Adaptation to Climate Change: Historical Evidence from the Indian Monsoon

Vis Taraz*

July 27, 2016

Abstract

Estimating the potential impacts of climate change requires understanding the ability of agents to adapt to changes in their climate. This paper uses panel data from India spanning from 1956 to 1999 to investigate the ability of farmers to adapt. To identify adaptation, I exploit persistent, multidecadal monsoon regimes, during which droughts or floods are more common. These regimes generate medium-run variation in average rainfall, and there is spatial variation in the timing of the regimes. Using a fixed effects strategy, I test whether farmers have adapted to the medium-run rainfall variation induced by the monsoon regimes. I find evidence that farmers adjust their irrigation investments and their crop portfolios in response to the medium-run rainfall variation. However, adaptation only recovers a small fraction of the profits farmers have lost due to adverse climate variation.

1 Introduction

Climate scientists broadly agree that the global climate is changing and that these changes will accelerate in coming decades (Christensen and Hewitson, 2007). However, estimates of the economic

*Smith College, Department of Economics, Pierce Hall, 21 West Street, Northampton, MA 01063-6317 (Email: vtaraz@smith.edu). I would like to thank Christopher Udry, Mark Rosenzweig, Dean Karlan and Mushfiq Mobarak for their guidance and support. I thank Victoria Beckley for assistance in creating maps. I thank Juan Pablo Rud for sharing data. I thank Sophie D'Arcy and Emily Zhou for excellent research assistance. David Atkin, Reena Badiani, David Childers, Simon Halliday, Rachel Heath, Richard Hornbeck, Namrata Kala, Dan Keniston, Melanie Morten, Elizabeth Savoca, Susan Sayre and participants of the Yale Development Lunch, the NEUDC 2012, and the Southern Economic Association Conference 2013 provided helpful comments and suggestions. All remaining errors are my own.

impacts of climate change vary widely, in large part due to uncertainty about adaptation (Mendelsohn et al., 1994; Adams et al., 1998; Schlenker et al., 2005; Deschênes et al., 2007; Schlenker and Roberts, 2009; Tol, 2014). Rapid adaptation may curb economic damages, but slower adaptation will likely magnify them. Understanding adaptation is particularly crucial in developing countries and in the agricultural sector, as both are especially vulnerable to climate change (Parry, 2007).

Recent scholarship has typically estimated climate change damages using year-to-year weather variation to compare economic outcomes under hotter versus cooler temperatures. This climate–economy relationship is then extrapolated to future climate change to estimate impacts (Deschênes et al., 2007; Schlenker and Roberts, 2009; Guiteras, 2009; Dell et al., 2012; Burgess et al., 2014).[1] Since these calculations rely on annual weather variability, they do not account for possible adaptations that agents may undertake in response to sustained climate change. Therefore, to assess the accuracy of these estimates, it is vital to predict the likely extent of future adaptation.

In this paper, I exploit historical rainfall variation in India to estimate adaptation. Rather than analyzing year-to-year weather deviations, I focus on climate fluctuations that last several decades. The Indian monsoon undergoes multidecadal phases during which droughts or floods are more common. These monsoon phases induce persistent deviations in rainfall from decade to decade. I test whether farmers adapt their irrigation investments and crop portfolios in response to these persistent rainfall deviations.[2]

Figure 1 shows a moving average of India's summer rainfall, highlighting the monsoon phases. These phases induce persistent rainfall deviations and, hence, lagged rainfall provides information about future rainfall. Therefore, forward-looking farmers should adjust their agricultural decisions in response to recent weather.

I test for adaptation by analyzing whether agricultural decisions respond to lagged weather, looking specifically at irrigation investments and crop choice. I exploit the fact that the return to irrigation investment varies across wet versus dry growing seasons and that, similarly, the relative

[1] Another methodology uses cross-sectional climate variation to link climate and the economy, but this work suffers from potential omitted variable bias (Mendelsohn et al., 1994; Schlenker et al., 2005; Sanghi and Mendelsohn, 2008).

[2] The monsoon regimes don't cause variation in temperature, so I do not analyze adaptation to temperature changes.

yields of different crops vary across wet versus dry growing seasons. My empirical strategy is to regress irrigation assets and crop portfolios on rainfall from the past decade, while controlling for current rainfall, wealth, household fixed effects, and year fixed effects. Regional variation in the timing of the decadal rainfall regimes, displayed in Figure 2, allows me to include year fixed effects in my regressions and, hence, I can separate adaptation to rainfall from unrelated temporal changes in irrigation and crop choice.

Analyzing two agricultural data sets, I find evidence of both irrigation adaptation and crop adaptation. Each additional dry year in the past decade increases the probability that a farmer will invest in irrigation by 1.2 percentage points, relative to a baseline 5% probability of investing.[3] Each additional dry year in the past decade reduces the average daily water need of a farmer's monsoon season crop portfolio by 0.2 mm/day, relative to an average water need of 8 mm/day.[4] In addition to testing for the presence of adaptation, I also measure the extent to which adaptation prevents profit losses. I find that farmers are able to recover only a limited amount of their lost profits by adapting. Specifically, I estimate that in the face of sustained adverse weather conditions adaptation recovers, at most, 19% of lost profits and, more likely, only 9%.

This paper contributes to a rapidly growing literature on climate change adaptation.[5] Researchers have used a variety of techniques to identify the magnitude and efficacy of adaptation, including the Ricardian method or hedonic valuation method (Mendelsohn et al., 1994; Fleischer et al., 2008; Seo et al., 2010; Kurukulasuriya et al., 2011; da Cunha et al., 2014) and variants of the Ricardian method that incorporate panel data (Luis and Orlando, 2015) and structural agro-economic models (Kurukulasuriya and Mendelsohn, 2008). Researchers have also analyzed adaptation by looking at long-run responses to one-time environmental shocks (Hornbeck, 2012; Deryugina, 2013; Hornbeck and Naidu, 2015), applying instrumental variables approaches that address the endogeneity of adaptation (Di Falco and Veronesi, 2013, 2014), using economic models that integrate biophysical modeling (Finger et al., 2010), employing multinomial logit choice

[3] I define a dry year to be a year in which rainfall is below the 20th percentile of the rainfall distribution for a particular location.

[4] The crops with lower water needs have lower expected yields, which is why farmers do not plant them exclusively.

[5] Dell et al. (2014) present a helpful synthesis of this literature.

models (Seo and Mendelsohn, 2008; Wang et al., 2010), testing whether new technologies have changed weather impacts over time (Barreca et al., 2015), analyzing differential weather impacts by the long-run frequency of the event (Deschênes and Greenstone, 2011; Hsiang and Narita, 2012), estimating correlations between farmer behavior and their perceptions of changes in climate (Bryan et al., 2009), and, lastly, using a "long-difference" approach that compares short-run weather impacts with long-run impacts (Dell et al., 2012; Burke and Emerick, 2015).

My paper contributes to the adaptation literature in multiple ways. First, my study is unique because I use a household data set that spans several decades. The existing literature on adaptation uses either administrative data (Dell et al., 2012), cross-sectional household data (Bryan et al., 2009; Mukherjee and Schwabe, 2015), or a short panel of household data that spans less than 10 years (Luis and Orlando, 2015). My paper is also unique because I estimate adaptation to large-scale, cyclical, decadal variation in climate that exhibits both spatial and temporal variation. The bulk of the existing literature on adaptation exploits either cross-sectional (spatial) variation in climate, a one-time shock to climate, or perceived changes in climate that are measured at a single point in time.

My unique data set and source of climate variation allow me to make a methodological contribution to the literature. Specifically, I can estimate how farmers adapt to medium-run (10-20 year) changes in climate that are occurring over the span of my data set, while controlling for unobserved heterogeneity. Put differently, my data allow me to look at how the behavior of a household changes across several decades, in response to time-varying changes in climate. My estimates of decade-to-decade adaptation are an important complement to the long-run adaptation estimates that are generated by methods that rely on purely spatial climate variation.[6] Conversely, my estimates are also a complement to studies that estimate how farmers respond to recent perceived changes in climate. Typically, these studies use cross-sectional household data and focus only on behavior and climate perceptions from the past 10-20 years. My multidecadal household

[6]When adaptation is estimated using cross-sectional climate variation, the relationship between farmer behavior and climate is based on the long-run climate of each location. As a result, these estimates are best thought of as estimates of how farmers will adapt to climate change in a long-run, or steady state, setting.

panel, on the other hand, allows me to control for unobserved farmer heterogeneity and to analyze adaptation over several decades. In addition, I have built a model that allows me to disentangle the effects of wealth and expectations. This is a methodological contribution because it allows me to directly test whether farmers are updating their beliefs about future rainfall in response to past rainfall, even in the absence of explicit data on farmers' perceptions about climate change.

There are several important limitations of my study to acknowledge. First, this study only analyzes irrigation and crop choice. Data limitations do not permit me to study other potential adaptations, such as adjusting fertilizer and agricultural inputs (Duflo et al., 2011), shifting sowing dates (Giné et al., 2009), purchasing crop insurance (Di Falco et al., 2014), switching out of agriculture (Rose, 2001), or migrating (Viswanathan and Kavi Kumar, 2015). Second, since the monsoon regimes affect only precipitation, I do not analyze adaptation to temperature changes. Third, since my household data set spans several decades, there is substantial, non-random attrition, which causes my analyzed sample to include households that are, on average, wealthier than a representative sample would be.[7] Fourth, there are potential threats to the exclusion restriction for my instrumental variables strategy, which I discuss in greater detail in Section 5. Fifth, due to data limitations, I am not fully able to rule out the possibility that depletion of water supplies or confounding factors, such as changes in agricultural technology or policies, are driving my results.[8]

The paper is organized as follows. Section 2 describes the monsoon phases in greater detail. Section 3 presents a model of climate, irrigation, and crop choice. Section 4 describes the data, and Section 5 proposes the empirical strategy. Section 6 presents the main results. In Section 7, I discuss several robustness tests that I perform in a separate, supplementary file. Section 8 calculates the fraction of lost profits farmers recovered by adapting. Section 9 concludes.

[7] Appendix A discusses the attrition in more detail and its implications for my study.

[8] Appendix C discusses these issues in greater detail.

2 Background on Interdecadal Rainfall Variability

Indian agriculture depends heavily on the summer monsoon, which occurs during June, July, August, and September (Krishna Kumar et al., 2004). Because India's climate is semi-arid, wetter monsoons increase agricultural output, and drier monsoons decrease it (Das, 1995; Jayachandran, 2006). Monsoon rainfall exhibits high interannual variability, as shown in Figure 3. The monsoon also undergoes interdecadal variability, in the form of wet and dry phases that typically each last for about three decades (Pant and Kumar, 1997). Meteorologists refer to these as meditional and zonal regimes, respectively. Figure 3 shades the wet regimes gray; Figure 1 smoothes annual rainfall with a moving average filter, to further highlight the regimes.[9]

The monsoon regimes cause average rainfall to vary more from decade to decade than it would if rainfall was independent and identically distributed (i.i.d.).[10] This persistent decadal variation means that lagged rainfall has predictive value for future rainfall. If rainfall were i.i.d., then lagged rainfall would not have this predictive element. Rational farmers should notice these persistent rainfall variations and update their future rainfall expectations in response. This updating could occur even if farmers were not aware of the existence of the monsoon regimes, per se. On the other hand, if rainfall were i.i.d., lagged rainfall would have no predictive value, and it would be irrational for farmers to update their rainfall expectations in response to it. The statistical significance of the decadal variations allows me to interpret a farmer's response to lagged rainfall as evidence of rational adaptation, rather than an indicator of irrational behavior.[11]

The monsoon regimes are not geographically homogeneous. There is significant spatial vari-

[9]Meteorologists widely agree upon the existence of the monsoon regimes (Subbaramayya and Naidu, 1992; Kripalani and Kulkarni, 1997; Pant and Kumar, 1997; Pant, 2003; Varikoden and Babu, 2014). The precise mechanisms that generate the regimes are not well understood, in part due to a lack of good quality data for a sufficiently long period. One theory is that an atmospheric-oceanic feedback mechanism induces the regimes (Wang, 2006).

[10]Mooley and Parthasarathy (1984) and Kripalani and Kulkarni (1997) perform statistical analysis demonstrating that the monsoon regimes are statistically significant. That is, they demonstrate that the interdecadal rainfall variability is greater than what we would expect under an i.i.d process. In Section B of the supplementary file, I describe their analysis in greater detail and also run an additional test that further verifies the monsoon's non-stationarity.

[11]I have not been able to find descriptive survey data regarding the question of whether farmers in India are aware of the monsoon regimes or the decadal rainfall variation that they induce. However, Palanisami et al. (2014) note several surveys that find that, more recently, farmers have noticed changes in temperature and rainfall that have been induced by anthropogenic climate change, which are comparable in magnitude to the changes that I analyze in my study.

ation in the length and timing of the regimes (Subbaramayya and Naidu, 1992). In particular, rainfall in the southern peninsula and the easternmost region tends to be out of phase with the rest of the country (Wang, 2006). Figure 2 displays smoothed rainfall graphs for India's five meteorological regions, highlighting the spatial variation. Providing more detail, Figures 4 and 5 map district rainfall from the previous decade, for the three survey years of the REDS data set and at four decade intervals for the WB data set.[12] The spatial variation in recent rainfall allows me to include year fixed effects in my regressions and, hence, distinguish rainfall adaptation from time trends in irrigation and crop choice.

3 Theoretical Framework

I now derive tests for farmer adaptation. Sections 3.1 and 3.2 outline the climate and agricultural models, respectively. Section 3.3 shows the farmer's maximization problem, and Section 3.4 presents the adaptation tests.

3.1 Climate Model

I model the monsoon regimes as a hidden Markov process. Let s_t indicate the monsoon regime in year t, with $s_t = 0$ denoting a dry regime and $s_t = 1$ denoting a wet regime. Year t rainfall can be written as:

$$r_t = \theta_0 + \delta s_t + u_t, \tag{1}$$

where θ_0 is the average rainfall during a dry regime, $\theta_0 + \delta$ is the average wet regime rainfall, and u_t represents year-to-year rainfall variability. The monsoon regimes are persistent but not permanent, and they switch according to a Markov process. During a dry regime, the probability of switching to a wet regime during the next period is p_0. During a wet regime, the probability of

[12] I choose rainfall from the previous decade as a rough measure of the current monsoon regime (Kripalani and Kulkarni, 1997).

switching to a dry regime is p_1. Each year, farmers observe r_t and use this information to update their belief about the current regime state, which they do not observe. A farmer's belief about the current regime state determines his expectation of the next period's rainfall.

3.2 Agricultural Model

In my model, each farmer lives for two periods. In each period t, the farmer allocates his wealth, w_t, between an irrigation asset, i_t, and another agricultural asset, a_t, such that $a_t + i_t = w_t$.[13] The farmer also chooses a crop portfolio each period. The farmer has one unit of land, which he divides between a drought-tolerant crop and a crop that is relatively more sensitive to drought.[14] Let ρ_t be the area planted with the drought-tolerant crop, and let $1 - \rho_t$ be the drought-sensitive crop area.

Profits are determined by the asset mix, the crop portfolio, and rainfall r_t. I assume a quadratic profit function of the form:

$$\begin{aligned} \pi_t = &\beta_a a_t + \beta_i i_t + \beta_\rho \rho_t + \frac{1}{2}\delta_{aa} a_t^2 + \frac{1}{2}\delta_{ii} i_t^2 + \frac{1}{2}\delta_{\rho\rho}\rho_t^2 + \delta_{\rho i}\rho_t i_t + \delta_{ir} i_t r_t + \\ &+ \delta_{\rho r}\rho_t r_t + \delta_r r_t + \epsilon_t \end{aligned} \tag{2}$$

where π_t is profits per acre and ϵ_t is a mean zero productivity shock.[15] To establish my adaptation tests, I assume that:

1. Profits are increasing in rainfall ($\delta_r > 0$). This assumption is consistent with earlier work on India (Jayachandran, 2006; Cole et al., 2012), and I verify it in Section 6.1.

2. The return to irrigation is higher during periods of low rainfall ($\delta_{ir} < 0$). This assumption, while intuitive, is also verified in Section 6.1.

3. The drought-tolerant crop is less profitable, on average, than the drought-sensitive crop ($\beta_\rho < 0$). This assumption is necessary to ensure that farmers do not plant all their land

[13] Examples of other agricultural assets include tractors, tillers, ploughs, threshers, and livestock. I abstract away from the possibility of credit markets and non-agricultural assets.

[14] I will test my model with data that includes a large number of crops, with a range of different water needs, but for clarity in my theoretical model, I assume there are only two different crops.

[15] I assume this reduced form expression for profits for tractability purposes.

with the drought-tolerant crop.

4. Low rainfall reduces the profitability of the drought-tolerant crop less than it reduces the profitability of the drought-sensitive crop ($\delta_{\rho r} < 0$). This assumption comprises my definition of the drought-tolerant crop.

3.3 Maximization Problem

Each farmer maximizes:

$$u(c_1) + \beta E_1[u(c_2)] \tag{3}$$

subject to:

$$c_1 = w_1 + \pi_1 - w_2 \text{ and } c_2 = w_2 + \pi_2, \tag{4}$$

where $0 < \beta < 1$. For tractability, I assume constant absolute risk aversion utility of the form:

$$u(c_t) = -e^{-\eta c_t}. \tag{5}$$

The timing of the model is as follows. To begin, the farmer chooses his first-period assets and crop portfolio, based on initial wealth and rainfall expectations. Next, first-period rainfall occurs and first-period profits are determined. With these profits in hand, the farmer chooses how much to consume in the first period and how much wealth to bring into the second period. The farmer also chooses his second-period asset mix and crop portfolio. Lastly, second-period rainfall occurs, and second-period profits are determined.

3.4 Tests for Adaptation

I now derive tests to determine whether farmers are updating their rainfall expectations in response to past rainfall and whether they are adapting their agricultural decisions accordingly. I lack data on farmer rainfall expectations, but the structure of my model allows me to test for adaptation, even without explicit data on expectations.

To clarify the analysis, I introduce the following notation. Let $\mu_1 = E_0(r_1)$ and $\mu_2 = E_1(r_2)$ denote rainfall expectations. Let w_2^* denote the optimal amount of wealth to bring into second-period wealth. Let i_2^* and ρ_2^* denote the optimal second-period irrigation and crop choice decisions. Note that i_2^* and ρ_2^* depend solely on μ_2 and w_2^*. Furthermore, w_2^* itself is a function of w_1, μ_1, r_1 and μ_2.

3.4.1 Tests for Irrigation Adaptation

To begin, note that the total derivative of second-period irrigation with respect to first period rainfall is:

$$\begin{aligned}\frac{di_2^*}{dr_1} &= \frac{\partial i_2^*}{\partial w_2}\frac{dw_2^*}{dr_1} + \frac{\partial i_2^*}{\partial \mu_2}\frac{d\mu_2}{dr_1} \\ &= \frac{\partial i_2^*}{\partial w_2}\left[\frac{\partial w_2^*}{\partial r_1} + \frac{\partial w_2^*}{\partial \mu_2}\frac{d\mu_2}{dr_1}\right] + \frac{\partial i_2^*}{\partial \mu_2}\frac{d\mu_2}{dr_1}.\end{aligned}$$

Rearranging terms, we get:

$$\frac{di_2^*}{dr_1} = \underbrace{\frac{\partial i_2^*}{\partial w_2}\frac{\partial w_2^*}{\partial r_1}}_{\text{wealth effect}} + \underbrace{\left[\frac{\partial i_2^*}{\partial w_2}\frac{\partial w_2^*}{\partial \mu_2} + \frac{\partial i_2^*}{\partial \mu_2}\right]\frac{d\mu_2}{dr_1}}_{\text{expectations effect}}.$$

I have written the response of second-period irrigation to first-period rainfall as the sum of a wealth effect and an expectations effect. In Section C of the supplementary file, I demonstrate that the signs of the partial derivatives in this expression are:

$$\frac{\partial i_2^*}{\partial w_2} > 0, \frac{\partial i_2^*}{\partial \mu_2} < 0, \frac{\partial w_2^*}{\partial r_1} > 0, \frac{\partial w_2^*}{\partial \mu_2} < 0.$$

Taken together, the signs of these partial derivatives imply that, for irrigation, the wealth effect is positive, and the expectations effect is negative.[16]

Having separated the influences of wealth and expectations, I present two tests for whether

[16] I have used a CARA utility function for tractability purposes. I am not able to prove the signs of the wealth effect and the expectations effect for a broader range of utility functions.

farmers are adapting their irrigation in response to expected rainfall.

Proposition 3.1 *If farmers increase their irrigation investment after low rainfall, this demonstrates adaptation:* $\frac{di_2^*}{dr_1} < 0$ *implies* $\frac{d\mu_2}{dr_1} > 0$

Proposition 3.2 *If, conditional on wealth, farmers increase their irrigation investment after low rainfall, this also demonstrates adaptation:* $\left.\frac{di_2^*}{dr_1}\right|_{w_2=constant} < 0$ *implies* $\frac{d\mu_2}{dr_1} > 0$

Proposition 3.1 is an unconditional test that does not require accounting for wealth. Proposition 3.1 is useful because it allows me to test for adaptation, even in data sets that lack information on wealth. This is relevant to this study because my two data sets differ in this regard. As Section 4 will explain, my household data set includes data on wealth, but my district data set does not. On the other hand, Proposition 3.2 is a conditional test that incorporates a measure of wealth. It is a more powerful test than Proposition 3.1. If farmers are adapting, but the size of the wealth effect dominates the expectation effect, then Proposition 3.2 will detect the presence of adaptation but Proposition 3.1 will not. Additionally, because Proposition 3.2 separates out the wealth and expectation effects, the empirical analog of Proposition 3.2 can more accurately estimate the magnitude of the expectation effect. Proposition 3.1, on the other hand, conflates the wealth and expectation effects and so, when estimated, it will understate the size of the expectation effect.

3.4.2 Test for Crop Adaptation

Lastly, I derive a test for crop adaptation. I take the derivative of the second-period drought-tolerant crop area with respect to first-period rainfall. Rearranging terms, I get:

$$\frac{d\rho_2^*}{dr_1} = \underbrace{\frac{\partial \rho_2^*}{\partial w_2}\frac{\partial w_2^*}{\partial r_1}}_{\text{wealth effect}} + \underbrace{\left[\frac{\partial \rho_2^*}{\partial w_2}\frac{\partial w_2^*}{\partial \mu_2} + \frac{\partial \rho_2^*}{\partial \mu_2}\right]}_{\text{expectations effect}}\frac{d\mu_2}{dr_1}$$

In Section C of the supplementary file, I demonstrate the following signs of the partial derivatives:

$$\frac{\partial \rho_2^*}{\partial w_2} < 0, \frac{\partial \rho_2^*}{\partial \mu_2} < 0, \frac{\partial w_2^*}{\partial r_1} > 0, \frac{\partial w_2^*}{\partial \mu_2} < 0.$$

Substituting in these partial derivatives, I find that, for crop choice, the wealth effect is negative and the sign of the expectation effect is ambiguous.[17] Therefore, it is not possible to test for crop adaptation without controlling for wealth. On the other hand, if I hold wealth constant, this removes the wealth effect *and* makes the sign of the expectations effect unambiguously negative. This generates the following test for adaptation:

Proposition 3.3 *If, conditional on wealth, farmers plant a greater area of drought-tolerant crops after low rainfall, this demonstrates adaptation to climate:* $\left.\frac{d\rho_2^*}{dr_1}\right|_{w_2=constant} < 0$*, then* $\frac{d\mu_2}{dr_1} > 0$

The necessity of controlling for wealth means that I can test for crop adaptation in my household data set but not in my district data set. Without a wealth control, a negative correlation between lagged rainfall and drought-tolerant crop areas could be occurring solely through a wealth channel and, hence, would not provide evidence of adaptation.

