Eliminating Uncertainty in Market Access: The Impact of New Bridges in Rural Nicaragua*

Wyatt Brooks
University of Notre Dame

Kevin Donovan
University of Notre Dame

March 2017

Preliminary and Incomplete Draft – Comments Welcome

Abstract

We estimate the impact of new infrastructure in rural Nicaraguan villages facing seasonal floods that unpredictably eliminate access to outside markets. We build footbridges designed to eliminate this risk. Identification exploits small engineering requirements that preclude construction in some villages, despite their need for a bridge. We collect detailed annual household surveys over three years and weekly telephone followups with a subset of households for sixty-four weeks, both before and after construction. Bridges eliminate uncertainty in market access driven by floods: during flood episodes in control villages labor market earnings decrease by 30 percent, while there is no change in treatment villages. This translates into substantial reallocation of activities between farming and wage work, increased fertilizer spending and yields on farms, and lower savings. In a model of occupational choice and risky farm investment, we show that these results are a rationalize response to lower income risk induced by the bridges. In particular, the bridge decreases distortions in the agricultural sector through its change in access to labor markets.

*This paper has benefited from comments by Nilesh Fernando, Manuel García-Santana, Lakshmi Iyer, Joe Kaboski, Tim Lee, Jack Rossbach and participants at the Chicago Fed Development Workshop, the Notre Dame Development Lunch, and Michigan State Development Day. Thanks also to Avery Bang, Maria Gibbs, and Brandon Johnson for help coordinating the project, and to Leo Castro, Katie Lovvorn, and Justin Matthews for managing the project in Nicaragua. Giuliana Carozza, Jianyu Lu, Andrea Ringer, and Melanie Wallskog provided outstanding research assistance. We thank the Helen Kellogg Institute for International Studies for financial support. *Contact Info:* 434 Flanner Hall, Notre Dame, IN 46556. Brooks: wbrooks@nd.edu; Donovan: kdonova6@nd.edu

1 Introduction

A large fraction of households in the developing world live in rural areas that are less productive than urban areas in the same countries (Restuccia, Yang and Zhu, 2008; Gollin, Lagakos and Waugh, 2014). While a growing literature has pointed out large gains from structural reallocation of factors from rural to urban areas, recent work suggests important constraints within rural economies as well.[1] Much of this work, however, utilizes what Restuccia and Rogerson (2016) refer to as the *indirect approach*, backing out inefficient resource allocations from model assumptions. Direct evidence on how these distortions are generated is critical for both understanding rural poverty and designing policies to alleviate it.

In this paper we take a step in that direction by directly considering the impact of new infrastructure in rural Nicaragua. Every year between May and October, flooding occurs unpredictably. These floods routinely eliminate access to outside food, product, and labor markets, in large part driven by the lack of suitable infrastructure to traverse flooded rivers.[2] We build footbridges in villages that face this seasonal flood risk and assess their impact. We conduct household-level surveys over three years before and after bridge construction, along with 64 weeks of phone surveys with a subset of households. This allows us to focus on multiple margins affected by the bridge, critical given the heterogeneity of income-generating activities in rural areas (Foster and Rosenzweig, 2007; World Bank, 2008*b*).

Our identification strategy is based on the fact that there are many villages that need bridges, but some cannot be completed due to engineering requirements. These requirements are small from the perspective of households in the village, but critical for safely constructing a footbridge. We discuss this further in Section 2 and show that these features are orthogonal to any relevant household or village characteristics. A major barrier to studying transportation infrastructure as an intervention is the high construction cost of these interventions.[3] This is true in our context as well, where

[1] See, for instance, work by Restuccia and Santaeulalia-Llopis (2017) on land misallocation and Adamopoulos et al. (2017) on inefficient selection across sectors.

[2] This issue is considered a major rural development hurdle by both international policy organizations and citizens of Nicarauga (World Bank, 2008*a*). More broadly, tropical geography and its seasonal flooding and monsoons have long been discussed as a contributor to poverty. See Kamarck (1973) for an early study on agriculture and health issues in the tropics.

[3] This generally implies a difficult identification issue, as these expensive projects are generally targeted toward

bridge costs approximately $40,000. As such, our study includes 500 households from 15 villages surveyed over three years. Since we have a small number of clusters, we use the wild bootstrap cluster-t procedure from Cameron, Gelbach and Miller (2008) throughout.

Despite the small number of clusters, we find economically and statistically significant effects. We first show using high frequency data that the bridge eliminates market access risk during flood episodes. During a flood, average labor market earnings decrease by 15 percent in the absence of a bridge. In villages with a bridge, there is no change in average labor market earnings during a flood. Similar results hold when we consider the likelihood of earning no income in a given period. We confirm these results in the annual surveys, as households earn nearly 30 percent more from wage work in response to a bridge. The result is entirely driven by changes in days worked outside the village, while wages remain the same. These results show that (1) floods generate uncertain access to labor markets and (2) bridges eliminate this uncertainty and thus increase earnings in the labor market. This labor market access has important effects that spill over into the agricultural sector. First, we find that farmers spend nearly 50 percent more on intermediate inputs like fertilizer and pesticide in response to a bridge. Moreover, yields increase by 60 percent on staple crops.

One explanation for these results is selection. That is, a bridge increases the attractiveness of the labor market for all households, and some households shift away from farming. Those who remain are high productivity, and therefore average investment and yield increase. We test this result and find that there are indeed substantial changes in (broad) occupations in response to a bridge. However, these flows occur in both directions: farming households are significantly more likely to exit farming and begin wage work after a bridge, but working households are also significantly more likely to begin farming.

