



When it Rains, it Pours: Estimating the Spatial Spillover Effect of Rainfall

Farzana Hossain¹ · Reshad N. Ahsan²

Accepted: 22 March 2022 / Published online: 16 April 2022
© The Author(s) 2022

Abstract

A large fraction of the world's poor rely on rain-fed agriculture, which makes them vulnerable to changes in rainfall patterns. In this paper, we examine whether spatial correlation in rainfall results in these households also being vulnerable to an adverse spatial-spillover effect. In particular, we use household-level panel data from India along with high-resolution meteorological data to show how rural household consumption varies with own-region rainfall as well as rainfall in neighboring areas. We find that while greater own rainfall has a positive effect on rural household consumption, greater rainfall in neighboring regions has an adverse spatial-spillover effect. Our results suggest that when this spillover effect is taken into account, the positive impact of own-region rainfall on household consumption falls by 38 percent.

Keywords Rainfall · Rural households · Climate

JEL Classification Q12 · O13 · Q54

✉ Farzana Hossain
farzana.hossain@monash.edu; farzanah@student.unimelb.edu.au
Reshad N. Ahsan
rahsan@unimelb.edu.au

¹ Centre for Health Economics, Monash University, 900 Dandenong Road, Caulfield East, VIC 3145, Australia

² Department of Economics, University of Melbourne, Level 3, FBE Building, 111 Barry Street, Melbourne, VIC, Australia

1 Introduction

In 2015, more than 735 million people globally lived in extreme poverty, with 79 percent of them in rural areas (World Bank 2018).¹ Many of these rural households rely on rain-fed agriculture, making them especially vulnerable to production and price risk due to changes in rainfall patterns. In this paper, we examine whether spatial correlation in rainfall can result in these households also being vulnerable to an adverse spatial-spillover effect.

To understand why spatial correlation in rainfall matters, consider a farmer residing in region d . Greater rainfall in d will increase this farmer's crop output and, for a given crop price, increase her income. But the crop price received by farmers in d will likely depend on rainfall in its neighboring districts. With greater neighboring rainfall, production of the same crop will increase, which will lower crop prices, and with inelastic demand, also lower farm incomes. On the other hand, for a given income, the lower crop price will increase a farmer's purchasing power. Thus, the overall impact of this spatial-spillover effect is an empirical question.

To examine this spatial-spillover effect empirically, we use household-level panel data from India along with high-resolution meteorological data. Our choice of India as a setting for this analysis provides us with two important benefits. As a geographically large country, India experiences significant spatial and temporal variation in weather patterns. As we document below, this ensures that we have sufficient variation in rainfall to identify our key results. Second, agricultural production in India is mainly rain-fed and the sector plays a dominant role in the overall economy. For instance, agriculture accounts for 49 percent of India's total employment and 52 percent of agricultural land is rain-fed (Economic Survey 2018). Thus, any adverse spatial spillover effect of rainfall is likely to be of first-order importance here.

To capture rainfall in district d 's neighboring regions, we calculate the cumulative weighted rainfall in all districts $j \neq d$, where the weights are the inverse straight-line distance between d and j . This flexible approach places greater weight on rainfall in nearby districts without requiring us to define which neighbors matter. We then examine whether a household's consumption depends on rainfall in neighboring districts as well as rainfall in its own district. Our identification strategy incorporates household fixed effects, which means that our results are identified from within-district deviations in own and neighboring district's rainfall from its long-term average. These deviations are likely to be orthogonal to unobserved determinants of rural household consumption.

We find that a one standard deviation increase in own-district rainfall increases household consumption by 8.46 percent. However, we also find that rainfall shocks in neighboring districts attenuate this positive effect. Indeed, when we account for this adverse spatial spillover effect, we find that the same increase in own-district rainfall increases household consumption by only 5.23 percent. This is approximately 38 percent lower than the benchmark case with no spatial spillovers. These results, therefore, suggest that spatial correlation in rainfall creates economically meaningful general equilibrium effects that attenuate the overall consumption gains from rainfall.

Our results are robust to controlling for a district's average temperature and to adjusting standard errors for spatial correlation following (Conley 1999). To further guard against

¹ According to World Bank (2018), people who live on less than USD 1.90 a day (2011 PPP) are considered to be in extreme poverty. The latest year for which these estimates are available globally is 2015.

spurious spatial correlation, we show that our results are robust to including a spatial lag of average household consumption and to a falsification test where we regress crop yields in a district d on neighboring rainfall. Our hypothesis is that rainfall in neighboring districts can have adverse effects on rural households in d via a reduction in crop prices. Rainfall in neighboring districts should not directly affect crop yields in d . Our results indicate that this is indeed the case and suggest that our core result is not being driven by spurious spatial correlation.

To empirically examine the key mechanisms that explain the spatial-spillover effect, we first use district-level crop-price data to confirm that a rainfall shock in neighboring districts lowers crop prices received by farmers. Next, we show that households that experience a higher neighbor's rainfall shock earn lower income from selling crops at the market and are also less likely to participate in markets. We find no such effect on non-agricultural income as well as remittances, which suggests that our results are not being driven by other shocks that are correlated with rainfall.

Our paper is related to a small but growing literature that documents the spatial-spillover effect of weather shocks. For instance, Harari and La Ferrara (2018) examine the spillover effect of weather shocks on conflict in Africa. Elliott et al. (2019) examine the spillover effect of typhoons on manufacturing firms in China while Boustan et al. (2020) construct a measure of natural disasters for U.S. counties that account for both own-county disasters as well as disasters that occur in nearby counties. Ours is the first paper to examine the spatial spillover effect of weather shocks on rural agricultural households. Individuals in such households constitute the majority of people living in extreme poverty globally and are especially vulnerable to weather shocks due to climate change. Our approach is also related to past efforts at estimating spatial spillovers of other shocks and interventions (see e.g. Miguel and Kremer 2004).

Our paper is also related to a literature that examines the effect of climate change induced variation in temperature and rainfall on agricultural outcomes using both simulation methods (Adams 1989) and regression analysis (Mendelsohn et al. 1994; Schlenker et al. 2006; Deschênes and Greenstone 2007; Dell et al. 2012, 2014). Our paper is also related to a literature that documents the welfare consequences of weather shocks in developing countries. These studies find that weather shocks affect agricultural production, employment, and wages (Jayachandran 2006; Emerick 2018; Kaur 2019; Colmer 2021) as well as human capital (Maccini and Yang 2009; Shah and Steinberg 2017). We contribute to these literatures by showing that rainfall can have economically significant adverse spatial spillover effects.²

Indeed, our results also have an important methodological implication for the literature, where the typical approach is to regress an outcome of interest on rainfall in its own region. Thus, rainfall in neighboring districts is implicitly included in the error term. Our results suggest that in the presence of spatial correlation, such an econometric model is misspecified. This point is also relevant when estimating the effects of other weather and environmental shocks with substantial spatial correlation.

We structure the rest of the paper as follows. In Sect. 2, we discuss the mechanisms that can lead to rainfall in neighboring districts lower rural household welfare. In Sect. 3,

² Finally, our paper is related to a literature that examines agricultural households in developing countries with an emphasis on their market participation decision. Singh et al. (1986), de Janvry et al. (1991), Key et al. (2000), Taylor and Adelman (2003), Fafchamps and Hill (2005), and Barrett (2008), among others, are key contributions to this literature.

we describe our household-level panel data as well as our rainfall data. In this section, we also describe how we construct our own-rainfall and neighbor's rainfall shock variables. In Sect. 4, we describe the empirical strategy we use to identify the impact of rainfall shocks on household consumption. In Sect. 5, we present our baseline results and address key econometric issues while in Sect. 6 we provide supporting evidence for the mechanisms that are driving the spatial spillover effect. In Sect. 7, we explore additional results and robustness checks while in Sect. 8 we provide a conclusion.

2 Conceptual Framework

In this section, we use an agricultural household model to show that increased rainfall in neighboring districts will have an ambiguous effect on rural household consumption. The distinguishing feature of an agricultural household is that it is both a consumer as well as a producer of agricultural products (Singh et al. 1986; de Janvry et al. 1991).

2.1 Production

To begin, consider a risk-neutral farmer that produces a homegrown crop, H , which we refer to as a home crop from hereon. At the beginning of the growing season, the farmer must decide how much of the home crop to produce. Her household's welfare in this initial period, $U_1(Q_H, T_L)$, depends on her crop yield at harvest time, Q_H , as well as the leisure time enjoyed during the growing season, T_L . To produce output, the farmer can use her fixed arable land as well as family labor, T_N . We normalize the arable land to one and assume that the farmer cultivates this land in its entirety. We also assume that the household has a time endowment of T with $T_L + T_N = T$.

Suppose her output at harvest time is determined by $Q_H = \hat{R}^O f(T_N)$, where f is a production function and \hat{R}^O is the farmer's expectation of the rainfall in her own district during the growing season. Thus, \hat{R}^O serves as a Hicks-neutral productivity shifter.³ The farmer's objective is to pick the family labor, T_N , that maximizes $U_1(Q_H, T_L)$ subject to the time-endowment constraint and the production technology. Let the resulting optimal family labor choice be

$$T_N^* = \theta T,$$

where $0 < \theta < 1$ is the fraction of the family's time endowment spent on farm work.

