Adaptation to climate change: historical evidence from the Indian monsoon

VIS TARAZ

*Smith College, Department of Economics, Pierce Hall, 21 West Street,
Northampton, MA 01063-6317, USA.*
Email: vtaraz@smith.edu

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, the author exploits 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, she tests whether farmers have adapted to the medium-run rainfall variation induced by the monsoon regimes. The author finds 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 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 and Greenstone *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 and Greenstone *et al.*, 2007; Guiteras, 2009; Schlenker and Roberts, 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, we exploit historical rainfall variation in India to estimate adaptation. Rather than analyzing year-to-year weather deviations, we 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. We test whether farmers adapt their irrigation investments and crop portfolios in response to these persistent rainfall deviations. The monsoon regimes do not cause variation in temperature, so we do not analyze adaptation to temperature changes.

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.

Adaptation is tested for by analyzing whether agricultural decisions respond to lagged weather, looking specifically at irrigation investments and crop choice. We exploit the fact that the return to irrigation investment varies across wet versus dry growing seasons and that, similarly, the relative yields of different crops vary across wet versus dry growing seasons. The empirical strategy here 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 for the inclusion of year fixed effects in the regressions and, hence, adaptation to rainfall can be separated from unrelated temporal changes in irrigation and crop choice.

---

[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).
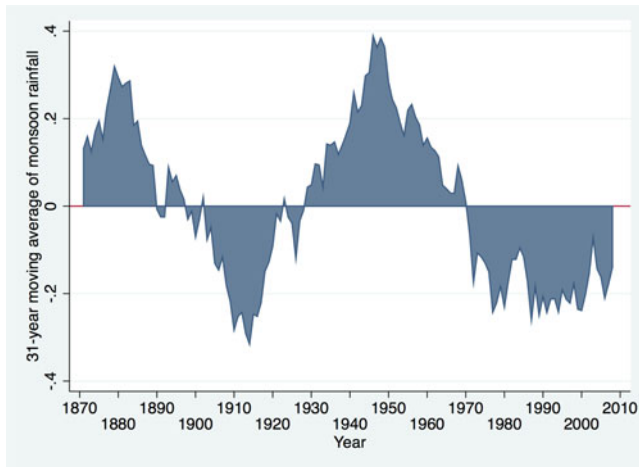
Figure 1. *Interdecadal variability of the Indian monsoon*
Notes: This figure displays the 31-year moving average of India's summer monsoon rainfall, measured as a z-score deviation from the historical mean.
Source: The rainfall data are from the India Institute of Tropical Meteorology's Homogeneous Indian Monthly Rainfall Data Set (1871–2008). The figure is constructed based on the author's calculations.

Analyzing two agricultural data sets, evidence is found 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 per cent probability of investing.[2] 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.[3] In addition to testing for the presence of adaptation, the extent to which adaptation prevents profit losses is also measured. It is found that farmers are able to recover only a limited amount of their lost profits by adapting. Specifically, it is estimated that in the face of sustained adverse weather conditions, adaptation recovers no more than 9 per cent of lost profits.

In addition to contributing to the literature on climate change and agriculture in India (Guiteras, 2009; Krishnamurthy, 2012; Fishman, 2016), this paper contributes to a rapidly growing literature on climate change adaptation. Dell *et al.* (2014) present a helpful synthesis of this literature. 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)

---

[2] A dry year is defined as a year in which rainfall is below the 20th percentile of the rainfall distribution for a particular location.
[3] The crops with lower water needs have lower expected yields, which is why farmers do not plant them exclusively.
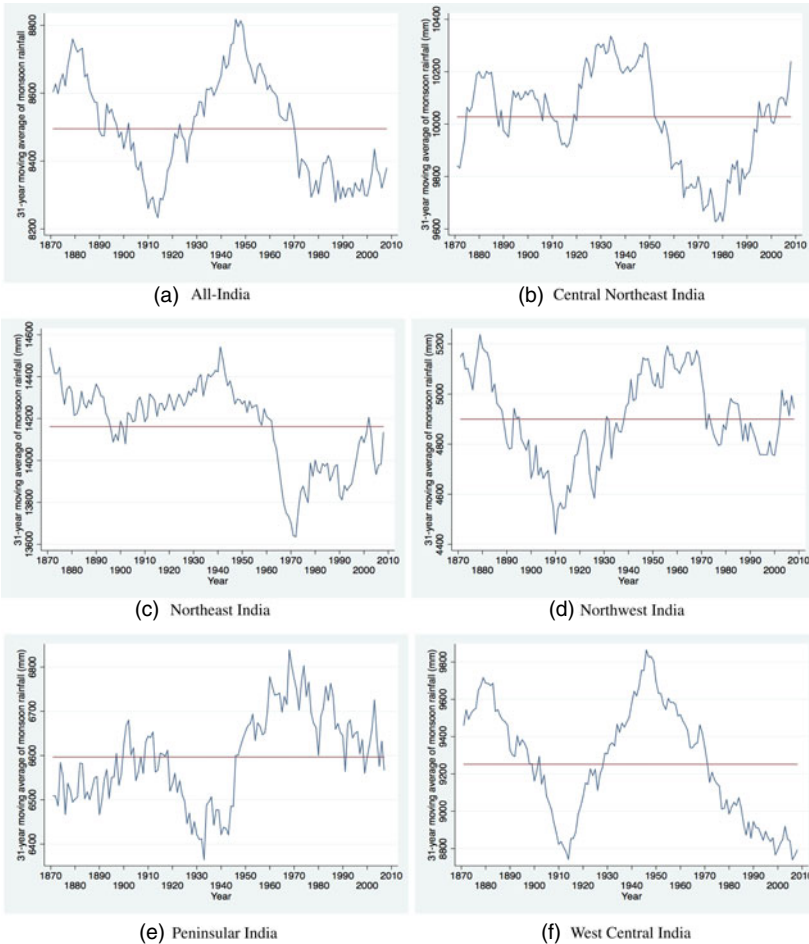
Figure 2. *Spatial variation of the interdecadal variability of the Indian monsoon*
Notes: This figure graphs the 31-year moving average of the summer monsoon rainfall,
measured in millimeters for India's five meteorological regions. The horizontal line rep-
resents mean rainfall for that region.
Source: The rainfall data are from the India Institute of Tropical Meteorology's Homoge-
neous Indian Monthly Rainfall Data Set (1871–2008). The figure is constructed based
on the author's calculations.