4 Data Sources and Summary Statistics

I test my model with two agricultural data sets: a household panel and a district panel. The household panel—the Rural Economic and Demographic Survey—was collected by the National Council of Applied Economic Research (NCAER).[18]

The data covers three rounds (1970/71, 1981/82, and 1998/99) and 259 villages across the 17 major states of India. In 1971, 4,527 households were surveyed. In 1982, the original villages were revisited, and 4,979 households were surveyed, of which roughly two-thirds were the same households from the 1971 round. The 1999 round covers 7,474 households in the same villages, including all the households from 1982, any households that split off from the original 1982 households, as well as a small random sample of new households. Figure 2 in Appendix A displays a

[17] I have used a CARA utility function for tractability purposes. I am not able to prove the signs of the wealth effect and the expectation effect for a broader range of utility functions.

[18] The data can be downloaded from http://adfdell.pstc.brown.edu/arisreds_data/.

map with the locations of the REDS villages. I restrict my analysis to households that either were surveyed in multiple rounds or split off from a previously surveyed household.[19] The REDS survey includes detailed data on irrigation, crop areas, assets, wealth, profits, and inherited assets.

The district panel—the India Agriculture and Climate Data Set—was compiled by a World Bank research group and covers 271 districts across 14 states for each year between 1956 and 1987 (Sanghi et al., 1998). Figure 1 in Appendix A displays the districts covered in the World Bank data set. The data set includes information on irrigated areas, crop areas, crop yields, and prices, but does not include information about assets, wealth, or profits.

Panel A in Table 1 presents the summary statistics for agricultural variables of both data sets. For the household data set, agricultural profits per acre are measured as crop receipts minus crop expenses, divided by the area of land cultivated. The World Bank data set lacks information on crop expenses. Instead, I use crop revenue per acre of land cultivated.[20]

For the household data, I define irrigation investment as a dummy variable that is equal to one if the household invested in irrigation during the recall period, which is defined as the 12 months prior to the survey interview. Investing in irrigation is defined as purchasing materials, hiring labor, or using family labor to construct new irrigation assets, purchase new irrigation assets, or improve existing irrigation assets.[21,22] The district data lacks direct information on irrigation investment, so I define irrigation investment as the log of the 1-year change in the area of irrigated land.[23] For the household data set, I measure wealth as the sum of the value of irrigation assets, farm equipment, livestock, non-farm assets, housing, durable goods, farm inventory, and financial assets minus debts.[24]

---

[19]Unfortunately, due to the long-time span between the survey rounds, the REDS data set suffers from non-negligible attrition. In Appendix A, I discuss this attrition and its implications for my study.

[20]Table 1 also includes information about inherited irrigated land, which is used in my empirical strategy and is discussed in greater detail in Section 5.

[21]Irrigation assets include wells, Persian wheels, and irrigation channels.

[22]The survey contains information on the dollar value of irrigation investment. However, I choose to use a dummy for irrigation investment, because a substantial fraction of the dollar value of investment is the value of family labor, which is imputed and appears to suffer from substantial measurement error.

[23]In order to address a few negative values of the change in the area of irrigated land, I first do a linear transformation where I add a small amount to the 1-year change in the area of irrigated land, so that log function is defined for all districts in all years.

[24]I do not include the value of land because land markets in India are inactive, and land prices are unreliable. I

To analyze the crop choices of farmers, I construct a measure of the water need of the crop portfolio planted. To do this, I use the daily crop water need values (in centimeters per day) provided on the Agriinfo.in (2015) website. These values, which are specific to the way that crops are grown in India, are presented in Table 2. These numbers represent the average daily amount of water each crop needs, over the course of its growing season, in order to achieve optimal growth. For each farmer or district, I construct an area-weighted average of the water need across all of the crops grown. I also construct a separate water need index that is only for the crops that are primarily grown during the monsoon season. This index drops the crops of wheat, barley, mustard, oilseeds, and potato, which are primarily grown during the dry season.[25]

I merge the agricultural data with gridded weather data from the Terrestrial Precipitation: Monthly Time Series (1900–2008), version 2.01, and the companion Terrestrial Air Temperature data set.[26] The weather data for each 0.5-degree latitude–longitude grid point measure combines information from 20 nearby weather stations, using an interpolation algorithm based on the spherical version of Shepard's distance-weighting method. To merge the weather data with my household data set, I use the rainfall from the weather grid point nearest to each village. Because some villages are closest to the same weather grid point, this method yields 163 unique grid points that provide weather data.[27] For the district data set, I use the rainfall from the grid point nearest to the district center. This results in a sample of 268 unique grid points that provide weather data.

I use several different rainfall measures, all of which are based on growing season rainfall.[28] I measure current year rainfall as the z-score deviation from that location's historical mean, where

deflate wealth values to 1971 rupees. Data on farm inventory is only available for the 1999 REDS round, so I only include it for that round. My results are unchanged if I drop farm inventory from my measure of wealth. Table 1 also includes information about inherited wealth, which is used in my empirical strategy and is discussed in greater detail in Section 5.

[25]I do not have season-specific crop planting data in my household or district datasets. Therefore, I make this distinction based on which crops are typically grown in the monsoon versus non-monsoon seasons.

[26]Kenji Matsuura and Cort J. Willmott, at the Center for Climatic Research, University of Delaware, constructed the data sets with support from IGES and NASA.

[27]Of these 163 grid points, 112 are used for a single village, 34 grid points are used for two villages each, 12 grid points are used for three villages each, and 7 grid points are used for four or more villages each.

[28]Based on the state-specific monthly rainfall charts in Pant and Kumar (1997), the growing season is defined as June through September for most of the country, and June through December for the peninsular region (located in the south).

I use the entire range of my rainfall data (from 1900–2008) to construct the historical means and standard deviations. I use two different specifications to capture the decadal variability of the monsoon. In the first specification, I calculate the simple average of the rainfall z-scores from the past decade. Figures 4 and 5 display the values of these lagged average rainfall for each district, for each year of the REDS survey and for representative years from the WB survey.

I also use a specification where I measure lagged rainfall as the number of especially wet or dry years over the past decade. Following Jayachandran (2006), I use the 20th percentile as the cutoff for a dry year and the 80th percentile as the cutoff for a wet year.[29] I choose these lagged decadal rainfall measures as a rough indicator of the current monsoon regime.[30] In Figures 3 and 4 in Appendix B, I present graphs of the number of dry shocks for the relevant years of the REDS and World Bank data sets in order to demonstrate the variation in this rainfall measure.[31]

Panel B of Table 1 gives the means and standard deviations for the rainfall variables for the relevant years of the household and district surveys. In addition, since the variation in decadal rainfall is critical for my empirical strategy, I present additional discussion of this variation in Appendix B. Specifically, I present a more detailed summary of the statistics of the rainfall variables. As noted above, I present maps of the number of dry shock sin the previous decade for relevant survey years. I also construct spatial correlograms, which allow me to measure the spatial autocorrelation of my decadal lagged rainfall variable.[32]

[29]As with the z-score, I use the entire span of the rainfall data, from 1900–2008, to construct the percentiles.

[30]In Section A.1 of the supplementary file, I test whether my regressions are robust to using an alternate 5- or 15-year rainfall window.

[31]The corresponding graphs of wet shocks from the past decades are available upon request from the author.

[32]Decadal lagged rainfall exhibits substantial spatial autocorrelation. In Appendix C.1, I adjust my regression specifications to be robust to this spatial autocorrelation.

5 Empirical Strategy

5.1 The Returns to Irrigation

I begin by estimating the effects of irrigation and rainfall on profits. To verify the assumptions from Section 3.2, I need to demonstrate that higher rainfall both increases profits and reduces the returns to irrigation. I run the following regression for agricultural households:

$$\begin{aligned}\pi_{ijt} =& \beta_1 rain_{jt} + \beta_2 propirr_{ijt} + \beta_3 rain_{jt} * propirr_{ijt} + \beta_4 wealth_{ijt} + \\ &+ \beta_5 temperature_{jt} + \delta_t + \kappa_{ij} + \epsilon_{ijt}.\end{aligned} \tag{6}$$

The dependent variable π_{ijt} represents agricultural profits per acre for household i, in village j, in year t. The explanatory variables are current rainfall $rain_{jt}$, the proportion of irrigated land $propirr_{ijt}$, wealth $wealth_{ijt}$, temperature $temperature_{jt}$, a year fixed effect δ_t, a household fixed effect κ_{ij}, and an error term ϵ_{ijt} that includes all (non-weather) productivity shocks. The household survey follows households after household splits and after changes of the household head. Therefore, my household fixed effect is common to all parts of the household dynasty that have broken off from the original surveyed household, and can be thought of as a dynasty fixed effect.

Despite the dynasty fixed effect, there are two potential sources of endogeneity for $propirr_{ijt}$ in equation 6. The first is that, despite the dynasty fixed effect, $propirr_{ijt}$ may be correlated with ϵ_{ijt} if households can adjust their irrigation investments in response to the current productivity shocks. A second source of concern is that $propirr_{ijt}$ may be correlated with (unobserved) farmer ability. The dynasty fixed effect controls for average farmer ability across all households within a common dynasty. However, if there is variation in farmer ability amongst the households within a dynasty, then ϵ_{ijt} may be correlated with $propirr_{ijt}$, despite the dynasty fixed effect.

Similarly, there are two potential sources of endogeneity for $wealth_{ijt}$. First, if the current period's productivity shocks are correlated with lagged productivity shocks, then wealth will be endogenous (because lagged productivity shocks affect wealth).[33] Second, as above, if there is

[33]For the 1999 round of REDS I have data on the value of the current farm inventory and I can include this in my measure of wealth, which eliminates the concern for that survey round. However, the earlier rounds of REDS do not

variation in farming ability within a dynasty, and farming ability is correlated (within the dynasty) with wealth, this will also cause endogeneity.

I employ an instrumental variables strategy that attempts to mitigate these endogeneity concerns. The REDS survey contains information, for each household and split-off household, about the amounts of wealth, land, and irrigated land that were inherited at the time of household formation. Typically in India, at the time of a father's death, each son in the household will inherit land and become head of his own separate household (Fernando, 2014). I instrument for $propirr_{ijt}$ and $wealth_{ijt}$ with $inhpropirr_{ijt}$—the proportion of inherited land that was irrigated at the time of inheritance—and $inhwealth_{ijt}$. Earlier work has used the same instrumental variables strategy (Foster and Rosenzweig, 1995, 2001, 2010).

Due to household splits, many dynasties include multiple household heads, which gives me variation in these inheritance variables, even in the presence of the dynasty fixed effect.[34] I now discuss the exclusion restriction for my instruments and the extent to which these instruments reduce the endogeneity issues outlined above. First, consider the endogeneity that arises due to transient, current-period productivity shocks (that are not due to unobserved farmer ability). Because inheritances occurred in an earlier period, we should expect that inherited wealth and inherited irrigated land should be less correlated with transient current period productivity shocks than current period wealth and current period irrigated land. Thus, we should expect the instruments to significantly reduce this source of endogeneity bias.

Second, consider the endogeneity that arises due to variations in farmer ability within a dynasty. Relative to this source of endogeneity, there is potential concern about whether the exclusion restriction holds. For example, if a son with higher farming ability inherits more wealth and more irrigated land, then the exclusion restriction would be violated. However, there is evidence that inheritances may not be strongly correlated with variations in sons' ability (Foster and Rosenzweig, 2002; Fernando, 2014). Fernando (2014) finds that the amount of inherited land is very

include farm inventory in the survey; so, for those rounds, concerns about lagged productivity shocks are still an issue.

[34]Amongst the households that I analyze, 42% come from dynasties with multiple household heads. My estimate of the coefficients of wealth and the proportion of irrigated land will be a local average treatment effect based on these households.

strongly dictated by the number of sons in the family (and the total amount of land), and states that "equal division amongst sons [at the time of a father's death] is the norm."[35,36] This suggests that household inheritances may not be strongly correlated with farmer ability. I cannot fully prove that inheritances are not correlated with unobserved household ability, but to the extent that these inheritances are *less* correlated than current wealth and current irrigated land, my instrumental variables strategy should at least *reduce* the endogeneity bias.[37]

Due to data limitations, my district regression is a modified version of equation 6. The unit of observation for the regression is district j in year t. I use agricultural revenue per acre, $revenue_{jt}$, as the dependent variable. I do not control for wealth. I include $propirr_{jt}$, but do not instrument for it. The household fixed effect becomes a district fixed effect κ_j.

For both data sets, finding $\beta_1 > 0$ and $\beta_3 < 0$ will confirm the assumptions of Section 3.2, namely that higher rainfall increases profits and also reduces the returns to irrigation.

5.2 Tests for Irrigation Adaptation

I next analyze how irrigation investment responds to lagged rainfall:

$$irr\_inv_{ijt} = \alpha_1 decaderain_{jt} + \alpha_2 rain_{jt} + \alpha_3 rain_{jt-1} + \lambda_t + \mu_{ij} + \zeta_{ijt}. \tag{7}$$

In the household specification, irr_inv_{ijt} is a dummy variable equal to 1 if, during the recall period, a household purchased irrigation equipment or used labor to create/improve irrigation assets.[38] The explanatory variables are past decade rainfall $decaderain_{jt}$, current year rainfall $rain_{jt}$, 1-year lagged rainfall $rain_{jt-1}$, a year fixed effect λ_t, a household fixed effect μ_{ij}, and an error term ζ_{ijt}.[39] I measure $decaderain_{jt}$ in two ways. The first measure is a simple average of the rainfall

[35] I haven't found a similar analysis of irrigated land but it seems plausible that it would follow the same pattern.

[36] Note that even if there is purely equal division of inheritances amongst sons, changes across generations of household heads provides variation of the inheritances within each dynasty.

[37] In Appendix C.1, I also rerun my regressions without instrumenting to see how this affects my estimates.

[38] I do not use the rupee value of investment, because a large component of it is family labor, the value of which is measured with a lot of noise.

[39] I include lagged rainfall because residual impacts of last year's rainfall may influence this year's irrigation and cropping decisions directly, independent of an expectations/adaptation effect. In addition, my household data set does not include the specific interview date for each household, so including lagged rainfall is important because, for

z-scores from the past decade. The second measure tabulates the proportion of years in the past decade that were especially wet or dry. Following Jayachandran (2006), I use the 20th percentile as the cut-off for a dry year and the 80th percentile as the cut-off for a wet year.

The coefficient of interest in this regression is α_1. My model demonstrates that the sign of α_1 is ambiguous and must be determined empirically. If the wealth effect dominates, then α_1 will be positive. Irrigation investment will increase after wet decades, due to an accumulation of wealth and increased investment in all assets. On the other hand, if farmers are adapting to expected rainfall *and* the size of this effect is larger than the wealth effect, then we will find $\alpha_1 < 0$. Irrigation investment will increase after dry decades, due to farmers expecting more dry years in the future. Thus, finding $\alpha_1 < 0$ provides evidence of adaptation.[40]

I control for current year rainfall because farmers can invest in irrigation at any time during the year. Thus, a farmer's observation of current year rainfall (based on, say, the first half of the growing season) might directly affect his decision to invest in irrigation that period. This response would not indicate adaptation to expected future year rainfall, but would simply reveal within-season adjustment to current year rainfall.[41]

Propositions 3.1 and 3.2 demonstrate that I can test for irrigation adaptation with or without a wealth control. Thus, for completeness, I run a second household specification where I control for wealth. Once I have isolated the wealth effect, my model predicts that $\alpha_1 = 0$ if farmers are not adapting. On the other hand, if farmers are adapting, then $\alpha_1 < 0$. The variable $wealth_{ijt}$ is endogenous in this regression, and so I instrument it with $inhwealth_{ijt}$. The validity of the instrument follows the same logic as for equation 6.

For my district regression, I define irr_inv_{jt} as the log of the 1-year change in the district's irrigated area, I use a district fixed effect, and I do not control for wealth. Proposition 3.1 demonstrates that I can test for irrigation adaptation, even in the absence of a wealth control. As with the

households interviewed early in survey year, rainfall from the previous calendar year may be the most relevant.

[40]Finding a positive coefficient would be inconclusive; it would neither demonstrate nor rule out the possibility of adaptation.

[41]I also control for rainfall from the previous year because the exact date of the REDS survey for each household is unknown, but all households use a 12-month recall period for their answers. Therefore, for some households, the actual relevant rainfall year may be earlier than the year of the survey.

household regression, finding $\alpha_1 < 0$ provides evidence of adaptation.

5.3 Test for Crop Adaptation

Lastly, I test for crop adaptation. I only perform this test with my household data set, and my regression is of the form:

$$water\_need_{ijt} = \gamma_1 decaderain_{jt} + \gamma_2 rain_{jt} + \gamma_3 rain_{jt-1} + \gamma_4 wealth_{ijt} + \tau_t + \phi_{ij} + \psi_{ijt} \quad (8)$$

where $water\_need_{ijt}$ is the area-weighted water need of the farmer's crop portfolio. As mentioned above, I control for current year rainfall because farmers may have some knowledge of the current year's rainfall before they sow all of their crops. As in the irrigation regression, a response of crop choice to current year rainfall would indicate a within-season adjustment to rainfall, but would not provide evidence of adaptation to the expected future year rainfall.[42]

I control for $wealth_{ijt}$ because, as demonstrated in Section 3.4.2, without a control for $wealth_{ijt}$, I could not interpret γ_1 as evidence of adaptation. As in the equation 6, wealth is endogenous and I instrument for it with inherited wealth. The validity of the instrument follows the same argument as its validity in equation 6. Finding $\gamma_1 = 0$ demonstrates that farmers are not adapting their crop portfolios. Conversely, in the presence of adaptation, I expect to find $\gamma_1 > 0$.

6 Results

6.1 The Returns to Irrigation

Table 3 tests the impacts of rainfall and irrigation on profits. In the household regressions (shown in columns 1 and 2) the dependent variable is profits per acre. In column 1, I deduct the value of family labor, and in column 2, I do not. I measure rainfall using quintiles to allow for non-linear effects. I instrument for the proportion of irrigated land with the proportion of inherited irrigated

[42] I also control for rainfall from the previous year because the exact date of the REDS survey for each household is unknown, but all households use a 12-month recall period for their answers. Therefore, for some households, the actual relevant rainfall year may be earlier than the year of the survey.

land.[43] In the district regression (shown in column 3) the dependent variable is revenue per acre, and I do not instrument for irrigation. For both data sets, the coefficients demonstrate that higher rainfall increases profits and that the returns to irrigation rise during dry years, thus confirming the assumptions of Section 3.2. In Table 3, and all the tables below, standard errors are clustered at the rainfall grid point level to account for correlation in error terms between adjacent households (or districts) that share the same value of rainfall due to the resolution of the rainfall data.[44]

6.2 Tests for Irrigation Adaptation

Table 4 tests whether farmers are adapting their irrigation investments in response to lagged rainfall. Recall that I can test for irrigation adaptation either with, or without, a wealth control. Columns 1 through 4 use the household data and, in columns 2 and 4, I control for wealth, which is instrumented for with inherited wealth. Columns 5 and 6 use the district data and do not control for wealth. In all columns, I find the coefficient of lagged rainfall is negative, which provides evidence of adaptation. In terms of magnitudes, column 4 demonstrates that a dry year in the preceding decade increases the probability of irrigation investment during the recall period by 1.2 percentage points. The baseline probability of investing in irrigation during the recall period is 5%.[45]

6.3 Test for Crop Adaptation

In Table 5, I test for crop adaptation using the household data set. I control for wealth in all columns and instrument for it with inherited wealth.[46] Columns 1 and 2 look at the daily water need of all crops grown in the year, and columns 3 and 4 focus on the daily water need of monsoon season crops only. The columns that use the average rainfall specification are not significant. However,

[43] The F-statistics, presented at the bottom of the table, indicate that the first-stage regressions are sufficiently strong.

[44] In Appendix C.1, I test the robustness of my adaptation results to using spatially correlated standard errors.

[45] The F-statistics for columns 2 and 4, presented at the bottom of the table, indicate that the first-stage regressions for wealth are sufficiently strong. For concision, I don't display the first-stage regression coefficients, but they are available upon request.

[46] The F-statistics, presented at the bottom of the table, indicate that the first-stage regressions are sufficiently strong. For concision, I do not display the first-stage regression coefficients, but they are available upon request.

the columns that use the wet/dry shock rainfall specification are significant at the 1% level for both all-season crops and monsoon season crops. The coefficient in column 2 indicates that each additional dry year in the past decade reduces the average water need of the crops planted by 0.16 mm/day, relative to an average water need of 7.3 mm/day. Restricting the analysis to monsoon season crops, I find that each additional dry year in the past decade reduces the average daily water need of a farmer's monsoon season crop portfolio by 0.2 mm/day, relative to an average water need of monsoon season crops of 8 mm/day.[47]

To further understand the crop adaptation results, I run individual crop regressions. In each column of Table 6, the dependent variable is the proportion of the farmer's land that is planted with a specific crop for the top eight crops by area in the REDS data set. These crops are (in order of area): rice, wheat, pulses, millet, cotton, sorghum, groundnut, and maize). Consistent with the average crop portfolio results in Table 5, I find that after decades with more dry years, farmers plant less proportional area with rice (1.075 mm/day water need) and more with pulses (3.50 mm/day water need). These are the crops with the highest and lowest water needs, respectively. I also find that after wet shocks, farmers plant less millet (5.75 mm/day) and sorghum (5.75 mm/day), possibly because they are switching toward planting more area with the higher water need crop of rice. I also find that after decades with more wet shocks, farmers plant more wheat (4.25 mm/day) and more cotton (5.25 mm/day). These results are counterintuitive because these crops have slightly lower water needs than millet and sorghum. However, the coefficient in the wheat regression is only significant at the 10% level, and is of smaller magnitude than the coefficients for the other crops. Lastly, the coefficients for peanut and maize are not statistically significant, but this may be because they represent a smaller total area.

Taken together, the irrigation and crop regressions suggest evidence that farmers are adapting in response to recent decadal rainfall. However, it is important to note that due to attrition, the household sample in REDS is not nationally representative.[48] Specifically, the famers in my sample

[47]The full list of crops analyzed is shown in Table 2. For the monsoon season crops, I drop barley, mustard, oilseeds, potato, and wheat, all of which are primarily grown during the dry season.

[48]This attrition is discussed in greater detail in Appendix A.

have, on average, higher land areas, higher proportions of irrigated land, and higher levels of non-land wealth than would be found in a representative sample. Thus, my adaptation results are an accurate representation of the behavior of this particular population, but my results may not hold more broadly.

7 Robustness

In the supplementary file, I investigate the robustness of my results.

First, I re-run my adaptation regressions using a fixed effects specification that drops my instrumental variables strategy. In light of potential issues with my instruments, these results provide another set of estimates, which may be of interest. In addition, using the non-IV specification allows me to estimate a new set of standard errors that allow for spatial correlation. This is important because, as indicated in Appendix B, my lagged rainfall variable demonstrates significant spatial correlation. To implement these standard errors, I use code from Hsiang and Solow (2010) and Fetzer (2014). Guided by the autocorrelograms in Appendix B, I allow for a correlation within 800 kilometers using a Bartlett (triangular) kernel. My results are robust to these changes.[49]

Second, I re-estimate my regressions using rainfall lag windows of 5 or 15 years, to verify that the choice of a 10-year window is not driving my results. My district irrigation regressions and my household crop regressions are robust to using alternate rainfall windows. In my household irrigation regressions, however, the signs of the coefficients of interest are preserved but are no longer statistically significant. I also present regressions in which I control for lagged rainfall separately for each year (rather than as an average). In this case, the coefficients are no longer individually statistically significant (likely because they are correlated with each other), but the set of rainfall lags is jointly significant.

