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

Farzana Hossain†
University of Melbourne

Reshad N. Ahsan‡
University of Melbourne

October, 2018

Abstract

The Intergovernmental Panel on Climate Change projects that global rainfall levels will increase by 8 percent by the second half of the 21st century. The existing literature generally finds that increases in rainfall either have no effects or actually raises agricultural profits and economic growth. Using household-level, panel data from India along with high-resolution meteorological data, we show that these average effects mask the fact that greater rainfall creates both winners and losers. Central to this novel finding is our focus on identifying the spatial spillover effect of rainfall. We show that while greater own-district rainfall raises rural household consumption, greater rainfall in neighboring districts actually has a negative effect on such consumption. While this spatial spillover effect generally attenuates the positive effect of own-district rainfall, households in districts with a low-to-moderate own rainfall shock and a large rainfall shock in neighboring districts may be made worse off from increases in rainfall.

Keywords: Rainfall, Rural Households, Climate
JEL Codes: O13, O12, Q54

*We thank Leslie Martin, Victoria Baranov, Paras Kharel, Yanos Zylberberg, and various seminar participants at the University of Melbourne for helpful comments and discussions. Farzana Hossain gratefully acknowledges funding from the Faculty of Business and Economics Doctoral Program Scholarship for this research.

†Department of Economics, Level 4, FBE Building, 111 Barry Street, University of Melbourne. Email: farzanah@student.unimelb.edu.au.

‡Department of Economics, Level 3, FBE Building, 111 Barry Street, University of Melbourne. Email: rahsan@unimelb.edu.au.

1 Introduction

Climate change has had a significant impact on weather patterns throughout the world. According to the first Intergovernmental Panel on Climate Change (IPCC) report, this impact is likely to continue in to the near future with a projected increase in mean global surface temperatures of 5°F and a projected increase in mean precipitation of 8 percent by the second half of the 21st century. Given the implications for human welfare, fully understanding the impact of these changes in climate on agricultural outcomes is of first-order importance. In the case of precipitation, an emerging consensus in the current literature is that higher precipitation either has no effect on agricultural output (Dell, Jones, and Olken, 2012) or that it actually increases agricultural profits (Deschênes and Greenstone, 2007) and economic growth (Barrios, Bertinelli, and Strobl, 2010).

In this paper, we show that this average effect masks the fact that increases in rainfall can create both winners and losers. Central to this novel finding is our focus on identifying the spatial spillover effect of rainfall. To better understand why this matters, consider an agricultural household in district d. Greater rainfall in d will increase this household's crop output and, for a given crop price, increase its income. However, as we show below, patches of high or low rainfall tend to span multiple districts. Thus, if district d experiences greater rainfall then so will other neighboring districts. The resulting rainfall-induced supply shock in neighboring areas will lower crop prices, which in turn will lower farm incomes in district d. This adverse spatial spillover effect will generally attenuate the positive effect of own-district rainfall and in some cases may even dominate it.

To explore this spillover effect empirically, we use 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. For a given district d, the former is defined as the deviation in a district's rainfall in a given year from its long-term average, which is then normalized by its long-term standard deviation.[1] In contrast, the latter is defined as the weighted average of own-district rainfall shocks in all other

[1] Such an approach to calculating rainfall shocks has been used by Barrios, Bertinelli, and Strobl (2010), Cole, Healy, and Werker (2012), and Emerick (2016), among others.

districts j in the sample, where the weights are the inverse of the distance between d and j.[2]

Our identification strategy incorporates household fixed effects, which allows us to purge the effect of any unobserved, time-invariant household and district characteristics. Thus, our results are identified from within-district variation in own rainfall and neighbor's rainfall from its long-term average. 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 rainfall shocks. Further, our choice of India as a setting for this analysis provides us with an important benefit in implementing this identification strategy. 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.[3]

Our results indicate that both own-district rainfall shocks and neighbor's rainfall shocks have a statistically and economically significant effect on rural household consumption. We find that greater own-district rainfall results in an increase in household consumption. This confirms the positive effect of own-region rainfall that has been documented by Deschênes and Greenstone (2007), among others. We also find that a rainfall shock in neighboring districts has a large negative effect on household consumption. This suggests that a neighbor's rainfall shock significantly attenuates the consumption benefits of own-district rainfall.

Further, our estimates indicate that this attenuating effect is economically significant. For instance, if we ignore the spatial spillover effect of neighbor's rainfall shocks, we find that a one-standard deviation increase in a district's own rainfall shock raises household consumption by 7.72 percent. However, after accounting for the adverse spatial spillover effect, we find that a one-standard deviation increase in a district's own rainfall shock raises household consumption by just 2.54 percent. These results provide empirical support to the view posited by Dell, Jones,

[2] These inverse distance weights ensure that rain shocks in nearby districts play a greater role in determining the size of the neighbor's rainfall shock.

[3] A second benefit of studying this issue in India is that the agricultural production there is mainly un-irrigated and 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 un-irrigated and rain-fed (Economic Survey, 2018). Thus, any adverse spatial spillover effect of rainfall is likely to be of first-order importance in India.

and Olken (2014) that the spatial spillover effects of weather shocks are likely to be of first-order importance. Further, these results also suggest that greater rainfall can create both winners and losers. More precisely, households that reside in districts with a low-to-moderate own-rainfall shock and a large neighbor's rainfall shock may be made worse off from increases in precipitation.

As mentioned above, our finding of a negative spatial spillover effect can be explained by a decrease in crop prices and agricultural income as a result of greater rainfall in neighboring districts. We empirically examine whether these channels are supported by the data. To do so, we first use ICRISAT crop-price data to examine whether a rainfall shock in neighboring districts leads to a reduction in crop prices. These data, which are at the crop and district level, allow us to regress the natural logarithm of crop prices on a district's own rainfall shock and its neighbor's rainfall shock. We further control for crop, district, and year fixed effects in these regressions. Our results suggest that a higher neighbor's rainfall shock does indeed lower crop prices in a district.

Next, we use our baseline household data to regress a household's agricultural income on its own-district rainfall shock as well as neighbor's rainfall shock. Consistent with our story above, we find that households that experience a higher neighbor's rainfall shock experience a reduction in their agricultural income. We find no such effect on non-agricultural salaries and wages as well as on remittances. This suggests that a neighbor's rainfall shock affects rural household welfare through its adverse effect on the agricultural sector. To further confirm this, we show that the negative effect of a neighbor's rainfall shock on household consumption only holds for households that report agriculture as their main source of income. All of these results confirm that the price-based mechanism outlined above is a plausible explanation for our spatial spillover effect.

Our paper is related to a sparse literature that documents the spatial spillover effects of weather shocks. For instance, Burgess and Donaldson (2012) examine whether openness to trade reduces or exacerbates the sensitivity of real incomes to productivity shocks in India. As part of their overall analysis, they show that crop prices in a particular district are negatively related to rainfall in neighboring districts. However, they do not examine the effects of rainfall in neighboring districts on rural household consumption. This is important because even when rainfall in neighboring areas lowers crop prices, its effect on rural household consumption is theoretically ambiguous. As consumers, these households gain from the lower crop prices. However, as producers, these households lose from the lower crop prices since it lowers their farm income. Thus, the net effect of such a crop price reduction on rural household welfare is theoretically ambiguous and is therefore an empirical question. A key advantage of our paper is that we have the disaggregated data necessary to examine the household welfare implications of a neighbor's rainfall shock.[4]

Our paper is also related to a literature that examines the effect of climate change induced variation in temperature and rainfall on agricultural outcomes. This literature can be roughly divided in to three categories. The first relies on experimental data to calibrate an agricultural production function and then uses the latter to simulate the effect of climate change on crop outcomes (Adams, 1989; Adams et al., 1995). A second category of studies uses cross-sectional regressions to estimate the effects of climate change on agricultural outcomes (Mendelsohn, Nordhaus, and Shaw, 1994; Schlenker, Hanemann, and Fisher, 2006). Lastly, a third category uses panel approaches to estimate the effect of climate change on agricultural outcomes (Deschênes and Greenstone, 2007; Dell, Jones, and Olken, 2012). Our contribution to this literature is that in addition to examining the effect of own-region rainfall with panel data, we also examine the spatial spillover effect of rainfall in neighboring regions.

