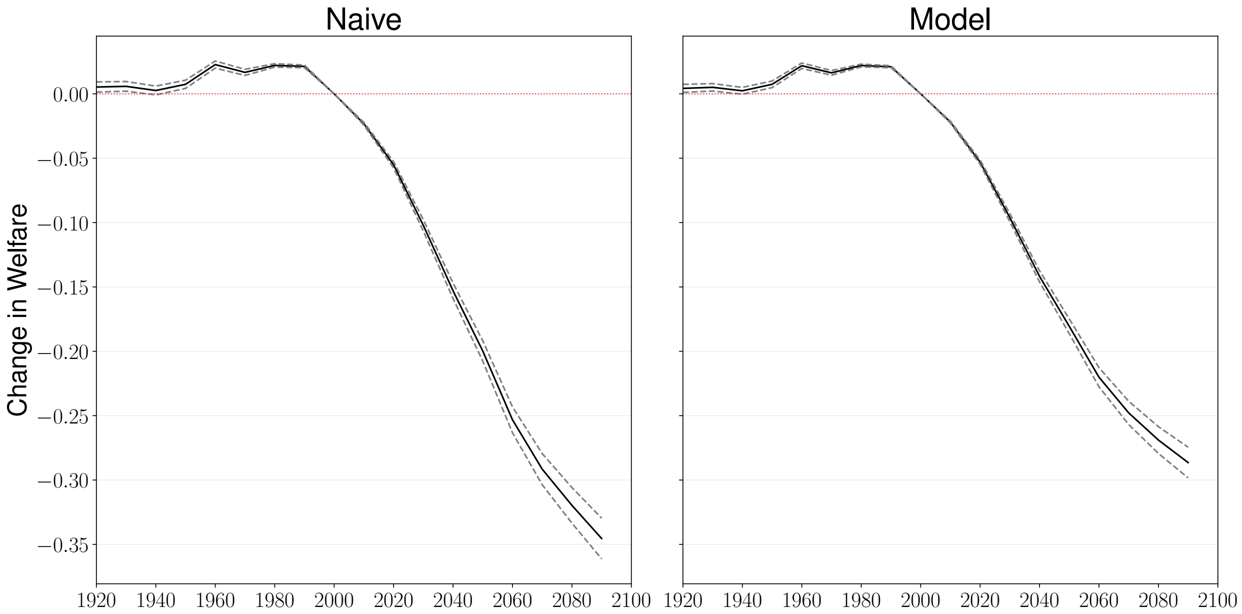

(a) Temperature Bins

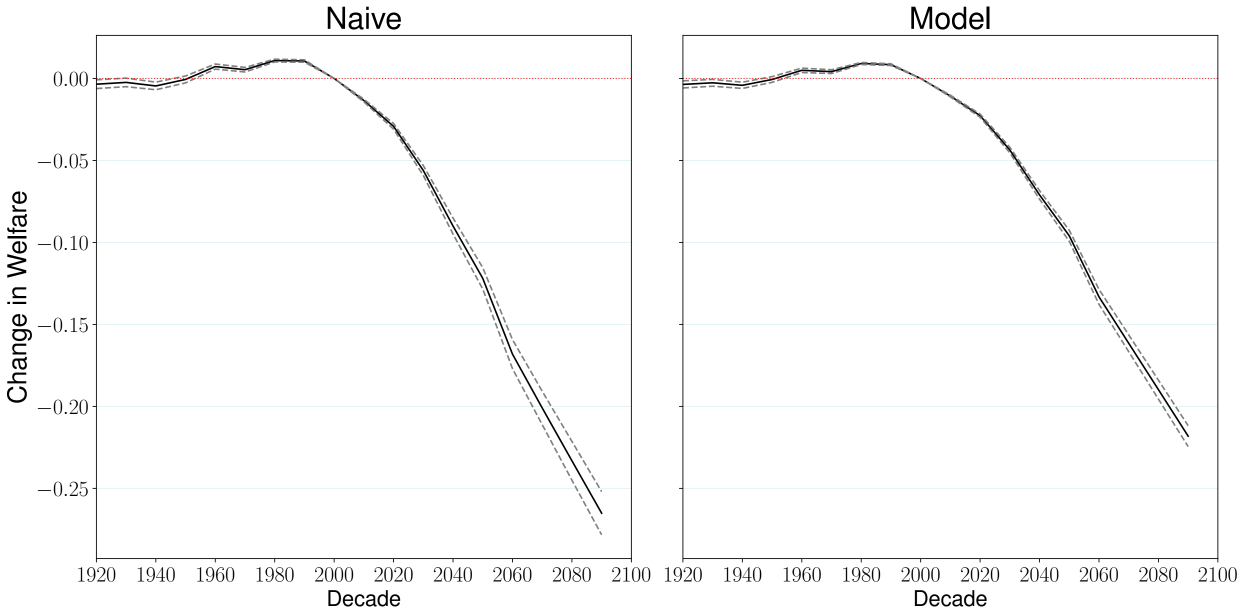

(b) Degree Days

Figure 1-12: 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 1.5. 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.

real returns in the welfare formula.[48]

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

1.6 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

---

[48] Results using projections based on our other yield-weather specifications may be seen in Appendix Figures 1-C.17 to 1-C.24. 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.

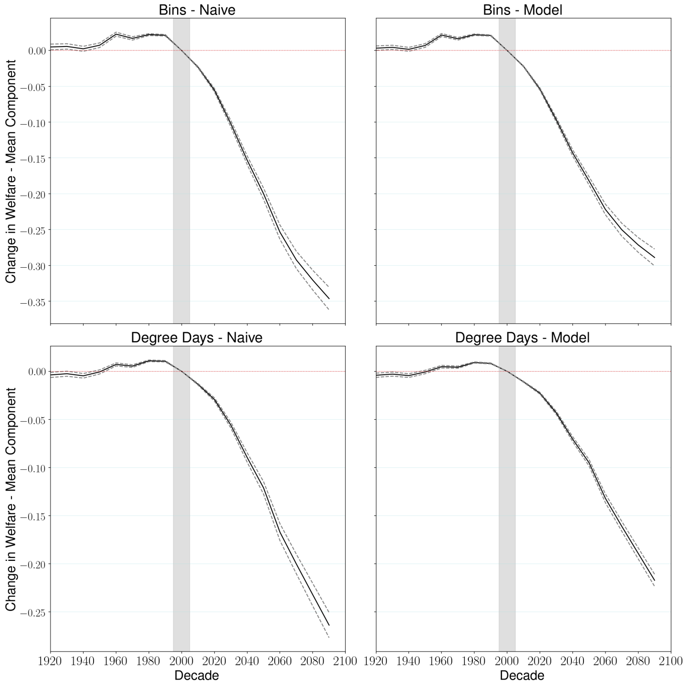

(a) Mean component

(b) Variance component

Figure 1-14: 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 1.5. 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.

practice, technological changes could be a crucial mode of adaptation.[49]

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 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.

Appendix

1-A 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

---

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

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.

1-B Details of model and parameterization

In this appendix, we discuss the setup and key implications of the (4) model, and then explain how the model is parameterized in our empirical application. Note that this model is sufficiently complex that a full exposition would require us to replicate a substantial part of the paper by Allen and Atkin. Here, we therefore aim to provide enough information to make clear the model's main components and their intuitive interpretations. A reader interested in understanding the model in depth should refer to the original paper.

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

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

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

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

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

Trade between districts is governed by price differentials and trade costs across districts, implying that goods will flow from low-cost producers to high-cost ones, with the intensity of the effects mediated by the magnitude of trade costs. Specifically, trade between districts is governed by the following log-linear arbitrage condition relating consumption, quantity produced and goods prices:[50]

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

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

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

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

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

Farmers in a given district choose the share θ_{ig} of land to be planted with crop g to maximize their (log) expected utility under the restriction that $\sum_g \theta_{ig} = 1$. After some manipulation, (4) show that the first-order conditions for each district and good can be

---

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

written in the following way:[51]

$$\mu_{ig}^{z} - \rho \sum_{h} \theta_{ih} \Sigma_{igh}^{z} = \lambda_{i}. \tag{1-B.2}$$

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

For counterfactual analysis using this model, we require sufficient information to calculate μ_{ig}^{z} and Σ_{igh}^{z} for each district and crop, given a particular distribution of planting choices, as well as an estimate of ρ. These are then used to identify a set of crop choices that satisfy the first-order conditions above. In our counterfactuals of interest, μ_{ig}^{z} and Σ_{igh}^{z} are functions of the distribution of our projected yields (based on the CESM-LENS data and our yield-weather regressions) in each decade. Several parameters are needed for these calculations; we next provide details on how we calibrate or estimate each of these parameters.

1-B.1 Parameterization

We assign α_g to be equal to the share of the value of each crop in total India-wide agricultural revenue according to the VDSA data for the year 2000. We use an estimate of 2.38 for σ from (4), 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

---

[51] See Proposition 1 in (4).

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 (1-B.2)):

$$\left(\mu^z_{igd} + c_{igd}\right) = \rho \sum_h \theta_{ihd} \Sigma^z_{ighd} + \delta_{id} + \delta_{gd} + \delta_{ig} + \xi_{igd}$$

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 (1-B.2)). As in (4), 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 (4), 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.

1-C Additional figures and tables

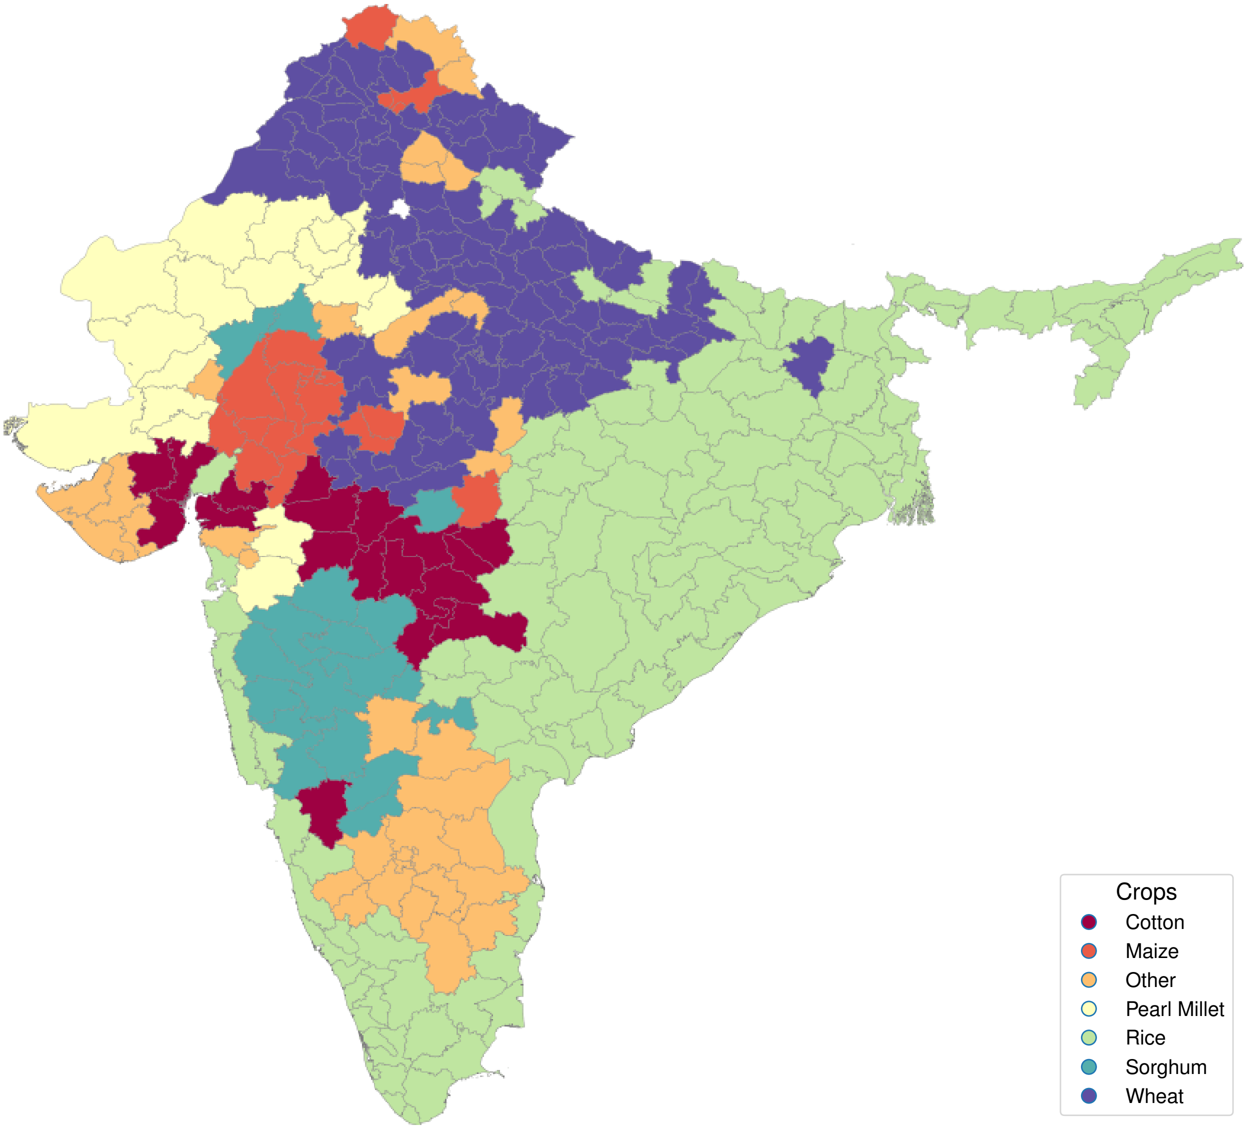

Figure 1-C.1: Most-planted crop (by land area) by district in 2000

Table 1-C.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).

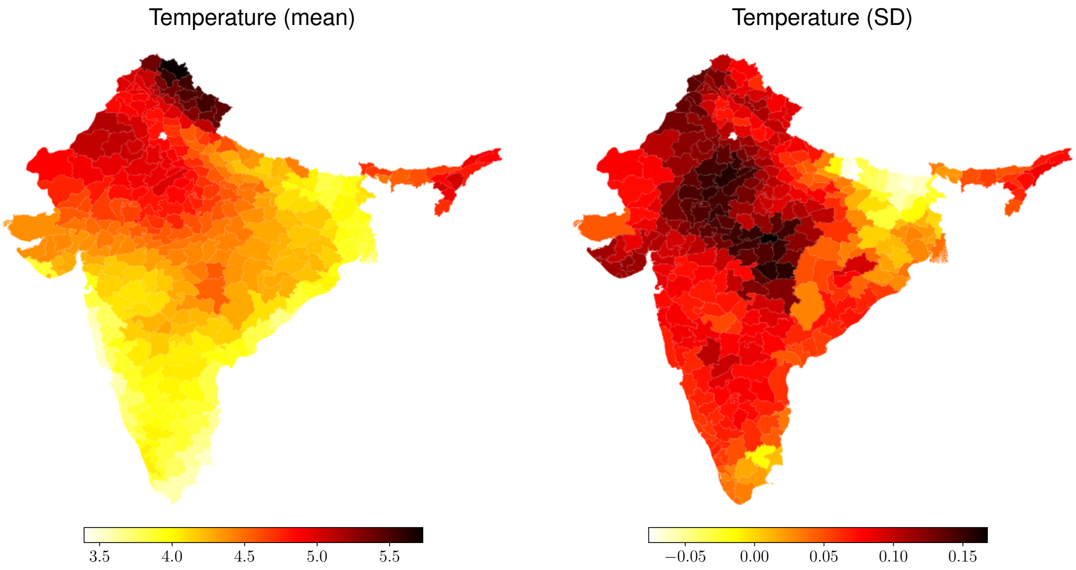

(a) Temperature

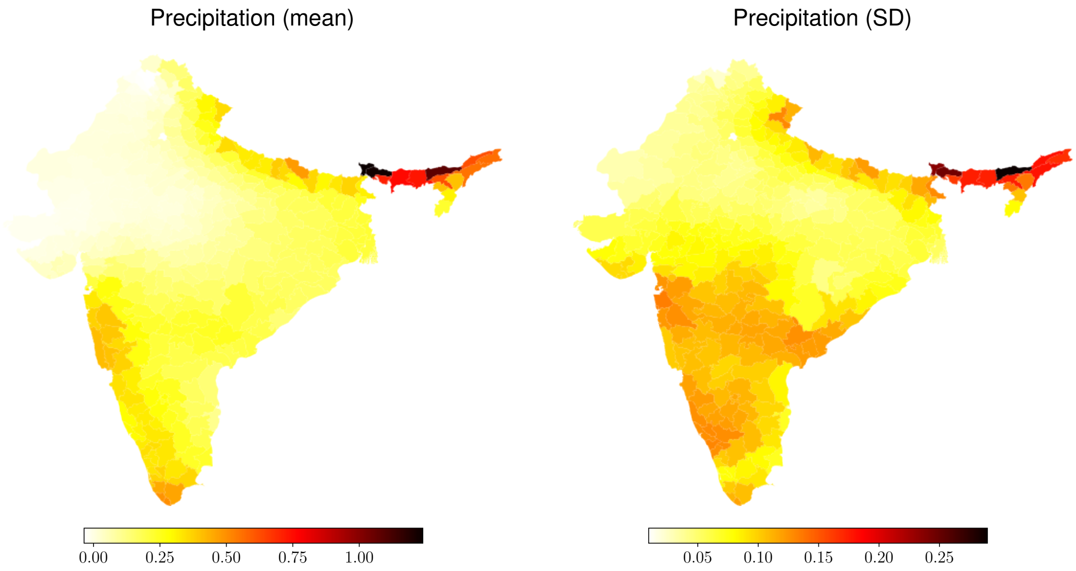

(b) Precipitation

Figure 1-C.3: 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

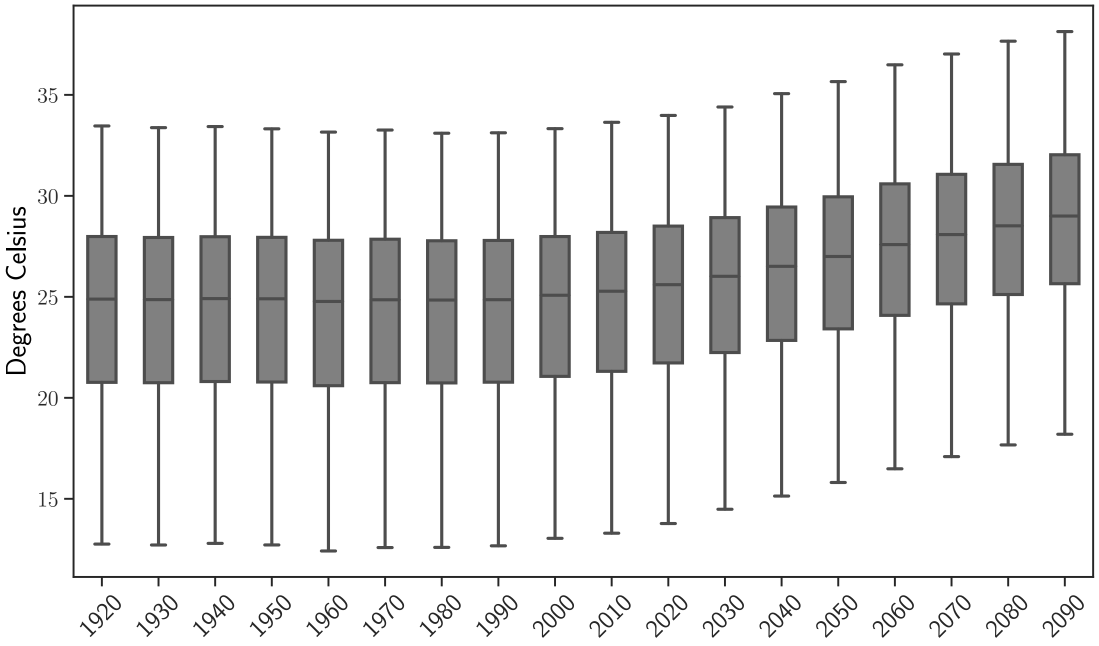

(a) Temperature

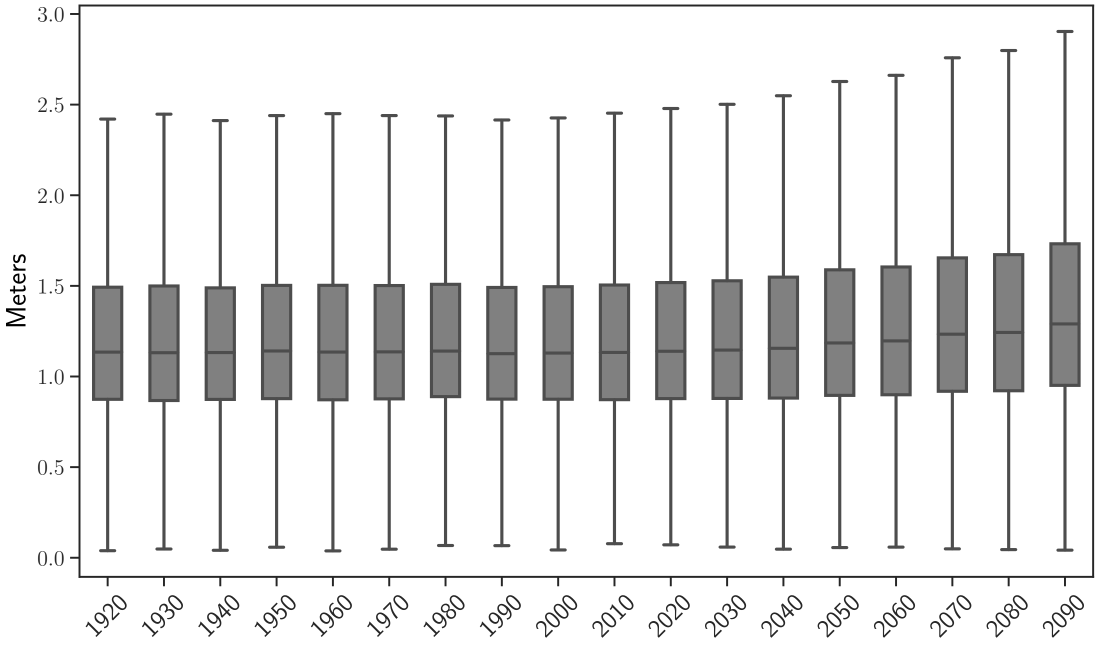

(b) Precipitation

Figure 1-C.5: 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.

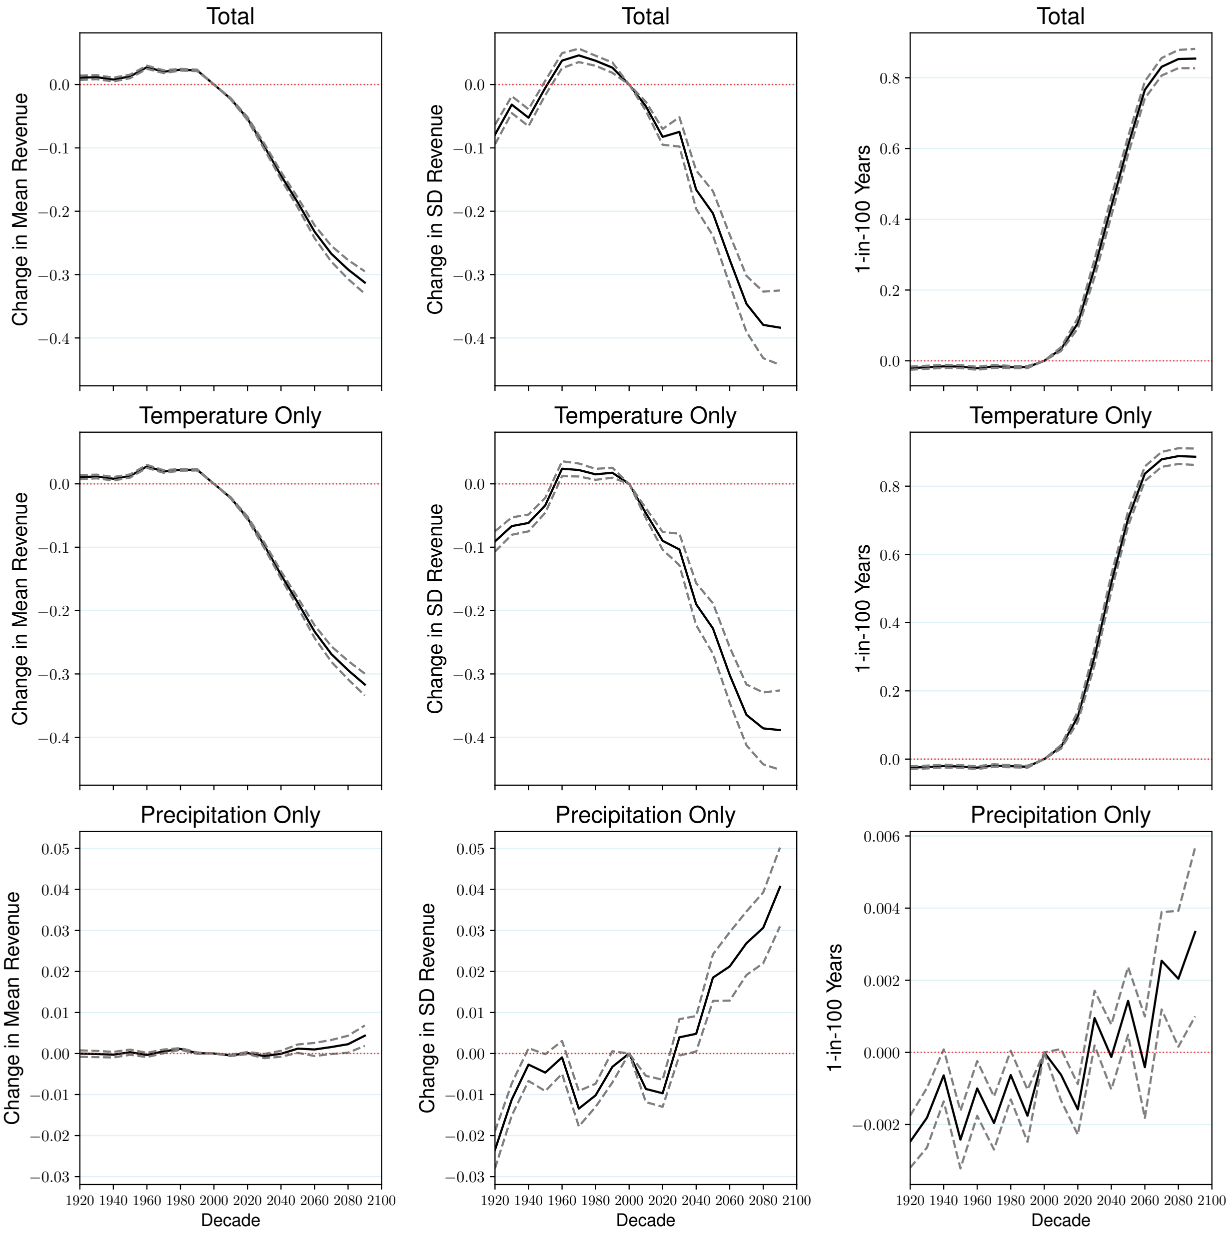

Figure 1-C.6: Projected changes in agricultural revenue – three-degree temperature bins specification without additional precipitation variables

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.

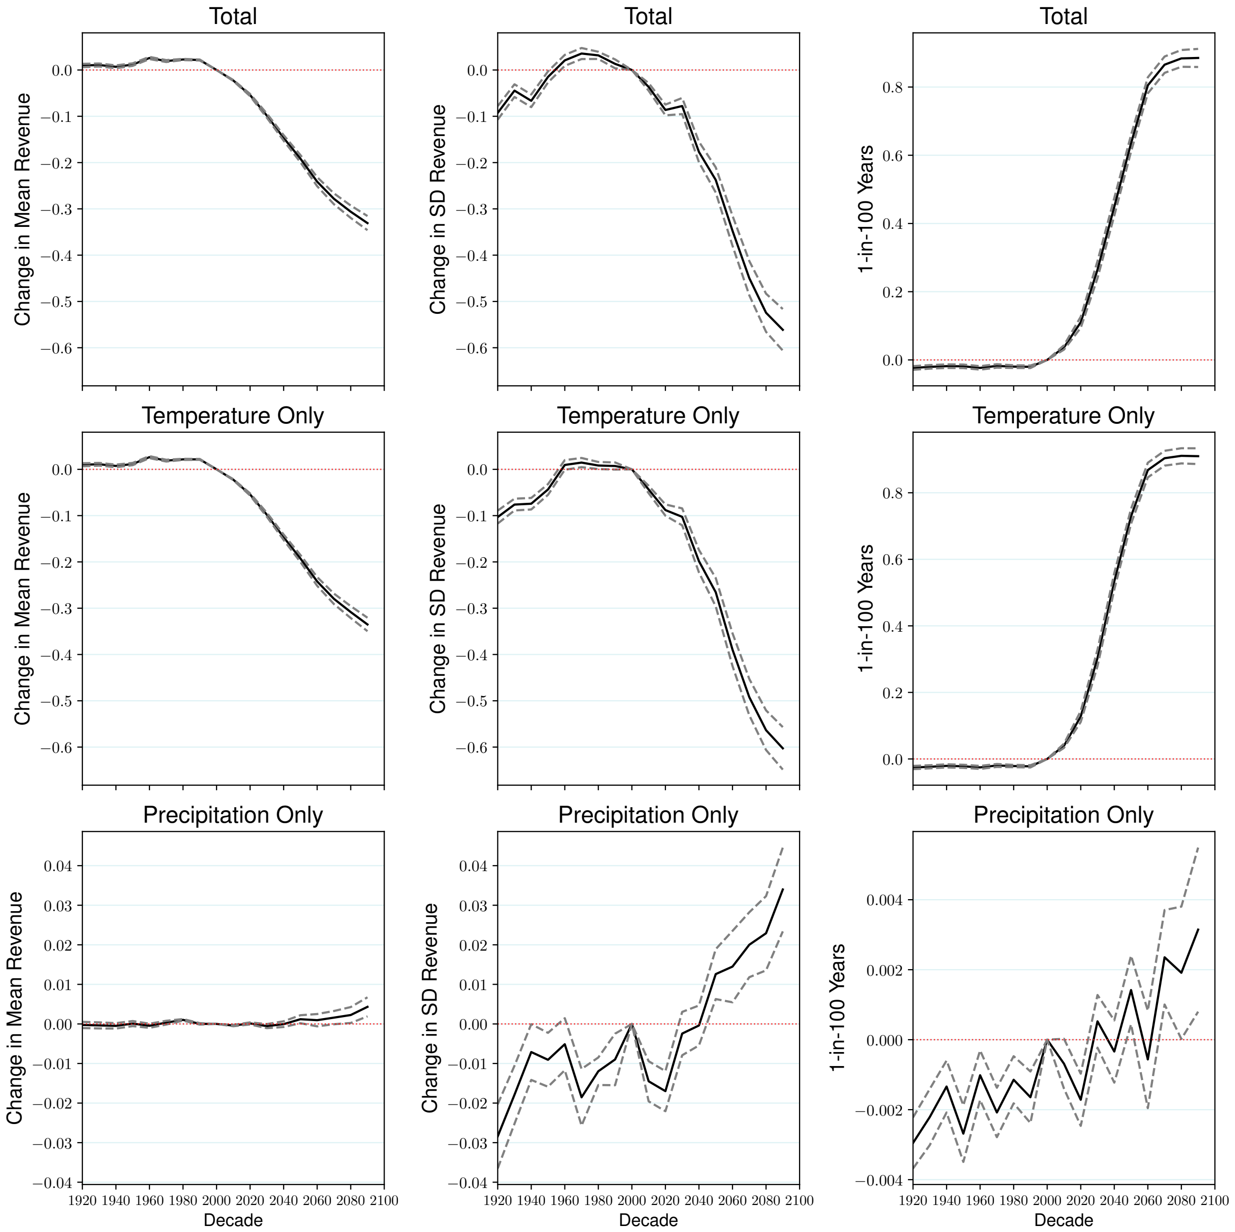

Figure 1-C.7: Projected changes in agricultural revenue – decile temperature bins specification including HDD variable without additional precipitation variables

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.