In light of the empirical results, we develop a theoretical model to shed light on the relationship between these results. In particular, we show that these results imply

areas with the largest impact. We sidestep this by using engineering-related requirements, similar in style to Dinkelman (2011). Moreover, our close involvement with the data collection and construction allows us to collect detailed data along a number of dimensions before and after construction.

that the bridge significantly lowers agricultural distortions through its affect on the labor market. The model starts from the assumption that bridges allow better access to labor markets, as implied by the empirical results. We then build a model in which this labor market access also implies benefits in the agricultural sector. The link from from the labor market to agriculture is predicated on the idea that there is substantial risk – unaffected by the bridge – that negatively affects agriculture.[4] In particular, farmers are required to choose fertilizer investment before the realization of the shock. The absence of insurance implies that larger *ex ante* investment generates lower *ex post* consumption in the event of a low shock realization. Farmers internalize this fact and limit their exposure to *ex post* consumption risk by limiting *ex ante* investment.

We show that this model is consistent with our findings for wage earnings, agricultural outcomes, and occupational churn. The first is assumed. The second comes from the bridge's ability to limit the impact of bad farm shocks. In particular, farmers can adjust labor in response to bad farm shocks, in the spirit of empirical work by Kochar (1999). This encourages *ex ante* fertilizer expenditures in treatment villages, as it limits the downside risk of investment. On average, this implies the insurance channel made accessible by the bridge increases agricultural investment and yields. We show theoretically that this intuition is indeed the outcome of our model, consistent with our empirical evidence. Moreover, the model is consistent with the substantial shift of households from farming to wage work (as the wage increases), but also wage work to farming (as previously constrained households can better utilize their agricultural skill). The latter effect follows a similar logic to work on financial frictions and occupational choice (e.g. Banerjee and Newman, 1993; Buera, Kaboski and Shin, 2011; Midrigan and Xu, 2014). Constrained households may be forced to exit agriculture if their wealth is too low to properly utilize their agricultural skill. By allowing for an additional income stream from labor markets, a bridge allows those households to return to agricultural production. This result does not follow from a model in which there are no agricultural distortions.

[4] Rainfall variation, for example, limits fertilizer spending (Mobarak and Rosenzweig, 2012; Karlan et al., 2014). A bridge has no direct impact on rainfall variation.

1.1 Related Literature

The study of infrastructure benefits is large and varied. A recent literature has combined quantitative models with detailed data to provide evidence on the impact of trade costs from major infrastructure projects (Donaldson, 2013; Alder, 2017; Asturias, Garcia-Santana and Ramos, 2016). More closely related are those papers who explicitly highlight the rural-urban link in their study of trade, such as Adamopoulos (2011), Gollin and Rogerson (2014), Van Leemput (2016), and Sotelo (2016). Recently, a number of important papers have taken advantage of policy changes and natural experiments to identify the effects of infrastructure development, including Casaburi, Glennerster and Suri (2013) and Asher and Novosad (2016). The latter is closest to our work, as they find that new Indian roads generate movement out of agriculture. Dinkelman (2011) finds similar results, due to electrification in rural South Africa, and uses a similar "engineering-related" identification strategy based on land gradients. Relative to these papers, our close involvement in the actual construction of these bridges allows us to conduct detailed household-level surveys before and after construction to provide additional insight into the underlying mechanisms and multiple channels through which the bridge affects households. Moreover, while a number of these papers equate the rural economy with the agricultural economy, we show an important relationship between on- and off-farm outcomes within rural villages.

In particular, on-farm productivity increases because bridges allow for increases consumption smoothing through labor markets. This is consistent with a growing literature linking consumption risk to farm investments, including experimental evidence from Mobarak and Rosenzweig (2012) and Karlan et al. (2014), while Donovan (2016) highlights the importance of this channel for aggregate income differences. We show that self-insurance through labor market access can also generate increased fertilizer use and yields, qualitatively similar to results from formal insurance contracts highlighted in this literature. Moreover, our results have the policy implication that this self-insurance channel can be improved through better infrastructure. Relatedly, Bryan, Chowdhury and Mobarak (2014) and Bryan and Morten (2015) also highlight constraints to the spatial allocation of labor as a component of this agricultural productivity gap based on the misallocation of talent across sectors.

2 Background and Description of Intervention

2.1 Flooding Risk

Around the world, flooding – and especially river flooding – affects a disproportionate number of people relative to other natural issues. Using the EM-DAT (2017) International Disaster Database, we compile worldwide disasters from 2000 to 2016.[5] Figure 1a shows the fraction of occurrences accounted for by various types of emergencies, along with the fraction of people affected by each type. They are ordered according to affected population. Flooding accounts for over 40 percent of the people affected by disasters since 2000, followed only by drought. Figure 1b breaks “Floods” into four categories: river floods, flash floods, coastal floods, and uncategorized floods and reproduces Figure 1a in finer detail. Here, river floods in particular are the second largest factor affecting individuals, and lags only slightly behind droughts.