Finally, our paper is also related to an extensive literature that documents the welfare consequences of weather-induced productivity shocks in developing countries. Using weather extremes as an exogenous shock to agricultural productivity, these studies find that weather shocks in an area have a significant impact on agricultural production, employment, and wages (Jayachandran, 2006; Kaur, 2018; Emerick, 2016). Weather shocks have also been shown to have a significant effect on human capital formation in the developing world (Maccini and Yang, 2009; Shah and Steinberg, 2017). We add to this literature by considering the spatial spillover effect of rainfall shocks

[4]Boustan, Rhode, Kahn, and Yanguas (2017) 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. Unlike our analysis, they do not separately estimate the effect of an own-county disaster shock and a neighbor's disaster shock. As a result, they are unable to examine whether the latter effect differs substantially from the former.

in neighboring areas on rural household welfare.

We structure the rest of the paper as follows. In section 2, we use a conceptual framework to highlight the theoretically ambiguous relationship between rainfall in neighboring districts and rural household welfare. In section 3, 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 section 4, we describe the empirical strategy we use to identify the impact of rainfall shocks on household consumption. In section 5, we present our baseline results and address key econometric issues. In section 6, we explore the potential channels through which rainfall shocks in neighboring districts affect household consumption. Finally, in section 7 we provide a conclusion.

2 Conceptual Framework

In this section, we use the canonical agriculture household model described in Singh, Squire, and Strauss (1986) to show that increased rainfall in neighboring districts will have an ambiguous effect on rural household consumption. An important benefit of our discussion below is that it will highlight the mechanisms that drive the relationship between neighbor's rainfall shocks and household consumption. This will allow us to examine whether these mechanisms are supported by our data in section 6.1. To begin, consider a risk-neutral farmer that consumes a staple, agricultural crop, A. This is a crop that the farmer can either purchase at the retail market or produce on his own. We consider this farmer's optimal behavior in two stages.

Stage 1: Production

We assume that at the beginning of the growing season, the farmer must decide how much of the staple crop to produce. He does so using his fixed amount of arable land as well as his endowment of family labor, N. We normalize the fixed unit of arable land to one and assume that the farmer cultivates this land in its entirety,. The farmer's output of the staple crop is then determined by the following production function: $Q_A = \hat{R}^O f(N)$, where Q_A is the quantity of the staple crop produced and f is a production function. $\hat{R}^O$ is the farmer's expectation of the amount of rainfall in his own district during the growing season. Thus, $\hat{R}^O$ serves as a productivity shock for the farmer with greater own-district rainfall resulting in a greater yield of the staple crop, holding N constant.[5]

The farmer's production problem is to pick the level of household farm labor, N, that maximizes his profits. The optimal labor choice, N^*, results in the following farm income: $Y^* = P_A \hat{R}^O f(N^*) - wN^*$, where P_A is the price of the staple crop, which we discuss in more detail below, and w is the prevailing wage rate in the district. We assume that the farmer takes this wage as given. Note here that Y^* is monotonically increasing in P_A and $\hat{R}^O$.

Stage 2A: Price Determination

In the second stage, the actual own-district rainfall during the growing season is realized. This value of R^O along with the optimal labor chosen in the first stage determines the farmer's staple-crop yield, Q_A^*. The farmer must now decide where to sell his crop, which will determine the price he receives for them. He can sell his crop at a nearby *mandi*, which is a government-regulated wholesale agricultural market. While *mandis* are open-outcry auctions where farmers are less vulnerable to being exploited by unscrupulous traders, they also tend to be relatively costly to get to. Goyal (2010) reports that the typical farmer may have to travel 30 to 40 kilometers to reach the closest *mandi*. This may explain why most sales at *mandis* are made by large farmers while small farmers sell mostly to local intermediaries (Chatterjee and Kapur, 2016).

Due to this high transportation cost, the farmer can alternatively sell his crop to a local trader.[6] If he does so, then the price at which the farmer sells to the trader is determined by a negotiation process. We assume that both parties can observe the quality of the crop, which means that the negotiation occurs under complete information. If the negotiation with the trader is successful, the farmer is able to sell his crop at a price P_A. If it is unsuccessful, his outside option is to take his

[5] As in the canonical version of this agriculture household model, we exclude other intermediate inputs and also abstract from the possibility that the household may cultivate multiple crops.

[6] See Chau, Goto, and Kanbur (2009) and Bardhan, Mookherjee, and Tsumagari (2013) for recent models of transactions between farmers and intermediaries in a developing country context and Spulber (1996) for a review of the general literature on intermediation.

crop from his district d to the nearest *mandi* (m) at a transport cost of $\tau_{dm} > 0$. At the *mandi*, the farmer is a price taker and will receive the following price:

$$\underline{P}_A = \frac{P_A^m(R^N)}{\tau_{dm}} < P_A \tag{1}$$

where the *mandi* price P_A^m is decreasing in the output of neighboring districts and hence also decreasing in the amount of rainfall in neighboring districts, R^N.

From the trader's perspective, if the negotiation is successful then he buys the Q_A^* units of the crop at a price of P_A and sells it to processors at a price $P_A^r > P_A$. If negotiations fail, then he gets zero. We follow Chatterjee (2018) and assume that the farmer and the trader engage in Nash bargaining where the farmer's bargaining weight is δ. This results in the following price for the staple crop:

$$P_A = (1-\delta)\underline{P}_A(R^N) + \delta P_A^r \tag{2}$$

Equation (2) suggests that greater rainfall in neighboring districts, R^N, will lower the price received by the farmer for his crop, P_A, by weakening his outside option, $\underline{P}_A$.[7]

Stage 2B: Consumption

Having negotiated a price for his staple crop in stage 2A, the final step for the farmer is to decide his family's total consumption. We assume that the farmer's utility depends on his family's consumption of the staple crop, C_A, as well as his family's total leisure time, L. Leisure is simply the total time endowment, T, minus the hours of farm labor, N. Given crop price P_A and his farm income Y^*, we know that his optimal consumption of the staple, C_A, can be written as $C_A = C(P_A(R^N), Y^*(R^O, R^N))$. In turn, this allows us to write the effect of greater rainfall in

[7]We are abstracting here 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_A 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 MSP's 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.

neighboring districts, R^N, on household consumption as

$$\frac{\mathrm{d}C_A}{\mathrm{d}R^N} = \underbrace{\frac{\partial C_A}{\partial P_A} \times \frac{\partial P_A}{\partial R^N}}_{\text{Own-Price Effect}} + \underbrace{\frac{\partial C_A}{\partial Y^*} \times \frac{\partial Y^*}{\partial P_A} \times \frac{\partial P_A}{\partial R^N}}_{\text{Farm-Income Effect}} \tag{3}$$

Equation (3) suggests that greater rainfall in neighboring districts will affect a farmer's consumption through two distinct channels. The first term on the right-hand-side indicates that by lowering the price of the farmers staple crop, greater rainfall in neighboring districts will raise a farmer's consumption. This is the own-price effect. On the other hand, the second term on the right-hand-side indicates that by lowering the income earned from its staple crop, greater rainfall in neighboring districts will lower a farmer's consumption. This is the farm-income effect. Thus, the net effect of greater rainfall in neighboring districts on this farmer's consumption is ambiguous. Which of these two channels will dominate is therefore an empirical question.