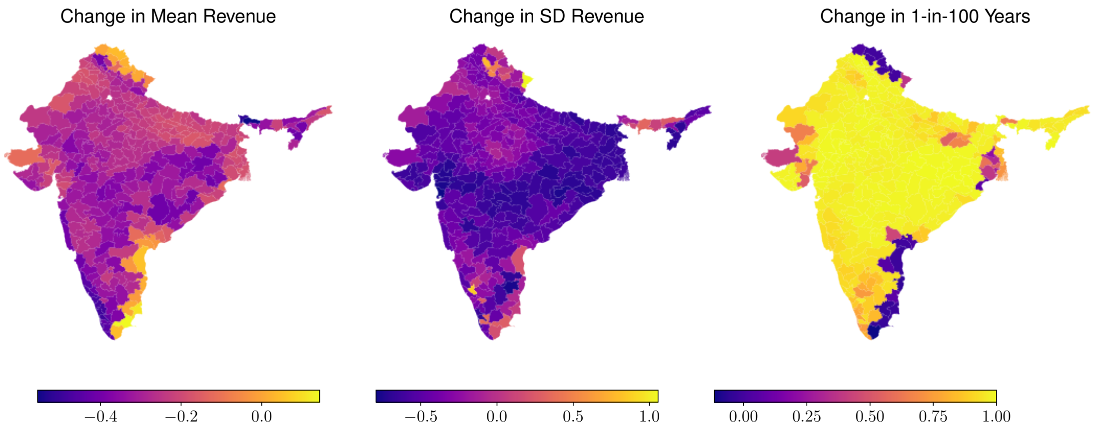

(a) Temperature Bins

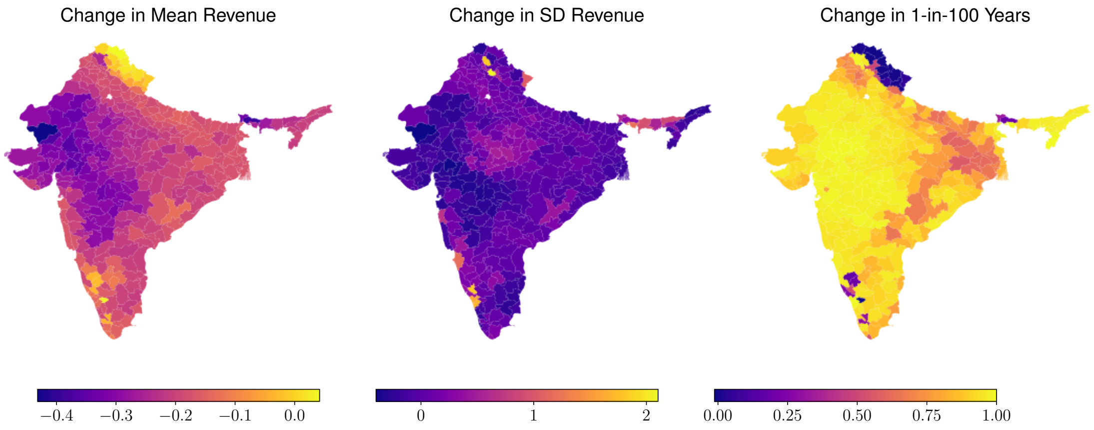

(b) Degree Days

Figure 1-C.9: 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.

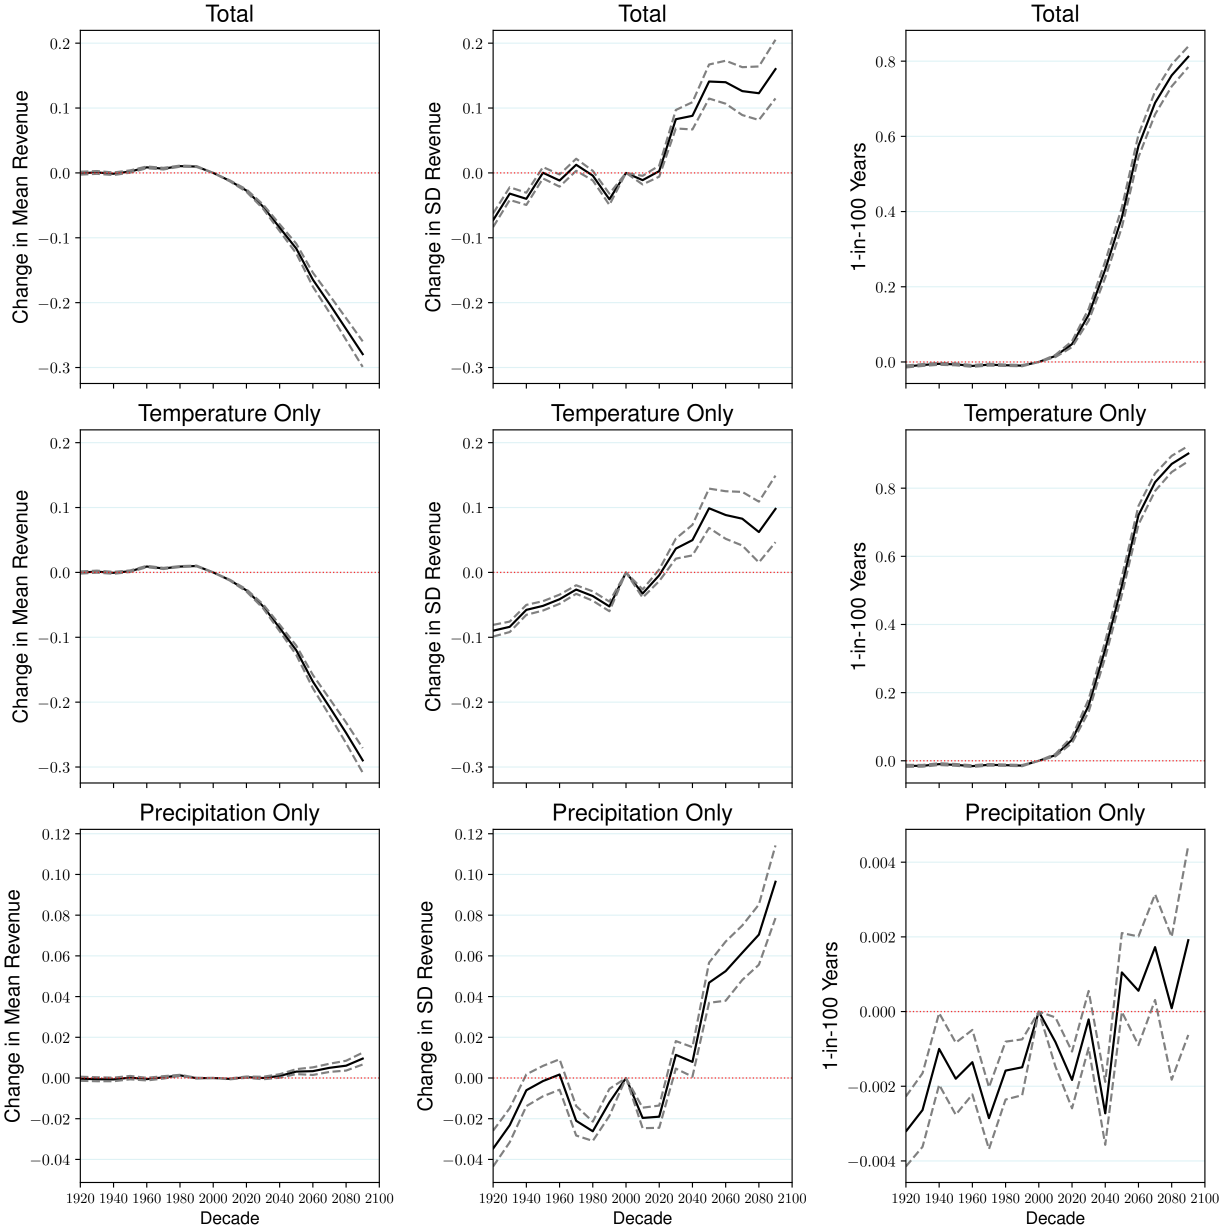

Figure 1-C.10: Projected changes in agricultural revenue – quadratic degree-days specification without additional precipitation variables

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.

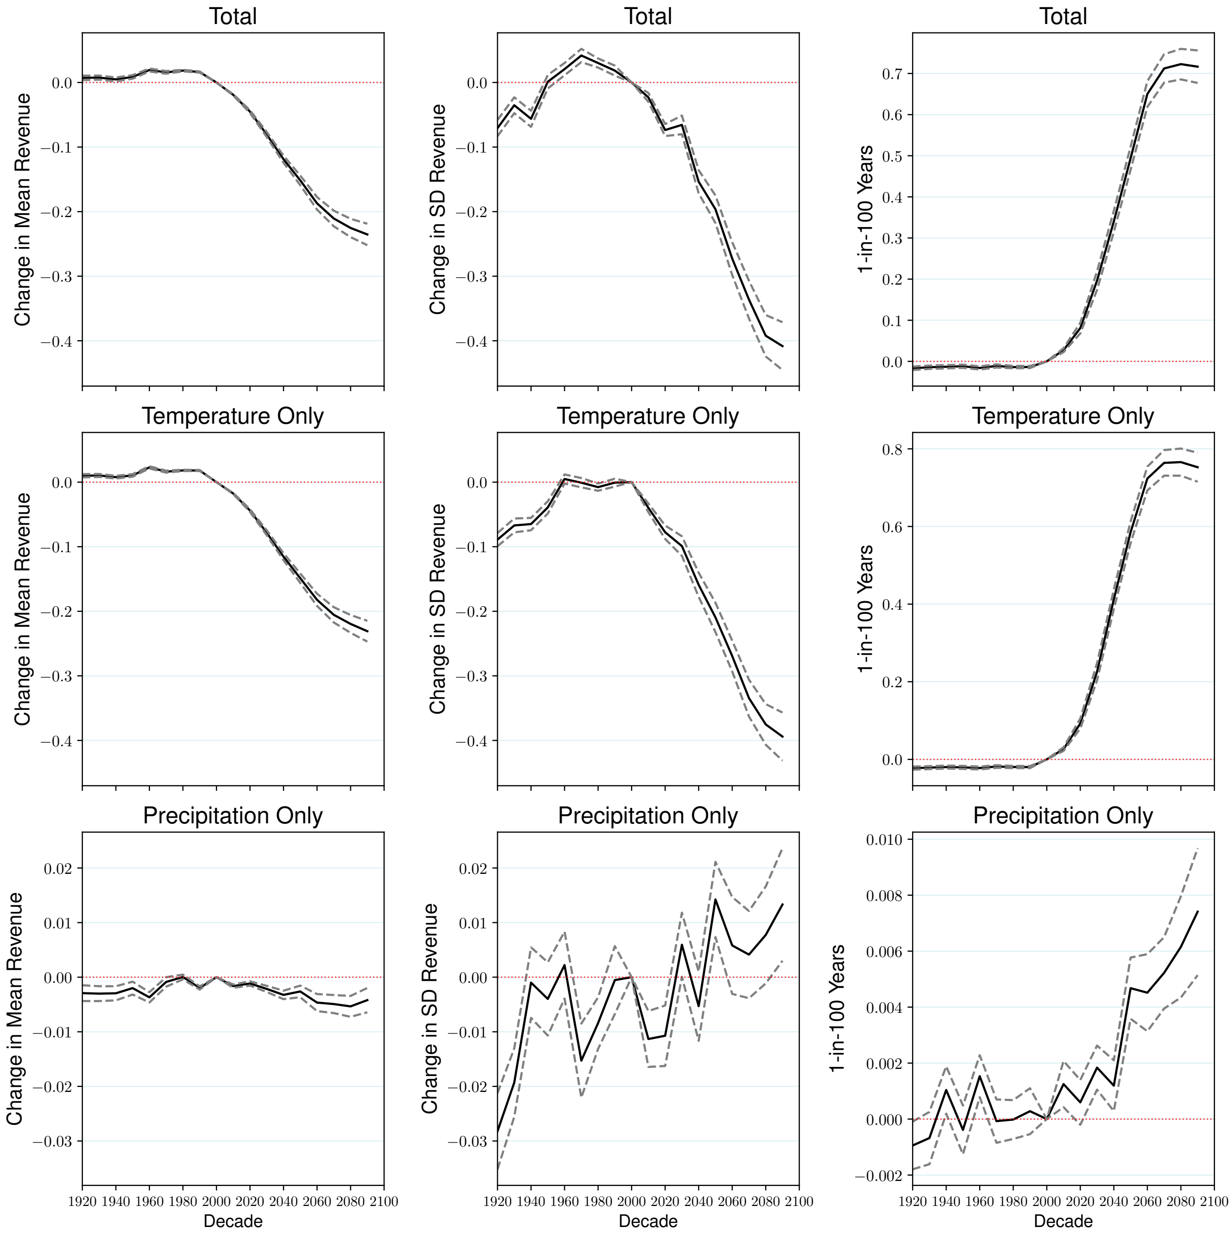

Figure 1-C.11: Projected changes in agricultural revenue – decile temperature bins specification with additional precipitation variables

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.

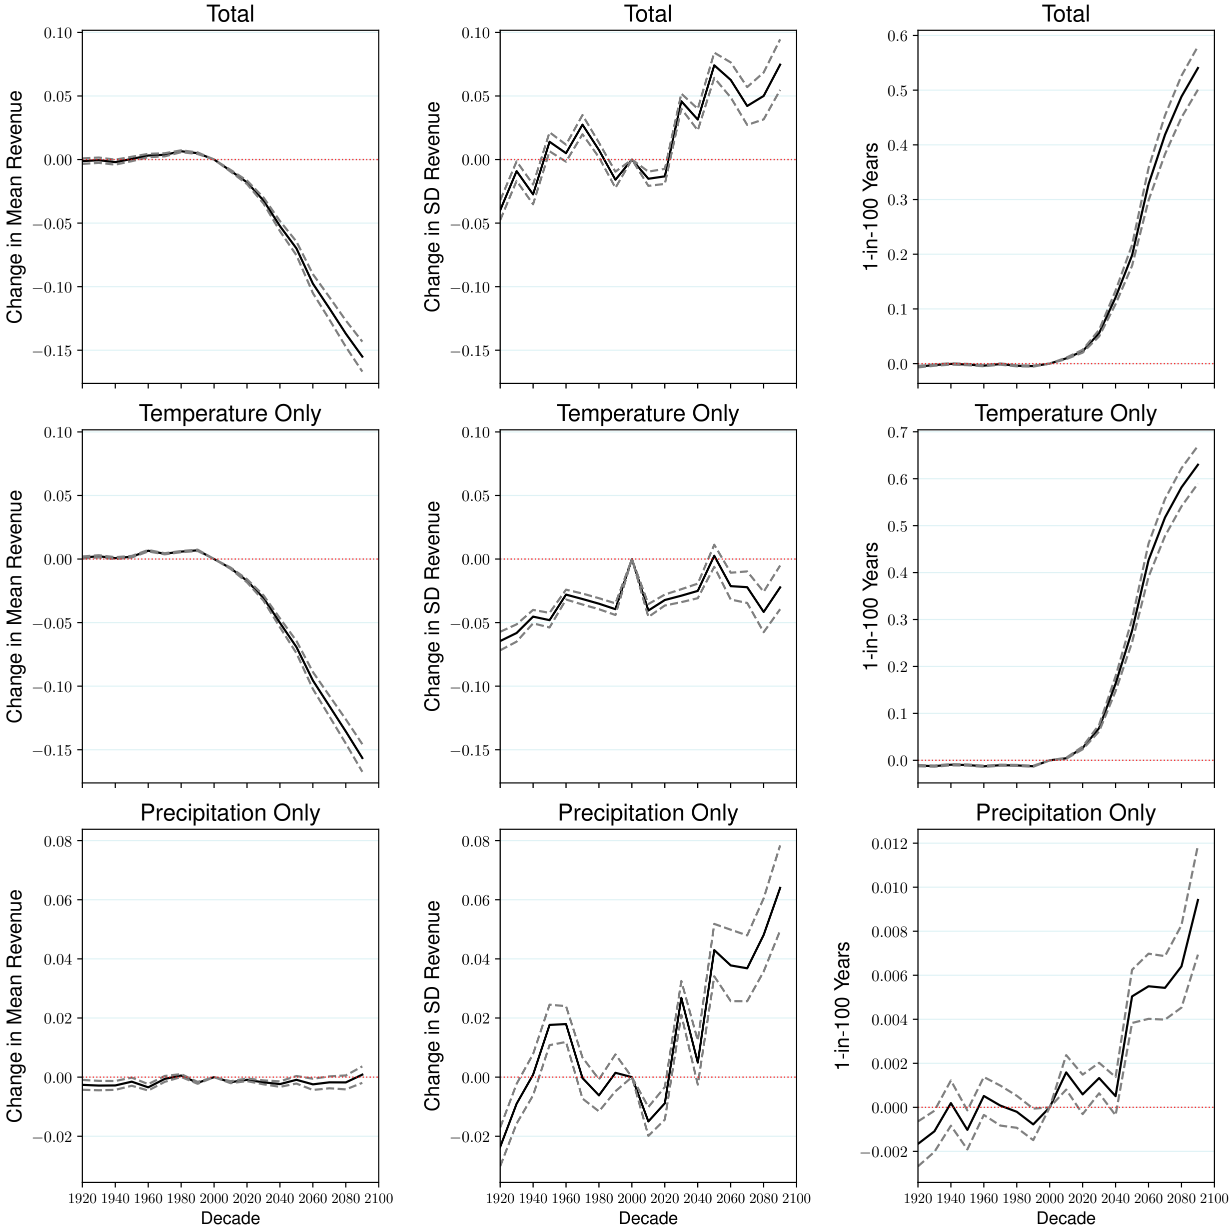

Figure 1-C.12: Projected changes in agricultural revenue – linear degree-days specification with additional precipitation variables

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.

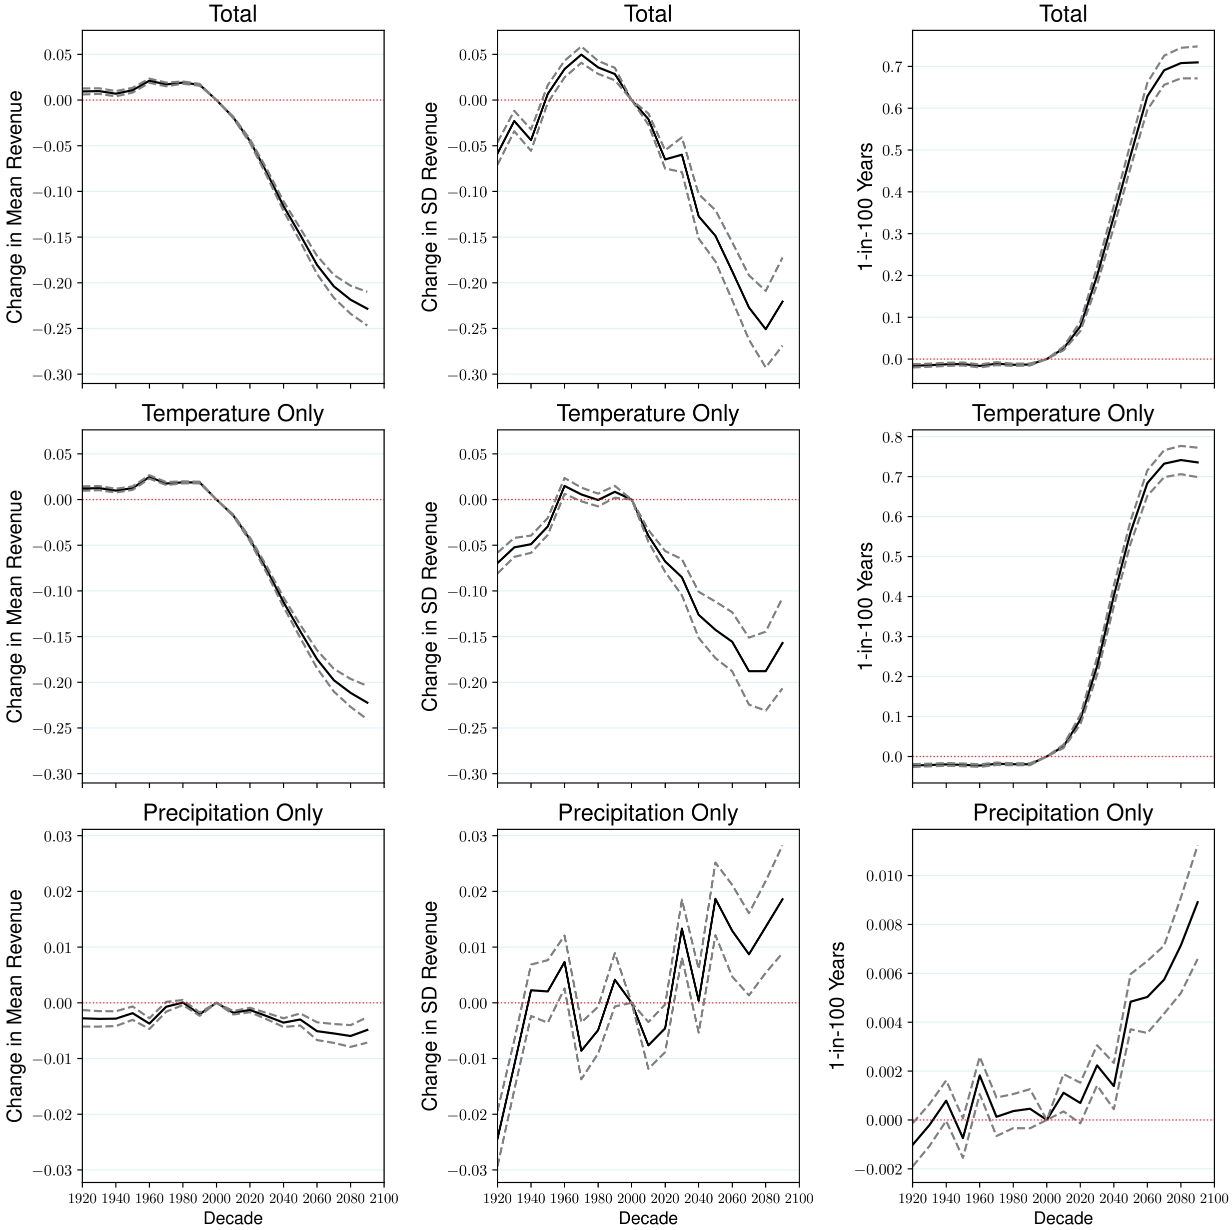

Figure 1-C.13: Projected changes in agricultural revenue – three-degree temperature bins specification with additional precipitation variables

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.

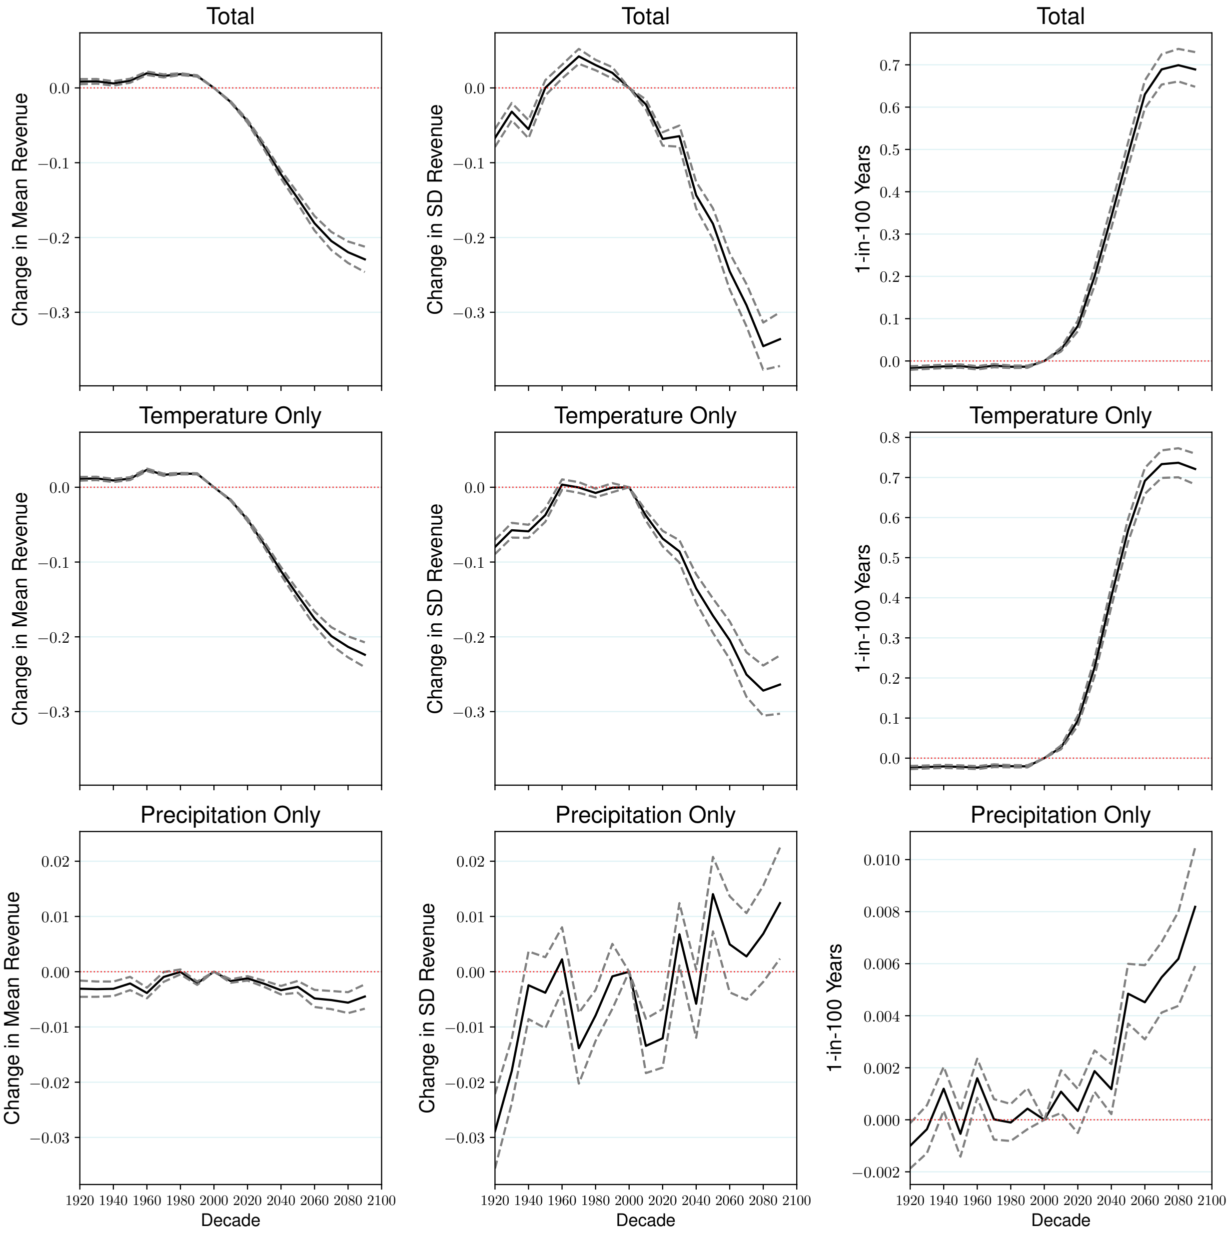

Figure 1-C.14: Projected changes in agricultural revenue – decile temperature bins specification including HDD variable with additional precipitation variables

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.

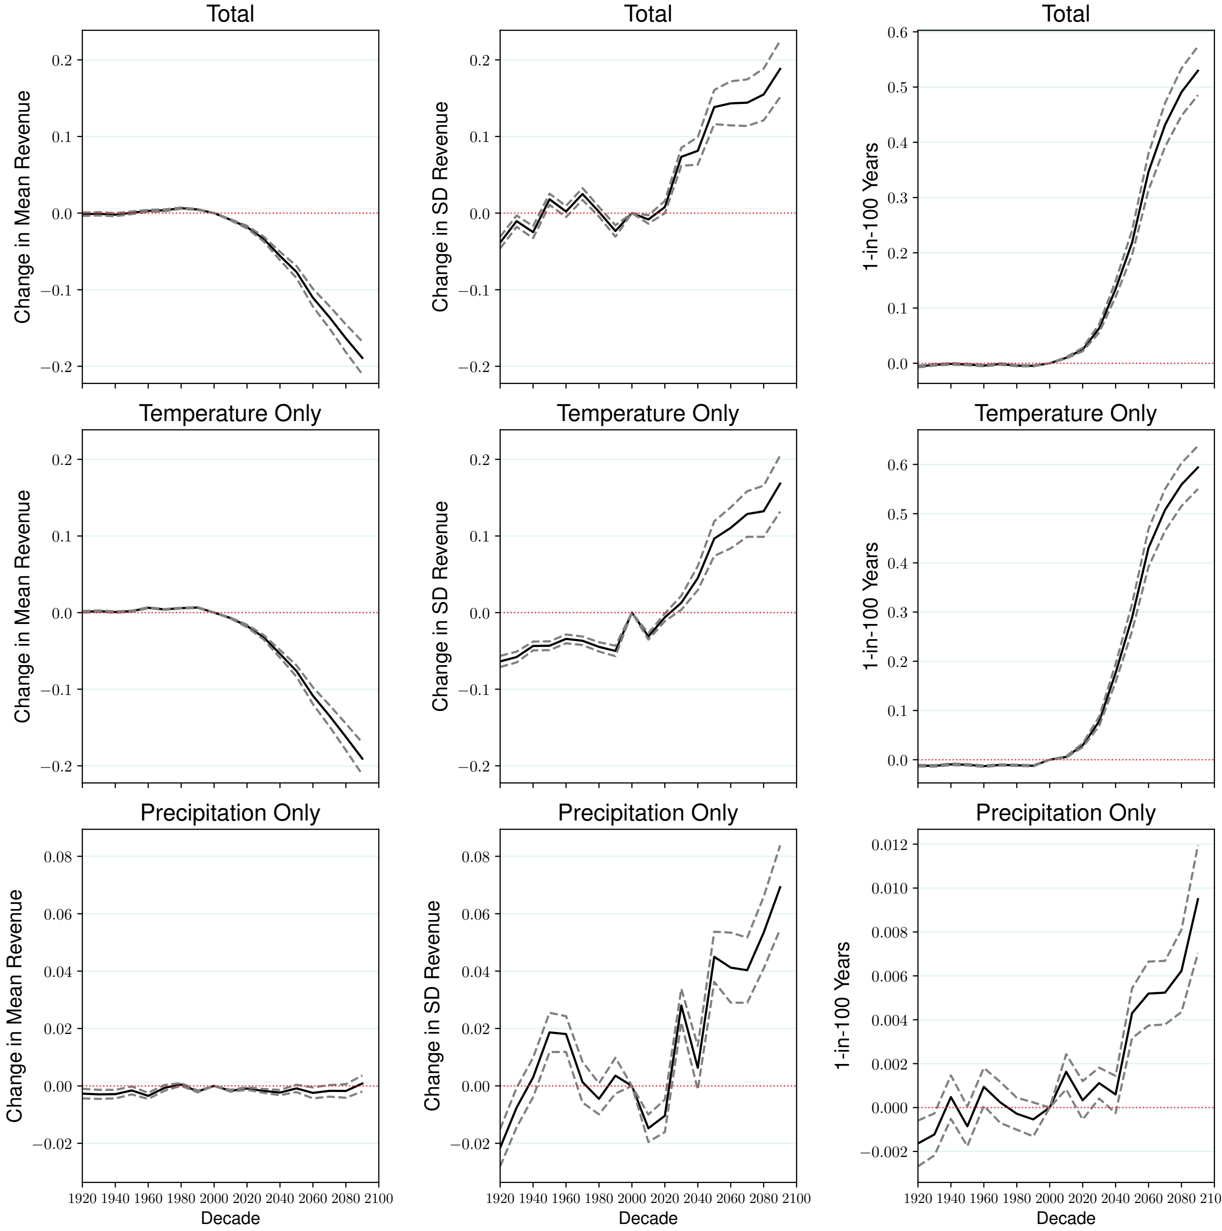

Figure 1-C.15: Projected changes in agricultural revenue – quadratic degree-days specification with additional precipitation variables

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.