and variants of the Ricardian method that incorporate panel data (Luis
and Orlando, 2015) and structural agro-economic models (Kurukula-
suriya and Mendelsohn, 2008). Researchers have also analyzed adapta-
tion by looking at long-run responses to one-time environmental shocks
(Hornbeck, 2012; Deryugina, 2013; Hornbeck and Naidu, 2015), apply-
ing instrumental variables approaches that address the endogeneity of
adaptation (Di Falco and Veronesi, 2013, 2014), using economic mod-
els that integrate biophysical modeling (Finger *et al.*, 2010), employing

multinomial logit choice 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, 2016).

This paper contributes to the adaptation literature in multiple ways. First, this study is unique because it uses 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). This paper is also unique because it estimates 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.

The unique data set and source of climate variation lead to a methodological contribution to the literature. Specifically, we can estimate how farmers adapt to medium-run (10–20 year) changes in climate that are occurring over the span of our data set, while controlling for unobserved heterogeneity. Put differently, the data allow us to look at how the behavior of a household changes across several decades, in response to time-varying changes in climate. The 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.[4] Conversely, the estimates in this study 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. The multidecadal household panel used in this paper, on the other hand, allows us to control for unobserved farmer heterogeneity and to analyze adaptation over several decades. In addition, a model is built that allows us to disentangle the effects of wealth and expectations. This is a methodological contribution because it allows us to test directly 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.

Several important limitations of this study should be acknowledged. First, this study only analyzes irrigation and crop choice. Data limitations

---

[4] 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 (steady-state) setting.

do not permit the study of 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 (Fishman, 2015; Viswanathan and Kavi Kumar, 2015). Secondly, since the monsoon regimes affect only precipitation, adaptation to temperature changes is not analyzed. Thirdly, since the household data set spans several decades, there is substantial, non-random attrition, which causes the analyzed sample to include households that are, on average, wealthier than a representative sample would be. Fourthly, there are potential threats to the exclusion restriction for the instrumental variables strategy, which are discussed in greater detail in section 5. Fifthly, due to data limitations, we are 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 the results.

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. Section 7 discusses several robustness tests that were performed and which are reported in an online appendix, available at https://doi.org/10.1017/S1355770X17000195. Section 8 calculates the fraction of lost profits farmers recovered by adapting, and section 9 concludes.

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 A1 of online appendix 1. 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). Figure 1 in the main text smoothes national, annual rainfall with a moving average filter, to highlight the regimes.

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.), as discussed in greater detail in online appendix 1. 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 us to interpret a farmer's

response to lagged rainfall as evidence of rational adaptation, rather than as an indicator of irrational behavior.[5]

The monsoon regimes are not geographically homogeneous. There is significant spatial variation 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, online appendix 4 presents maps of 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. We choose rainfall from the previous decade as a rough measure of the current monsoon regime (Kripalani and Kulkarni, 1997). The spatial variation in recent rainfall allows us to include year fixed effects in the regressions and, hence, distinguish rainfall adaptation from time trends in irrigation and crop choice.

3. Theoretical framework
We present climate and agricultural models and derive tests for adaptation.

3.1. Climate model
We 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 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 this 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

---

[5] We 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, surveys demonstrate that farmers have noticed recent changes in temperature and rainfall that have been induced by anthropogenic climate change, which are comparable in magnitude to the changes that we analyze in this study (Palanisami *et al.*, 2014).

agricultural asset, a_t, such that $a_t + i_t = w_t$.[6] 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. 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. For tractability, we assume a quadratic profit function of the form:

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

where π_t is profits per acre and ϵ_t is a mean zero productivity shock. To establish our adaptation tests, we 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 is verified in online appendix 5.
(2)   The return to irrigation is higher during periods of low rainfall ($\delta_{ir} < 0$). This assumption is verified in online appendix 5.
(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 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 our 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, we assume constant absolute risk aversion utility of the form $u(c_t) = -e^{-\eta c_t}$. 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.

---

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

3.4. Tests for adaptation

We now derive tests to determine whether farmers update their rainfall expectations in response to past rainfall and whether they adapt their agricultural decisions accordingly. We lack data on farmer rainfall expectations, but the structure of the model allows us to test for adaptation, even without explicit data on expectations. To clarify the analysis, we 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 derive a test for irrigation adaptation, we take the derivative of second-period irrigation with respect to first-period rainfall. Rearranging terms, we get:

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

The wealth effect term captures the impact that first-period rainfall has on second-period irrigation investment that occurs strictly via the impact of first-period rainfall on second-period wealth. The expectations effect term captures the impact that first-period rainfall has on second-period irrigation investment that occurs due to farmers updating their expectations of second-period rainfall. In online appendix 7, it is demonstrated that, for irrigation, the wealth effect is positive, and the expectations effect is negative.[7] Having separated the influences of wealth and expectations, two tests for irrigation adaptation are presented.