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 and are available for two rounds, 2004–05 and 2011–12 (Desai et al., 2005; Desai and Vanneman, 2012). The raw data cover households in 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.[8] This results in a working sample that consists of 28,087 households in 283 districts across India.[9]

Our key outcome variable is each household's total consumption expenditure per capita.

[8] The only exception to this restriction are new households in 2011–2012 that split off from households in 2004–2005. We retain these households even thought they appear in only IHDS2.

[9] 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.

IHDS constructs this using expenditure on a a series of food and non-food items. They use a mixed recall period - a 30-day recall period for more frequently purchased items and a 365-day recall period for infrequently purchased non-food items. To calculate monthly per capita consumption expenditure, they divide expenditures on infrequently-purchased items by 12 before adding it to the total expenditure on frequently-purchased items. They then divide total household consumption expenditure by the number of household members. This yields each households nominal consumption per capita. To convert these to real values, we use the deflator provided in the raw data. This results in final consumption values that are in constant 2005 Rupees.

Unlike other commonly used household surveys in India, the IHDS data have the advantage that it follows households over time. This provides us with two benefits. First, we can include household fixed effects in our baseline econometric specification. This allows us to control for unobservable, time-invariant household characteristics that may otherwise bias our results. It also allows us to account for unobservable, time-invariant district characteristics such as topographical features that may be correlated with both rainfall and agricultural production. Second, the panel nature of the data allows us to use a balanced sample of households that appear in both survey rounds in our baseline specifications. This ensures that our key results are not being driven by attrition, household migration, or other endogenous compositional changes in the sample.

Table 1 provides descriptive statistics of the households in our IHDS sample. The average household has monthly consumption expenditure of approximately 876 Rupees per person and 20.41 percent of these households are below the official poverty line.[10] The households in our sample are also highly dependent on agricultural production, with 55.17 percent of households reporting agriculture as their main source of income. In addition, the average household in our sample has 5.48 members and 1.75 children. On average, 88.50 percent of households have a male head with an average age of 48.95 and 83.02 percent of households are Hindu.

Our analysis rests on the assumption that households in our sample produce crops for sale to traders and/or agricultural markets. If the households in our sample are mainly subsistence

[10]The monthly expenditure is equivalent to 20.14 U.S. dollars per person in 2005. The poverty indicator is as provided in the IHDS data. This indicator is calculated using the Indian Planning Commission poverty line.

farmers with little or no connection to markets then 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.[11]

To explore these issues further, we report summary statistics on crop sales and market access indicators in Table 2. In Panel A, we examine whether households in our sample are active participants in markets. The data we use to construct the summary statistics in Panel A were only collected during the 2004–2005 IHDS round. Despite this, they provide a useful snapshot of the degree of market participation among the households in our sample. The data suggest that only 3 percent of households in our sample are sharecroppers. Further, 59.84 percent of the households in our sample sell their crops with these sales representing, on average, 34.39 percent of their total production. These numbers suggest that, while the households in our sample are poor, they nonetheless actively sell their crops to traders and/or in agricultural markets.

In Panel B of Table 2, we use various village-level market access measures to examine how isolated the households in our sample are. Unfortunately we do not have household-level data on the distance to the nearest wholesale market, so instead we use several village-level proxies of market access instead. 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 agricultural retail market and 14.26 kilometers away from the nearest town.[12] 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.

[11] Recall from equations (1) and (2) that

$$\frac{\mathrm{d}P_A}{\mathrm{d}R^N} = \left(\frac{1-\delta}{\tau_{dm}}\right) \times \frac{\mathrm{d}P_A^m}{\mathrm{d}R^N}$$

Thus, the pass-through of a neighbor's rainfall shock on to local crop prices is inversely proportional to how far a household is to the regional market, τ_{dm}. Extremely isolated households, i.e. high τ_{dm}, may have a very small pass through.

[12] Note that these agricultural retail markets are not the same as the wholesale markets (i.e. the *mandis* mentioned in section 2).

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 $0.25^o \times 0.25^o$ 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, Hsiang, Schlenker, and Sobel, 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, 2016).

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 Figure 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, Bertinelli, and Strobl (2010), Cole, Healy, and Werker (2012), and Emerick (2016) 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

$$R^O_{dt} = \frac{R_{dt} - \overline{R}_d}{S_d} \tag{4}$$

where R_{dt} is the total monsoon rainfall in a district in year t and $\overline{R}_d$ is each 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^O_{dt} indicates that a district received total monsoon rainfall in a year that was above its long-term average.[13]

In Figure 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 Figure 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:

$$R^N_{dt} = \sum_{j \neq d} \left(\frac{1}{\omega_{dj}} \times R^O_{jt} \right) \tag{5}$$

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^O_{jt} is the own rainfall anomaly shock in district j and year t. Thus, for each district d in year t, equation (5) 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^N_{dt}.[14] The correlation coefficient between a district's own rainfall shock,

[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 it's 80th percentile rainfall value over the period 1979 to 2015. All other districts have an own rainfall shock value of 0 (no rainfall shock)

[14] This measure of neighbor's rainfall shock builds on measures of market access that is frequently used in the trade

R^O_{dt}, and its neighbor's rainfall shock, R^N_{jt}, is 0.77. Such a high correlation follows naturally from the spatial clustering of rainfall evident in Figure 2. Summary statistics for all rainfall variables used in the paper are reported in Table 3.

4 Econometric Strategy

In this section we describe the econometric specification we use to examine the effect of both own rainfall shocks and neighbor's rainfall shocks on household consumption. In section 2, we showed that an agricultural household's total consumption will depend on the price of crops and its total farm income. Farm income will depend on the rainfall in the household's own district while both crop prices and farm income will depend on rainfall in neighboring districts. To capture these effects, we posit the following reduced-form specification:

$$\ln\left(C_{hdt}\right) = \alpha + \beta_1 R^O_{dt} + \beta_2 R^N_{dt} + \gamma_1 X_{hdt} + \theta_h + \theta_t + \epsilon_{hdt} \tag{6}$$

where C_{hdt} is the total consumption for household h in district d and year t, R^O_{dt} is a district d's own district rainfall shock, and R^N_{dt} is the rainfall shock in neighboring districts. Our coefficient of interest is β_2. If a positive rainfall shock in neighboring districts has a negative effect on a household's income, then we would expect β_2 to be negative.

X_{hdt} is a set of household-level variables that is likely to determine its total 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. θ_t is a year fixed effect that captures any country-wide shocks that might be related to household consumption while ϵ_{hdt} is a classical error term.

Importantly, our baseline specification incorporates household fixed effects, θ_h. This provides us with several 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-

and economic geography literature. See Donaldson and Hornbeck (2016) for a recent example.

yield crops tend to 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 factors such as its topographic features and other locational features such as distance to the nearest *mandi* that might affect a household's total consumption.

Another advantage of including household fixed effects is that it allows us to purge the effect of other unobserved, time-invariant household characteristics such as a household's religion, caste status etc. that might impact its consumption. Finally, these fixed effects will also allow us to capture a household's bargaining power in negotiations with a local trader, which we showed in section (2) will affect how rainfall in neighboring districts affect the price a household receives for its crops.

While the inclusion of household fixed effects has key advantages, it is worth noting that our rainfall shock measures, R^O_{dt} and R^N_{dt}, 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 Figure 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 4. In column (1) we estimate a parsimonious version of (6) where we exclude household fixed effects. The coefficient of the own rainfall shock is positive and statistically significant. This confirms the findings of an earlier literature that document the positive effects of rainfall on rural household welfare (Deschênes and Greenstone, 2007). However, the coefficient of the neighbor's rainfall shock variable suggests that having greater rainfall in nearby districts lowers 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.