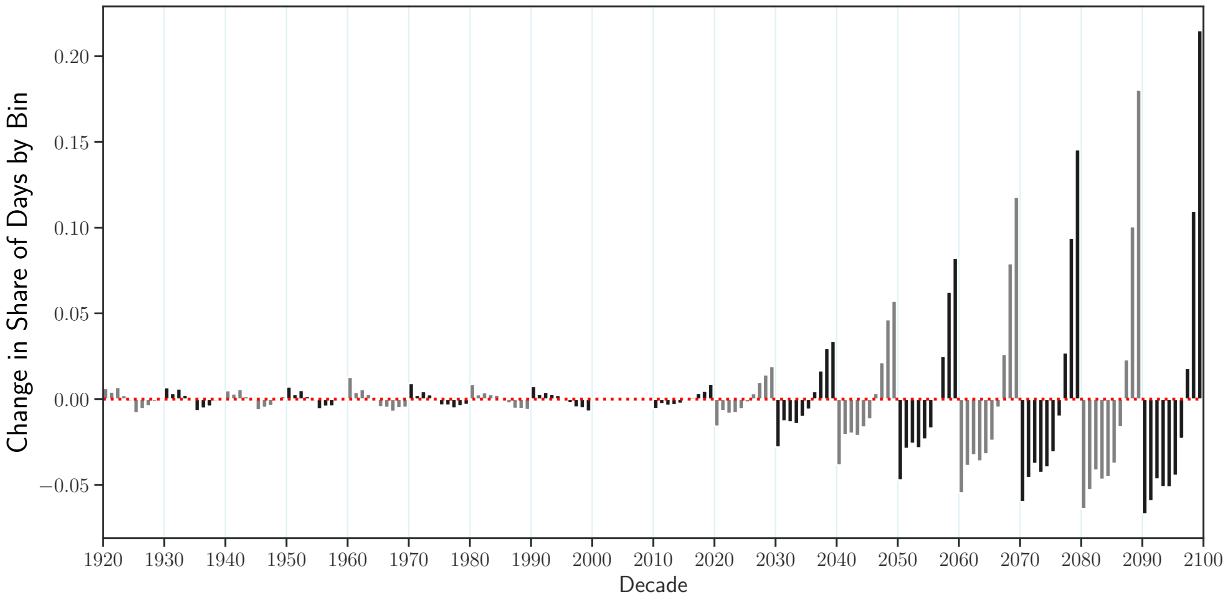

Figure 1-C.16: Share of days in each decile temperature bin relative to the 2000s

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 1-C.2 for the limits of each bin in degrees Celsius.

Table 1-C.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.

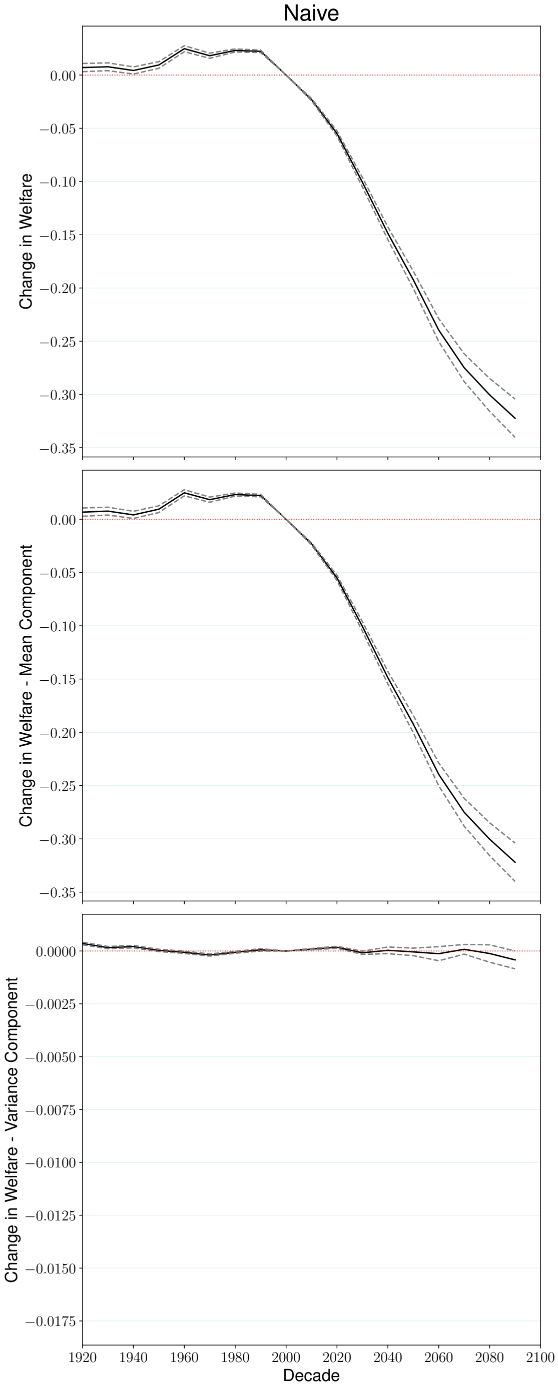

Figure 1-C.17: Projected changes in welfare measures – three-degree temperature bins specification without additional precipitation variables

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 1.5. 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.

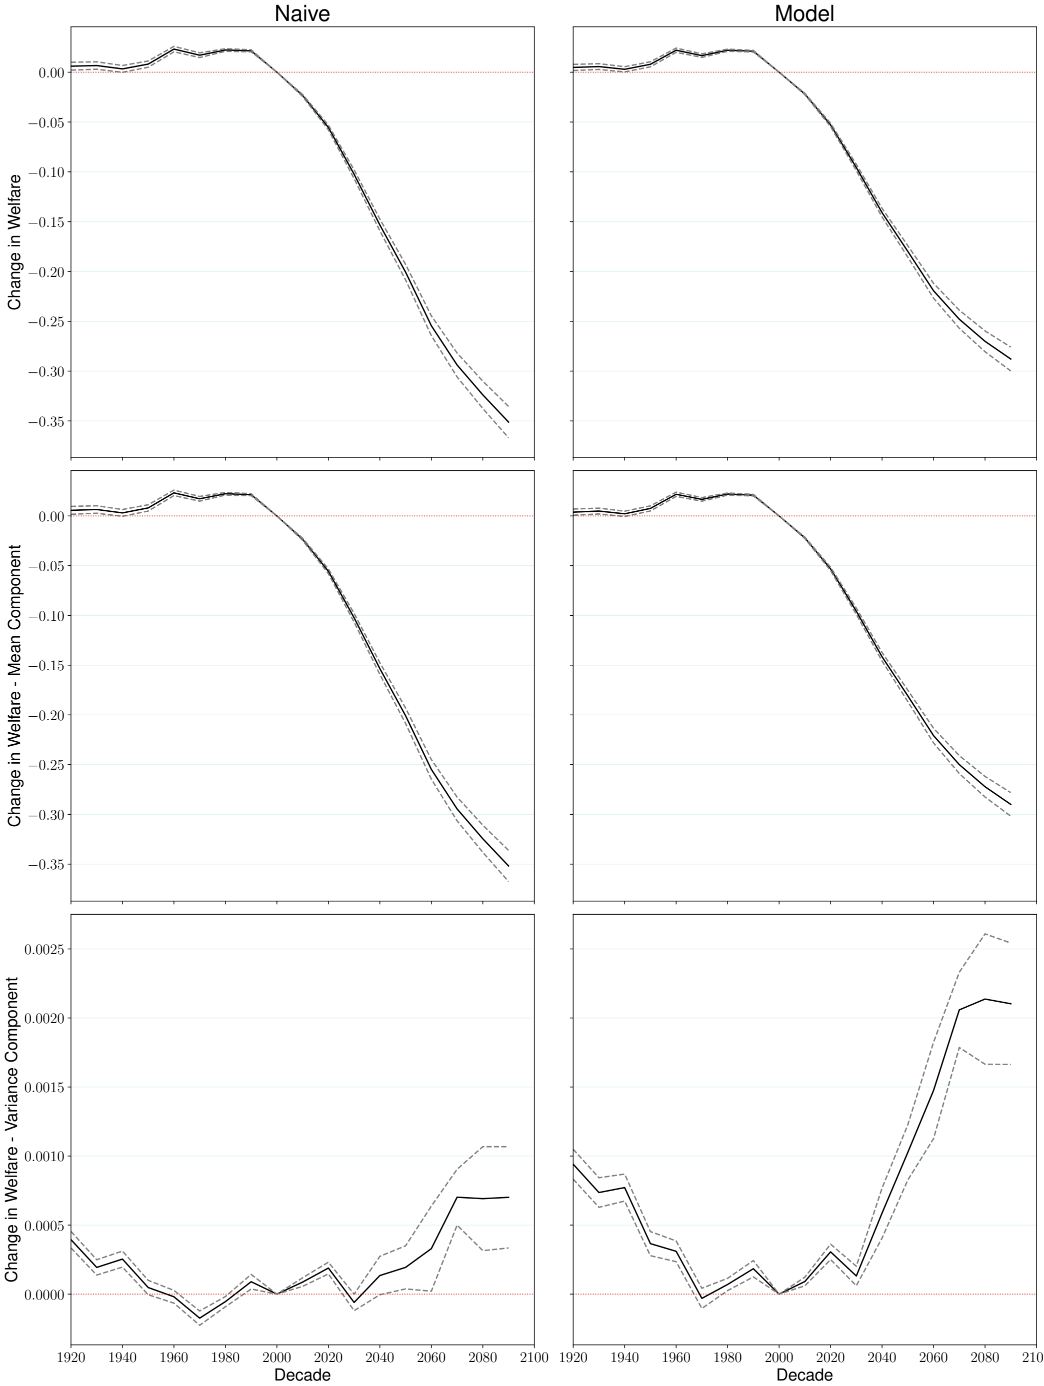

Figure 1-C.18: Projected changes in welfare measures – decile temperature bins specification including HDD variable without additional precipitation variables

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 1.5. 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.

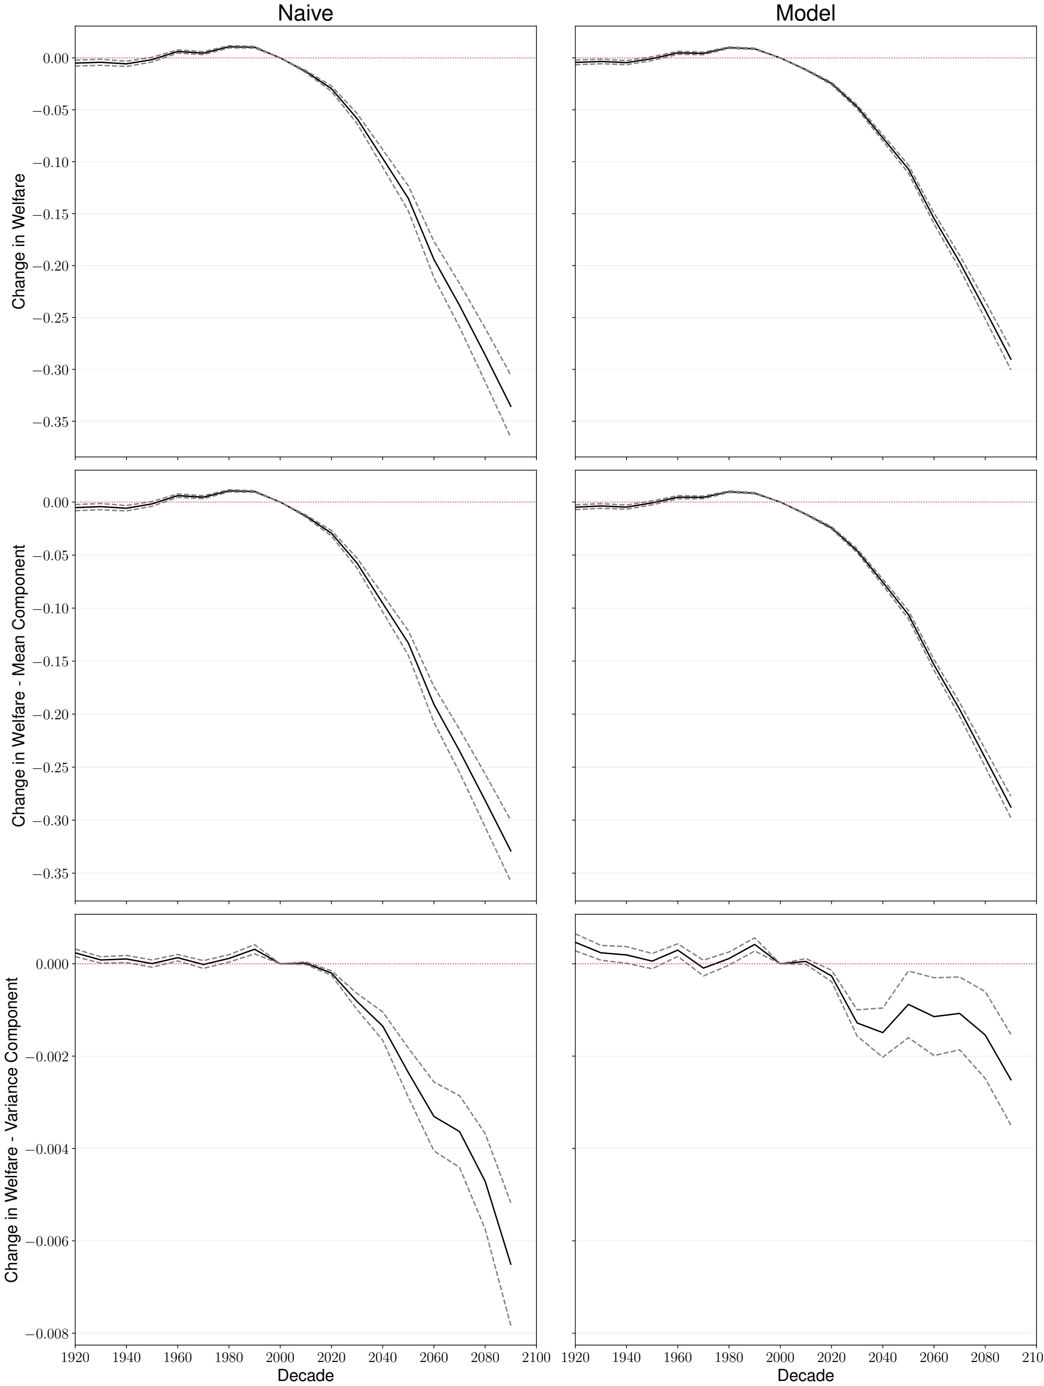

Figure 1-C.19: Projected changes in welfare measures – quadratic degree-days specification without additional precipitation variables

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 1.5. 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.

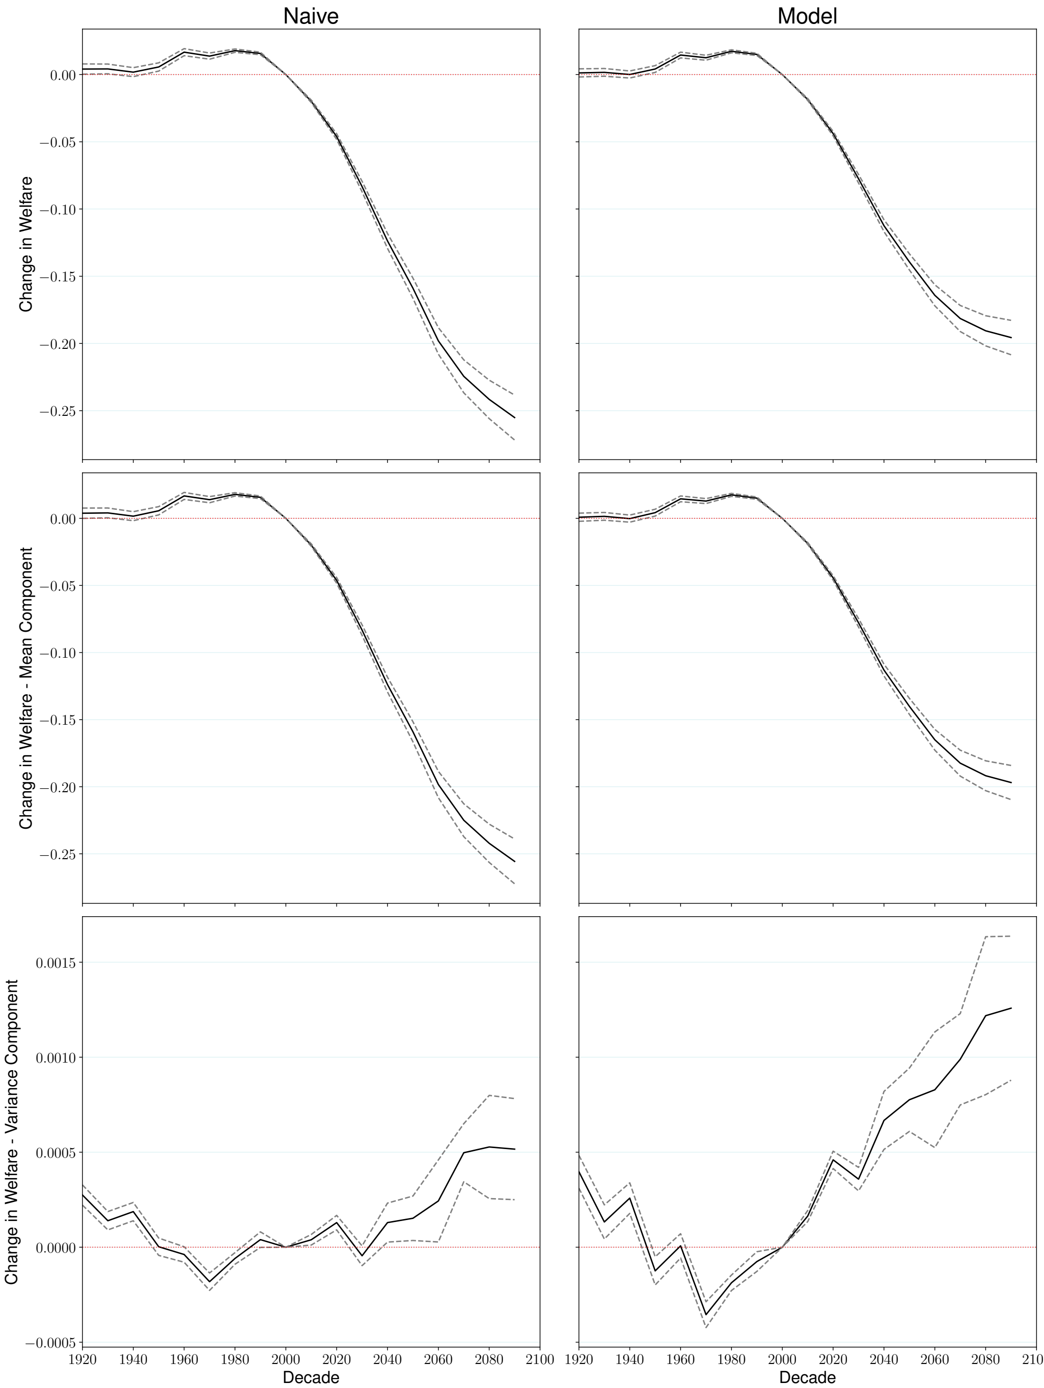

Figure 1-C.20: Projected changes in welfare measures – decile temperature bins specification with additional precipitation variables

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 1.5. 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.

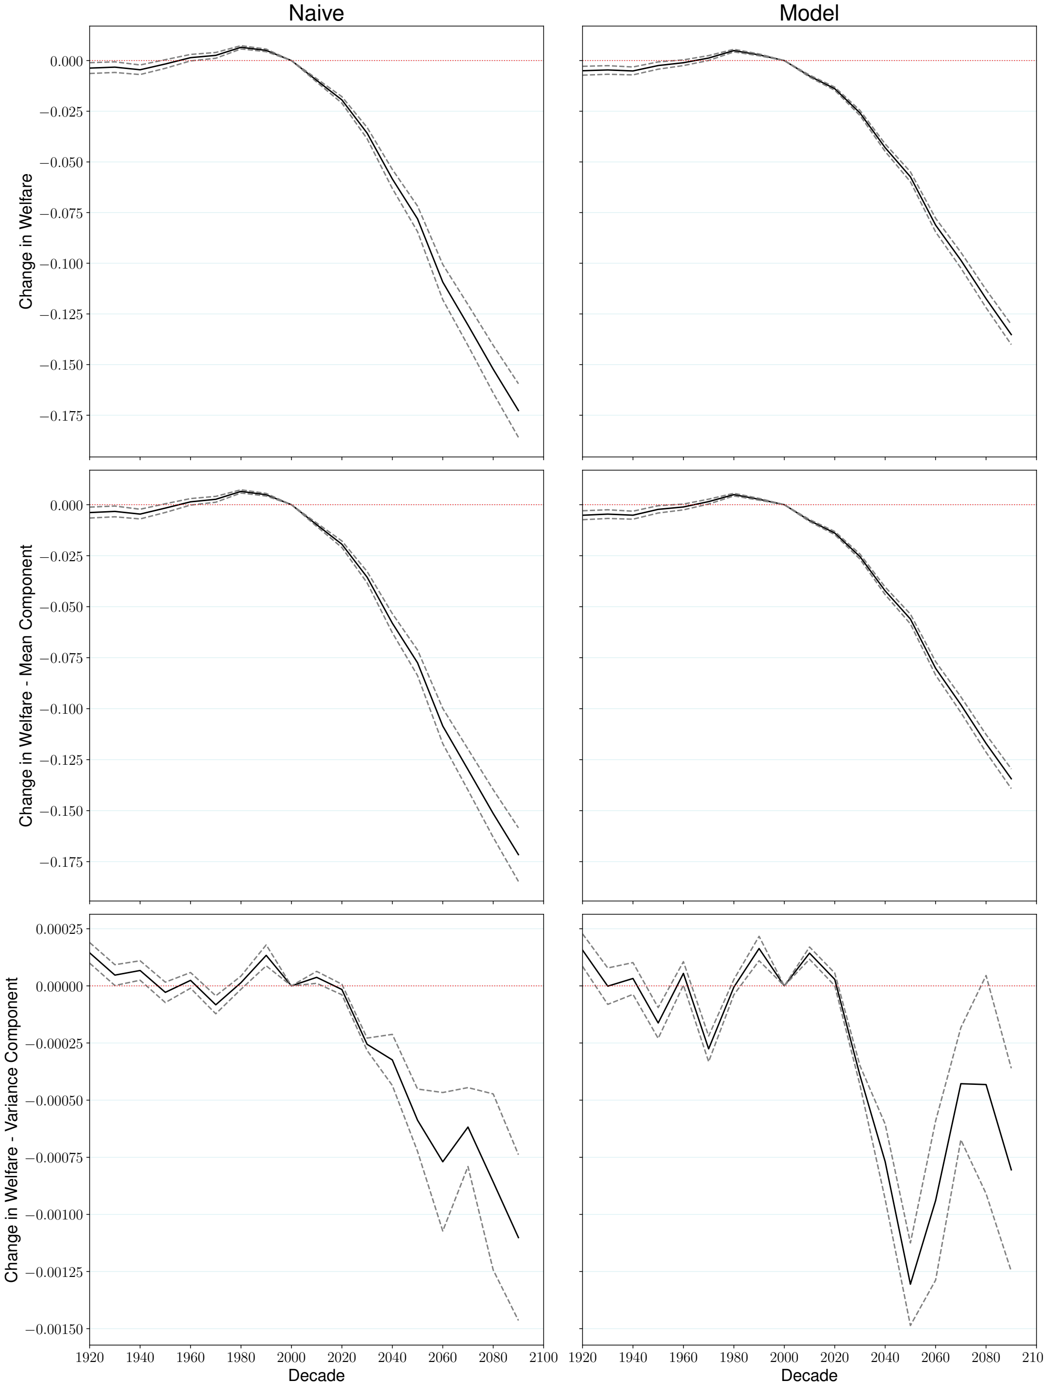

Figure 1-C.21: Projected changes in welfare measures – linear degree-days specification with additional precipitation variables

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 1.5. 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.

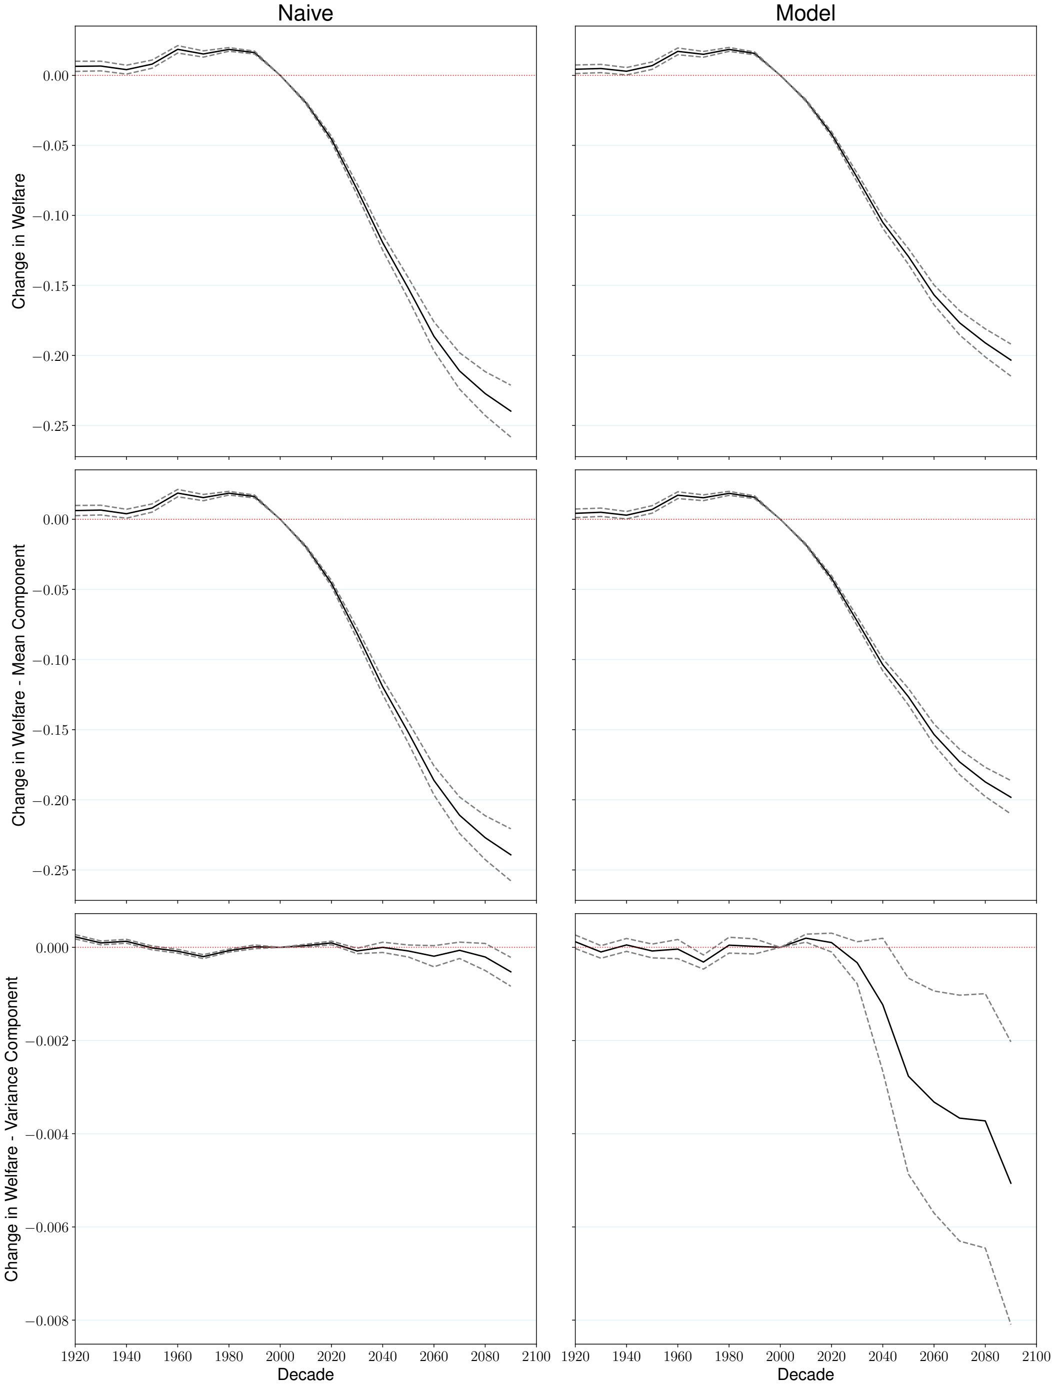

Figure 1-C.22: Projected changes in welfare measures – three-degree temperature bins specification with additional precipitation variables

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 1.5. 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.

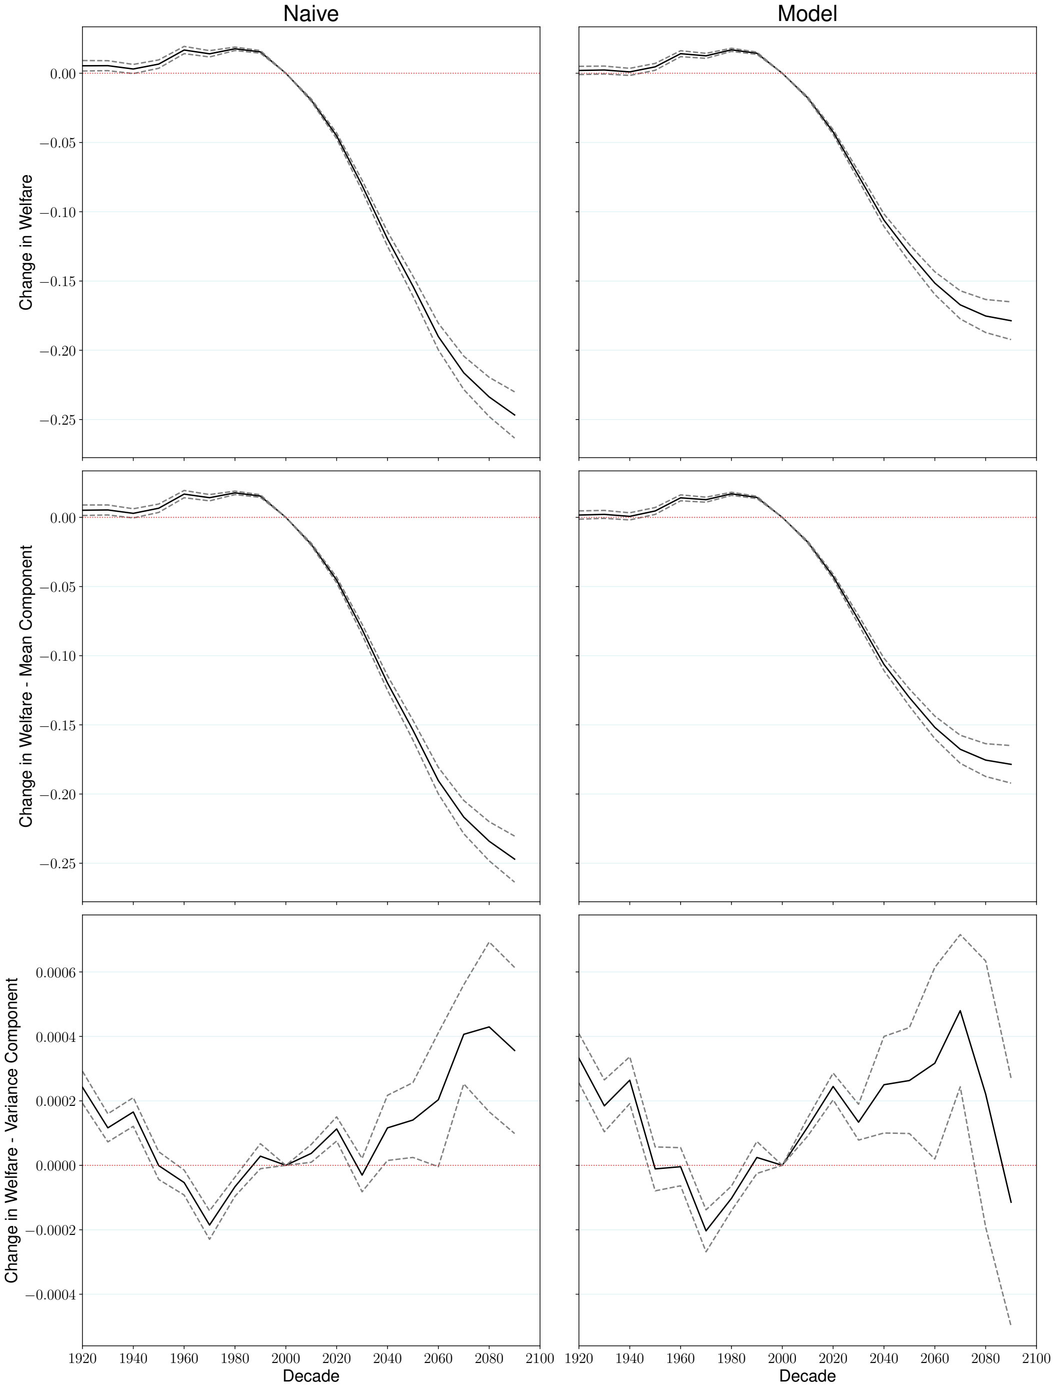

Figure 1-C.23: Projected changes in welfare measures – decile temperature bins specification including HDD variable with additional precipitation variables

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 1.5. 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.

