The Monsoon Shock in Rural Nepal

Panel Evidence from the Household Risk and Vulnerability Survey

Hanan G. Jacoby[*] and Thomas Walker[**]

November 20, 2019

ABSTRACT

Monsoon rainfall is a key driver of economic life in rural Nepal as well as a major source of income variability. In this paper, we use a newly collected 3-year panel data set, representative of rural Nepal, merged with global monthly precipitation data to investigate the nature of the monsoon shock and to quantify household vulnerability to it. We find that the impact of the monsoon shock is concentrated in communities where water-intensive paddy dominates wet season cultivation and, coincidently, where groundwater irrigates dry season cultivation. In these communities, household size, area cultivated, agricultural and non-agricultural income, and household per capita food consumption measured nine months after the wet season harvest, all decline in response to a negative monsoon rainfall shock. A one standard deviation fall in monsoon precipitation is estimated to reduce total income by 3.8% and lead to a 0.8% drop in food consumption for the average rural household, but these figures rise to 11.5% and 3.3%, respectively, for households in the most paddy-intensive communities. These results have implications for social protection policies, especially in the lowlands of Nepal.

[*] Development Research Group, World Bank (hjacoby@worldbank.org). Thanks to Srinivasan Vasudevan and Qiuyan Yin for able assistance.

[**] Social Protection & Jobs Global Practice, World Bank (twalker@worldbank.org).

1. Introduction

Monsoon rainfall is a key driver of economic life in rural Nepal as well as a major source of income variability. Roughly 80% of annual precipitation falls during the summer months of June to September. Agricultural production, especially the principal wet-season crop, paddy, is highly dependent upon the vagaries of the monsoon. The extent to which the resulting output fluctuations lead to corresponding changes in consumption or welfare depends upon households' exposure to the shock and their ability to smooth. Since the monsoon shock is by its very nature highly covariate, households that can smooth likely must do so largely by self-insuring through savings or dis-saving, or through more costly mechanisms. Up to now, however, lack of high frequency household-level panel data has limited our understanding of household responses to monsoon rainfall shocks in Nepal.

In this paper, we use a newly collected 3-year panel data set, representative of rural Nepal, to investigate the nature of the monsoon shock and quantify the extent to which households are vulnerable to it. The Nepal Household Risk and Vulnerability Survey (NHRVS) is a nationally representative multi-topic rural survey, carried out in three waves from 2016-18, covering 6000 households in 50 districts. Four hundred primary sampling units corresponding roughly to rural municipalities (VDCs) were randomly chosen, spanning Mountains, Hills, and the Terai lowlands on the Indo-Gangetic plain, regions representing starkly different agro-ecological zones. To construct monsoon rainfall, we merged gridded monthly precipitation from the global Climate Hazards Group InfraRed Precipitation with Station data (CHIRPS) to each household using their GPS coordinates.

There is ample evidence, most of it from India, that monsoon rainfall shocks have significant income, consumption, and human capital impacts (e.g., Jacoby and Skoufias 1997, 1998; Rose 1999; Kochar 1999; Shah and Steinberg 2017). For rural Nepal, Tiwari et al (2016) show that young children gain weight in relatively good monsoon years, presumably through an income channel, although these shocks have little or no lasting impact on stature later in life. A goal of this study is to explore the income channel in more detail, specifically the extent to which rural incomes in Nepal are vulnerable to drought conditions and how such income shortfalls translate into household welfare changes.

We find that the impact of the monsoon shock is concentrated in communities where water-intensive paddy dominates wet season cultivation. In these communities, household size, area cultivated, agricultural and non-agricultural income, and household per capita food consumption, the latter measured nine months after the wet season harvest, all decline in response to a negative monsoon rainfall shock. A one standard deviation fall in monsoon precipitation is estimated to reduce total income by 3.8% and lead to a 0.8% drop in food consumption for the average rural household (excluding those in the Mountains), but these figures rise to 11.5% and 3.3%, respectively, for households in the most paddy-intensive communities (i.e., those in the top 5%). These findings suggest that provision of free or heavily subsidized rainfall insurance, public works employment guarantees (such as India's NREGA), or similar social protection policies, may be warranted, especially in the Terai, not only because this is the poorest of Nepal's three regions, but because it is the most dependent on paddy cultivation and hence the most vulnerable to the monsoon shock.

After describing the data and context in Section 2, we discuss the empirical approach in Section 3, followed by the estimation results in Section 4. Section 5 draws together the conclusions and implications of the empirical analysis.

2. Data and context

2.1 Nature of the Monsoon in Nepal

During the June-September monsoon period, the plains of the Terai receive around 1400 mm of precipitation, around 84% of this region's annual total, whereas the Mountains and Hills receive 1200 and 1266 mm, respectively. Moreover, from 1981-2018, the coefficient of variation (ratio of standard deviation to the means) of monsoon rainfall was highest in the Terai (0.20) than in the Mountains (0.16) or Hills (0.17). Figure 1 gives a district-level picture of the standard deviation of log monsoon rainfall for the 50 HRVS districts, showing the most variability in the Terai, along Nepal's border with India, with the slight exception of the extreme southeast.

Figure 2-4 trace out the geographic pattern of monsoon rainfall over the period of our study corresponding to calendar years 2015-17 (and HRVS rounds 2016-18). The southeast corner of Nepal consistently receives the highest amount of precipitation, although its position relative to the rest of the country fluctuates from year to year. More generally, the maps reveal considerable year-to-year rainfall variability across districts of Nepal over the study period.

2.2 Features of Agricultural Production

Perhaps the most striking feature distinguishing the three geographical regions of Nepal economically is the extent of paddy cultivation. As seen in Figure 5, the wet season in the Terai lowlands is dominated by rice cultivation, with all but about 5% of area devoted to paddy. By contrast, in the Hills and Mountains, cultivation of paddy runs second to that of maize, and, in the latter region, even Millet is not

far behind rice. In the dry season, wheat is the dominant crop, with roughly equal importance across regions. Figure 6 provides district level estimates of the relative importance of wet season paddy.

Wet season cultivation, to be sure, contributes more to income than dry season cultivation, partly because fewer households can cultivate in the latter period due to lack of irrigation. However, the dominance of the wet season is not overwhelming. Figure 7 compares two ratios across years and geography. The first is the ratio of net revenue from cultivation in the wet season to total net revenue from cultivation.[1] The second is the ratio of cultivated area in the wet season to total cultivated area. The net revenue ratio is consistently higher than the cultivated area ratio, suggesting that wet season crops are more remunerative than dry season crops on average. Both ratios, however, show consistent patterns. Despite the importance of high value paddy, which is grown exclusively in the wet season, the relative economic importance of the wet season is actually *lower* in the Terai than in the mountains or Hills (see Figure 8 for a more disaggregated picture). There also appears to have been a broad, albeit modest, decline in the importance of the wet season across the three waves of the HRVS, although this trend is not so pronounced in the area ratio, which suggests that it is not an extensive (i.e., cultivation) margin phenomenon.

2.3 Sources of household income

Looking at income more broadly, in Figure 9, we see that over the three years of the HRVS the share derived from agriculture has declined slightly, from just under 30% in 2016 to just above 25% in 2018, as the nonagricultural income share has risen. Note that, in addition to net revenue from cultivation,

[1] Total revenue is computed as the sum of harvest quantities times median prices for the region and year by crop. Total cost is computed at the household level as the sum of several purchased input costs by season. However, as with most household surveys a direct estimate of seasonal net revenues (TR – TC) produces an inordinate number of negative values, which are likely due to measurement error. To guarantee positive net revenue (NR), we calculate $NR_i = TR_i(I - \frac{TC_r}{TR_r})$, where TC_r is the average total cost in region *r* and TR_r is the average total revenue.

agricultural income includes three relatively minor components, net revenue from livestock, agricultural wage earnings, and agricultural asset rental income.[2] The income components gaining in importance have been from non-agricultural sources, namely non-agricultural wage earnings + enterprise net revenues (relatively minor) and remittances received from household out-migrants. These general patterns for Nepal largely hold for each of the three regions individually (Figure 11), albeit with more ups and downs.

Broad regional patterns, however, mask a great deal of sub-regional heterogeneity in income sources. Figure 12-14 map the three income sources across the 50 districts of the HRVS. There are several districts in the Hills and Mountains with a non-agricultural income share as high as in some of the Terai districts. Not surprisingly, Hills districts in and around the Kathmandu valley have the highest remittance share of total income anywhere in Nepal. Cutting the data by per capita total expenditure quintile (Figure 15-17), we see a pattern of rising remittance shares, but no consistent pattern in the agricultural and non-agricultural shares across years. Indeed, it is striking how stable the agriculture income share is across quintiles and years.

3. Empirical Methodology

To understand the empirical approach, consider the timeline of monsoon, agricultural seasons, and HRVS survey in Nepal (see Figure 18). As noted the summer monsoon starts in June with the main wet season crop planting following in July. After the retreat of the monsoon in September, wet crops are typically harvested in October. Dry season crops are planted in the final quarter of the calendar year and harvested by around March of the next year. For each of its three waves, the HRVS was fielded between June and early August with annual recall for both agricultural and non-agricultural income.