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 and may also be correlated with rainfall in a district. These controls include a district's latitude and longitude, which are taken from Allen and Atkin (2016). In addition, we also include 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 results in column (2) suggest that the addition of these additional district controls does not appreciably change our results.

Next, in column (3) of Table 4 we report the results from estimating equation (6). 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 be biasing 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 Figure 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, note that we can use (6) to write the effect of an own-district shock (R^O) on household consumption (C) as

$$\frac{\mathrm{dln}\,(C)}{\mathrm{d}dR^O} = \hat{\beta}_1 + \hat{\beta}_2 \frac{\mathrm{d}R^N}{\mathrm{d}R^O} \tag{7}$$

where the second term on the right-hand-side accounts for the fact that rainfall tends to fall in clusters that span across districts (see Figure 2). Thus, greater rainfall in a district is likely to result in greater rainfall in neighboring districts (R^N). Now, in the absence of such a spatial spillover effect, the effect of an own-district rainfall shock is simply $\hat{\beta}_1$. Given the estimates in column (3) of Table 4, it follows that a household that experiences a one-standard deviation increase in its own rainfall shock will experience a 7.72 percent increase in its consumption per capita.

If instead we were to account for the spatial spillover effect, we need to calculate the entire effect given in (7). To implement this, we first need to estimate $\mathrm{d}R^N/\mathrm{d}R^O$. To do so, we aggregate our data to the district-year level. We then regress a district's neighbor's rainfall shock on its own-district shock, district fixed effects, and year fixed effects. The resulting coefficient of the own-district shock is 0.152. This is our estimate of $\mathrm{d}R^N/\mathrm{d}R^O$. Then, if we again use the estimates in column (3) of Table 4, we find that a a household that experiences a one-standard deviation increase in its own rainfall shock will experience a 2.54 percent increase in its consumption per capita. This lower value reflects the fact that $\hat{\beta}_2$ is negative, which means that households are made worse off by greater rainfall in neighboring areas.

These results demonstrate that accounting for this spatial spillover effect gives us a more conservative estimate of the gains from rainfall. In fact, it also suggests that greater rainfall can create both winners and losers. If a household resides in a district that experiences a moderate-to-high own rainfall shock and a low-to-moderate neighbor's rainfall shock, then they will be better off due to greater rainfall. In contrast, if a household resides in a district that experiences a low-to-moderate own rainfall shock and a moderate-to-high neighbor's rainfall shock, then they will be worse off due to greater rainfall. These findings suggest that the average effects of rainfall on various agricultural outcomes found in the earlier literature (Dell, Jones, and Olken, 2012; Deschênes and Greenstone, 2007; Barrios, Bertinelli, and Strobl, 2010, and others) fail to capture this important distributional consequence of greater rainfall.

Finally, we examine the non-linear effects of rainfall shocks. More precisely, we augment our baseline specification in (6) by including a squared term in both own-rainfall and neighbor's rainfall shocks. We report the resulting estimates in column (4) of Table 4. Interestingly, we find that both the level and squared own-district rainfall coefficients are positive and statistically significant. The latter is positive, which suggests that the benefits of own rainfall are increasing in the level of rainfall itself. In the case of the neighbor's rainfall shock, we find that only the squared term is statistically significant. In fact, while the coefficients of both the level and squared terms are negative, the latter is considerably larger. This suggests that the negative effect of a neighbor's rainfall shock that we have documented thus far is being driven by very large neighboring shocks. Relatively modest rainfall shocks in neighboring areas does not have a statistically significant effect on household consumption.

5.2 Econometric Issues

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 equation (6). To the extent that this is the case, the standard errors we report in Table 4 are incorrect even if our estimate of β_2 is unbiased. Our default approach in Table 4 has been to cluster the standard errors at the district-year level to allow household consumption within each district and year to be correlated with each other. We now explore an alternate approach to account for spatial correlation in our error term.

A common approach to adjusting for such spatial correlation is to the use the Conley (1999) standard-error correction. This approach requires the construction of a spatial variance-covariance matrix that incorporates the distance between all observations i and j. In fact, one can combine this spatial-correlation correction with a standard heteroskedastic and auto-correlation correction (HAC) to create spatial-HAC standard errors (Hsiang, 2016). We use this approach to estimate our standard errors in column (1) of Table 5.[15] 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. While the Conley (1999) approach is popular, it 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, pp. 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."

While our baseline specification allows rainfall to be spatially correlated, there may be other channels through which rural household consumption is correlated across space. For instance, nearby districts are likely to have similar farm production technology and soil types (Schenkler and Roberts, 2009; Chen, Chen, and Xu, 2016). These channels could result in spatial correlation in C_{hdt}. One way to account for this is to include a spatial lag (LeSage and Pace, 2009). 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 geocoordinates of each household. Such information are 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 vari-

[15] We use the STATA .ado file `reg2hdfespatial` created by Thiemo Fetzer and used in Fetzer (2014) to implement this.

able to our baseline specification (6). We report the results from estimating this new specification in column (2) of Table 5. 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.

As mentioned above, our baseline specification in (6) incorporates household fixed effects. This allows us to purge the effect of any unobserved, time-invariant household and district characteristics. However, there could be unobservable, time-varying district shocks that threaten our identification strategy. While rainfall shocks themselves are unanticipated, the timing of these shocks can coincide with other time-varying, district-level shocks. For instance, it could be the case that the districts in our sample that received a large neighbor's rainfall shock also experienced a negative productivity shock at the same time. Both of these shocks will lower household consumption, which means that the latter shocks will confound the effects of a neighbor's rainfall shock. To account for these time-varying district shocks, we include in (6) 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. We report the results from this augmented regression in column (3) of Table 5. As these results demonstrate, our coefficient of interest remains highly robust.

6 Additional Results

6.1 Mechanisms

Our results thus far suggest that a positive rainfall shock in neighboring districts will lead to a reduction in a household's consumption. What could explain such an effect? In section 2, we argued that one possible mechanism is that that greater rainfall in nearby districts leads to a positive supply shock, which in turn lowers the price of crops in regional wholesale markets. To the extent that this regional price acts as an outside option for farmers when they negotiate with local traders, the reduction in this regional price will pass-through to the price received by the farmer. In turn, this will lower the income earned by the farmer from selling his crops. Thus, the essential elements of this mechanism is that a positive neighbor's rainfall shock leads to (i) a reduction in crop prices and (ii) a reduction in household agricultural income. We now explore whether these elements are supported by our data.

To test whether a positive neighbor's rainfall lowers crop prices, we use crop price data from the ICRISAT Village Dynamics in South Asia Macro-Meso Database (henceforth ICRISAT). This dataset includes information on 16 major crops in 311 districts across India for the period 1966-67 to 2011-12.[16] 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:

$$\ln\left(P_{cdt}\right) = \alpha_c + \delta_1 R^O_{dt} + \delta_2 R^N_{dt} + \theta_d + \theta_c \times \theta_t + \nu_{cdt} \tag{8}$$

where P_{cdt} is the farm-gate price for crop c in district d and year t. R^O_{dt} and R^N_{dt} are the rainfall shock measures defined above while θ_c, θ_d, and θ_t are crop, district, and year fixed effects respectively. Lastly, ν_{cdt} is a classical 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 equation (8) in column (1) of Table 6. The coefficient of the neighbor's rainfall shock is negative and statistically significant. This confirms that our proposed mechanism of rainfall in neighboring areas lowering household consumption through a reduction in crop prices is a plausible one.