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 can be used even without data on wealth. Proposition 3.2, on the other hand, requires data on wealth but is a more powerful test than Proposition 3.1. If farmers are adapting, but the wealth effect dominates the expectation effect, then Proposition 3.2 will detect adaptation but Proposition 3.1 will not.

---

[7] A CARA utility function has been used for tractability purposes. We are not able to prove the signs of the wealth effect and the expectations effect for irrigation or crop choice for a broader range of utility functions.

3.4.2. Test for crop adaptation

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

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

The wealth effect term captures the impact that first-period rainfall has on second-period crop choice that occurs strictly via the impact of first-period rainfall on second-period wealth. The expectations effect term captures the impact that first-period rainfall has on second-period crop choice that occurs due to farmers updating their expectations of second-period rainfall. In online appendix 7, it is demonstrated that, for crop choice, the wealth effect is negative and the sign of the expectation effect is ambiguous. Therefore, it is not possible to test for crop adaptation without controlling for wealth. 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. On the other hand, if wealth is held constant, this removes the wealth effect *and* makes the sign of the expectations effect unambiguously negative, thereby generating 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:* $\frac{d\rho_2^*}{dr_1}\big|_{w_2=constant} < 0$, *then* $\frac{d\mu_2}{dr_1} > 0$.

4. Data sources and summary statistics

We test our model with two agricultural data sets: a household panel and a district panel. The household panel–the Rural Economic and Demographic Survey (REDS)–was collected by the National Council of Applied Economic Research (NCAER). The data cover three rounds (1970/71, 1981/82 and 1998/99) and 259 villages across the 17 major states of India. The second and third rounds of the survey include data on the original households and households that split off from the original households, as well as a sample of new households. The analysis is restricted to households that either were surveyed in multiple rounds or split off from a previously surveyed household. Online appendix 2 displays a map with the locations of the REDS villages and gives more information on the sampling. Online appendix 3 discusses attrition in the REDS data set and its implications for our study. 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). Online appendix 3 includes a map of the districts covered in the World Bank data

Table 1. *Summary statistics*

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

Notes: 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 as to how the variables are constructed.

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 households, agricultural profits per acre are measured as crop receipts minus crop expenses, divided by the area of land cultivated. For the district data set, which lacks information on crop expenses, crop revenue per acre of land cultivated is used. For households,

Table 2. *Daily water requirements of common crops grown in India*

Crop	Daily water requirement	Grown during monsoon
Barley	0.400	No
Cotton	0.525	Yes
Oilseeds	0.350	No
Maize	0.450	Yes
Millet	0.575	Yes
Peanut	0.525	Yes
Potato	0.750	No
Pulses	0.350	Yes
Rice	1.075	Yes
Sorghum	0.575	Yes
Soybean	0.525	Yes
Sugarcane	0.650	Yes
Wheat	0.425	No

Notes: The daily water requirement is measured in centimeters per day. Monsoon crops are crops who are grown during the summer monsoon season.
Sources: Agriinfo.in (2015).

a dummy variable is created for irrigation investment that is equal to one if the household purchased materials, hired labor or used family labor to construct new irrigation assets, purchase new irrigation assets or improve existing irrigation assets during the 12 months prior to the survey interview. For the district data set, which lacks direct information on irrigation investment, irrigation investment is defined as the log of the 1-year change in the area of irrigated land. For households, wealth is measured 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, deflated to 1971 rupees. Online appendix 2 gives additional details about the construction of the agricultural variables.

For both data sets, the water need of a crop portfolio is measured as the area-weighted average of the crop water need of the crops grown in a given year. We 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. An important caveat is that these values refer to optimal amounts, not necessarily what farmers apply in practice.

There are two main growing seasons in India: the kharif (or monsoon) season, which is the primary growing season and runs roughly from June to September, and the rabi (or winter) season, which is secondary and runs from October through February. The bulk of the rainfall comes during the monsoon season, but both kharif and rabi crops are affected by the quantity of monsoon rainfall. Rabi crops are affected because the monsoon

affects soil moisture levels over the winter months, as well as surface water and groundwater supplies that are used to irrigate some winter crops. To address this seasonality, two different measures of crop water need are used: one that includes all crops grown during the year, and an alternate version that only includes the crops grown during the monsoon season. This captures the fact that monsoon crops are likely to be more strongly affected by monsoon rainfall than non-monsoon season crops. Due to the lack of season-specific crop planting data, crops are classified based on the season in which they are typically grown. The monsoon crops are cotton, maize, millet, peanut, pulses, rice, sorghum, soybean and sugarcane and the non-monsoon season crops are barley, mustard, oilseeds, potato and wheat.

The agricultural data is merged with gridded weather data from the Terrestrial Precipitation: Monthly Time Series (1900–2008), version 2.01, and the companion Terrestrial Air Temperature data set. The weather data for each 0.5-degree latitude–longitude grid point combines information from 20 nearby weather stations, using an interpolation algorithm based on the spherical version of Shepard's distance-weighting method. Current year growing season rainfall is measured as the *z*-score deviation from that location's historical mean.[8]

To capture the decadal variability of the monsoon, a decadal rainfall variable is created that is the simple average of the rainfall *z*-scores from the past decade. We also use a specification where we measure lagged rainfall as the number of especially wet or dry years over the past decade. Following Jayachandran (2006), the 20th percentile is used as the cutoff for a dry year and the 80th percentile as the cutoff for a wet year. 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. Online appendix 2 gives additional details about the construction of the weather variables and online appendix 4 presents a detailed discussion of the variation of the decadal rainfall variables.