[2] See Figure 10 for a breakdown of agricultural and nonagricultural income by sub-components

Food consumption expenditures, however, are based on a last 7-days recall, which means that they refer to the current early wet season period. Insofar as households are consuming from their own production at this time, it is likely to be from stocks held over from previous harvests and not from the current (standing) crop. Likewise, households purchasing rice and other local foodstuffs are largely consuming other farmers' stocks held over from the previous harvest. Current wet season rainfall, consequently, should play little role.

Now, let y_{it} be an outcome, such as income, for household *i* in year *t*, R_{it} be the June-Sept. rainfall at that household's geospatial grid point in year *t*, and P_{vdc} be the proportion of area planted to rice in the VDC. We consider regressions of the form

$$logy_{it} = \alpha \log R_{it} + \beta\, logR_{it} \times P_{vdc} + \omega_t + \mu_i + \varepsilon_{it} \tag{1}$$

where ω_t is a year (survey wave) dummy, μ_i is a household fixed effect, and ε_{it} is a random error term. Equation (1) allows the impact of monsoon rainfall on agricultural production to vary according to the local importance of paddy. Note that our specification is more general, in some sense, than one in which the proportion of *household* area devoted to paddy is interacted with the rainfall variable. In particular, equation (1) captures general equilibrium effects within the VDC such as those arising from improved employment prospects on neighboring farms generating higher agricultural wage earnings for the household.

In the case of (food) consumption expenditures c_{it}, we also consider

$$logc_{it} = \gamma\, logy_{it} + \xi_t + \nu_i + u_{it} \tag{2}$$

where equation (2), with household (v_i) and time (ξ_t) fixed effects, is estimated by Two-stage Least Squares (2SLS) using $\log R_{it}$ and $logR_{it} \times P_{vdc}$ as instrumental variables. 2SLS is warranted by (i) measurement errors in income and (ii) simultaneity between consumption and income, such as preference shocks (e.g., an illness in the family) that both reduce household food expenditures and labor supply (and hence income). The parameter γ may be interpreted as the elasticity of welfare with respect to income, a measure of vulnerability. Note, however, that we do *not* partial out community level income shocks in equation (2) by including VDC × year dummies as in Townsend's (1994) test of full insurance (see also Morduch, 1995; Ravallion and Chaudhuri, 1997; Jalan and Ravallion, 1999). Removing covariate risk in this way would be tantamount to throwing the baby out with the bathwater as the monsoon is by and large an aggregate shock.

4. Estimation Results

Our regression analysis of monsoon shocks focuses on households in the Hills and Terai, which account for 90% of the HRVS sample and about the same percentage of agricultural households (across all three years). Dropping the Mountains from our analysis is motivated not only by the relative paucity of paddy production there, but also by its extreme climate, agro-ecology, and economic isolation, making it difficult to compare households in this region to those in the rest of Nepal.[3]

4.1 Household demographics

As a starting point, we consider how household demographic composition responds to the monsoon shock. Interestingly, average household size *declined* by around a quarter of a person over the three waves of the HRVS, mostly due to out-migration, presumably to form new (out-of-sample) households.

Table 1 presents fixed effects regressions analogous to equation (1) of log household size, as

[3] Another issue is that precipitation in mountain areas is notoriously difficult to estimate accurately from satellite data, since altitude and orientation of the household or village on the mountain range is critical.

well as the proportion of adult males (age 16-65) in the household, on $\log R$ and its interaction with the paddy ratio. Columns 1-2 consider the full sample of households (those with non-missing income data for at least two years), columns 3-4 focus on households with nonzero agricultural income data for at least two years, and columns 5-6 on households with nonzero wet season cultivated area for at least two years. All specifications control for year dummies, which net out the general decline in household size mentioned above. Regardless of sample, household size increases on average in "good" monsoons and declines in "bad" monsoons, with no significant change in the proportion of prime-age males. This evidence implies that in good (bad) monsoon years individuals are entering (leaving) our sample households for extended periods from (to) somewhere outside of the sample (such as the Mountains, India, or the Middle East).

That this demographic monsoon response is somewhat greater in households directly engaged in wet season cultivation than it is for households more generally (compare cols. 1 and 5) suggests that it has something to do with agricultural employment opportunities. This conclusion is bolstered by the positive estimated coefficients on the interaction between monsoon rainfall and the proportion of cultivated land in the VDC devoted to paddy (see Figure 20 for the sample distribution of this proportion). Household size increases by more in communities growing more paddy in the wet season. Indeed, where no paddy is grown, there is no significant response of household size to monsoon rainfall. Similarly, the proportion of prime-age males increases during a good monsoon year only in VDCs where paddy is an important crop. In sum, the evidence points toward individuals, especially men, moving in and out of rice-growing areas (and our sample) in response to labor productivity shocks.

4.2 Cultivated area

We next consider how household cultivated area responds to monsoon rainfall. Focusing on the wet season, when virtually all agricultural households are cultivating, we see in col. 1 of **Table 2** that higher

precipitation leads to greater area cultivated. Specifically, a 10% increase in monsoon rainfall leads to a precisely estimated 3.2% increase in wet season area. While in paddy dominated communities there is a greater cultivation response to the monsoon, this effect is not statistically significant.

Recall from the previous subsection that household size increases in good monsoons, although we have no idea from the survey when the influx occurs; i.e., at planting time, at harvest, or even after the harvest of wet season crops is finished. We can ask, however, how much of the area response is associated with having more workers in the household. In col. 2 of Table 2 we control for (log) household size, with only a negligible diminution of the rainfall coefficients. Thus, most of the acreage response to better rains has nothing to do with the (endogenously) increased availability of family labor but, instead, appears to reflect the increased productive potential afforded by greater precipitation.

The third and fourth columns of Table 2 replicate the above regressions using total area cultivated, the sum of wet and dry season acreage. Although local monsoon rainfall might matter in the dry season insofar as cultivable area is dependent on groundwater recharge,[4] we would not expect total area to be as responsive to precipitation as wet season area alone. And, indeed, this is precisely what we find: Total cultivated area increases by only 1.3% (compared to 3.2%) in response to a 10% increase in monsoon rainfall. Once again, controlling for household size barely alters this marginal effect.

4.3 Crop revenue

We continue to trace out the impact of monsoon rainfall by moving on to agricultural production. Calculating net crop revenue as discussed in section 2.2, we report alternative specifications of equation (1) in **Table 3** . Columns 1-3 use log net wet season revenue as the dependent variable, whereas cols. 4-6 use log net total revenue, which includes both wet and dry season. Crop income is higher in good

[4] To the extent that dry season cultivation is dependent on surface irrigation, local monsoon rainfall likely matters less than mountain snowpack.

monsoon years, but mainly in the wet season. As with cultivated area, there is weak evidence that wet season crop revenue responds more positively to rainfall in paddy-intensive VDCs. In the case of total revenue, however, the coefficient on the interaction term $logR_{it} \times P_{vdc}$ is not only positive but very large and significant. A plausible explanation for this effect is as the result of aquifer recharge that allows groundwater irrigation in the dry season. In the paddy-dominated Terai region, about 40% of land area is irrigated by wells as compared to a negligible fraction in the Hills (and Mountains). Thus, in VDCs growing mainly paddy, which are concentrated in the Terai, the impact of improved groundwater recharge on dry season crop revenue is necessarily stronger.

In cols. 2 and 4, we control for area cultivated in, respectively, the wet season and wet + dry season. Essentially the entire (average) monsoon effect on crop income operates through the extensive margin; i.e., it is explained by changes in land area cultivated. By contrast, controlling for log household size in cols. 3 and 6 has virtually no impact on the marginal effect of monsoon rainfall, consistent with the pattern observed in **Table 2**.

4.3 Household income

Turning to the broader aggregates discussed in section 2.3, we now consider the impact of monsoon rainfall on total income from agriculture inclusive of livestock, wage, and rental earnings. For the sake of comparability with previous regressions, we focus here on households that cultivated land in at least two wet seasons. We see in col. 1 of **Table 4** that more plentiful monsoon rainfall increases agricultural income, most of which effect is concentrated in paddy-dominated VDCs. The effect also persists after controlling for household demographics and so, as before, is robust to the positive co-movement of monsoon rainfall and household size. Comparing the marginal effect of monsoon rainfall on total agricultural income of 0.20 (0.08) with the marginal effect on total net crop revenue of 0.08 (0.70) from the previous table, the key difference is the inclusion of agricultural wage income in the former. These

earnings, especially important in the Terai (see Figure 10), are likely to be particularly sensitive to agricultural productivity shocks driven by weather (see, e.g., Jayachandran 2006 for evidence from India).

We can interpret the elasticity of income with respect to monsoon rainfall of 0.20 using the standard deviation of log monsoon rainfall over the 1981-2018 period (in the Terai and Hills combined) of 0.19. A positive shock equivalent to a one standard deviation increase in log monsoon rainfall will increase agricultural incomes in the *average* VDC by 0.20 x 0.19 x 100 = 3.8%. However, at the 75th percentile of the VDC paddy ratio this figure rises to 7.0%, and at the 95th percentile to 8.9%; meanwhile, at the 25th percentile of the VDC paddy ratio, the agricultural income response is effectively zero.