Figure 1: World Disasters (2000-2016)

(a) Different Disaster Types (b) Finer Decomposition of *Flood*

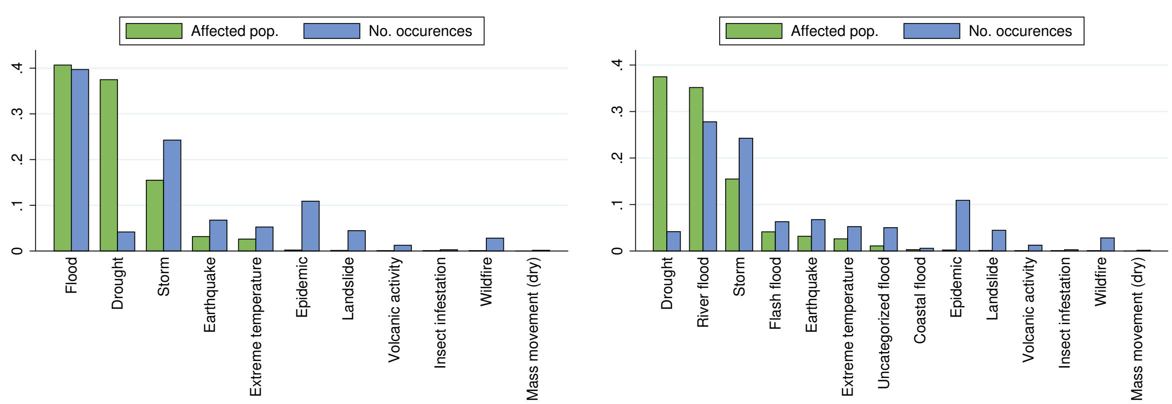

This issue is also salient in Northern Nicaragua, where our study takes place. Both policy makers and citizens cite flooding and the resulting isolation when combined with poor infrastructure as a critical development constraint (World Bank, 2008*a*). These villages are located in mountainous areas that face seasonal flooding during the rainy season each year (May to November). During these periods, streams and rivers that

[5] A disaster is included in the dataset if it meets one of the following conditions: 10 or more dead, 100 or more people affected, declaration of state of emergency, or a call for international assistance. See EM-DAT (2017) for further details.

are usually passable on foot rise very rapidly and may stay high for days or weeks. This flooding is unpredictable in its timing or intensity. Rainfall in the same location is a poor predictor of flooding, as rains at higher altitudes may be the cause of the flooding, a feature of flooding in other parts of the world as well (e.g. Guiteras, Jina and Mobarak, 2015, in Bangladesh). Moreover, this period is also the main cropping season. Crops are planted at the beginning of the rainy season in May, and harvested in late October and early November.

During these periods, some villages are cut off from access to outside markets. In particular, many villages have a river located between themselves and a larger, nearby city where agricultural markets and labor markets operate. When the river rises substantially, market access would require swimming across the river, which may be prohibitively dangerous and inhibit transportation of goods, or a long journey on foot to reach the market by another route.

2.2 Intervention and Identification Strategy

We investigate the impact of building footbridges that traverse these rivers. We do so by partnering with the non-governmental organization Bridges to Prosperity (B2P), that works to construct footbridges in these rural communities to solve some of the problems associated with flooding risk. Bridges to Prosperity provides engineering design, construction materials, and skilled labor to the village, as well as training in bridge maintenance. They ask members of the village to provide unskilled labor for construction, such as digging out the foundation of the bridge deck.

Bridges to Prosperity takes requests from local village organizations and governments, then evaluates these requests on two sets of criteria. First, they determine whether the village has sufficient need. That is, are there enough people that live in the village and that would use the bridge to justify the expense of the project. These decisions are made by an in-country manager employed by the organization who inspects each site.

If the village passes the needs assessment, the country manager personally goes to the site to do an engineering assessment. The purpose of this assessment is to determine if a bridge can, in fact, be built at the proposed site. To be considered feasible, the required bridge cannot exceed a maximum span of 30 meters, and the banks of the river on each side must be of similar height (a differential not exceeding 3 meters). Moreover, the estimated high water mark (maximum height of the river when flooded) must be at least two meters below the proposed bridge deck. The assessment makes other considerations as well: bridges cannot cross power lines, and they avoid building in places where the river bends (as river bends may indicate a river changing its course).

We compare communities that passed both the feasibility and the needs assessments, and therefore received a bridge, to those that passed the needs assessment, but failed the feasibility assessment. The second group makes for an ideal comparison group for two reasons. First, the fact that both groups have similar levels of need is crucial, as need is both unobservable and is likely to be highly correlated with the treatment effects. Second, failure of the feasibility assessment is very unlikely to be correlated with any relevant village characteristics. For observable differences, we show that villages that do and do not receive bridges are balanced.

Because of the expense of the bridges ($40,000) and the fact that a large part of the construction materials must be funded internally by Bridges to Prosperity, the number of bridges that can be constructed each year is limited. We study a total of fifteen villages. Of these, six passed both the needs and feasibility assessments, and therefore received bridges. The other nine passed only the needs assessment and did not receive a bridge. These villages are located in the provinces of Estelí and Matagalpa in northern Nicaragua.[6]

3 Data Collection and Design Validity

3.1 Data Collected

We collect two types of data. First, we conducted in-person household-level surveys with all households in each of the fifteen villages. The first such wave took place in May 2014, just as that year's rainy season was beginning. This survey was to

[6]One might be concerned that a control village may be treated if they are sufficiently close to a treatment village. That is, if the control villagers are sufficiently close to a bridge to access it. This is not the case in any of the fifteen villages. They are all sufficiently far from one another to eliminate this issue.

designed to give us an early indication of balance, and also to sign households up for the high frequency survey (discussed below). In this May survey, for those that agreed to participate, we conducted followups every two weeks by phone. The more critical surveys covering the main rainy season were conducted in November 2014, November 2015, and November 2016. Bridges were constructed in Spring of 2015. Therefore for all villages we observe three rainy seasons. For those that receive a bridge, we observe one rainy season without a bridge and two rainy season with a bridge. We will primarily focus on these three surveys, as the first survey in May 2014 covers the dry season which is not a major cropping period.[7] We do include it when we consider the validity of our identification strategy.