Third, I discuss the possibility that depletion of groundwater and/or surface water might be

[49] I do not implement spatially correlated standard errors in the specification presented in the paper because there is not code available to run spatially correlated standard errors for a regression that has instrumental variables and fixed effects.

causing the relationship between irrigation investment and lagged rainfall that I have found. Using irrigated area (rather than irrigation investment) as my dependent variable, I find that the area of irrigated land *increases* after dry decades. As I discuss in greater detail in Appendix C.3, this test provides some reassurance that water depletion is not driving my results, but it does not fully rule out that possibility.

Fourth, I test whether my irrigation adaptation results might be due to public (government) investments rather than private (farmer) investments. In India, the bulk of direct public irrigation investments are large-scale dams.[50] When I control for the presence of these dams, my irrigation adaptation results are preserved. However, as I discuss in Appendix C.4, government irrigation investment is an outcome variable and may be endogenous to the household investment decision. For this reason, my results are suggestive that the adaptation results are not driven by public investment, but do not definitively rule out that possibility.

Fifth, I test whether changes in agricultural technology or policies might be confounding my results. I add controls for high-yielding variety crops, electrification rates, fertilizer prices, financial institutions, agricultural extension services, transportation infrastructure, and government intervention in output markets. My irrigation and crop adaptation results are robust to adding these controls. However, as discussed in Appendix C.5, these results are suggestive only, not definitive, since I am only able to control for a subset of possible confounders. Furthermore, the confounders that I do control for are potentially endogenous to the household irrigation and crop decisions.

Lastly, I re-run my regressions with region-by-year fixed effects. Due to the large-scale spatial correlation of the monsoons, it is possible that unobserved, time-varying confounding factors might be correlated with my lagged rainfall variables. Including region-by-year fixed effects is a flexible way to control for this. My district irrigation results and my household crop results are robust to the addition of these controls. In my household irrigation regressions, however, the signs of the coefficients of interest are preserved but are no longer statistically significant. The lack of

[50]Although most *direct* government irrigation investment is via large-scale dams, the government does subsidize groundwater irrigation through credit programs and electricity subsidies. Controlling for dams does not address these channels. However, I do control for some of these policies in Appendix C.5.

robustness of my irrigation results to this change in specification is a major limitation of my study.

The supplementary file provides more details on these tests.

8 Effectiveness of Adaptation

The preceding text has found evidence of adaptation; this section quantifies its efficacy. What fraction of profits were farmers able to protect from adverse climate variations? To answer this question, I use the household data set to estimate the extent to which irrigation adaptation increased profits during 1971–1999.[51] Rainfall during this period was below average (as shown in Figure 1), and this reduced profits. On the other hand, my adaptation regressions indicate that farmers noticed that rainfall was below average during this period and increased their investment in irrigation. This investment in irrigation should have at least partially offset some of the losses due to the drier than average climate. I seek to estimate how beneficial or effective this adaptation was.

To calculate the efficacy of adaptation, I do two things. First, I estimate the percentage of profits that were lost due to the drier than average rainfall that occurred during 1971–1999. Next, I estimate the percent of these losses that were recovered due to increased investment in irrigation. To calculate these percentage changes requires that I estimate profits from three different scenarios: actual profits, counterfactual profits under a scenario where the dry regime did not occur, and counterfactual profits under a scenario where the dry regime occurred but farmers did not adapt. I now describe in detail how I calculate these three quantities.

First, I estimate what the actual profits were for each farmer over the period from 1971–1999, given the rainfall that actually occurred and the actual irrigation decisions of farmers. Since I only have survey data for three points in time (1971, 1982, and 1999), in order to estimate the profits in the intervening years, I use interpolation. I use the actual wealth and irrigation in each survey

[51]The analysis focuses on irrigation adaptation because the efficacy of crop adaptation is not calculable. Specifically, the data do not permit an unbiased estimate of the impact of crop portfolio on profits. Unobserved shocks, such as health shocks, may be correlated with both profits and drought-tolerant crop areas, and hence a regression of profits on drought-tolerant areas will be biased. For irrigation, in contrast, I can instrument for irrigated land with inherited irrigated land and remove, or at least reduce, this bias.

round and interpolate them, to have a predicted level of irrigation and wealth for each farmer for every year between 1971 and 1999. Next, I combine these estimates with the actual rainfall that occurred for each of these intermediate years. Then, I use the regression coefficients from Table 3, column 2, to estimate the profits per acre for each farmer. I then sum these up over all the years between 1971-1999 to have an estimate, for each farmer, of what his total profits were over this period, given the actual rainfall that occurred and the actual wealth and irrigation that he had.

Second, I calculate what the profits would have been for each farmer over the period from 1971–1999 if the drier-than-average rainfall period had not occurred. In order to do this, I calculate expected annual profits, using a 20% chance of each rainfall quintile occurring. This calculation effectively projects what expected profits would have been if rainfall were at its historical mean distribution. I interpolate irrigation and wealth for non-survey years for this counterfactual scenario as well, and I again use the regression coefficients from Table 3, column 2. I then sum these up over all the years between 1971 and 1999 to have an estimate, for each farmer, of what his expected profits would have been had rainfall not been below average. I then compare these "expected profits" to the actual profits computed above, in order to calculate what fraction of profits each farmer lost due to the drier rainfall regime.

Lastly, I want to calculate what fraction of these theoretical "lost profits" farmers were able to recover due to their increased investment in irrigation which, at least partially, offset these losses. In order to do this, I want to look at the path of irrigation over the period from 1971–1999 and determine what fraction of this irrigation investment was due to adaptation to lagged rainfall. I then will subtract away this quantity of "adapted irrigation" in order to calculate what level of irrigation each farmer would have had in the absence of adaptation. Then I will calculate the profits that farmers would have had, if they experienced the drier than average rainfall and had this lower level of irrigation that did not incorporate adaptation. This gives me a measure of counterfactual profits under a scenario where the dry regime occurred but farmers did not adapt their irrigation.

To do this, I need to compute a counterfactual value of what the proportion of irrigated land would have been for each farmer in the absence of adaptation. I use the coefficients from column 2 of Table 7 to calculate the adaptive response of irrigation to lagged rainfall.[52] This table is analogous to my baseline irrigation adaptation specification (Table 4) but uses the proportion of irrigated land as the dependent variable (rather than an irrigation investment dummy). The irrigation investment dummy captures precisely how the household is adjusting its irrigation this year. However, using it requires knowing what fraction of the farmer's land becomes irrigated when he invests in irrigation, since profits depend on the proportion of land irrigated. Thus, I instead use the proportion of irrigation, which is a coarser measure of adaptation. This allows me to subtract a quantity of "adapted irrigation" from the interpolated irrigation, to calculate what irrigation would have been in the absence of adaptation. I combine this lower irrigation level with the actual weather realizations and interpolated wealth and the profit regression coefficients to get an estimate of counterfactual profits under a scenario where the dry regime occurred but farmers did not adapt their irrigation.[53]

Using these profit measures, I find that on net the dry regime decreased farmers' profits by 0.4%. However, there is substantial heterogeneity among households, and for households with losses, the average loss was 3.1%. Furthermore, I estimate that farmers with profit losses recovered only 9% of their losses on average. However, it is important to note that my coefficient for adaptation in the alternate specification in Table 7 is estimated with less precision than my preferred estimate in Table 3, and is only statistically significant at the 10% level. In order to counterbalance this, I calculate a 95% confidence interval for that coefficient, and then estimate what fraction of lost profits would have been recovered if the true adaptation coefficient was at the upper or lower end of this range. I find that, at most, 19% of lost profits would have been recovered, and potentially as little as zero. Altogether my estimates suggest that farmer adaptation to persistent rainfall deviations appears to have had limited efficacy. This is suggestive that adaptation to future anthropogenic climate change may be limited. However, extrapolating my results directly

[52]Note that I use the adaptation specification where I control for wealth. This ensures that all of the response that I see for irrigation in response to lagged rainfall is due to the expectations and is not simply due to the wealth effect.

[53]A more straightforward calculation would be to compare actual profits over the period 1971–1999 based on actual irrigation to profits for that same period if irrigation had remained at its 1971 levels. However, this would assume that all of the growth in irrigation was due to responses to drier rainfall, whereas in fact a large part of it was likely due to over all trends in irrigation that were driven by changes in technology, etc.

to future climate change is problematic, since future climate change will affect both rainfall and temperature.

In addition, as discussed above, due to sample selection, the farmers in my sample have, on average, higher land areas, a higher fraction of irrigated land, and higher levels of non-land wealth than would be found in a representative sample. Nevertheless, I argue that we would expect wealthier households to be more likely to adapt to climate change than poorer households; so that if anything, it means that the 9% that I estimate is actually an upper bound on the true fraction of profits recovered via adaptation. And, hence, my study still provides useful, although more limited, information.

9 Conclusion

To accurately predict future climate change damages requires an accurate understanding of the ability of agents to adapt to changes in climate. In this paper, I exploit persistent rainfall variations in India over the past 50 years to test whether farmers adjust their irrigation and crop choice decisions in response to recent rainfall. I find evidence of both irrigation adaptation and crop adaptation. However, analysis suggests that the efficacy of adaptation is limited; I estimate that adaptation recovers at most only 19% of lost profits and, more likely, only 9%.

There are several important caveats to my study. First, I look at adaptation to rainfall changes only (not temperature). Second, I analyze only two possible adaptations, when in fact a much broader array of adaptations are possible. Third, my instrumental variables strategy (instrumenting for current wealth and irrigated land with inherited wealth and inherited irrigated land) most likely reduces the endogeneity bias in my regressions but, to the extent that inheritances within a dynasty are non-random, does not fully eliminate this bias. Fourth, data limitations prohibit me from fully controlling for all potential confounding changes in agricultural technology or policies. When I try to control for potential confounders flexibly (by including region-by-year fixed effects), my crop adaptation results are preserved but my household irrigation results are not. Fifth, I am not able to completely rule out the possibility that depletion of water supplies could be driving my irrigation adaptation results. Lastly, my household data set exhibits substantial, non-random attrition and, as a result, the sample I analyze is wealthier, on average, than a representative sample would be. If we expect wealthier farmers to adapt more readily than poorer farmers, this means that my adaptation estimates may be an upper bound. However, since I find that the efficacy of adaptation is limited, even for the households that I analyze, the results of my study are still of interest.

Despite these caveats, my results suggest that, in the context of the historical rainfall deviations that I have analyzed, there are barriers to adaptation. My work does not elucidate the precise nature of these barriers. Other work, summarized by Jack (2011), indicates that credit and information constraints, as well as inefficiencies in input, output, land, labor, and risk markets, inhibit agricultural adaptation in a variety of situations. The specific barriers to climate change adaptation and, importantly, the institutions, technologies, and policies that might remove these barriers, call for further exploration.

References

Adams, R. M., Hurd, B. H., Lenhart, S., and Leary, N. (1998). Effects of global climate change on agriculture: an interpretative review. *Climate Research*, 11(1):19–30.

Agriinfo.in (2015). Water Requirement and Irrigation Requirement. http://www.agriinfo.in/?page=topic&superid=1&topicid=16.

Barreca, A., Clay, K., Deschênes, O., Greenstone, M., and Shapiro, J. S. (2015). Adapting to Climate Change: The Remarkable Decline in the U.S. Temperature-Mortality Relationship over the 20th Century. pages 1–68.

Bryan, E., Deressa, T. T., Gbetibouo, G. A., and Ringler, C. (2009). Adaptation to climate change in Ethiopia and South Africa: options and constraints. *Environmental Science & Policy*, 12(4):413–426.

Burgess, R., Deschenes, O., Donaldson, D., and Greenstone, M. (2014). The unequal effects of weather and climate change: Evidence from mortality in India.

Burke, M. and Emerick, K. (2015). Adaptation to Climate Change: Evidence from US Agriculture. pages 1–74.

Christensen, J. H. and Hewitson, B. (2007). Regional climate projections. In *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment*

Report of the Intergovernmental Panel on Climate Change, pages 847–940. Cambridge University Press, Cambridge, UK.

Cole, S., Healy, A., and Werker, E. (2012). Do voters demand responsive governments? Evidence from Indian disaster relief. *Journal of Development Economics*, 97(2):167–181.

da Cunha, D. A., Coelho, A. B., and Féres, J. G. (2014). Irrigation as an adaptive strategy to climate change: an economic perspective on Brazilian agriculture. *Environment and Development Economics*, 20(01):57–79.

Das, P. K. (1995). *The Monsoons*. National Book Trust, New Delhi, India, 3rd edition.

Dell, M., Jones, B. F., and Olken, B. A. (2012). Temperature Shocks and Economic Growth: Evidence from the Last Half Century. *American Economic Journal: Macroeconomics*, 4(3):66–95.

Dell, M., Jones, B. F., and Olken, B. A. (2014). What Do We Learn from the Weather? The New Climate-Economy Literature. *Journal of Economic Literature*, 52(3):740–98.

Deryugina, T. (2013). The Role of Transfer Payments in Mitigating Shocks: Evidence From the Impact of Hurricanes. pages 1–64.

Deschênes, O., Deschênes, O., Greenstone, M., and Greenstone, M. (2007). The Economic Impacts of Climate Change: Evidence from Agricultural Profits and Random Fluctuations in Weather. *The American Economic Review*, 97(1):354–385.

Deschênes, O. and Greenstone, M. (2011). Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US. *American Economic Journal: Applied Economics*, 3(4):152–185.

Di Falco, S., Adinolfi, F., Bozzola, M., and Capitanio, F. (2014). Crop Insurance as a Strategy for Adapting to Climate Change. *Journal of Agricultural Economics*, 65(2):485–504.

Di Falco, S. and Veronesi, M. (2013). How Can African Agriculture Adapt to Climate Change? A Counterfactual Analysis from Ethiopia. *Land Economics*, 89(4):743–766.

Di Falco, S. and Veronesi, M. (2014). Managing Environmental Risk in Presence of Climate Change: The Role of Adaptation in the Nile Basin of Ethiopia. *Environmental and Resource Economics*, 57(4):553–577.

Duflo, E., Kremer, M. R., and Robinson, J. (2011). Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya. 101(6):2350–2390.

Fernando, A. N. (2014). Shackled to the Soil: The Long-Term Effects of Inherited Land on Labor Mobility and Consumption. pages 1–81.

Fetzer, T. (2014). Can workfare programs moderate violence? Evidence from India.

Finger, R., Hediger, W., and Schmid, S. (2010). Irrigation as adaptation strategy to climate change—a biophysical and economic appraisal for Swiss maize production. *Climatic Change*, 105(3-4):509–528.

Fleischer, A., Lichtman, I., and Mendelsohn, R. (2008). Climate change, irrigation, and Israeli agriculture: Will warming be harmful? *Ecological Economics*, 65(3):508–515.

Foster, A. and Rosenzweig, M. (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy*.

Foster, A. and Rosenzweig, M. (2001). Imperfect commitment, altruism, and the family: Evidence from transfer behavior in low-income rural areas. *Review of Economics and Statistics*, 83(3):389–407.

Foster, A. D. and Rosenzweig, M. R. (2002). Household Division and Rural Economic Growth. *Review of Economic Studies*, 69(4):839–869.

Foster, A. D. and Rosenzweig, M. R. (2010). Is There Surplus Labor in Rural India? pages 1–39.

Giné, X., Townsend, R., and Vickery, J. (2009). Forecasting when it matters: evidence from semi-arid India.

Guiteras, R. (2009). The impact of climate change on Indian agriculture. *Manuscript, Department of Economics, University of Maryland, College Park, Maryland*.

Hornbeck, R. (2012). The Enduring Impact of the American Dust Bowl: Short- and Long-Run Adjustments to Environmental Catastrophe. *The American Economic Review*, 102(4):1477–1507.

Hornbeck, R. and Naidu, S. (2015). *When the Levee Breaks: Black Migration and Economic Development in the American South*. Forthcoming in American Economic Review.

Hsiang, S. M. and Narita, D. (2012). Adaptation to Cyclone Risk: Evidence from the Global Cross-Section. *Climate Change Economics*, 03(02):1–28.

Hsiang, S. M. and Solow, R. M. (2010). Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America. *Proceedings of the National Academy of Sciences of the United States of America*, 107(35):15367–15372.

Jack, B. K. (2011). Constraints on the adoption of agricultural technologies in developing countries. *White paper. Agricultural Technology Adoption Initiative, J-PAL (MIT) and CEGA (UC Berkeley).*

Jayachandran, S. (2006). Selling labor low: Wage responses to productivity shocks in developing countries. *Journal of Political Economy*, 114(3):538–575.

Kripalani, R. H. and Kulkarni, A. (1997). Climatic impact of El Niño/La Niña on the Indian monsoon: A new perspective. *Weather*, 52(2):39–46.

Krishna Kumar, K., Rupa Kumar, K., Ashrit, R. G., Deshpande, N. R., and Hansen, J. W. (2004). Climate impacts on Indian agriculture. *International Journal of Climatology*, 24(11):1375–1393.

Kurukulasuriya, P., Kala, N., and Mendelsohn, R. (2011). Adaptation And Climate Change Impacts: A Structural Ricardian Model of Irrigation and Farm Income in Africa. *Climate Change Economics (CCE)*, 2(02):149–174.

Kurukulasuriya, P. and Mendelsohn, R. (2008). Crop switching as a strategy for adapting to climate change. *African Journal of Agricultural and Resource Economics*, 2(1):105–125.

Luis, M. G. and Orlando, R. (2015). Climate change, irrigation and agricultural activities in Mexico: A Ricardian analysis with panel data. *Journal of Development and Agricultural Economics*, 7(7):261–272.

Mendelsohn, R., Nordhaus, W. D., and Shaw, D. (1994). The impact of global warming on agriculture: a Ricardian analysis. *The American Economic Review*, pages 753–771.

Mooley, D. A. and Parthasarathy, B. (1984). Fluctuations in All-India Summer Monsoon Rainfall During 1871-1978. *Climatic Change*, 6(3):287–301.

Mukherjee, M. and Schwabe, K. (2015). Irrigated Agricultural Adaptation to Water and Climate Variability: The Economic Value of a Water Portfolio. *American Journal of Agricultural Economics*, 97(3):809–832.

Palanisami, K., Ranganathan, C. R., and Nagothu, U. S. (2014). Climate change and agriculture in India: studies from selected river basins.

Pant, G. B. (2003). Long-Term Climate variability and change over monsoon Asia. *J. Ind. Geophys. Union*, 7(3):125–134.

Pant, G. B. and Kumar, K. R. (1997). *Climates of South Asia*. Wiley.

Parry, M. L. (2007). *Climate Change 2007: Impacts, adaptation and vulnerability*, volume 4 of *contribution of Working Group II to the fourth assessment report of the Intergovernmental Panel on Climate Change*. Cambridge University Press.

Rose, E. (2001). Ex ante and ex post labor supply response to risk in a low-income area. *Journal of Development Economics*, 64(2):371–388.

Sanghi, A., Kumar, K. K., and McKinsey, Jr, J. W. (1998). India agriculture and climate dataset. Technical report.

Sanghi, A. and Mendelsohn, R. (2008). The impacts of global warming on farmers in Brazil and India. *Global Environmental Change*, 18(4):655–665.

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. *The American Economic Review*, 95(1):395–406.

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.

Seo, S. N., McCarl, B. A., McCarl, B. A., and Mendelsohn, R. (2010). From beef cattle to sheep under global warming? An analysis of adaptation by livestock species choice in South America. *Ecological Economics*, 69(12):2486–2494.

Seo, S. N. and Mendelsohn, R. (2008). An analysis of crop choice: Adapting to climate change in South American farms. *Ecological Economics*, 67(1):109–116.

Staiger, D. and Stock, J. H. (1997). Instrumental Variables Regression with Weak Instruments. *Econometrica*, page 45.

Stock, J. H. and Yogo, M. (2002). Testing for weak instruments in linear IV regression.

Subbaramayya, I. and Naidu, C. (1992). Spatial variations and trends in the Indian monsoon rainfall. *International Journal of Climatology*, 12(6):597–609.

Tol, R. S. J. (2014). Correction and Update: The Economic Effects of Climate Change. *Journal of Economic Perspectives*, 28(2):221–226.

Varikoden, H. and Babu, C. A. (2014). Indian summer monsoon rainfall and its relation with SST in the equatorial Atlantic and Pacific Oceans. *International Journal of Climatology*.

Viswanathan, B. and Kavi Kumar, K. S. (2015). Weather, agriculture and rural migration: evidence from state and district level migration in India. *Environment and Development Economics*, pages 1–22.

Wang, B. (2006). *The Asian Monsoon*. Springer Praxis Books. Springer Science & Business Media.

Wang, J., Mendelsohn, R., and Dinar, A. (2010). How Chinese farmers change crop choice to adapt to climate change. *Climatic Change*, 1(3):167–185.

Figures

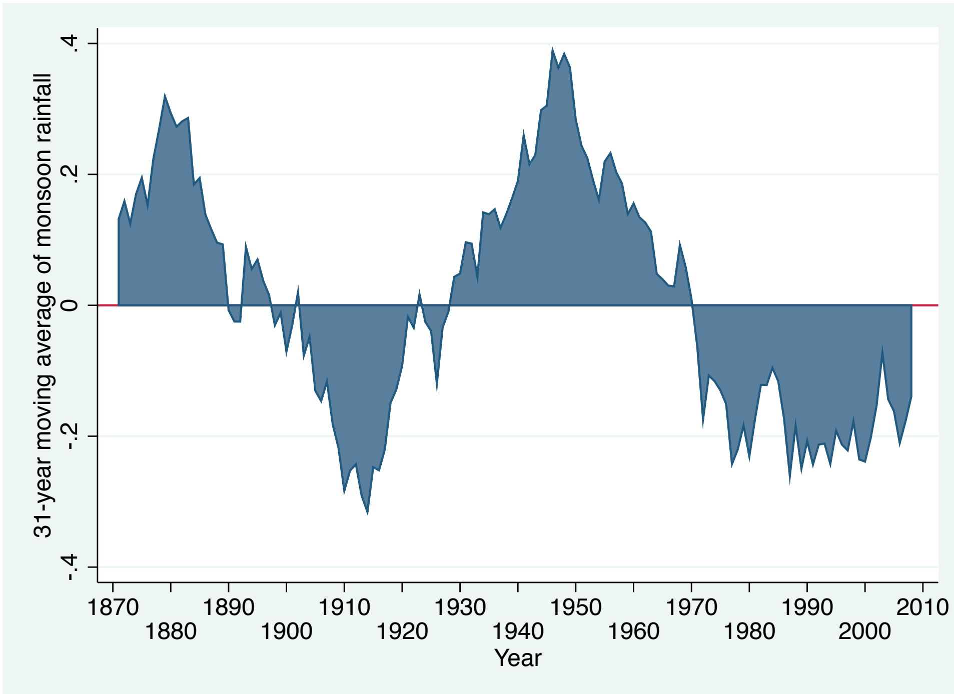

Figure 1: Interdecadal Variability of the Indian Monsoon

Note: This figure displays the 31-year moving average of India's summer monsoon rainfall, measured as a z-score deviation from the historical mean. *Source*: The rainfall data are from the India Institute of Tropical Meteorology's Homogeneous Indian Monthly Rainfall Data Set (1871–2008). The figure is constructed based on the author's calculations.