2.2 Price Determination and Consumption

At harvest time, the actual own-district rainfall during the growing season is realized. This value of R^O along with the optimal labor chosen earlier, T_N^* , pins down the farmer's crop output, $Q_H = R^O f(T_N^*)$. From hereon, we only consider farmers for whom this crop output exceeds their own desired consumption, so that they have a surplus to sell. Farmers that

³ To concentrate on the mechanisms we can analyze with our data, we consider a stylized model where we abstract from the use of intermediate inputs and the option to cultivate multiple crops. However, when discussing the household's consumption below, we introduce a market crop that the household does not produce and can only purchase from the market.

do not have a surplus, and hence do not sell, will not be impacted by the price decline due to increasing neighboring rainfall. Indeed, we verify that this is the case in our empirical analysis below.

With her surplus in hand, the farmer can transport her crop to the nearest *mandi*, which is a government-regulated wholesale market. *Mandis* are open auctions and are attractive to farmers as it minimizes vulnerability to unscrupulous buyers. However, for the typical farmer in India, *mandis* are costly to get to (Goyal, 2010). To capture this travel cost, we assume that the farmer incurs an iceberg transport cost of $\tau > 1$ to sell her crops at the market. Thus, if she sells at the market, she will receive a price of P/τ , where P is the prevailing market price that the farmer takes as given.

If the market price is low or if the transport cost is high, the farmer can alternatively sell her surplus to local traders at the farm gate (Fafchamps and Hill 2005). We assume that there are no transport costs associated with selling to traders and thus the price that the farmer will receive in this scenario is $P_T < P$. It follows that the farmer will participate in the market if and only if $P \geq \tau P_T$ and therefore, the equilibrium price she will receive is

$$P_H = \begin{cases} P/\tau & \text{if } P \geq \tau \bar{P}_H, \\ P_T & \text{otherwise.} \end{cases} \tag{1}$$

Given the above, the farmer’s income will be $P_H(\bar{Q}_H - C_H)$, where $\bar{Q}_H - C_H > 0$ is her marketed surplus. Suppose that her household has preferences over the consumption of the home crop and a large variety of market crops. We assume that the farmer does not grow the market crop varieties and must purchase them at a market. Let her preferences for these crops be given by

$$U_2 = C_M^\eta C_H^{1-\eta}, \tag{2}$$

where C_M is the aggregate consumption of market crops.⁴ The farmer’s objective is to maximize (2) subject to the following budget constraint:

$$P_M C_M = P_H(\bar{Q}_H - C_H),$$

where P_M is a price index that captures the price of market crops. Importantly, the household does not have any market expenditure on C_H as it consumes from its own production. The tradeoff it faces is that while higher C_H consumption increases its utility, it lowers its farm income and hence consumption of the market crop. The household’s optimization problem results in the following demand for the home crop:

$$C_H = (1 - \eta)\bar{Q}_H.$$

That is, C_H only depends on crop yield, \bar{Q}_H , and is not directly impacted by changes in the market price. Given her optimal consumption of C_H , the residual income left over to spend on the market crop varieties is $Y_M = \eta P_H \bar{Q}_H$. This results in the following aggregate demand for market crops:

⁴ C_M could be a CES aggregate over a large number of market crop varieties, C_M^k , where k indexes varieties. C_M could also include non-food items that are purchased at the market.

$$C_M = \frac{\eta P_H \bar{Q}_H}{P_M}. \quad (3)$$

Notice that the consumption of the market crop depends on both its own price as well as the level of farm income, which in turn is a function of P_H .⁵

2.3 Impact of Neighboring Rainfall

Now consider the impact of greater rainfall in neighboring districts, R^N . For farmers in these neighboring districts, greater rainfall will act as a positive productivity shock and result in an increase in regional crop output. In turn, this will reduce both P and P_M (Burgess and Donaldson 2012).^{6,7} To see how this will affect her household consumption, consider the case where the farmer continues to sell at the market after the fall in P . For the market crop, we can use (3) to decompose the effect of the lower prices into the following channels:

$$\frac{dC_M}{dR^N} = \underbrace{-\left(\frac{\eta P_H \bar{Q}_H}{P_M^2}\right)\left(\frac{\partial P_M}{\partial R^N}\right)}_{\text{Own-Price Effect}} (> 0) + \underbrace{\left(\frac{\eta \bar{Q}_H}{P_M}\right)\left(\frac{\partial P_H}{\partial R^N}\right)}_{\text{Farm-Income Effect}} (< 0). \quad (4)$$

The first term on the right-hand-side says that by lowering the own price, greater R^N will raise C_M . This is the own-price effect. The second term says that by lowering the income earned from its home crop, greater R^N will lower C_M .⁸ This is the farm-income effect (Singh et al. 1986). Thus, the net effect of greater rainfall in neighboring districts on this farmer's overall consumption is ambiguous and is ultimately an empirical question.

Before moving on to the empirical analysis, it is worth pointing out that rainfall in neighboring districts can affect farm income through both the intensive margin above as well as the extensive margin. To see the latter, note from (1) that if greater neighboring rainfall lowers the market price P , then the transportation-cost inclusive market price received by farmers may fall below the price received from traders. If so, the household will prefer not to sell to the market. We will empirically explore this extensive margin effect in Sect. 6 below.

⁵ Abidoye and Cali (2021) also discusses the contrasting effects of price changes on consumption and income respectively in the context of conflict in Nigeria.

⁶ We assume that the farmer's district is small in the regional economy. This means that the market price of both the home and market crops will not depend on her own district's rainfall. We report crop-price regressions below that support this assumption.

⁷ We are abstracting from the presence of minimum support prices (MSP) that the Indian government uses to place a floor on the price of certain agricultural commodities. In principle, such a policy will attenuate the effect of R^N on P by setting a lower bound on the latter. Thus, incorporating this in our framework above will not qualitatively change our predictions. We chose to abstract from this as there is evidence that these MSPs are not fully effective. For instance, Aditya et al. (2017) show that less than 25 percent of farmers in their data are even aware of what the MSP is for their crops.

⁸ Even when the output of the home crop is not fixed, lower P_H will lower farm income as long as the demand for the home crop is price inelastic. Kumar et al. (2011) shows that this is indeed the case for staple crops in India.

3 Data

3.1 Household Data

We use household data from the Indian Human Development Survey (IHDS). IHDS is a nationally representative longitudinal household survey that is available for two rounds, 2004–05 and 2011–12 (Desai and Vanneman 2005, 2012). The raw data cover 1,503 villages and 971 urban areas across India. However, given that we are interested in the effect of rainfall on agricultural household consumption, we restrict our sample to rural households that are observed in both periods. The restriction to rural households ensures that we only consider rainfall shocks in neighboring districts that grow crops in our analysis.⁹ This results in a working sample that consists of 28,087 households in 283 districts across India.¹⁰

Our key outcome variable is each household's total annual consumption expenditure per capita. IHDS constructs this by dividing each household's total expenditure on a series of food and non-food items by the number of household members. We converted these nominal values to real ones using the deflator provided by IHDS, which gives us annual consumption expenditure per capita in constant 2005 Rupees.

Unlike other commonly used household surveys in India, the IHDS data have the advantage that it follows households over time. Apart from enabling us to control for time-invariant household characteristics, the panel nature of the data allows us to use a balanced sample of households that appear in both survey rounds. This ensures that our key results are not being driven by compositional changes in the sample. A second advantage of our household-level data is that it avoids the attenuation bias that is present in more aggregated analysis of weather shocks and agricultural outcomes (Fezzi and Bateman 2015).

In panel A of Table 1, we provide descriptive statistics of the households in our IHDS sample. The average household has monthly consumption expenditure of approximately 824 Rupees per person, which is equivalent to 18.94 U.S. dollars per person in 2005. In addition, the average household has 5.43 members and 1.77 children. On average, 88.40 percent of households have a male head with an average age of 48.70.¹¹

Our analysis rests on the assumption that households in our sample produce crops for sale in agricultural markets. Without such market participation, we would not expect changes in market prices to impact household consumption. Similarly, if the households in our sample reside in isolated areas that are far removed from local markets, then rainfall-induced price changes in neighboring markets may have little impact on local prices. To explore these issues further, we report summary statistics on crop sales and market access indicators in Table 2.¹² In panel A, we show that 55 percent of households in our sample report agriculture as their main source of income while 10 percent of households are

⁹ The restriction that only households appearing in both survey rounds are included is relaxed for new households in 2011–2012 that split off from households in 2004–2005. We retain these households even though they appear in only IHDS2.

¹⁰ To minimize measurement error in our data, we eliminate households that report negative values for consumption per capita, educational expenses, and medical expenses. We also omit households that report negative values for whether or not they own/cultivate land.

¹¹ Summary statistics by survey year are provided in Table 10 in the Appendix.

¹² The data we use to construct the summary statistics in this table were only collected during the 2004–2005 IHDS round.

sharecroppers. Further, 47 percent of the households in our sample sell their crops with these sales representing, on average, 31 percent of their total production. These numbers suggest that while the households in our sample are poor, they are nonetheless actively involved in agricultural markets.

In Panel B of Table 2, we examine how isolated the households in our sample are from nearby markets. Unfortunately we do not have household-level data on the distance to the nearest market, so instead we use several village-level proxies of market access. These results suggest that 94.32 percent of villages in our sample are accessible by road. Further, on average, the villages in our sample are 6.37 kilometers away from the nearest retail market and 14.26 kilometers away from the nearest town. This suggests that the households in our sample are not so isolated that we can dismiss the pass through of rainfall-induced price changes in neighboring markets on to local prices.