To collect the in-person household surveys, we employed local Nicaraguan enumerators. Our strategy was to survey all households within three kilometers of the proposed bridge site on the side of the river that was intended to be connected. In many villages, this implied a census of village households.

Participation in the first round of the survey was very high in general, with 97% of households agreeing to participate. This is true even though we offered no incentive for participation. Enumerators and participants were told that the purpose of the study was to understand the rural economy. We did not disclose our interest in the bridges because we suspected that would bias their answers, or may make them feel they are compelled to answer the survey when they would not otherwise want to participate. The number of households identified in each village varied widely, from a maximum of 80 to a minimum of 24.

Survey questions covered household composition, education, health, sources of income, consumption, farming choices (including planting, harvests, equipment and inputs), and business activities.

The second data collection was high frequency surveys. Because the floods are a high frequency and short term event, we also wanted to include these surveys to provide supporting evidence to the more detailed annual surveys and also validate the fact that flooding (and the bridge) affects income generating activities. We therefore carried out these surveys for 64 weeks, covering the rainy season before construction,

[7]Anticipating the results somewhat, none of the empirical results change if we include this first survey in our regressions.

along with the first dry and rainy seasons after construction. During the first wave, we solicited participation in cell phone followup interviews. Each household was called every other week, so that the maximum number of responses per household is 32. This high frequency survey covered income-generating activities, livestock purchases and sales, and food security questions over the past two weeks.

3.2 Balance and Validity of Design

As discussed above, we base our analysis on a comparison of villages that pass both the needs and feasibility assessment with those that pass only the needs assessment. The identification assumption is that the features required to pass the feasibility test are independent of any relevant household or village-level statistics. Using the first two waves of data, we run the regression

$$y_{ivt} = \alpha + \beta B_{vt} + \eta_t + \varepsilon_{ivt}$$

where $B_{vt} = 1$ if village v gets a bridge between $t = 2$ and $t = 3$. We consider a number of different outcomes, and show that households show no observable differences across the two groups. Table 1 produces the results, and we find no difference across households in build and no-build villages.

3.3 High Frequency Sample Selection

Because the high frequency data was collected over the phone, two issues are worth highlighting before turning to the empirical results. First, the high frequency data is not representative of the villages under study as not every individual has a cell phone. However, the households that participate are extremely close to population averages except for household head age. As one might suspect with a cell phone-based survey, those that agreed were slightly younger. The average age of a household that agreed to participate was 37 years old, compared to the average of 43 in the population as a whole. On other margins – occupation, farming, etc. – there is no statistical difference between those that participated and did not. Second, it is an unbalanced panel as not everyone answered the phone each time. We discuss the importance of this, and

Table 1: Pre-Bridge Differences

	Constant	Bridge
Household Composition		
HH head age	43.34***	1.39
	(0.00)	(0.18)
HH head yrs. of education	6.40***	0.33
	(0.00)	(0.22)
No. of children	1.30***	-0.03
	(0.00)	(0.70)
HH size	4.18***	0.15
	(0.00)	(0.19)
Occupational Choice		
Agricultural production	0.47***	0.01
	(0.00)	(0.76)
Off farm work	0.58***	0.03
	(0.00)	(0.54)
Total wage earnings (C$)	865.14***	46.94
	(0.00)	(0.74)
Farming		
Corn harvest	16.66***	0.43
	(0.00)	(0.88)
Bean harvest	12.09***	-1.79
	(0.00)	(0.26)
Plant corn?	0.17***	0.01
	(0.00)	(0.62)
Plant beans?	0.16***	-0.03
	(0.00)	(0.23)

Table notes: p-values in parentheses. We do no clustering procedure here as to give the regression the greatest chance of finding a statistically significant difference between the two groups.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

provide robustness checks, when discussing the results in Section 4.

4 The Impact of Bridges on Labor Market Earnings

We begin by assessing the direct impact of the bridge on labor market access. We first do so in both the high frequency data (Section 4.1), where we can assess the relationship between flooding and contemporaneous income realizations. In Section 4.2, we then ask whether the larger annual surveys also show higher labor market earnings. In both cases, we find that the bridge increases access to labor markets, and thus increases earnings.

4.1 High Frequency Effects of a Bridge

We begin by assessing the immediate affect of flooding and the impact of a bridge. To do so, we use the high frequency data to considering income realizations during floods. Figure 2 plots the histogram of the number of observations per household in the high frequency data. The minimum is 1, the maximum is 32 (also the maximum possible number of responses), and the average is 12.

Figure 2: Number of Observations per Household

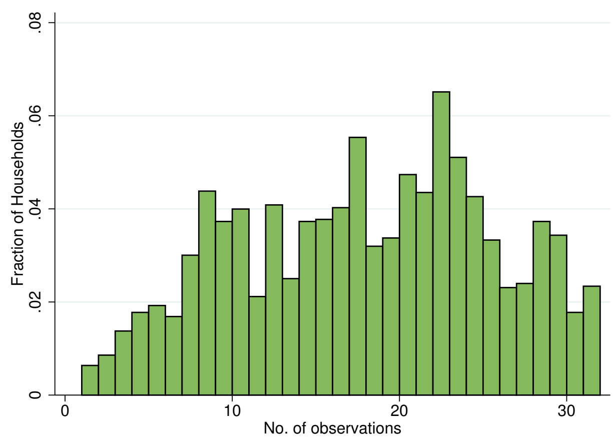

To assess the impact of flooding on different outcomes, we run regressions of the form

$$y_{ivt} = \alpha + \beta B_{vt} + \gamma\Big(B_{vt} \times F_{vt}\Big) + \theta\Big(NB_{vt} \times F_{vt}\Big) + \eta_t + \delta_i + \varepsilon_{ivt}. \tag{4.1}$$

The variable $B_{vt} = 1$ if village v has a bridge in week t, while $NB_{vt} = 1 - B_{vt}$ is the "no bridge" variable. The variable $F_{vt} = 1$ if village v is flooded at week t, while η_t and δ_i are week and individual fixed effects. Throughout, we use a wild bootstrap cluster at the village level.