Columns 3-4 of **Table 4** repeat the two regressions in cols. 1-2 using total income inclusive of all non-agricultural earnings and remittances. While the overall marginal rainfall effect remains at 0.20, the coefficient on the rainfall/paddy share interaction increases by 50%, which means that the total incomes of households in paddy dominant areas are even more sensitive to monsoon rainfall than are their agricultural incomes. Specifically, a one standard deviation increase in log monsoon rainfall will increase total incomes at the 25th percentile of the VDC paddy ratio by, again, essentially zero, at the 75th percentile by 8.7%, and at the 95th percentile by 11.5%. Note also that the household size effect, previously positive and significant, goes to zero in col. 4. The reason for this is mechanical: An increase in household size, all else equal, coincides with the return of a remittance-sender, who, even if engaged in local non-agricultural employment, is likely earning much less than as a migrant. Thus, the correlation between remittance income and household size, conditional on the household fixed effect, is negative.

To unpack what is going on with total income, we include analogous regressions in cols. 5-6 with the dependent variable being an indicator that takes on a value of one if anyone in the household is employed outside of agriculture. Here we see the reason for the large increase in the coefficient on the

interaction term $logR_{it} \times P_{vdc}$ between col. 1 to 3: The corresponding interaction coefficient in the nonagricultural employment regression is similarly positive and highly significant. In other words, a good monsoon greatly enhances off-farm employment opportunities in paddy-intensive areas, but not in areas where paddy is of medium or low importance. This employment effect is also not explained by changes in household size (col. 6) induced by monsoon rainfall. Thus, there appear to be significant sectoral spillovers in the rural economy whereby temporarily higher (lower) agricultural productivity induces higher (lower) non-agricultural productivity. As a result, the monsoon shock in paddy dominant areas increases *both* farm and non-farm incomes at the same time.

4.4 Food consumption

Our final step is to consider the implications of monsoon-induced income fluctuations for household welfare, again restricting the sample to households that cultivate in at least two wet seasons. We first present "reduced form" estimates of equation (1) with per capita food expenditures as the dependent variable. Columns 1-2 of **Table 5** report these estimates, with and without controlling for household size. In either case, we find that higher monsoon rainfall leads to significantly higher total food consumption in more rice intensive areas, but the average marginal effect of rainfall is only weakly positive (col. 1) and even nonexistent once household size is controlled for (col. 2). These results are largely consistent with the patterns observed for income in **Table 4**. Thus, the col. 1 estimates imply that a one standard deviation increase in log monsoon rainfall will increase total food consumption at the 25th percentile of the VDC paddy ratio by roughly zero, at the 75th percentile by 2.3%, and at the 95th percentile by 3.3%.

The remainder of ***Table 5*** reports estimates of equation (2) with two alternative definitions of income, agricultural income (cols. 3-4) and total income (cols. 5-6). The 2SLS estimates of the vulnerability parameter γ are always positive and significant, albeit substantially attenuated once log

household size is included (cols. 4 and 6).[5] Moreover, the Kleibergen-Paap weak instruments statistics are all well above the rule-of-thumb critical value of 10, whereas the over-identification tests all fail to reject. Thus, the instruments appear to be adequate and there is no evidence that monsoon rainfall affects consumption other than through income.

In comparing the food consumption elasticity with respect to agricultural income of $\gamma_A = 0.36$ with the corresponding elasticity with respect to total income of $\gamma_T = 0.27$ note that, since agricultural income constitutes between 25-30% of total income (depending on the year; see Figure 9), we should expect $\frac{\gamma_A}{\gamma_T}$ to be in the range of 0.25 to 0.3.[6] In fact, the ratio of coefficients is greater than one. At least part of the explanation for the discrepancy may be that the instruments perform better for total income than for agricultural income (compare the weak identification test statistic between cols. 3 and 5), and hence γ_T is less biased toward OLS than γ_A. At any rate, based on γ_T, a 10% fall in income (regardless of the source of this decline) would translate into a 2.7% fall in food consumption.[7] This is likely an understatement of the elasticity of consumption with respect to *contemporaneous* income given that expenditures here are measured roughly nine months after the wet season harvest and at least 3 months after the dry season harvest, allowing households a considerable adjustment period.[8]

[5] The OLS estimates of γ (not reported) are much smaller in magnitude suggesting that measurement error may be the more serious endogeneity concern than simultaneity of income and consumption.

[6] That is, by the mathematical identity $\frac{\partial \log c}{\partial \log y_A} = \frac{y_A}{y_T}\frac{\partial \log c}{\partial \log y_T}$, where the subscripts A refers to agricultural income and T to total income.

[7] We prefer the col. 5 specification to the one in col. 6 that includes log household size because of our evidence in Table 1 showing that household size responds endogenously to monsoon rainfall.

[8] Jalan and Ravallion (1999) obtain a comparable consumption-income elasticity of about 0.20 in rural China using panel data from the second half of the 1980's, at which time it was at roughly the same level of development as Nepal. Their analysis uses both annual consumption and income data.

5. Conclusions and policy implications

We have shown that households in rural Nepal are not well-insured against the monsoon shock. This is what we would expect in the case of such a large covariate risk, and one which also appears to have general equilibrium repercussions in local labor markets (as suggested by the off-farm employment response). Our evidence further shows that the brunt of a bad monsoon is borne by households in paddy-dominant communities, most of which are in the Terai. While rainfall has an obvious link to rice productivity, given paddy's water-intensiveness, there is a dry-season effect as well, most likely operating through aquifer recharge. This latter effect is limited to areas of significant groundwater exploitation, which in Nepal also happen to be in the Terai.

Social protection policy in Nepal should be attentive to this significant uninsured risk and its geographic incidence. India's approach of using a public works employment guarantee (100 days per year under the NREGA program) as a hedge against drought is one possible policy response. In the case of Nepal, public works schemes should be targeted to the Terai. Alternatively, or as a complementary policy, the government could subsidize the purchase of rainfall index insurance to farmers, or, perhaps more equitably, purchase such insurance itself on a large scale, distributing the payout as an unconditional means-tested cash transfer during bad years. Once again, the results of this study suggest that such efforts should be targeted to the rice-intensive lowlands of Nepal.

References

Jacoby, H. G., & Skoufias, E. (1997). Risk, financial markets, and human capital in a developing country. *The Review of Economic Studies*, *64*(3), 311-335.

Jacoby, H. G., & Skoufias, E. (1998). Testing theories of consumption behavior using information on aggregate shocks: Income seasonality and rainfall in rural India. *American Journal of Agricultural Economics*, *80*(1), 1-14.

Jalan, J., & Ravallion, M. (1999). Are the poor less well insured? Evidence on vulnerability to income risk in rural China. *Journal of development economics*, *58*(1), 61-81.

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

Kochar, A. (1999). Smoothing consumption by smoothing income: hours-of-work responses to idiosyncratic agricultural shocks in rural India. *Review of Economics and Statistics*, *81*(1), 50-61.

Morduch, J. (1995). Income smoothing and consumption smoothing. *Journal of economic perspectives*, *9*(3), 103-114.

Rose, E. (1999). Consumption smoothing and excess female mortality in rural India. *Review of Economics and statistics*, *81*(1), 41-49.

Shah, M., & 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.

Tiwari, S., Jacoby, H. G., & Skoufias, E. (2017). Monsoon Babies: Rainfall Shocks and Child Nutrition in Nepal. *Economic Development and Cultural Change*, *65*(2), 167-188.

Townsend, R. M. (1994). Risk and insurance in village India. *Econometrica*, *62*, 539-539.

Figure 1. Standard deviation of logarithm of monsoon rainfall 1981-2018

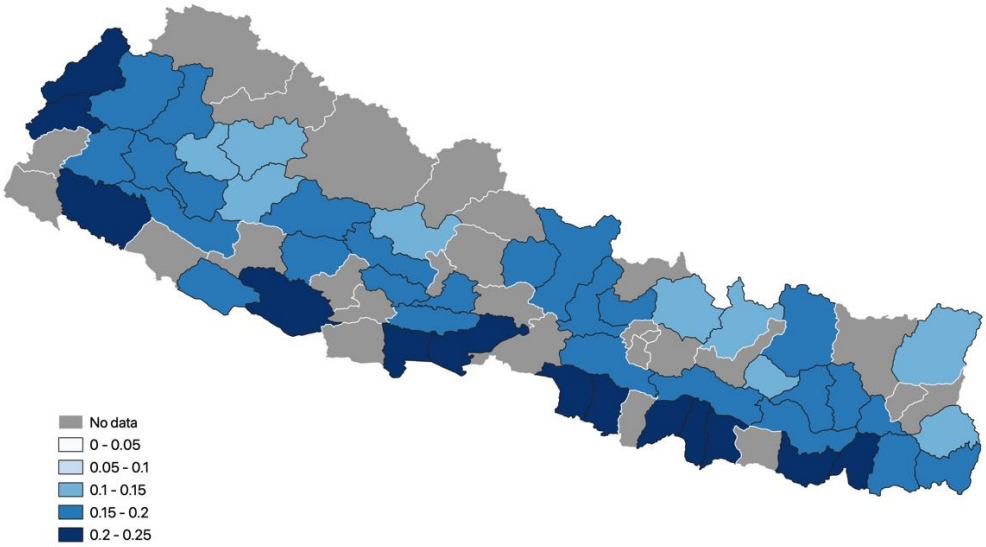

Notes: Logarithm the level of June-Sept. rainfall for each household in each year and average all household-level standard deviations of logarithm rainfall across years in district.