5. Empirical strategy

5.1. Tests for irrigation adaptation

In order to test for irrigation adaptation, we analyze how irrigation investment responds to lagged rainfall. For the household data set, a regression is run of the form

$$irr\_inv_{ijt} = \alpha_1 decaderain_{jt} + \alpha_2 rain_{jt} + \alpha_3 rain_{jt-1} + \lambda_t + \mu_{ij} + \zeta_{ijt}, \quad (5)$$

where the subscripts represent household i, village j, and year t. The dependent variable irr_inv_{ijt} is a dummy variable equal to 1 if, during the recall period, a household purchased irrigation equipment or used labor

---

[8] 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 (southern) region.

to improve irrigation assets. The explanatory variables are past decade rainfall $decaderain_{jt}$, current year rainfall $rain_{jt}$, one-year lagged rainfall $rain_{jt-1}$, a year fixed effect λ_t, a household fixed effect μ_{ij}, and an error term ζ_{ijt}.[9] The household survey follows households after household splits and after changes of the household head. Therefore, our 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.

$decaderain_{jt}$ is measured in two ways. The first measure is a simple average of the rainfall 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), the 20th percentile is used as the cut-off for a dry year and the 80th percentile as the cut-off for a wet year. The coefficient of interest is α_1. Our 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.[10]

Current year rainfall is controlled for 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 during that period. This response would not indicate adaptation to expected future year rainfall, but would simply reveal within-season adjustment to current year rainfall.[11]

Propositions 3.1 and 3.2 demonstrate that we can test for irrigation adaptation with or without a wealth control. Thus, for completeness, we run a second household specification where we control for wealth. Once we have isolated the wealth effect, our model predicts that $\alpha_1 = 0$ if farmers are not adapting. On the other hand, if farmers are adapting, then $\alpha_1 < 0$.

Despite the dynasty fixed effect, there are two potential sources of endogeneity for $wealth_{ijt}$. First, if the current period's productivity shocks are

[9] Lagged rainfall is included because residual impacts of last year's rainfall may influence this year's irrigation and cropping decisions directly, independently of an expectations/adaptation effect. In addition, the household data set does not include the specific interview date for each household, so including lagged rainfall is important because, for households interviewed early in the survey year, rainfall from the previous calendar year may be the most relevant.

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

[11] We 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.

correlated with lagged productivity shocks, then wealth will be endogenous (because lagged productivity shocks affect wealth). Secondly, as above, if there is variation in farming ability within a dynasty, and farming ability is correlated (within the dynasty) with wealth, this will also cause endogeneity.

We 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). We instrument for $wealth_{ijt}$ with $inhwealth_{ijt}$, the amount of inherited wealth. 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 us variation in $inhwealth_{ijt}$, even in the presence of the dynasty fixed effect.[12] We now discuss the exclusion restriction for our 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 should be less correlated with transient current period productivity shocks than current period wealth. Thus, we should expect the instruments to significantly reduce this source of endogeneity bias.

Secondly, 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, 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 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'. If land is evenly divided, this is suggestive that non-land inheritances may also be divided roughly evenly. Note that even if there is purely equal division of inheritances among sons, changes across generations of household heads provide variation of the inheritances within each dynasty.

We cannot fully prove that contemporaneous variations in inheritances are uncorrelated with unobserved household ability, but to the extent that these inheritances are *less* correlated than current wealth and current irrigated land, our instrumental variables strategy should at least *reduce* the endogeneity bias. Due to potential concerns about the validity of the

---

[12] Amongst the households that are analyzed, 42 per cent come from dynasties with multiple household heads. Our estimate of the coefficient of wealth will be a local average treatment effect based on these households.

instrument, a non-instrumented version of the regression is run in online appendix 6.

For the district regression, irr_inv_{jt} is defined as the log of the one-year change in the district's irrigated area, we use a district fixed effect, and we do not control for wealth. Proposition 3.1 demonstrates that we can test for irrigation adaptation, even in the absence of a wealth control. As with the household regression, finding $\alpha_1 < 0$ provides evidence of adaptation.

5.2. Test for crop adaptation

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

$$water\_need_{ijt} = \gamma_1 decaderain_{jt} + \gamma_2 rain_{jt} + \gamma_3 rain_{jt-1} + \gamma_4 wealth_{ijt}$$
$$+ \tau_t + \phi_{ij} + \psi_{ijt} \tag{6}$$

where $water\_need_{ijt}$ is the area-weighted water need of the farmer's crop portfolio. In one specification, we look at the area-weighted water need of all crops grown during the year. In a second specification, we look at the area-weighted water need of monsoon-season crops only. As mentioned above, we 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.[13]

We control for $wealth_{ijt}$ because, as demonstrated in section 3.4.2, without a control for $wealth_{ijt}$, we could not interpret γ_1 as evidence of adaptation. Wealth is endogenous and we instrument for it with inherited wealth. The validity of the instrument follows the same argument as above. Due to potential concerns about the validity of the instrument, we also run a non-instrumented version of the equation in online appendix 6. Finding $\gamma_1 = 0$ demonstrates that farmers are not adapting their crop portfolios. Conversely, in the presence of adaptation, we expect to find $\gamma_1 > 0$.

5.3. Standard error clustering

Standard errors are clustered at the weather grid-point level, which allows for serial correlation within dynasties and within districts (because any dynasty or district is always associated with the same weather grid point). In addition, this clustering allows for correlation in the error terms between adjacent households (or districts) that share the same value of rainfall due to the resolution of the rainfall data.

In online appendix 6, we verify that the results are robust to standard errors that allow for spatial correlation. This is important because lagged decadal rainfall exhibits substantial spatial correlation, as discussed in online appendix 4. To implement spatial clustering, we use code from

---

[13] We 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.

Hsiang and Solow (2010) and Fetzer (2014). Guided by the autocorrelograms in online appendix 4, we allow for a spatial correlation within 800 km with a Bartlett (triangular) kernel and temporal correlation within a 30-year window.