As discussed section 2, a reduction in crop prices due to rainfall in neighboring districts has contrasting effects on rural household consumption. First, by lowering overall food prices, rainfall in neighboring districts will increase household consumption. In contrast, by lowering crop prices, rainfall in neighboring districts will lower farm income and hence household consumption. We now examine whether this second effect is supported by our data. The IHDS data provides a

[16]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 prices that are above the 95th percentile of the crop-price distribution.

breakdown of each households income by source. We use this to examine whether own-district rainfall and rainfall in neighboring districts affect farm income in a manner that is consistent with our hypothesis above.

To implement this, we first choose agricultural wage income as our proxy for farm income. This measure of income is less vulnerable to measurement error compared to more direct measures of agricultural profits.[17] We then estimate a version of (6) where the dependent variable is the natural logarithm of a household's agricultural wage income. We report the results from estimating this regression in column (2) of Table 6. These result suggest that a greater own rainfall shock raises farm wage income while a greater neighbor's rainfall shock lowers farm wage income. These results are both fully consistent with our hypothesis above.

The results above show that both own rainfall shocks and neighbor's rainfall shocks affect agricultural households in a manner that is consistent with our hypothesis. We now check whether these shocks affect agricultural households differently than non-agricultural households. One would expect both the positive effects of greater own rainfall shocks and the negative effects of greater neighbor's rainfall shocks to be larger in magnitude for agricultural households. To test this, we define an agricultural household as one that reports agricultural income (either farm profits or farm wages) as their primary source of income.[18] We classify all other households as being non-agricultural. We then estimate the following version of equation (6):

$$\begin{aligned}\ln\left(C_{hdt}\right) &= \alpha_1+\beta_1^A R_{dt}^O\times A_h+\beta_1^{NA} R_{dt}^O\times NA_h+\beta_2^A R_{dt}^N\times A_h+\beta_2^{NA} R_{dt}^N\times NA_h \\ &\quad +\gamma_1 X_{hdt}+\theta_h+\theta_t+\varepsilon_{hdt} \end{aligned} \tag{9}$$

where A_h takes the value of one if household h is an agricultural household while NA_h takes the value of one if household h is a non-agricultural household. All other variables are as described above and ε_{hdt} is a classical error term. In equation (9) we are decomposing the effects of own

[17]The raw farm wage data does have significant outliers. To ensure that our results are not driven by these outliers, we omit from our sample households that report farm wages that is equal to or above the 95th percentile.

[18]More precisely, agricultural households are ones who report that their principal income is from either cultivation, allied agriculture, or agricultural wages. Thus, these households include both cultivators and agricultural laborers.

rainfall and neighbor's rainfall into an effect that is agricultural-household specific and one that is non-agricultural-household specific. If our proposed mechanism is correct, then we should expect $\hat{\beta}_1^A > \hat{\beta}_1^{NA}$ and $\hat{\beta}_2^A < \hat{\beta}_2^{NA}$.

We report the results from estimating equation (9) in column (3) of Table 6. These results confirm that the positive effects of greater own rainfall shocks and the negative effects of greater neighbor's rainfall shocks are larger in magnitude for agricultural households. Thus, taken together, the results in Table 6 are fully consistent with our hypothesis that greater rainfall in neighboring districts can lower household consumption through a reduction in crop prices and a reduction in agricultural income.

6.2 Alternate Channels

In this section we examine alternate channels through which a neighbor's rainfall shock can affect rural household consumption. The mechanism we proposed in section 2 was one where neighbor's rainfall shocks affected rural households through its effect on agricultural income. If these shocks also affect these households through non-agricultural channels, then it would suggest that our discussion in section 2 was incomplete. We begin by examining the effect of these shocks on income and wages from non-agricultural sources. This is motivated by an existing literature that shows that own-district weather fluctuations in rural areas can lead to a reallocation of economic activity from agriculture to non-agriculture (Emerick, 2016; Santangelo, 2016; Colmer, 2017). We now examine whether a similar effect exists for neighbor's rainfall shocks. To do so, we first estimate a version of equation (6) where we change the dependent variable to salary income from non-farm sources. These results are reported in column (1) of Table 7. They suggest that both the effect of own rainfall shocks and neighbor's rainfall shocks are statistically insignificant. In column (2) we repeat the analysis above, but use a household's non-farm wage income as the dependent variable.[19] As in column (1), we find that both the effect of own rainfall shocks and neighbor's rainfall shocks are statistically insignificant.

[19] 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-far wages.

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 (3) of Table 7 where the dependent variable is now the natural logarithm of each household's remittance earnings per capita. As with columns (1) and (2), we find that both the effect of own rainfall shocks and neighbor's rainfall shocks on a household's remittance earnings are statistically insignificant. Together, the results in Table 7 suggest that the household-consumption effects we've document thus far are not being driven by changes in non-agricultural industries or due to out migration.

6.3 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 as well as on types of expenditure such as schooling and medical that are likely to have long-term consequences. We begin in columns (1) and (2) of Table 8 by decomposing total household consumption in to food consumption and non-food consumption. 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 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 these data are available. The impact of rainfall on schooling is both theoretically ambiguous and empirically contested.[20] 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.

6.4 Alternate Rainfall and Consumption Measures

We next examine whether our main findings are robust to using alternate measures of rainfall and consumption. We report the results from this exercise in Table 9. In column (1), we follow Jayachandran (2006) and construct categorical measures of rainfall shocks. More precisely, for each district we create an own positive shock variable that takes the value of one if a district's annual

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

monsoon rainfall is above the 80th percentile of that district's monsoon rainfall over the period 1979 to 2015. Recall that this is the entire period for which we have rainfall data. Similarly, for each district, we construct a neighbor's positive shock measure that takes the value of one if a district's annual neighbor's monsoon rainfall is above the 80th percentile of that district's neighbors rainfall over the period 1979 to 2015. 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 these categorical measures 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 4.

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 (2016) 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 (2) 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 (3) 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 4.

Finally, in column (4) 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 (4) of Table 9. As the results confirm, the effect of both own and neighbor's rainfall shocks are very similar to the baseline.

7 Conclusion

In this paper, we showed that greater rainfall can create both winners and losers. Central to this novel conclusion is our focus on estimating the effect of both own-district rainfall and rainfall in neighboring districts on rural household consumption. This is contrast to the typical approach in the literature that only examines the effect of an own-region rainfall shocks. This literature concludes that greater rainfall either has no effect on agricultural output (Dell, Jones, and Olken, 2012) or that it actually increases agricultural profits (Deschênes and Greenstone, 2007) and economic growth (Barrios, Bertinelli, and Strobl, 2010).

In contrast, we focus on the spatial spillover effect of rainfall. This is motivated by the observation that rainfall patches tend to span multiple districts. This means that if a district receives greater rainfall then it is likely that neighboring districts will also receive greater rainfall. To examine the implication of this, we first described a conceptual framework where a farmer receives greater yield due to greater rainfall in his own district. However, greater rainfall in neighboring districts results in a positive supply shock that drives 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. Thus, when we consider both own-district rainfall as well as neighboring-district's rainfall, the overall effect of rainfall on household welfare is theoretically ambiguous. The adverse effect of neighbor's rainfall will generally attenuate the positive effect of own-district rainfall and in some cases may even dominate it.

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. Conditional on including household fixed effects, 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. Further, they suggested that neighbor's rainfall shocks, on their own, lowered rural household consumption and therefore significantly attenuated the benefit of own-district rainfall shocks. For instance, if we ignored the spatial spillover effect of neighbor's rainfall shocks, we found that a one-standard deviation increase in a district's own rainfall shock raised household consumption by 7.72 percent. However, after accounting for this spatial spillover effect, we found a one-standard deviation increase in a district's own rainfall shock raised household consumption by just 2.54 percent.