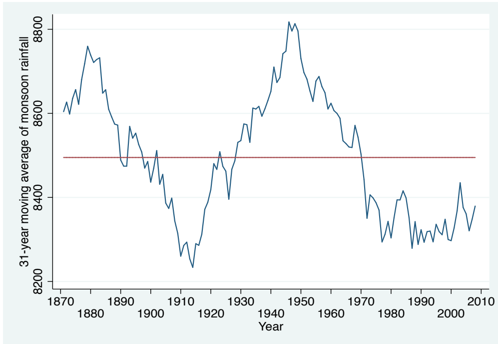

(a) All-India

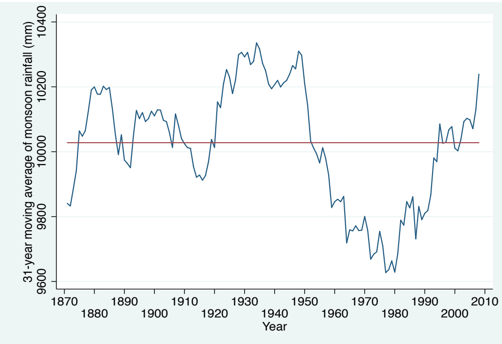

(b) Central Northeast India

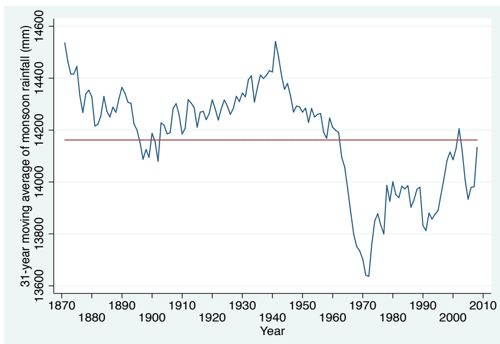

(c) Northeast India

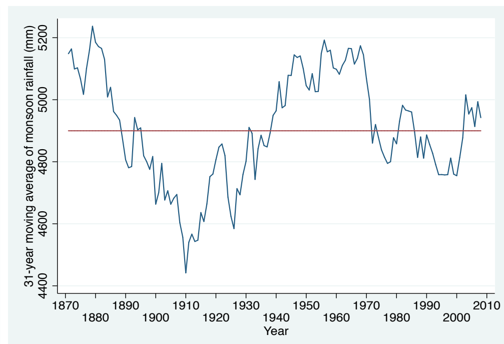

(d) Northwest India

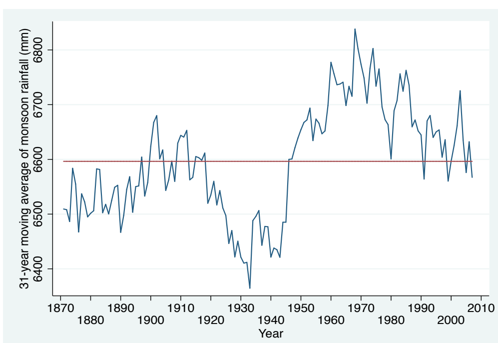

(e) Peninsular India

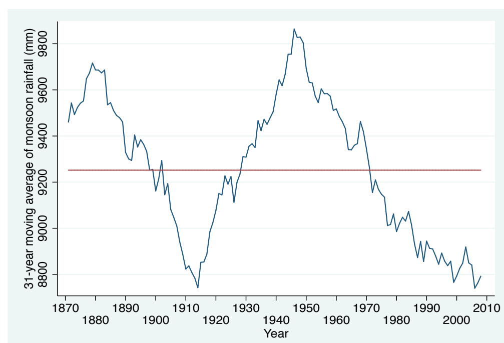

(f) West Central India

Figure 2: Spatial Variation of the Interdecadal Variability of the Indian Monsoon

Note: This figure graphs the 31-year moving average of the summer monsoon rainfall, measured in millimeters for India's five meteorological regions. The horizontal line represents mean rainfall for that region. *Source*: The rainfall data are from the India Institute of Tropical Meteorology's Homogeneous Indian Monthly Rainfall Data Set (1871–2008). The figure is constructed based on the author's calculations.

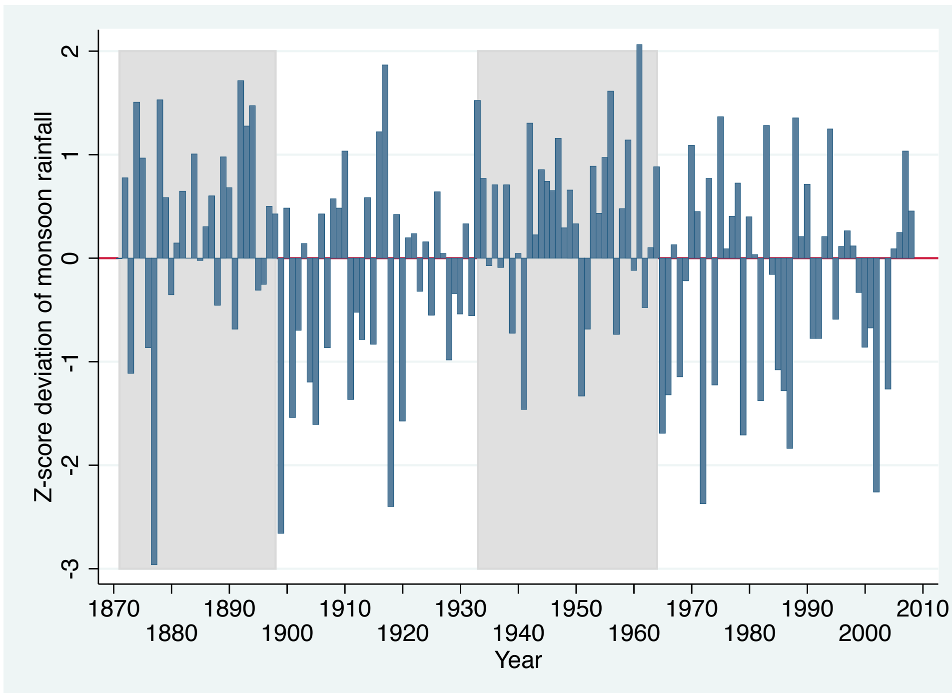

Figure 3: Annual Variability of the Indian Monsoon

Note: The y-axis graphs the All-India summer monsoon rainfall, expressed as a z-score deviation from its historical mean. *Source*: The rainfall data are from the India Institute of Tropical Meteorology's Indian Monthly Rainfall Data Set (1871–2008). The figure is constructed based on the author's calculations.

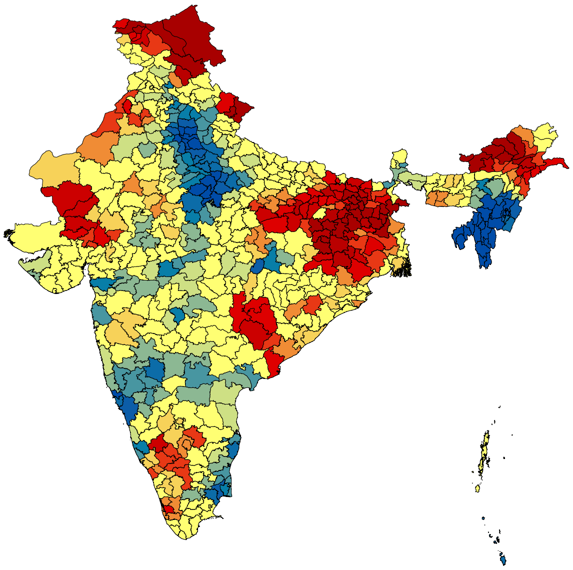

(a) 1971

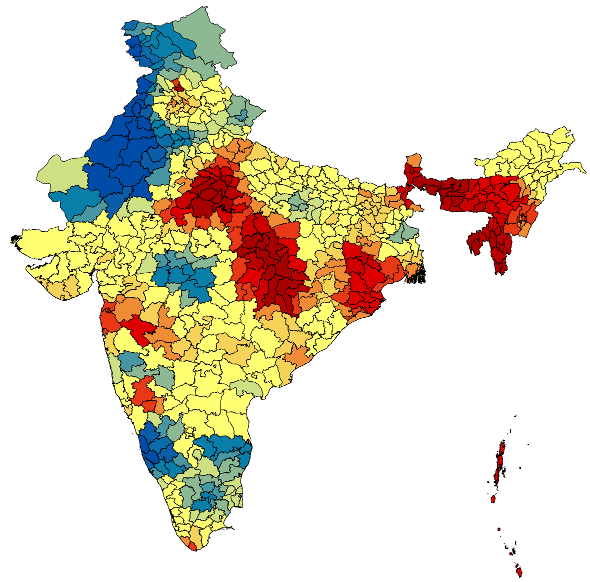

(b) 1982

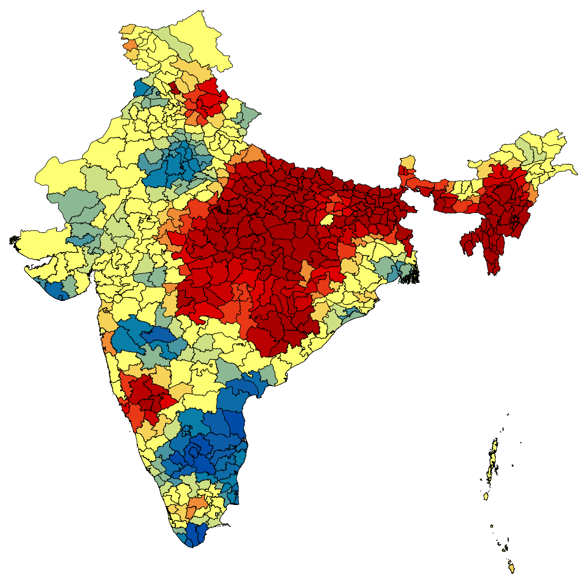

(c) 1999

Figure 4: Spatial Variation in Decadal Rainfall: REDS Survey Years

Note: The map displays average (z-score) summer rainfall for each district over the previous decade. Blue represents higher rainfall, and red represents lower rainfall. *Source*: The rainfall data are from the India Institute of Tropical Meteorology's Homogeneous Indian Monthly Rainfall Data Set (1871–2008). The figure is constructed based on the author's calculations.

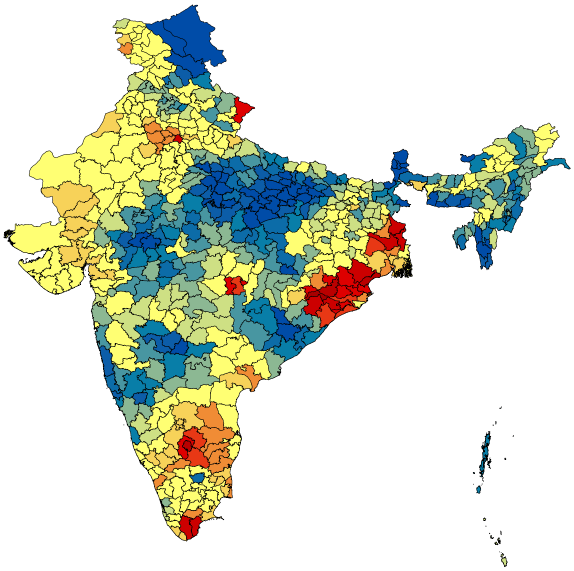

(a) 1956

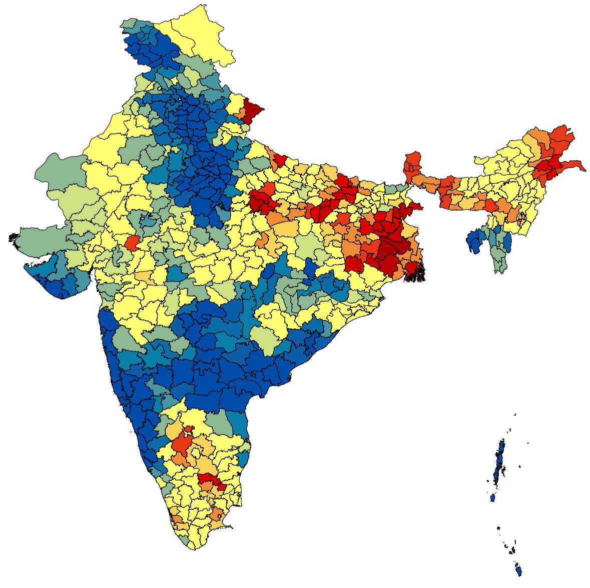

(b) 1966

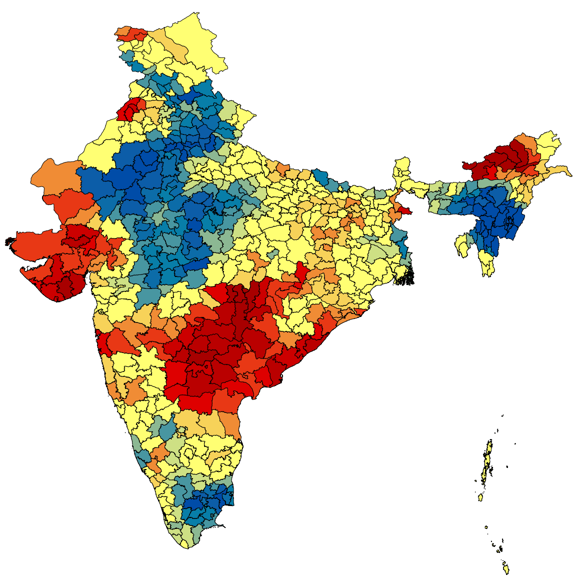

(c) 1976

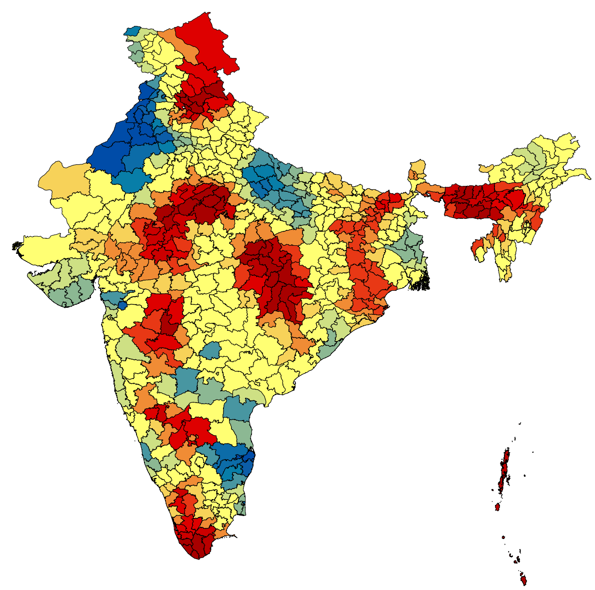

(d) 1986

Figure 5: Spatial Variation in Decadal Rainfall: WB Years

Note: The map displays average (z-score) summer rainfall for each district over the previous decade. Blue represents higher rainfall, and red represents lower rainfall. *Source*: The rainfall data are from the India Institute of Tropical Meteorology's Homogeneous Indian Monthly Rainfall Data Set (1871–2008). The figure is constructed based on the author's calculations.

Tables

Table 1: Summary Statistics

	Household			District		
	1971	1982	1999	1971	1956	1986
Panel A: Agricultural variables						
Agricultural profits per acre (1971 Rs.)	502.96	586.6	741.7	-	-	-
	(440.9)	(654.9)	(940.0)			
Agricultural profits per acre, deducting the value of family labor (1971 Rs.)	-	375.3	425.3	-	-	-
		(530.9)	(819.2)			
Agricultural revenue per acre	-	-	-	4425.6	1439.5	15340.0
				(2070.2)	(637.3)	(4796.9)
Proportion of land irrigated	0.378	0.414	0.483	0.234	0.178	0.321
	(0.437)	(0.455)	(0.466)	(0.203)	(0.175)	(0.256)
Irrigation investment during the recall period (dummy)	0.0767	0.0724	0.0116	-	-	-
	(0.266)	(0.259)	(0.107)			
Log non-land wealth (1971Rs.)	8.065	7.040	9.123	-	-	-
	(1.081)	(1.406)	(1.228)			
Proportion of inherited land irrigated	0.329	0.407	0.416			
	(0.380)	(0.456)	(0.468)			
Log non-land inherited wealth (1971Rs.)	7.133	2.959	5.848	-	-	-
	(0.962)	(2.789)	(3.690)			
Average crop water need (centimeters)	-	0.706	0.736	-	-	-
		(0.224)	(0.229)			
Average crop water need (centimeters) (monsoon crops)	-	0.754	0.820	-	-	-
		(0.239)	(0.249)			
Panel B: Weather variables						
Current year rainfall	0.313	0.208	0.279	0.436	0.579	-0.400
	(0.929)	(0.772)	(0.723)	(1.007)	(0.883)	(0.748)
Ten-year lagged average rainfall	-0.000634	0.0653	-0.0303	0.000608	0.108	-0.0353
	(0.328)	(0.251)	(0.326)	(0.288)	(0.294)	(0.234)
Ten-year lagged average of dry shock	0.196	0.183	0.166	0.203	0.176	0.191
	(0.125)	(0.0925)	(0.150)	(0.122)	(0.111)	(0.106)
Ten-year lagged average of wet shock	0.177	0.220	0.167	0.185	0.224	0.163
	(0.122)	(0.130)	(0.124)	(0.106)	(0.133)	(0.115)

Note: The table displays mean coefficients, with standard deviations in parentheses. The household sample is restricted to farmers who cultivate land. See Section 4 for details on how the variables are constructed.

Table 2: Daily Water Requirements of Common Crops Grown in India

Crop	Daily Water Requirement
Barley	0.400
Cotton	0.525
Oilseeds	0.350
Maize	0.450
Millet	0.575
Peanut	0.525
Potato	0.750
Pulses	0.350
Rice	1.075
Sorghum	0.575
Soybean	0.525
Sugarcane	0.650
Wheat	0.425

Note: The daily water requirement is measured in centimeters per day. *Source*: Agriinfo.in (2015).

Table 3: The Impacts of Irrigation and Rainfall on Profits

Data set:	Household	Household	District
Specification:	FE-IV	FE-IV	FE
Dependent variable:	Profits per Acre	Profit per Acre	Revenue per Acre
	(1)	(2)	(3)
Rainfall below 20th percentile (dummy)	10.67	-46.64	-471.6***
	(136.2)	(149.1)	(122.7)
Rainfall between 20th and 40th percentiles	87.95	82.35	-272.1**
	(93.76)	(101.6)	(121.5)
Rainfall between 60th and 80th percentiles	154.7*	89.57	108.0
	(81.28)	(87.40)	(106.5)
Rainfall above 80th percentile	312.4***	335.3***	127.9
	(82.46)	(85.01)	(114.3)
Proportion of irrigated land	364.1***	430.2***	3031.7***
	(126.1)	(141.5)	(900.3)
Propirr*Rainfall below 20th percentile	-226.0	-135.7	1001.0**
	(170.5)	(189.1)	(468.9)
Propirr*Rainfall between 20th and 40th percentiles	-250.3	-186.9	721.5
	(167.5)	(173.6)	(476.9)
Propirr*Rainfall between 60th and 80th percentiles	-154.3	-84.58	-295.1
	(139.1)	(150.6)	(390.1)
Propirr*Rainfall above 80th percentile	-447.8**	-462.8**	-227.4
	(198.5)	(223.3)	(397.6)
Temperature	-15.16	-31.42	-174.2***
	(32.44)	(39.76)	(46.58)
Log non-land wealth (1971 Rs.)	67.96	68.59	
	(58.62)	(63.35)	
Fixed effects	Household	Household	District
Year fixed effects	Yes	Yes	Yes
Observations	6828	6828	8384
First stage			
F statistic (Proportion of irrigated land)	92.51	92.51	
F statistic (Log non-land wealth)	19.96	19.96	

Notes: Standard errors, in parentheses below the coefficients, allow for clustering within a latitude–longitude grid point. Column 1 deducts the value of family labor from profits and column 2 does not. In columns 1 and 2, I instrument for the proportion of irrigated with the proportion of inherited land that was irrigated, and I instrument for wealth with inherited wealth. The first-stage F-statistics are reported in the table. Full first-stage regressions are also available from the author. F-test: The Staiger and Stock (1997) rule of thumb is that instruments are "weak" if the first-stage F is less than 10, and the Stock and Yogo (2002) Weak ID test critical value for 2SLS bias being less than 10% of OLS bias is 16.38. See Section 4 for details on how the variables are constructed.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Testing for Irrigation Adaptation

Data set:	Household	Household	Household	Household	District	District
Specification:	FE	FE-IV	FE	FE-IV	FE	FE
Dependent variable:	Irrigation Investment (Dummy)	Irrigation Investment (Dummy)	Irrigation Investment (Dummy)	Irrigation Investment (Dummy)	Log of the One-Year Change of Irrigated Area	Log of the One-Year Change of Irrigated Area
	(1)	(2)	(3)	(4)	(5)	(6)
Ten-year lagged average rainfall	-0.0501**	-0.0463*			-0.00739***	
	(0.0238)	(0.0238)			(0.00254)	
Ten-year lagged average of dry shock			0.111*	0.115**		0.0197***
			(0.0587)	(0.0554)		(0.00652)
Ten-year lagged average of wet shock			-0.0538	-0.0278		-0.00231
			(0.0543)	(0.0561)		(0.00421)
Current year rainfall	0.00728	0.00881	0.00559	0.00723	0.00273***	0.00276***
	(0.00680)	(0.00671)	(0.00696)	(0.00688)	(0.000900)	(0.000903)
One year lagged rainfall	-0.00368	-0.00660	-0.00349	-0.00675	-0.000166	-0.000348
	(0.00615)	(0.00674)	(0.00612)	(0.00676)	(0.00104)	(0.000978)
Log non-land wealth (1971 Rs.)		0.0481***		0.0472***		
		(0.0129)		(0.0127)		
Fixed effects	Household	Household	Household	Household	District	District
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
First stage						
F statistic (Log non-land wealth)		109.73		107.72		
Observations	12003	11759	12003	11759	8130	8130

Note: Standard errors, in parentheses below the coefficients, allow for clustering within a latitude–longitude grid point. A dry shock is defined rainfall below the 20th percentile and a wet shock is defined as rainfall above the 80th percentile. In columns 2 and 4, I instrument for wealth with inherited wealth. The first-stage F-statistics are reported in the table. Full first-stage regressions are also available upon request. F-test: The Staiger and Stock (1997) rule of thumb is that instruments are "weak" if the first-stage F is less than 10, and the Stock and Yogo (2002) Weak ID test critical value for 2SLS bias being less than 10% of OLS bias is 16.38. See Section 4 for details on how the variables are constructed.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Testing for Crop Adaptation

Data set:	Household	Household	Household	Household
Specification:	FE-IV	FE-IV	FE-IV	FE-IV
Dependent variable:	Crop Water Need	Crop Water Need	Crop Water Need (Monsoon)	Crop Water Need (Monsoon)
	(1)	(2)	(3)	(4)
Ten-year lagged average rainfall	0.0269		0.0379	
	(0.0233)		(0.0287)	
Ten-year lagged average of dry shock		-0.159***		-0.204***
		(0.0478)		(0.0675)
Ten-year lagged average of wet shock		0.0193		0.0834
		(0.0454)		(0.0572)
Current year rainfall	0.0145**	0.0159**	0.0145	0.0166*
	(0.00715)	(0.00686)	(0.0103)	(0.00981)
One year lagged rainfall	-0.0160**	-0.0186**	-0.0163*	-0.0215**
	(0.00800)	(0.00765)	(0.00945)	(0.00935)
Log non-land wealth (1971 Rs.)	-0.00428	-0.00522	-0.00368	-0.00512
	(0.0120)	(0.0109)	(0.0192)	(0.0171)
Fixed effects	Household	Household	Household	Household
Year fixed effects	Yes	Yes	Yes	Yes
First stage				
F statistic (Log non-land wealth)	69.22	71.49	66.94	69.38
Observations	5577	5577	5462	5462