6. Results

6.1. Tests for irrigation adaptation

Table 3 tests whether farmers adapt their irrigation investments in response to lagged rainfall. Recall that we 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, we 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, we 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 per cent.[14] It is interesting to note that most of the irrigation adaptation effect is driven by dry shocks, rather than wet shocks (as seen in columns 3 and 4). Our theoretical model predicts that farmers should update their rainfall expectations equally much in response to wet or dry rainfall shocks. The larger observed response to dry shocks suggests that these shocks may be more salient to farmers. In addition, it suggests that farmers may be particularly concerned with protecting their farm profits (via irrigation investment) when they expect harmful dry years in the future.

6.2. Test for crop adaptation

In table 4, we test for crop adaptation using the household data set. We control for wealth in all columns and instrument for it with inherited wealth.[15] 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 monsoon season crops are cotton, maize, millet, peanut, pulses, rice, sorghum, soybean and sugarcane. The non-monsoon season crops are barley, mustard, oilseeds, potato and wheat. The columns that use the average rainfall specification are not significant. However, the columns that use the wet/dry shock rainfall specification are significant at the 1 per cent 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

---

[14] 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, the first-stage regression coefficients are not displayed, but they are available upon request.

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

Table 3. *Testing for irrigation adaptation*

Data set Specification	Household FE	Household FE-IV	Household FE	Household FE-IV	District FE	District 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	12,003	11,759	12,003	11,759	8,130	8,130

Notes: Standard errors, in parentheses below the coefficients, allow for clustering within each latitude-longitude rainfall grid point. A dry shock is defined as 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 per cent 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 crop adaptation*

Data set Specification	Household FE-IV	Household FE-IV	Household FE-IV	Household 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 (0.0233)		0.0379 (0.0287)	
Ten-year lagged average of dry shock		−0.159*** (0.0478)		−0.204*** (0.0675)
Ten-year lagged average of wet shock		0.0193 (0.0454)		0.0834 (0.0572)
Current year rainfall	0.0145** (0.00715)	0.0159** (0.00686)	0.0145 (0.0103)	0.0166* (0.00981)
One-year lagged rainfall	−0.0160** (0.00800)	−0.0186** (0.00765)	−0.0163* (0.00945)	−0.0215** (0.00935)
Log non-land wealth (1971 Rs.)	−0.00428 (0.0120)	−0.00522 (0.0109)	−0.00368 (0.0192)	−0.00512 (0.0171)
Fixed effects Year fixed effects	Household Yes	Household Yes	Household Yes	Household Yes
First stage *F*-statistic (Log non-land wealth)	69.22	71.49	66.94	69.38
Observations	5,577	5,577	5,462	5,462

Notes: Standard errors in parentheses allow for clustering within each latitude–longitude rainfall grid point. A dry shock is rainfall below the 20th percentile and a wet shock is rainfall above the 80th percentile. In all columns, we 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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, we 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.[16]

---

[16] Note that our average crop water need regressions use the *daily* amount of water that each crop requires over its growing season. However, different crops have

Table 5. *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	−0.209**	0.0401	0.125**	−0.0257	−0.0554	−0.0843	0.0724	0.0324
of dry shock	(0.0814)	(0.0663)	(0.0540)	(0.0681)	(0.0456)	(0.0774)	(0.0522)	(0.0405)
Ten-year lagged average	0.0789	0.0726*	−0.0172	−0.155**	0.133**	−0.152***	0.0464	−0.00858
of wet shock	(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	−0.00625	0.00616	−0.00346	0.00523	−0.00395	0.0191**	−0.00921	−0.00725
(1971 Rs.)	(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	5,622	5,622	5,622	5,622	5,622	5,622	5,622	5,622

Notes: Standard errors, in parentheses below the coefficients, allow for clustering within each latitude–longitude rainfall 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, we 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

To further understand the crop adaptation results, individual crop regressions are run. In each column of table 5, the dependent variable is the proportion of land that is planted with a specific crop for the top eight crops by area in the REDS data set. These crops are (by area): rice, wheat, pulses, millet, cotton, sorghum, groundnut and maize. Consistent with the table 4 results, we find that after dry decades, farmers plant less rice (1.075 mm/day water need) and more pulses (3.50 mm/day water need). These are the crops with the highest and lowest water needs, respectively. In addition, after wet decades, farmers plant less millet (5.75 mm/day) and sorghum (5.75 mm/day), possibly because they are switching toward planting more (high water need) rice.

We also find that after wet decades, farmers plant more wheat (4.25 mm/day) and more cotton (5.25 mm/day). At first glance, these results seem counterintuitive because wheat and cotton have slightly lower water needs than millet and sorghum. However, wheat is grown during the rabi season, unlike millet and sorghum which are monsoon season crops. Compared to the other main rabi crops (barley and oilseeds), wheat has a higher water need. Thus our results indicate that after wet decades farmers shift their rabi crop portfolio towards higher water need crops. In a similar vein, cotton has a long growing season that spans both kharif and rabi. Famers may be growing more cotton after wet decades because they are planting more cotton, which is in the ground in both seasons, and growing less of the low water need rabi crops of barley and oilseeds.

Taken together, the irrigation and crop regressions suggest evidence that farmers are adapting in response to recent decadal rainfall. However, there are two important caveats to consider. First, our analysis treats irrigation investment and crop choice as two independent decisions, but in reality the two choices are interrelated. Our analysis is not able to completely disentangle the irrigation and crop choice decisions. A second caveat is that, due to attrition, the household sample in REDS is not nationally representative. Specifically, the famers in our sample 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, the adaptation results are an accurate representation of the behavior of this particular population, but our results may not hold more broadly.