3.2 Rainfall Data

We pair our household data with rainfall data from the ERA-Interim Reanalysis Archive. These daily data are available at a $0.25^\circ \times 0.25^\circ$ grid level for the period 1979 to 2015 (Dee et al. 2011). These reanalysis data combine ground station and satellite data with results from global climate models to create consistent measures of precipitation at a spatially granular level (Auffhammer et al. 2013). When compared to standard rainfall data from ground stations, using such reanalysis data has the advantage that we do not need to worry about the endogenous placement of ground stations as well as spatial variation in the quality and quantity of rainfall data that is available (Colmer, forthcoming).

To merge these data with our IHDS household survey data, we first overlay the GIS boundaries of each district in our IHDS sample on the gridded climate data. We then calculate the total rainfall in each district by using the weighted average across all grids that fall within a district. The weights are the inverse distance between each district's centroid and each grid point. Finally, we sum the daily rainfall data over the period June to September to calculate total monsoon rainfall for each district in our sample in a given year. In Fig. 1, we plot the trend in average monsoon rainfall in our sample over the period 1979 to 2011. As is evident from this figure, average rainfall in India has been increasing during this period. Further, there is also substantial year-to-year variation in monsoon rainfall.

To capture a district's own rainfall shock, we follow Barrios et al. (2010) and Emerick (2018) and create a rainfall anomaly measure for each district. This anomaly measure captures the deviation in a district's monsoon rainfall in any given year from the long-term monsoon average and is normalized by the long-term standard deviation. More precisely, for a district d in year t , we define its own rainfall shock as

$$R_{dt}^O = \frac{R_{dt} - \bar{R}_d}{S_d}, \quad (5)$$

where R_{dt} is the total monsoon rainfall in a district in year t and \bar{R}_d is the district's average monsoon rainfall over the entire period for which we have data (1979 to 2015). Similarly, S_d is each district's monsoon rainfall standard deviation during the 1979 to 2015 period. Thus, a higher value of R_{dt}^O indicates that a district received total monsoon rainfall in a year that was above its long-term average.¹³

¹³ In addition to using this rainfall anomaly shock, we also follow Jayachandran (2006) and construct a categorical variable that takes the value of one if a district's rainfall in year t is above its 80th percentile

Table 1 Summary statistics of IHDS households

	(1)	(2)	(3)	(4)	(5)
	Obs.	Mean	SD	Min.	Max
<i>Panel A: household characteristics</i>					
Household consumption per capita	55,192	823.97	853.83	4.00	42,913.67
Indicator for households in poverty	55,205	0.22	0.41	0.00	1.00
Number of household members	55,205	5.43	2.82	1.00	38.00
Number of children	55,205	1.77	1.70	0.00	18.00
Male household head	55,205	0.88	0.32	0.00	1.00
Household head's age	55,205	48.70	13.63	1.00	100.00
<i>Panel B: weather</i>					
Total rainfall (in metres)	54,554	0.13	0.08	0.01	0.63
Own rainfall shock	54,554	0.16	0.89	- 1.81	3.21
Neighbor's rainfall shock	54,554	0.14	0.57	- 0.73	0.95
Own positive shock	54,554	0.26	0.44	0.00	1.00
Neighbor's positive shock	54,554	0.26	0.23	0.01	0.63
Average temperature	53,880	25.57	2.54	5.86	28.89

In panel A, we report summary statistics for various household indicators. These averages were taken over both survey rounds. In panel B, we report the summary statistics for various rainfall shock measures and average monthly temperature. Household consumption per capita is calculated using monthly household expenditure while the poverty indicator is as provided by the IHDS and is calculated using the Indian Planning Commission poverty line. Household consumption per capita is reported in constant 2005 Indian rupees. 1 U.S. dollar is approximately equal to 43.5 Indian rupees in 2005. Total rainfall is reported in meters. Own rainfall shocks and neighbor's rainfall shocks are calculated using the anomaly approach described in the text. Note that the numbers reported above incorporate sampling weights to ensure that they reflect population averages

In Fig. 2, we illustrate the spatial variation in rainfall in India by plotting rainfall anomaly shocks at the district level by year. These maps yield two important insights. First, it highlights the inter-temporal variation in rainfall during our sample period. For instance, we observe that 2005 was a relatively dry year compared to 2011. This figure also makes clear the significant within-district variation in the data.

The second important insight is that rainfall is highly spatially clustered. From Fig. 2 we can see that in 2004 the low rainfall shocks were clustered in the north and south-west regions of India. In 2011, the higher rainfall shocks were concentrated in the central and south-west regions of the country. This spatial clustering of rainfall reinforces the point that if a household's own district receives a high (low) rainfall shock, then nearby districts are also highly likely to receive a high (low) rainfall shock. This suggests that to correctly account for the overall effect of rainfall on household welfare, one must also account for rainfall in nearby areas.

To examine this spatial spillover effect, we use the following measure of rainfall in neighboring districts:

Footnote 13 (continued)

rainfall value over the period 1979 to 2015. All other districts have an own rainfall shock value of 0 (no rainfall shock)

Table 2 Agricultural production and market access in the 2004–2005 IHDS sample

	(1)
<i>Panel A: household's agricultural production</i>	
Fraction with agriculture as main income	0.552
Fraction that are sharecroppers	0.102
Fraction that sell crops	0.472
Share of output sold	0.310
	[0.377]
<i>Panel B: village's access to markets</i>	
Indicator for road-accessible villages	0.943
	[0.232]
Distance to retail market	6.368
	[6.881]
Distance to nearest town	14.261
	[11.226]

In panel A, we report the sample averages and standard deviation in brackets for various household agricultural production characteristics while in panel B, we report the sample averages and standard deviation in brackets for various aspects of a village's access to markets. The fraction with agriculture as main income variable is averaged over both IHDS rounds while the remaining data are only available in the 2004–2005 IHDS round. Note that the numbers in panel A incorporate sampling weights to ensure that they reflect population averages. The numbers in panel B are calculated at the village level. All distances reported in the table are in kilometers




Fig. 1 Trends in average annual rainfall in India (1979–2011)

$$R_{dt}^N = \sum_{j \neq d} \left(\frac{1}{\omega_{dj}} \times R_{jt}^O \right) \tag{6}$$

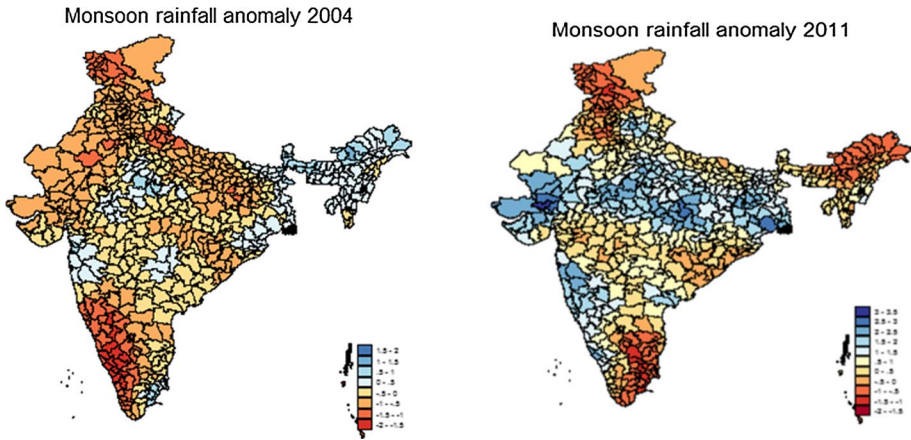


Fig. 2 Spatial variation in rainfall anomalies in India. Anomalies are defined as the difference between a district’s rainfall during June to September in a year and its average rainfall between 1979 to 2015 divided by the standard deviation of its rainfall over the same period. Thus, a higher value (bluer color) represents greater than average rainfall

where j indexes all other districts in the sample and ω_{dj} is the straight-line distance (in kilometers) between the centroids of d and j . We normalize this distance to ensure that the ratio $1/\omega_{dj}$ sum to one. Finally, R_{jt}^O is the own rainfall shock in neighboring district j in year t . Note that we exclude district d ’s own rainfall, R_{dt}^O , when calculating R_{dt}^N .

Thus, for each district d in year t , Eq. (6) provides us with a weighted average of rainfall shocks experienced by all other districts in the sample, where the weights are the inverse of the distance between d and j . These inverse distance weights ensure that rain shocks in nearby districts play a greater role in determining the size of R_{dt}^N .¹⁴ An advantage of measuring neighbor’s rainfall using (6) is that it includes all other districts j in the calculation with faraway districts having a low weight due to the greater distance. Importantly, with R^N defined as in Eq. (6), we do not have to make an *ad hoc* decision on which neighbor’s to include.

The correlation coefficient between a district’s own rainfall shock, R_{dt}^O , and its neighbor’s rainfall shock, R_{jt}^N , is 0.77. Such a high correlation follows naturally from the spatial clustering of rainfall evident in Fig. 2. This is further confirmed by the Moran’s I statistic for own-district rainfall, which yields a z-score of 10.81 and is statistically significant at the 1 percent level. Thus, we can comfortably reject the null of no spatial autocorrelation in rainfall. Summary statistics for all rainfall variables used in the paper are reported in panel B of Table 1.

4 Econometric Strategy

To examine the effect of both own rainfall shocks and neighbor’s rainfall shocks on household consumption, we use the following specification:

¹⁴ This measure of neighbor’s rainfall shock builds on measures of market access that is frequently used in the trade and economic geography literature. See Donaldson and Hornbeck (2016) for a recent example.

$$\ln(C_{hdt}) = \alpha + \beta_1 R_{dt}^O + \beta_2 R_{dt}^N + \gamma_1 X_{hdt} + \theta_h + \theta_t + \epsilon_{hdt} \quad (7)$$

where C_{hdt} is the total consumption for household h in district d and year t , R_{dt}^O is a district d 's own district rainfall shock, and R_{dt}^N is the rainfall shock in neighboring districts. Our coefficient of interest is β_2 , which will be negative if a positive rainfall shock in neighboring districts has an adverse effect on a household's consumption.