These results support the view that one must account for spatial spillover effects to correctly estimate the welfare effects of own-district rainfall shocks. They also suggest that the findings of the previous literature that greater rainfall has either benign or positive effects on agricultural outcomes should be interpreted with caution. Instead, we showed that depending on how much rainfall fell in neighboring districts, greater rainfall can make rural households better off or worse off. While this novel result 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 mitigation strategies adopted by the households in our sample. For instance, we did not examine whether households adjust their production choices during periods of positive or negative neighbor's rainfall shocks and whether these households adjusted their non-monsoon crops due to these monsoon-season neighbor's rainfall shocks. Exploring these mitigation strategies is a fruitful avenue for future research.

References

[1] Adams, R.M., 1989. "Global Climate Change and Agriculture: An Economic Perspective." *American Journal of Agricultural Economics*, 71(5): 1272–1279.

[2] Adams, R.M., Fleming, R.A., Chang, C., McCarl, B.A., and Rosenzweig, C., 1995. "A Reassessment of the Economic Effects of Global Climate Change on U.S. Agriculture." *Climatic Change*, 30(2): 147–167.

[3] Aditya, K.S., Subash, S.P., Praveen, K.V., Nithyashree, M.L., Bhuvana, N., and Sharma, A., 2017. "Awareness about Minimum Support Price and Its Impact on Diversification Decision of Farmers in India." *Asia & the Pacific Policy Studies*, 4(3): 514–526.

[4] Allen, T., and Atkin, A., 2016. "Volatility and the Gains from Trade." Dartmouth University, mimeograph.

[5] Auffhammer, M., Hsiang, S.M., Schlenker, W., and Sobel, A., 2013. "Using Weather Data and Climate Model Output in Economic Analyses of Climate Change." *Review of Environmental Economics and Policy*, 7(2): 181–198.

[6] Bardhan, P., Mookherjee, D., and Tsumagari, M., 2013. "Middlemen Margins and Globalization." *American Economic Journal: Microeconomics*, 5(4): 81–119

[7] Barrios, S., Bertinelli, L., and Strobl, E., 2010. "Trends in Rainfall and Economic Growth in Africa: A Neglected Cause of the African Growth Tragedy." *Review of Economics and Statistics*, 92(2): 350–366.

[8] Burgess, R., and Donaldson, D., 2012. "Railroads and the Demise of Famine in Colonial India." MIT, mimeograph.

[9] Boustan, L.P., Kahn, M.E., Rhode, P.W., and Yanguas, M.L., 2017. "The Effect of Natural Disasters on Economic Activity in U.S. Counties: A Century of Data." National Bureau of Economic Research Working Paper No. 23410.

[10] Chatterjee, S., 2018. "Market Power and Spatial Competition in Rural India." Princeton University, mimeograph.

[11] Chatterjee, S., and Kapur, D., 2016. "Understanding Price Variation in Agricultural Commodities in India: MSP, Government Procurement, and Agriculture Markets." *India Policy Forum*.

[12] Chau, N., Goto, H., and Kanbur, R., 2009. "Middlemen, Non-Profits, and Poverty." IZA Discussion Paper No. 4406.

[13] Chen, S., Chen, X., and Xu, J., 2016. "Impacts of Climate Change on Agriculture: Evidence from China." *Journal of Environmental Economics and Management*, 76: 105–124.

[14] Cole, S., Healy, A., and Werker, E., 2012. "Do Voters Demand Responsive Governments? Evidence from Indian Disaster Relief." *Journal of Development Economics*, 97(2): 167–181.

[15] Colmer, J., 2016. "Weather, Labor Reallocation, and Industrial Production: Evidence from India" University of Virginia, mimeograph.

[16] Conley T., 1999. "GMM estimation with cross sectional dependence." *Journal of Econometrics*, 92(1): 1–45

[17] Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P., Kobayashi, S., Andrae, U., Balmaseda, M. A., Balsamo, G., Bauer, P., Bechtold, P., Beljaars, A. C. M., van de Berg, L., Bidlot, J., Bormann, N., Delsol, C., Dragani, R., Fuentes, M., Geer, A. J., Haimberger, L., Healy, S. B., Hersbach, H., Hlm, E. V., Isaksen, L., Kllberg, P., Khler, M., Matricardi, M., McNally, A. P., Monge-Sanz, B. M., Morcrette, J.-J., Park, B.-K., Peubey, C., de Rosnay, P., Tavolato, C., Thpaut, J.-N. and Vitart, F., 2011. "The ERA-Interim Reanalysis: Configuration and Performance of the Data Assimilation System." *Quarterly Journal of the Royal Meteorological Society*, 137: 553–597.

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

[19] Dell, M., Jones, B.F., Olken, B.A., 2014. "What Do We Learn from the Weather? The New Climate Economy Literature." *Journal of Economic Literature*, 52(3): 740–798.

[20] Desai, S., and Vanneman, R., National Council of Applied Economic Research, New Delhi, 2005. India Human Development Survey (IHDS), 2005. ICPSR22626-v11. Ann Arbor, MI: Inter-university Consortium for Political and Social Research.

[21] Desai, S., and Vanneman, R., 2012. India Human Development Survey-II (IHDS-II), 2011-12. ICPSR36151-v5. Ann Arbor, MI: Inter-university Consortium for Political and Social Research.

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

[23] Donaldson, D., and Hornbeck. R., 2016. "Railroads and American Economic Growth: A "Market Access" Approach." *Quarterly Journal of Economics*, 131(2): 799–858.

[24] Economic Survey, 2018. "Climate, Climate Change, and Agriculture." Ministry of Finance, Govenment of India.

[25] Goyal, A., 2010. "Information, Direct Access to Farmers, and Rural Market Performance in Central India." *American Economic Journal: Applied Economics*, 2(3): 22–45.

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

[27] Emerick, K., 2016. "Agricultural Productivity and the Sectoral Reallocation of Labor in Rural India." Tufts University, mimeograph.

[28] Fetzer, T., 2014. "Can Workfare Programs Moderate Violence? Evidence from India." STICERD Working Paper.

[29] Jayachandran, S., 2006. "Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries." *Journal of Political Economy*, 114(3): 538–575.

[30] Kaur, S., 2018. "Nominal Wage Rigidity in Village Labor Markets." *American Economic Review*, forthcoming.

[31] LeSage, J., and Pace, R.K., 2009. *Introduction to Spatial Econometrics*. Boca Raton: CRC Press.

[32] Maccini, S., and Yang, D., 2009. " Under the Weather: Health, Schooling, and Economic Consequences of Early-Life Rainfall." *American Economic Review*, 99(3): 1006–1026.

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

[34] Santangelo, G., 2016. "Firms and Farms: The Impact of Agricultural Productivity on the Local Indian Economy." University of Cambridge, mimeograph.

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

[36] Shah, M., and Steinberg, B.M., 2017. "Drought of Opportunities: Contemporaneous and Long-Term Impacts of Rainfall Shocks on Human Capital." *Journal of Political Economy*, 125(2): 527–561.

[37] Singh, I., Squire, L., and Strauss, J., 1986. *Agricultural Household Models: Extensions, Applications and Policy*. Baltimore: Johns Hopkins University Press.

[38] Spulber, D., 1996. "Market Microstructure and Intermediation." *Journal of Economic Perspectives*, 10(3): 135–152.