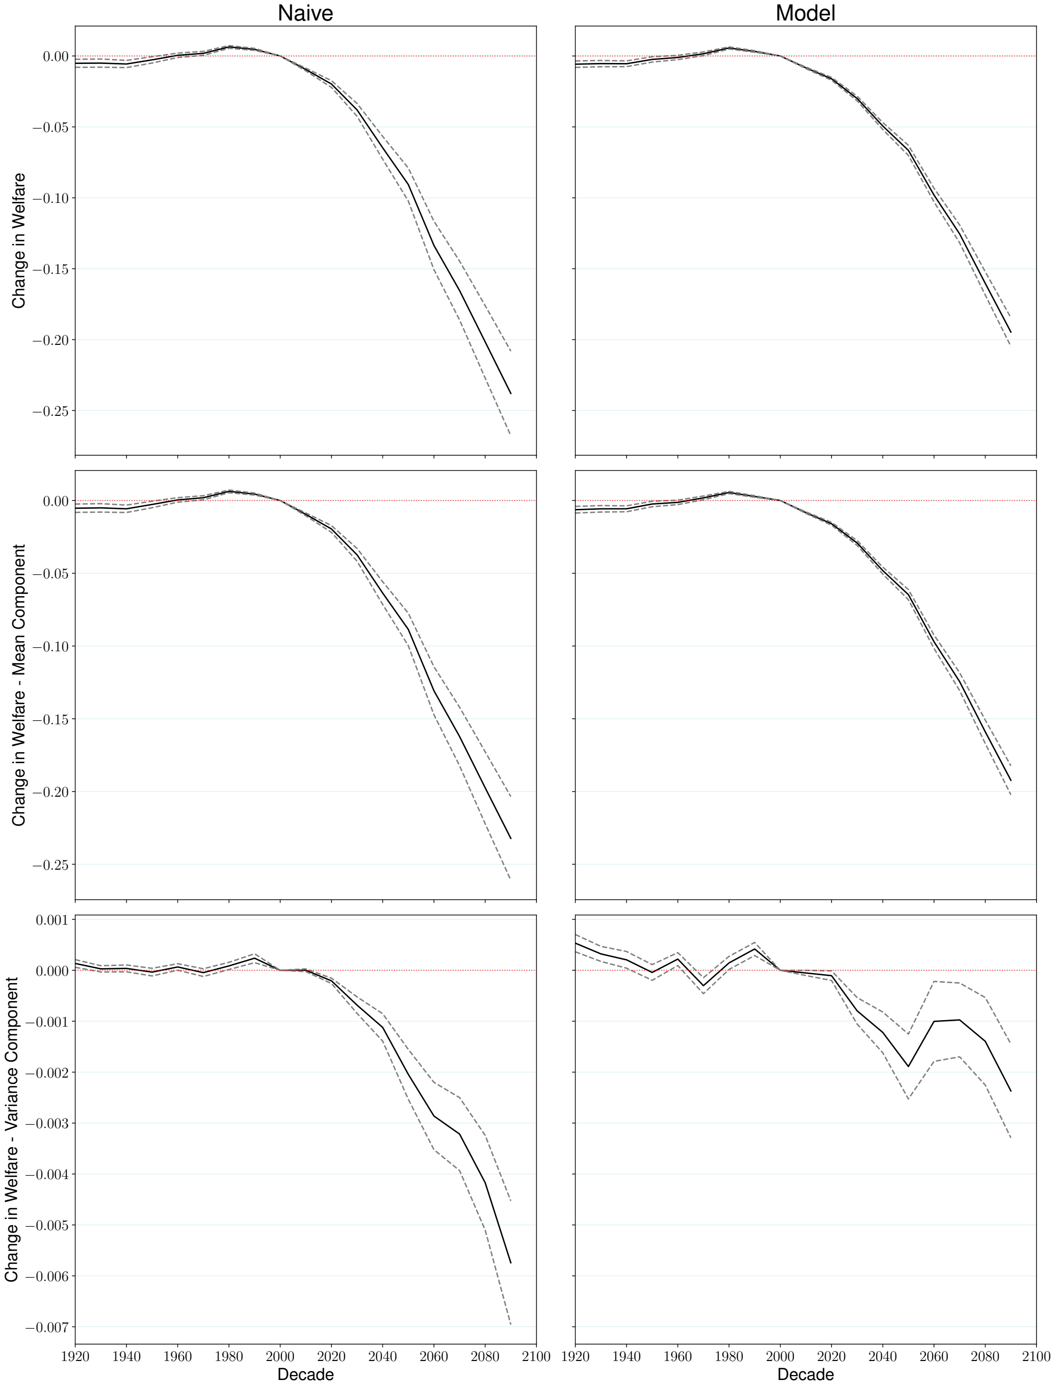

Figure 1-C.24: Projected changes in welfare measures – quadratic degree-days specification with additional precipitation variables

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 1.5. 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 1-C.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 (1.1) using *decile temperature bins without additional precipitation variables*. Definitions of the variables may be found in Section 3. See also Appendix Table 1-C.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 1-C.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 $(-\infty,0)$			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 $[39,\infty)$	-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 (1.1) 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 1-C.5: Yield-weather estimation – decile temperature bins including HDD variable

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	Castor	Cotton	Finger Millet	Groundnut	Maize	Pearl Millet	Pigeon Pea	Rice	Sesame	Sorghum	Barley	Chickpea	Linseed	Rapeseed/Mustard	Wheat	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 (1.1) 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 1-C.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 1-C.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 (1.2) 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 1-C.7: Yield-weather estimation – quadratic degree days

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	Castor	Cotton	Finger Millet	Groundnut	Maize	Pearl Millet	Pigeon Pea	Rice	Sesame	Sorghum	Barley	Chickpea	Linseed	Rapeseed/Mustard	Wheat	Sugarcane
Degree days	-0.000370	-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 (1.2) 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 1-C.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 (1.1) using *decile temperature bins with additional precipitation variables*. Definitions of the variables may be found in Section 3. See also Appendix Table 1-C.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 1-C.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 (1.2) 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 1-C.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 $[39,\infty)$	-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 (1.1) 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 1-C.11: Yield-weather estimation – decile temperature bins including HDD variable with additional precipitation variables

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
	Castor	Cotton	Finger Millet	Groundnut	Maize	Pearl Millet	Pigeon Pea	Rice	Sesame	Sorghum	Barley	Chickpea	Linseed	Rapeseed/Mustard	Wheat	Sugarcane
Bin [min,p10)	0.345	-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 (1.1) 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 1-C.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 1-C.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 (1.2) 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 $^{***}\, p < 0.01$.

Chapter 2

CCTs and Fertility: Long-Term Impacts Across Two Generations

By Fabien Forge

Abstract

This paper investigates the relationship between income, education and fertility by looking at the long term impact of the Mexican conditional cash transfer program. To do so, I define two cohorts of women that were exposed to the program in two different ways. Treated women, from the older cohort, received cash transfers on the condition that they sent their children to school. The younger cohort, first received extra education, and then faced the same treatment conditions as the older cohort. I observe the fertility of these two cohorts 13 years after the program started using the Integrated Public Use Microdata Series (IPUMS) data. The difference-in-difference identification strategy relies on cross-census variations between 2010 and 1990 and spatiotemporal variations in the roll-out of the program at the municipality level. Controlling for time varying effects of belonging to a poor municipality, my results suggest that far from creating an incentive to have more children, women exposed to *Progresa* lowered their fertility in a small but significant way. Unlike what a quality-quantity trade-off model would suggest, I fail to find an additional effect of education on the fertility of the younger cohort.

2.1 Introduction

Starting with Malthus' seminal work, the link between fertility and development is one of the oldest issues in the development literature, yet both the direction of the relationship and its microeconomic determinants remain unclear. Far from the Malthusian view of the world, higher fertility rates may be the consequence rather than the cause of underdevelopment (85; 50). Conditional Cash Transfer (CCTs) programs present a unique opportunity to revisit this relationship. Twenty-three years ago, Mexico implemented such a program, *Progresa*, that combines immediate poverty relief via cash transfers with longer term goals of poverty reduction via mandatory schooling for children whose parents were enrolled.

In this paper, I investigate how income and education can influence fertility using the Mexican program as a treatment. I test whether this program has tended to reduce fertility for treated women and use the long period of time since the creation of the program to define two types of treated groups. The first group, labelled the *older cohort*, corresponds to women who were potentially enrolled in the program when it started but, because of their age, were not eligible to receive extra schooling. The second group, labelled the *younger cohort*, differ from the first by the fact that these women were young enough, when the program started, to first be treated as school-attending children, and then as parents in the same way as the *older cohort*. Thus, the *younger cohort* potentially captures the additional effect of being more educated when making fertility decisions and speaks to the inter-generational evolution of poverty.

I first motivate theoretically why and how a program such as *Progresa* may reduce fertility in the context of a standard quality-quantity trade-off model commonly used in the literature. I embed the cash transfer and the mandatory education components of the program into the overlapping generations model developed in (76). I show that the model predicts a reduction in fertility, compared with a situation with no program, for both the *younger* and the *older cohorts* provided that the cash transfer is small enough relative to the impact of the "mandatory" investment in education. I then analyse the conditions under which members of the *younger cohort* could further decrease their fertility compared to the *older cohort* as a result of their increased schooling level.

Empirically, I use a difference-in-difference strategy which takes advantage of spatiotemporal variations in the program's expansion. I link this municipality level exposure measure to the Integrated Public Use Micro Data (IPUMS) to observe the fertility levels of women across decennial censuses 13 years after the onset of the program for the post-treatment group and 7 years before for the pre-treatment group. This use of cross-census variation is dictated by the impossibility of using within census variation, especially for the *younger cohort*, since age jointly determines exposure to the treatment and the outcome variable.[1] Instead, I compare how fertility of women within a given age bracket differentially evolved between 1990 and 2010 for treated compared to control cohorts depending on the proportion of households treated changes in the municipality these women belong to. I construct this ratio using information from the Mexican evaluation organization and the Mexican National Institute of Statistics and Geography for the 2,392 municipalities in my data.

An important threat to the identification comes from the initial treatment of the program being particularly targeted at the poorest areas of Mexico. These areas being the poorest, there is a risk that fertility was already converging towards the Mexican mean, even in the absence of a treatment. This would violate the parallel trends assumption needed for the difference-in-difference estimate to be identified. To circumvent this issue I follow the strategy proposed by Parker and Vogl (83) and use a unique feature in the roll-out of the program. Specifically, the timing of the expansion was shaped by Mexican national elections during which the law forbids the extension of social programs, resulting in three distinct waves of expansion. The first two waves were separated by 5 years and targeted similar households according to the "marginality index" used to chose the first recipients, while the third wave to urban and richer areas. I argue that controlling for exposure intensity in 2005 and the marginality index allows me to account for fertility dynamics linked to poverty and thus capture the true effect of the treatment.

My findings suggest that far from creating an incentive to bear more children in order to receive additional transfers, *Progresa* reduced fertility for both the *younger* and the *older cohorts* in a way predicted by my model. A one percentage point increase in the 1999 coverage of households in a municipality is associated with a decrease in fertility by 0.002 to 0.003 fewer children for the *older cohort*. In the data, these are women, aged 16 to 36 when the program started in 1997, who are observed when they were aged 29–49 in the 2010 census, compared with women of the same age in the 1990 census. For the *younger cohort* aged 7 to 15 in 1997 and observed when they are 20 to 28 in 2010, this number ranges from 0.0001 to 0.002. These results are robust to placebo tests in which the same cross-census comparisons are made for women who already reached the end of their fertility window when the program started. I also provide suggestive evidence this reduction in fertility is mainly coming from adjustments along the intensive margins.

---

[1] This is not true of other outcome variables, such as education, which no longer depends on age once the schooling years are over.

I do not find strong evidence of an additional effect of the *young cohort*'s own education on fertility, using both a within census difference-in-difference strategy in which the *older cohort* constitutes the control group and the *young cohort* the treated group and a triple-difference strategy across censuses. This is suggestive evidence that the fertility decline experienced by both cohorts was in response to changes in the incentives relative to investment in education. The extra education received by the *younger cohort* as a result of their enrollment in the program does not seem to have additionally affected their fertility in a way that would be suggested by an inter-generational quality-quantity trade-off model. It is worth noting that since the *younger cohort* is observed, at still a relatively young age, there is a possibility that the full effect of additional education could be observed closer to the end of the fertility window of these women.

This paper builds on several strands of the literature. First, a large microeconomic literature has linked both monetary and non-economic factors such as female empowerment (access to contraceptives, education) (42), level of income (15), health (child mortality and life expectancy) (Doepke 41; Jayachandran and Lleras-Muney 62), and demand for human capital (51) with fertility. Second, many impact evaluation papers[2] have documented direct effects of CCT programs such as improved school indicators for boys and girls (Coady and Parker 29; Schultz 93; Behrman et al. 17; Todd and Wolpin 102) or indirect effects such as decreased child labor participation (Skoufias et al. 95; Rubio-Codina 87; Schultz 93), improved health for children (Gertler 53; Fernald et al. 44; Fernald et al. 45; Barham 12) and better natal care (Urquieta et al. 105, Barber 11).

Several papers have looked at the impact of conditional cash transfers on fertility. Considering the short run impacts of the CCT programs of Mexico, Nicaragua and Honduras, Stecklov et al. (96) found no effect on fertility within the 24 months following the treatment for the first two and a positive effect in Honduras. Both Baird et al. (10) and Schultz (93) found that teenage pregnancy was reduced a year after the program started in Malawi and Mexico. Taking a longer term perspective Barham et al. (13) documented a reduction in teenagers' fertility for 42 communities in Nicaragua 10 years after the program started while Laszlo et al. (70) found no effect on fertility for the Peruvian program.

Finally, this paper contributes to a nascent strand of the literature that gauges whether the identified short-run positive effects of *Progresa* has persisted since the program started in 1997 (Parker and Vogl 83, Kugler and Rojas 69). It also responds to the predictions made by Todd and Wolpin (102) who predicted that treated Mexican women would keep their fertility at the same level or slightly higher.

---

[2] For a thorough review of the Mexican program see Parker and Todd (82)

The rest of the paper is organized as follows. Section 2.2 reviews the main features of the program, defines the *older* and *younger cohorts* and documents the main determinants of the roll-out. It then frames the link between the program and fertility using a variant of the workhorse quality-quantity trade-off model. Section 2.3 describes the data used to define the fertility and enrollment ratio measures used in the main specification. Section 2.4 presents the results for both cohorts and provides suggestive evidence on the potential mechanisms.

2.2 Context and Theoretical Motivation

This section provides a brief overview of an already widely studied program and links it to the workhorse quality-quantity trade-off model traditionally used when thinking about the link between poverty and fertility. In Section 2.2.1, I highlight the main features of *Progresa* and describe how, from a long term perspective, two analytically distinct treated groups emerge which I label the *younger* and *older cohorts*. I also describe the determinants of the roll-out of the program that inspire the identification strategy. In Section 2.2.2, I build on the overlapping generations model proposed by Moav (76) and include the two main components of the program, income transfers and conditionality, to predict potential changes in fertility for both the *younger* and the *older cohorts*.

2.2.1 Context

Main features – The Mexican conditional cash transfer program *Progresa* started in 1997 and is, with the program *Bolsa Familia* in Brazil, one of the oldest CCT programs. This anti-poverty intervention tries to combine short-term financial support to the poor with better long-term investment in their children.[3] Specifically, the program is meant to encourage families to use existing schooling capacity for the education of their children. Thus, conditional on their children attending school regularly and receiving some health care, mothers would receive cash transfers from the government for each child enrolled.[4] The value of these transfers were designed to match the opportunity cost of child labor.

At the heart of the program, women play a very important role. First, they are the only ones allowed to receive the transfers. Second, they were supported in their ability to care for their children's health and nutrition. They, themselves, were recipients of a series of health care interventions, including pre-natal and post-natal care. Female students also benefited

---

[3] Parker and Todd (82) provide an extensive review of the program and the numerous papers written about it. This section will therefore focus on the main components of the program and how they relate to the research question this paper tries to address.

[4] A detailed breakdown of the grants as of 2003 can be found in table 2-A.1.

from preferential treatment in the form of more generous grants. Girls continuing school after grade 7 were encouraged not to drop out by receiving higher grants than their male counterparts.

The program was also designed to prevent parents from abusing it by having additional children receive benefits. In particular, a cap of three was placed on the number of children that could be enrolled in the program. In addition, in the first years after entering the program it was not possible to add more children. But the key monitoring part of the program has remained the conditionality. Initial take up of the program was very close to 100% suggesting that this program was both understood and in demand.

Expansion – Because of its success, *Progresa* has expanded widely since 1997.[5] Indeed, the program started with 219,944 beneficiary households in that year and counted more than 6 million beneficiaries as of 2010. Which localities would be treated first was determined using a marginality index. This index, using information from the 1990 census, corresponds, for individuals above 15, to the normalized first principal component of: the share of illiterates, without secondary school education and earning less than twice the minimum wage, the share with no access to toilets, electricity or water supply, the share with house crowding (according to the number of rooms per person) and dirt floor, as well as the share living in communities with less than 5,000 inhabitants. Table 2-A.2, reports the summary statistics of these components overall for Mexico.

Importantly for the identification strategy used in this paper, the expansion of the program followed three main waves. Figure 2-1 displays the number of households enrolled yearly in the program at the national level. These waves appear as a result of the Mexican anti-vote buying law that forbids the extension of such program during election years. Thus, the first wave ranges from 1997 to 1999. Its expansion was stopped by the 2000 Presidential election. The second wave began in 2000 and was stopped in 2005 because of the 2006 Presidential election. Finally, the last wave resumed after the 2009 midterm elections.[6]

The first wave ending in 1999 targeted the poorest rural areas. The second ending in 2005, followed the same pattern though it extended to relatively richer areas. Finally, the third wave, in 2010, extended to both richer and more urban areas. Table 2-A.2 presents the correlation between the marginality index and its components with the changes in the enrollment proportion by municipalities.[7] The results presented in this table suggest that,

---

[5] In 2002, the program was rename *Oportunidades*.

[6] Presidential elections occurred in 2000, 2006, and 2012 while midterms occurred in 1997, 2003, and 2009.

[7] As will be described in more details in Section 2.3.1, the enrollment proportion is a measure of the cumulative number of households enrolled in the program in a given year over the total number of households in that municipality.

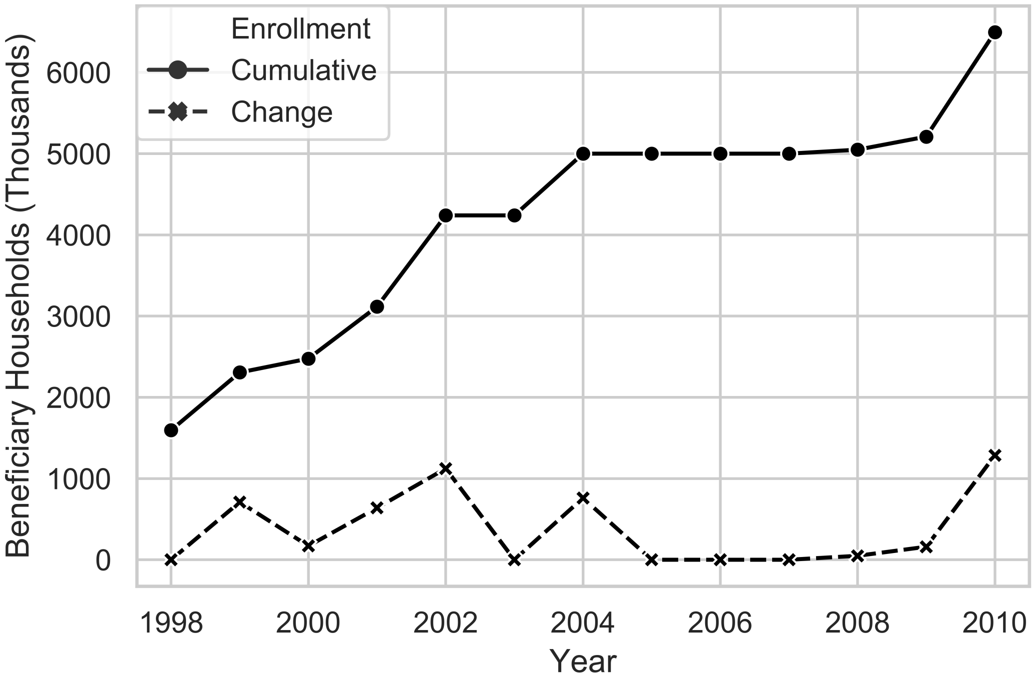

Figure 2-1: Treatment Expansion Over Time

This figure displays the evolution of the cumulative sum (solid line) and yearly changes (dashed line) in the number of beneficiary households in Mexico. The pattern of this roll-out shows three waves: 1997–1999, 2000–2005 and 2010 onward. The pauses are explained by the national elections of 2000, 2003, 2006 and 2009.

Source: *Progresa/Oportunidades* external examiner.

while richer, the municipalities that saw the largest expansion during the second wave were closer to those of the first wave in terms of poverty profile than the municipalities in the third expansion. Thus, municipalities that benefited most from the first and second waves are presumably more alike.

This pattern can also be seen in Figure 2-2. This figure plots the count of enrolled households in 1999 against 2005 (left) and the count of enrolled households in 2005 against 2010 (right). The size and color of the points are representative of the total number of households in a municipality. We can see from the left panel that while a fraction of the new enrollments went to municipalities with more than 150,000 households, the second wave mainly targeted small municipalities. As can be seen from the right panel, this is no longer true with the third wave. Indeed, the expansion principally occurred in larger municipalities that had relatively low levels of enrollment in 2005.

Younger and Older cohorts – Taking a long term perspective on the program suggests the definition of two treated cohorts that are treated differently. Consider the cohort of

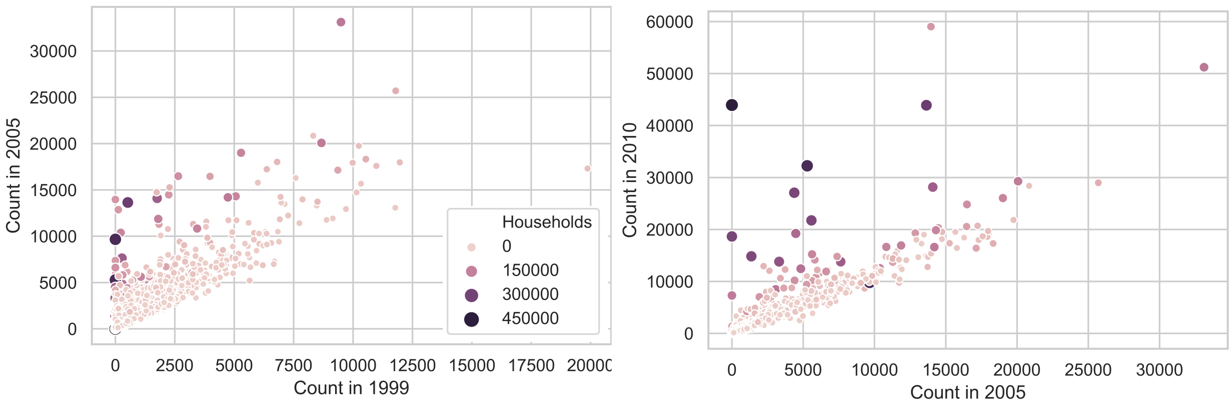

Figure 2-2: Program Expansion Towards More Urban Areas
This figure plots the count of households treated in 1999 against the number of households treated in 2005 (left) and the count of households treated in 2005 against the number of households treated in 2010 (right). The color and size of the points represent the total number of households in that municipality. This figure shows that the first two waves were mainly targeted towards similarly small municipalities while the third expansion disproportionately favored larger municipalities.
Source: INEGI and *Progresa/Oportunidades* external examiner.

women aged 16 or higher when the program started in 1997. Provided that they were in a treated area, they had the opportunity to enroll in the program and receive a cash transfer proportional to the number of children they had at the time, conditional on respecting the school and health requirements for their children. This cohort of women, labeled the *older cohort* throughout this paper, was thus treated on two fronts: they received cash transfers and were required to invest in the human capital of their children through school attendance.[8] I define this *older cohort* to be composed of women aged 16 to 36 in 1997. A rationale for the definition of the upper bound is that women aged 36 when the program started will be 49 in 2010 when I observe their fertility. As can be seen in Figure 2-A.1, 49 is the last age where the observed probability for a woman to have a child is non zero. In other words, 36 is the latest age for which fertility (total number of children ever born) can still meaningfully change in 2010.

Now consider the cohort of women aged 11 or lower in 1997. Unlike the *older cohort*, by design of the program, this group of women were more likely to be enrolled in school because of the program when it started.[9] Yet, after these school years, they faced the same program eligibility and conditionality as the *older cohort*. In other words, this group, labelled the *younger cohort*, first received extra schooling and then had children for which they would potentially receive cash transfers from the government under the condition that they go to

---

[8] I will also consider effects child health and mortality in Table 2-4. For now I abstract from this dimension to simplify exposition.

[9] This effect has been documented multiple times, see for instance: (93)

school, the same way the *older cohort* did. Therefore, to the extent that these three elements have a potential effect on fertility, the *older cohort* can be used to capture the joint effect of the handout and the mandatory investment in their children's education on fertility. The *younger cohort* adds to this list the effect for a woman of being more educated.

I define this *younger cohort* to be composed of women aged 7 to 11 in 1997. In some specifications, I extend the range to include women aged 7 to 15. This is because the eligibility for entering the program as a child was 15 or lower. Still, these women aged 12 to 15 were at best partially treated by the program. A rationale for the definition of the lower bound is that women aged 7 when the program started will be 20 in 2010 when I observe their fertility which, as can be seen in Figure 2-A.2, is still very early in the fertility profile of a Mexican woman. In my empirical analysis, I also report results for a *teenagers* category composed of women aged 2 to 6 in 1997 whose fertility I observe when they are 15 to 19 in 2010. A summary of these age thresholds is displayed in Figure 2-3.

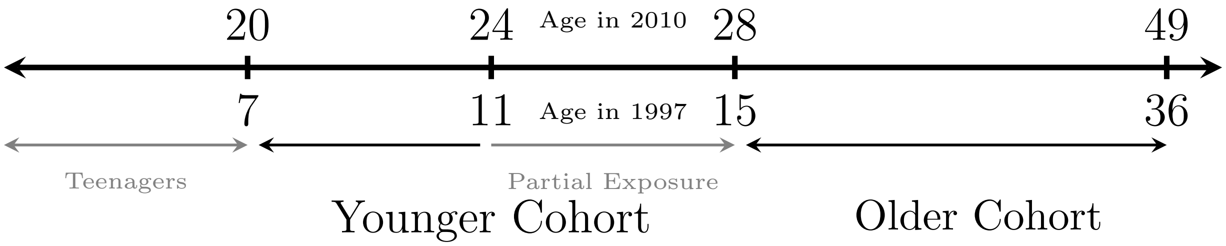

Figure 2-3: Cohorts Definitions

This figure summarizes the definitions of the *younger* and *older cohorts* depending on their age when the program started in 1997 and when they are observed in the IPUMS data in 2010. The *older cohort* is composed of women aged 16 to 36 in 1997. The *younger cohort*, in its narrow definition, is composed of women age 7 to 11. These women correspond to the first cohort who benefited from the additional schooling component of the program for the maximum number of years possible. A wider definition of the *younger cohort* also includes women aged 12 to 15. These women were partially exposed in the sense that they were eligible to receive additional schooling but not for the full length of their education. The teenagers cohort is composed of women 6 or younger in 1997 whose fertility is observed when they are still teenagers in 2010.

2.2.2 Theoretical Motivation

This section details how the three main treatment components (cash transfer, conditionality and schooling for the *younger cohort)* are expected to influence fertility from a theoretical standpoint. The effect of *Progresa* on fertility is not *ex-ante* self-evident. First, although there is a large literature documenting the negative correlation between income and fertility, the program may instead create an incentive for parents to bear more children since the program is an increasing function of the number of children. Second, it is not obvious whether the two different treatments received by the *older* and the *younger cohorts* (defined in the previous section) should result in different fertility changes. If the *younger cohort* reduces fertility further compared to the *older cohort*, this would be in line with *Progresa*'s ambition to alleviate poverty in the short run via cash transfers and to eradicate poverty in the long run thanks to its mandatory schooling component.

In order to frame the research question and capture both the static effect of cash transfer and mandatory schooling and the inter-generational dynamic coming from women's own education, I use an overlapping-generations model developed in (76). I first present the mechanisms by which fertility may be reduced in this *quality-quantity* trade-off model in the absence of treatment. I then show successively i) how fertility may be impacted by the cash transfer alone, ii) the effect of the introduction of conditionality which may reverse the first prediction and iii) the conditions under which the *younger cohort* would further decrease its fertility compared with the *older cohort*.

Benchmark model - no *Progresa*

Consider the following overlapping generations model inspired by Moav (76) in which a representative parent lives two periods. In the first period, she acquires human capital. In the second period, she consumes, decides the number of children to have and decides whether and how much to (uniformly) invest in her children's education. She maximizes the log-linear utility function[10] over her own consumption (c_t), her number of children (n_t) and her children's future income (h_{t+1}).

$$U_t(c_t, n_t, h_{t+1}) = (1-\beta)\log(c_t) + \beta[\log(n_t) + \theta\log(h_{t+1})] \tag{2-B.1}$$

Here $\beta \in (0,1)$ is the weight determining the parent's preference for consumption over the number and income of her children, while $\theta \in (0,1)$ drives the preferences between quantity and quality captured by the child's income, h_{t+1}, in the next period. Wages are normalized to 1 in this economy so that income only depends on human capital h.

---

[10] Many of the results derived in this model do not demand log-linear utility. Yet, as noted by Jones et al. (64), the utility function does need to be separable in consumption, quantity and quality of children for the negative relationship between income and fertility to hold across generations.