Note: Standard errors, in parentheses below the coefficients, allow for clustering within a latitude-longitude grid point. A dry shock is rainfall below the 20th percentile and a wet shock is rainfall above the 80th percentile. In all columns, I instrument for wealth with inherited wealth. The first-stage F-statistics are reported in the table. Full first-stage regressions are also available from the author. F-test: The Staiger and Stock (1997) rule of thumb is that instruments are “weak” if the first-stage F is less than 10, and the Stock and Yogo (2002) Weak ID test critical value for 2SLS bias being less than 10% of OLS bias is 16.38. See Section 4 for details on how the variables are constructed. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Testing for Crop Adaptation: Individual Crops

Data set:	Household	Household	Household	Household	Household	Household	Household	Household
Specification:	FE-IV	FE-IV	FE-IV	FE-IV	FE-IV	FE-IV	FE-IV	FE-IV
Dependent variable:	Rice	Wheat	Pulses	Millet	Cotton	Sorghum	Peanut	Maize
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ten-year lagged average of dry shock	-0.209**	0.0401	0.125**	-0.0257	-0.0554	-0.0843	0.0724	0.0324
	(0.0814)	(0.0663)	(0.0540)	(0.0681)	(0.0456)	(0.0774)	(0.0522)	(0.0405)
Ten-year lagged average of wet shock	0.0789	0.0726*	-0.0172	-0.155**	0.133**	-0.152***	0.0464	-0.00858
	(0.0778)	(0.0431)	(0.0448)	(0.0734)	(0.0604)	(0.0543)	(0.0419)	(0.0362)
Current year rainfall	0.0191	-0.00764	-0.0217*	0.0150	-0.00655	0.00498	-0.00212	-0.000965
	(0.0122)	(0.00984)	(0.0112)	(0.0148)	(0.00629)	(0.0115)	(0.00893)	(0.00595)
One year lagged rainfall	-0.0344***	0.00427	0.00682	0.00666	-0.0113*	0.00343	0.00930	0.00234
	(0.0129)	(0.00608)	(0.00673)	(0.00780)	(0.00590)	(0.00563)	(0.00728)	(0.00399)
Log non-land wealth (1971 Rs.)	-0.00625	0.00616	-0.00346	0.00523	-0.00395	0.0191**	-0.00921	-0.00725
	(0.0197)	(0.0119)	(0.0112)	(0.0158)	(0.0121)	(0.00970)	(0.00821)	(0.00641)
Fixed effects	Household	Household	Household	Household	Household	Household	Household	Household
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
First stage								
F statistic (Log non-land wealth)	70.50	70.50	70.50	70.50	70.50	70.50	70.50	70.50
Observations	5622	5622	5622	5622	5622	5622	5622	5622

Note: Standard errors, in parentheses below the coefficients, allow for clustering within a latitude-longitude grid point. The dependent variable in each column is the proportion of total cultivated area planted with that crop. A dry shock is rainfall below the 20th percentile and a wet shock is rainfall above the 80th percentile. In all columns, I instrument for wealth with inherited wealth. The first-stage F-statistics are reported in the table. Full first-stage regressions are also available from the author. F-test: The Staiger and Stock (1997) rule of thumb is that instruments are "weak" if the first-stage F is less than 10, and the Stock and Yogo (2002) Weak ID test critical value for 2SLS bias being less than 10% of OLS bias is 16.38. See Section 4 for details on how the variables are constructed. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Testing for Irrigation Adaptation: Dependent Variable is the Proportion of Irrigated Land

Data set: Specification:	Household FE	Household FE-IV	Household FE	Household FE-IV
Dependent variable:	Proportion of Irrigated Land	Proportion of Irrigated Land	Proportion of Irrigated Land	Proportion of Irrigated Land
	(1)	(2)	(3)	(4)
Ten-year lagged average rainfall	-0.0659**	-0.0564*		
	(0.0325)	(0.0330)		
Ten-year lagged average of dry shock			0.0139	0.0144
			(0.0735)	(0.0756)
Ten-year lagged average of wet shock			-0.102	-0.0570
			(0.0675)	(0.0708)
Current year rainfall	-0.0202	-0.0184	-0.0220	-0.0200
	(0.0131)	(0.0137)	(0.0134)	(0.0140)
One year lagged rainfall	-0.00852	-0.0132	-0.0110	-0.0161
	(0.00983)	(0.0106)	(0.00940)	(0.0103)
Log non-land wealth (1971 Rs.)		0.0736***		0.0715***
		(0.0162)		(0.0163)
Fixed effects	Household	Household	Household	Household
Year fixed effects	Yes	Yes	Yes	Yes
First stage				
F statistic (Log non-land wealth)		116.51		113.72
Observations	11856	11759	11856	11759

Notes: Standard errors, in parentheses below the coefficients, allow for clustering within a latitude-longitude grid point. A dry shock is rainfall below the 20th percentile and a wet shock is rainfall above the 80th percentile. In columns 2 and 4, I instrument for wealth with inherited wealth. The first-stage F-statistics are reported in the table. Full first-stage regressions are also available from the author. F-test: The Staiger and Stock (1997) rule of thumb is that instruments are "weak" if the first-stage F is less than 10, and the Stock and Yogo (2002) Weak ID test critical value for 2SLS bias being less than 10% of OLS bias is 16.38. See Section 4 for details on how the variables are constructed.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Adaptation to Climate Change: Historical Evidence from the Indian Monsoon (Supplementary File)

Vis Taraz

July 27, 2016

A Study Area and Attrition Issues

Figure 1 shows the districts that are included in the World Bank data set and Figure 2 shows the locations of the villages that are included in the REDS data set.

In my analysis of the REDS data set, I restrict my analysis to households that were interviewed in at least two of the survey rounds. I now discuss potential sample selection issues that arise due to attrition. First, it is important to note that if households split into multiple households between survey rounds, the split-off households are included in the sample and therefore not dropped from the sample. Consequently, the only cause of attrition is if a household could not be located in the follow-up survey round or if a household refused to participate in the follow-up round of the survey. Because of the long time periods between the different rounds (1971, 1982, and 1999), attrition is fairly high. Of the 1971 households, 69% were located and participated in the 1982 survey round (Vashishtha, 1989). Of the 1982 households, 77% were located and participated in the 1999 survey round (Deininger et al., 2009b). The survey design attempts to partially address this attrition by adding an equal number of randomly selected households from the same village to balance out any households that could not be located in the follow-up survey round. These randomly selected households are stratified by wealth categories, in order to better match the characteristics of the attrited households (Deininger et al., 2009a). Although households added this way in 1999 are not included in my sample (because they were only interviewed during a single round), households that were added in 1982 are included in my sample if they were interviewed again in 1999.

Earlier work has found that attrition appears to be random relative to many household demographic characteristics. Specifically, one cannot reject the hypothesis that the demographic attributes of caste, household size, number of earners, age of the household head, and education of the household head are the same amongst attriting and non-attriting households (Deininger et al., 2009a,b; Jha et al., 2013; Jha, 2013). Unfortunately, earlier work has found that attrition is non-random relative to land size. Specifically, landless households are more likely to attrite (Deininger et al., 2009b). Furthermore, comparing attriting to non-attriting households (for both the 1971 and 1982 rounds), I find that non-attriting households had statistically significantly higher land areas, higher fractions of irrigated land, and higher levels of non-land wealth. Therefore, the sample that I analyze is not nationally representative of the Indian population as a whole. For this reason, my results are an accurate representation of the behavior of the population that my sample represents, but are not necessarily an accurate characterization of all farmers in India. Nevertheless, as I argue in Section 8 of the main paper, I believe that my core results are still of interest. Specifically, when I analyze adaptation, I find that adaptation appears to recapture a relatively small fraction of the profits that are lost due to sustained adverse climate (at most, 19% of profits; most likely only 9%). If we believe that wealthier farmers are more likely to adapt than poorer farmers (and have less barriers to adaptation), then this would suggest that my estimate is actually an upper bound on the efficacy of adaptation and, hence, still of interest.

B Rainfall Variation

In this section, I provide a more detailed discussion of the variation in my key lagged rainfall variables. The basic means and standard deviations of my lagged rainfall measures, for relevant years of each survey, are given in the main text in panel B of Table 1. In Appendix Table 1, I present a more detailed summary of statistics for the decadal lagged z-score rainfall measure, specifically showing its percentile distribution.[1]

Appendix Table 1, however, only captures the temporal variation in decadal rainfall and does not capture any of the spatial variation. To demonstrate the spatial variation of lagged rainfall, I use maps. Figures 4 and 5 in the main text show the spatial variation in decadal lagged average rainfall—for each round of REDS survey and each decade of the WB survey. To complement these maps, Appendix Figures 3 and 4 present district maps of the number of lagged dry shocks in the past decade, for representative years of the REDS and WB data sets. Recall that a dry shock is defined as a year that is below the historical 20th percentile of rainfall for that district.[2] On the maps, the districts shaded in the lightest color of pink had no dry shocks in the past decade, whereas the districts shaded in the darkest red had five or more dry shocks in the past decade.[3] Focusing on the maps for the REDS survey years, we can see that there is variation in the number of dry shocks in the past decade, with a substantial number of districts having no dry shocks in the past decade and, conversely, a substantial number having five dry shocks or more in the past decade.

These maps provide a qualitative description of the spatial variation in lagged rainfall, but do not capture these spatial variation in a precise, quantitative manner. To explore the spatial variations of lagged rainfall in a complementary, more rigorous way, I construct graphs that display the spatial autocorrelation of lagged decadal average rainfall. Specifically, I calculate Moran's *I*, a measure of spatial autocorrelation, which is given by the following formula:

$$I = \frac{N}{\Sigma_i \Sigma_j w_{ij}} \frac{\Sigma_i \Sigma_j w_{ij}(X_i - \bar{X})(X_j - \bar{X})}{\Sigma_i (X_i - \bar{X})^2},$$

where N is the number of spatial units indexed by i and j, X is the variable of interest; $\bar{X}$ is

[1]Note: the means for the REDS data are slightly different here than in Table 1 in the main paper, because the main paper Table takes the household as the unit of observation whereas the Appendix Table uses the village as the unit of observation.

[2]I use the entire range of the rainfall data (1900–2008) to construct the percentiles.

[3]For concision, I do not present the corresponding maps with the number of lagged wet shocks, but they are available upon request.

the mean of X, and w_{ij} is an element of a matrix of spatial weights (Moran, 1950). The value of I ranges from –1 to 1, where negative values indicate negative spatial autocorrelation, positive values indicate positive spatial autocorrelation, and a zero value indicates a random spatial pattern. Due to the large-scale nature of the monsoon decadal variations, I expect to see a positive autocorrelation of lagged decadal average rainfall. Calculating Moran's I allows me to measure the magnitude of this autocorrelation, and also to see how much it decays with distance.

Figures 5 and 6 plot the autocorrelation coefficient as a function of distance between pairs of observations, for each year of the REDS survey, and for representative years from the WB data, respectively.[4] The graphs were created in Stata using the "spatcorr" command (Pisati, 2012). The spatial correlation coefficient weights each pair of points as an inverse linear function of the distance between the two points. The graphs indicate that spatial autocorrelation is very high for points that are within 200 kilometers of each other (ranging from 0.5 correlation to 0.95 correlation, depending on the year). However, the autocorrelation decays roughly linearly as a function of distance, reaching a correlation of close to zero once the distance between two points is roughly 600 or 800 kilometers (again, depending on the year). I use this information about the spatial autocorrelation to guide my analysis in Appendix C.1, where I test the robustness of my results to the use of standard errors that allow for spatial correlation. In AppendixC.1 , I use an 800-kilometer radius for the potential spatial correlation and allow it to decay at a linear rate, in accordance with the pattern of spatial correlation that the Moran's I graphs have indicated. The results of this exercise are discussed in Appendix C.1.

C Robustness

In this section, I investigate the robustness of my results to several changes in the regression specification.

[4] For the REDS data set, I use each village in the data set as the unit of analysis for constructing Moran's I. For the WB data set, I use each district in the data set as the unit of analysis for constructing Moran's I.

C.1 Non-IV Specification and Spatially Correlated Standard Errors

I first re-run my adaptation regressions using a fixed effects specification, in which I do not instrument for household wealth. I report this specification with two sets of standard errors. The first set of standard errors is clustered at the rainfall grid point level, in keeping with the specifications in the main paper. The second set of standard errors allow for spatial correlation. I present these specifications for two reasons. First, because there are some potential issues with my instrumental variables strategy, running the regression without the IV may be informative for comparative purposes. Second, as discussed in Section B my lagged rainfall measure exhibits substantial spatial correlation. Therefore, it is important to make sure that my results are robust to a specification that allows for this spatial correlation. To implement these standard errors, I use code from Hsiang and Solow (2010) and Fetzer (2014). Guided by the autocorrelograms in Section B, I allow for a spatial correlation within 800 kilometers with a Bartlett (triangular) kernel and temporal correlation within a 30-year window. I do not implement spatially correlated standard errors in the specifications presented in the main paper because I was not able to find Stata code that could calculate spatially correlated standard errors for a regression using instrumental variables and fixed effects.

I present the results of these specifications in Appendix Tables 2 and 3, which are the analogues of Tables 4 and 5 from the main text. The tables present the fixed effects specification in all columns. Cluster robust standard errors are presented below the coefficient estimates in parentheses, and spatially correlated standard errors are presented below that in brackets.

There are several important things to note. First, in comparing the fixed effects regressions to the IV fixed effects regressions from the main paper, we see that the coefficients on wealth change substantially. This indicates that wealth is indeed likely to be endogenous in theses regressions. However, the coefficients of interest on lagged rainfall are essentially unchanged across the IV and the non-IV specifications. This suggests that the endogeneity of wealth may not be affecting my estimates of adaptation.

Next, I compare the cluster robust standard errors to the errors that allow for spatial correlation. In all cases, the coefficients of interest that are statistically significant in the regressions from the main paper are still significant with the new set of standard errors. Looking at the district irrigation regression results in columns 5 and 6 of Appendix Table 2, we see that the standard errors increase by 15% to 50% but the coefficients on lagged rainfall are still statistically significant. This demonstrates that there is non-negligible spatial correlation in the errors for my district regressions, but that my results are robust to allowing for this spatial correlation.

Interestingly, for the household irrigation and crop regressions, the standard errors change very little when I allow for spatial correlation and, in some cases, become slightly smaller. As discussed in Cameron and Miller (2015), situations in which clustered errors are smaller are typically either due to negatively correlated standard errors or, more frequently, simply due to noise. I believe the latter is more likely to be the case for my data. The small change in the standard errors demonstrates that there is not substantial correlation in the standard errors of my household regressions. This suggests that there may be substantially more idiosyncratic variability from household to household in the REDS data set (compared to the districts in the WB data set), and this may be why the standard errors change so little.

C.2 Sensitivity to Rainfall Specification

My baseline regressions use average rainfall from the past decade as an approximation of the current monsoon regime. I now verify that the choice of a 10-year rainfall window is not driving my results. First, I explore a specification where I control individually for rainfall from each year from the past decade (rather than using the decadal average). I do this to test if farmers give greater weight to more recent rainfall when making their irrigation and crop decisions. The results, presented in Appendix Table 4, are inconclusive. Very few of the individual rainfall lags are (individually) statistically significant. Furthermore, there is no discernible pattern in the magnitudes of the coefficients, probably due to the lack of precision with which the coefficients are estimated. It is likely that the single-year rainfall lags are not individually significant because they are collinear with each other—due precisely to the persistent nature of the monsoon decadal variation. However, in all columns except column 2, the lags from years 2 to 10 are jointly significant at the 5% level or higher .

Next, I reestimate my adaptation regressions using rainfall lag windows of 5 and 15 years. Appendix Table 5 presents the irrigation results for both data sets. The district irrigation results are robust to this change in specification. In household irrigation regressions, the coefficient signs are unchanged but the coefficients (in three out of four columns) are no longer statistically significant.[5] Appendix Table 6 presents the household crop adaptation regressions with five and fifteen year rainfall lag windows. For the specifications that use average decadal rainfall, the signs of the coefficients are unchanged, but the coefficients are no longer statistically significant. For the specifications that use the wet/dry shock measure, however, the results remain statistically significant. Taken as a whole, I interpret that Appendix Tables 5 and 6 demonstrate that my household irrigation results are not robust to the use of alternate rainfall windows, but that my district irrigation results and my crop results are robust to this change.

C.3 Depletion of Groundwater and Surface Water Supplies

Thus far, I have interpreted the response of irrigation investments to lagged rainfall as evidence of adaptation. However, other mechanisms are possible. For example, suppose that a dry decade reduces groundwater and/or surface water availability. Lower groundwater levels could induce farmers to deepen their existing wells or perhaps switch to an investment in surface water irrigation. Conversely, a decrease in surface water resources might prompt farmers to invest in wells.

In an attempt to address this concern, I adjust my regressions to use irrigated area, rather than irrigation investment *per se*, as the dependent variable. The dependent variable in my baseline district irrigation regressions is already the log of the 1-year change in irrigated area (see columns 5 and 6 of Table 4 in the main text). We would expect diminished groundwater (or surface water) supplies to reduce, not increase, the irrigated area, so this regression specification would seem to rule out a water depletion story, and hence I do not modify it. For my household data, my core

[5] For concision, I do not report columns for the household irrigation adaptation regressions that control for wealth, but the results are comparable to those shown here.

specification uses irrigation investment over the recall period as the dependent variable. I like this measure because it captures precisely how the household is adjusting its irrigation each year. Nevertheless, in Table 7 in the main text, I present an alternate specification with the proportion of irrigated land as the dependent variable.[6] The results are consistent with my baseline results, although the coefficients are estimated with less precision. Specifically, the proportion of irrigated land increases after drier decades. Note that the proportion of irrigation land is a coarser measure of adaptation, since it is a stock measure that is the sum of all irrigation investments made to date. In contrast, irrigation investment (from the past year) is a flow variable and one which we would expect to respond more readily to recent rainfall. This distinction may cause the reduction in precision of the coefficients in this table.

However, it is important to note that the strategy of using irrigated area as the dependent variable does not fully rule out the possibility that water depletion is driving my results. For example, depletion of surface water supplies might prompt a farmer to switch to an investment in a well, and this newly dug well might allow the farmer to irrigate a larger area than before. Unfortunately, due to data limitations, I am not able to fully rule out this possible mechanism.

C.4 Government Investment in Irrigation

In India, private entities (farmers) and public entities (the government) both invest in irrigation. I now test whether government investment, rather than farmer investment, might be driving my results. In India, most direct investment in groundwater irrigation is private, whereas most surface water investments are public (Shah, 1993). The bulk of the government's investments are large dams (Thakkar, 1999; Vaidyanathan, 2010).[7] Hence, large dams are a good measure of the government's direct irrigation investment.[8] I use data from the World Registry of Large Dams, which

[6] Since the gaps between household survey rounds last 10 years or more, I am unable to analyze 1-year changes in irrigation.

[7] A large dam is defined as a dam that has a height of 15 meters from the foundation or a reservoir capacity of more than 3 million cubic meters (Thakkar, 1999, p. 103).

[8] Although most of the government's direct irrigation investment is via large-scale dams, the government does subsidize groundwater irrigation through credit programs and electricity subsidies. Controlling for dams does not rule out this mechanism. The government also subsidizes electricity, which is a complement to groundwater irrigation. I

lists all large dams in India, by district and year.[9] Dams provide surface water supplies to downstream districts, so I control for the number of upstream dams as my measure the government's irrigation investment. The results, shown in Appendix Table 7, are consistent with my baseline results, suggesting that my baseline results are not solely driven by public irrigation investments. However, it is important to note that the government's irrigation investment is an outcome variable and may be endogenous to the household investment decision. For this reason, this regressions must be seen as suggestive or descriptive, but not conclusive.

C.5 Changes in Agricultural Technology and Policies

Next, I explore whether changes in agricultural technology or policies might be confounding my results. A major change in technology during this period was the Green Revolution, which introduced high-yielding varieties (HYVs) of rice and wheat. Irrigation and HYVs are complements (McKinsey and Evenson, 1999), so controlling for their availability is important. Furthermore, HYV seeds are better suited to certain agro-climatic zones, causing regional variation in Green Revolution impacts (Evenson, 2003) that will not be captured by my year fixed effects. In addition, the government subsidizes many agricultural inputs, offers extension services, and intervenes in output prices (Gulati, 1989; Fan et al., 2000, 2008). Many of these policies are implemented at the state level and will not be controlled for by year fixed effects(Fan et al., 2000, 2008; Birner et al., 2011). Therefore, I explore controlling for agricultural technology and policies directly.

With the district data, I control for electrification rates, fertilizer prices, and HYV suitability. I measure electrification rates as the percentage of electrified villages in each state, using data from Rud (2012). I draw the prices of nitrogen, phosphorus and potassium fertilizers from Sanghi et al. (1998).[10] To control for HYV suitability, I use Foster and Rosenzweig (2003)'s strategy by ex-

discuss electricity subsidies in the next section.

[9]The World Registry of Large Dams is analyzed by Pande and Duflo (2007). The data is publicly available at http://hdl.handle.net/1902.1/IOJHHXOOLZ (Duflo and Pande, 2006).

[10]Fertilizer prices are plausibly exogenous because they are determined at the national level; the only cross-sectional price variation arises from the cost of transportation from the railhead to the field (Sanghi et al., 1998).

ploiting variation in the timing of wheat and rice yield advances.[11,12] I proxy for HYV suitability with the proportions of wheat and rice planted in the first year of the survey, interacted with year dummies. This captures which districts were initially more likely to plant HYVs, and which districts became more likely to do so as the technology progressed. HYV seeds were also specifically promoted in some districts as part of the Intensive Agricultural District Program (IADP).[13] I include an IADP dummy, interacted with the year dummies, to capture the higher HYV usage in these districts and the evolution of this usage over time.

Columns 3 and 4 in Appendix Table 8 present the district irrigation results, using the wet/dry rainfall shocks specification. Column 3 controls for fertilizer prices and HYV suitability, with data spanning from 1956 to 1986. In column 4, I add the electrification control, which truncates the panel to 1965–1984. The results are robust to the addition of these controls.[14]

I include a broader array of controls for the household data. For the regressions that use all three of the survey rounds, I control for village financial institutions (credit cooperatives, moneylenders, and/or banks), the presence of agricultural extension services, and a dummy for whether the village is electrified. I also use the HYV suitability measures discussed above. In the regressions that use only the last two survey rounds, I control for village-level measures of transportation infrastructure, government intervention in output markets, government irrigation assets and subsidies, and the proportion of groundwater versus surface water irrigation.[15] I include the last control because surface water is more likely to be publicly funded.

Columns 1, 2, 5, and 6 of Appendix Table 8 present the household results. The household

[11] I use HYV suitability, rather than the area planted with HYVs because the HYV area is endogenous.

[12] Specifically, advances in wheat seeds preceded those for rice seeds and the agro-climatic suitability for growing rice versus wheat varies across districts.

[13] The IADP was initiated in the late 1960s in one district in each Indian state to diffuse technical know-how, credit and agricultural technology to accelerate the adoption of the HYVs.

[14] For concision, I do not report results using the average rainfall regression specification, but the results are consistent with those results presented.