Figure 2. 2015 monsoon rainfall (mm)

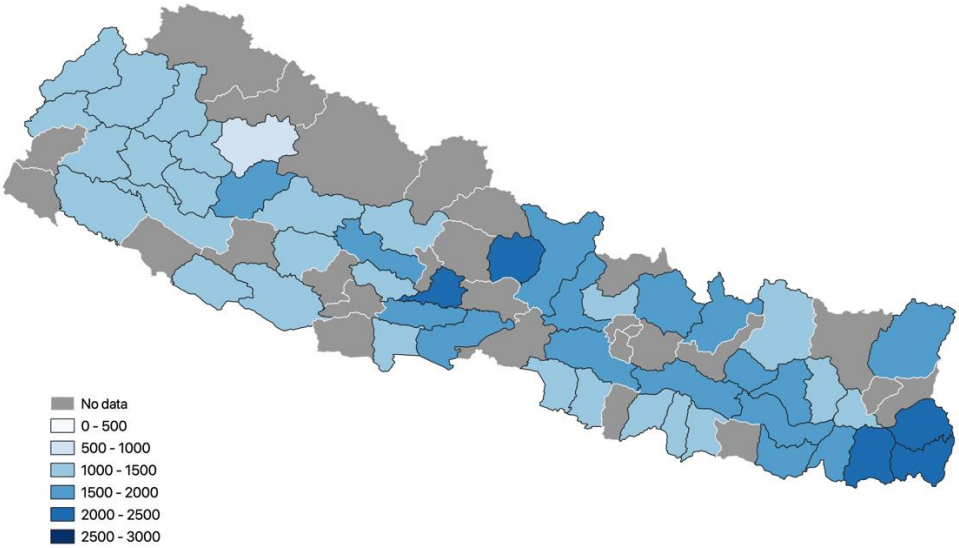

Notes: Average 2015 June-Sept. rainfall by district.

Figure 3. 2016 monsoon rainfall (mm)

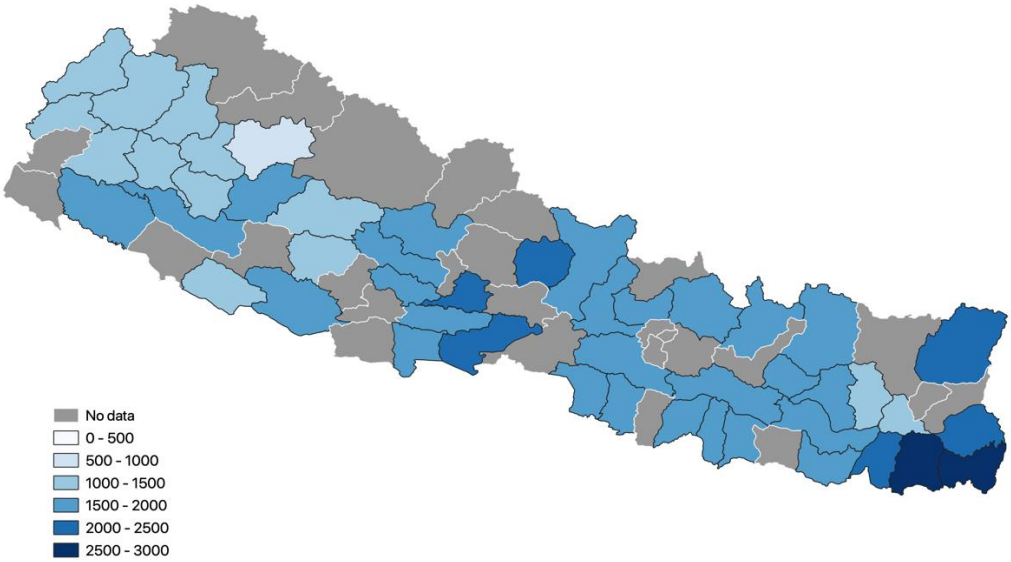

Notes: Average 2016 June-Sept. rainfall by district.

Figure 4. 2017 monsoon rainfall (mm)

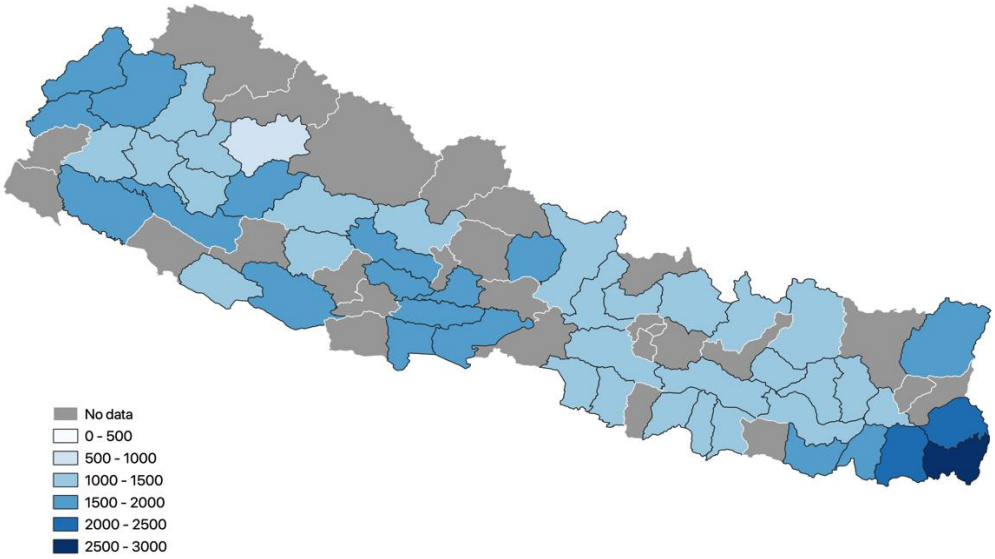

Notes: Average 2017 June-Sept. rainfall by district.

Figure 5. Seasonal crop composition by geography

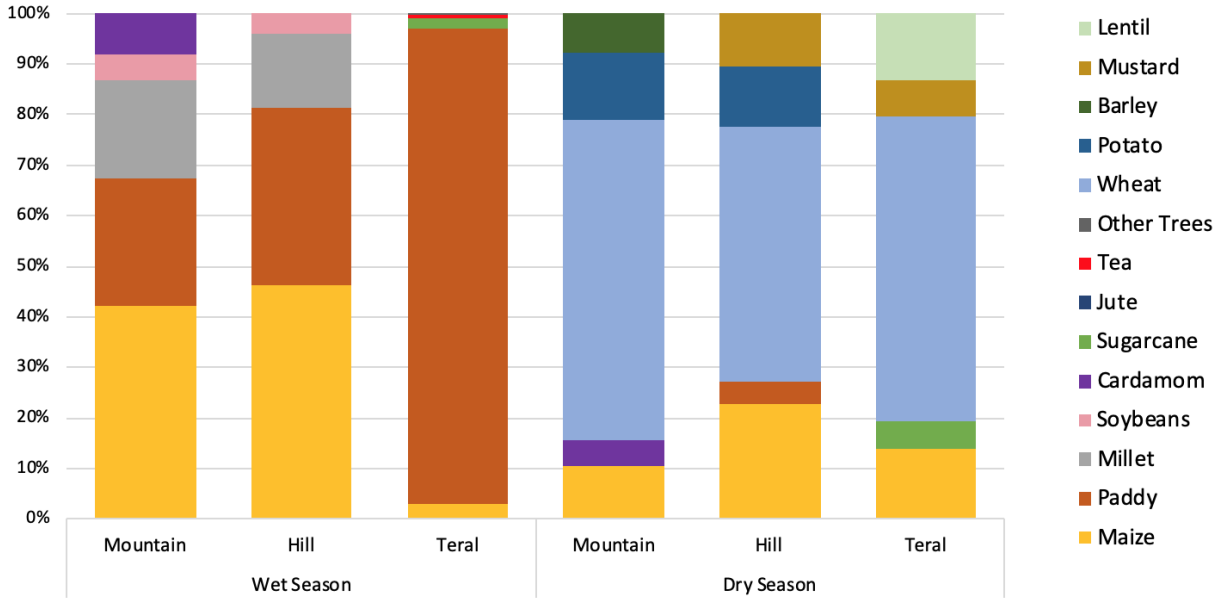

Notes: Top five crops in each geography according to total area planted across all households.