4.1.1 Changes in Income Realizations

We begin by considering changes in income. This is the critical assumption our model makes, and is required for the results. We show here that labor market income does indeed change. Moreover, the bridge eliminates short-term market inaccessibility from flooding.

We first consider the high frequency data. Figure 3 is a histogram counting the share of weeks each household receives positive labor market income. Despite the fact that about half of households farm some kind of crop, most are also active in the labor market. When we rank households by the share of periods we observe positive income, even the fifth percentile household receives labor market income in 21 percent of the periods we observe it.[8]

Figure 3: Fraction of weeks with labor market income

(a) All households

(b) Only households with ≥ 10 observations

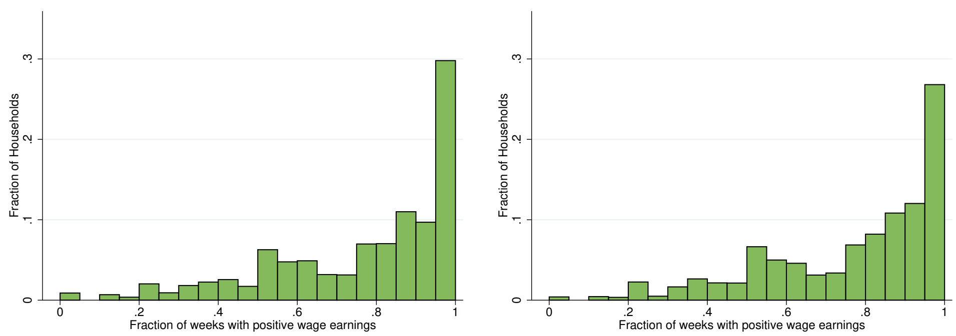

We therefore ask how income realizations change during flooding episodes, and how the bridge changes the results. We use two measures of income in regression (4.1): amount earned in the past two weeks and an indicator equal to one if no income was earned.

Table 2 illustrates the effects of flooding on contemporaneous income realizations. First, having a bridge in the absence of a flood does not increase income relative to households in villages without a bridge. This is shown by the insignificant effect on the bridge variable. However, when there is a flood, this changes. Income drops by C$184 ($p = 0.03$) during a flood in the absence of a bridge, a decrease of nearly 20 percent of its no-flood baseline.[9] This effect is not present in villages with a bridge.

[8] One possibility is that survey non-response is correlated with realizations of zero income, thus biasing our results toward observing positive income. This would be the case if heavy rains strongly reduced cell coverage, for example. In Appendix B we show that there is no relationship between flooding and the likelihood of response to surveys. Moreover, we take an extreme stance and assume every missed call implies zero income. This naturally affects the intensive margin of periods with income, but not the extensive margin.

[9] The Nicaraguan currency is the córdoba, denoted C$. The exchange rate is approximately C$29 = 1 USD.

Table 2: Effects of Flooding on Income

	Household Income	No Income Earned
Flood × No Bridge	-141.518*	0.070**
	(0.097)	(0.041)
Flood × Bridge	65.453	-0.038*
	(0.592)	(0.082)
Bridge	186.278**	0.061*
	(0.033)	(0.082)
Control mean	934.244	0.249
Observations	6756	6756
Individual F.E.	Y	Y
Week F.E.	Y	Y

Table notes: p-values in parentheses computed using the wild cluster bootstrap-t with 1000 simulations, clustered at the village level. Control mean is average dependent variable over entire time horizon for households in villages that never receive a bridge. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Here, a flood has no statistical effect on the average household income realization. That is, the flood has no effect on average income realizations in the presence of a bridge, but a negative effect without one.

The same pattern holds when one considers the fraction of people who earn no income in the preceding two weeks. The likelihood of earning no income increases by 13 percentage points ($p = 0.00$) when a flood occurs in villages without a bridge, from 0.21 to 0.33. In villages with a bridge, the fraction is 0.21 regardless of whether or not there is a flood. This seems the critical margin that the bridge affects. Figure 4 plots the density of income realizations in villages without a bridge (left panel) and with a bridge (right panel) during periods of flooding and no flooding. Among villages without a bridge, flooding shifts the distribution closer to zero. Once a bridge is constructed, the distributions track either other closely, regardless of flooding.

4.2 Longer Run Impacts from Annual Surveys

We ask how the short-run change in income and consumption risk generated by bridges translates into longer-term effects on labor market income among rural households. For that, we utilize our larger, annual surveys. Throughout, we use the three surveys conducted at the end of the rainy season from 2014 to 2016. We refer to them as

Figure 4: Density of Income Realizations

(a) No Bridge (b) Bridge

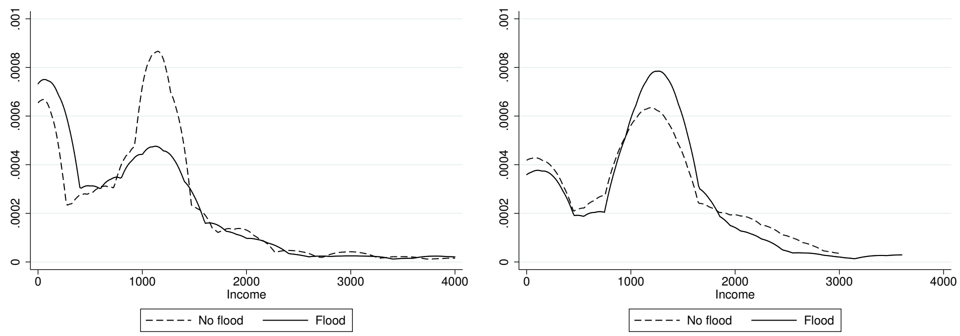

Figure notes: Figure 4a includes all village-weeks without a bridge, including those villages that eventually receive a bridge. Figure 4b includes all village-weeks post-construction.