Let human capital production be a function of the parent's investment in her children,[11] so that :

$$h_{t+1} = h(e_{t+1}) := \delta_0 + \delta_1 e_{t+1} \tag{2-B.2}$$

Here $\delta_0 > 0$ is the human capital endowment which is independent of investment in quality and $\delta_1 > 0$ is the return to investment on children e_{t+1}.

The representative parent is endowed with a unit of time and a constant fraction τ is needed to rise each child n_t regardless of their quality.[12] The cost of quantity is therefore given by τh_t due to forgone income and increases with the parent's human capital h_t while the cost of quality is given by e_{t+1} and does not depend on h_t. The budget constraint is given by:

$$c_t + n_t(\tau h_t + e_{t+1}) \leq h_t \tag{2-B.3}$$

Combining equations (2-B.1),(2-B.2) and (2-B.3) gives the optimization problem of the representative parent. First, note that the optimal consumption of a parent is a fixed fraction of the parent's income:

$$c_t^* = (1 - \beta)h(e_t) \tag{2-B.4}$$

This depends on the investment made by her parent in her education e_t in the previous period. In other worlds, consumption is increasing in fixed proportion, across generations, provided that education is increasing as well.

The non-linear interaction between n_t and e_{t+1} in the budget constraint generates the usual quality-quantity trade-off first introduced by Becker (14). In this model, the shadow price of the number of children is increasing with the representative parent's level of capital h_t and the investment in the quality of her children e_{t+1}. The marginal rate of substitution

---

[11] This is a common functional form first given by Becker and Tomes (16). Although Moav (76) does not define the human capital functional form, the author requires for the elasticity $\frac{e_{t+1}h'(e_{t+1})}{h(e_{t+1})}$ to be increasing in e_{t+1}. $h_{t+1} := \delta_0 + \delta_1 e_{t+1}$ represents a special case also used by Vogl (107). This functional form allows both for fertility to decline in parental skills in an interior solution, and for a corner solution with no investment in quality if parental skills are too low. Vogl (107) also includes a subsistence level of consumption which implies that the demand for children is first increasing in h_t for low levels of h_t before decreasing in this argument. This stylized fact does not seem to hold in the Mexican context. Indeed, childlessness or more broadly a small number of children for low levels of income do not seem to be important. This will be apparent in the empirical part below. In particular, I find that there is no extensive margin response to *Progresa* which would confirm this stylized fact.

[12] Another assumption is that $\beta > \tau$. This is because a corner solution of no investment in quality e_{t+1} implies a constant and positive number of children given by β/τ.

between quality and quantity for an interior solution is thus given by:

$$\frac{h(e_{t+1})}{\theta n_t} = \frac{\tau h(e_t) + e_{t+1}}{n_t / h'(e_{t+1})} \tag{2-B.5}$$

We can see that the cost of quantity $[\, h(e_t) + e_{t+1}]$ is increasing in e_{t+1} but also in the parent's own education e_t. The dynamic implications are that the price of quantity increases across generations if education is increased over time.

Finally, the optimal number of children as a function of investment decision includes a possible corner solution with no investment and is given by:[13]

$$n_t(e_{t+1}) = \begin{cases} \beta/\tau & \text{, for } e_{t+1} = 0 \\ \beta h(e_t)/(\tau h(e_t) + e_{t+1}) & \text{, for } e_{t+1} > 0 \end{cases} \tag{2-B.6}$$

Fertility decisions with *Progresa*

I now introduce, in turn, two of the main components of *Progresa*: the cash transfer and mandatory schooling. Results derived here are meant to represent *Progresa*'s treatment effect in a context where the circumstances of the baseline model from Section 2.2.2 represent the counterfactual. As noted in Parker and Todd (82) take up of treatment were very high so I do not model decisions to enroll in the program, nor do I model when to exit it.

Unconditional cash transfer – First consider the effect of the introduction of a cash transfer without mandatory schooling. Let γ represent the average per child cash transfer.[14] For simplicity, the model abstracts from the fact that transfers were capped to a maximum of three children. The new budget constraint is now given by:

$$c_t + n_t(\tau h_t + e_{t+1}) \leq h_t + \gamma n_t \tag{2-B.7}$$

Equations (2-B.1),(2-B.2) and (2-B.7) now form the optimization problem. Note first that optimal consumption is unchanged compared to the benchmark model (2-B.4) which means that the relaxation of the budget constraint offered by the cash transfer only influences spending on quantity and/or quality.[15] The optimal fertility decision from (2-B.6) is now

---

[13] The closed form solution for the optimal investment in education is given by: $e^*_{t+1} = \frac{\theta}{1-\theta}(\tau h_t - \frac{\delta_0}{\delta_1})$, which, for an interior solution, yields the following optimal number of children: $n^*_t = \frac{(1-\theta)\beta h_t}{\tau h_t - \delta_0/\delta_1}$.

[14] Transfers amount are increasing with age and larger for female students in practice but I abstract from these dimensions here.

[15] This is a property of the model used. Because γ multiplies n_t it follows that $h_t - c_t = n_t(\tau h_t + e_{t+1} - \gamma)$. If γ does not depend on n_t, the optimal consumption is given by $c^*_t = (1-\beta)(h_t + \gamma)$.

given by:

$$n_t(e_{t+1}) = \begin{cases} \beta h(e_t)/(\tau h(e_t) - \gamma) & \text{, for } e_{t+1} = 0 \\ \beta h(e_t)/(\tau h(e_t) + e_{t+1} - \gamma) & \text{, for } e_{t+1} > 0 \end{cases} \tag{2-B.8}$$

In the absence of investment in the quality of children ($e_{t+1} = 0$), an unconditional cash transfer will increase the number of children and the magnitude of this effect depends on the distance between the parent's income $h(e_t)$ and the transfer amount γ.[16] It is easy to see that the positive impact of the cash transfer on fertility is dampened when the parent does not opt for a corner solution and invests in quality ($e_{t+1} > 0$).

Conditional cash transfer – Now consider the effect of the cash transfer when coupled with mandatory schooling. Define the new investment in quality $\tilde{e}_{t+1}$. I can assume that $\tilde{e}_{t+1} > 0$ because investment in education is mandatory under the program and cannot be zero. $\tilde{e}_{t+1} > e_{t+1}$ implies that the parent does not reduce her own investment compared to a situation without the program. Thus, I can rewrite (2-B.8) dropping the corner solution as:

$$n_t(\tilde{e}_{t+1}) = \frac{\beta h(e_t)}{\tau h(e_t) + \tilde{e}_{t+1} - \gamma} \tag{2-B.9}$$

The relationship between $\tilde{e}_{t+1}$ and e_{t+1} (investment in the absence of the program) depends in part on the response of the parent; e.g. whether mandatory schooling and the parent's own investment are complements instead of substitutes. Equation (2-B.9) suggests that, at the margin, fertility will be reduced for all treated cohorts if the total induced investment in quality, $\tilde{e}_{t+1} - e_{t+1}$, exceeds the benefit from the transfer γ such that:

$$\tilde{e}_{t+1} - \gamma > e_{t+1}$$

Younger vs older cohort – Finally, there is a possibility that, out of the two treated groups, the *younger cohort* will experience an even larger reduction in fertility.[17] To see this let the superscripts $\{Y, O\}$ represent respectively the *younger* and *older cohort* and recall from equation (2-B.5) that the cost of quantity is given by: $[\,h(e_t) + e_{t+1}]$. Therefore, this cost of quantity will increase for the younger cohort if the following inequality is respected:

$$h(\tilde{e}_t^Y) + \tilde{e}_{t+1}^Y > h(e_t^O) + \tilde{e}_{t+1}^O$$

---

[16] The special case $\gamma = 0$ yields the original corner solution in (2-B.6)

[17] Remember that the *younger cohort* first receives extra schooling as a result of exposure to the program and then make fertility decisions otherwise facing the same conditions as their parents.

Which, assuming that the education investment under mandatory schooling is the same for both cohorts ($\tilde{e}^Y_{t+1} = \tilde{e}^O_{t+1}$), yields the following inequality:

$$h(\tilde{e}^Y_t) > h(e^O_t) \iff \delta_1 \tilde{e}^Y_t > \delta_1 e^O_t$$

Given the investment in her education, the model predicts that an individual in the *younger cohort* will further decrease her fertility and that the increase in the cost of quantity is driven by $\delta_1 > 0$. Thus, overall, the program is expected to reduce the fertility of both cohorts as a result of the increase in education investment, provided that this boost to investment is large enough compared to the amount of the transfer. Furthermore, the *younger cohort* is expected to further reduce fertility as a consequence of a higher education level which increases income and the cost of quantity.

2.3 Empirical Strategy and Data

Identifying the potential effect *Progresa* has had on the *younger* and *older cohorts* presents a double challenge: i) the difference of treatment intensities between treatment and control need to be large enough for an effect to actually be detected, ii) both treatment exposure and fertility levels are jointly determined by age. This section motivates the difference-in-difference strategy used in this paper, which relies on cross-census variation to identify pre and post treatment cohorts, and the characteristics of the program's roll-out to determine areas treated and account for pretrends. Section 2.3.1 describes the construction of the intensive margin treatment variable based on the different waves characterizing the expansion of the program. It also describes the method used to account for pre-trends. Section 2.3.2 uses the timing of *Progresa*'s implementation to define the two treated cohorts and a control cohort in the IPUMS data.

2.3.1 Treatment Exposure

I identify the effect of *Progresa* on fertility using a difference-in-difference strategy inspired by Parker and Vogl (83). This strategy relies on two measures of treatment intensity, interacted with a post treatment indicator, varying at the municipality level. The first measure represents the share of households enrolled in the program in 1999 and the second represents the analogous share for the year 2005. This yields a specification in which the coefficient on the interaction term for the 1999 enrollment measure captures an effect conditional on enrollment levels in 2005. In other words, conditional on having a high level of enrollment in 2005, the impact of a high level of enrollment in 1999 is given by the coefficient on the first interaction term.[18] This strategy retains variation in the timing and intensity of the treatment by municipality while also accounting for heterogeneity between municipalities. Specifically, the second interaction term aims at capturing the time varying effect of being in a poor municipality, allowing for the first interaction term to measure the effect of earlier high intensity enrollment.[19] The identification relies on the assumption that earlier treatment is not correlated with relevant unobservables.[20]

These two exposure measures are defined by the following ratios:

$$enroll_m^y = \frac{\text{Treated Households}_m^y}{\text{Total Number of Households}_m^y} \text{ , for } y \in 1999, 2005 \tag{2-C.10}$$

Enrollment in municipality m in year y is defined as the total of the cumulative sum of treated households at the end of that year divided by the number of households in the same location in that year.[21] I obtain geostatistical information from the Mexican external evaluation organization which provides yearly information on the cumulative number of beneficiary households by municipality. For the total number of households in a municipality, I use the household count provided by the Mexican national institute of statistics INEGI.[22]

Table 2-1 presents the summary statistics for the municipalities included in my analysis. The average municipality has an enrollment ratio of 0.26 in 1999 and 0.48 in 2005 reflecting the large increase in beneficiaries between the first and the second waves. Notice that while the bottom 75% of municipalities have exposure measures below .70 in 2005, this variable is not bounded above by 1. This is because the number of treated households corresponds to

---

[18] Recall from Section 2.2.1 that municipalities relatively more intensively treated during the second wave of the program's expansion were closer, as defined by marginality indicators, to municipalities treated during the first wave.

[19] See equation (2-C.11) in Section 2.3.2 for more details.

[20] Note that looking at the long term effect of *Progresa* one could instead use the initial randomization of the program. The effect captured would thus be the effect of being exposed to the program for up to 18 months ahead of the control group. While this may be enough to capture differences in market outcomes such as education or income it may not be enough for fertility. To see this, consider a women who is 20 when the program starts in 1997. When her fertility is observed in 2010, she will have been exposed to the program for 13 years while a woman of the same age in the control group will have been exposed for 11.5–12 years. Therefore, by 2010, both treatment and control will have made fertility decisions facing very close conditions thus making the detection of an effect unlikely.

[21] Although this ratio is not bounded above by 1 (because over time, cumulatively, there can be more beneficiary households than the total number of households) allowing the denominator to vary over time helps with the interpretability of the results.

[22] In their paper, Parker and Vogl (83) use the the household count provided by the IPUMS census survey. An issue with this method is that households belonging to small municipalities are more likely to be undercounted as noted in the IPUMS documentation.

both current and formerly enrolled households.[23]

	1999 exposure ratio	2005 exposure ratio
mean	0.26	0.48
std	0.24	0.28
min	0.00	0.00
25%	0.02	0.23
50%	0.21	0.49
75%	0.46	0.70
max	1.05	1.26

Table 2-1: Exposure Measures

This table displays summary statistics of the 1999 and 2005 enrollment share measures defined in equation (2-C.10) across Mexican municipalities. Between 1999 and 2005, the average Mexican municipality ratio went from 0.26 to 0.48. Measures above 1 are indicative of a cumulative number of enrolled households larger than the total number of households in that municipality.
Source: *Progresa*/*Oportunidades* external examiner and INEGI.

2.3.2 Cross-Census Variation

I obtain information on fertility and other socio-economic features such as schooling and income, using the Integrated Public Use Micro Sample (IPUMS) census surveys for Mexico for the years 2010, 2000 and 1990.[24] The main dependent variable used throughout this study is the number of children ever born. This corresponds to the total number of children a woman ever gave birth to. Because it includes both surviving and dead children the effects found using this variable should be interpreted as the change in the total number of births.[25]

The use of this dependent variable implies that some of the fertility decision can potentially be made prior to the beginning of the treatment and are therefore unrelated to it. Depending on when the treatment occurs with respect to one's fertility timing, a woman can either stop having additional children or plan to have fewer children compared to what she would have had in the absence of the program. It is only by the end of a woman's fertility window that one can observe the total chosen number of children. In this study, complete fertility will be

---

[23] This is the case for 3 municipalities out of 2,392. The definition of a household in *Progresa* may also be looser than the statistical definition used by INEGI. Therefore it is possible for *Progresa* to report more households treated than the total number of statistical households even early on. To the extent this would represent measurement error I replicated the main regression having top-coded these ratios. Results remain quantitatively similar.

[24] I use the 10% sample from 1990, the 10.6% sample from 2000 and the 10% sample from 2010.

[25] The IPUMS censuses also include those two variables although the information on the number of dead children is absent from the 1990 census.

observed only for some women belonging to the *older cohort*. It is unlikely to be observed for the *younger cohort* since the oldest are 28 in 2010. Therefore, an important assumption of this work is that reductions in fertility can be at least partially observed before a woman's fertility window closes. In other words, the decision to reduce fertility due to the treatment can be observed only if, on average within a cohort, there is reductions of the number of children throughout the fertility window.

In order to measure the total effect of the treatment, I follow the methodology introduced by Aaronson et al. (1) and use cross-census variation between the 2010 and the 1990 census surveys instead of within 2010 census variation across age groups. For the *older cohort*, I include women, from the 2010 and 1990 IPUMS, aged 29 to 49 at the time of the census. For the *younger cohort*, I include women, from the 2010 and 1990 IPUMS, aged 20 to 24 (or 28) at the time of the census. Thus, in each case, women appearing in the 2010 census belongs to the post-treatment category while women appearing in the 1990 census belongs to the pre-treatment category.[26]

I merge this individual level information with information on the treatment exposure intensity, detailed in Section 2.3.1. This measure, which enters the interaction term of interest, varies at the municipality level. When necessary these municipalities are merged into a single, time-invariant, municipality.

Thus, the main identification strategy takes advantage of the repeated cross Section offered by the different decennial censuses and applies the difference-in-difference strategy inspired by Parker and Vogl (83). Combined, these two components yield the following empirical specification:

$$y_{imt} = \beta(enroll_m^{1999} \times post_t) + \gamma(enroll_m^{2005} \times post_t) + \Psi(X_m^{1990} \times post_t) + \delta_m + \eta_t + \epsilon_{imt} \tag{2-C.11}$$

This model relates the total number of children y_{imt} for woman i, in municipality m, aged t, to a municipality fixed effect δ_m, a year of census fixed effect η_t and the interaction of the municipal enrollment share in 1999 and an indicator equals to 1 if the observation belongs to the 2010 census, as well as the interaction between the post indicator and 2005 enrollment. I also include an interaction term between the post indicator and the marginality index described in Section 2.2.1. This strategy retains the timing and intensity variations

---

[26] I will also report the results from an exercise in which I use the 2000 census as pre-treatment. The advantage of using the 2000 census is that it makes post and pre-treatments more comparable and the requirement on pretends less likely to be violated. Nevertheless, this does not constitute the main specification because the program started in 1997 which means that the 2000 census does not truly represent a pre-treatment group but a short term treatment instead.

in the treatment captured by β while accounting for heterogeneity between municipalities in the data set. Thus, γ and Ψ, jointly capture the time varying effect of being in a poor municipality. This allows β to measure the effect of earlier high intensity enrollment without this cofounder.[27]

2.4 Results

This section presents the results of estimating the empirical model described in Section 2.3. Section 2.4.1 displays the main results of the analysis and shows the extent to which fertility declined in areas treated by *Progresa*. Section 2.4.2, tries to isolate the effect of schooling for the *younger cohort* by making comparisons to the *older cohort*. Finally, Section 2.4.3 attempts to shed light on the mechanisms by which fertility was reduced.

2.4.1 Main Results

Table 2-2 reports the results from estimating equation (2-C.11), for the different cohorts, using women of the same age in the 1990 census as comparison group. I also report the average fertility for each cohort-census combination – the difference between which is captured by the year dummy. Throughout, I report the estimate for the 2005 enrollment share interacted with the post indicator. Recall that, jointly with the marginality index, this interaction term should capture the time varying effect of belonging to a poor locality.

Column (1) presents the results for a teenager cohort composed of women aged 2 to 6 in 1997 when the program started and whose fertility I observe when they were 15 to 19 in the 2010 census. The coefficient on the interaction term between the 1999 enrollment share and the post indicator is negative but small and insignificant. This is unsurprising given the average fertility for Mexico in this age category, displayed at the top of the table for each cohort, which remained stable at low levels between 1990 and 2010. Columns (2)–(4) present the results for different definitions of the *younger cohort*. The negative effect of the program on fertility is not significant in column (2) for the fully treated cohort (aged 7 to 11 in 1997). When considering the partially treated, either in isolation in column (3) or in combination with the fully treated cohort in column (4) this negative effect becomes significant. Thus, for the extended definition of the *younger cohort* presented in column (4) a 1 percentage point increase in the enrollment share in 1999 is associated with a fertility decline of 0.002 fewer children. For the average Mexican municipality this corresponds to a decline in fertility by

---

[27] In appendix 2-A.1, I present elements defending both the relevance of the research question and the validity of the cross-census specification.

	Teenagers	*Younger Cohort*			*Older Cohort*	*All Cohorts*	*Placebo Cohorts*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Age in 1997	2-6	7-11	11-15	7-15	16-36	7-36	40-47
	-	-	-	-	-	-	-
Age in 2010	15-19	20-24	20-28	15-28	29-49	20-49	53-60
Fertility							
2010 Mean of Cohort	0.17	0.87	1.65	1.19	3.14	2.44	5.13
2010 SD of Cohort	(0.45)	(1.06)	(1.42)	(1.28)	(2.22)	(2.15)	(3.22)
1990 Mean of Cohort	0.17	1.03	2.09	1.46	4.25	3.08	6.45
1990 SD of Cohort	(0.53)	(1.28)	(1.73)	(1.57)	(2.95)	(2.82)	(4.02)
Enrollment 1999	-0.0124	-0.0904	-0.283***	-0.198***	-0.377***	-0.308***	-0.291
×*Post*	(0.0215)	(0.059)	(0.0914)	(0.0695)	(0.128)	(0.0911)	(0.210)
Enrollment 2005	-0.112***	-0.439***	-0.436***	-0.331***	-0.453***	-0.455***	0.937***
×*Post*	(0.0242)	(0.0668)	(0.0787)	(0.0542)	(0.141)	(0.102)	(0.234)
Marginality	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Municipality	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	961,253	834,974	1,422,252	2,383,505	2,294,703	3,716,955	484,997
R^2	0.017	0.055	0.063	0.032	0.140	0.071	0.123

All regressions are estimated using ordinary least squares. Standard errors clustered at the municipality level are in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2-2: Effects on Fertility

This table displays estimated effects of being exposed to *Progresa* on fertility for different cohorts. The dependent variable is the total number of children ever born and varies at the individual-census level. The effect of the treatment is captured by the coefficient on the interaction of a 2010 census dummy with the 'Enrollment 1999' share, which varies at the municipality level and corresponds to the cumulative sum of households enrolled in the program divided by the total number of households in that municipality. The variable 'Enrollment 2005' is the equivalent for the year 2005. The variable 'Marginality' is the index that determined the location of the initial treatment. Combined, these two variables aim to capture the time varying effect of belonging to a poor municipality. 'Municipality' represents fixed effects for the 2,392 municipalities. The 'Year' fixed effect corresponds to a dummy equal to 1 if an observation belongs to the 2010 census and 0 if it belongs to the 1990 census.

0.052 children per woman.[28] This number seems sensible compared with the average 0.33 decrease in children in Mexico for this age group over the same period.

Column (5) reports the change in fertility experienced by the *older cohort*, whose fertility I observe when they were aged 29 to 49 in the 2010 census. The coefficient on the 1999 enrollment share suggests a reduction in fertility by 0.004 fewer children for each 1 percentage point increase in the 1999 enrollment ratio. This corresponds to a decline of 0.099 children per woman, for the average Mexican municipality, compared with a decline in fertility by 1.21 children that this cohort experienced on average between 1990 and 2010. Column (6) pools the *younger* and the *older cohorts* together and confirms the individual results: the program did not act as an incentive to have more children and instead negatively affected

---

[28] As can be seein in Table 2-1, the mean of the enrollment measure in 1999 is 0.26.

women's fertility, a result compatible with the quality-quantity trade-off model developed in Section 2.2.2.

Finally, column (7) provides the results of a placebo regression including a cohort of older women. These are women aged 40 to 47 when the program started and whose total fertility I observe when they were 53 to 60 in 2010 the census. Because of their age at the time, these women are expected to have been at the end of their fertility window when the program started. Reassuringly, the effect for this cohort of woman is no longer significant.[29] In Table 2-A.3, I reproduce the same regressions using the 2000 census as pre-treatment.[30] The main conclusions remain unchanged.

2.4.2 Additional Effect of Schooling

Given that the program is estimated to have impacted women's fertility negatively, including for the *younger cohort*, I now ask whether we can detect an extra effect of schooling on fertility. Recall from Section 2.2.2, that the *younger cohort* is expected to further decrease its fertility in comparison with the *older cohort* as a result of their increased education level which translates into a greater cost of quantity. The first cohort who received a full dose of extra education because of the program were women aged 7 to 11 in 1997. The youngest cohort that did not receive the education treatment in 1997 were aged 16 to 20 in 1997.

Because these cohorts are close in age, I first display results of estimating a model using within census comparison under the assumption that the key issue with such a model – the fact that age jointly determines treatment exposure and fertility – is mitigated by the age proximity of the two cohorts.[31] I also present results from a triple-difference strategy, similar to the main model presented in equation (2-C.11, in which the comparison between the younger and older cohorts constitutes the third difference. Both methods attempt to isolate the effect of having received the education treatment.

Table 2-3 presents the results from both strategies. All regressions include controls for the interaction measure of the 2005 enrollment share and marginality index with a post dummy. These are once more included in order to control for the time varying effect of belonging to a poor municipality. Panel A displays the results using the narrow definition of the *younger cohort* limited to the fully treated category of women aged 7 to 11 in 1997. They are compared

---

[29] Note that from a biological standpoint these women could still modify marginally their fertility. The issue with including older women as placebo would then be age expectancy related. To the extent that the poorest women are both more likely to bear more children and to die younger this would constitute a wrong placebo group.

[30] This, of course, is not a true pre-treatment since *Progresa* started in 1997.

[31] See Appendix 2-A.1 for a full discussion of this issue.

	(1) Enrollment 1999 $\times$ Educated (Within 2010)	(2) Enrollment 1999 $\times Post \times Educated$ (2010 vs 1990)
Panel A: [7-11] vs [16-20] in 1997		
Interaction 1999	-0.0730	0.0820
	(0.0625)	(0.0617)
Interaction 2005	-0.229***	0.203***
	(0.0579)	(0.0443)
Double Interactions	No	Yes
Age	Yes	No
Marginality	Yes	Yes
Municipality	Yes	Yes
Observations	915,661	1,521,826
R^2	0.278	0.277
Panel B: [7-15] vs [16-24] in 1997		
Interaction 1999	-0.151***	-0.0147
	(0.0577)	(0.0530)
Interaction 2005	-0.229***	0.195***
	(0.0577)	(0.0394)
Double Interactions	No	Yes
Age	Yes	No
Marginality	Yes	Yes
Municipality	Yes	Yes
Observations	1,576,745	2,576,501
R^2	0.292	0.259

All regressions are estimated using ordinary least squares. Standard errors clustered at the municipality level are in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2-3: Extra Effect of Schooling

This table displays estimated effects of being exposed to the schooling component of *Progresa* on fertility. The dependent variable is the total number of children ever born. Panel A presents the results where the educated cohort was 7 to 11 while the control group was 16 to 20 in 1997. Panel B presents the results where the educated cohort was 7 to 15 while the control group was 16 to 24 in 1997. Results in column (1) use treatment and control groups from the 2010 census only. Results in column (2) use the 2010 and 1990 censuses. The effect of interest is captured by the variable 'Interaction 1999' which corresponds to the interaction of the 1999 enrollment share with an educated dummy variable equal to 1 if an individual belongs to the educated cohort in column (1). In column (2), 'Interaction 1999' captures the triple interaction of the 1999 Enrollment, with the educated dummy and a 'Post' dummy equal to 1 if an individual is observed in the 2010 census. The variable 'Interaction 2005' follows the same definition using the 2005 Enrollment ratio. The variable 'Marginality' is the index that determined the location of the initial treatment. Combined, these two variables aim to capture the time varying effect of belonging to a poor municipality. 'Municipality' represents fixed effects for the 2,392 municipalities. The 'Year' fixed effect corresponds to a dummy equal to 1 if an observation belongs to the 2010 census. The variable 'Double Interactions' corresponds to the set of double-interactions of the control variables and fixed effects for the triple-difference estimation. .

to women from the *older cohort* aged 16 to 20 in 1997. These are the youngest women from the *older cohort* not to have been treated by the schooling component of the program and thus the closest comparison group available. Panel B displays results extending the definition of the *younger cohort* to partially treated women and includes women aged 7 to 15 in 1997. The comparison group is also extended to match the same number of age cohorts and includes women from the *older cohort* aged 16 to 24 in 1997.

Column (1) presents the results for the within census strategy. Using the narrow definition of the *younger cohort* composed of women aged 20 to 24 in 2010 yields a small negative and insignificant estimated additional effect of schooling on fertility, which appears in Panel A. The extended definition of the *younger cohort*, composed of women aged 20 to 28 in 2010, appearing in Panel B, suggests a significant negative effect. This effect is not robust to the alternative specifications using the triple-difference strategy presented in column (2). Overall these results suggest no detectable additional effect on the fertility of the younger cohort. This could be explained by a return on investment of schooling[32] that is too low. Alternatively, I note that when observed in 2010 between 20 and 28, most of the women belonging to *younger cohort* are at a relatively early stage of their fertility window. Thus, these results do not rule out an effect of schooling on fertility manifesting itself in later years.

2.4.3 Mechanisms

In this section, using the same cross-census specification as in equation (2-C.11), I test whether several of the assumptions and my interpretation of the results made so far are in line with what can be found in the data. Columns (1)–(4) of Table 2-4 present results for the narrow definition of the *younger cohort* (women aged 7 to 11 in 1997). Columns (5)–(8) present the results for the same dependent variables using the *older cohort*. All results displayed use the cross-census variation between 2010 and 1990 except for columns (4) and (8). Because the count of dead children is not available for in 1990, I instead use the 2000 census as pre-treatment.

Column (1) presents the effect of the program on the number of years of education for the *younger cohort*. As documented in the literature and in Appendix 2-A.1, *Progresa* positively impacted the education of the young cohort. As expected, I find no such effect of the program on the education of the *older cohort* as can be seen in column (6). These results support two key assumptions of this paper. First, this suggests that the *younger* and *older cohorts* were indeed treated differently by the program thus confirming that they represent two distinct cohorts. Second, columns (1) and (5) are suggestive evidence that the reduced form results obtained in this paper are consistent with the theory in Section 2.2.2. Specifically, that fertility may have been reduced in response to changes in investment in the education of children, which reduced the optimal number of children.

Columns (2), (3), (6) and (7) decompose the effect on fertility along the intensive and

---

[32] δ_1 in Section 2.2.2.

	younger cohort				*older cohort*			
	(1) Schooling	(2) Extensive	(3) Intensive	(4) Children Dead	(5) Schooling	(6) Extensive	(7) Intensive	(8) Children Dead
Enrollment 1999	0.480*	-0.0002	-0.245**	-0.0194	0.0808	-0.0103	-0.336***	-0.0394
	(0.209)	(0.0172)	(0.0633)	(0.0108)	(0.191)	(0.00760)	(0.127)	(0.0271)
Enrollment 2005	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Marginality	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Municipality	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	865,015	878,085	439,467	489,470	2,317,792	2,337,470	2,076,353	2,402,181
R^2	0.233	0.034	0.067	0.017	0.282	0.012	0.175	0.059

All regressions are estimated using ordinary least squares.
Standard errors clustered at the municipality level are in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 2-4: Potential Mechanisms

This table displays estimated effects of being exposed to *Progresa* on range of outcome variables for difference cohorts. These dependent variables are 'Schooling' (number of years of education), 'Extensive' (dummy variable equal if a woman has any child), 'Intensive' (number of children conditional on having at least one child), 'Children dead' (total number of children dead). The effect of the treatment is captured by the 'Enrollment 1999' ratio which varies at the municipality level and corresponds to the cumulative sum of households enrolled in the program divided by the total number of households in that municipality. This variable is interacted with the 'Year' dummy variable which equals 1 if an observation belongs to the 2010 census. The variable 'Enrollment 2005' is the equivalent for the year 2005. The variable 'Marginality' corresponds to the index that determined the location of the initial treatment. Combined, these two variables aim to capture the time varying effect of belonging to a poor municipality. The variable 'Municipality' is a dummy variable for the time invariant 2,392 municipalities.

extensive margins. Column (2) reports a precisely estimated zero on the extensive margin for the *younger cohort*. The effect on the extensive margin of the older cohort in column (6) is also estimated to be zero. These results, in combination with the significant intensive margin results in columns (3) and (7), suggest that the effect on fertility was entirely driven by the intensive margin. This may be surprising for the *younger cohort* as one could have expected that spending more time at school might translate into a delay in the arrival of the first child. But these results are consistent with Figure 2-A.3 which shows that by age 20 close to 50% of Mexican woman have had at least one child.

Finally columns (4) and (8) show the effect of the program on child mortality. The inclusion of these regressions is motivated by the fact that, if a woman does not maximize the total number of children ever born but the total number of children alive instead, then total fertility could also decline as a result of improvements in the health of children enrolled.[33] The estimates in columns (4) and (8) suggest a small negative and insignificant effect of the

---

[33] As discussed in Section 2.2.1, *Progresa*'s payments were also conditional on receiving some health care.

program on the number of deceased children. However, because the results use variation between the 2010 and 2000 censes, these point estimates may be upward biased, since the program had existed for 3 years by 2000.

2.5 Conclusion

In this paper I have studied the impact that one of the world's oldest conditional cash transfer programs has had on fertility in the long run, in Mexico. I defined two cohorts of women that were treated differently by Mexico's *Progresa* program based on their age, and motivated theoretically why such a program may reduce each group's fertility via its impact on education investment, income and cost of quantity. The results of this paper suggest that, far from creating an incentive for women to have more children in order to receive more transfers, *Progresa* reduced the total number of children to a small but significant extent. Perhaps surprisingly, the increase in human capital encouraged by the program does not seem to have had any detectable differential impact on the fertility of the *younger cohort*. The results also offer suggestive evidence that the effect on fertility corresponds less to a change in the timing of births (first births don't seem to be postponed), but rather to women tending to have fewer children as they progress towards the end of their fertility window. If true, then I may observe the fertility of the *younger cohort* too early in their fertility cycle to capture the true additional effect of schooling.