Figure 6. Importance of paddy area in wet season (% of area)

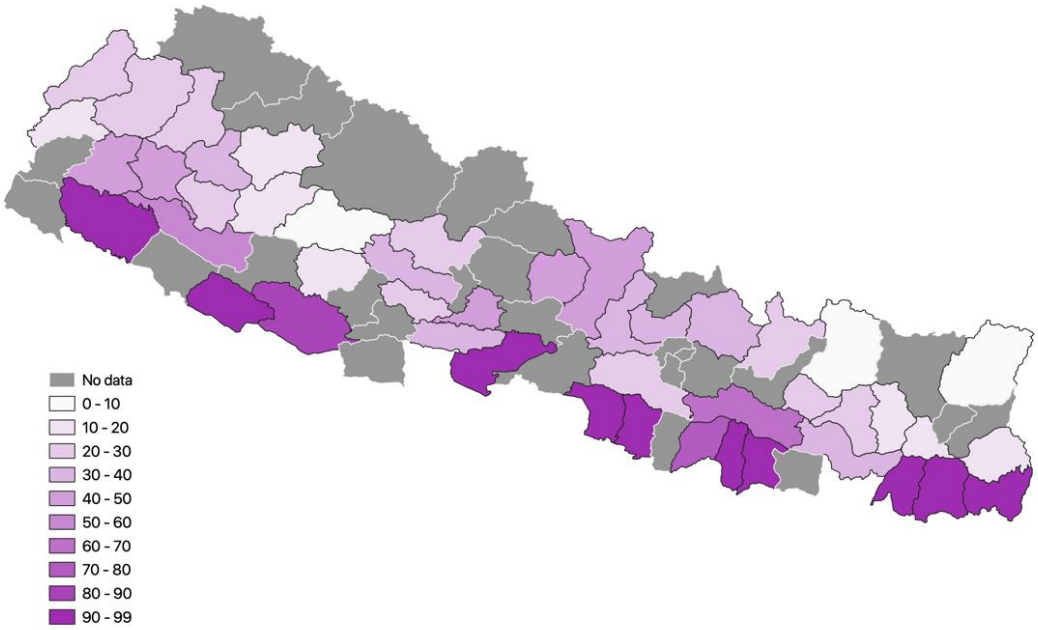

Notes: Percentage of total area cultivated paddy in wet season. Total paddy area in wet season across all households and years as a proportion of total wet season cropped area by district.

Figure 7. Proportions of wet crop net income and area by year and geography

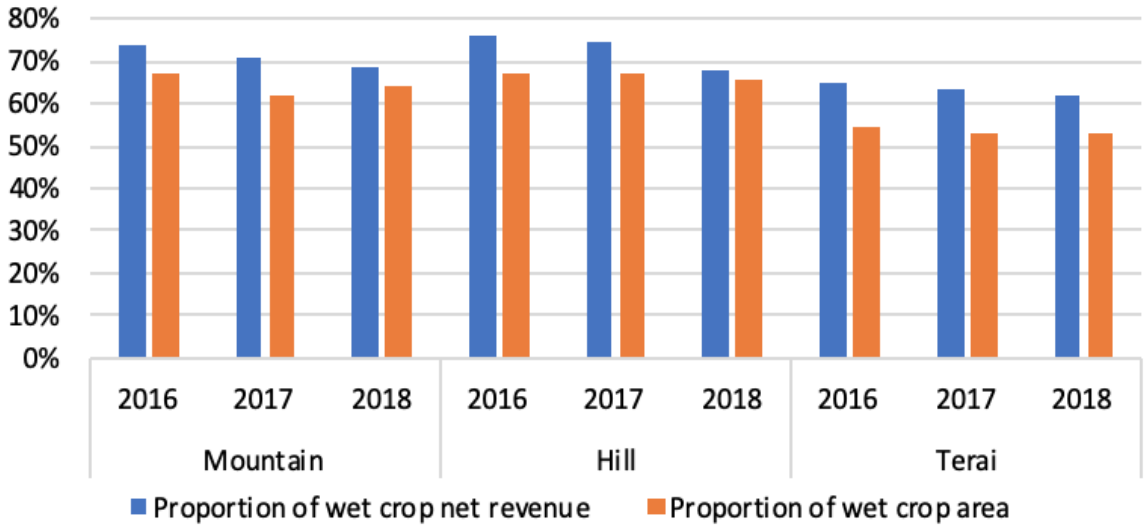

Notes: Proportion of wet crop net revenue is the proportion of annual net crop revenue earned in wet season. Proportion of wet crop area is the proportion of annual total crop area cultivated in wet season. Average of all household-level proportions of wet crop net revenue and cultivated area in year and geography.

Figure 8. Proportion of crop net revenue in wet season (percentage, %)

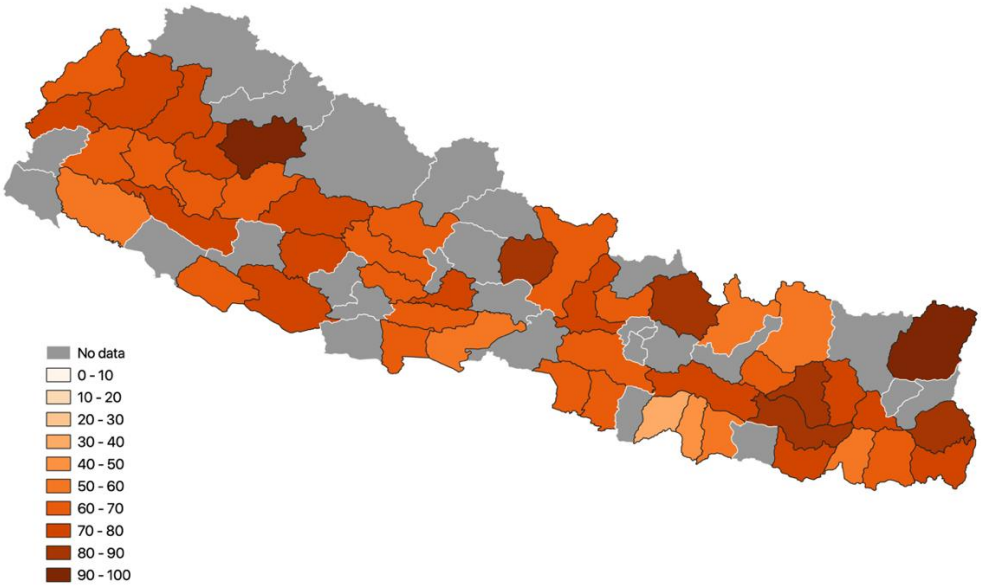

Notes: Proportion of annual net crop revenue earned in wet season by district.

Figure 9. Income shares by year

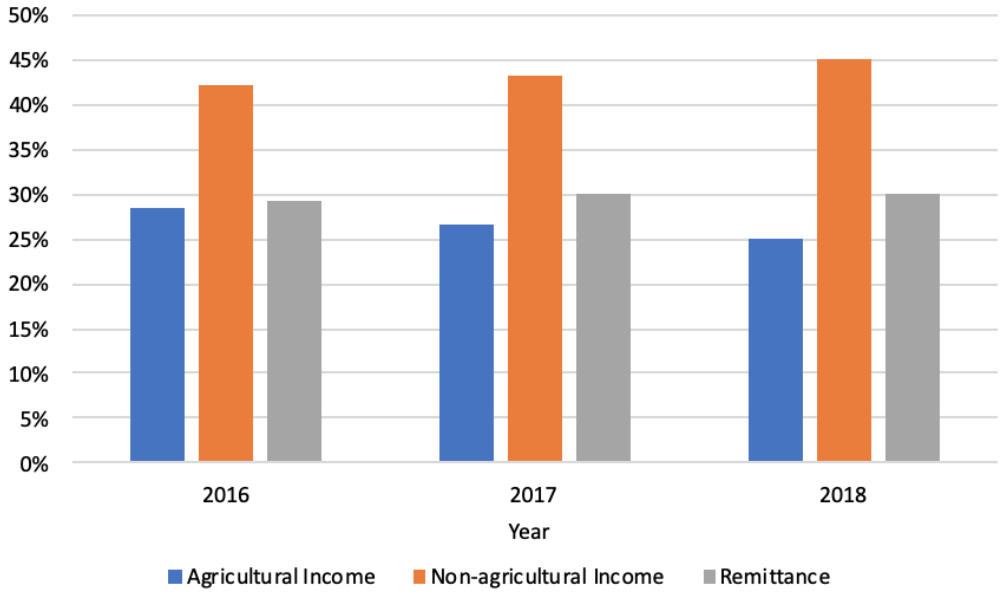

Notes: Agricultural income is the sum of crop net revenue, livestock net revenue, rental income, and agricultural wage income. Non-agricultural income is the sum of business net income and non-agricultural wage Income. Average of all household-level income shares by year.