$t = 0, 1, 2$ throughout this section. Our baseline regression specification is

$$y_{ivt} = \alpha + \beta B_{vt} + \eta_t + \delta_i + \varepsilon_{ivt} \tag{4.2}$$

where $B_{vt} = 1$ if a bridge is built, η_t and δ_i are time and household fixed effects, and standard errors are clustered at the village level using a wild cluster bootstrap-t. Panel A of Table 3 shows the results for total labor market income, along with its components of the daily wage rate and days worked. First, earnings increase by C\$308 ($p = 0.06$). This is almost entirely accounted for by an increase in income earned outside the village, consistent with the bridge providing better access to outside markets. Earnings outside the village increase by C\$295 ($p = 0.00$), while earnings inside the village decrease slightly by C\$42 ($p = 0.72$). These results are accounted for by changes in days worked, not by changes in the daily wage rate. Households work 1.25 extra days outside the village ($p = 0.00$), and 0.33 fewer days inside the village ($p = 0.41$), though the latter cannot be statistically distinguished from zero. We find no statistically significant effects on realized wages either within or outside the village.

Panel B of Table 3 distinguishes between intensive and extensive margin changes by interacting the bridge indicator with an indicator for positive earnings at baseline. In terms of total earnings, we see a significant movement of households into the labor market. Households with no baseline labor market earnings see an increase of C\$405

($p = 0.01$) compared to a statistically insignificant increase of C$221 ($p = 0.38$) among households with positive earnings. Again, this is driven by changes in days worked. Those with no baseline earnings increase days worked by 1.60 ($p = 0.00$), while those with baseline earnings increase days worked by a statistically insignificant 0.45 ($p = 0.53$). These results are consistent with households shifting from labor markets inside the village to outside the village. Indeed, among those with positive baseline earnings, we see a C$362 ($p = 0.00$) increase in earnings and a 1.36 increase in days worked outside the village, but also a decrease in earnings of C$205 ($p = 0.31$) and 1.02 days ($p = 0.10$) within the village. On the other hand, new entrants into the labor market more strongly move toward earnings outside the village, where we find an increase of C$295 ($p = 0.00$) and 1.72 days ($p = 0.00$), and smaller statistically insignificant changes within the village.

Table 3: Effects on Market Income, by Source

Panel A: Earnings	Total Earnings			Earnings Outside Village			Earnings Inside Village		
	Earnings	Wages	Days	Earnings	Wages	Days	Earnings	Wages	Days
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Build	307.59*	-21.25	1.00*	295.24***	-24.84	1.25***	-41.76	-54.75	-0.33
	(0.064)	(0.359)	(0.062)	(0.000)	(0.361)	(0.000)	(0.717)	(0.293)	(0.405)
Constant	1025.73***	275.77***	3.52***	295.00***	168.36***	1.72***	661.11***	263.43***	1.65***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Panel B: Intensive and Extensive Margins	Total Earnings			Earnings Outside Village			Earnings Inside Village		
	Earnings	Wages	Days	Earnings	Wages	Days	Earnings	Wages	Days
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Build × Pos. Earnings	221.12		0.45	362.44***		1.36***	-205.18		-1.02*
	(0.380)		(0.532)	(0.000)		(0.006)	(0.305)		(0.098)
Build × Zero Earnings	404.65**		1.60***	220.33**		1.12***	140.04		0.45
	(0.010)		(0.002)	(0.022)		(0.002)	(0.151)		(0.107)
Constant	1025.73***		3.52***	295.00***		1.72***	661.11***		1.65***
	(0.000)		(0.000)	(0.000)		(0.000)	(0.000)		(0.000)

Table notes: Pos. Earnings is an indicator for positive baseline labor market earnings, either inside or outside village. *Zero Earnings* is 1-*Pos. Earnings*. Wages are not included in Panel B since the zero earnings group has no defined wages in baseline. p-values in parentheses are clustered using the wild cluster bootstrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

These results constitute the immediate effects of the bridge. Taken together, they show that bridges increase access to labor markets. In particular, the high frequency results show that bridges eliminate the uncertainty related to flash flooding, and thus allow households to access the market even during floods. Moreover, using our more comprehensive annual surveys, we confirm that bridges generate an increase in labor market income.[10] The remaining question is how these changes in market access generate treatment effects in other aspects of the economy, including occupational choice and agricultural decisions.

5 Impact on Agricultural Outcomes

Bridges allow households to access labor markets. However, nearly half of household economic activity in the survey is accounted for by agricultural production. We therefore ask whether the bridge has any impact on agricultural outcomes. We begin with intermediate input use on farms, using regression (4.2), with results presented in Table 4. We consider intermediate input (fertilizer plus pesticide) expenditures, and also the two components individually. In odd columns, we provide the average effect of the bridge, while in even columns we decompose the treatment effects based on whether or not the household is operating a farm at baseline.