We include in (7) a set of household- and district-level controls, X_{hdt} , that are likely to affect consumption. This set includes an indicator for whether the household head is male, the household head's age and its square, and the number of children in the household.¹⁵ In addition to rainfall, there may be other channels, such as temperature, through which rural household consumption is correlated across space. To account for this, we also include a district's average monthly temperature in X_{hdt} . Lastly, θ_h and θ_t are household and year fixed effects respectively while ϵ_{hdt} is an error term.

The inclusion of household fixed effects in our specification provides us with two key advantages. First, a negative β_2 could reflect the impact of differential crop choices. For instance, it could be the case that households that grow higher-priced or higher-yield crops endogenously locate in districts with a lower probability of a large neighbor's rainfall shock. In other words, households in these districts cultivate different crops compared to households in districts that tend to receive larger neighbor's rainfall shocks. To the extent that these crop choices are time invariant, our household fixed effects will capture this confounding effect. Second, these fixed effects will also account for time-invariant, district-specific and household characteristics that might impact its consumption.

While the inclusion of household fixed effects has key advantages, it is worth noting that our rainfall shock measures, R_{dt}^O and R_{dt}^N , vary by district and year and not by household. Thus, the inclusion of household fixed effects means that our results are identified from within-district variation in own rainfall and neighbor's rainfall from its long-term average. As we argued above, conditional on including household fixed effects, these deviations are likely to be orthogonal to unobserved determinants of rural household consumption and allow us to identify the causal effects of own rainfall shocks as well as rainfall shocks in neighboring districts. In addition, as is clear from Fig. 2, there is significant within-district, temporal variation in our rainfall data. This allows us to identify β_1 and β_2 . Nonetheless, we show below that our results are robust to excluding household fixed effects.

5 Results

5.1 Baseline Results

We report our baseline results in Table 3. In column (1), we estimate a parsimonious version of (7) where we exclude household fixed effects. The coefficient of the own rainfall shock is positive and statistically significant. However, the coefficient of the neighbor's rainfall shock variable suggests that having greater rainfall in nearby districts lowers

¹⁵ Some of these control variables may be correlated with rainfall and thereby introduce confounding effects. To guard against this, we show in Table 9 that our coefficients of interest - β_1 and β_2 - are robust to excluding all household controls from our regressions.

a household's consumption. In other words, while rainfall in a household's own district raises its consumption, rainfall in nearby districts has the opposite effect.¹⁶

In column (2), we add a set of district controls to the specification in column (1) to account for district-level factors that are correlated with a household's consumption. These controls include indicators for a district's elevation and slope, the natural logarithm of a district's population, the share of workers in a district that are in agriculture, and the share of literate workers in a district. To ensure that these latter variables are not endogenous to current rainfall, we use National Sample Survey Organization data from 1987 to construct them. The coefficient of both own-rainfall shock and neighbor's rainfall shock remain robust, albeit the magnitude of the latter falls.

Next, in column (3) of Table 3, we report the results from estimating Eq. (7). That is, we now include household fixed effects in our regression. The inclusion of these fixed effects account for all time invariant, omitted household and district characteristics that may bias our estimates of the own rainfall shock and the neighbor's rainfall shock. As the results in column (3) demonstrate, the effects we have identified thus far remain robust to the inclusion of household fixed effects. That is, we continue to find that experiencing a greater own rainfall shock raises household consumption while experiencing a greater neighbor's rainfall shock lowers household consumption. In column (3), with the inclusion of household fixed effects, we are relying on within-district variation in rainfall to identify our rainfall shock effects. As is clear from Fig. 2, our data does exhibit significant within-district variation in rainfall. Nonetheless, it is reassuring that our key result remains robust regardless of whether we include household fixed effects.

To gauge how important the spatial spillover effect of rainfall is, consider first a case where we ignore rainfall in neighboring districts. In this benchmark case, the estimates in column (3) of Table 3 suggest that a one-standard deviation increase in a district's own rainfall will result in an 8.46 percent increase in household consumption per capita. To see how this effect changes when we incorporate the spatial spillover effect, note that we can use (7) to write the effect of R^O on household consumption (C) as

$$\frac{d\ln(C)}{dR^O} = \hat{\beta}_1 + \hat{\beta}_2 \frac{dR^N}{dR^O}, \quad (8)$$

where the second term on the right-hand-side captures the attenuating effect of rainfall in neighboring districts.

To implement this, we aggregate our data to the district-year level and regress R_{dt}^N on R_{dt}^O , district fixed effects, and state and year interaction fixed effects. The resulting coefficient of the own-district shock is 0.086, which is our estimate of dR^N/dR^O . Combining this with our estimates of $\hat{\beta}_1$ and $\hat{\beta}_2$ from column (3) of Table 3, we find that a one-standard deviation increase in the own-rainfall shock now increases a household's per-capita consumption by 5.23 percent. That is, accounting for the spatial spillover effect reduces the consumption gains from own-district rainfall by approximately 38 percent.

The consumption expenditure reported by IHDS includes both expenditure on market-purchased items as well as the value of homegrown crops consumed. For

¹⁶ One may be concerned that the IHDS survey is conducted throughout the year and thus the consumption data may not overlap with the rainfall data, which is collected over the monsoon period. Fortunately, only 0.38 percent of households in IHDS1 and 0.09 percent of households in IHDS2 were interviewed during or prior to the monsoon season of 2004 and 2011 respectively. Thus, our consumption data are almost entirely in the post-monsoon period.

Table 3 Spatial-spillover effect of rainfall on rural households

Dependent variable	(1)	(2)	(3)	(4)	(5)
	ln (consumption per capita)				
Own rainfall shock	0.143*** (0.031)	0.076*** (0.026)	0.084*** (0.023)	0.103*** (0.026)	0.084*** (0.025)
Neighbor's rainfall shock	- 1.254*** (0.139)	- 0.436*** (0.159)	- 0.379*** (0.136)	- 0.433*** (0.152)	- 0.379** (0.151)
Constant	6.059*** (0.186)	54.293 (146.422)	7.233*** (0.448)	7.584*** (0.470)	- 0.000 (0.008)
Time-invariant controls	No	Yes	-	-	-
Household fixed effects	No	No	Yes	Yes	-
Homegrown growers excluded	No	No	No	Yes	No
Conley standard errors	No	No	No	No	Yes
Observations	53,792	46,514	53,791	37,843	53,782
R-squared	0.210	0.274	0.248	0.276	0.091

The dependent variable in all columns is the natural logarithm of a household's monthly per capita consumption expenditure. The construction of the own rainfall shock and neighbor's rainfall shock variables are described in the text. In column (2), we control for a household's religion, caste, indicators for its district's elevation and slope, the district's total population in 1987, its share of agricultural workers in 1987, its share of literate individuals in 1987, and state fixed effects. These additional controls are all time invariant and are absorbed by the household fixed effects in column (3). In column (4), we exclude households that report consuming only homegrown staples (rice, wheat, and cereals). In column (5), we report Conley (1999) spatial correlation-adjusted standard errors with a distance cutoff of 250 kilometers. All regressions control for the number of children in a household, the household head's age, age squared, whether the household head is male, and the average monthly temperature during the year. All regressions also include year fixed effects and incorporate sampling weights to ensure that the sample reflects the population. Robust standard errors in parenthesis are clustered at the district-year level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

homegrown crops, the reported quantities consumed by each household were multiplied by the market price and added to total consumption expenditure. Unfortunately, total expenditure on homegrown crops is not separately reported in the IHDS data, which means that we cannot subtract it from total consumption expenditure and isolate the expenditure on market-purchased items only. Recall that the latter is the ideal proxy for C_M in our conceptual framework in Sect. 2.

Instead, we identify households in our data that only consume homegrown staples and exclude them from our sample.¹⁷ Because the remaining households are ones who are less reliant on homegrown staples, their total consumption expenditure will largely reflect expenditure on market-purchased items. In column (4) of Table 3, we report the results using this sub-sample. As the results demonstrate, our key findings remain highly robust to excluding households that only consume homegrown staples.

¹⁷ We err on the side of caution and classify a household as a homegrown consumer if it consumes only homegrown rice or wheat or cereals. This is cautious in the sense that some of these households could consume only homegrown rice and a mixture of homegrown and market-purchased wheat or cereals.

5.2 Alternate Inference Approach and Specifications

Our econometric approach above controls for spatial correlation in rainfall by including a neighbor's rainfall shock measure. However, there could also be spatial correlation in the error term itself in Eq. (7). To the extent that this is the case, the standard errors we report in Table 3 are incorrect even if our estimate of β_2 is unbiased. To address this, we report standard errors in column (5) of Table 3 that adjust for spatial correlation following Conley (1999) as well as a standard heteroskedastic and auto-correlation correction (HAC) following Hsiang (2016).¹⁸ As these results show, our baseline findings are largely unaffected when we use the spatial-HAC correction. We still find that a higher neighbor's rainfall shock has a negative and statistically significant effect on a household's consumption.

The Conley (1999) approach, while popular, is also computationally intensive as one must account for distances between every pair of observations when constructing the spatial variance-covariance matrix. Given our relatively large, household-level sample, this is an especially acute computational challenge. In light of this, our choice of district-year level clustering as the baseline approach follows the advice of Hsiang (2016, p. 66), who argues that it is "reasonable to estimate approximate standard errors using simpler techniques, verifying that spatial-HAC adjustments do not alter the result substantively."