Figure 10. Income sub-components by geography

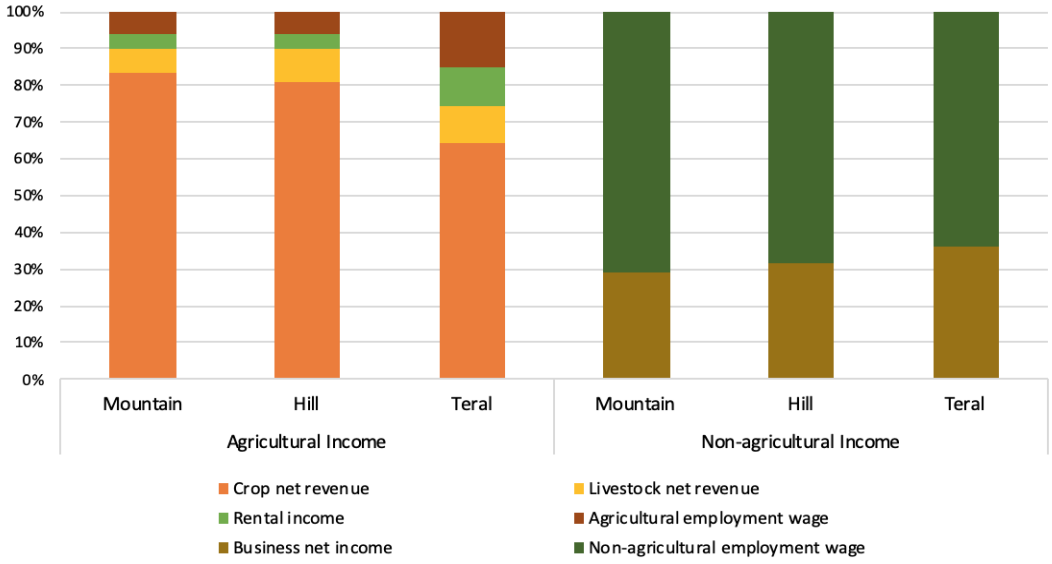

Notes: Sum of all household-level observations across years by geography for each income source divided by total across all households.

Figure 11. Income shares by year and geography

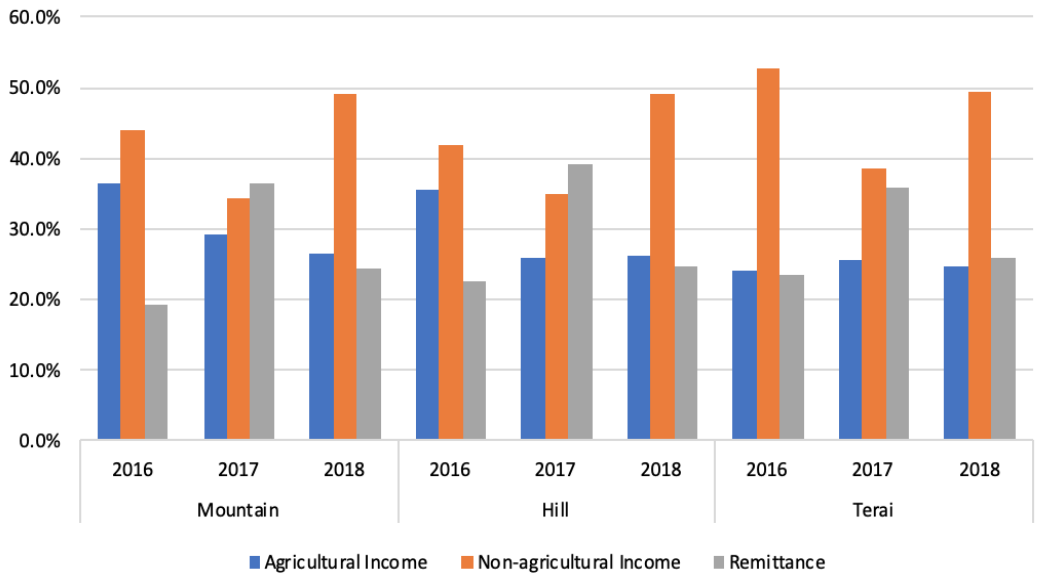

Notes: Agricultural income is the sum of crop net revenue, livestock net revenue, rental income, and agricultural wage income. Non-agricultural income is the sum of business net income and non-agricultural wage income. Average of all household-level income shares by year for each geography.

Figure 12. Agricultural income share (%)

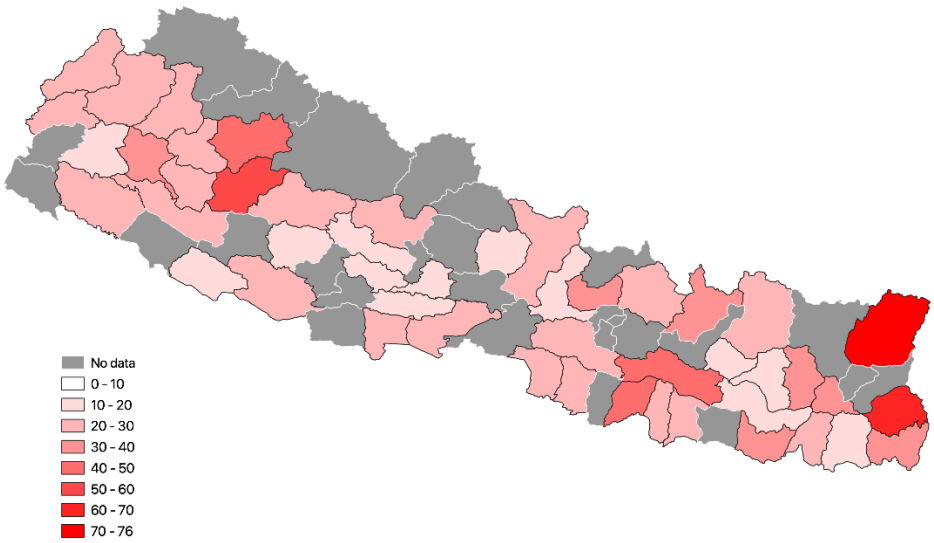

Notes: The proportion of agricultural income in total income by district. Agricultural income is the sum of crop net revenue, livestock net revenue, rental income, and agricultural wage income.

Figure 13. Non-agricultural income share (%)

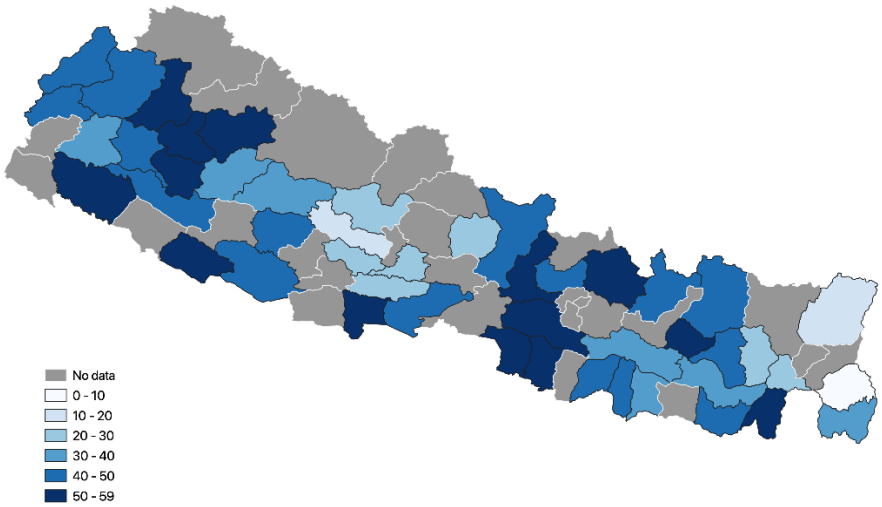

Notes: Proportion of non-agricultural income in total income by district. Non-agricultural income is the sum of business net income and non-agricultural wage income.

Figure 14. Remittance income share (%)

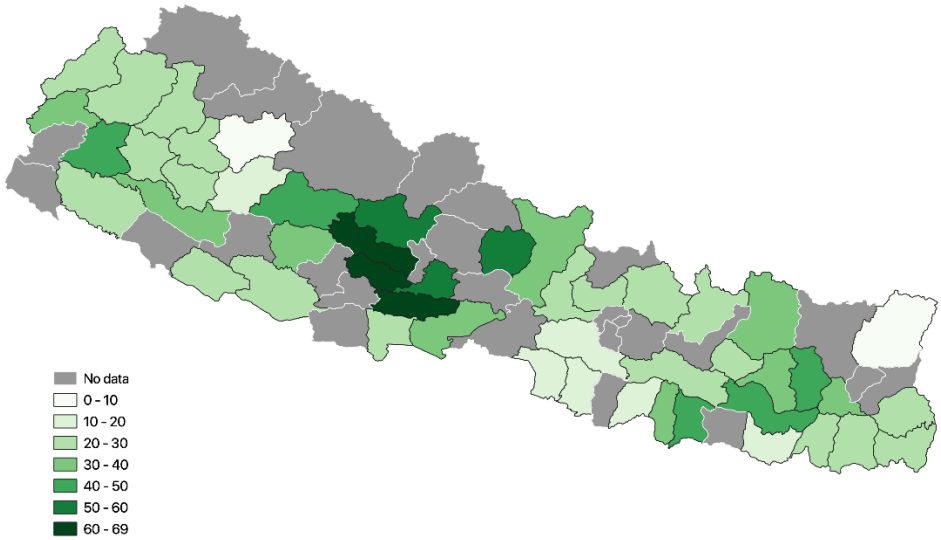

Notes: The proportion of remittance in total income by district.