First, we see a substantial increase in intermediate expenditure, mostly driven by changes in fertilizer. Intermediate expenditures increase by C\$646 ($p = 0.013$) on a baseline of C\$934, and its components fertilizer and pesticide increase by C\$438 ($p = 0.005$) and C\$153 ($p = 0.286$) respectively. The even-numbered regressions decompose the results by continuing farmers and those who did not farm at baseline. We see that the results are roughly evenly split between the two groups. New farmers increase intermediate spending by C\$614 ($p = 0.042$) and continuing farmers increase by C\$675 ($p = 0.134$), with the caveat that we slightly miss statistical significance at the 10 percent level among the latter. Both groups see similar increases in fertilizer expenditures, while new farmers have higher pesticide expenditures in response to the bridge.

[10] Of course, income risk only matters to the extent it translates into consumption risk. In Appendix B.2, we use our high frequency data to show that households are more likely to be constrained from purchasing sufficient food during a flood in villages without a bridge. Similar to the results on income realizations, the bridge eliminates this risk.

Table 4: Farm Input Usage

	Intermediate Spending		Fertilizer Spending		Pesticide Spending	
	(1)	(2)	(3)	(4)	(5)	(6)
Build	646.48**		437.81***		152.94	
	(0.013)		(0.005)		(0.286)	
Build × Farming		674.72		464.66**		65.02
		(0.134)		(0.029)		(0.777)
Build × No farming		614.18**		407.22**		253.95**
		(0.042)		(0.034)		(0.045)
Control mean	889.56	889.56	607.43	607.43	303.48	303.48
Observations	1,601	1,601	1,601	1,601	1,601	1,601
Time F.E.	Y	Y	Y	Y	Y	Y
Household F.E.	Y	Y	Y	Y	Y	Y

Table notes: Farming = 1 if the household is engaged in any crop production at baseline. p-values in parentheses are clustered using the wild cluster bootstrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

We therefore next consider changes in harvest for maize and beans, measured in total quintales (100 kilograms) harvested.[11] The results are in Table 5. While the point estimates imply that harvest amounts increase in response to the bridge, we do not observe any statistically significant effects. This could be due to the relatively small number of clusters or unobserved aggregate shocks that limit the return to investment, a point recently emphasized in Rosenzweig and Udry (2016).

These results, coupled with the impact on labor market earnings, show that bridges have important effects on the two critical income generating processes in rural Nicaragua, wage work and agricultural production.

6 Occupational Choice

Since the bridge increases labor market earnings and agricultural production, a natural explanation is selection out of agriculture. If wage work becomes uniformly more profitable, then remaining farmers must be extremely productive. On average, this would generate an increase in agricultural productivity. We therefore ask how households switch between wage work and farming in response to a bridge.

[11] In Appendix B.3, we show that the bride has no effect on crop selection by farmers, hence our focus directly on yields here.

Table 5: Harvest and Yield of Staple Crops

	Maize		Beans	
	Harvest Quantity	Yield	Harvest Quantity	Yield
	(1)	(2)	(3)	(4)
Build	1.63	13.47**	1.06	3.26***
	(0.234)	(0.019)	(0.123)	(0.003)
Control mean	2.49	12.29	1.50	4.59
Observations	1,601	313	1,601	324
Time F.E.	Y	Y	Y	Y
Household F.E.	Y	Y	Y	Y

Table notes: Harvest quantity is measured in pounds harvested. Harvest quantity equals zero for any non-farming households. Yield is quantity harvested per manzana (1.73 acres) of land cropped, and is therefore not defined for non-farmers. p-values in parentheses are clustered using the wild cluster bootstrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

In our surveys, we asked individuals about their primary and secondary occupations, and use them to categorize households into four broad economic activities. Households are consider agricultural households if they only operate a farm, wage work households if they only have wage income (either on someone's farm or in a non-agricultural firm), both, or neither.

We begin by assessing the impact of the bridge on the persistence of sectoral employment. Figure 5 plots the simple averages of households engaged in agricultural production and labor market work in treatment and non-treatment villages.[12] The results show a remarkably stable aggregate fraction of households in both types of work across both treatment and control villages. Moreover, there is no obvious treatment effect from bridges.

To assess this more formally, we run a series of regressions. We define $O^j_{ivt} = 1$ if a household i in village v is engaged in activity $j = 1, 2$ (agricultural production and wage-earning activities, respectively) at year t. We interact the treatment with baseline activities.

[12] Note that these are not mutually exclusive categories, as households can be engaged in both agricultural production and earn wages in the labor market.

Figure 5: Fraction of Households Engaged in Agricultural Production and Labor Market Earnings

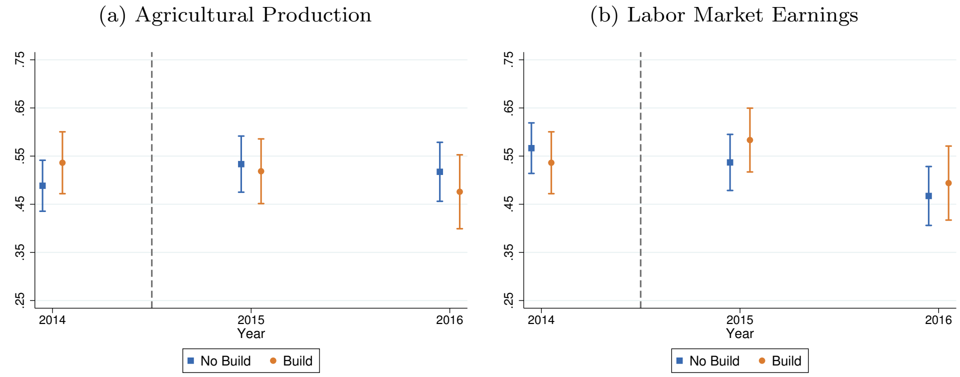

Figure notes: The dashed line indicates when bridges were constructed. The 95 percent confidence intervals are denoted by the bars surrounding the point estimates.