[15] I measure transportation infrastructure as the distance to the nearest blacktop road, the nearest bus stand, and the nearest railroad station, as well as dummies for whether the roads to the bus stand and railroad station are blacktop roads. Government intervention in output markets is a dummy for whether or not most of the village's produce is sold to government agencies. Government irrigation assets and subsidies are the number of government irrigation sources (defined as the total number of government-owned tanks, wells, pumps and other irrigation assets) and as a dummy for the presence of public irrigation subsidies or loans.

results are robust to the addition of these controls.[16]

There are two important limitations of this robustness exercise. First, due to data limitations, I only control for a subset of the technology and policy changes that might influence the decisions of farmers. Second, the additional controls are all potentially endogenous to the irrigation and crop choice decisions of the farmers. Therefore, the results must be interpreted as suggestive and descriptive only. It appears that changes in policies and technology are not the mechanisms through which rainfall shocks affect irrigation and crop choice, but I cannot definitively rule out that possibility.

C.6 Controls for Region-by-Year Fixed Effects

In a final robustness exercise, I attempt to control for possible confounding factors more flexibly. The large-scale spatial correlation of the decadal shifts of the monsoon raises the concern that my baseline results could be induced by confounding factors. For example, irrigation investment may increase in a certain region of the country following a decade with below-average rainfall. But perhaps that same region instituted a policy that made irrigation investment more attractive. In this case, this policy itself might have been what drove the increase in irrigation, and my adaptation result would be a spurious correlation.

Previously, I attempted to address this problem by controlling for some specific changes in agricultural policies and technology to verify that these changes were not driving my results. Now, in an attempt to control for unobserved confounders more flexibly, I add region-by-year fixed effects to my regressions.[17] This allows me to control for any time-varying unobservables that are common to a given region. These results are shown in Appendix Tables 9 and 10.

For the irrigation tables, the REDS regressions are no longer significant, although the signs of the coefficients are preserved. The district irrigation regressions are still significant. For the crop adaptation regressions, the coefficients are still significant; in fact, the results are stronger.

[16]For concision, I do not report specifications that use the average rainfall regression specification, but the results are consistent with those presented.

[17]I use the six meteorological regions of India.

Specifically, in the baseline specification, only the dry shocks are significant; however, once I add region-by-year fixed effects, the lagged z-score columns are also significant.

In addition, I estimate versions of the adaptation regressions that include state-by-year fixed effects. This allows me to control for any time-varying unobservables that are common to a given state.[18] With these added controls, neither the household nor the district irrigation regressions are statistically significant, although the signs of the coefficients on lagged rainfall are preserved. My crop regressions, however, are robust to these controls, and the coefficients remain statistically significant at the 1% level in 3 out of the 4 columns.

Taken together, these results indicate that my crop choice regressions are substantially more robust to controlling for unobserved confounders than my irrigation results. The lack of robustness of my irrigation results to the region-by-year and state-by-year controls is a major limitation of my study.

[18]For concision, these tables are not reported, but they are available from the author upon request.

List of Figures

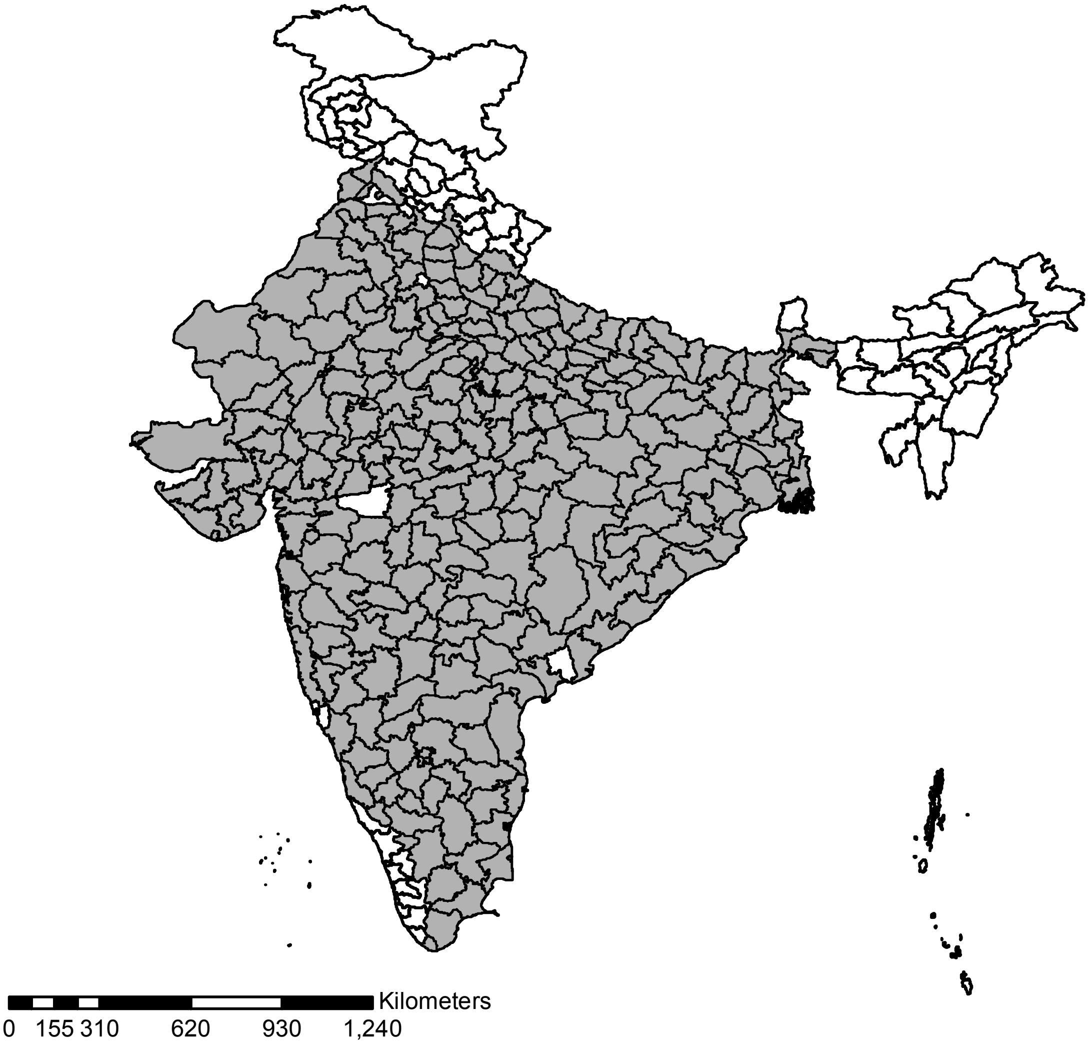

Figure 1: Districts Included in the World Bank Dataset

Note: The districts included in the World Bank data set are shaded gray.

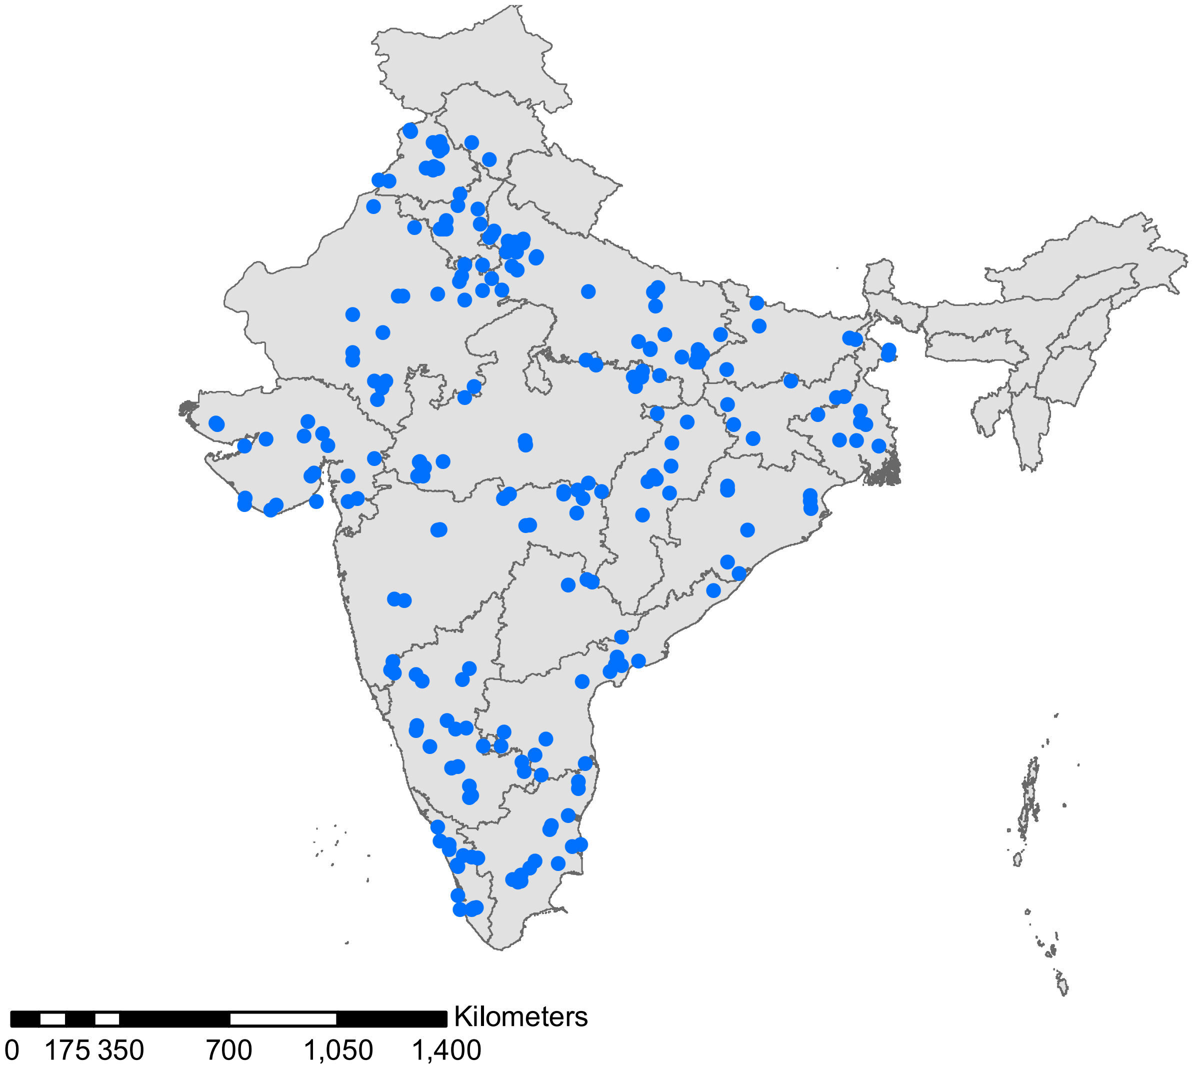

Figure 2: Villages Included in the Household Dataset

Note: The points on the map show the locations of the villages included in the REDS data set.

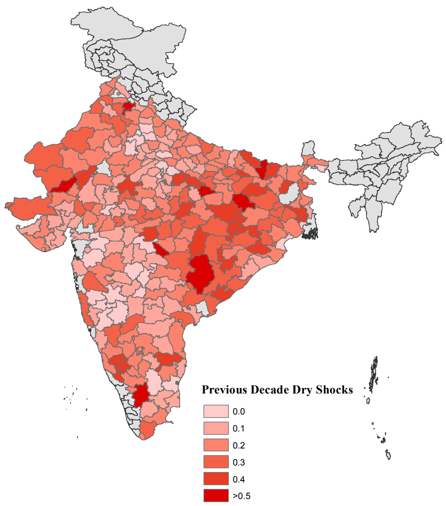

(a) 1971

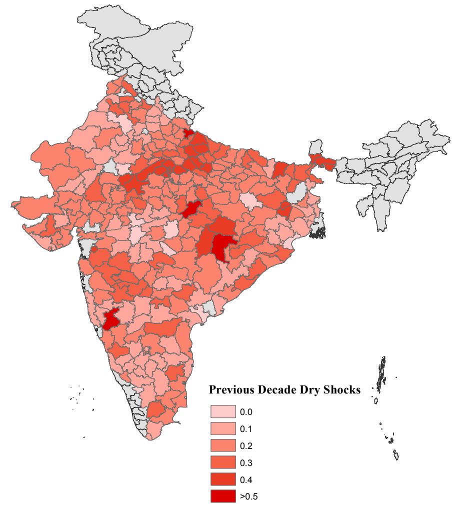

(b) 1982

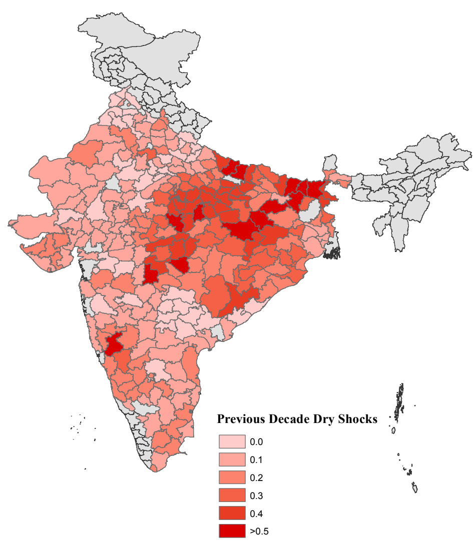

(c) 1999

Figure 3: Spatial Variation in Decadal Dry Rainfall Shocks: REDS Survey Years

Note: A dry shock is defined as a year that is below the historical 20th percentile of rainfall for that district. The districts shaded in the lightest color of pink had no dry shocks in the past, whereas the districts shaded in the darkest red had five or more dry shocks in the past decade. The years shown correspond to the survey years of the REDS dataset.

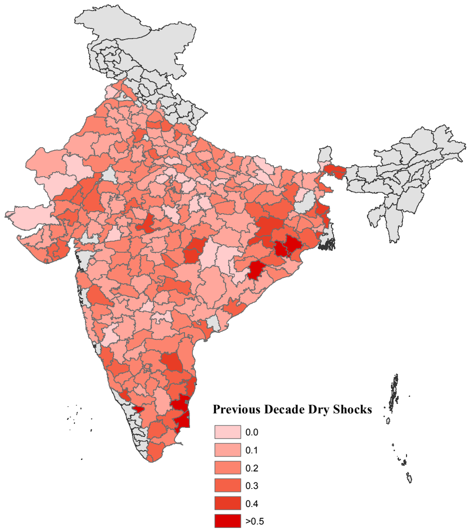

(a) 1956

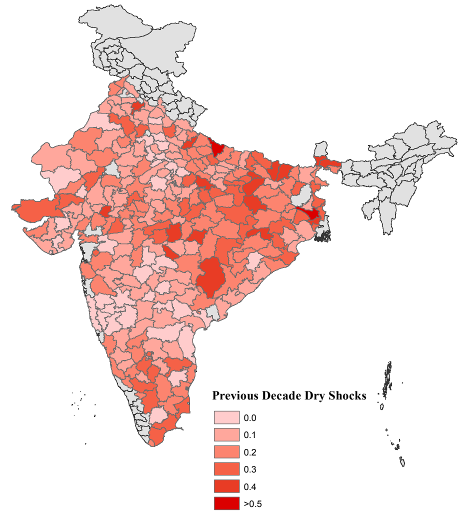

(b) 1966

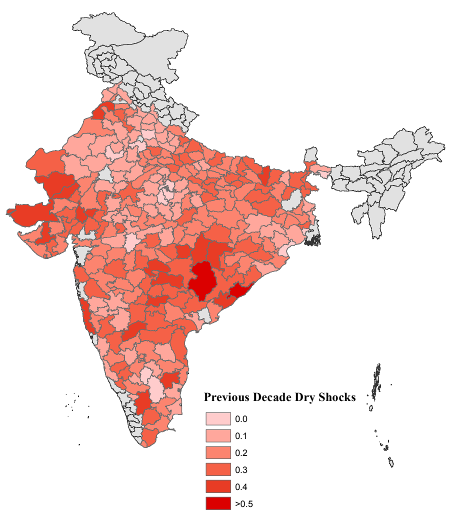

(c) 1976

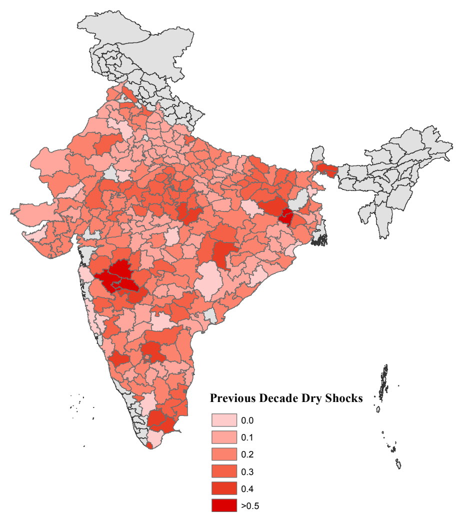

(d) 1986

Figure 4: Spatial Variation in Decadal Dry Rainfall Shocks: WB Years

Note: A dry shock is defined as a year that is below the historical 20th percentile of rainfall for that district. The districts shaded in the lightest color of pink had no dry shocks in the past, whereas the districts shaded in the darkest red had five or more dry shocks in the past decade. The years shown correspond to the time range of the WB dataset.

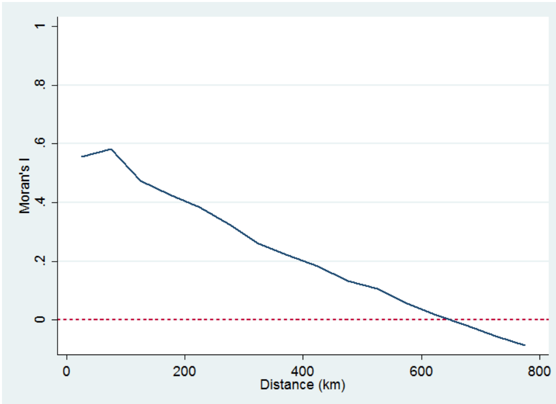

(a) 1971

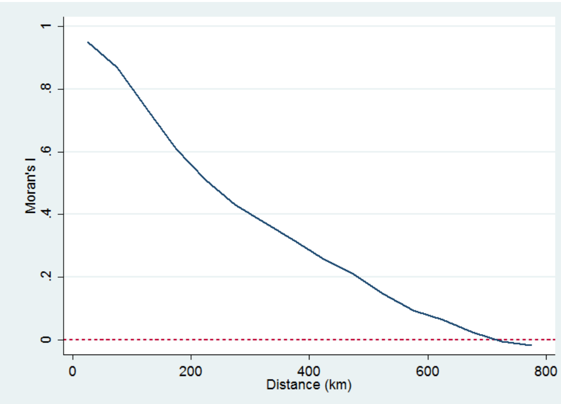

(b) 1982

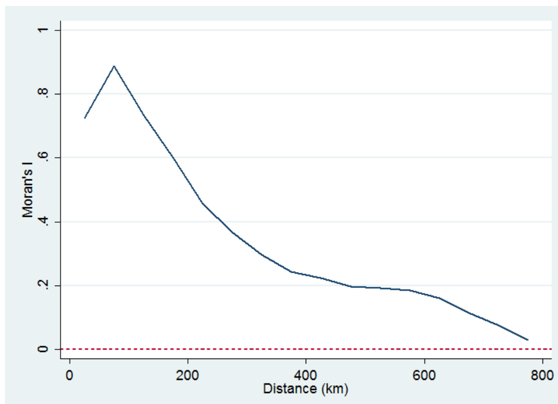

(c) 1999

Figure 5: Spatial Correlation of Lagged Rainfall: Moran's I for the REDS Survey Years

Note: These graphs plot the autocorrelation coefficient (Moran's I) as a function of distance between coefficient as a function of distance between pairs of observations, for each year of the REDS survey, using weights that are a linear function of distance. The graphs were created in Stata using the "spatcorr" command (Pisati, 2012). Refer to the text for more details about Moran's I.

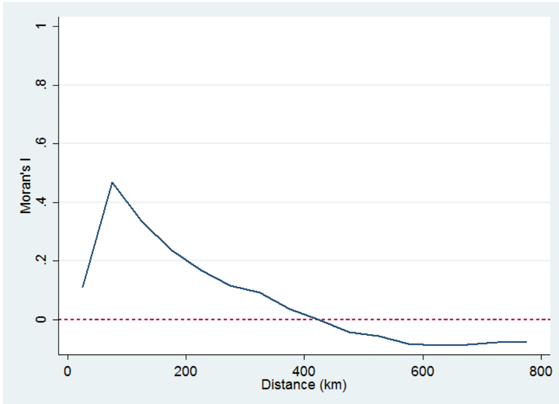

(a) 1956

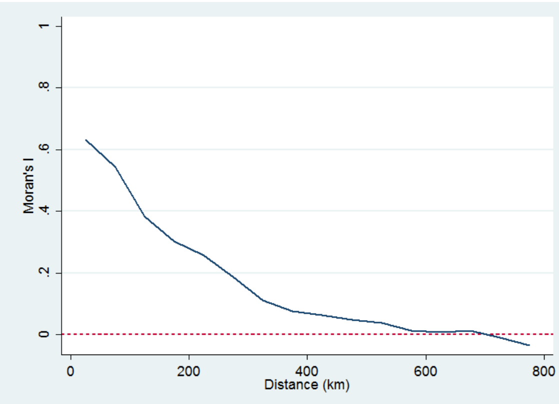

(b) 1966

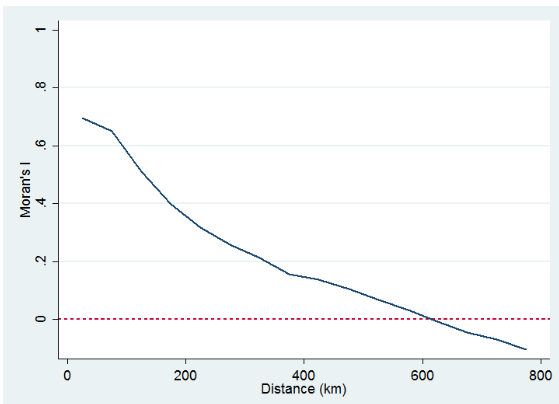

(c) 1976

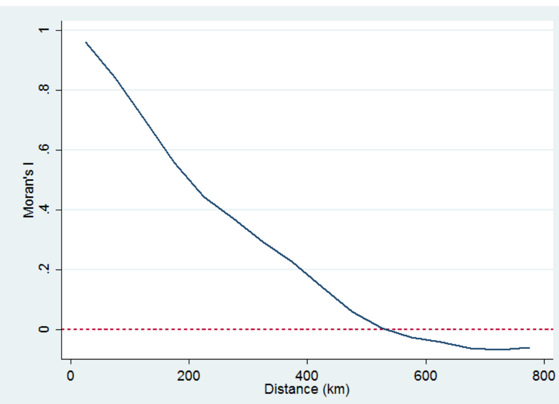

(d) 1986

Figure 6: Spatial Correlation of Lagged Rainfall: Moran's I for WB Dataset Years

Note: These graphs plot the autocorrelation coefficient (Moran's *I*) as a function of distance between coefficient as a function of distance between pairs of observations, for representative years from the WB data, using weights that are a linear function of distance. The graphs were created in Stata using the "spatcorr" command (Pisati, 2012). Refer to the text for more details about Moran's *I*.