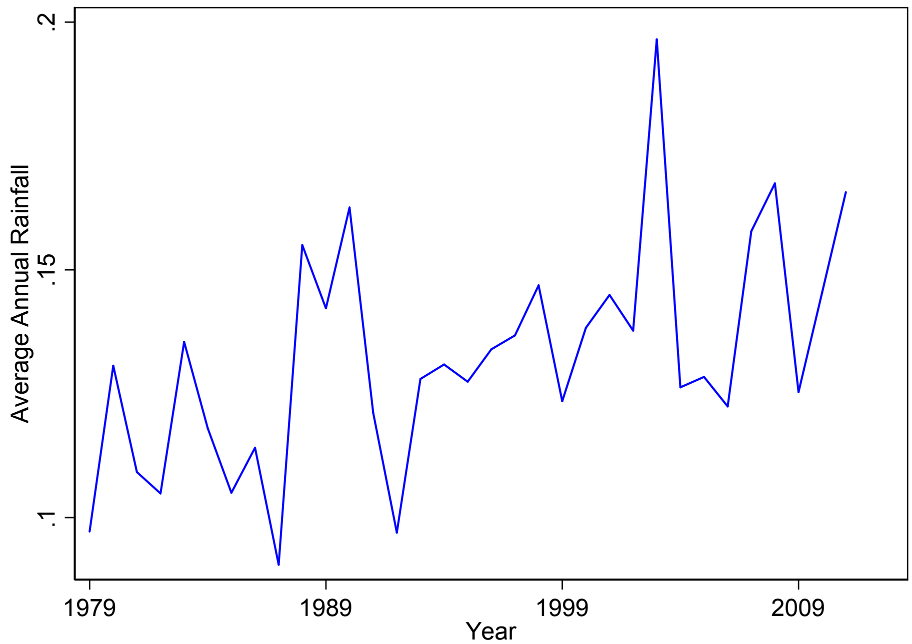

Figure 1: Trends in Average Annual Rainfall in India (1979–2011)

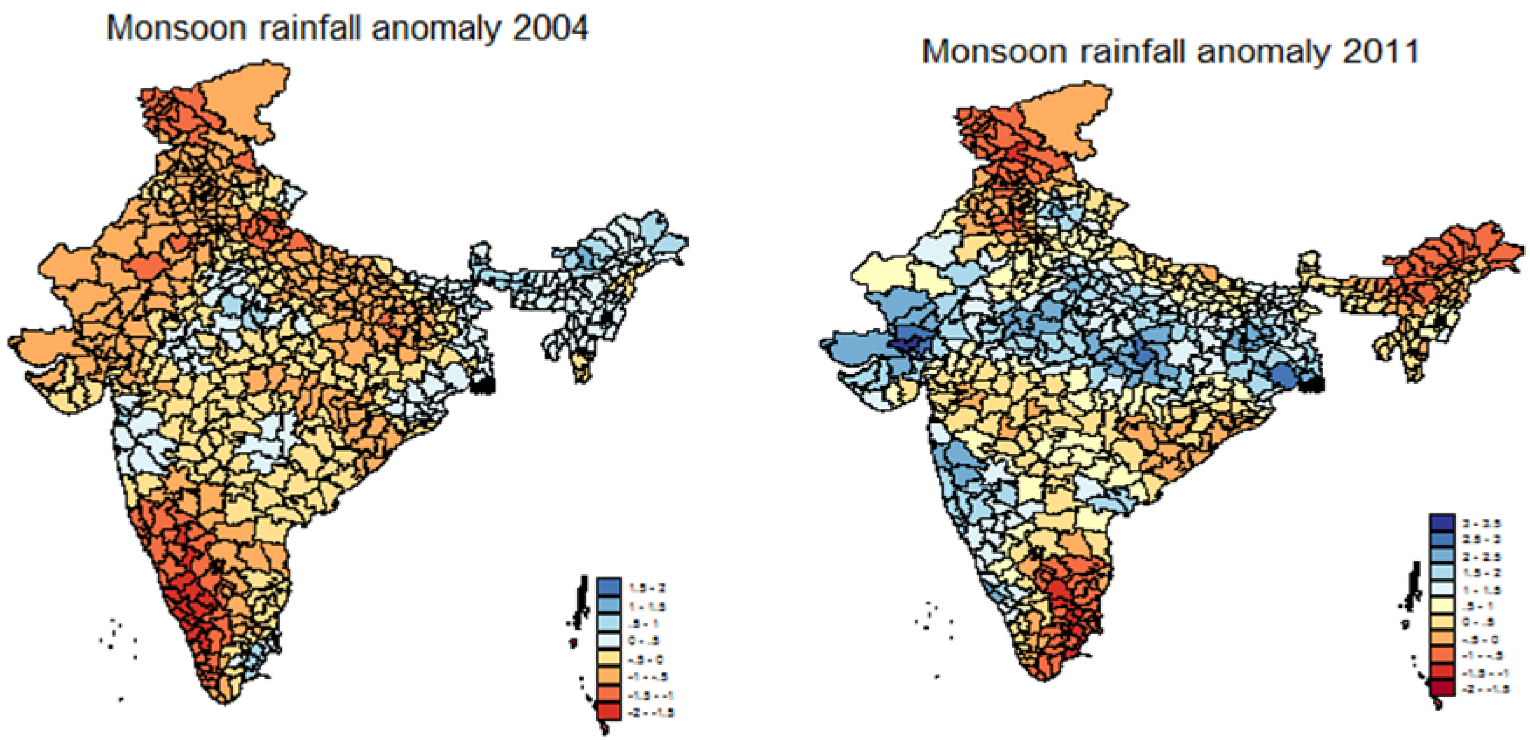

Figure 2: Spatial Variation in Rainfall in India

Table 1: Descriptive Statistics of IHDS Households

	(1)
Household Consumption per Capita	876.096
	[958.283]
Indicator for Households in Poverty	0.204
	[0.403]
Main Income Source is Agriculture	0.552
	[0.497]
Number of Household Members	5.475
	[2.868]
Number of Children	1.748
	[1.698]
Male Household Head	0.885
	[0.319]
Household Head's Age	48.950
	[13.715]

Notes: this table reports average values for various household indicators. These averages were taken over both survey rounds. We report the standard deviation for each variable in the square brackets. 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 was approximately equal to 43.5 Indian rupees in 2005.

Table 2: Agricultural Production and Market Access in the 2004–2005 IHDS Sample

	(1)
Panel A: Household's Agricultural Production	
Fraction of Households that are Sharecroppers	0.030
	[0.172]
Fraction of Households that Sell Crops	0.598
	[0.490]
Share of Output Sold	0.344
	[0.363]
Panel B: Village's Access to Markets	
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]

Notes: this table reports summary statistics for various household production characteristics and village-level access to markets indicators from the 2004–2005 IHDS sample. Panel A reports the sample averages and standard deviation in brackets for various household agricultural production characteristics while Panel B reports the sample averages and standard deviation in brackets for various aspects of a village's access to markets. Note that the numbers in Panel B are calculated at the village level. All distances reported in the table are in kilometers.

Table 3: Rainfall Summary Statistics

	(1)
Total Rainfall	0.124
	[0.086]
Own Rainfall Shock	0.058
	[0.914]
Neighbor's Rainfall Shock	0.106
	[0.574]
Own Positive Shock	0.235
	[0.424]
Neighbor's Positive Shock	0.245
	[0.218]

Notes: this table reports the sample averages for various rainfall and rainfall shock measures along with its standard deviation in brackets. Total rainfall is reported in meters. Own rainfall shocks and neighbor's rainfall shocks are calculated using the anomaly approach described in the text. In contrast, own positive shock takes the value of one if a district's monsoon rainfall was above the 80th percentile for that district during 1979 to 2015. Similarly, neighbor's positive shock takes the value of one if a district's neighbor's rainfall was above the 80th percentile for that district during 1979 to 2015.