2-A Appendix

2-A.1 Support for the Identification Strategy

This Section aims to present various arguments in support of the framing of the research question and the specification used. Specifically, I provide suggestive evidence that the addition of the 2005 enrollment measure and marginality index suggested by Parker and Vogl (83) is convincing in accounting for pre-trends. I also provide evidence that schooling indeed increased in targeted areas. Finally, I point out the challenges with using within census variation when looking at fertility.

In Section 2.2.2, I motivated the impact of *Progresa* on fertility by the changes it imposed on education investment by the *older cohort* for the *younger cohort*. In order to provide evidence that this investment in education is indeed occurring, I keep the methodology from Parker and Vogl (83) in equation (2-C.11) but instead use within census comparisons using the data from the 2010 census:

$$\begin{aligned} y_{imj} = \alpha_j + \delta_m + \sum_{j=13}^{59}(enroll_m^{1999} \times d_j)\beta_j \\ + \sum_{j=13}^{59}(enroll_m^{2005} \times d_j)\gamma_j + \sum_{j=13}^{59}(X_m^{1990} \times d_j)\Psi_j + \epsilon_{imj} \end{aligned} \tag{2-A.1}$$

In this model, each coefficient β_j multiplies the interaction of the enrollment measure in 1999 with a specific age dummy d_j where the omitted category is 60 in 2010 (47 in 2000). These age dummies are also interacted with the 2005 enrollment ratio and the marginality index to account, once more, for the time varying effect of being in a poor municipality.

The orange line in Figure 2-A.4 plots the estimates $\hat{\beta}_j$ and their associated 95% confidence intervals from the estimation of equation (2-A.1) in which the dependent variable is number of years of education. As expected if the identification strategy is successful, these estimates are close to zero and insignificant for women aged 29 to 59 in the 2010 census. Recall that women aged 16 or higher in 1997 were not eligible for the program and should therefore not have benefited from the program in terms of their own education. It can also be seen in Figure 2-A.4 that the effect on schooling starts to go up for women between the age of 11 and 15 in 1997 and becomes positive and significant for women age 10 or younger in 1997.[34] These are the expected results.[35] The *younger cohort* seem to have indeed benefited from the program in terms of schooling while the *older cohort* did not. Also, the stability of the estimates before the 16 year old threshold is encouraging vis-à-vis the ability of the identification strategy to account for pre-trends.[36]

Figure 2-A.4 also plots the $\hat{\beta}_j$ and their associated 95% confidence interval for the estimation of equation (2-A.1) when the dependent variable is instead the total number of children. This corresponds to the purple line. The results seem to confirm two key elements of the model presented in equation (2-C.11). First, they seem to indicate that the age 36 in 1997 is a relevant threshold for the definition of the *older cohort*. Indeed, past this age we can see that there is no significant difference with the omitted category which is in line with the fact that, by that age, the number of children is far less likely to change.

Second, the shape of the estimates as decreases indicates the issue with using a within census strategy for fertility, in which the pre-treatment group could for instance be defined by women aged 46 or higher in 1997. To see why, consider the simple case where exposure

---

[34] Recall from Section 2.2.1 that the 11 to 15 in 1997 cohort corresponds to a partially treated cohort in terms of schooling.

[35] These results are also quantitatively similar to those found in Parker and Vogl (83).

[36] Results from estimation of equation (2-A.1) without the inclusion of the 2005 enrollment ratio display a pattern in which schooling is also increasing for the *older cohort*.

to the treatment is defined by a dummy variable $T = \{0, 1\}$ and let $J = o$ represent the omitted category composed of older women (in our example women aged 60 in 2010). Then by definition of the difference-in-difference estimator, β_j captures the following differences in conditional expectations:

$$\hat{\beta}_j = E\left[Y|T=1, J=j\,\right] - E\left[Y|T=0, J=j\,\right] - \left(E\left[Y|T=1, J=o\right] - E\left[Y|T=0, J=o\right]\right)$$

Note that the fertility expectations conditional on treatment are the same as the unconditional expectation and converge towards zero when age tends to zero, so that:

$$\lim_{j\to 0} E\left[Y|J=j\,\right] = \lim_{j\to 0} E\left[Y|T=1, J=j\,\right] = \lim_{j\to 0} E\left[Y|T=0, J=j\,\right] = 0$$

This means that the within census estimation of a difference-in-difference estimator for fertility will tend towards the single difference represented by the omitted category:

$$\lim_{j\to 0} \hat{\beta}_j = 0 - \left(E\left[Y|T=1, J=o\right] - E\left[Y|T=0, J=o\right]\right)$$

2-A.2 Supplemental Figures

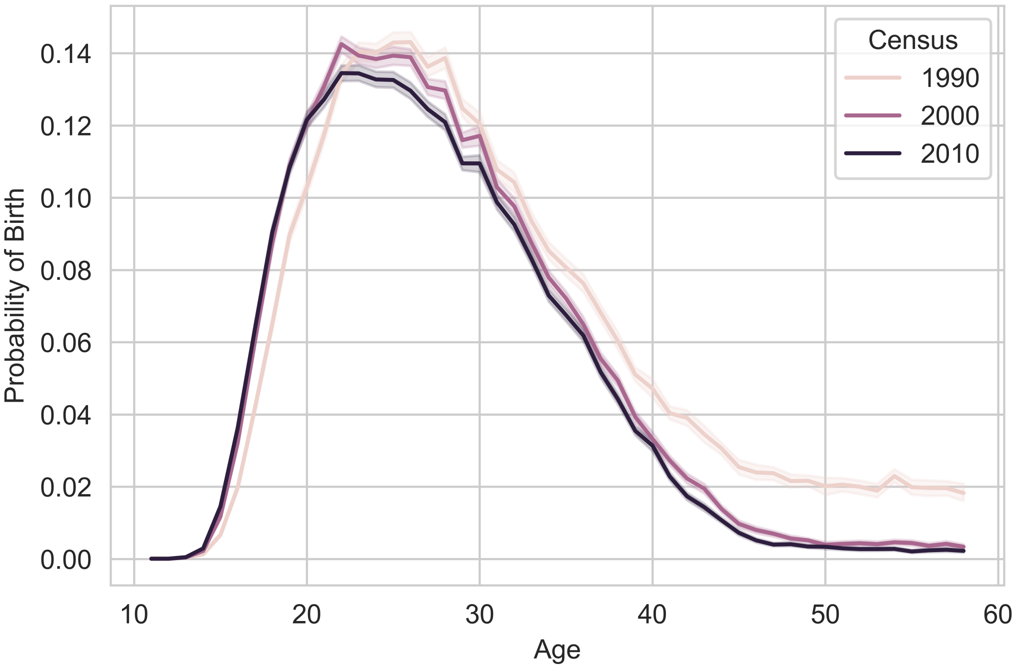

Figure 2-A.1: Probability of Giving Birth by Age in Each Census

This figure displays the probability, by age, for a woman, between 12 and 60 to have a child aged 0. Each line represents the same exercise for a different census.
Source: IPUMS

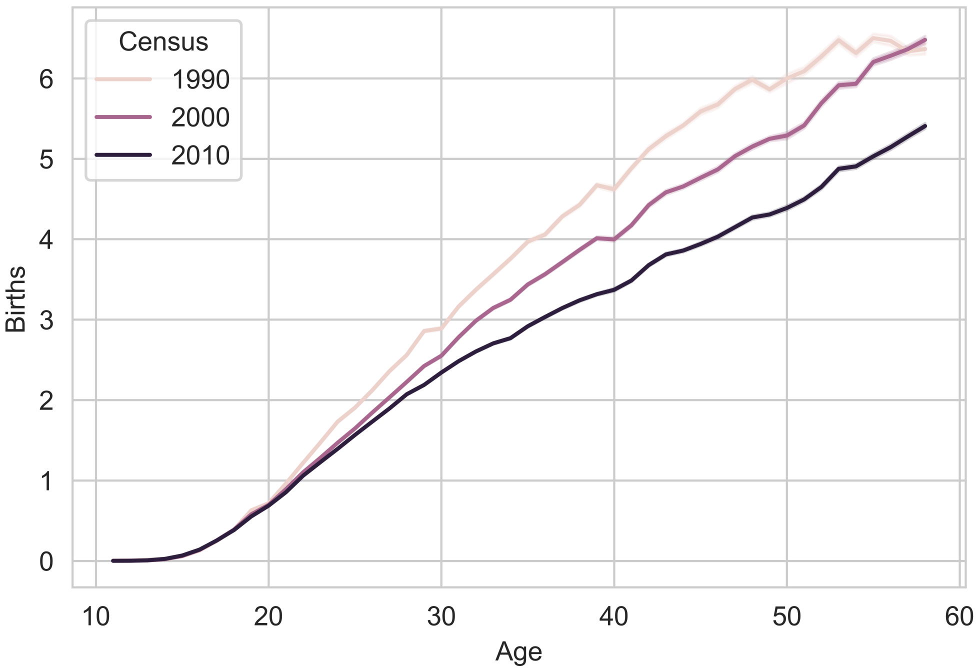

Figure 2-A.2: Cumulative Fertility by Age and Census

This figure displays the cumulative number of children on average by age from 12 to 60. Each line represents the same exercise for a different census.
Source: IPUMS

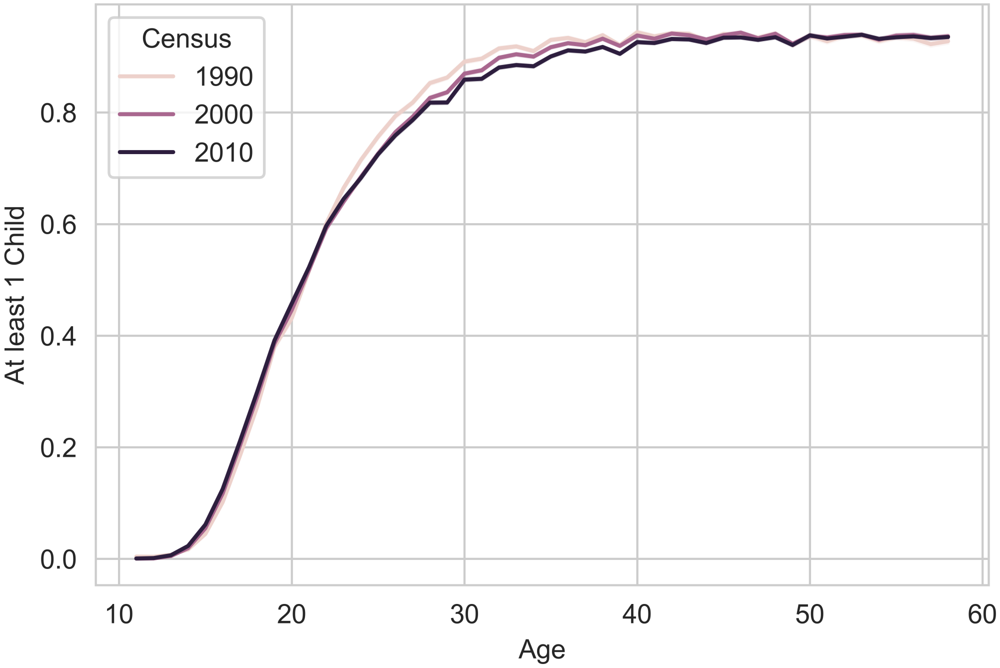

Figure 2-A.3: Probability of First Child

This figure displays the probability of already having at least one child, on average by age, from 12 to 60. Each line represents the same exercise for a different census.
Source: IPUMS

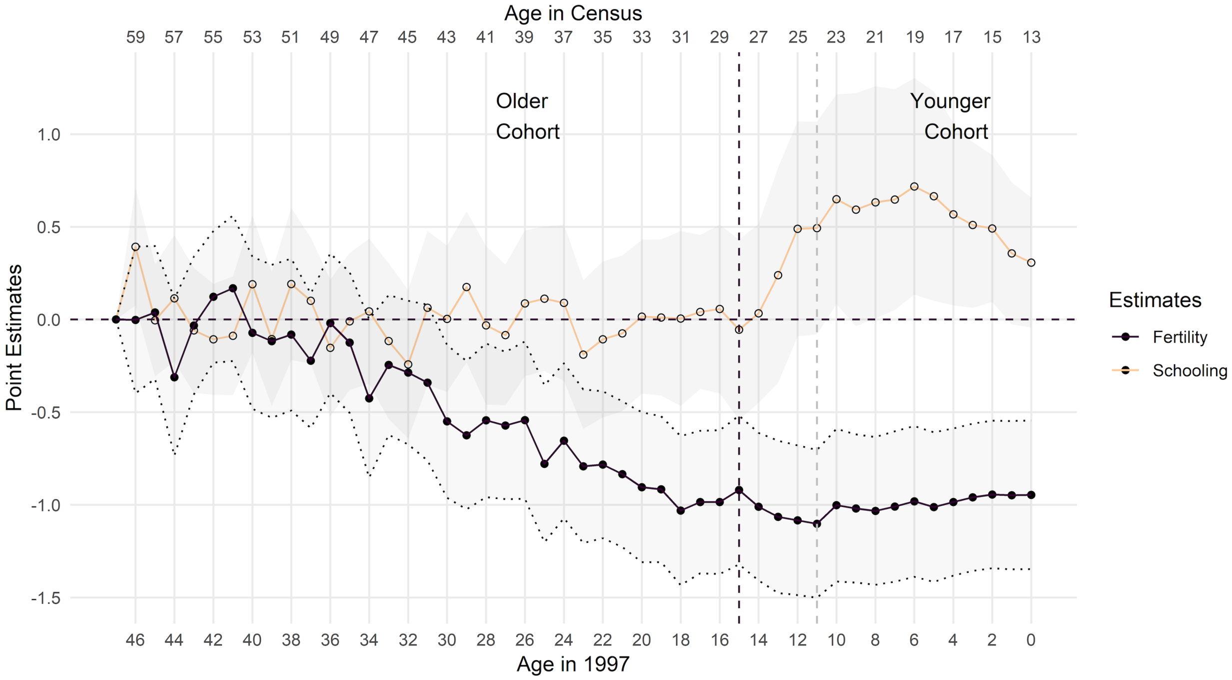

Figure 2-A.4: Within Census Comparisons, Schooling and Fertility

This figure displays the coefficients on interactions between age cohort dummies and the 1999 enrollment intensity, from the estimation of the model detailed in equation (2-A.1), for the two dependent variables: years of education (purple) and fertility (orange). Data includes women aged 13 to 60 in the 2010 census. Women aged 60 constitute the omitted category. The 95% confidence intervals are calculated using of standard errors clustered at the municipality level.

2-A.3 Supplemental Tables

Cash Benefits of Progresa/Oportunidades: Monthly Pesos, 2003 Second Semester		
	Boys	Girls
Primary School		
Grade 3	105	105
Grade 4	120	120
Grade 5	155	155
Grade 6	210	210
Middle School		
Grade 7	305	320
Grade 8	320	355
Grade 9	335	390
High School		
Grade 10	510	585
Grade 11	545	625
Grade 12	580	600
Fixed monthly nutrition grant per household	155	
Maximum household monthly transfer with no children in senior high school	950	
Maximum household monthly transfer with children in senior high school	1,610	

Table 2-A.1: Amount of Transfers
(100 pesos is approximately 6.5 CAD in 2018)

	1997-1999		1999-2005		2005-2010	
	Univariate (1)	Multivariate (2)	Univariate (3)	Multivariate (4)	Univariate (5)	Multivariate (6)
Marginality index	0.183		0.0487		0.0168	
(mean=0, s.d.=1)	(0.00325)		(0.00360)		(0.00203)	
Illiterate	1.111	0.411	0.210	-0.146	0.139	0.317
(mean=0.22, s.d.=0.15)	(0.0283)	(0.119)	(0.0275)	(0.126)	(0.0145)	(0.0623)
No toilet	0.559	0.0589	0.188	0.0251	0.0236	-0.0284
(mean=0.47, s.d.=0.26)	(0.0157)	(0.0233)	(0.0160)	(0.0257)	(0.00815)	(0.0125)
No electricity	0.600	0.119	0.106	-0.0142	0.0704	0.0301
(mean=0.24, s.d.=0.22)	(0.0181)	(0.0254)	(0.0181)	(0.0280)	(0.00929)	(0.0138)
No running water	0.542	0.124	0.0858	-0.0752	0.0300	-0.0155
(mean=0.34, s.d.=0.25)	(0.0167)	(0.0218)	(0.0168)	(0.0235)	(0.00804)	(0.0115)
With dirt floor	0.607	0.129	0.196	0.0740	0.0639	0.0607
(mean=0.41, s.d.=0.27)	(0.0138)	(0.0312)	(0.0143)	(0.0327)	(0.00737)	(0.0155)
Share $\leq 2\times$ min. wage	1.211	0.238	0.606	0.672	0.0676	-0.0733
(mean=0.79, s.d.=0.12)	(0.0319)	(0.0481)	(0.0293)	(0.0513)	(0.0156)	(0.0256)
Share below primary school	1.114	-0.0756	0.184	-0.158	0.135	-0.130
(mean=0.24, s.d.=0.14)	(0.0296)	(0.111)	(0.0289)	(0.117)	(0.0154)	(0.0588)
Household crowding	-0.856	-0.216	-0.255	-0.105	-0.0259	0.103
(mean=0.70, s.d.=0.18)	(0.0234)	(0.0305)	(0.0201)	(0.0313)	(0.0118)	(0.0180)
Municipalities	2392	2392	2392	2392	2392	2392

All regressions are estimated using ordinary least squares. Standard errors clustered at the municipality level are in parentheses

Table 2-A.2: Marginality Index and Roll-Out

This table displays the marginality index and its components and how they correlate with changes in exposure shares over the three waves. Odd numbered columns present the correlation between each variable (including the full marginality index) and the change in enrollment ratios. Even columns present the estimates for each component from a multivariate regression. The 'Marginality index' regressor varies at the municipality level and corresponds to the normalized first principal component of the other variables appearing in this table. These variables are shares of individual above 15 in a municipality with the following characteristics: 'Illiterate' (share of illiterates); 'No toilet' (share without access to a toilet); 'No electricity' (share without electricity); 'No running water' (share without access to running water); 'With dirt floor' (share with a dirt floor in their house); 'Share $\leq 2\times$ min. wage' (share earning less than twice the minimum wage); 'Share below primary school' (share with less than primary school education); 'Household crowding' (share with crowding as measured by number of rooms divided by household size).

	Teenagers	*Younger Cohort*			*Older Cohort*	*All Cohorts*	*Placebo Cohorts*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Age in 1997	2-6	7-11	11-15	7-15	16-36	7-36	40-47
	-	-	-	-	-	-	-
Age in 2010	15-19	20-24	20-28	15-28	29-49	20-49	53-60
Fertility							
2010 Mean of Cohort	0.17	0.87	1.65	1.19	3.14	2.44	5.13
2010 SD of Cohort	(0.45)	(1.06)	(1.42)	(1.28)	(2.22)	(2.15)	(3.22)
2000 Mean of Cohort	0.17	0.93	1.78	1.29	3.69	2.74	6.31
2000 SD of Cohort	(0.48)	(1.14)	(1.52)	(1.38)	(2.63)	(2.51)	(3.74)
Enrollment 1999	0.00985	-0.0773*	-0.126**	-0.0946**	-0.227***	-0.209***	0.115
$\times Post$	(0.0142)	(0.0396)	(0.0581)	(0.0462)	(0.0872)	(0.0695)	(0.148)
Enrollment 2005	-0.0435**	-0.102**	-0.167***	-0.127***	-0.370***	-0.226***	0.409***
$\times Post$	(0.0176)	(0.0446)	(0.0611)	(0.0485)	(0.0957)	(0.0682)	(0.158)
Marginality	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Municipality	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes	Yes	Yes	Yes

All regressions are estimated using ordinary least squares. Standard errors clustered at the municipality level are in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

This table displays estimated effects of being exposed to *Progresa* on fertility for different cohorts. The dependent variable is the total number of children ever born and varies at the individual-census level. The effect of the treatment is captured by the coefficient on the interaction of a 2010 census dummy with the 'Enrollment 1999' share, which varies at the municipality level and corresponds to the cumulative sum of households enrolled in the program divided by the total number of households in that municipality. The variable 'Enrollment 2005' is the equivalent for the year 2005. The variable 'Marginality' is the index that determined the location of the initial treatment. Combined, these two variables aim to capture the time varying effect of belonging to a poor municipality. 'Municipality' represents fixed effects for the 2,392 municipalities. The 'Year' fixed effect corresponds to a dummy equal to 1 if an observation belongs to the 2010 census and 0 if it belongs to the 2000 census.

Table 2-A.3: Effects on Fertility 2010 vs 2000

Chapter 3

When are Tariff Cuts Not Enough? Heterogeneous Effects of Trade Preferences for the Least Developed Countries

By Fabien Forge, Jason Garred and Kyae Lim Kwon

Abstract

We study a large set of recent reforms to trade preferences for the least developed countries (LDCs), examining their impact on export diversification. In the late 1990s and early 2000s, OECD countries greatly expanded their LDC trade preference programs, cutting tariffs on imports from LDCs for a wide range of products. We show that these tariff reductions had a positive average effect, but also that the breadth of their impact was limited. New export relationships (by importer-exporter-product) were stimulated only in cases with previous ‘export experience’, where countries already exported the same product to another OECD country or exported a similar product to the same importer. Tariff cuts for LDCs thus led to increases in ‘sequential exporting’, as well as making existing trade flows more likely to survive, but otherwise made little progress in reducing the overwhelming prevalence of zero trade flows in LDCs’ export portfolios.

3.1 Introduction

For decades, international development policy has promoted the idea of stimulating exports from developing countries, with the ultimate goal of sparking export-led growth. Among the most commonly used instruments in attempting to achieve this goal have been trade preferences, which reduce tariffs and other import barriers for goods produced in less developed economies.

A recent focus of such initiatives has been the least developed countries (LDCs), whose historical export performance has been especially poor. In the mid-1990s, the conclusion of the Uruguay Round of trade negotiations included a call for improved trade preferences targeted specifically at LDCs. This was followed by the WTO Plan of Action for the LDCs in 1996, which again advocated for the improvement of market access for LDCs. This push for policy reforms was successful, as the late 1990s and early 2000s saw many wealthy countries initiating or greatly expanding schemes granting LDCs preferential market access. A key component of these reforms was the reduction of tariffs on LDC exports to zero across a large range of products.

Yet by the 2010s, many policymakers were suggesting that trade preferences had not been successful in inducing export-led growth in LDCs (see e.g. UNCTAD 2015). Indeed, in a 2013 document notifying the WTO of a revised approach to its system of preferences, the European Union referred to "the disappointing performance of the poorest, both in terms of total export growth and of diversification".[1] Three years earlier, the chairperson of the WTO's subcommittee on LDCs drew a similarly pessimistic picture of LDC export performance, stating that "[w]hile there had been some growth in the LDCs, it was not adequate to make a structural change in LDC economies. The vast majority of LDCs continued to be dependent on a limited number of export products with little value addition. Diversification of their production and export base had not taken place."[2] This gloomy assessment suggests that the wide-ranging tariff cuts of the late 1990s and early 2000s may have failed to boost export diversification in LDCs.

In this study, we undertake a quantitative evaluation of this hypothesis. To do so, we use variation by importer, exporter, product and year in a quadruple-difference framework suggested by a model based on (59). The model makes a simple and intuitive prediction: for a particular importer, exporter and product, new trade flows are likely to emerge in response

---

[1] See WTO document WT/COMTD/N/4/Add.6.

[2] See WTO document WT/COMTD/LDC/M/56.

to tariff cuts only in cases that were already on the margin of exporting. However, these will generally be countries already exporting the same product to another destination, or similar products to the same destination; i.e. situations in which countries already have relevant 'export experience'.

We test this prediction by considering the impact of the product-specific tariff cuts for LDCs from thirteen trade policy reforms by nine OECD members, using annual product-level trade data for 1996 to 2011. We find a positive average effect of these reforms: they significantly boost the probability that a treated product is exported to the importer supplying trade preferences. But we show that this effect is entirely driven by importer-exporter-product categories with pre-existing export experience as defined above, as well as an increased probability of survival of existing trade flows.

Our findings therefore suggest that trade preferences for LDCs have resulted in export diversification, in the sense that new trade flows have emerged as a consequence of the reforms we study. But this diversification is of a limited scope, since these new flows consist almost entirely of cases where LDCs had already established closely related trade relationships. So even these wide-ranging cuts to import tariffs were, predictably, not enough to generate a set of entirely new export flows.

This study contributes to a surprisingly small literature on the impact of trade preferences for developing countries, recently reviewed in (80). (49) employ a triple-difference approach to assess the United States' African Growth and Opportunity Act (AGOA), whereby the difference in export growth of treated and control products before and after the policy change for treated countries is compared to that of other countries. They find that on average, treatment by AGOA increased African exports to the US by 13%.[3] (54) examine a wide set of preference programs for developing countries, including various countries' Generalized System of Preferences (GSP) programs, finding generally positive impacts. (57) finds that the temporary expiration of the United States' GSP program in 2011 had a significant adverse effect on developing countries' exports to the US.

We uncover systematic heterogeneity in the effects of trade preference regimes, in line with our hypothesis that their benefits flow disproportionately to exporters with particularly favourable initial conditions. A few previous studies have suggested the possibility of such heterogeneous effects. Soon after the implementation of the EU's Everything But Arms program (EBA), (19) suggested that it was likely to have limited effects on a group of countries

---

[3] (99) replicate Frazer and Van Biesebroeck (2010) for the EU's Everything but Arms program (EBA) for the least developed countries, and also find positive and statistically significant effects. Other studies of AGOA and EBA include (30), (86), (55) and (84).

whose exports were concentrated in products whose tariffs were not reduced by the program. Ten years later, (34) found that benefits from AGOA's liberalization of apparel trade had flowed disproportionately to just seven African countries.[4] Our study differs from the above papers by jointly considering the impacts of the full set of reforms to tariff preferences for LDCs implemented across the OECD in the late 1990s and early 2000s, using an empirical framework suggested by theory. In this sense, it is most similar to (81), who consider the impacts of all trade preference programs for developing countries from 1950 to 2009, using country-pair-level data, and find that their effects vary according to the poverty level and WTO membership status of recipients.

Our paper also contributes to the literature on 'sequential exporting' (Albornoz et al. 3, Schmeiser 92, Chaney 27, Araujo et al. 6, Morales et al. 77). Existing work demonstrates that the initiation of exports to one destination can lead to new export flows to additional destinations: i.e. that acquisition of export experience can result in new trade relationships. We find a related result: that tariff cuts for LDCs are only effective when export experience is present, so that the impact of these policies is essentially limited to the stimulation of additional episodes of sequential exporting.[5]

The remainder of the paper proceeds as follows. In Section 2, we present the theoretical framework that underpins our predictions regarding the heterogeneity of the effects of LDC tariff preferences. We next describe our empirical strategy, which is based on this conceptual framework, in Section 3. Section 4 provides information on the reforms we study, as well as discussing our data sources and presenting summary statistics. We present and discuss our results in Section 5, and offer a brief conclusion in Section 6.

3.2 Theoretical framework

To motivate the paper's analysis and shape our empirical strategy, we present a theoretical framework based on the model of Helpman, Melitz and Rubenstein (2008). There are a finite number of countries i, a finite set of products p across which a representative consumer in each country i has Cobb-Douglas preferences, and a continuum of varieties v_p of each product

---

[4] Notably, (Olarreaga) find that the size of tariff rents accruing to African firms from AGOA has also been heterogeneous, with exporters in smaller and poorer countries receiving smaller rents. Also, while (Cadot and Iacovone) do not study tariff preferences, their results are also related to our study of heterogeneity across exporters: they present firm-level evidence that the probability of survival in export markets upon entry is higher for African firms when more other firms from the same country already export the same product to the same importer.

[5] Similarly, Carrère and Strauss-Kahn (2017) show that prior export experience by developing countries (with non-OECD partners) increases the probability of survival of new export relationships with the OECD. Our results suggest that such survival probabilities are also enhanced by tariff cuts.

for which each consumer has CES preferences with elasticity of substitution σ. That is, the utility function of the representative consumer in country i is (where α_{ip} sums to one across p for each i):

$$U_i = \prod_p \left(\int_{v_p} c_i(v_p)^{\frac{\sigma-1}{\sigma}} dv_p \right)^{\frac{\sigma}{\sigma-1}\alpha_{ip}}$$

Assume that firms each produce a single product, that each firm participating in product market p produces a distinct variety, and that each country i has a potentially different measure N_{ip} of firms producing a given product p. Firms sell under conditions of monopolistic competition. Production of a unit of output requires firms to expend a cost $c_i a$ on inputs, where the country-specific cost c_i is set outside the model, and a is a firm-specific (inverse) productivity parameter. As in (59), the distribution $F(a_{ip})$ of a across firms producing product p in country i is bounded, with support $[a^L_{ip}, a^H_{ip}]$. Also, for firms to ship their output to an importer j requires a fixed cost f_{ij} and a product-specific variable cost of the iceberg type, τ_{ijp}.

The assumption of bounded productivity is particularly appropriate for this empirical setting because it allows for situations in which none of the N_{ip} firms active in the domestic market of country i is sufficiently productive to profitably export its output to some (or all) potential importers $j \neq i$. Specifically, a firm in i with productivity a profitably exports to j if the following condition holds:

$$\frac{1}{\sigma} \left(\frac{\sigma P_{jp}}{(\sigma-1)\tau_{ijp} c_i a} \right)^{\sigma-1} E_{jp} - f_{ij} > 0 \tag{3-B.1}$$

where P_{jp} is the ideal price index for product p prevailing in importer j and E_{jp} is the total expenditure of country j on product p. If there are no trade imbalances, the latter equals the national income of country j multiplied by the Cobb-Douglas share α_{jp}.

The total value of exports of product p from country i to country j is then:

$$X_{ijp} = \left(\frac{\sigma P_{jp}}{(\sigma-1)\tau_{ijp} c_i} \right)^{\sigma-1} E_{jp} N_{ip} V_{ijp}$$

where V_{ijp} is as follows:

$$V_{ijp} = \begin{cases} \int_{a^L_{ip}}^{a_{ijp}} a^{1-\sigma} dF(a_{ip}) \text{ if } a_{ijp} \geq a^L_{ip} \\ 0 \text{ if } a_{ijp} < a^L_{ip} \end{cases}$$

Here a_{ijp} represents the cutoff productivity for which the profits from exporting product p

from i to j equal zero.

A reduction in the tariff on imports of p from i to j corresponds to a fall in τ_{ijp}. It should be clear from the above that if initial imports equal zero, this tariff cut will not necessarily have any effect on the value of trade between i and j. Specifically, if the distance between a_{ip}^L and a_{ijp} is sufficiently large, the fall in τ_{ijp} will not be sufficient to bring the most productive firm producing p in country i to the export threshold. Instead, the country-product pairs most likely to begin exporting to j due to such a tariff cut are those where a_{ip}^L is initially only slightly below a_{ijp}.

We can use this model to predict which exporter-product pairs might fit these criteria, and therefore respond on the extensive margin to such a tariff cut. We can rewrite equation (1) in terms of the cutoff productivity level a_{ijp} to clarify the scenario in which there are no firms from i exporting product p to j:

$$\begin{aligned} a_{ip}^L &> a_{ijp} \\ &= \frac{\sigma P_{jp}}{(\sigma-1)\tau_{ijp}c_i}\left(\frac{E_{jp}}{\sigma f_{ij}}\right)^{\frac{1}{\sigma-1}} \end{aligned}$$

Countries responding on the extensive margin to a small cut in τ_{ijp} will initially satisfy this inequality only marginally. Other than τ_{ijp}, the inequality varies in three parameters specific to the exporter i: c_i, f_{ij} and a_{ip}^L. All three of these parameters must be sufficiently small for j to be a marginal case.