To account for other channels of spatial spillovers such as similar farm production technology and soil types (Chen et al. 2016), we include a spatial lag (LeSage and Pace 2009) to our baseline specification. Given that our unit of observation is a household, a spatial lag in our case is a weighted average of household consumption in nearby areas, where the weights are the bilateral distance between households.

Unfortunately, to construct such a spatial lag at the household level, we need the geo-coordinates of each household. Such information is not available. Instead, we adopt an alternate approach where we calculate a district-level spatial lag of the dependent variable. That is, for each household in our sample, we calculate the weighted average district-level consumption per capita in all other districts. The weights are the bilateral distance between a household's district of residence and all other districts. We then add this district-level spatial lag as an explanatory variable to our baseline specification (7). We report the results from estimating this new specification in column (1) of Table 4. As these estimates demonstrate, our coefficient of interest remain highly robust. We continue to find that a higher neighbor's rainfall shock has a negative and statistically significant effect on a household's consumption.

To further guard against our spillover effect being driven by spurious spatial correlation, we conduct a placebo test in column (2). The rationale for this placebo test is the idea that rainfall in neighboring districts should not have any effect on a district's crop yields. To test this, we use the ICRISAT Village Dynamics in South Asia Macro-Meso Database (henceforth ICRISAT) to construct crop-district-year-specific measures of yields for the period 2004 to 2011. This dataset includes information on 16 major crops in 311 districts across India.¹⁹ We then regress the natural logarithm of these crop yields on both

¹⁸ We use the STATA .ado file `reg2hdfespatial` created by Thiemo Fetzer and used in Fetzer (2020) to implement this procedure and assume a distance cutoff of 250 kilometers.

¹⁹ The 16 crops are rice, wheat, sorghum, pearl millet, maize, finger millet, barley, chickpea, pigeon-pea, sugarcane, groundnut, sesame, rape and mustard, linseed, castor, and cotton. As we use monsoon rainfall data in our baseline analysis, we restrict the ICRISAT data to crops that are primarily grown during the monsoon months of June to September. Further, to account for outliers, we omit crop yields that are above the 95th percentile of the crop yield distribution.

own-district and neighboring-district rainfall. The results in column (2) suggest that while higher own-district rainfall increases crop yields, greater rainfall in neighboring districts does not have an effect on crop yields.

Lastly, while the household fixed effects in our baseline specification purges the effect of any time-invariant district characteristics, there could be unobservable, time-varying district shocks that threaten our identification strategy. For instance, the timing of rainfall shocks may coincide with other time-varying agricultural productivity shocks. To account for this, we include in Eq. (7) the interaction between a district's share of agricultural workers in 1987 and year fixed effects respectively. These interaction terms will allow us to flexibly capture these time-varying, location-specific agricultural shocks. As these results in column (3) of Table 4 demonstrate, our coefficient of interest remains highly robust.

Thus far, we have estimated a parsimonious baseline specification with linear own-district and neighbor's rainfall shocks. We now explore alternate specifications in Table 5. In column (1), we include both a squared own-rainfall shock and neighbor's rainfall shock. Interestingly, we find that both the level and squared own-district rainfall coefficients are positive and statistically significant. The latter suggests that the benefits of own rainfall are increasing in the level of rainfall itself. In the case of the neighbor's rainfall shock, while the coefficients of both the level and squared terms are negative, the latter is considerably larger and statistically significant. This suggests that at higher values of R^N , the marginal effect of R^N on household consumption increases in magnitude.

Next, in column (2) of Table 5, we estimate an alternate specification where we include an interaction between own-district and neighbor's rainfall shocks. The coefficients of interest remain robust while the interaction term itself is statistically insignificant. Lastly, in column (3), we include both the squared terms and the interaction term to our baseline specification. We continue to find the squared neighbor's rainfall shocks is negative and significant. Thus, the results in Table 5 suggest that even with alternate specifications, the spatial spillover effect of rainfall remains robust. The key additional insight from this table is that the spillover effect is being driven by very large neighboring shocks.

6 Mechanisms

In Sect. 2, we hypothesized that rainfall in neighboring districts can have adverse general-equilibrium effects on rural households by (a) lowering crop prices and (b) lowering farm income via the intensive margin (market sales) as well as the extensive margin (market participation). We now examine whether these channels are supported by the data. To test whether greater neighbor's rainfall lowers crop prices, we use crop price data from ICRI-SAT. For each district, this dataset provides farm-gate prices of crops in Indian rupees per quintal (100 kg). For our analysis, we use annual data for the period 2004 to 2011. With these data in hand, we examine whether greater rainfall in neighboring districts lower the price of crops in a given district by estimating the following econometric specification:

$$\ln(P_{c dt}) = \alpha_c + \delta_1 R_{dt}^O + \delta_2 R_{dt}^N + \theta_d + \theta_c \times \theta_t + v_{c dt} \quad (9)$$

where $P_{c dt}$ is the farm-gate price for crop c in district d and year t . R_{dt}^O and R_{dt}^N are the rainfall shock measures defined above while θ_c , θ_d , and θ_t are crop, district, and year fixed effects respectively. Lastly, $v_{c dt}$ is an error term. If the mechanism we propose above is correct, then we would expect δ_2 to be negative. We report the results from estimating Eq. (9) in column (1) of Table 6. The coefficient of the neighbor's rainfall shock is indeed negative

Table 4 Addressing confounding effects

Dependent variable	(1)	(2)	(3)
	ln (copc)	ln (yield)	ln (copc)
Own rainfall shock	0.073*** (0.023)	0.014** (0.007)	0.070*** (0.025)
Neighbor's rainfall shock	- 0.298** (0.128)	- 0.021 (0.024)	- 0.314** (0.137)
Constant	- 5.850** (2.479)	- 0.099*** (0.004)	7.198*** (0.443)
District spatial lag included	Yes	No	No
Agriculture share \times year FE	No	No	Yes
Observations	53,791	24,267	52,666
R-squared	0.256	0.606	0.250

The dependent variable in columns (1) and (3) is the natural logarithm of a household's monthly per capita consumption expenditure (copc) while in column (2) it is the natural logarithm of crop yield. The unit of observation in column (2) is crop-district-year, which is why the sample size is smaller. In column (1), we include a district-level spatial lag of average household consumption per capita. In column (3), we include the interaction between a district's share of agricultural employment in 1987 and year fixed effects. The regressions in columns (1) and (3) control for the household head's age, age squared, whether the household head is male, the average monthly temperature during the year, the number of children, household fixed effects, and year fixed effects. These regressions also incorporate sampling weights to ensure that the sample reflects the population. The regression in columns (2) includes district fixed effects and crop and year interaction fixed effects. The standard errors in parenthesis in all columns are robust and clustered at the district-year level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

and statistically significant, which supports the idea that greater neighboring rainfall will lower crop prices.²⁰

Next, we examine how rainfall in neighboring districts affect a household's farm income through both the intensive and extensive margins. For the former, we use IHDS's household-level data on market participation to estimate each farming household's ratio of market sales to production. We then use this ratio as the dependent variable in Eq. (7). Note that the crop data are only available for 2005, which is why the farm-income regressions do not include household fixed effects. The results from estimating this intensive margin effect are reported in column (2) of Table 6. They show that greater own-district rainfall increases market sales and greater rainfall in neighboring districts lower market sales.

In addition to lowering the value of market sales, equation (1) suggests that greater rainfall in neighboring districts can also lower the likelihood that a household participates in markets. More precisely, if the decrease in crop price, P , due to higher neighboring rainfall is large enough, then it could be the case that $P < \tau P_T$, where τ is the transportation cost of

²⁰ The results in column (1) also suggest that own-district rainfall, on average, does not have a statistically significant effect on crop prices. This validates our assumption in Sect. 2 that each district is small in the regional economy, which means that crop output shocks due to own rainfall will not impact crop prices.

Table 5 Alternate specifications

Dependent variable	(1)	(2)	(3)
	ln (consumption per capita)		
Own rainfall shock	0.048* (0.026)	0.092*** (0.025)	0.038 (0.027)
Own rainfall shock squared	0.023** (0.011)		0.003 (0.023)
Neighbor's rainfall shock	- 0.109 (0.141)	- 0.374*** (0.136)	- 0.077 (0.142)
Neighbor's rainfall shock squared	- 0.207*** (0.062)		- 0.308** (0.121)
Own shock × Neighbor's shock		- 0.018 (0.019)	0.079 (0.078)
Constant	6.761*** (0.459)	7.094*** (0.459)	6.698*** (0.473)
Observations	53,791	53,791	53,791
R-squared	0.250	0.248	0.250

The dependent variable in all columns is the natural logarithm of a household's monthly per capita consumption expenditure. All regressions control for the household head's age, age squared, whether the household head is male, the number of children, the average monthly temperature during the year, household fixed effects, and year fixed effects. All regressions also incorporate sampling weights to ensure that the sample reflects the population. The standard errors in parenthesis in all columns are robust and clustered at the district-year level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 6 Mechanisms

Dependent variable	(1)	(2)	(3)	(4)
	ln (crop price)	Sales/production	Market indicator	ln (farm wages)
Own rainfall shock	0.011 (0.007)	0.128** (0.052)	0.089 (0.063)	0.097* (0.058)
Neighbor's rainfall shock	- 0.087*** (0.033)	- 1.047*** (0.238)	- 0.771*** (0.256)	- 1.696*** (0.387)
Constant	7.225*** (0.005)	- 146.565 (168.942)	- 194.601 (191.006)	12.573*** (1.356)
Observations	8,135	9,754	10,114	17,394
R-squared	0.794	0.188	0.121	0.105

The dependent variable in column (1) is the natural logarithm of crop prices and is constructed using ICRI-SAT data. The unit of observation in this column is crop-district-year. The dependent variable in column (2) is the ratio of a household's total market sales to its total production while in column (3) it is an indicator for whether a household sells its crop to the market. The market participation data are at the household level and only available for 2005. The dependent variable in column (4) is the natural logarithm of a household's farm wages per capita. The sample in this column is restricted to households that have positive farm wages. The regression in columns (1) includes district fixed effects and crop and year interaction fixed effects. The regressions in columns (2) and (3) include all controls in column (2) of Table 3 while the regression in column (4) includes all controls in column (3) of Table 3. The regressions in columns (2)–(4) also incorporate sampling weights to ensure that the sample reflects the population. Robust standard errors in parenthesis are clustered at the district-year level in columns (1) and (4) and at the district level in columns (2) and (3). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

selling to the market and $P_T < P$ is the price received by the farmer if she sells her surplus to traders at the farm gate. As we can see from (1), $P < \tau P_T$ would result in the farmer no longer participating in the market.