Figure 15. Income shares by quintiles in 2016

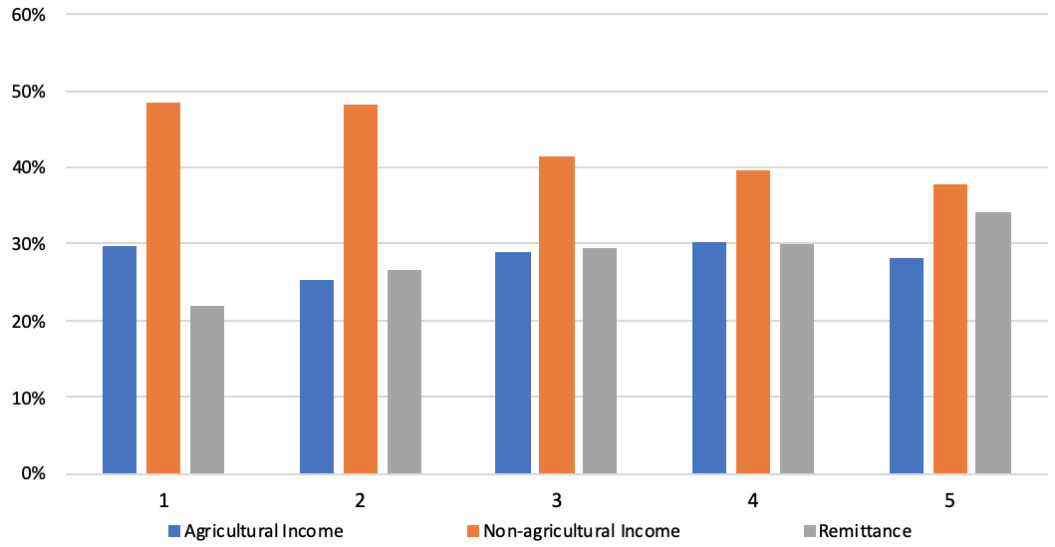

Notes: Agricultural income is the sum of crop net revenue, livestock net revenue, rental income, and agricultural wage income. Non-agricultural income is the sum of business net income and non-agricultural wage Income.

Figure 16. Income shares by quintiles in 2017

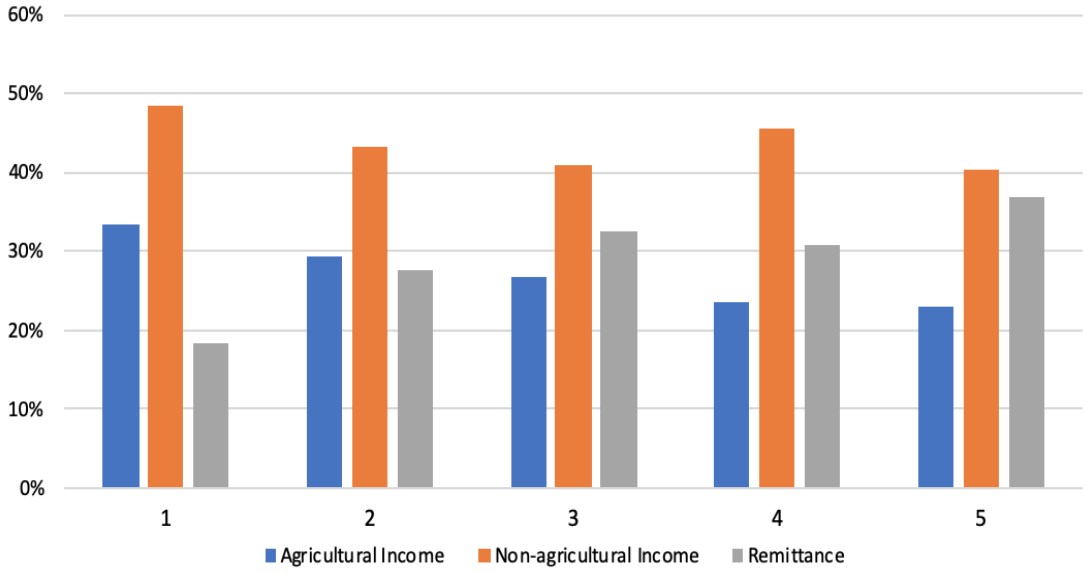

Notes: Agricultural income is the sum of crop net revenue, livestock net revenue, rental income, and agricultural wage income. Non-agricultural income is the sum of business net income and non-agricultural wage Income.

Figure 17. Income shares by quintiles in 2018

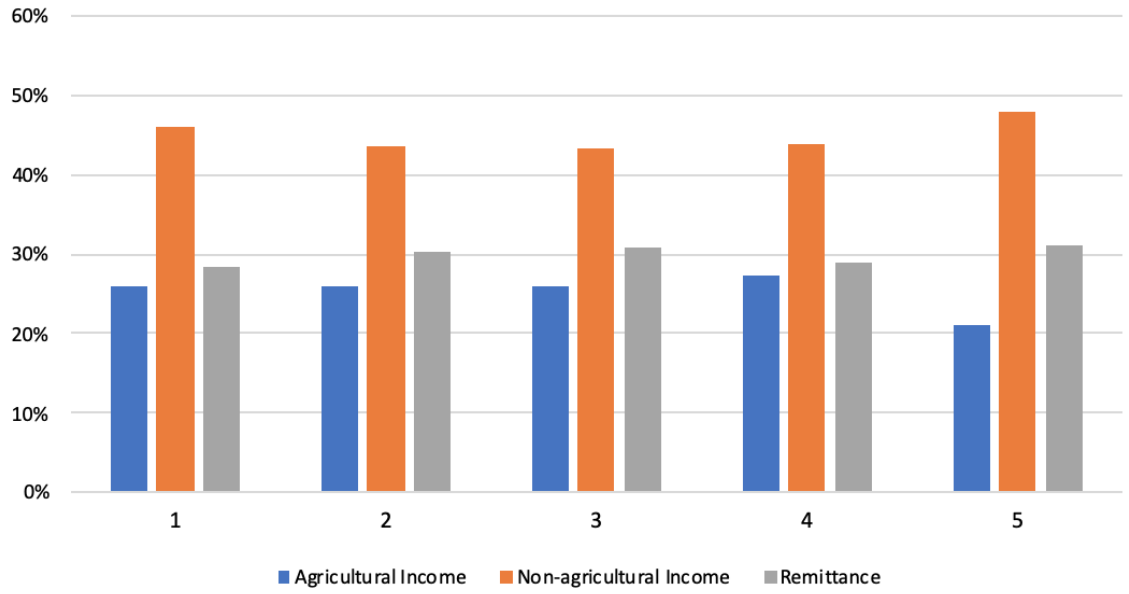

Notes: Agricultural income is the sum of crop net revenue, livestock net revenue, rental income, and agricultural wage income. Non-agricultural income is the sum of business net income and non-agricultural wage Income.

Figure 18. Annual timeline of events

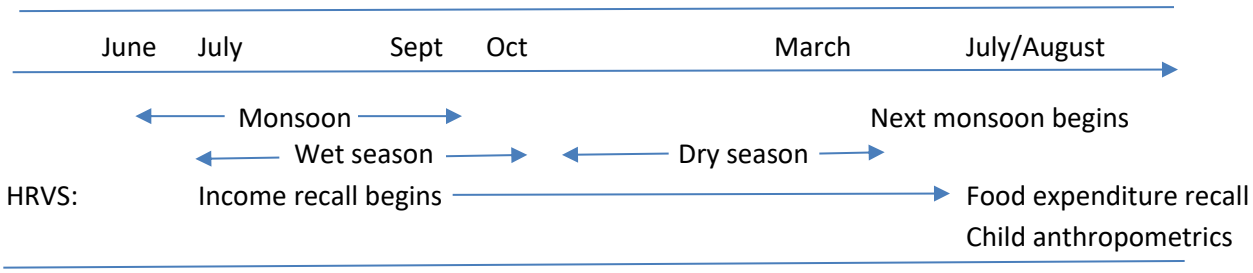

Figure 19. Household size by year and geography

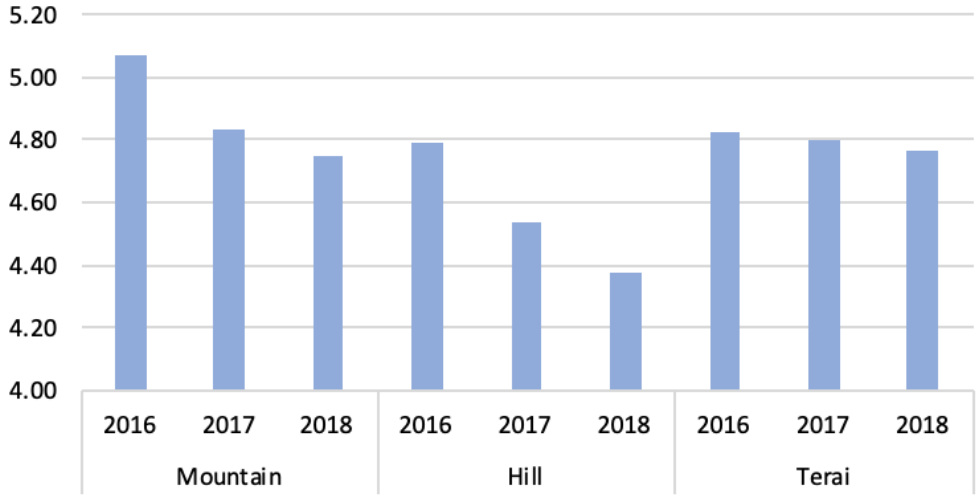

Notes: Average household sizes by year for each geographical region.