This leads us to some initial conditions on the pre-tariff-cut distribution of exports that may predict an extensive margin response. First, if f_{ij} is sufficiently small, then country i is more likely to already be exporting at least one other product to country j. Moreover, if a_{ip}^L is positively correlated across closely related products within country i, then positive exports by i of products similar to p to j are a potential predictor of an extensive margin response to a cut in τ_{ijp}. In other words, a country already exporting similar products to a given destination may be more likely to begin exporting new products to that destination when their tariffs are reduced. Second, if c_i and a_{ip}^L are both relatively small, then i is more likely to already be exporting the same product p to a destination other than j.

These observations suggest that the probability that a tariff cut leads to a new product-importer-exporter flow may be related to initial 'export experience', in the sense that the exporter already sends similar products to that importer, or sends the same product to another importer. This prediction is analogous to those of the 'sequential exporting' literature, in which export experience makes it more likely for a firm to enter additional markets, for example by allowing the firm to learn about potential export profitability (Albornoz et al. 2012) or by reducing its entry costs into other markets (Morales, Sheu and Zahler 2019). We build on the findings of this literature by considering export experience as a potential factor affecting the impact of a trade policy, rather than discussing its effect on trade dynamics in the absence of policy interventions.

3.3 Empirical strategy

In this section, we use the basic ingredients of the theoretical framework, along with our hypothesis on the potential relevance of export experience, to construct an empirical strategy. As noted by Helpman, Melitz and Rubenstein (2008), the condition under which the profits from exporting product p to importer j are positive for at least one firm in exporter i (equation (1)) can be rewritten in the following multiplicative way:

$$\frac{1}{\sigma}\left(\frac{\sigma}{\sigma-1}\right)^{\sigma-1} c_i^{-\sigma}\left(a_{ip}^{L}\right)^{1-\sigma} f_{ij}^{-1} P_{jp}^{\sigma-1} E_{jp} \tau_{ijp}^{1-\sigma} > 1$$

Allowing each of these variables to vary by time t, and defining a variable Z_{ijpt} equal to the left-hand side, gives:

$$\ln Z_{ijpt} = \zeta - \sigma \ln c_{it} + (1-\sigma)\ln a_{ipt}^{L} - \ln f_{ijt} + (\sigma-1)\ln P_{jpt} + \ln E_{jpt} + (1-\sigma)\ln \tau_{ijpt} \quad \text{(3-C.2)}$$

We decompose τ_{ijpt} into an importer-exporter-time-specific component τ_{ijt} (including factors such as time-varying distance effects), an importer-product-time-specific component τ_{ipt} (including factors such as time-varying most-favoured-nation tariffs), an importer-exporter-product-specific component τ_{ijp} (including factors such as time-invariant preferential tariffs), the treatment variable T_{ijpt} and an error component. Then replacing all variables except the treatment and the error term with these fixed effects yields the following equation:

$$\ln Z_{ijpt} = \beta T_{ijpt} + \nu_{ijp} + \phi_{ijt} + \theta_{ipt} + \psi_{jpt} + \epsilon_{ijpt}$$

In theory, equation (2) captures the average extensive-margin effect of the treatment via a threshold rule: once the right-hand side exceeds zero, i begins exporting product p to j. We consider the possible heterogeneity of this effect depending on observed initial 'export experience', as outlined in the previous section. Specifically, we define a dummy variable I_{ijp} that is equal to one if exports of p by i to any importer in the sample are positive in the first year of the sample period (1996), and/or if exports by i to j of products in the same two-digit category as p are positive in 1996. We then interact this new export experience variable with the treatment variable T_{ijpt}.

Because of the large number of observations in our sample and the high dimensionality of our fixed effects, we estimate a linear probability model using ordinary least squares. On the left-hand side of our estimating equation, we use a dummy variable W_{ijpt} equal to one if there is a positive trade flow in cell $ijpt$ and zero otherwise. The regression we run is therefore as follows:

$$W_{ijpt} = \beta T_{ijpt} + \gamma T_{ijpt} * I_{ijp} + \nu_{ijp} + \phi_{ijt} + \theta_{ipt} + \psi_{jpt} + \omega_{ijt} * I_{ijp} + \kappa_{ipt} * I_{ijp} + \lambda_{jpt} * I_{ijp} + \epsilon_{ijpt} \quad \text{(3-C.3)}$$

Because this regression includes a full set of three-dimensional fixed effects, the identification of treatment effects is through a quadruple-difference strategy. This builds naturally on the existing literature. Recent gravity-based estimates of the effects of trade agreements, such as Baier, Bergstrand and Feng (2014), rely on aggregate trade flows and include importer-exporter, importer-time and exporter-time fixed effects: a triple-difference identification strategy. We augment this strategy here by using variation in the treatment at the product level. Similarly, the study of the effects of the African Growth and Opportunity Act by Frazer and Van Biesebroeck (2010) used a triple-difference strategy with product-level variation, but for trade flows to a single importer (the US). They thus included fixed effects by exporter-time, exporter-product and product-time. We build on their strategy by estimating the effects of multiple program reforms across several importers, accounting for factors potentially correlated to the treatment by importer, product and time (such as the end of the Multifiber Arrangement).

As discussed in the introduction, this identification strategy captures effects of tariff cuts – since we compare products that are treated by tariff cuts to those that are not – but excludes any impacts of changes in fixed costs (such as rules of origin) or any other aspects of the reforms that vary at the importer-exporter-year level. In our regressions, these impacts are absorbed by fixed effects, and so our estimates may thus capture only part of the full effect of the reforms we study.[6]

A second (more standard) caveat to our strategy is that our estimates may be affected by the presence of trade diversion. For example, new trade preferences may have led beneficiaries

---

[6] Note that these reforms also usually cut product-level import quotas to zero for LDC beneficiaries. However, such changes in quantitative restrictions are unlikely to be relevant to our study, given that we focus on effects on the extensive margin, rather than growth in trade volume for existing flows.

to divert exports from one importer in the sample to another, or might alternatively have resulted in relocation of production between exporters in the sample (as suggested by Borchert (2009) in the case of the European GSP scheme). Our estimates might also be influenced in the opposite direction by the effects identified in the sequential exporting literature: the successful initiation of exports to one of the importers in our sample due to tariff cuts might increase the likelihood of nonzero exports of the same product to another destination.

3.4 Data and summary statistics

3.4.1 Importers and program reforms

We begin assembling our dataset by identifying reforms of OECD countries' trade preference programs for LDCs during our sample period of 1996 to 2011. This period begins just after the conclusion of the Uruguay Round agreements and the WTO Plan of Action for the LDCs, and ends before another major pivot in EU trade policy towards developing countries.[7] Also, we focus on the programs of OECD members because of the disproportionate importance of these importers in world trade, and because similar preferences offered by some non-OECD members (such as China and Russia) were generally introduced much later than OECD members' reforms. In practice, there was a wave of reforms of OECD countries' trade policy towards LDCs in the first several years after the end of the Uruguay Round.

Of the 26 countries who were members of the OECD as of the beginning of 1996, all but two made significant changes to an LDC trade preference program at least once during our sample period.[8] Fifteen of these countries are European Union members, and thus maintain a unified trade policy; we consider the EU as a single importer throughout the paper so as not to put disproportionate weight on EU trade preferences in our estimates. Also, we exclude Iceland from the analysis because of the unavailability of its product-level tariff schedule for most of the sample period, and its especially small size.[9] This leaves us with nine importers: Australia, Canada, the EU, Japan, New Zealand, Norway, Switzerland, Turkey and the US.

We identify the timing of each major reform by these importers using information from

---

[7] The EU outlined significant changes to its trade preference programs in 2012 (and, as noted earlier, notified the WTO of these changes in 2013). This included the removal of a large number of countries from its GSP program, in order to increase the preference margin enjoyed by LDCs relative to other developing countries.

[8] The two exceptions are the Czech Republic, which joined the EU later in the sample period, and Mexico.

[9] In 1996, the number of products imported by Iceland from individual exporters in our sample (i.e. nonzero flows by exporter-product) was less than one-third as many as the next smallest of our sample of importers (Norway). While WTO documents indicate that Iceland implemented an LDC trade preference program in 2002, we observe only its MFN tariffs and preferences for the European Economic Area in the WTO tariff data throughout the sample period.

Table 3-1: Reforms of OECD programs for LDCs, 1996-2011

Importer	Program	Year	Countries	Products
Australia	Expansion of GSP-LDC	2003	48	750
Canada	Expansion of GSP-LDC	2000	47	342
	Expansion of GSP-LDC (MAI)	2003	47	536
European Union	Expansion of GSP-LDC (EBA)	2001	47	155
Japan	Expansion of GSP-LDC	2001-07	48	496
New Zealand	Expansion of GSP-LDC	2001	48	209
Norway	Expansion of GSP-LDC	2008	14	269
Switzerland	Expansion of GSP-LDC	1998	47	119
	Expansion of GSP-LDC	2008	47	101
Turkey	Expansion of GSP-LDC	2004	47	1,094
United States	Expansion of GSP-LDC	1998	43	758
	Creation of AGOA	2001	38	260
	Expansion of AGOA	2007	38	436

This table displays a list of initiations and expansions of trade preference programs for least developed countries during the sample period by members (as of 1996) of the OECD. 'GSP-LDC' refers to the country's Generalized System of Preferences program for least developed countries. 'EBA' refers to the EU's 'Everything but Arms' program, and 'MAI' refers to Canada's Market Access Initiative, both of which represented expansions of existing GSP-LDC programs. 'AGOA' refers to the African Growth and Opportunity Act in the US. The column titled 'Year' lists the year(s) of implementation of each program's reforms. The column titled 'Countries' displays a count of exporters directly affected by the reform, restricted to countries in the dataset used in this paper. The column titled 'Products' displays a count of products directly affected by the reform, for six-digit products in the HS classification.

notifications to the WTO, as well as handbooks on each country's GSP scheme (including subschemes for LDCs) published by UNCTAD. We study the thirteen reforms listed in Table 3-1, almost all of which were expansions of existing LDC preference programs. The exception is the creation and later expansion of the United States' African Growth and Opportunity Act (AGOA), which is not an LDC-specific program, but whose beneficiaries are mainly LDCs because of its focus on Africa.[10]

3.4.2 Exporters

Most of the importers in our sample determine LDC programs' beneficiaries using a list of the least developed countries defined by the United Nations. However, there is some variation in the list of beneficiaries across importers. In order to collect data by importer on eligibility for treatment as an LDC, we again refer to WTO notifications and UNCTAD handbooks, as well as government documents from the importers in our sample (such as EU regulations). Table 3-1 displays the number of exporters that are defined as beneficiaries by each importer, while Table 3-A.1 provides the full list of countries in each case.[11]

While most of the entries in Table 3-1 share a very similar list of eligible exporters, two exceptions are notable. First, as previously mentioned, AGOA is targeted only to African countries. This includes several countries that are not on the UN list of LDCs. Also, Norway's 2008 LDC program reform significantly expanded its list of beneficiaries rather than broadening the range of products covered. While we include both of these sets of countries in our baseline regressions as beneficiaries of these reforms, we also perform robustness checks in which they are dropped from the sample. Finally, because individual countries in the South African Customs Union are treated differently by importers (e.g. Lesotho is usually considered an LDC while South Africa is not) but are not separately observed in the trade data we use, we exclude these countries from our sample of exporters.[12] This leaves 66 program beneficiaries in our sample, 18 of which were only included in AGOA or the 2008 Norwegian reform.

Because our identification strategy depends on a comparison of reforms' beneficiaries to nonbeneficiaries, we also include additional exporters in our sample. We add only coun-

---

[10] We do not include contemporaneous US programs for Caribbean countries in this list because very few of their beneficiaries are defined as LDCs by any LDC-specific program.

[11] Note that in practice, there were some changes in LDC eligibility over time. For example, Senegal was added as an LDC, and Cape Verde and the Maldives ceased to be defined as LDCs, by the United Nations during the sample period, though not all of the importers we study immediately adopted these changes. In our analysis, we use a time-invariant definition of beneficiaries and consider cases such as these to be beneficiaries.

[12] Specifically, we drop Botswana, Lesotho, Namibia, Swaziland and South Africa from the sample.

tries that are comparable to treated exporters on two key dimensions: they are developing countries and started the sample period with relatively undiversified exports. The former criterion excludes small but relatively wealthy countries such as Bahrain and Bermuda, while the latter rules out large developing countries such as China, India and Indonesia.

Specifically, we define developing countries as low-income and low-middle-income countries according to the World Bank classification as of 1996. Because the World Bank criteria differ from those of the UN, the LDCs in the UN list span both of these categories (e.g. Burundi is low-income while Djibouti is a low-middle-income country). However, the set of countries meeting this World Bank definition is wider, including 59 countries that are neither on the UN list nor benefit from AGOA or the Norwegian program expansion.

We measure export diversification at the beginning of the sample period by counting each exporter's nonzero flows by importer-product in 1996.[13] We then drop any exporters whose initial level of export diversification, by this measure, was larger than any exporter in our initial sample. That is, in order to be added to the sample, a country cannot have had a wider range of 1996 export flows than the most diversified beneficiary of the LDC program reforms in Table 3-1. In practice, Bangladesh was the LDC program beneficiary with the largest set of flows by importer-product as of 1996. We therefore drop all countries whose exports were more diversified than those of Bangladesh in 1996. This leaves us with 34 exporters to be added to the sample, for a total of 100 exporters. These are all listed in Table 3-A.1.

3.4.3 Affected products and tariffs

In order to identify the set of products affected by each reform, we use tariff data from the World Trade Organization, sourced via the WTO Tariff Download Facility. While the WTO notifications, UNCTAD handbooks and government documents mentioned above include lists of treated products for some (though not all) reforms, this WTO data gives us additional relevant information. Specifically, it provides each importer's annual tariff schedule, including most-favoured nation (MFN) tariffs, as well as preferential tariffs for GSP and LDC beneficiaries (and AGOA beneficiaries for the US).[14] While our main specifications use a dummy variable for treatment, this data allows us to also quantify the treatment in terms of the size of the relevant tariff cut, or the gap between the LDC tariff and the GSP or MFN tariff at the time of treatment. For the European Union, we instead use data from UNCTAD TRAINS, supplemented by information from government documents, because the WTO data does not

---

[13] For this exercise, we only include the nine importers in our sample, and use the extensive margin threshold of one thousand US dollars discussed below.

[14] This data also provides ad valorem equivalents of specific tariffs.

provide full information on EU preferential tariffs for our sample period.

In order to compare products across years, we need to take account of the fact that product classifications evolve over time. All of our tariff data over the sample period is defined at the six-digit product level using the Harmonized System (HS) classification, for which concordances between versions are publicly available. Using these, we construct a single concordance that accounts for all changes in six-digit product codes across the 1996, 2002 and 2007 HS versions. When multiple products are merged or divided, we consolidate these into a single code. However, because our export experience variable is defined using information on exports within the same two-digit HS category, we drop concorded product codes that span more than one such category.[15]

Next, to identify products that are treated by tariff cuts, we compare an importer's tariff schedules in the year of each reform and the previous year. We first quantify the tariff facing LDCs for each product in each of these two years. If there is no LDC-specific tariff, then we assume that LDC exporters instead faced the GSP tariff (or the MFN tariff if there is also no GSP tariff for that product). If the tariff facing LDCs declines between one year and the next for a given product, and also falls relative to the prevailing GSP (or MFN) tariff, then we classify that product as treated. We include the second criterion so that cases in which the GSP or MFN tariff is reduced are not mistakenly classified as LDC-specific tariff cuts.[16]

We display the number of six-digit products affected by each reform in the rightmost column of Table 3-1. Then, in Table 3-2, we provide more detailed information about the types of products whose tariffs were reduced by these reforms. We first tabulate, across the 66 exporters affected by at least one reform, the share of importer-exporter-product cells experiencing tariff reductions. This proportion, which appears at the bottom of column (1) of Table 3-2, is approximately 14%; the reforms therefore affected a substantial share of tradable goods.

We then tabulate this proportion in the same way for each of the 21 sections of the HS classification. That is, within each section, we calculate the share of treated importer-exporter-product cells. Column (1) of Table 3-2 shows that there was substantial heterogeneity in the products treated by the reforms; while only 1.4% of importer-exporter-product observations in the category 'Mineral products' (where most MFN tariffs were already equal to zero) were

---

[15] This results in the exclusion of approximately 5% of the six-digit products in the 1996 version of the HS classification. However, because of the large scale of the reclassification of these products (e.g. due to major changes in consumer electronics), our concordance had consolidated these products into just seven six-digit categories. Their omission is thus of minimal importance to our regression results.

[16] Note that we use the generally applicable LDC, GSP and MFN tariffs of each importer, disregarding cases in which specific beneficiaries are treated differently for a subset of products.

Table 3-2: Distribution of products affected by reforms

	Share treated within HS section (1)	Share of HS section in all treated (2)
Animal products	24.4%	7.6%
Vegetable products	17.5%	6.7%
Animal and vegetable oils	23.6%	1.6%
Foods, beverages, tobacco	27.1%	7.9%
Mineral products	1.4%	0.3%
Chemical products	4.9%	5.0%
Plastics and rubber	5.9%	1.6%
Leather and articles	10.6%	0.7%
Wood and articles	3.7%	0.3%
Wood pulp and paper	0%	0%
Textiles and apparel	42.5%	50.5%
Footwear, headgear, etc.	23.3%	1.7%
Ceramics, glass, etc.	7.1%	1.5%
Precious stones/metals	1.0%	0.1%
Base metals and articles	10.7%	9.2%
Machinery and equipment	2.1%	2.2%
Transport equipment	8.3%	1.6%
Instruments	2.7%	0.9%
Arms and ammunition	1.6%	0.04%
Miscellaneous manufactures	3.2%	0.6%
Works of art	0%	0%
All sections	14.2%	

This table displays information on the six-digit HS products affected by the initiations and expansions of trade preference programs studied in this paper, by the 21 sections of the HS classification. The first column tabulates, among exporters affected by at least one of these reforms, the share of importer-exporter-product observations for which tariffs were cut to zero by a reform, within each section. The second column tabulates the number of importer-exporter-product cells in each section for which tariffs were cut to zero by a reform, as a share of the total number of importer-exporter-product cells for which tariffs were cut to zero by a reform.

treated, around 40% of importer-exporter-product cells in the section covering textiles and apparel saw their tariffs reduced by a reform.

In column (2), we display the proportion of all treated importer-exporter-product cells falling into each section. This shows that around half of treated importer-exporter-product cells were textiles and apparel. Note that this is both because of the intensity of treatment within this section, and because this is one of the HS sections with the largest number of six-digit products. In our analysis below, we explore the extent to which our main results are driven by tariff cuts for textiles and apparel.

Finally, in the center of Table 3-3, we summarize the frequency of the main treatment variable used in our analysis below. This is a dummy variable equal to one in the years after a tariff is cut to zero by a particular importer for a given exporter and product due to an LDC program reform. Because all of the reforms we study occurred after 1996, this variable is equal to zero for all observations in the dataset in 1996. However, by 2011, approximately 6.3% of importer-exporter-product observations in the sample had been treated in this way. As already documented in Table 3-2, this proportion is 14% for exporters treated by at least one of the reforms studied here. Across all years in the sample period, the treatment variable is equal to one for 3.7% of observations (8.4% for treated exporters).

3.4.4 Nonzero trade flows

We take data on trade flows from the UN Comtrade database. Although the Comtrade data provides us with the value of goods trade as reported by both exporters and importers, we use only import-side information reported by the nine importers in our sample.[17] We observe the value of trade at the level of the importing country, exporting country, six-digit product and year. We calculate aggregate EU imports as the sum of the imports of the fifteen EU members as of 1996.[18] So that very small flows of a given product between a country pair are not counted as breaches of the extensive margin, we define a nonzero trade flow as a flow greater than or equal to one thousand US dollars. We also check the robustness of our results to alternative definitions below.

In Table 3-3, we show summary statistics for nonzero trade flows, across importer-exporter-product-year cells. This yields a few notable pieces of information about the sample. First,

---

[17] We do this both because import data is likely to be more closely monitored and thus more reliable, and because this reliability is probably further enhanced by the greater local monitoring capacity in OECD countries.

[18] These are Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, the Netherlands, Portugal, Spain, Sweden and the UK. Imports of countries that acceded to the EU during the sample period are not considered in our analysis.

Table 3-3: Summary statistics

	Treated countries (1)	Full sample (2)
Positive trade flows, all years	0.9%	1.4%
Positive trade flows, 1996	0.7%	1.0%
Positive trade flows, 2011	1.1%	1.6%
Treatment dummy, all years	8.4%	3.7%
Treatment dummy, 1996	0%	0%
Treatment dummy, 2011	14.2%	6.3%
Export experience	18.3%	23.4%
Similar product (same HS2), same importer	15.3%	18.7%
Same product, other importer	3.4%	5.6%
Both of the above	1.0%	1.8%
Existing flow	0.7%	1.0%
Treated products	15.0%	

This table displays summary statistics by importer-exporter-product-year observation, for the sample years 1996 to 2011 and for the initial and final year of the sample. Column (1) displays summary statistics for exporters who are subject to at least one of the reforms studied here, while column (2) includes all exporters in the sample. The first three rows show the share of nonzero trade flows across observations, defined using a threshold of 1000 US dollars. The 'Treatment dummy' rows display the share of observations equal to one for a dummy variable for years after a tariff is cut to zero for a particular importer-exporter-product cell. The 'Export experience' rows show the proportion of importer-exporter-product cells in which (i) a product in the same two-digit HS category was exported by the same country to the same importer in 1996; and/or (ii) the same product was exported by the same country to one of the nine importers in the sample in 1996. Shares for each of these two cases, and for importer-exporter-product cells satisfying both criteria, are shown in the next three rows. These three figures exclude cases where there was a positive flow in 1996 for the importer-exporter-product cell itself, which are instead tabulated separately in the next row ('Existing flow'). The final row shows the same proportion displayed in the 'Export experience' row, but calculated using the restricted sample of importer-exporter-product cells treated at some time during the period of study.

products with nonzero trade flows between these importers and exporters are rare: among the more than sixty million importer-exporter-product-year observations in the sample, all but 1.4% are zero trade flows. Second, the 66 countries who benefited from at least one of the reforms we study have only slightly less diversified exports than the sample as a whole. Third, there is an upward trend in export diversification across both the treated exporters and the full sample; the initial share of positive trade flows for the full sample was 1.0%, but by 2011, the share of positive observations had expanded to 1.6%.

3.4.5 Export experience

As discussed in the previous section, we define export experience as a time-invariant characteristic of an importer-exporter-product cell, based on conditions at the beginning of the sample period. It is equal to one if either of the following conditions is satisfied: first, if a product in the same two-digit HS category is exported by the same country to the same importer in 1996; second, if the same product is exported by the same country to one of the nine importers in the sample in 1996.

As displayed at the bottom of Table 3-3, such cases constitute around 20% of importer-exporter-product cells: 23.4% as a share of the full sample, and 18.3% among treated countries. The table also shows that the most common determinant of export experience is the first criterion (export of a similar product to the same importer), which is satisfied much more often than the second criterion. In approximately 1.8% of cases (1.0% for treated countries), a cell is classified as having export experience by satisfying both of the criteria above.

Note that the figures cited so far exclude cases where there is a positive flow in 1996 for the importer-exporter-product cell itself. However, we do count such existing trade flows as cases of export experience, since they satisfy both of the criteria above. As already noted in the previous subsection, trade flows were positive for approximately 0.7% of importer-exporter-product cells among treated countries as of 1996. We will later divide our export experience variable into cases of existing positive flows and all other cases, in order to separately identify the effects of tariff cuts on trade flow survival and sequential exporting.

Finally, in the last row of the table, we show that importer-exporter-product cells that are treated at some point in the sample are not substantially more or less likely to be cases in which treated exporters have export experience. Specifically, the proportion of cases in which export experience equals one for treated exporters is similar whether or not one restricts the sample to treated products only (15.0% vs. 18.3%).

3.5 Results

In this section, we present the findings of our empirical analysis, multiplying all estimated coefficients by 100 so that our tables display the estimated effects of the reforms in percentage points. As usual, it is appropriate to cluster standard errors at a cross-sectional level; we cluster by importer-exporter pair, as in Baier, Bergstrand and Feng (2014).

We first discuss our main findings: i.e. the results of estimating equation (3) with and without interaction terms with the export experience dummy I_{ijp}. Column (1) of Table 3-4 reveals that treatment by a reform has a statistically significant impact on whether a treated exporter sends any goods to the importer in the treated product category. This average effect of the reforms is approximately 0.14 percentage points, which is relatively large in terms of the initial average level of the dependent variable for treated exporters. Specifically, the estimated effect is equal to approximately 20% of the proportion of importer-exporter-product cells in which treated exporters had positive trade flows as of 1996 (see Table 3-3). The p-value of the estimated coefficient in column (1) is 0.096.

However, this effect is highly concentrated in cells where the export experience dummy equals one. According to the results in column (2), the mean effect of treatment given initial export experience is 1.23 percentage points (the sum of the two displayed coefficients). At the same time, the average impact of treatment in the absence of export experience is estimated to be substantially smaller in magnitude. This is not just a matter of statistical significance, although it is notable that the estimated coefficient on the interaction term is significant at the 1% level while the coefficient on the treatment dummy is insignificantly different from zero. The estimated size of the effect in cases without initial export experience is just 0.015 percentage points, almost ten times smaller than the average impact in column (1) and nearly a hundred times smaller than the effect conditional on initial export experience. The findings in column (2) are our main results, suggesting that the impact of these reforms has been limited to the encouragement of trade flow survival and additional episodes of 'sequential exporting', rather than stimulating wider export diversification.

Column (3) breaks down the effect for cells with export experience into these two sub-categories. To do so, we define separate variables for importer-exporter-product cells with existing trade flows in 1996 and all other cases of export experience, and include interaction terms between the treatment variable (and fixed effects) and each component separately. This reveals that that the average impact of treatment is to increase the probability of trade flow survival by approximately two percentage points, although this effect (relative to that for cells without export experience) is imprecisely estimated. At the same time, the effect of

Table 3-4: Baseline results and export experience decomposition

	(1)	(2)	(3)	(4)
Treatment dummy	0.141	0.015	0.015	0.015
	(0.085)	(0.026)	(0.026)	(0.027)
Treatment dummy × Export experience		1.217		
		(0.409)		
Treatment dummy × Existing flow			1.967	1.967
			(1.733)	(1.752)
Treatment dummy × Related export experience			0.523	
			(0.293)	
Treatment dummy × Similar product (same HS2)				0.036
				(0.289)
Treatment dummy × Other location (same HS6)				0.516
				(0.307)
Treatment dummy × Both				3.428
				(1.033)
Exporter-importer-product FEs	Yes	Yes	Yes	Yes
Exporter-importer-year FEs	Yes	Yes	Yes	Yes
Exporter-product-year FEs	Yes	Yes	Yes	Yes
Importer-product-year FEs	Yes	Yes	Yes	Yes
Export experience interactions	No	Yes	Yes	Yes
Clusters (exporter-importer)	900	900	900	900
Observations	63,403,200	63,403,200	63,403,200	63,403,200

This table displays estimated effects of reforms in OECD countries' trade preference programs for the least developed countries. The dependent variable is a dummy for positive exports, with a threshold of 1000 US dollars. Each observation is an importer-exporter-product-year; in the full sample, this includes nine importers, 100 exporters, 4,403 products at the six-digit level of the HS classification, and sixteen years from 1996 to 2011. The variable 'treatment dummy' equals one in years after the import tariff is cut for an exporter-product as part of an LDC trade preference program by an importer. The variable 'export experience' is a dummy that equals one for importer-exporter-product cells for which either the exporter had at least one positive trade flow to the same importer in the same two-digit HS category in 1996, or the exporter had at least one positive trade flow to any importer in the sample for the same six-digit product in 1996. In column (3), export experience is separated into two dummy variables, for importer-exporter-product cells with existing trade flows in 1996 ('existing flow') and all other cases of export experience ('related export experience'). In column (4), the latter variable is itself subdivided into dummies according to whether a similar product was exported to the same importer in 1996 ('similar product (same HS2)'), or the same product was exported to another importer ('other location (same HS6)'), or both. All specifications include exporter-importer-year, exporter-product-year and importer-product-year fixed effects. Columns (2), (3) and (4) also include interactions between the export experience variable(s) and the fixed effects. All regressions are estimated using ordinary least squares. All coefficients are multiplied by 100 so that they represent effects in percentage points. Standard errors, clustered by importer-exporter pair, are in parentheses.

a tariff cut for all other cases of export experience is estimated to be 0.52 percentage points larger than the effect in cases without experience, with a p-value of 0.075. This is direct support for the hypothesis that the reforms have had a substantial impact on the occurrence of sequential exporting.

We consider the determinants of this sequential exporting effect in column (4), where we further subdivide the export experience variable according to whether a similar product was exported to the same importer in 1996, or the same product was exported to another importer, or both. The largest treatment effect is apparent in cases where both of these conditions were in place, which presumably corresponds to situations where the exporter was most likely to be on the margin of exporting the product to the importer in question. The occurrence of a tariff cut due to an LDC reform made the exporter around 3.4 percentage points more likely to breach this margin (as compared to cases without export experience), according to the estimate in Table 3-4. Tariff cuts are also estimated to have had smaller, but statistically significant, effects for cells where the same product had already been exported to another importer in the sample, but no product in the same two-digit HS category had been exported to the same importer.

Notably, for cases satisfying the 'similar product, same importer' criterion only, the estimated treatment effect is very small and not significantly different from that of cells without export experience. Recall from Table 3-3 that these cases make up the large majority of importer-exporter-product cells with export experience. This implies that the intensive-margin effects of reforms to LDC programs have been even more concentrated than suggested by our initial results, as they are entirely driven by the relatively rarer situations of export experience in which the product itself was already being exported to at least one of the importers in the sample.

A possible concern with our baseline results is that the estimated heterogeneity might be driven by the presence of some very small countries, such as Pacific island nations, in the sample. These countries are likely to have a narrow range of initial exports, so that the export experience dummy is probably usually equal to zero in these cases. At the same time, their capacity for export diversification may be severely constrained by the simple fact of their size. However, when we re-estimate the first two regressions dropping the 21 sample exporters with a population under one million as of 1996, the results remain similar (see columns (1) and (2) of Table 3-5).[19]

---

[19] We drop Belize, Bhutan, Cape Verde, Comoros, Djibouti, Dominica, Equatorial Guinea, Fiji, Grenada, Guyana, Kiribati, Maldives, Marshall Islands, Micronesia, Samoa, Sao Tome and Principe, Seychelles, Solomon Islands, Suriname, Tuvalu and Vanuatu. Population data is sourced from the World Bank's World Development Indicators.

Table 3-5: Robustness checks - alternative country samples

	Excluding small countries		Excluding AGOA and Norway 2008	
	(1)	(2)	(3)	(4)
Treatment dummy	0.152	0.004	0.137	0.018
	(0.104)	(0.032)	(0.100)	(0.033)
Treatment dummy *		1.230		1.273
export experience		(0.448)		(0.491)
Exporter-importer-product FEs	Yes	Yes	Yes	Yes
Exporter-importer-year FEs	Yes	Yes	Yes	Yes
Exporter-product-year FEs	Yes	Yes	Yes	Yes
Importer-product-year FEs	Yes	Yes	Yes	Yes
Export experience interactions	No	Yes	No	Yes
Clusters	711	711	738	738
Observations	50,088,528	50,088,528	51,990,624	51,990,624

This table displays estimated effects of reforms in OECD countries' trade preference programs for the least developed countries. The dependent variable is a dummy for positive exports, with a threshold of 1000 US dollars. Each observation is an importer-exporter-product-year; in the full sample, this includes nine importers, 100 exporters, 4,403 products at the six-digit level of the HS classification, and sixteen years from 1996 to 2011. The 21 exporters with 1996 populations smaller than one million are excluded from the sample in columns (1) and (2). The eighteen exporters that are beneficiaries only of the African Growth and Opportunity Act (AGOA) and/or Norway's 2008 expansion of its GSP-LDC program are excluded from the sample in columns (3) and (4). The variable 'treatment dummy' equals one in years after the import tariff is cut for an exporter-product as part of an LDC trade preference program by an importer. The variable 'export experience' is a dummy that equals one for importer-exporter-product cells for which either the exporter had at least one positive trade flow to the same importer in the same two-digit HS category in 1996, or the exporter had at least one positive trade flow to any importer in the sample for the same six-digit product in 1996. All specifications include exporter-importer-year, exporter-product-year and importer-product-year fixed effects. Columns (2) and (4) also include interactions between the export experience variable and the fixed effects. All regressions are estimated using ordinary least squares. All coefficients are multiplied by 100 so that they represent effects in percentage points. Standard errors, clustered by importer-exporter pair, are in parentheses.