7. Robustness
In online appendix 6, we investigate the robustness of our results.

First, we re-run our adaptation regressions using a fixed-effects specification that drops our instrumental variables strategy. In light of potential issues with our instruments, these results provide another set

---

different length growing seasons, which thus affects the total amount of water they require over their total growing season. We can construct the area-weighted average of the total (rather than daily) water requirements of a given crop portfolio. If we use this measure as our dependent variable, we find that the total water requirements go up after wet decades and go down after dry decades, but the coefficients are not statistically significant.

of estimates which may be of interest. In addition, using the non-IV specification allows us to estimate a new set of standard errors that allow for spatial correlation. This is important because, as indicated in online appendix 4, our lagged rainfall variable demonstrates significant spatial correlation. To implement these standard errors, we use code from Hsiang and Solow (2010) and Fetzer (2014). Guided by the autocorrelograms in online appendix 4, we allow for a correlation within 800 km using a Bartlett (triangular) kernel. The results are robust to these changes.

Secondly, we re-estimate our regressions using rainfall lag windows of 5 or 15 years, to verify that the choice of a 10-year window is not driving the results. The district irrigation regressions and the household crop regressions are robust to using alternate rainfall windows. In the household irrigation regressions, however, the signs of the coefficients of interest are preserved but are no longer statistically significant. We also present regressions in which we 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.

Thirdly, we analyze how irrigation investment and crop choice are affected by lagged temperature shocks from the past decade. Temperature shocks affect profits and wealth. Therefore, we would expect to see wealth effects from lagged temperature on irrigation investment and crop choice. However, temperature does not undergo the persistent cycles that precipitation undergoes. Therefore, we would not expect to see an expectations effect for temperature. Our analysis corroborates this story: we find evidence of a temperature-induced wealth effect, but no evidence of a temperature-induced expectations effect.

Fourthly, we discuss the possibility that depletion of groundwater or surface water might be causing the relationship between irrigation investment and lagged rainfall that we have found. Using irrigated area (rather than irrigation investment) as the dependent variable, we find that the area of irrigated land *increases* after dry decades. As is discussed in greater detail in the appendix, this test provides some reassurance that water depletion is not driving the results, but it does not fully rule out that possibility.

Fifthly, we test whether the 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. When we control for the presence of these dams, the irrigation adaptation results are preserved. However, government irrigation investment is an outcome variable and may be endogenous to the household investment decision. For this reason, our results are suggestive that the adaptation results are not driven by public investment, but do not definitively rule out that possibility.

Sixthly, we test whether changes in agricultural technology or policies might be confounding our results. We add controls for high-yielding variety crops, electrification rates, fertilizer prices, financial institutions, agricultural extension services, transportation infrastructure and government intervention in output markets. The irrigation and crop adaptation results are robust to adding these controls. However, these results are

suggestive only, not definitive, since we are only able to control for a subset of possible confounders. Furthermore, the confounders that we do control for are potentially endogenous to the household irrigation and crop decisions.

Lastly, we run our 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 the lagged rainfall variables. Region-by-year fixed effects flexibly control for this. The district irrigation results and the household crop results are robust to the addition of these controls. The coefficients of interest are preserved in sign, magnitude and significance level. In the household irrigation regressions, however, the signs of the coefficients of interest are preserved but the magnitudes decrease by about 50 per cent and the coefficients are no longer statistically significant. The inclusion of region-by-year fixed effects is a stringent test since the household data set only spans three survey years. Therefore it is perhaps unsurprising that our household irrigation regressions are only weakly robust to the inclusion of these controls.

8. Effectiveness of adaptation

The preceding text finds evidence of adaptation; we now quantify its efficacy. What fraction of profits were farmers able to protect from adverse climate variations? To answer this question, we use the household data set to estimate the extent to which irrigation adaptation increased profits during 1971–1999.[17] Rainfall during this period was below average (as shown in figure 1), which reduced profits. On the other hand, farmers adapted to this below-average rainfall, by increasing their irrigation investments, which should have partially offset profit losses.

We take a two-step approach to estimate the efficacy of adaptation. First, we estimate the percentage of profits that were lost due to below-average rainfall during 1971–1999. Next, we estimate the percentage of these losses that were recovered due to increased irrigation investment. To calculate these percentage changes, we estimate profits π from three scenarios: a counterfactual scenario A, in which the drought did not occur; a counterfactual scenario B, in which the drought occurred but farmers did not adapt their irrigation; and the actual scenario C, in which the drought occurred and farmers adapted. The percentage of profits lost due to the below-average rainfall is $\frac{\pi^B - \pi^A}{\pi^A}$. The percentage of lost profits that were recovered by adaptation is $\frac{\pi^C - \pi^B}{\pi^B - \pi^A}$.

---

[17] 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, we can instrument for irrigated land with inherited irrigated land and remove, or at least reduce, this bias.

We observe π^C only in the survey years (1971, 1982 and 1999). To estimate π^C for the non-survey years, we regress survey-year profits on survey-year rainfall, temperature, irrigation and wealth:

$$\pi_{ijt} = \Sigma\beta_{1q} \times (rain_{jt} = q) + \beta_2 propirr_{ijt} + \Sigma\beta_{3q} \times (rain_{jt} = q) * propirr_{ijt}$$
$$+ \beta_4 wealth_{ijt} + \beta_5 temperature_{jt} + \delta_t + \kappa_{ij} + \epsilon_{ijt}. \tag{7}$$

The variable $rain_{jt} = 1, 2, 3, 4, 5$ is a categorical variable representing the rainfall quintile for village j in year t and $rain_{jt} = 3$ is the omitted category. The results of this regression are shown in table A2 of online appendix 5. We use the estimated coefficients and fixed effects to estimate non-survey-year profits:

$$\pi^C_{ijt} = \Sigma\beta_{1q} \times (rain_{jt} = q) + \beta_2 \overline{propirr_{ijt}} + \Sigma\beta_{3q} \times (rain_{jt} = q) * \overline{propirr_{ijt}}$$
$$+ \beta_4 \overline{wealth_{ijt}} + \beta_5 temperature_{jt} + \overline{\delta_t} + \kappa_{ij}, \tag{8}$$

where $\overline{propirr_{ijt}}$, $\overline{wealth_{ijt}}$ and $\overline{\delta_t}$ are linearly interpolated between survey rounds.