To examine this extensive margin effect, we construct a market participation indicator that takes the value of one if a household sells any crops in the market and is zero otherwise. The results from estimating this extensive margin effect are reported in column (3) of Table 6. As above, they show that while greater own-district rainfall increases market participation, greater rainfall in neighboring districts lowers market participation. Thus, through both the intensive and extensive margins, greater neighbor's rainfall lowers a rural household's market income.

Lastly, we examine whether rainfall in neighboring districts also has adverse effects on a household's income from agricultural labor. To the extent that such rainfall lowers farm revenue, it should also have an adverse effect on the wage received by farm workers. To explore this, we use a household's total agricultural wage income per capita.²¹ We then estimate a version of (7) where the dependent variable is the natural logarithm of a household's agricultural wage income. We report these results in column (4) of Table 6. They suggest that while greater own rainfall shock raises farm wage income, greater neighbor's rainfall shock lowers it. These results are consistent with the farm revenue effects above.

6.1 Alternate Mechanisms

In Sect. 2, the mechanisms we described to explain the impact of neighboring rainfall on rural household consumption only applied to households that sold their crops at the market. Households that either do not have a surplus to sell or choose not to participate in the market for other reasons should not be impacted by price changes due to neighboring rainfall. Indeed, this last point provides a useful falsification test that we can run to validate our mechanisms.

To explore these implications, we use the market participation data to identify households that report participating in the market in 2005 as well as households that do not. We then estimate our baseline regression separately for each sub-sample. If the mechanisms we discussed are valid, we should find that the negative spillover effect from neighboring rainfall only holds for the market-participant sample. In columns (1) and (2) of Table 7, we show that this is indeed the case. In column (1), where we restrict the sample to market participants, our baseline result for neighbor's rainfall is robust. In contrast, in column (2), where the sample is restricted to non-market participants, there is no neighbor's rainfall effect.

In column (3), we verify that this result holds in a regression where we interact both own-rainfall shock and neighbor's rainfall shock with the market participation indicator. The result confirms that both an own-rainfall shock and a neighbor's rainfall shock has a statistically significant effect only for households that participate in the market.²²

A concern with our headline result is that it could be driven by time-varying structural changes in a district's economy that happen to be correlated with rainfall shocks.

²¹ To account for outliers in the raw farm wage data, we omit from our sample in column (4) households that report farm wages that is equal to or above the 95th percentile.

²² While the result in column (3) validates the mechanisms we highlight in Sect. 2, the market participant indicator in this column is itself endogenous to rainfall (column (3) of Table 6). This is why we do not treat this interaction specification as our baseline.

To address this, we examine whether our rainfall shock measures are related to income changes among non-agricultural households. If our results are being spuriously driven by unobserved structural changes, then we should find negative effects of rainfall of comparable magnitude on non-agricultural households. In contrast, if there are weaker effects of rainfall on such households, then we can be confident that our headline result is indeed being driven by rainfall.²³

We explore this by examining whether our rainfall shocks affect income from non-agricultural sources. To do so, we first estimate a version of Eq. (7) where we change the dependent variable to the natural logarithm of a household's salary income per capita from non-farm sources. These results, which are reported in column (4) of Table 7, indicate that both the effect of own-rainfall shocks and neighbor's rainfall shocks are statistically insignificant. In column (5), we repeat the analysis above, but use the natural logarithm of a household's non-farm wage income per capita as the dependent variable.²⁴ As in the previous column, we find that both the effect of own-rainfall shocks and neighbor's rainfall shocks are statistically insignificant.

Next, we examine whether own and neighbor's rainfall shocks lead to out migration from the rural households in our sample. While the survey data we use do not measure temporary migration in both IHDS rounds, it does include each household's income from remittances. This allows us to use remittance values as proxies for the rate of out migration from a household. These results are reported in column (6) of Table 7 where the dependent variable is now the natural logarithm of each household's remittance earnings per capita. As with columns (4) and (5), we find that both the effect of own-rainfall shocks and neighbor's rainfall shocks on a household's remittance earnings are statistically insignificant. Taken together, the results in Table 7 suggest that the household-consumption effects we've documented thus far are not being driven by changes in the non-agricultural sector or due to out-migration.

7 Additional Results

7.1 Results by Expenditure Type

Up to this point, our default measure of household welfare has been total consumption per capita. We now examine the effect of own-district rainfall shocks as well as neighbor's rainfall shocks on various types of consumption expenditure. Our motivation for doing this is to examine the impact of these rainfall shocks on particularly important types of expenditure such as food, schooling, and medical expenses. We begin in columns (1) and (2) of Table 8 by decomposing total household consumption into food consumption and non-food consumption respectively. In column (1), we use the natural logarithm of a household's total food expenditure per capita as the dependent variable. The coefficient of own-district

²³ Note that an existing literature shows that own-district weather fluctuations in rural areas can lead to a reallocation of economic activity from agriculture to non-agriculture (Santangelo 2016; Emerick 2018, Colmer, forthcoming). Thus, in principle, the impact of rainfall of non-agricultural income need not be zero.

²⁴ The income data we use are as constructed by IHDS. They decomposed non-farm income into income from household members who received monthly salaries and income from household members who received daily wages. We define the former as non-farm salary while we treat the latter as non-farm wages.

Table 7 Alternate mechanisms

Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)
	ln (cons. per capita)			ln (nf salary)	ln (nf wages)	ln (remitt-ances)
Own rainfall shock	0.080*** (0.028)	0.038 (0.034)		-0.069 (0.061)	-0.023 (0.079)	0.101 (0.156)
Neighbor's rainfall shock	-0.540*** (0.180)	-0.075 (0.162)		0.410 (0.378)	0.280 (0.322)	-1.471 (0.967)
Own rainfall × Non-market participant			0.047 (0.047)			
Own rainfall × Market participant			0.128*** (0.045)			
Neighbor's rainfall × Non-market participant			-0.591 (0.377)			
Neighbor's rainfall × Market participant			-1.010** (0.399)			
Market participant	Yes	No	-	-	-	-
Observations	14,576	14,581	28,856	8,092	16,979	4,142
R-squared	0.225	0.216	0.764	0.127	0.092	0.335

The dependent variable in columns (1)–(3) is the natural logarithm of a household's consumption per capita. The dependent variable in column (4) is the natural logarithm of a household's total non-farm salary income per capita while in column (5) it is the natural logarithm of a household's non-farm wages per capita. Lastly, the dependent variable in column (6) is a household's total remittance earnings per capita. The sample in column (1) is restricted to households that sold crops at the market in 2005 while in column (2), the sample is restricted to households that did not sell at the market in 2005. The regressions in all columns control for the number of children in a household, the household head's age, age squared, whether the household head is male, the average monthly temperature during the year, household fixed effects, year fixed effects, and a constant that is not reported. All regressions incorporate sampling weights to ensure that the sample reflects the population. Robust standard errors in parenthesis are clustered at the district-year level in all columns. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

rainfall shock is positive and statistically significant while the coefficient of neighbor's rainfall shock is negative and statistically significant.

In column (2), we use the natural logarithm of a household's total non-food expenditure per capita as the dependent variable. Non-food items include rent, expenditure on electricity, telephone, entertainment and other miscellaneous items. Thus, compared to food, these items are comparatively durable in nature. The coefficients in column (2) suggest that both an own-district rainfall shock and a neighbor's rainfall shock has a statistically insignificant effect on rural household consumption. Taken together, the results in columns (1) and (2) of Table 8 indicate that households respond to a neighbor's rainfall shock by primarily lowering expenditure on food items and not by lowering expenditure on the relatively more durable, non-food items.