Figure 20. Importance of paddy at VDC level

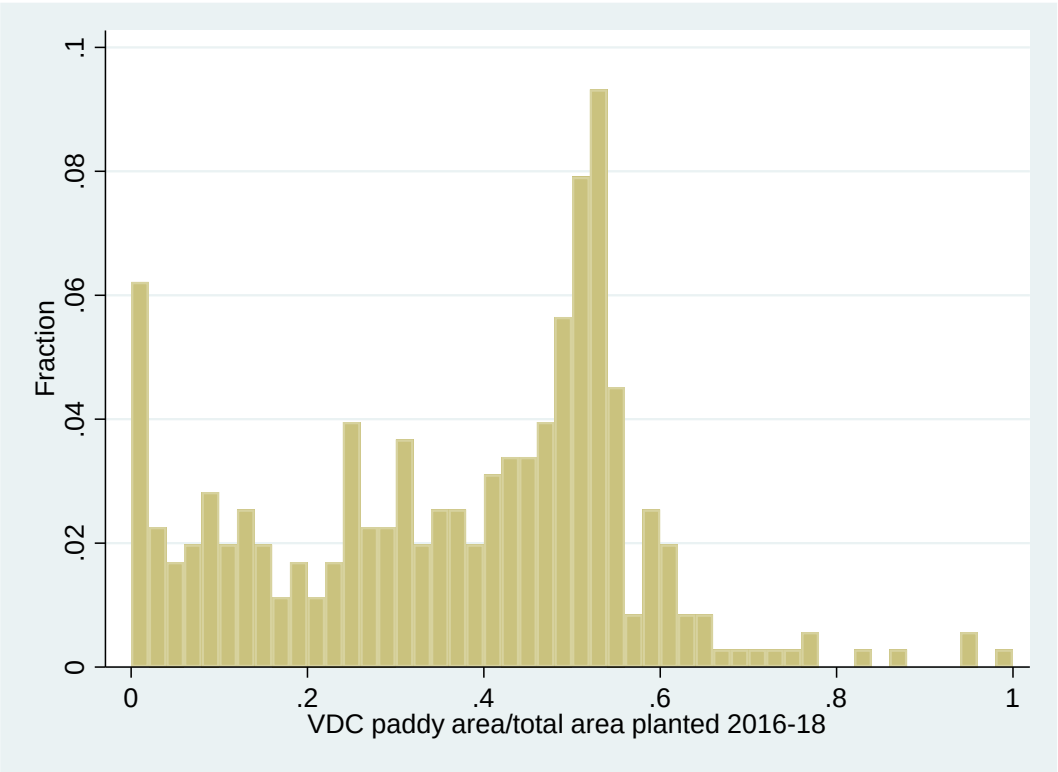

Notes: Hills and Terai only (Mountains excluded).

Table 1. Monsoon Shocks and Household Demographics

	All HHs		HHs with agricultural Income in 2+ years		HHs with cultivated area in 2+ wet season	
	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES	Log hhsize	Males 16-65/hhsize	Log hhsize	Males 16-65/hhsize	Log hhsize	Males 16-65/hhsize
$\log R$	-0.013	-0.042***	-0.0093	-0.039***	-0.0032	-0.040**
	(0.035)	(0.015)	(0.036)	(0.015)	(0.036)	(0.016)
$\log R \times P_{vdc}$	0.25***	0.086**	0.25***	0.085**	0.27***	0.086**
	(0.074)	(0.034)	(0.077)	(0.035)	(0.079)	(0.036)
$\log R$ (marginal effect)	0.077***	-0.011	0.079***	-0.0090	0.088***	-0.010
	(0.018)	(0.0080)	(0.019)	(0.0083)	(0.019)	(0.0087)
HH fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	15,738	15,738	14,387	14,387	13,028	13,028
R-squared	0.893	0.851	0.891	0.850	0.888	0.847
Number of clusters	1062	1062	1062	1062	1059	1059

Notes: Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1 adjusted for two-way clustering on household and vdc-year. *R* represents total precipitation from June to September, P_{vdc} is the proportion of cultivated land in VDC devoted to paddy from 2016 to 2018.

Table 2. Monsoon Shocks and Cultivated Areas for Cultivating Households

	Cultivated area in wet season		Cultivated area in wet and dry seasons	
VARIABLES	(1)	(2)	(3)	(4)
$\log R$	0.24*	0.24*	0.029	0.029
	(0.13)	(0.13)	(0.11)	(0.11)
$\log R \times P_{vdc}$	0.24	0.20	0.29	0.26
	(0.27)	(0.27)	(0.26)	(0.26)
$\log R$ (marginal effect)	0.32***	0.30***	0.13**	0.12**
	(0.058)	(0.058)	(0.053)	(0.053)
Log (household size)	---	0.15***	---	0.12***
		(0.037)		(0.035)
HH fixed effects	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes
Observations	12,546	12,546	12,577	12,577
R-squared	0.824	0.825	0.830	0.830
Number of clusters	1058	1058	1058	1058

Notes: Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1 adjusted for two-way clustering on household and vdc-year. *R* represents total precipitation from June to September, P_{vdc} is the proportion of cultivated land in VDC devoted to paddy from 2016 to 2018.

Table 3. Monsoon Shocks and Crop Net Revenue for Cultivating Households

	Crop net revenue in wet season			Crop net revenue in wet and dry seasons		
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
$\log R$	-0.0065	-0.21	-0.0061	-0.28*	-0.31**	-0.28*
	(0.19)	(0.17)	(0.19)	(0.15)	(0.13)	(0.15)
$\log R \times P_{vdc}$	0.68	0.51	0.63	1.06***	0.86***	1.02***
	(0.43)	(0.38)	(0.43)	(0.35)	(0.30)	(0.35)
$\log R$ (marginal effect)	0.22**	-0.042	0.21**	0.080	-0.015	0.067
	(0.089)	(0.079)	(0.088)	(0.070)	(0.064)	(0.071)
Log (area cult. wet)		0.83***				
		(0.026)				
Log (household size)	---	---	0.20***	---	---	0.14***
			(0.053)			(0.041)
Log (area cult. total)	---	---	---	---	0.75***	---
					(0.020)	
HH fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	12,452	12,450	12,452	12,589	12,589	12,589
R-squared	0.712	0.785	0.712	0.784	0.850	0.784
Number of clusters	1058	1058	1058	1058	1058	1058

Notes: Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1 adjusted for two-way clustering on household and vdc-year. *R* represents total precipitation from June to September, P_{vdc} is the proportion of cultivated land in VDC devoted to paddy from 2016 to 2018.

Table 4. Monsoon Shocks and Households Income Sources

	Total agricultural income		Total income including remittances		Any nonagricultural employment (0/1)	
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
$\log R$	-0.14	-0.14	-0.33*	-0.33*	-0.22***	-0.22***
	(0.17)	(0.17)	(0.19)	(0.19)	(0.070)	(0.069)
$\log R \times P_{vdc}$	1.00**	0.94**	1.54***	1.55***	0.60***	0.52***
	(0.40)	(0.40)	(0.44)	(0.44)	(0.16)	(0.16)
$\log R$ (marginal effect)	0.20**	0.18**	0.20**	0.20**	-0.016	-0.041
	(0.080)	(0.080)	(0.095)	(0.095)	(0.038)	(0.037)
Log (household size)		0.24***		-0.0089		0.28***
		(0.049)		(0.072)		(0.024)
HH fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	12,853	12,853	13,002	13,002	13,028	13,028
R-squared	0.709	0.710	0.637	0.637	0.565	0.574
Number of clusters	1058	1058	1059	1059	1059	1059

Notes: Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1 adjusted for two-way clustering on household and vdc-year. *R* represents total precipitation from June to September, P_{vdc} is the proportion of cultivated land in VDC devoted to paddy from 2016 to 2018.

Table 5. Monsoon Shocks and Food Expenditures

	Reduced form OLS		2-stage least squares			
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
$\log R$	-0.12	-0.12	---	---	---	---
	(0.10)	(0.10)				
$\log R \times P_{vdc}$	0.49**	0.37*	---	---	---	---
	(0.22)	(0.22)				
$\log R$ (marginal effect)	0.041	0.0041	---	---	---	---
	(0.048)	(0.048)				
Log (household size)		0.44***	---	0.38***	---	0.44***
		(0.028)		(0.038)		(0.029)
Log (agric Income)	---	---	0.36***	0.23**	---	---
			(0.11)	(0.10)		
Log (total Income)	---	---			0.27***	0.17**
					(0.079)	(0.069)
HH fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	12,852	12,852	12,852	12,852	12,852	12,852
R-squared	0.615	0.634	0.447	0.447	0.450	0.570
Number of clusters	1058	1058	4415	4415	4415	4415
Under-id. p-value			0.000	0.000	0.000	0.000
Weak id statistic			21.2	17.9	28.0	27.9
Over-id p-value			0.30	0.15	0.69	0.30

Notes: Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1 adjusted for two-way clustering on household and vdc-year in cols 1-2, adjusted for clustering on household in cols. 3-6. *R* represents total precipitation from June to September, P_{vdc} is the proportion of cultivated land in VDC devoted to paddy from 2016 to 2018. Instruments in cols. 3-6 are $\log R$ and $\log R \times P_{vdc}$.