Table 4: Spillover Effects of Rainfall on Household Consumption

	(1)	(2)	(3)	(4)
Dependent variable	Ln(Consumption Per Capita)			
Own Rainfall Shock	0.127***	0.073***	0.085***	0.048*
	(0.029)	(0.024)	(0.024)	(0.026)
Neighbor's Rainfall Shock	-1.148***	-0.424***	-0.374***	-0.096
	(0.130)	(0.159)	(0.140)	(0.139)
Own Rainfall Shock Squared				0.023**
				(0.011)
Neighbor's Rainfall Shock Squared				-0.218***
				(0.058)
Constant	6.056***	6.937***	6.370***	6.493***
	(0.068)	(0.726)	(0.066)	(0.065)
Time-Invariant Controls	No	Yes	–	–
Household Fixed Effects	No	No	Yes	Yes
Observations	54,519	52,657	54,541	54,541
R-squared	0.216	0.293	0.247	0.250

Notes: the dependent variable in all columns is a household's monthly per capita consumption expenditure in rural India. The construction of the own rainfall shock and neighbor's rainfall shock variables are described in the text. All regressions control for the number of children in a household, the household head's age, age squared, and whether the household head is male. In column (2), we control for a household's religion, caste, its district's latitude, longitude, total population in 1987, share of agricultural workers in 1987, the 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 columns (3) and (4). All regressions include year fixed effects. All regressions also incorporate sampling weights to ensure that our sample reflects the population. Robust standard errors in parenthesis are clustered at the district-year level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 5: Econometric Issues

	(1)	(2)	(3)
Dependent variable	Ln(Consumption Per Capita)		
Own Rainfall Shock	0.085***	0.073***	0.070***
	(0.025)	(0.023)	(0.025)
Neighbor's Rainfall Shock	-0.374**	-0.293**	-0.310**
	(0.164)	(0.128)	(0.142)
Constant	6.370***	-6.252***	6.383***
	(0.066)	(2.214)	(0.067)
District Spatial Lag Included	No	Yes	No
Additional District Controls	No	No	Yes
Observations	54,541	54,541	52,666
R-squared	0.247	0.256	0.249

Notes: the dependent variable in all columns is a household's monthly per capita consumption expenditure in rural India. The construction of the own rainfall shock and neighbor's rainfall shock variables are described in the text. In column (1) we report Conley (1999) spatial correlation-adjusted standard errors. In column (2) we include a district-level spatial lag of average household consumption per capita. Finally, in column (3) we include the interaction between a district's share of agricultural employment in 1987 and year fixed effects. All regressions control for the number of children in a household, the household head's age, age squared, and whether the household head is male, household fixed effects, and year fixed effects. All regressions also incorporate sampling weights to ensure that our sample reflects the population. The standard errors in parenthesis in columns (2) and (3) are robust and clustered at the district-year level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 6: Mechanisms

	(1)	(2)	(3)
Dependent variable	Ln(Crop Price)	Ln(Farm Wage)	Ln(Consumption Per Capita)
Own Rainfall Shock	0.011	0.143**	
	(0.007)	(0.060)	
Neighbor's Rainfall Shock	-0.087***	-1.963***	
	(0.033)	(0.399)	
Own Rainfall Shock × Non-Agri. Households			0.025
			(0.040)
Own Rainfall Shock × Agri. Households			0.133***
			(0.034)
Neighbor's Rainfall Shock × Non-Agri. Households			-0.019
			(0.310)
Neighbor's Rainfall Shock × Agri. Households			-1.060***
			(0.290)
Observations	8,135	17,497	37,658
R-squared	0.794	0.096	0.759

Notes: the dependent variable in column (1) is the natural logarithm of crop prices, which are drawn from the ICRISAT data. The unit of observation in column (1) is crop-district-year, which is why the sample size there is smaller. The dependent variable in column (2) is the natural logarithm of household-level farm wage earnings. The sample in this column is restricted to households that have positive farm wage earnings. Lastly, the dependent variable in column (3) is a household's consumption per capita in natural logarithm. Agricultural households are those that report agricultural income as their main source of income. All other households are classified as being non-agricultural households. The regressions in columns (2) and (3) control for the number of children in a household, the household head's age, age squared, and whether the household head is male. The regression in column (1) includes a district fixed effect and crop and year interaction fixed effects. The remaining regressions all include household fixed effects and year fixed effects. All regressions also include a constant that is not reported. The regressions in columns (2) and (3) also incorporate sampling weights to ensure that our sample reflects the population. Robust standard errors in parenthesis are clustered at the district-year level in all columns. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 7: Alternate Channels

	(1)	(2)	(3)
Dependent variable	Ln(Non-Farm Salary)	Ln(Non-Farm Wage)	Ln(Remitt ances)
Own Rainfall Shock	-0.072	0.013	0.092
	(0.061)	(0.086)	(0.155)
Neighbor's Rainfall Shock	0.401	0.019	-1.508
	(0.380)	(0.333)	(0.976)
Constant	7.682***	8.706***	5.765***
	(0.693)	(0.182)	(0.769)
Observations	8,234	17,074	4,145
R-squared	0.126	0.076	0.334

Notes: the dependent variable in column (1) is the natural logarithm of a household's total non-farm salary income per capita. The dependent variable in column (2) is the natural logarithm of a household's non-farm wages per capita. Lastly, the dependent variable in column (3) is a household's total remittance earnings per capita. The construction of the own rainfall shock and neighbor's rainfall shock variables are described in the text. All regressions control for the number of children in a household, the household head's age, age squared, and whether the household head is male. All regressions include household fixed effects and year fixed effects. All regressions also incorporate sampling weights to ensure that our 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.

Table 8: Spillover Effects of Rainfall – By Expenditure Type

	(1)	(2)	(3)	(4)
Dependent variable	Ln(Expenditure Per Capita on)			
	Food	Non-Food	Schooling	Medical
Own Rainfall Shock	0.064***	0.033	-0.020	0.085
	(0.021)	(0.039)	(0.042)	(0.058)
Neighbor's Rainfall Shock	-0.226*	0.027	-0.359	0.191
	(0.129)	(0.185)	(0.233)	(0.309)
Constant	5.860***	4.526***	1.470***	4.816***
	(0.060)	(0.096)	(0.179)	(0.188)
Observations	54,519	54,313	33,453	41,077
R-squared	0.127	0.328	0.245	0.055

Notes: 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, and whether the household head is male. All regressions also include household fixed effects and year fixed effects and incorporate sampling weights to ensure that our sample reflects the population. Robust standard errors in parenthesis are clustered at the district-year level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 9: Robustness Checks

	(1)	(2)	(3)	(4)
Dependent variable	Ln(Cons. Per Capita)		Winsorized Ln(Cons. PC)	Ln(Cons.)
Own Positive Shock	0.135***			
	(0.031)			
Neighbor's Positive Shock	-0.748***			
	(0.217)			
Own Rainfall Shock - UDEL		-0.003		
		(0.021)		
Neighbor's Rainfall Shock - UDEL		-0.153**		
		(0.071)		
Own Rainfall Shock			0.084***	0.082***
			(0.023)	(0.023)
Neighbor's Rainfall Shock			-0.368***	-0.365***
			(0.135)	(0.123)
Observations	54,541	52,137	54,541	54,541
R-squared	0.249	0.252	0.251	0.122

Notes: the dependent variable in columns (1) and (2) is a household's monthly per capita consumption expenditure in rural India. In column (3), 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. Lastly, in column (4), the dependent variable is a household's total monthly consumption. The construction of the own rainfall shock and neighbor's rainfall shock variables are described in the text. Own positive shock is a binary variable that takes the value of one if a district received a 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. Own rainfall shock - UDEL and neighbor's rainfall shock - UDEL are our baseline shock measures 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, and whether the household head is male, household fixed effects, year fixed effects, and a constant that is not reported. All regressions also incorporate sampling weights to ensure that our 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.