Next, we examine the impact of own and neighbor's rainfall shocks on components of consumption that may have long-term consequences. More precisely, in column (3) of Table 8 we use the natural logarithm of a household's total schooling expenditure over the previous 365 days as the dependent variable. This is the only recall period for which

Table 8 Spillover effects of rainfall—by expenditure type

Dependent variable	(1)	(2)	(3)	(4)
	ln (expenditure per capita on)			
	Food	Non-food	Schooling	Medical
Own rainfall shock	0.063*** (0.021)	0.033 (0.039)	- 0.019 (0.042)	0.085 (0.058)
Neighbor's rainfall shock	- 0.231* (0.124)	0.030 (0.186)	- 0.377 (0.230)	0.197 (0.307)
Constant	6.691*** (0.373)	4.149*** (0.712)	3.634*** (0.928)	5.946*** (1.206)
Observations	53,769	53,568	33,007	40,728
R-squared	0.128	0.328	0.246	0.055

The dependent variable in column (1) is the natural logarithm of household expenditure per capita on food items. Similarly, the dependent variables in column (2) to (4) are the the natural logarithm of household expenditure per capita on non-food items, schooling, and medical purposes respectively. All regressions controls for the number of children in a household, the household head's age, age squared, whether the household head is male, and the average monthly temperature during the year. All regressions also include household fixed effects and year fixed effects and incorporate sampling weights to ensure that the sample reflects the population. Robust standard errors in parenthesis are clustered at the district-year level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

these data are available. The impact of rainfall on schooling is both theoretically ambiguous and empirically contested.²⁵ Our results in column (3) suggest that both own-district rainfall shocks and neighbor's rainfall shocks have statistically insignificant effects on a household's expenditure on schooling.

Finally, in column (4) of Table 8, we explore the impact of rainfall shocks on a household's medical expenses. This is an alternate channel through which these shocks may have adverse long-term consequences. The dependent variables here is the natural logarithm of a household's total medical expenditure over the previous 365 days. The results in this column suggest that both a positive own-district rainfall shock and a positive neighbor's rainfall shock have statistically insignificant effects on a household's medical expenditure. Thus, the results in Table 8 indicate that the rural households in our sample respond to a neighbor's rainfall shock by primarily reducing food expenditure. We find no such effect on durable, non-food expenditure as well as on schooling and medical expenditures. These results are consistent with the idea that a neighbor's rainfall shocks mainly represent an adverse shock to a household's transitory income.

7.2 Alternate Rainfall and Consumption Measures

We next examine whether our main findings are robust to using alternate measures of rainfall and consumption. In column (1) of Table 9, we follow Jayachandran (2006) and construct categorical measures of rainfall shocks. More precisely, for each district we create

²⁵ See the discussion in Shah and Steinberg (2017) for further details on this literature.

Table 9 Robustness checks

Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)
	ln (cons. per capita)			Winsorized		
				ln (copc)	ln (cons.)	ln (copc)
Own positive shock	0.134*** (0.030)					
Neighbor's positive shock	-0.727*** (0.212)					
Own rainfall shock		0.009 (0.015)	-0.003 (0.021)	0.083*** (0.023)	0.082*** (0.023)	0.072*** (0.024)
Neighbor's shock (population)			-0.276** (0.125)			
Neighbor's rainfall shock			-0.153** (0.071)	-0.372*** (0.132)	-0.368*** (0.122)	-0.289** (0.142)
Rainfall data	ERA (default)		UDEL	ERA (default)		
Household Controls	Yes	Yes	Yes	Yes	Yes	No
Observations	53,791	51,862	52,137	53,791	53,791	54,541
R-squared	0.250	0.245	0.252	0.252	0.122	0.175

The dependent variable in columns (1)–(3) and (6) is the natural logarithm of a household's monthly per capita consumption expenditure (copc). In column (4), the dependent variable is the natural logarithm of a household's monthly per capita consumption expenditure that has been Winsorized at the 1 percent and 99 percent levels respectively. Lastly, in column (5), the dependent variable is a household's total monthly consumption. Own positive shock is a binary variable that takes the value of one if a district received an own rainfall shock above the 80th percentile. Similarly, neighbor's positive shock is a binary variable that takes the value of one if a district received a neighbor's rainfall shock above the 80th percentile. Neighbor's shock (population) uses population-weighted distance as weights when constructing the neighbor's rainfall shock. The rainfall shocks in column (3) are constructed using the University of Delaware's rainfall data. All regressions control for the number of children in a household, the household head's age, age squared, whether the household head is male, the average monthly temperature during the year, household fixed effects, year fixed effects, and a constant that is not reported. All regressions also incorporate sampling weights to ensure that the sample reflects the population. The standard errors in parenthesis in all columns are robust and clustered at the district-year level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

an own positive shock variable that takes the value of one if a district's annual monsoon rainfall is above the 80th percentile of that district's monsoon rainfall over the period 1979 to 2015. All other districts have a value of zero. Similarly, for each district, we construct a neighbor's positive shock measure that replaces R_{jt}^N in Eq. (6) with this categorical version.

In contrast to our default measure, these categorical measures do not use the full rainfall data and instead focus on extreme positive shocks (i.e. above the 80th percentile). Thus, we do not treat them symmetrically to our default baseline. Nonetheless, it is useful to check whether our core results are robust to this alternative way of capturing rainfall shocks. Indeed, the results in column (1) of Table 9 show that households in districts that received greater than 80th percentile own rainfall experience an increase in consumption. These results also show that households in districts that received greater than 80th percentile neighbor's rainfall experience a decrease in consumption. Both of these results are consistent with our baseline findings in Table 3.

In column (2), we use an alternate definition of neighbor's rainfall that takes into account the market size of neighboring districts. More precisely, we define

$$R_{dt}^{NP} = \sum_{j \neq d} \left(\frac{POP_j}{\omega_{dj}} \times R_{jt}^O \right),$$

where the weight now depends on both the other district j 's population, POP_j , as well as the distance between d and j , ω_{dj} .²⁶ The benefit of using this alternate definition is that it allows a neighbor's market size to influence how a rainfall shock there will impact regional prices and household consumption in d . As the results in column (2) demonstrate, the spatial spillover effect of neighboring rainfall remains robust to using this alternate definition.

In constructing our baseline sample, we used rainfall data from the ERA-Interim Reanalysis Archive. These re-analysis data combine ground-station and satellite data with results from global climate change models to create a consistent measure of rainfall across time and space. In contrast, alternate sources such as the University of Delaware's (UDEL) terrestrial precipitation data tends to rely more heavily on ground station data. This has the disadvantage that ground stations, especially in developing countries, are not uniformly distributed across space. Further, as Colmer (forthcoming) points out, the quality of ground stations in India has deteriorated over time. Nonetheless, for the sake of completeness, we examine the robustness of our findings to the use of the alternate UDEL data. We report the results from this robustness check in column (3) of Table 9. As the results demonstrate, the coefficient of the neighbor's rainfall shock remains negative and statistically significant. While the own-rainfall effect is not robust, these alternate data yield a neighbor's rainfall shock effect that is fully consistent with our baseline findings.

We next turn to whether our results are robust to our choice of dependent variable. Recall that our default dependent variable is the natural logarithm of a household's consumption per capita. We used this variable as provided by IHDS without excluding outliers. To examine whether our core results are driven by such outliers, we winsorize the consumption data at the 1 percent and 99 percent levels. The results in column (4) of Table 9 suggest that these potential outliers do not drive our results. Even after winsorizing the consumption data, our coefficient of interest remains highly robust with magnitudes that are similar to the baseline results in Table 3.

In column (5) of Table 9, we consider the effect of rainfall shocks on total household consumption rather than on consumption per capita. That is, we multiply our default measure of consumption per capita with a household's size to obtain each household's total consumption. We do so to account for the fact that our default consumption per capita measure captures both the effect of rainfall on consumption as well as its effect on household size. In Table 7, we showed the rainfall shocks do not have any effect on a household's remittance income. Thus, we do not believe that the effect of rainfall shocks on household size due to migration is a meaningful confounding effect. To verify this, we use as the dependent variable the natural logarithm of a household's total consumption in column (5)

²⁶ To calculate R_{dt}^{NP} , we use j 's population from the initial round of the IHDS. This means that the population we use is time invariant. Note that the population data are at the district level, where each district is classified as either being rural or urban. In other words, we do not observe the size of the rural population and urban population separately in each district. This prevents us from using each district's rural population only to construct R_{dt}^{NP} .

of Table 9. As the results confirm, the effect of both own and neighbor's rainfall shocks are very similar to the baseline.

Lastly, in column (6), we estimate a version of our baseline specification without any household controls. Recall that our baseline specification includes a household head's age, age squared, whether the household head is male, and the number of children. Of particular concern is the possibility that some of our control variables are correlated with rainfall and can therefore bias our coefficients of interest. However, as the results in column (6) demonstrate, our coefficients of interest remain highly robust to excluding all control variables from our sample.

8 Conclusion

In this paper, we showed that greater rainfall can have adverse spatial spillover effects on rural households. Central to this new conclusion was our focus on estimating the effect of both own-district rainfall and rainfall in neighboring districts on rural household consumption. In theory, the welfare effect of such spatial spillovers is ambiguous. To see this, consider a farmer that receives greater yield due to greater rainfall in his own district. To the extent that rainfall spans multiple districts, there will also be greater rainfall in neighboring districts, which will result in a positive supply shock. All else equal, this will drive down the regional price of agricultural crops. This reduction in price can create both welfare gains and losses for a farming household. As consumers, such a household gains from the lower prices. As producers, however, the lower prices result in lower farm income, given price inelastic demand. Thus, when we consider both own-district rainfall as well as neighboring-district's rainfall, the overall effect of rainfall on household welfare is ambiguous.

To explore this spillover effect empirically, we used household-level, panel data from India along with high-resolution meteorological data to examine whether rural household consumption depends on rainfall shocks in its own district as well as rainfall shocks in neighboring districts. Our identification strategy incorporated household fixed effects, which allowed us to purge the effect of any unobserved, time-invariant household and district characteristics. Thus, our results were identified from within-district variation in own rainfall and neighbor's rainfall from its long-term average. These deviations are orthogonal to unobserved determinants of rural household consumption and allow us to identify the causal effects of rainfall shocks.