In columns (3) and (4) of Table 3-5, we instead exclude countries that are not on the UN list of the least developed countries, but are included in AGOA by the US, and/or in Norway's 2008 expansion of its LDC trade preferences program. We perform this robustness exercise in order to establish whether the estimated effects remain similar when the treatment group includes only the 48 exporters that are the beneficiaries of most importers' LDC programs.[20] This sample restriction also has little effect on our estimates. Overall, the only notable change from Table 3-4 to Table 3-5 is that the estimated average impact of the reforms shifts from slightly below to somewhat above the 10% threshold for statistical significance. This is mainly due to a rise in the standard error rather than a change in the point estimate, which is unsurprising given the diminished size of the sample.

So far, we have defined the extensive margin in terms of a \$1000 cutoff by importer, exporter, product and year. In Table 3-6, we check whether the results are robust to changing this definition. We first allow any nonzero level of imports, even if below \$1000, to constitute a trade flow. As shown in columns (1) and (2), this makes little difference to our findings.

We next impose a stricter definition, requiring an observed import value of \$100,000 or more in order to qualify as a trade flow. This does not result in a substantial change to the two estimated coefficients in our regression with an interaction term (see column (4)). But because the share of importer-exporter-product cells with export experience is now smaller due to the higher threshold, the average effect (see column (3)) is approximately half the size of that in Table 3-4 column (1). In other words, among cases with experience at the higher threshold, the estimated impact of the reforms is similar to before, but this is now a smaller share of the sample. So there are fewer breaches of the (higher) extensive margin due to the reforms, and the overall average effect of the reforms is smaller.

Because a high proportion of products treated by reforms to LDC tariff preference programs are textiles or apparel (see Table 3-2), we next check the extent to which our main results are driven by this product category. In columns (1) and (2) of Table 3-7, we display estimates for a sample excluding all goods in the 'Textiles and apparel' section of the HS product classification. We find that the average effect in column (1) is slightly larger than the result for the full sample. It is also somewhat more precisely estimated, so that the effect is now statistically significant at the 1% level.

More notably, the heterogeneous effects documented for the full sample are still present for non-textile products, but are greatly diminished. This is for two reasons. First, the total effect of tariff cuts for these products on 'experienced' importer-exporter-product cells is

---

[20] Note that this includes Myanmar, which is a beneficiary only for three importers in our sample, but is on the list of least developed countries compiled by the UN.

Table 3-6: Robustness checks - varying extensive margin cutoff

	Cutoff = zero		Cutoff = $100,000	
	(1)	(2)	(3)	(4)
Treatment dummy	0.149	0.007	0.069	0.003
	(0.097)	(0.035)	(0.052)	(0.014)
Treatment dummy × Export experience		1.131		1.373
		(0.439)		(0.670)
Exporter-importer-product FEs	Yes	Yes	Yes	Yes
Exporter-importer-year FEs	Yes	Yes	Yes	Yes
Exporter-product-year FEs	Yes	Yes	Yes	Yes
Importer-product-year FEs	Yes	Yes	Yes	Yes
Export experience interactions	No	Yes	No	Yes
Clusters (exporter-importer)	900	900	900	900
Observations	63,403,200	63,403,200	63,403,200	63,403,200

This table displays estimated effects of reforms in OECD countries' trade preference programs for the least developed countries. The dependent variable is a dummy for positive exports, with a threshold of zero US dollars in columns (1) and (2) and 100,000 US dollars in columns (3) and (4). Each observation is an importer-exporter-product-year; in the full sample, this includes nine importers, 100 exporters, 4,403 products at the six-digit level of the HS classification, and sixteen years from 1996 to 2011. The variable 'treatment dummy' equals one in years after the import tariff is cut for an exporter-product as part of an LDC trade preference program by an importer. The variable 'export experience' is a dummy that equals one for importer-exporter-product cells for which either the exporter had at least one positive trade flow to the same importer in the same two-digit HS category in 1996, or the exporter had at least one positive trade flow to any importer in the sample for the same six-digit product in 1996. All specifications include exporter-importer-year, exporter-product-year and importer-product-year fixed effects. Columns (2) and (4) also include interactions between the export experience variable and the fixed effects. All regressions are estimated using ordinary least squares. All coefficients are multiplied by 100 so that they represent effects in percentage points. Standard errors, clustered by importer-exporter pair, are in parentheses.

Table 3-7: Additional results

	Excluding textiles/apparel		Tighter definition
	(1)	(2)	(3)
Treatment dummy	0.181	0.068	0.083
	(0.050)	(0.030)	(0.042)
Treatment dummy * export experience		0.349	0.611
		(0.238)	(0.198)
MFN tariff			
MFN tariff * export experience			
Exporter-importer-product FEs	Yes	Yes	Yes
Exporter-importer-year FEs	Yes	Yes	Yes
Exporter-product-year FEs	Yes	Yes	Yes
Importer-product-year FEs	Yes	Yes	Yes
Export experience interactions	No	Yes	Yes
Clusters	900	900	900
Observations	52,704,000	52,704,000	63,403,200

This table displays estimated effects of reforms in OECD countries' trade preference programs for the least developed countries. The dependent variable is a dummy for positive exports, with a threshold of 1000 US dollars. Each observation is an importer-exporter-product-year; in the full sample, this includes nine importers, 100 exporters, 4,403 products at the six-digit level of the HS classification, and sixteen years from 1996 to 2011. The 743 products in the 'textiles and apparel' section of the HS classification are excluded from the sample in columns (1) and (2). The variable 'treatment dummy' equals one in years after the import tariff is cut for an exporter-product as part of an LDC trade preference program by an importer. In columns (1) and (2), the variable 'export experience' is a dummy that equals one for importer-exporter-product cells for which either the exporter had at least one positive trade flow to the same importer in the same two-digit HS category in 1996, or the exporter had at least one positive trade flow to any importer in the sample for the same six-digit product in 1996. In column (3), the definition of export experience is adjusted so that the first criterion is instead that an exporter had at least one positive trade flow to the same importer in the same four-digit HS category in 1996. All specifications include exporter-importer-year, exporter-product-year and importer-product-year fixed effects. Columns (2) and (3) also include interactions between the export experience variable and the fixed effects. All regressions are estimated using ordinary least squares. All coefficients are multiplied by 100 so that they represent effects in percentage points. Standard errors, clustered by importer-exporter pair, are in parentheses.

approximately one-third as large as the full-sample estimate. Second, the impact on importer-exporter-products lacking experience is almost five times larger (and is now statistically significant with a p-value of 0.025). So the sharper results we observe in the full sample are driven by both textiles and non-textile products. But they are reinforced by an especially strong role for export experience in driving the effects of tariff cuts in the textiles and apparel sector.

Column (3) of Table 3-7 displays the results from a regression using a modified definition of export experience, which defines 'similar products' as goods in the same four-digit (rather than two-digit) HS category. Although column (4) of Table 3-4 showed that importer-exporter-product cells satisfying only the 'similar product, same importer' criterion have seen little benefit from tariff cuts, this restriction nonetheless makes a substantial difference to the estimates. As shown in Table 3-7, the estimated coefficient for cases without experience increases, while the estimate for the interaction term is halved as compared to our baseline regression. This is due to situations where a country already exported the same product to a different importer, and also had pre-existing trade flows in the same two-digit category – but not the same four-digit category – to the same importer.[21] The findings of column (3) indicate that such cases are important drivers of our main results.

Finally, in Table 3-8 we consider the impact of modifying our treatment variable to reflect the intensity of treatment. In the first two columns, we estimate the effects of the tariff cuts in terms of the size of the cut in each case. To do this, we first calculate the year-to-year reduction in the tariff faced by LDCs for treated products at the time of each reform. We then multiply this by our original dummy variable for treatment, so that in each year after a given reform, our new variable is equal to the magnitude of the original product-level tariff cut in percentage points.

While the basic pattern of heterogeneous effects remains as in our baseline regressions, one change in the results is notable: the impact of the reforms is now statistically significant even in cases without export experience (with a p-value of 0.93; see column (2)). This finding makes sense when evaluated in terms of the theoretical framework. While the original dummy variable specification captured an average effect across all products whose tariffs were reduced, this regression identifies the differential effect of larger tariff cuts – i.e. those that are more likely to bring even an infra-marginal case over the threshold for exporting.

Columns (3) and (4) repeat this exercise, but instead multiplying the treatment dummy

---

[21] These could, for example, be cases where tariffs were reduced for knitted sweaters, and a country already exported knitted mittens but not knitted sweaters to that importer, while simultaneously sending knitted sweaters to other importers in our sample.

Table 3-8: Intensity of treatment

	Size of tariff cut		Gap with GSP/MFN	
	(1)	(2)	(3)	(4)
Tariff cut	0.205	0.044		
	(0.087)	(0.026)		
Tariff cut × Export experience		3.801		
		(1.166)		
Gap with GSP/MFN			0.148	0.024
			(0.094)	(0.020)
Gap with GSP/MFN × Export experience				4.017
				(1.479)
Exporter-importer-product FEs	Yes	Yes	Yes	Yes
Exporter-importer-year FEs	Yes	Yes	Yes	Yes
Exporter-product-year FEs	Yes	Yes	Yes	Yes
Importer-product-year FEs	Yes	Yes	Yes	Yes
Export experience interactions	No	Yes	Yes	Yes
Clusters (exporter-importer)	900	900	900	900
Observations	63,403,200	63,403,200	63,134,989	63,134,989

This table displays estimated effects of reforms in OECD countries' trade preference programs for the least developed countries. The dependent variable is a dummy for positive exports, with a threshold of 1000 US dollars. Each observation is an importer-exporter-product-year; in the full sample, this includes nine importers, 100 exporters, 4,403 products at the six-digit level of the HS classification, and sixteen years from 1996 to 2011. The variables 'tariff cut' and 'gap with GSP/MFN' take fixed positive values in years after the import tariff is cut for an exporter-product as part of an LDC trade preference program by an importer. Specifically, 'tariff cut' is equal to the size of the tariff cut for that product in the year of the reform relative to the year before, while 'gap with GSP/MFN' is equal to the difference between the LDC and GSP tariffs in the year of the reform, or the difference between the LDC and MFN tariffs if there is no GSP tariff. The variable 'export experience' is a dummy that equals one for importer-exporter-product cells for which either the exporter had at least one positive trade flow to the same importer in the same two-digit HS category in 1996, or the exporter had at least one positive trade flow to any importer in the sample for the same six-digit product in 1996. All specifications include exporter-importer-year, exporter-product-year and importer-product-year fixed effects. Columns (2) and (4) also include interactions between the export experience variable and the fixed effects. All regressions are estimated using ordinary least squares. All coefficients are multiplied by 100 so that they represent effects in percentage points. Standard errors, clustered by importer-exporter pair, are in parentheses.

by the gap between the LDC tariff and the GSP tariff (or the MFN tariff if there is no GSP tariff) in the year after a reform. In other words, if immediately after a tariff reduction there is a difference of ten percentage points between the LDC and GSP tariffs for a given product, this becomes our measure of treatment intensity for that product. In contrast to the previous result, we find that this measure is not significantly related with the presence of a nonzero trade flow for cells without export experience.

3.6 Conclusion

In this paper, we have studied the effects of a number of LDC trade preference programs initiated or expanded by OECD countries in the years following the conclusion of the Uruguay Round. We have first noted that the model of Helpman, Melitz and Rubenstein (2008) predicts that only a subset of importer-exporter-product cells that are close to an export threshold should benefit from widened trade preferences on the extensive margin. Using a quadruple-difference framework suggested by this theory, we have then found that the extensive-margin effect of LDC program reforms has indeed been limited to cases with previous ‘export experience’.

Our results suggest that tariff cuts have not been sufficient to achieve substantial gains in export diversification among LDCs. Further progress may depend on the success of complementary initiatives that bring local firms in untapped sectors closer to the margin of exporting, allowing trade preferences to bring them ‘over the top’. In fact, such ‘aid for trade’ programs – including direct support to trade infrastructure and production – have also long been an important part of the favoured policy approach to fostering export-led growth. The continued prevalence of zero trade flows in LDCs’ export portfolios suggest that this approach may still have a long road to travel.

3-A Appendix

Table 3-A.1: Sample exporters and treatment status by importer

	Aus	Can	EU	Jap	NZ	Nor	Swi	Tur	US LDC	US AGOA
Afghanistan	X	X	X	X	X		X	X	X	
Albania										
Algeria										
Angola	X	X	X	X	X		X	X	X	X
Armenia										
Azerbaijan										
Bangladesh	X	X	X	X	X		X	X	X	
Belarus										
Belize										
Benin	X	X	X	X	X		X	X	X	X
Bhutan	X	X	X	X	X		X	X	X	
Bolivia										
Bosnia-Herzegovina										
Burkina Faso	X	X	X	X	X		X	X	X	X
Burundi	X	X	X	X	X		X	X	X	X
Cambodia	X	X	X	X	X		X	X	X	
Cameroon						X				X
Cape Verde	X	X	X	X	X		X	X	X	X
Central Afr. Rep.	X	X	X	X	X		X	X	X	X
Chad	X	X	X	X	X		X	X	X	X
Comoros	X	X	X	X	X		X	X	X	X
Cuba										
Dem. Rep. Congo	X	X	X	X	X		X	X	X	X
Djibouti	X	X	X	X	X		X	X	X	X
Dominica										
Ecuador										
El Salvador										
Eq. Guinea	X	X	X	X	X		X	X	X	
Eritrea	X	X	X	X	X		X	X		
Ethiopia	X	X	X	X	X		X	X	X	X
Fiji										
FYR Macedonia										
Gabon										X
Gambia	X	X	X	X	X		X	X	X	X
Georgia										

Table 3-A.1: Sample exporters and treatment status by importer (continued)

	Aus	Can	EU	Jap	NZ	Nor	Swi	Tur	US LDC	US AGOA
Ghana						X				X
Grenada										
Guatemala										
Guinea	X	X	X	X	X		X	X	X	X
Guinea-Bissau	X	X	X	X	X		X	X	X	X
Guyana										
Haiti	X	X	X	X	X		X	X	X	
Honduras										
Iraq										
Ivory Coast						X				X
Jamaica										
Jordan										
Kazakhstan										
Kenya						X				X
Kiribati	X	X	X	X	X		X	X	X	
Kyrgyzstan						X				
Laos	X	X	X	X	X		X	X		
Lebanon										
Liberia	X	X	X	X	X		X	X	X	X
Madagascar	X	X	X	X	X		X	X	X	X
Malawi	X	X	X	X	X		X	X	X	X
Maldives	X	X	X	X	X		X	X		
Mali	X	X	X	X	X		X	X	X	X
Marshall Islands										
Mauritania	X	X	X	X	X		X	X	X	X
Mauritius										X
Micronesia										
Moldova						X				
Mongolia						X				
Mozambique	X	X	X	X	X		X	X	X	X
Myanmar	X			X	X					
Nepal	X	X	X	X	X		X	X	X	
Nicaragua						X				
Niger	X	X	X	X	X		X	X	X	X
Nigeria										X

Table 3-A.1: Sample exporters and treatment status by importer (continued)

	Aus	Can	EU	Jap	NZ	Nor	Swi	Tur	US	US
									LDC	AGOA
North Korea						X				
Panama										
Papua N. G.						X				
Paraguay										
Rep. of Congo						X				X
Rwanda	X	X	X	X	X		X	X	X	X
St. Vincent & Gr.										
Samoa	X	X	X	X	X		X	X	X	
S. T. & Principe	X	X	X	X	X		X	X	X	X
Senegal	X	X	X	X	X		X	X	X	X
Serbia & Montenegro										
Seychelles										X
Sierra Leone	X	X	X	X	X		X	X	X	X
Solomon Islands	X	X	X	X	X		X	X	X	
Somalia	X	X	X	X	X		X	X	X	
Sudan	X	X	X	X	X		X	X		
Suriname										
Syria										
Tajikistan						X				
Tanzania	X	X	X	X	X		X	X	X	X
Togo	X	X	X	X	X		X	X	X	X
Tonga										
Turkmenistan										
Tuvalu	X	X	X	X	X		X	X	X	
Uganda	X	X	X	X	X		X	X	X	X
Uzbekistan						X				
Vanuatu	X	X	X	X	X		X	X	X	
Yemen	X	X	X	X	X		X	X	X	
Zambia	X	X	X	X	X		X	X	X	X
Zimbabwe						X				

This table displays eligibility by exporter for the reforms of LDC trade preference programs that are studied in this paper, for all 100 exporters in the sample. 'US LDC' refers to the US expansion of its GSP-LDC program, while 'US AGOA' refers to the African Growth and Opportunity Act.

Bibliography

[1] Aaronson, D., Lange, F., and Mazumder, B. (2014). Fertility transitions along the extensive and intensive margins. *The American Economic Review*, 104(11):3701–3724.

[2] 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.

[3] Albornoz, F., Pardo, H. F. C., Corcos, G., and Ornelas, E. (2012). Sequential exporting. *Journal of International Economics*, 88(1):17–31.

[4] Allen, T. and Atkin, D. (2016). Volatility and the Gains from Trade. *NBER working paper*, 22276.

[5] Aragón, F. M., Oteiza, F., and Rud, J. P. (2020). Climate Change and Agriculture: Subsistence Farmers' Response to Extreme Heat. *American Economic Journal: Economic Policy*.

[6] Araujo, L., Mion, G., and Ornelas, E. (2016). Institutions and export dynamics. *Journal of International Economics*, 98:2–20.

[7] Arkolakis, C., Costinot, A., and Rodríguez-Clare, A. (2012). New Trade Models, Same Old Gains? *American Economic Review*, 102(1):94–130.

[8] 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.

[9] Ault, T. R., Mankin, J. S., Cook, B. I., and Smerdon, J. E. (2016). Relative Impacts of Mitigation, Temperature, and Precipitation on 21st-Century Megadrought Risk in the American Southwest. *Science Advances*, 2(10):e1600873.

[10] Baird, S., Chirwa, E., McIntosh, C., and Özler, B. (2010). The short-term impacts of a schooling conditional cash transfer program on the sexual behavior of young women. *Health economics*, 19(S1):55–68.

[11] Barber, S. L. (2009). Mexico's conditional cash transfer programme increases cesarean section rates among the rural poor. *European journal of public health*, 20(4):383–388.

[12] Barham, T. (2011). A healthier start: the effect of conditional cash transfers on neonatal and infant mortality in rural mexico. *Journal of Development Economics*, 94(1):74–85.

[13] Barham, T., Macours, K., and Maluccio, J. A. (2018). Experimental evidence of exposure to a conditional cash transfer during early teenage years: Young women's fertility and labor market outcomes.

[14] Becker, G. S. (1960). An economic analysis of fertility. In *Demographic and economic change in developed countries*, pages 209–240. Columbia University Press.

[15] Becker, G. S. and Lewis, H. G. (1973). On the interaction between the quantity and quality of children. *Journal of political Economy*, 81(2, Part 2):S279–S288.

[16] Becker, G. S. and Tomes, N. (1976). Child endowments and the quantity and quality of children. *Journal of political Economy*, 84(4, Part 2):S143–S162.

[17] Behrman, J. R., Parker, S. W., and Todd, P. E. (2009). Schooling impacts of conditional cash transfers on young children: Evidence from mexico. *Economic development and cultural change*, 57(3):439–477.

[18] 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.

[19] Brenton, P. (2003). Integrating the least developed countries into the world trading system: The current impact of eu preferences under everything but arms. Working paper, World Bank Policy Research.

[20] Burgess, R., Deschenes, O., Donaldson, D., and Greenstone, M. (2017). Weather, Climate Change and Death in India. *Mimeo*.

[21] Burke, M., Dykema, J., Lobell, D. B., Miguel, E., and Satyanath, S. (2015). Incorporating Climate Uncertainty into Estimates of Climate Change Impacts. *Review of Economics and Statistics*, 97(2):461–471.

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

[Cadot and Iacovone] Cadot, O. and Iacovone, L. Martha denisse pierola and ferdinand rauch (2013). "success and failure of african exporters". *Journal of Development Economics*, 101:284–296.

[24] Cattaneo, C. and Peri, G. (2016). The Migration Response to Increasing Temperatures. *Journal of Development Economics*, 122:127–146.

[25] Chaijaroen, P. (2019). Long-Lasting Income Shocks and Adaptations: Evidence from Coral Bleaching in Indonesia. *Journal of Development Economics*, 136:119–136.

[26] 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.

[27] Chaney, T. (2014). The network structure of international trade. *American Economic Review*, 104(11):3600–3634.

[28] 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.

[29] Coady, D. P. and Parker, S. W. (2004). Cost-effectiveness analysis of demand-and supply-side education interventions: the case of progresa in mexico. *Review of Development Economics*, 8(3):440–451.

[30] Collier, P. and Venables, A. J. (2007). Rethinking trade preferences: How africa can diversify its exports. *World Economy*, 30(8):1326–1345.

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

[32] 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.

[33] Da Mata, D. and Resende, G. (2020). Changing the Climate for Banking: The Economic Effects of Credit in a Climate-Vulnerable Area. *Journal of Development Economics*, 146:102459.

[34] de Melo, J. and Portugal-Perez, A. (2013). Preferential market access design: Evidence and lessons from african apparel exports to the united states and european union. *World Bank Economic Review*, 28(1):74–98.

[35] Deryugina, T. and Hsiang, S. (2017). The Marginal Product of Climate. *NBER working paper*, 24072.

[36] 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.

[37] 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.

[38] Diffenbaugh, N. S., Singh, D., Mankin, J. S., Horton, D. E., Swain, D. L., Touma, D., Charland, A., Liu, Y., Haugen, M., Tsiang, M., et al. (2017). Quantifying the Influence of Global Warming on Unprecedented Extreme Climate Events. *Proceedings of the National Academy of Sciences*, 114(19):4881–4886.

[39] Diffenbaugh, N. S., Swain, D. L., and Touma, D. (2015). Anthropogenic Warming has Increased Drought Risk in California. *Proceedings of the National Academy of Sciences*, 112(13):3931–3936.

[40] 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.

[41] Doepke, M. (2005). Child mortality and fertility decline: Does the barro-becker model fit the facts? *Journal of population Economics*, 18(2):337–366.

[42] Duflo, E. (2012). Women empowerment and economic development. *Journal of Economic Literature*, 50(4):1051–1079.

[43] 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.

[44] Fernald, L. C., Gertler, P. J., and Neufeld, L. M. (2008). Role of cash in conditional cash transfer programmes for child health, growth, and development: an analysis of mexico's oportunidades. *The Lancet*, 371(9615):828–837.

[45] Fernald, L. C., Gertler, P. J., and Neufeld, L. M. (2009). 10-year effect of oportunidades, mexico's conditional cash transfer programme, on child growth, cognition, language, and behaviour: a longitudinal follow-up study. *The Lancet*, 374(9706):1997–2005.

[46] 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.

[47] Fishman, R. (2016). More Uneven Distributions Overturn Benefits of Higher Precipitation for Crop Yields. *Environmental Research Letters*, 11(2):024004.

[48] Fishman, R. (2018). Groundwater Depletion Limits the Scope for Adaptation to Increased Rainfall Variability in India. *Climatic Change*, 147(1-2):195–209.

[49] Frazer, G. and Biesebroeck, J. V. (2010). Trade growth under the african growth and opportunity act. *Review of Economics and Statistics*, 92(1):128–144.

[50] Galor, O. (2012). The demographic transition: Causes and consequences. cliometrica, 6, 1-28.

[51] Galor, O. and Weil, D. N. (1999). From malthusian stagnation to modern growth. *The American Economic Review*, 89(2):150–154.

[52] Garg, T., Jagnani, M., and Taraz, V. (2020). Temperature and Human Capital in India. *Journal of the Association of Environmental and Resource Economists*.

[53] Gertler, P. (2004). Do conditional cash transfers improve child health? evidence from progresa's control randomized experiment. *American economic review*, 94(2):336–341.

[54] Gil-Pareja, S., Llorca-Vivero, R., and Martínez-Serrano, J. A. (2014). Do nonreciprocal preferential trade agreements increase beneficiaries' exports? *Journal of Development Economics*, 107:291–304.

[55] Gradeva, K. and Martínez-Zarzoso, I. (2016). Are trade preferences more effective than aid in supporting exports? evidence from the 'everything but arms' preference scheme. *World Economy*, 39(8):1146–1171.

[56] Guiteras, R. (2009). The Impact of Climate Change on Indian Agriculture. *Mimeo.*

[57] Hakobyan, S. (2020). Gsp expiration and declining exports from developing countries. *Canadian Journal of Economics*, 53(3):1132–1161.

[58] 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.

[59] Helpman, E., Melitz, M., and Rubinstein, Y. (2008). Estimating trade flows: Trading partners and trading volumes. *Quarterly Journal of Economics*, 123(2):441–487.

[60] Henderson, J. V., Storeygard, A., and Deichmann, U. (2017). Has Climate Change Driven Urbanization in Africa? *Journal of Development Economics*, 124:60–82.

[61] Hsiang, S. (2016). Climate Econometrics. *Annual Review of Resource Economics*, 8:43–75.

[62] Jayachandran, S. and Lleras-Muney, A. (2009). Life expectancy and human capital investments: Evidence from maternal mortality declines. *The Quarterly Journal of Economics*, 124(1):349–397.

[63] 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.

[64] Jones, L. E., Schoonbroodt, A., and Tertilt, M. (2008). Fertility theories: can they explain the negative fertility-income relationship? Technical report, National Bureau of Economic Research.

[65] Kala, N. (2017). Learning, Adaptation, and Climate Uncertainty: Evidence from Indian Agriculture. *MIT Center for Energy and Environmental Policy Research Working Paper*, 23.

[66] 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.

[67] Kleemans, M. and Magruder, J. (2018). Labour Market Responses to Immigration: Evidence from Internal Migration Driven by Weather Shocks. *The Economic Journal*, 128(613):2032–2065.

[68] 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.

[69] Kugler, A. D. and Rojas, I. (2018). Do ccts improve employment and earnings in the very long-term? evidence from mexico. Technical report, National Bureau of Economic Research.

[70] Laszlo, S., Farhan Majid, M., and Renée, L. (2019). Conditional cash transfers, women's empowerment and reproductive choices.

[71] 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.

[72] 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.

[73] Macours, K., Premand, P., and Vakis, R. (2012). Transfers, Diversification and Household Risk Strategies: Experimental Evidence with Lessons for Climate Change Adaptation. *CEPR Discussion Paper*, 8940.

[74] 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.

[75] Ministry of Agriculture and Farmers Welfare (2017). *Agricultural Statistics at a Glance 2016*. New Delhi: Government of India.

[76] Moav, O. (2004). Cheap children and the persistence of poverty. *The economic journal*, 115(500):88–110.

[77] Morales, E., Sheu, G., and Zahler, A. (2019). Extended gravity. *Review of Economic Studies*, 86(6):2668–2712.

[78] 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.

[Olarreaga] Olarreaga, M. a. and Çaglar Özden (2005). "agoa and apparel: Who captures the tariff rent in the presence of preferential market access?". *World Economy*, 28(1):63–77.

[80] Ornelas, E. (2016). Special and differential treatment for developing countries. In Bagwell, K. and Staiger, R., editors, *Handbook of Commercial Policy*, pages 369–432. Vol. 1 Part B.

[81] Ornelas, E. and Ritel, M. (2020). The not-so-generalized effects of the generalized system of preferences. *World Economy*, 43(7):1809–1840.

[82] Parker, S. W. and Todd, P. E. (2017). Conditional cash transfers: The case of progresa/oportunidades. *Journal of Economic Literature*, 55(3):866–915.

[83] Parker, S. W. and Vogl, T. (2018). Do conditional cash transfers improve economic outcomes in the next generation? evidence from mexico. Technical report, National Bureau of Economic Research.

[84] Persson, M. and Wilhelmsson, F. (2016). Eu trade preferences and export diversification. *World Economy*, 39(1):16–53.

[85] Potter, J. E., Schmertmann, C. P., and Cavenaghi, S. M. (2002). Fertility and development: Evidence from brazil. *Demography*, 39(4):739–761.

[86] Rotunno, L., Vézina, P.-L., and Wang, Z. (2013). The rise and fall of (chinese) african apparel exports. *Journal of Development Economics*, 105:152–163.

[87] Rubio-Codina, M. (2010). Intra-household time allocation in rural mexico: Evidence

from a randomized experiment. In *Child labor and the transition between school and work*, pages 219–257. Emerald Group Publishing Limited.

[88] 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.

[89] 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.

[90] Schlenker, W. and Lobell, D. B. (2010). Robust Negative Impacts of Climate Change on African Agriculture. *Environmental Research Letters*, 5(1):014010.

[91] 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.

[92] Schmeiser, K. N. (2012). Learning to export: Export growth and the destination decision of firms. *Journal of International Economics*, 87(1):89–97.

[93] Schultz, T. P. (2004). School subsidies for the poor: evaluating the mexican progresa poverty program. *Journal of development Economics*, 74(1):199–250.

[94] Shrader, J. (2017). Expectations and Adaptation to Environmental Risks. *Mimeo.*

[95] Skoufias, E., Parker, S. W., Behrman, J. R., and Pessino, C. (2001). Conditional cash transfers and their impact on child work and schooling: Evidence from the progresa program in mexico [with comments]. *Economia*, 2(1):45–96.

[96] Stecklov, G., Winters, P., Todd, J., and Regalia, F. (2007). Unintended effects of poverty programmes on childbearing in less developed countries: experimental evidence from latin america. *Population Studies*, 61(2):125–140.

[97] 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.

[98] Taraz, V. (2018). Can Farmers Adapt to Higher Temperatures? Evidence from India. *World Development*, 112:205–219.

[99] Thelle, M. H., Jeppesen, T., Gjødesen-Lund, C., and Biesebroeck, J. V. (2015). *Assessment of Economic Benefits Generated by the EU Trade Regimes Towards Developing Countries.* Directorate-General for International Cooperation and Development, EU Development Policy and International Cooperation, Policy and Coherence, Economic Analysis Team.

[100] Thirumalai, K., DiNezio, P. N., Okumura, Y., and Deser, C. (2017). Extreme Temperatures in Southeast Asia Caused by El Niño and Worsened by Global Warming. *Nature Communications*, 8:15531.

[101] 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.

[102] Todd, P. E. and Wolpin, K. I. (2006). Assessing the impact of a school subsidy program in mexico: Using a social experiment to validate a dynamic behavioral model of child schooling and fertility. *American economic review*, 96(5):1384–1417.

[103] UNEP (2019). *Emissions Gap Report 2019*. Nairobi: United Nations Environment Programme.

[104] 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.

[105] Urquieta, J., Angeles, G., Mroz, T., Lamadrid-Figueroa, H., and Hernandez, B. (2009). Impact of oportunidades on skilled attendance at delivery in rural areas. *Economic Development and Cultural Change*, 57(3):539–558.

[106] 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.

[107] Vogl, T. S. (2015). Differential fertility, human capital, and development. *The Review of Economic Studies*, 83(1):365–401.

[108] Yoon, J.-H., Wang, S. S., Gillies, R. R., Kravitz, B., Hipps, L., and Rasch, P. J. (2015). Increasing Water Cycle Extremes in California and in Relation to ENSO Cycle Under Global Warming. *Nature Communications*, 6:8657.