To estimate π^B, we compute $\widehat{propirr_{ijt}}$, a counterfactual value of the proportion of irrigated land in the absence of adaptation. We use the coefficients from column 2 of table A10 in online appendix 6 to calculate the adaptive response of irrigation to lagged rainfall.[18] This table is analogous to the baseline irrigation adaptation specification (table 3) 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, we instead use the proportion of irrigation, which is a coarser measure of adaptation. We use this regression to subtract a quantity of 'adapted irrigation' from $\overline{propirr_{ijt}}$, which gives us $\widehat{propirr_{ijt}}$, i.e., what irrigation would have been in the absence of adaptation.[19] We estimate

$$\pi^B_{ijt} = \Sigma\beta_{1q} \times (rain_{jt} = q) + \beta_2 \widehat{propirr_{ijt}} + \Sigma\beta_{3q} \times (rain_{jt} = q) * \widehat{propirr_{ijt}}$$
$$+ \beta_4 \overline{wealth_{ijt}} + \beta_5 temperature_{jt} + \overline{\delta_t} + \kappa_{ij}. \tag{9}$$

---

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

[19] 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 overall trends in irrigation that were driven by changes in technology, etc.

Lastly, to estimate π^A, we calculate expected annual profits, using a 20 per cent chance of each rainfall quintile occurring. This calculation projects what expected profits would have been if rainfall were at its historical mean distribution. We estimate

$$\pi^A_{ijt} = \Sigma\beta_{1q} \times (0.2) + \beta_2\widehat{propirr}_{ijt} + \Sigma\beta_{3q} \times (0.2) * \widehat{propirr}_{ijt}$$
$$+ \beta_4\overline{wealth}_{ijt} + \beta_5temperature_{jt} + \overline{\delta_t} + \kappa_{ij}. \tag{10}$$

We use π^A, π^B and π^C to calculate profit losses and recoveries. For each farmer, we sum total profits during 1971–1999, under each scenario. Using these profit estimates, we find that, on net, the dry regime decreased profits by 0.3 per cent. However, there is substantial heterogeneity and, for households with losses, the average loss was 7.9 per cent. We estimate that farmers with losses recovered, on average, only 9.4 per cent of their losses. Note that this estimate is an upper bound on the effectiveness of irrigation adaptation, because it does not incorporate the cost of the irrigation investments.

Our results demonstrate that farmer adaptation to persistent rainfall deviations was of limited efficacy. This suggests that adaptation to anthropogenic climate change may also be limited. However, there are four caveats to note. First, we only analyze irrigation adaptation. Farmers are likely to adapt via other channels as well, which will allow them to recapture additional profits. Secondly, extrapolating our results directly to future climate change is problematic, since future climate change will affect both rainfall and temperature. Thirdly, in addition to changing the demand for irrigation water, climate change will also affect the supply of irrigation water (Zaveri *et al.*, 2016). In some regions, the available supply of irrigation water will decrease, potentially limiting the ability of farmers to increase their irrigation. Lastly, rich farmers with above-average land areas are over-represented in our sample. However, if we expect wealthier households to adapt more easily than poor households, then our estimate would overstate the efficacy of irrigation adaptation. Our results indicate that, even among richer households, irrigation adaptation was of limited benefit.

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, we 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. Evidence is found of both irrigation adaptation and crop adaptation. However, analysis suggests that the efficacy of adaptation is limited; we estimate that adaptation recovers at most 9 per cent of lost profits. These results suggest that, in the context of the historical rainfall deviations that were analyzed, there are barriers to adaptation. This 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.

Supplementary material and methods

To view supplementary material for this article, please visit https://doi.org/10.1017/S1355770X17000195.

References

Adams, R.M., B.H. Hurd, S. Lenhart, and N. Leary (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', [Available at] http://www.agriinfo.in/default.aspx?page=topic&superid=7&topicid=16.

Barreca, A., K. Clay, O. Deschênes, M. Greenstone, and J.S. Shapiro (2015), 'Adapting to climate change: the remarkable decline in the U.S. temperature-mortality relationship over the 20th Century', NBER working Paper No. 18692, Cambridge, MA.

Bryan, E., T.T. Deressa, G.A. Gbetibouo, and C. Ringler. (2009). 'Adaptation to climate change in Ethiopia and South Africa: options and constraints', *Environmental Science & Policy* **12**(4): 413–426.

R. Burgess, O. Deschênes, D. Donaldson, and M. Greenstone (2014), 'The unequal effects of weather and climate change: evidence from mortality in India', Unpublished working paper.

Burke, M. and K. Emerick (2016), 'Adaptation to climate change: evidence from US agriculture', *American Economic Journal: Economic Policy* **8**(3): 106–140.

Christensen, J.H. and B. Hewitson (2007), 'Regional climate projections', in S. Solomon, D, Qin, M. Manning, Z. Chen, M. Merquis, K.B. Areryt, M. Tignor and H.L. Miller (eds), *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge: Cambridge University Press (pp. 847–940).