List of Tables

Table 1: Detailed Summary Statistics for Lagged Rainfall

Year	Mean	Std. Dev.	10th percentile	25th percentile	50th percentile	75th percentile	90th percentile
1971	-0.037	0.292	-0.389	-0.240	-0.090	0.191	0.360
1982	0.088	0.273	-0.312	-0.068	0.100	0.294	0.396
1999	0.055	0.331	-0.407	-0.194	0.089	0.345	0.458
1956	0.108	0.294	-0.292	-0.105	0.105	0.315	0.484
1966	0.128	0.340	-0.301	-0.098	0.093	0.335	0.576
1976	0.026	0.262	-0.298	-0.162	0.018	0.229	0.375
1986	-0.035	0.234	-0.322	-0.182	-0.044	0.136	0.242

Note: The table displays summary statistics for lagged average z-score rainfall from the previous decade for each year of the REDS data set (consisting of 259 villages), and for four different points of the WB data set (consisting of 271 districts).

Table 2: Testing for Irrigation Adaptation, FE and FE with Spatial Standard Errors

Data set:	Household	Household	Household	Household	District	District
Specification:	FE	FE	FE	FE	FE	FE
Dependent variable:	Irrigation Investment (Dummy)	Irrigation Investment (Dummy)	Irrigation Investment (Dummy)	Irrigation Investment (Dummy)	Log of the One-Year Change of Irrigated Area	Log of the One-Year Change of Irrigated Area
	(1)	(2)	(3)	(4)	(5)	(6)
Ten-year lagged average rainfall	-0.0501	-0.0497			-0.00739	
	(0.0239)	(0.0239)			(0.00254)	
	[0.0225]	[0.0228]			[0.00384]	
Ten-year lagged average of dry shock			0.111	0.115		0.0197
			(0.0589)	(0.0566)		(0.00652)
			[0.0508]	[0.0508]		[0.00772]
Ten-year lagged average of wet shock			-0.0538	-0.0435		-0.00231
			(0.0545)	(0.0550)		(0.00421)
			[0.0402]	[0.0401]		[0.00698]
Current year rainfall	0.00728	0.00819	0.00559	0.00651	0.00273	0.00276
	(0.00682)	(0.00676)	(0.00698)	(0.00693)	(0.000900)	(0.000903)
1-yr lagged rain	-0.00368	-0.00495	-0.00349	-0.00494	-0.000166	-0.000348
	(0.00617)	(0.00639)	(0.00614)	(0.00638)	(0.00104)	(0.000978)
	[0.00566]	[0.00613]	[0.00517]	[0.00563]	[0.00132]	[0.00126]
Log non-land wealth (1971 Rs.)		0.0219		0.0219		
		(0.00355)		(0.00350)		
		[0.00426]		[0.00430]		
Fixed effects	Household	Household	Household	Household	District	District
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	12005	11856	12005	11856	8130	8130

Note: Standard errors, in parentheses below the coefficients, allow for clustering within a latitude–longitude grid point. A dry shock is defined rainfall below the 20th percentile and a wet shock is defined as rainfall above the 80th percentile. See Section 4 for details on how the variables are constructed.

Table 3: Testing for Crop Adaptation, FE and FE with Spatial Standard Errors

Data set:	Household	Household	Household	Household
Specification:	FE	FE	FE	FE
Dependent variable:	Crop Water Need	Crop Water Need	Crop Water Need (Monsoon)	Crop Water Need (Monsoon)
	(1)	(2)	(3)	(4)
Ten-year lagged average rainfall	0.0267		0.0374	
	(0.0235)		(0.0289)	
	[0.0232]		[0.0318]	
Ten-year lagged average of dry shock		-0.160		-0.206
		(0.0475)		(0.0661)
		[0.0276]		[0.0569]
Ten-year lagged average of wet shock		0.0175		0.0781
		(0.0449)		(0.0570)
		[0.0418]		[0.0658]
Current year rainfall	0.0140	0.0156	0.0132	0.0156
	(0.00698)	(0.00673)	(0.00965)	(0.00927)
	[0.00512]	[0.00475]	[0.00813]	[0.00808]
L1gz	-0.0157	-0.0183	-0.0155	-0.0208
	(0.00809)	(0.00763)	(0.00946)	(0.00919)
	[0.00480]	[0.00415]	[0.00655]	[0.00664]
Log non-land wealth (1971 Rs.)	-0.00687	-0.00698	-0.0107	-0.0104
	(0.00289)	(0.00278)	(0.00369)	(0.00341)
	[0.00242]	[0.00198]	[0.00327]	[0.00250]
Fixed effects	Household	Household	Household	Household
Year fixed effects	Yes	Yes	Yes	Yes
Observations	6389	6389	6305	6305

Note: Standard errors, in parentheses below the coefficients, allow for clustering within a latitude-longitude grid point. A dry shock is rainfall below the 20th percentile and a wet shock is rainfall above the 80th percentile. See Section 4 for details on how the variables are constructed.

Table 4: Irrigation And Crop Adaptation, Individual Rain Lags, Z-scores

Data set:	Household	Household	District	Household	Household
Specification:	FE-IV	FE-IV	FE	FE-IV	FE-IV
Dependent variable:	Irrigation Investment (Dummy)	Irrigation Investment (Dummy)	Log of the One-Year Change of Irrigated Area	Crop Water Need	Crop Water Need (Monsoon)
	(1)	(2)	(3)	(4)	(5)
Current year rainfall	0.00633	0.00858	0.00280***	0.0146	0.0103
	(0.00755)	(0.00736)	(0.000896)	(0.00961)	(0.0135)
One-year lagged rainfall	-0.00476	-0.00856	-0.000791	-0.0222***	-0.0203*
	(0.00610)	(0.00671)	(0.000961)	(0.00745)	(0.0111)
Two-year lagged rainfall	0.00216	0.00378	-0.000185	0.00556	0.0161
	(0.00796)	(0.00769)	(0.000655)	(0.00868)	(0.0112)
Three-year lagged rainfall	-0.0138*	-0.0124	-0.00114	0.0135	0.0288**
	(0.00780)	(0.00820)	(0.000769)	(0.00895)	(0.0112)
Four-year lagged rainfall	0.000822	0.00378	-0.00214	-0.0000231	0.000670
	(0.00697)	(0.00740)	(0.00154)	(0.00666)	(0.00865)
Five-year lagged rainfall	-0.0119	-0.0136*	0.00117*	-0.000194	0.000969
	(0.00749)	(0.00781)	(0.000623)	(0.00769)	(0.0107)
Six-year lagged rainfall	0.0147	0.00789	0.0000650	-0.00298	-0.00574
	(0.00932)	(0.00924)	(0.000822)	(0.00924)	(0.0105)
Seven-year lagged rainfall	-0.0110	-0.00801	0.000585	0.0263***	0.0188
	(0.00833)	(0.00877)	(0.000824)	(0.00915)	(0.0127)
Eight-year lagged rainfall	-0.00995	-0.00790	-0.00231	0.0103	0.0249***
	(0.00944)	(0.00971)	(0.00170)	(0.00655)	(0.00930)
Nine-year lagged rainfall	-0.00486	-0.00821	-0.00141**	-0.00875	-0.0120
	(0.00774)	(0.00773)	(0.000561)	(0.00743)	(0.00963)
Ten-year lagged rainfall	-0.00266	0.00218	-0.00101	-0.00534	-0.00161
	(0.00796)	(0.00803)	(0.000719)	(0.00531)	(0.00671)
Eleven-year lagged rainfall	0.00659	0.00791	0.000680	0.000137	0.00607
	(0.00510)	(0.00550)	(0.000827)	(0.00457)	(0.00533)
Twelve-year lagged rainfall	0.0143**	0.00985	0.000951	0.00337	-0.00274
	(0.00657)	(0.00705)	(0.000613)	(0.00675)	(0.00853)
Thirteen-year lagged rainfall	-0.00262	-0.00240	-0.000214	-0.0164**	-0.0182*
	(0.00577)	(0.00603)	(0.000726)	(0.00811)	(0.00972)
Fourteen-year lagged rainfall	-0.00376	-0.00355	0.000955	0.00399	0.00742
	(0.00593)	(0.00611)	(0.000741)	(0.00595)	(0.00803)
Fifthteen-year lagged rainfall	-0.0160***	-0.0255***	-0.000824	-0.0187**	-0.0251**
	(0.00587)	(0.00634)	(0.000914)	(0.00861)	(0.00989)
Log non-land wealth (1971 Rs.)		0.0483***		0.000876	0.00674
		(0.0128)		(0.0134)	(0.0198)
Fixed effects	Household	Household	District	Household	Household
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	12003	11759	8130	5577	5462

Notes: Standard errors, in parentheses below the coefficients, allow for clustering within a latitude–longitude grid point. In columns 1, 2, 4 and 5, I instrument for wealth with inherited wealth. See Section 4 in the main text for details on how the variables are constructed. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Testing for Irrigation Adaptation: Alternative Rainfall Lag Windows

Data set: Specification:	Household FE	Household FE	Household FE	Household FE	District FE	District FE	District FE	District FE
Dependent variable:	Irrig. Invest. (Dummy)	Irrig. Invest. (Dummy)	Irrig. Invest. (Dummy)	Irrig. Invest. (Dummy)	Log of the One-Year Change of Irrigated Area	Log of the One-Year Change of Irrigated Area	Log of the One-Year Change of Irrigated Area	Log of the One-Year Change of Irrigated Area
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Five-year lagged average rainfall	-0.0418**				-0.00240			
	(0.0192)				(0.00190)			
Five-year lagged average of dry shock		0.0340				0.0129***		
		(0.0426)				(0.00358)		
Five-year lagged average of wet shock		-0.0522				0.00153		
		(0.0385)				(0.00329)		
Fifteen-year lagged average rainfall			-0.0362				-0.00483**	
			(0.0246)				(0.00245)	
Fifteen-year lagged average of dry shock				-0.0505				0.0185***
				(0.0811)				(0.00713)
Fifteen-year lagged average of wet shock				-0.0836				0.00601
				(0.0673)				(0.00575)
Current year rainfall	0.00690	0.00590	0.00687	0.00698	0.00289***	0.00294***	0.00283***	0.00287***
	(0.00698)	(0.00738)	(0.00690)	(0.00684)	(0.000904)	(0.000909)	(0.000898)	(0.000903)
One-year lagged rain	-0.000342	-0.00347	-0.00716	-0.00755	-0.000275	-0.000195	-0.000528	-0.000626
	(0.00609)	(0.00601)	(0.00549)	(0.00535)	(0.00116)	(0.00108)	(0.000957)	(0.000933)
Fixed effects	Household	Household	Household	Household	District	District	District	District
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	12003	12003	12003	12003	8130	8130	8130	8130

Notes: Standard errors, in parentheses below the coefficients, allow for clustering within a latitude-longitude grid point. A dry shock is rainfall below the 20th percentile and a wet shock is rainfall above the 80th percentile. See Section 4 in the main text for details on how the variables are constructed.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Testing for Crop Adaptation: Alternative Rainfall Lag Windows

Data set:	Household	Household	Household	Household	Household	Household	Household	Household
Specification:	FE-IV	FE-IV	FE-IV	FE-IV	FE-IV	FE-IV	FE-IV	FE-IV
Dependent variable:	Crop Water Need	Crop Water Need	Crop Water Need	Crop Water Need	Crop Water Need (Monsoon)	Crop Water Need (Monsoon)	Crop Water Need (Monsoon)	Crop Water Need (Monsoon)
	(1)	(2)	(3)	(4)	(5)	(6)		
Five-year lagged average rainfall	0.0244				0.0414			
	(0.0210)				(0.0290)			
Five-year lagged average of dry shock		-0.0476				-0.105**		
		(0.0380)				(0.0531)		
Five-year lagged average of wet shock		0.0561*				0.0866*		
		(0.0338)				(0.0450)		
Fifteen-year lagged average rainfall			0.00608				0.0179	
			(0.0316)				(0.0388)	
Fifteen-year lagged average of dry shock				-0.140**				-0.195**
				(0.0699)				(0.0976)
Fifteen-year lagged average of wet shock				-0.00569				0.0841
				(0.0533)				(0.0667)
Current year rainfall	0.0130*	0.0133*	0.0155**	0.0143**	0.0116	0.0107	0.0156	0.0127
	(0.00735)	(0.00746)	(0.00726)	(0.00728)	(0.0107)	(0.0101)	(0.0106)	(0.0104)
L1gz	-0.0179*	-0.0196**	-0.0138*	-0.0124*	-0.0202*	-0.0233**	-0.0134	-0.0132
	(0.00914)	(0.00884)	(0.00722)	(0.00726)	(0.0108)	(0.0108)	(0.00881)	(0.00914)
Log non-land wealth (1971 Rs.)	-0.00173	0.000557	-0.00313	-0.00655	-0.000145	0.00238	-0.00278	-0.00785
	(0.0117)	(0.0121)	(0.0130)	(0.0126)	(0.0183)	(0.0181)	(0.0203)	(0.0194)
Fixed effects	Household	Household	Household	Household	Household	Household	Household	Household
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5577	5577	5577	5577	5462	5462	5462	5462

Notes: Standard errors, in parentheses below the coefficients, allow for clustering within a latitude-longitude grid point. A dry shock is rainfall below the 20th percentile and a wet shock is rainfall above the 80th percentile. In all columns, I instrument for wealth with inherited wealth. See Section 4 in the main text for details on how the variables are constructed.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Testing for Irrigation Adaptation: Controls for Government Dams

Data set:	Household	Household	Household	Household	District	District
Specification:	FE	FE-IV	FE	FE-IV	FE	FE
Dependent variable:	Irrigation Investment (Dummy)	Irrigation Investment (Dummy)	Irrigation Investment (Dummy)	Irrigation Investment (Dummy)	Log of the One-Year Change of Irrigated Area	Log of the One-Year Change of Irrigated Area
	(1)	(2)	(3)	(4)	(5)	(6)
Ten-year lagged average rainfall	-0.0465*	-0.0480*			-0.00789***	
	(0.0259)	(0.0262)			(0.00300)	
Ten-year lagged average of dry shock			0.134**	0.127**		0.0188**
			(0.0627)	(0.0606)		(0.00803)
Ten-year lagged average of wet shock			-0.0295	-0.0149		-0.00173
			(0.0568)	(0.0609)		(0.00488)
Upstream dams	0.153**	0.135**	0.148***	0.130**	0.000955	0.000903
	(0.0617)	(0.0619)	(0.0549)	(0.0565)	(0.00369)	(0.00373)
Current year rainfall	0.0137	0.0145*	0.0116	0.0123	0.00328***	0.00332***
	(0.00895)	(0.00872)	(0.00899)	(0.00880)	(0.000954)	(0.000954)
One-year lagged rain	-0.00431	-0.00606	-0.00453	-0.00707	-0.000561	-0.000810
	(0.00624)	(0.00691)	(0.00622)	(0.00705)	(0.00123)	(0.00114)
Log non-land wealth (1971 Rs)		0.0528***		0.0525***		
		(0.0134)		(0.0131)		
Fixed effects	Household	Household	Household	Household	District	District
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10632	10428	10632	10428	7046	7046

Notes: Standard errors, in parentheses below the coefficients, allow for clustering within a latitude-longitude grid point. A dry shock is rainfall below the 20th percentile and a wet shock is rainfall above the 80th percentile. In columns 2 and 4, I instrument for wealth with inherited wealth. See Section 4 in the main text for details on how the variables are constructed. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Testing for Adaptation: Controls for Changing Technology and Policies

Data set:	Household	Household	District	District	Household	Household
Specification:	FE-IV	FE-IV	FE	FE	FE-IV	FE-IV
Dependent variable:	Irrigation Investment (Dummy)	Irrigation Investment (Dummy)	Log of the One-Year Change of Irrigated Area	Log of the One-Year Change of Irrigated Area	Crop Water Need	Crop Water Drought (Monsoon)
	(1)	(2)	(3)	(4)	(5)	(6)
Ten-year lagged average of dry shock	0.104^{*}	0.219^{**}	0.0218^{***}	0.0355^{**}	-0.138^{***}	-0.173^{***}
	(0.0600)	(0.0855)	(0.00762)	(0.0154)	(0.0438)	(0.0582)
Ten-year lagged average of wet shock	-0.00392	0.00974	-0.00301	0.00632	0.0367	0.101^{*}
	(0.0542)	(0.0656)	(0.00379)	(0.00617)	(0.0444)	(0.0576)
Current year rainfall	0.0173^{**}	0.00921	0.00303^{***}	0.00396^{***}	0.00761	0.00164
	(0.00747)	(0.0144)	(0.000976)	(0.00126)	(0.00704)	(0.00947)
L1gz	-0.0110	-0.0113	-0.000716	-0.00122	-0.0121^{*}	-0.00756
	(0.00698)	(0.00821)	(0.00102)	(0.00130)	(0.00692)	(0.00802)
Log non-land wealth (1971 Rs.)	0.0383^{***}	0.0465^{***}			-0.00158	0.0107
	(0.0112)	(0.0155)			(0.0107)	(0.0151)
Fixed effects	District	District	Household	Household	Household	Household
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	11759	9367	8130	4808	5577	5462

Notes: Standard errors, in parentheses below the coefficients, allow for clustering within a latitude-longitude grid point. For details of the control variables included in each column, please refer to Section C.5. A dry shock is rainfall below the 20th percentile and a wet shock is rainfall above the 80th percentile. In columns 1, 2, 5 and 6, I instrument for wealth with inherited wealth. Column 1 controls for access to financial services, the presence of agricultural extension services, electrification rates and HYV suitability. Columns 2, 5 and 6 use the same controls as column 2 and add controls for transportation infrastructure, government intervention in village output markets, and government village-level irrigation subsidies, credit and irrigation assets. Column 3 controls for fertilizer prices and HYV suitability. Column 4 uses these same controls and adds a control for electrification rates. See Section 4 in the main text for details on how the variables are constructed.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 9: Testing for Irrigation Adaptation, Region*Year FE

Data set:	Household	Household	Household	Household	District	District
Specification:	FE-IV	FE-IV	FE-IV	FE-IV	FE	FE
Dependent variable:	Irrigation Investment (Dummy)	Irrigation Investment (Dummy)	Irrigation Investment (Dummy)	Irrigation Investment (Dummy)	Log of the One-Year Change of Irrigated Area	Log of the One-Year Change of Irrigated Area
	(1)	(2)	(3)	(4)	(5)	(6)
Ten-year lagged average rainfall	-0.0334	-0.0180			-0.00467*	
	(0.0273)	(0.0265)			(0.00263)	
Ten-year lagged average of dry shock			0.0596	0.0509		0.0167**
			(0.0654)	(0.0629)		(0.00713)
Ten-year lagged average of wet shock			-0.0613	-0.0233		0.000802
			(0.0500)	(0.0507)		(0.00382)
Current year rainfall	0.0154**	0.0198***	0.0143*	0.0189***	0.00245**	0.00244**
	(0.00713)	(0.00664)	(0.00729)	(0.00677)	(0.000954)	(0.000958)
1-yr lagged rain	-0.00407	-0.00575	-0.00371	-0.00546	0.000211	0.0000984
	(0.00687)	(0.00706)	(0.00660)	(0.00684)	(0.00120)	(0.00113)
Log non-land wealth (1971 Rs.)		0.0541***		0.0537***		
		(0.0137)		(0.0135)		
Fixed effects	Household	Household	Household	Household	District	District
Region-year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	12003	11759	12003	11759	8130	8130

Notes: Standard errors, in parentheses below the coefficients, allow for clustering within a latitude-longitude grid point. A dry shock is rainfall below the 20th percentile and a wet shock is rainfall above the 80th percentile. In columns 2 and 4, I instrument for wealth with inherited wealth. See Section 4 in the main text for details on how the variables are constructed. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 10: Testing for Crop Adaptation, Region*Year FE

Data set:	Household	Household	Household	Household
Specification:	FE-IV	FE-IV	FE-IV	FE-IV
Dependent variable:	Crop Water Need	Crop Water Need	Crop Water Need (Monsoon)	Crop Water Need (Monsoon)
	(1)	(2)	(3)	(4)
Ten-year lagged average rainfall	0.0450*		0.0783**	
	(0.0237)		(0.0314)	
Ten-year lagged average of dry shock		-0.199***		-0.244***
		(0.0499)		(0.0655)
Ten-year lagged average of wet shock		0.0288		0.132**
		(0.0436)		(0.0546)
Current year rainfall	0.0183**	0.0238***	0.0198*	0.0259**
	(0.00791)	(0.00764)	(0.0118)	(0.0114)
L1gz	-0.0136*	-0.0143**	-0.0141	-0.0170*
	(0.00789)	(0.00702)	(0.0104)	(0.00986)
Log non-land wealth (1971 Rs.)	0.00489	0.00387	0.0158	0.0158
	(0.0128)	(0.0120)	(0.0195)	(0.0179)
Fixed effects	Household	Household	Household	Household
Region-year fixed effects	Yes	Yes	Yes	Yes
Observations	5577	5577	5462	5462

Note: Standard errors, in parentheses below the coefficients, allow for clustering within a latitude-longitude grid point. A dry shock is rainfall below the 20th percentile and a wet shock is rainfall above the 80th percentile. In all columns, I instrument for wealth with inherited wealth. See Section 4 in the main text for details on how the variables are constructed. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

D Statistical Test of the Non-Stationarity of the Indian Monsoon

There is a consensus among meteorologists that the Indian monsoon undergoes multi-decadal wet and dry regimes (Mooley and Parthasarathy, 1984; Parthasarathy et al., 1991; Subbaramayya and Naidu, 1992; Pant and Kumar, 1997; Kripalani and Kulkarni, 1997; Naidu et al., 1999; Torrence and Webster, 1999; Pant, 2003; Varikoden and Babu, 2014). There are a few meteorological papers that demonstrate that these monsoon regimes generate statistically significant variation in rainfall. Mooley and Parthasarathy (1984) find evidence of statistically significant rainfall via several different statistical analyses: a Cramer's t-test, a low pass binomial filter, and residual mass curve analysis. Kripalani and Kulkarni (1997) use a Cramer's t-test to demonstrate the statistical significance of the regimes. Specifically, they apply this test to the 11-year running means of India's summer rainfall. The test indicates that there are statistically significant persistent deviations of this running mean from the historical mean. This provides evidence that the summer rainfall does not follow an i.i.d. process, but rather does demonstrate epochal behavior. Kripalani and Kulkarni (1997) also analyze the statistical significance of the relationship between the monsoon regimes and the El Niño Southern Oscillation.

In this appendix, I run an additional test that verifies the statistical significance of the monsoon regimes. Specifically, I compute a quasi-likelihood ratio statistic for a mixture model to test the null hypothesis of one regime versus the alternative of two regimes in a Markov regime-switching context, following the approach developed by Cho and White (2007). The distribution of the test statistic is nonstandard due to nuisance parameters that only exist under the alternative hypothesis; however I am able to use the critical values tabulated in Steigerwald and Carter (2011) for this purpose. I calculate the test statistic to be 9.61, which is greater than the tabulated 5% critical value of 5.54, and therefore I reject the null hypothesis of a single rainfall regime. This suggests that, in this context, farmer adaptation to recent rainfall can be interpreted as a rational response to persistent rainfall variations that are greater than what would be expected under i.i.d. rainfall.