Our results indicated that both own-district rainfall shocks and neighbor's rainfall shocks have a statistically and economically significant effect on rural household consumption. Indeed, we found that accounting for the spatial spillover effect resulted in a 38 percent decrease in the consumption benefits of an increase in own-district rainfall. These results suggest that one must account for spatial spillover effects to correctly estimate the welfare effects of own-district rainfall shocks. While this adds important nuance to our understanding of the effects of rainfall shocks, the lack of appropriate data meant that we were unable to examine the adaptation strategies adopted by the households in our sample. Exploring these adaptation strategies is a fruitful avenue for future research.

Appendix A: Additional Tables

See Table 10.

Table 10 Summary statistics by survey year

Survey year	(1) 2005	(2) 2012	(3) Both
<i>Panel A: household characteristics</i>			
Household consumption per capita	691.3 (632.8)	961.5 (1016.2)	824.0 (853.8)
Indicator for households in poverty	0.237 (0.425)	0.202 (0.401)	0.220 (0.414)
Number of household members	5.974 (3.093)	4.872 (2.385)	5.433 (2.822)
Number of children	2.030 (1.826)	1.492 (1.507)	1.766 (1.699)
Male household head	0.911 (0.285)	0.856 (0.351)	0.884 (0.320)
Household head's age	48.19 (13.27)	49.23 (13.97)	48.70 (13.63)
<i>Panel B: weather</i>			
Total rainfall (in metres)	0.106 (0.0759)	0.151 (0.0781)	0.128 (0.0802)
Own rainfall shock	-0.397 (0.503)	0.741 (0.831)	0.158 (0.889)
Neighbor's rainfall shock	-0.393 (0.133)	0.707 (0.197)	0.143 (0.575)
Daily average temperature	25.82 (2.546)	25.32 (2.511)	25.57 (2.541)

In panel A, we report summary statistics for various household indicators. In panel B, we report the summary statistics for various rainfall shock measures and average monthly temperature. Household consumption per capita is reported in constant 2005 Indian rupees. 1 U.S. dollar was approximately equal to 43.5 Indian rupees in 2005. Total rainfall is reported in meters. Note that the numbers reported above incorporate sampling weights to ensure that they reflect population averages

Acknowledgements We thank Victoria Baranov, John Gibson, Paras Kharel, Leslie Martin, Yanos Zylberberg, and various seminar and conference participants at the University of Melbourne, the Australasian Development Economics Workshop, and NEUDC for helpful comments and discussions. Farzana Hossain gratefully acknowledges funding from the Faculty of Business and Economics Doctoral Program Scholarship and the Faculty of Business and Economics Graduate Research Abroad Travelling Scholarship for this research.

Funding Open Access funding enabled and organized by CAUL and its Member Institutions.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abidoye B, Cali M (2021) Income shocks and conflict: evidence from Nigeria. *J Afr Econ*
- Adams RM (1989) Global climate change and agriculture: an economic perspective. *Am J Agr Econ* 71(5):1272–1279
- Aditya KS, Subash SP, Praveen KV, Nithyashree ML, Bhuvana N, Sharma A (2017) Awareness about minimum support price and its impact on diversification decision of farmers in India. *Asia Pacific Policy Stud* 4(3):514–526
- Auffhammer M, Hsiang SM, Schlenker W, Sobel A (2013) Using weather data and climate model output in economic analyses of climate change. *Rev Environ Econ Policy* 7(2):181–198
- Barrett CB (2008) Smallholder market participation: concepts and evidence from Eastern and Southern Africa. *Food Policy* 33(4):299–317
- Barrios S, Bertinelli L, Strobl E (2010) Trends in rainfall and economic growth in Africa: A neglected cause of the African growth tragedy. *Rev Econ Stat* 92(2):350–366
- Boustan LP, Kahn ME, Rhode PW, Yanguas ML (2020) The effect of natural disasters on economic activity in U.S. counties: a century of data. *J Urban Econ* 118
- Burgess R, Donaldson D (2012) Railroads and the demise of famine in colonial India. MIT, Mimeograph
- Chen S, Chen X, Xu J (2016) Impacts of climate change on agriculture: evidence from China. *J Environ Econ Manag* 76:105–124
- Colmer J (2021) Weather, labor reallocation, and industrial production: evidence from India. *Am Econ J Appl Econ*
- Colmer J (2021) Rainfall variability, child labor, and human capital accumulation in rural Ethiopia. *Am J Agr Econ* 103(3):858–877
- Conley T (1999) GMM estimation with cross sectional dependence. *J Econometr* 92(1):1–45
- de Janvry A, Fafchamps M, Sadoulet E (1991) Peasant household behaviour with missing markets: some paradoxes explained. *Econ J* 101(409):1400–1417
- Dee DP, Uppala SM, Simmons AJ, Berrisford P, Poli P, Kobayashi S, Andrae U, Balmaseda MA, Balsamo G, Bauer P, Bechtold P, Beljaars ACM, van de Berg L, Bidlot J, Bormann N, Delsol C, Dragani R, Fuentes M, Geer AJ, Haimberger L, Healy SB, Hersbach H, Hólm EV, Isaksen L, Kållberg P, Köhler M, Matricardi M, McNally AP, Monge-Sanz BM, Morcrette J-J, Park B-K, Peubey C, de Rosnay P, Tavolato C, Thépaut J-N, Vitart F (2011) The ERA-Interim reanalysis: configuration and performance of the data assimilation system. *Q J R Meteorol Soc* 137:553–597
- Dell M, Jones BF, Olken BA (2012) Temperature shocks and economic growth: evidence from the last half century. *Am Econ J Macroecon* 4(3):66–95
- Dell M, Jones BF, Olken BA (2014) What do we learn from the weather? The new climate economy literature. *J Econ Lit* 52(3):740–798
- Desai S, Vanneman R (2005) India human development survey (IHDS-I), 2004–2005. ICPSR22626-v11. Inter-university Consortium for Political and Social Research, Ann Arbor, MI
- Desai S, Vanneman R (2012) India human development survey-II (IHDS-II), 2011–12. ICPSR36151-v5. Inter-university Consortium for Political and Social Research, Ann Arbor, MI

- Deschênes O, Greenstone M (2007) The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather. *Am Econ Rev* 97(1):354–385
- Donaldson D, Hornbeck R (2016) Railroads and American economic growth: a “market access” approach. *Quarterly J Econ* 131(2):799–858
- Economic Survey (2018) Climate, climate change, and agriculture. Ministry of Finance, Government of India
- Elliott RJR, Liua Y, Strobl E, Tong M (2019) Estimating the direct and indirect impact of typhoons on plant performance: evidence from Chinese manufacturers. *J Environ Econ Manag* 98:1–19
- Emerick K (2018) Agricultural productivity and the sectoral reallocation of labor in rural India. *J Dev Econ* 135:488–503
- Fafchamps M, Hill RV (2005) Selling at the farmgate or traveling to market. *Am J Agr Econ* 87(3):717–734
- Fetzer T (2020) Can workfare programs moderate violence? Evidence from India. *J Eur Econ Assoc* 18(6):3337–3375
- Fezzi C, Bateman IJ (2015) The impact of climate change on agriculture: nonlinear effects and aggregation bias in Ricardian models of farmland value. *J Assoc Environ Resour Econ* 2(1):57–92
- Harari M, La Ferrara E (2018) Conflict, climate, and cells: a disaggregated analysis. *Rev Econ Stat* 100(4):594–608
- Hsiang S (2016) Climate econometrics. *Ann Rev Resour Econ* 8:43–75
- Jayachandran S (2006) Selling labor low: wage responses to productivity shocks in developing countries. *J Polit Econ* 114(3):538–575
- Key N, Sadoulet E, de Janvry A (2000) Transactions costs and agricultural household supply response. *Am J Agr Econ* 82(2):245–259
- Kaur S (2019) Nominal wage rigidity in village labor markets. *Am Econ Rev* 109(10):3585–3616
- Kumar P, Kumar A, Parappurathu S, Raju SS (2011) Estimation of demand elasticity for food commodities in India. *Agric Econ Res Rev* 24:1–14
- LeSage J, Pace RK (2009) Introduction to spatial econometrics. CRC Press, Boca Raton
- Maccini S, Yang D (2009) Under the weather: health, schooling, and economic consequences of early-life rainfall. *Am Econ Rev* 99(3):1006–1026
- Mendelsohn R, Nordhaus WD, Shaw D (1994) The impact of global warming on agriculture: a Ricardian analysis. *Am Econ Rev* 84(4):753–771
- Miguel E, Kremer M (2004) Worms: identifying impacts on education and health in the presence of treatment externalities. *Econometrica* 72(1):159–217
- Santangelo G (2016) Firms and farms: the impact of agricultural productivity on the local Indian economy. University of Cambridge, Mimeograph
- Schlenker W, Hanemann WM, Fisher AC (2006) The impact of global warming on U.S. agriculture: an econometric analysis of optimal growing conditions. *Rev Econ Stat* 88(1):113–125
- Shah M, Steinberg BM (2017) Drought of opportunities: contemporaneous and long-term impacts of rainfall shocks on human capital. *J Polit Econ* 125(2):527–561
- Singh I, Squire L, Strauss J (1986) Agricultural household models: extensions, applications and policy. Johns Hopkins University Press, Baltimore
- Taylor JE, Adelman I (2003) Agricultural household models: genesis, evolution, and extensions. *Rev Econ Household* 1:33–58
- World Bank (2018) Piecing together the poverty puzzle. World Bank Publications, Washington D.C

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.