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

da Cunha, D.A., A.B. Coelho, and J.G. Féres (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* (3rd edn), New Delhi: National Book Trust.

Dell, M., B.F. Jones, and B.A. Olken (2012), 'Temperature shocks and economic growth: evidence from the last half century', *American Economic Journal: Macroeconomics* **4**(3): 66–95.

Dell, M., B.F. Jones, and B.A. Olken (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', University of Illinois.

Deschênes, O. and M. Greenstone (2007), 'The economic impacts of climate change: evidence from agricultural profits and random fluctuations in weather', *American Economic Review* **97**(1): 354–385.

Deschênes, O. and M. Greenstone (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. and M. Veronesi (2013), 'How can African agriculture adapt to climate change? A counterfactual analysis from Ethiopia', *Land Economics* **89**(4): 743–766.

Di Falco, S. and M. Veronesi (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.

Di Falco, S., F. Adinolfi, M. Bozzola, and F. Capitanio (2014), 'Crop insurance as a strategy for adapting to climate change', *Journal of Agricultural Economics* **65**(2): 485–504.

Duflo, E., M.R. Kremer, and J. Robinson (2011), 'Nudging farmers to use fertilizer: theory and experimental evidence from Kenya', *American Economic Reiew* **101**(6): 2350–2390.

Fernando, A.N. (2014), 'Shackled to the soil: the long-term effects of inherited land on labor mobility and consumption', Harvard University.

Fetzer, T. (2014). 'Can workfare programs moderate violence? Evidence from India', STICERD Discussion Paper No. 53, London.

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

Fishman, R. (2015), *When Water Runs Out: Scarcity, Adaptation and Migration in Gujarat*, 4th IGC-ISI India Development Policy Conference.

Fishman, R. (2016), 'More uneven distributions overturn benefits of higher precipitation for crop yields', *Environmental Research Letters* **11**(2): 1–7.

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

Foster, A. and M. Rosenzweig (1995), 'Learning by doing and learning from others: human capital and technical change in agriculture', *Journal of Political Economy* **103**(6): 1176–1209.

Foster, A. and M. Rosenzweig (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 M.R. Rosenzweig (2002), 'Household division and rural economic growth', *Review of Economic Studies* **69**(4): 839–869.

Foster, A.D. and M.R. Rosenzweig (2010), 'Is there surplus labor in rural India?', Discussiion Paper No. 991, Yale University.

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

Guiteras, R. (2009), 'The impact of climate change on Indian agriculture', Unpublished manuscript, Department of Economics, University of Maryland.

Hornbeck, R. (2012), 'The enduring impact of the American Dust Bowl: short- and long-run adjustments to environmental catastrophe', *American Economic Review* **102**(4): 1477–1507.

Hornbeck, R. and S. Naidu (2014), 'When the levee breaks: black migration and economic development in the American South', *American Economic Review* **104**(3): 963–990.

Hsiang, S.M. and D. Narita (2012), 'Adaptation to cyclone risk: evidence from the global cross-section', *Climate Change Economics* **3**(2): 1–28.

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

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

Krishnamurthy, C.K.B. (2012), 'The distributional impacts of climate change on Indian agriculture. A quantile regression approach', Working Paper No. 69/2012, Madras School of Economics.

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

Kurukulasuriya, P., N. Kala, and R. Mendelsohn (2011), 'Adaptation and climate change impacts: a structural ricardian model of irrigation and farm income in Africa', *Climate Change Economics (CCE)* **2**(2): 149–174.

Luis, M.G. and R. Orlando (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., W.D. Nordhaus, and D. Shaw (1994), 'The impact of global warming on agriculture: a Ricardian analysis', *American Economic Review* **84**(4): 753–771.

Mukherjee, M. and K. Schwabe (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., C.R. Ranganathan, and U.S. Nagothu (2014), 'Climate Change and Agriculture in India: Studies from Selected River Basins. New Delhi: Routledge Taylor & Francis.

Pant, G.B. and K.R. Kumar (1997), *Climates of South Asia*. Hoboken, NJ: 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: 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. and R. Mendelsohn (2008), 'The impacts of global warming on farmers in Brazil and India', *Global Environmental Change* **18**(4): 655–665.

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

Schlenker, W. and M.J. Roberts (2009), 'Nonlinear temperature effects indicate severe damages to US crop yields under climate change', *Proceedings of the National Academy of Sciences* **106**(37): 15594–15598.

Schlenker, W., W.M. Hanemann, and A.C. Fisher (2005), 'Will US agriculture really benefit from global warming? Accounting for irrigation in the hedonic approach', *American Economic Review* **95**(1): 395–406.

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

Seo, S.N., B.A. McCarl, B.A. McCarl, and R. Mendelsohn (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.

Staiger, D. and J.H. Stock (1997), 'Instrumental variables regression with weak instruments', *Econometrica* **65**(3): 557–586.

Stock, J.H. and M. Yogo (2002), 'Testing for weak instruments in linear IV regression', NBER working paper No. 284, Cambridge, MA.

Subbaramayya, I. and C. Naidu (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.

Viswanathan, B. and K.S. Kavi Kumar (2015), 'Weather, agriculture and rural migration: evidence from state and district level migration in India', *Environment and Development Economics* **20**(4): 469–492.

Wang, B. (2006), *The Asian Monsoon*. Chichester: Springer Praxis Books.

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

Zaveri, E., D.S. Grogan, Fisher-Vanden, K., S. Frolking, R.B. Lammers, D.H. Wrenn, A. Prusevich, and R.E. Nicholas (2016), 'Invisible water, visible impact: groundwater use and Indian agriculture under climate change', *Environmental Research Letters* **11**(8): 1–13.