E Proof of Signs of Partial Derivatives

In order to have the desired signs for the wealth and expectation effects that are derived in Section 3.4 of the main text, we need to prove the following signs for these partial derivatives:

$$\frac{\partial i_2^*}{\partial w_2^*} > 0, \frac{\partial i_2^*}{\partial \mu_2} < 0, \frac{\partial \rho_2^*}{\partial w_2^*} < 0, \frac{\partial \rho_2^*}{\partial \mu_2} < 0, \frac{\partial w_2^*}{\partial r_1} > 0, \frac{\partial w_2^*}{\partial \mu_2} < 0$$

E.1 Solving for $\frac{\partial i_2^*}{\partial w_2^*}, \frac{\partial i_2^*}{\partial \mu_2}, \frac{\partial \rho_2^*}{\partial w_2^*}, \frac{\partial \rho_2^*}{\partial \mu_2}$

We can re-write the profit function so it is a function of irrigation, total wealth, drought-tolerant crop area, and rainfall:

$$\begin{aligned}\pi(i_2, w_2, \rho_2, r_2) =& \beta_a(w_2 - i_2) + \beta_i i_t + \beta_\rho \rho_t + \frac{1}{2}\delta_{aa}(w_2 - i_2)^2 + \frac{1}{2}\delta_{ii} i_t^2 + \\ &+ \frac{1}{2}\delta_{\rho\rho}\rho_t^2 + \delta_{\rho i}\rho_t i_t + \delta_{ir} i_t r_t + \delta_{\rho r}\rho_t r_t + \delta_r r_t \end{aligned} \tag{1}$$

I want to solve for the first and second order conditions that define $i_2^*(w_2, \mu_2)$ and $\rho_2^*(w_2, \mu_2)$.

Note that the farmer chooses second period irrigation and crop choice in order to maximize expected second period utility.

$$\max \mathrm{E}_1[u(w_2 + \pi(i_2, w_2, \rho_2, r_2)] \text{ w.r.t. } i_2 \text{ and } \rho_2$$

Since we are assuming CARA utility and normally distributed rainfall, we can make use the fact that if $r_t \sim N(\mu, \sigma)$, then $\mathrm{E}(e^{\gamma r_t}) = e^{\gamma\mu + \frac{1}{2}\gamma^2\sigma^2}$ (Bolton and Dewatripont (2005), p138). Substituting in our expression for the utility function and applying the above identity, we get that the farmer is solving $\max -e^{-f(i_2, w_2, \rho_2, \mu_2)}$, where

$$f(i_2, w_2, \rho_2, \mu_2) = \eta(w_2 + \pi(i_2, w_2, \rho_2, \mu_2)) - \frac{1}{2}\eta^2\sigma^2(\delta_{ir} i_2 + \delta_{\rho r}\rho_2 + \delta_r)^2 \tag{2}$$

Since the exponential function is monotonic, the farmer's maximization problem is equivalent to:

$$\max f(i_2, w_2, \rho_2, \mu_2) \text{ w.r.t. } i_2 \text{ and } \rho_2$$

The first order conditions for this maximization problem are $f_i = 0$ and $f_\rho = 0$. This system of equations defines $i_2^*(w_2, \mu_2)$ and $\rho_2^*(w_2, \mu_2)$. The second order conditions for this maximization problem are $f_{ii} < 0$, $f_{\rho\rho} < 0$ and $f_{ii}f_{\rho\rho} - f_{i\rho}f_{\rho i} > 0$. In order to solve for $\frac{\partial i_2^*}{\partial w_2^*}$ and $\frac{\partial \rho_2^*}{\partial w_2^*}$, I take the derivative of the first order conditions with respect to w_2, and get the resulting system of equations:

$$f_{ii}\frac{\partial i_2^*}{\partial w_2} + f_{i\rho}\frac{\partial \rho_2^*}{\partial w_2} + f_{iw} = 0$$
$$f_{\rho i}\frac{\partial i_2^*}{\partial w_2} + f_{\rho\rho}\frac{\partial \rho_2^*}{\partial w_2} + f_{\rho w} = 0$$

Solving this system of equations, I get the following expressions:

$$\frac{\partial i_2^*}{\partial w_2} = -\frac{1}{det}\left(f_{\rho\rho}f_{iw} - f_{i\rho}f_{\rho w}\right)$$
$$\frac{\partial \rho_2^*}{\partial w_2} = -\frac{1}{det}\left(f_{ii}f_{\rho w} - f_{\rho i}f_{iw}\right)$$

where $det = f_{ii}f_{\rho\rho} - f_{i\rho}f_{\rho i}$. Note that by the second order conditions, we have $det > 0$.

Similarly, in order to solve for $\frac{\partial i_2^*}{\partial \mu_2}$ and $\frac{\partial \rho_2^*}{\partial \mu_2}$, I take the derivative of the first order conditions with respect to μ_2, and solve the resulting system of equations, getting the following expressions:

$$\frac{\partial i_2^*}{\partial \mu_2} = -\frac{1}{det}\left(f_{\rho\rho}f_{i\mu} - f_{i\rho}f_{\rho\mu}\right)$$
$$\frac{\partial \rho_2^*}{\partial \mu_2} = -\frac{1}{det}\left(f_{ii}f_{\rho\mu} - f_{\rho i}f_{i\mu}\right)$$

Therefore in order to determine the signs of the comparative statics, it is sufficient to calculate the second-order partial derivatives and second-order mixed derivatives of $f(i_2, w_2, \rho_2, \mu_2)$.

Using the expression for $f(i_2, w_2, \rho_2, \mu_2)$ given in Equation 2 and substituting in the expression for the profit function given in Equation 1, we get that the first order partial derivatives of f are:

$$f_i = \eta(-\beta_a + \beta_i + \delta_{aa}(i_2 - w_2) + \delta_{ii}i_2 + \delta_{\rho i}\rho_2 + \delta_{ir}\mu_2) - \eta^2\sigma^2\delta_{ir}(\delta_{ir}i_2 + \delta_{\rho r}\rho_2 + \delta_r)$$
$$f_\rho = \eta(\beta_\rho + \delta_{\rho\rho}\rho_2 + \delta_{\rho i}i_2 + \delta_{\rho r}\mu_2) - \eta^2\sigma^2\delta_{\rho r}(\delta_{ir}i_2 + \delta_{\rho r}\rho_2 + \delta_r)$$
$$f_w = -\eta(1 + \beta_a + \delta_{aa}w_2 - \delta_{aa}i_2)$$
$$f_\mu = \eta(\delta_{ir}i_2 + \delta_{\rho r}\rho_2 + \delta_r)$$

Furthermore, using the assumed signs of the coefficients from the profit function given in Section 3.2 of the main text, we get the following expressions and signs for the second order partial derivatives:

$$f_{ii} = \eta\delta_{ii} + \eta\delta_{aa} - \eta^2\sigma^2\delta_{ir}^2 < 0$$
$$f_{iw} = -\eta\delta_{aa} > 0$$
$$f_{i\rho} = f_{\rho i} = \eta\delta_{\rho i} - \eta^2\sigma^2\delta_{\rho r}\delta_{ir} < 0$$
$$f_{i\mu} = \eta\delta_{ir} < 0$$

$$f_{\rho\rho} = \eta\delta_{\rho\rho} - \eta^2\sigma^2\delta_{\rho r}^2 < 0$$

$$f_{\rho w} = 0$$

$$f_{\rho\mu} = \eta\delta_{\rho r} < 0$$

Using these signs and expressions for the partial derivatives, we get that

$$\frac{\partial i_2^*}{\partial w_2} = -\frac{1}{det}\left(f_{\rho\rho}f_{iw} - f_{i\rho}f_{\rho w}\right) = -\frac{1}{det}(f_{\rho\rho}f_{iw}) > 0$$

$$\frac{\partial \rho_2^*}{\partial w_2} = -\frac{1}{det}\left(f_{ii}f_{\rho w} - f_{\rho i}f_{iw}\right) = \frac{1}{det}\left(f_{\rho i}f_{iw}\right) < 0$$

$$\begin{aligned}\frac{\partial i_2^*}{\partial \mu_2} &= -\frac{1}{det}\left(f_{\rho\rho}f_{i\mu} - f_{i\rho}f_{\rho\mu}\right)\\ &= -\frac{1}{det}[(\eta\delta_{\rho\rho} - \eta^2\sigma^2\delta_{\rho r}^2)(\eta\delta_{ir}) - (\eta\delta_{\rho i} - \eta^2\sigma^2\delta_{\rho r}\delta_{ir})(\eta\delta_{\rho r})]\\ &= -\frac{1}{det}(\eta\delta_{\rho\rho}\delta_{ir} - \eta^3\sigma^2\delta_{\rho r}^2\delta_{ir} - \eta^2\delta_{\rho i}\delta_{\rho r} - \eta^3\sigma^2\delta_{\rho r}\delta_{ir}\delta_{\rho r}) < 0,\end{aligned}$$

as long as $\left|\delta_{\rho i}\right|$ is not too large.

$$\begin{aligned}\frac{\partial \rho_2^*}{\partial \mu_2} &= -\frac{1}{det}\left(f_{ii}f_{\rho\mu} - f_{\rho i}f_{i\mu}\right)\\ &= -\frac{1}{det}[(\eta\delta_{ii} + \eta\delta_{aa} - \eta^2\sigma^2\delta_{ir}^2)(\eta\delta_{\rho r}) - (\eta\delta_{\rho i} - \eta^2\sigma^2\delta_{\rho r}\delta_{ir})(\eta\delta_{ir})]\\ &= -\frac{1}{det}(\eta^2\delta_{ii}\delta_{\rho r} + \eta^2\delta_{aa}\delta_{\rho r} - \eta^3\sigma^2\delta_{ir}^2\delta_{\rho r} - \eta^2\delta_{\rho i}\delta_{ir} + \eta^3\sigma^2\delta_{\rho r}\delta_{ir}^2)\\ &= -\frac{1}{det}(\eta^2\delta_{ii}\delta_{\rho r} + \eta^2\delta_{aa}\delta_{\rho r} - \eta^2\delta_{\rho i}\delta_{ir}) < 0,\end{aligned}$$

as long as $\left|\delta_{\rho i}\right|$ is not too large.

All of the signs are as desired.

E.2 Solving for $\frac{\partial w_2^*}{\partial r_1}$ and $\frac{\partial w_2^*}{\partial \mu_2}$

The farmer chooses second period wealth to maximize the following expression

$$g(w_2, r_1, \mu_2) = u(w_1 + \pi(i_1, w_1, \rho_1, r_1) - w_2) + \mathrm{E}_1[u(w_2 + \pi^*(w_2, \mu_2, r_2)]$$

where

$$\pi^*(w_2, \mu_2, r_2) = \pi(i_2^*(w_2, \mu_2), w_2, \rho_2^*(w_2, \mu_2), r_2)$$

The first order condition for this maximization problem is $g_w = 0$. This implicitly defines $w_2^*(r_1, \mu_2)$. The second order condition for this maximization problem is $g_{ww} < 0$. In order to solve for $\frac{\partial w_2^*}{\partial r_1}$, I take the derivative of first order condition with respect to r_1, and get the resulting equation:

$$g_{ww}\frac{\partial w_2^*}{\partial r_1} + g_{wr} = 0$$

We get $\frac{\partial w_2^*}{\partial r_1} = -\frac{g_{wr}}{g_{ww}}$. We have that $g_{ww} < 0$ by the second order conditions. Therefore, in order to demonstrate that $\frac{\partial w_2^*}{\partial r_1} > 0$, it is sufficient to show that $g_{wr} > 0$. Note that $g(w_2, r_1, \mu_2)$ has two pieces (first period utility and expected second period utility), and first period rainfall only enters in via first period utility. Therefore, using the expression for the profit function in Equation 1, we get the following derivatives for $g(w_2, r_1, \mu_2)$:

$$g_r = u'(w_1 + \pi_1 - w_2)(\delta_{ir} i_1 + \delta_{\rho r}\rho_1 + \delta_r)$$

$$g_{wr} = u''(w_1 + \pi_1 - w_2)(-1)(\delta_{ir} i_1 + \delta_{\rho r}\rho_1 + \delta_r)$$

Therefore, by the concavity of utility, we get

$$g_{wr} > 0,$$

as long as,

$$\delta_{ir} i_1 + \delta_{\rho r} \rho_1 + \delta_r > 0$$

i.e. as long as first-period irrigation and drought-tolerant crop area are not so large such that higher rainfall is bad for profits, which is a reasonable assumption. Therefore, we have demonstrated that

$$\frac{\partial w_2^*}{\partial r_1} > 0,$$

as desired.

In order to solve for $\frac{\partial w_2^*}{\partial \mu_2}$, I take the derivative of first order condition with respect to μ_2, and get the resulting equation:

$$g_{ww} \frac{\partial w_2^*}{\partial \mu_2} + g_{w\mu} = 0$$

We get $\frac{\partial w_2^*}{\partial \mu_2} = -\frac{g_{w\mu}}{g_{ww}}$. We have that $g_{ww} < 0$ by the second order conditions. Therefore, in order to demonstrate that $\frac{\partial w_2^*}{\partial \mu_2} < 0$, it is sufficient to show that $g_{w\mu} < 0$. Note that $g(w_2, r_1, \mu_2)$ has two pieces (first period utility and second period expected utility), and second period expected rainfall only enters via expected second period utility. Furthermore, since we have CARA utility, we can write:

$$\mathrm{E}_1[u(w_2 + \pi^*(w_2, \mu_2, r_2)] = u(w_2 + \pi^*(w_2, \mu_2, \mu_2)) * e^{h(w_2, \mu_2)}$$

where

$$h(w_2,\mu_2) = \frac{1}{2}\eta^2\sigma^2(\delta_{ir}i_2^*(w_2,\mu_2) + \delta_{\rho r}\rho_2^*(w_2,\mu_2) + \delta_r)^2$$

Taking the derivative with respect to μ_2 and applying the product rule, we get

$$\begin{aligned} g_\mu =& u'(w_2 + \pi^*(w_2,\mu_2,\mu_2)) * \frac{d\pi * (w_2,\mu_2,\mu_2)}{d\mu_2} * e^{h(w_2,\mu_2)} + \\ &+ u(w_2 + \pi^*(w_2,\mu_2,\mu_2)) * e^{h(w_2,\mu_2)} * h_\mu \end{aligned}$$

Note that this expression is the sum of two products, and that each of the products have three terms, with one term common to both products, e.g.

$$g_\mu = abc + cde$$

where

$$\begin{aligned} a &= u'(w_2 + \pi^*(w_2,\mu_2,\mu_2)) \\ b &= \frac{d\pi^*(w_2,\mu_2,\mu_2)}{d\mu_2} \\ c &= e^{h(w_2,\mu_2)} \\ d &= u(w_2 + \pi^*(w_2,\mu_2,\mu_2)) \\ e &= h_\mu \end{aligned}$$

In order to calculate $g_{w\mu}$, I apply the product rule for three terms and get that

$$g_{w\mu} = abc_w + ab_wc + a_wbc + cde_w + cd_we + c_wde$$

Therefore, in order to compute the sign of $g_{w\mu}$, I compute the sign of each of these subcomponents. Before computing the sub-components, I note two important regularity conditions that must hold in order for me to get the desired signs:

$$\delta_{ir} i_2^*(w_2, \mu_2) + \delta_{\rho r} \rho_2^*(w_2, \mu_2) + \delta_r > 0 \tag{3}$$

and

$$\delta_{ir} \frac{\partial i_2^*}{\partial w_2} + \delta_{\rho r} \frac{\partial \rho_2^*}{\partial w_2} \tag{4}$$

The first condition states that second period irrigation and drought-tolerant crop area must not be so high that higher rainfall is bad for profits. The second condition states that the responsiveness of irrigation to wealth must be greater than the responsiveness of drought-tolerant crop area to wealth.

We are now ready to compute the signs of the subcomponents.

$$\begin{aligned}
a =& u'(w_2 + \pi^*(w_2, \mu_2, \mu_2)) > 0, \text{ because utility is increasing} \\
a_w =& u''(w_2 + \pi^*(w_2, \mu_2, \mu_2))(1 + \frac{\partial \pi^*(w_2, \mu_2, \mu_2)}{\partial w_2}) < 0, \text{ because of decreasing} \\
& \text{marginal utility and because profits are increasing in wealth} \\
b =& \frac{d\pi^*(w_2, \mu_2, \mu_2)}{d\mu_2} = \delta_{ir} i_2^*(w_2, \mu_2) + \delta_{\rho r} \rho_2^*(w_2, \mu_2) + \delta_r > 0, \text{ by regularity} \\
& \text{condition 3} \\
b_w =& \delta_{ir} \frac{\partial i_2^*}{\partial w_2} + \delta_{\rho r} \frac{\partial \rho_2^*}{\partial w_2}, \text{ by regularity condition 4} \\
c =& e^{h(w_2, \mu_2)} > 0 \\
c_w =& e^{h(w_2, \mu_2)} * h_w \\
=& e^{h(w_2, \mu_2)} * \eta^2 \sigma^2 (\delta_{ir} i_2^*(w_2, \mu_2) + \delta_{\rho r} \rho_2^*(w_2, \mu_2) + \delta_r) \left(\delta_{ir} \frac{\partial i_2^*}{\partial w_2} + \delta_{\rho r} \frac{\partial \rho_2^*}{\partial w_2} \right)
\end{aligned}$$

$$
\begin{aligned}
&< 0, \text{ by regularity conditions 3 and 3.} \\
d =& u(w_2 + \pi^*(w_2, \mu_2, \mu_2)) < 0, \text{ based on the form of the utility function} \\
d_w =& u'(w_2 + \pi^*(w_2, \mu_2, \mu_2))(1 + \frac{\partial \pi^*(w_2, \mu_2, \mu_2)}{\partial w_2}) > 0 \text{ because utility is} \\
&\text{increasing and because profits are increasing in wealth.} \\
e =& h_\mu = \eta^2 \sigma^2 (\delta_{ir} i_2^*(w_2, \mu_2) + \delta_{\rho r} \rho_2^*(w_2, \mu_2) + \delta_r)(\delta_{ir} \frac{\partial i_2^*}{\partial \mu_2} + \delta_{\rho r} \frac{\partial \rho_2^*}{\partial \mu_2}) > 0, \text{ by} \\
&\text{regularity condition 4} \\
e_w =& \eta^2 \sigma^2 (\delta_{ir} i_2^*(w_2, \mu_2) + \delta_{\rho r} \rho_2^*(w_2, \mu_2) + \delta_r) \left(\delta_{ir} \frac{\partial i_2^*}{\partial \mu_2} + \delta_{\rho r} \frac{\partial \rho_2^*}{\partial \mu_2} \right) * \\
&* \left(\delta_{ir} \frac{\partial i_2^*}{\partial w_2} + \delta_{\rho r} \frac{\partial \rho_2^*}{\partial w_2} \right) \\
&< 0, \text{ by regularity conditions 3 and 4}
\end{aligned}
$$

Note that in deriving the expression for e_w from e, I used the fact that the expressions for $\frac{\partial i_2^*}{\partial \mu_2}$ and $\frac{\partial \rho_2^*}{\partial \mu_2}$, derived above in Section A1, do not depend on w_2.

Therefore, based on the signs of these components, we find that $g_{w\mu} < 0$ as desired. Furthermore, this implies that

$$\frac{\partial w_2^*}{\partial \mu_2} < 0,$$

as desired.

References

Birner, R., Gupta, S., Sharma, and Neeru (2011). *The Political Economy of Agricultural Policy Reform in India: Fertilizers and Electricity for Irrigation.*

Bolton, P. and Dewatripont, M. (2005). *Contract Theory*, volume 1. The MIT Press, 1 edition.

Cameron, A. C. and Miller, D. L. (2015). A Practitioner's Guide to Cluster-Robust Inference. *Journal of Human Resources*, 50(2):317–372.

Cho, J. and White, H. (2007). Testing for regime switching. *Econometrica*, 75(6):1671–1720.

Deininger, K., Jin, S., and Nagarajan, H. K. (2009a). Determinants and Consequences of Land Sales Market Participation: Panel Evidence from India. 37(2):410–421.

Deininger, K., Jin, S., and Nagarajan, H. K. (2009b). Land Reforms, Poverty Reduction, and Economic Growth: Evidence from India. *Journal of Development Studies*, 45(4):496–521.

Duflo, E. and Pande, R. (2006). Dams, Poverty, Public Goods and Malaria Incidence in India. Jameel Poverty Action Lab.

Evenson, R. E. (2003). Assessing the Impact of the Green Revolution, 1960 to 2000. *Science*, 300(5620):758–762.

Fan, S., Gulati, A., and Thorat, S. (2008). Investment, subsidies, and pro-poor growth in rural India. *Agricultural Economics*, 39(2):163–170.

Fan, S., Hazell, P., and Thorat, S. (2000). Government Spending, Growth and Poverty in Rural India. *American Journal of Agricultural Economics*, 82(4):1038–1051.

Fetzer, T. (2014). Can workfare programs moderate violence? Evidence from India.

Foster, A. and Rosenzweig, M. (2003). Economic growth and the rise of forests. *The Quarterly Journal of Economics*.

Gulati, A. (1989). Input Subsidies in Indian Agriculture: A Statewise Analysis. *Economic & Political Weekly*, 24(25).

Hsiang, S. M. and Solow, R. M. (2010). Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America. *Proceedings of the National Academy of Sciences of the United States of America*, 107(35):15367–15372.

Jha, R. (2013). Household Coping Strategies and Welfare: Does Governance Matter? pages 1–23.

Jha, R., Kang, W., and Nagarajan, H. K. (2013). Vulnerability and Responses to Risks in Rural India. *Available at SSRN* . . . .

Kripalani, R. H. and Kulkarni, A. (1997). Climatic impact of El Niño/La Niña on the Indian monsoon: A new perspective. *Weather*, 52(2):39–46.

McKinsey, J. W. and Evenson, R. E. (1999). Technology-Climate Interactions in the Green Revolution in India. Economic Growth Center, Yale.

Mooley, D. A. and Parthasarathy, B. (1984). Fluctuations in All-India Summer Monsoon Rainfall During 1871-1978. *Climatic Change*, 6(3):287–301.

Moran, P. (1950). Notes on continuous stochastic phenomena. *Biometrika*, 37(1/2):17.

Naidu, C. V., Rao, B., and Rao, D. (1999). Climatic trends and periodicities of annual rainfall over India. *Meteorological Applications*, 6(4):395–404.

Pande, R. and Duflo, E. (2007). Dams. *Quarterly Journal of Economics*, pages 1–46.

Pant, G. B. (2003). Long-Term Climate variability and change over monsoon Asia. *J. Ind. Geophys. Union*, 7(3):125–134.

Pant, G. B. and Kumar, K. R. (1997). *Climates of South Asia*. Wiley.

Parthasarathy, B., Kumar, K. R., and Munot, A. A. (1991). Evidence of secular variations in Indian monsoon rainfall-circulation relationships. *Journal of Climate*, 4(9):927–938.

Pisati, M. (2012). Exploratory spatial data analysis using Stata. pages 1–91.

Rud, J. P. (2012). Electricity provision and industrial development: Evidence from India. *Journal of Development Economics*.

Sanghi, A., Kumar, K. K., and McKinsey, Jr, J. W. (1998). India agriculture and climate dataset. Technical report.

Shah, T. (1993). *Groundwater markets and irrigation development*. political economy and practical policy. Oxford University Press.

Steigerwald, D. and Carter, A. (2011). *Markov Regime-Switching Tests: Asymptotic Critical Values*.

Subbaramayya, I. and Naidu, C. (1992). Spatial variations and trends in the Indian monsoon rainfall. *International Journal of Climatology*, 12(6):597–609.

Thakkar, H. (1999). Assessment of Irrigation Options. World Commission of Dams.

Torrence, C. and Webster, P. J. (1999). Interdecadal changes in the ENSO-monsoon system. *Journal of Climate*, 12(8):2679–2690.

Vaidyanathan, A. (2010). *Agricultural Growth in India* . The Role of Technology, Incentives, and Institutions. Oxford University Press India.

Varikoden, H. and Babu, C. A. (2014). Indian summer monsoon rainfall and its relation with SST in the equatorial Atlantic and Pacific Oceans. *International Journal of Climatology*.

Vashishtha, P. S. (1989). *Changes in the Structure of Investment of Rural Households*, volume 112. Journal of Income and Wealth.