$$
\begin{aligned}
O^j_{ivt} &= \alpha + \beta B_{vt} + \eta_t + \delta_i + \varepsilon_{ivt} \\
O^j_{ivt} &= \alpha + \beta(B_{vt} \times O^j_{iv0}) + \gamma(B_{vt} \times (1 - O^j_{iv0})) + \eta_t + \delta_i + \varepsilon_{ivt}
\end{aligned}
$$

The results of both regressions are in Table 6. Regressions 1 and 3 show slight movement toward labor market earnings and away from agriculture, though neither effect can be statistically distinguished from zero. However, this does not imply that households do not switch economic activities. When we interact the bridge treatment with baseline activity, we find substantial movement across both types of work. Baseline farmers are 0.31 percentage points less likely to farm ($p = 0.000$), while baseline wage earners are 13 percentage points less likely to earn wages ($p = 0.012$) once a bridge is constructed. Thus, the aggregated results in Figure 5 are a result of roughly offsetting movements into and out of each occupation. This result is consistent with the idea that a bridge indirectly eliminates some agricultural distortion, as previously constrained individuals move into agriculture.

To assess this reallocation in more detail, we decompose the occupational space into the four mutually exclusive groups – only agricultural production, only labor

Table 6: Effects on Persistence of Activities

	Agriculture	Agriculture	Labor Market	Labor Market
	(1)	(2)	(3)	(4)
Build	-0.076		0.062	
	(0.223)		(0.227)	
Build × Engaged		-0.311***		-0.134**
		(0.000)		(0.012)
Build × Not engaged		0.202**		0.286***
		(0.013)		(0.001)
Control: Fraction of HH engaged	0.488		0.538	
Control: Engaged – Engaged		0.799		0.853
Control: Not engaged – Engaged		0.192		0.193
Observations	1,347	1,347	1,347	1,347
Time F.E.	Y	Y	Y	Y
Household F.E.	Y	Y	Y	Y

Table notes: Engaged = 1 if the household is engaged in the relevant activity at baseline, and *Not engaged* = 1 if the household is not engaged in the relevant activity at baseline. p-values in parentheses are clustered using the wild cluster bootstrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

market earnings, both, and neither – and run the regressions

$$o_{j,ivt} = \alpha + \sum_{j=1}^{4} \beta_j \Big(B_{vt} \times o_{j,iv0} \Big) + \eta_t + \delta_i + \varepsilon_{ivt} \qquad \text{for } j \in \{1,2,3,4\}.$$

Here $o_{j,ivt}$ is an indicator that household i is engaged in activity $j \in \{1,2,3,4\}$ defined above. The results are in Table 7, and a number of results emerge. First, the bridge induces households to engage in market economic activity. For those currently engaged in no economic activity the bridge has a strong positive effect on engaging in either agriculture or wage work, and a strong negative effect (-0.547, $p = 0.00$) on engaging in no market activity. Second, the bridge allows households to specialize, whether it be in farming or wage work. For households engaged in both farming and wage work (e.g. "both"), there is a strong positive effect of the bridge on the likelihood of engaging in *only* farming or wage work. Moreover, the effect of the bridge on engaging in both is negative and significant (-0.505 with $p = 0.00$). Lastly, as in the previous set of results in Table 6, the bridge generates substantial switching across these categories. To see this, one can simply read the negative, statistically significant effects off the diagonal of Table 7. For any current economic activity, a bridge makes it significantly less likely that a household is engaged in that same activity post-treatment.

Table 7: Effects on Persistence of Activities, Mutually Exclusive Categories

	Agriculture only	Wage work only	Both	None
	(1)	(2)	(3)	(4)
Build × Agr only	-0.366***	0.225***	0.062*	0.080*
	(0.000)	(0.001)	(0.133)	(0.069)
Build × Wages only	0.044	-0.218***	0.126*	0.049
	(0.224)	(0.001)	(0.066)	(0.175)
Build × Both	0.178***	0.275***	-0.505***	0.052
	(0.009)	(0.000)	(0.000)	(0.417)
Build × None	0.268**	0.204***	0.076	-0.547***
	(0.024)	(0.000)	(0.201)	(0.000)
Control: From Agr	0.775	0.171	0.032	0.023
Control: From Wage	0.069	0.751	0.122	0.059
Control: From Both	0.168	0.189	0.613	0.029
Control: From None	0.185	0.147	0.014	0.653
Observations	1,601	1,601	1,601	1,601
Time F.E.	Y	Y	Y	Y
Household F.E.	Y	Y	Y	Y

Table notes: The interaction terms are the activity of the household at baseline. p-values in parentheses are clustered using the wild cluster bootstrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

7 Model and Rationale for Findings

We have shown three main results: labor market earnings increase, farm expenditures and yields increase, and there is substantial reallocation of activity between wage work and farming. In this section, we develop a model that links these results together. Critically, we show that occupational switching – in particular movement *toward* agriculture – is consistent with the idea that the bridge decreases distortions in the agricultural sector through its effect on labor markets. This has the additional implication of increasing fertilizer investment and yields, consistent with our results.

Households are infinitely lived and consume a single good. They are endowed with one unit of time. The good is storable between periods with exogenous return

R (which may be less than one), and households do not have access to any state-contingent assets, and therefore implies that savings are subject to the constraint $S \geq 0$. They maximize expected utility

$$\mathbb{E}_0\Big[\sum_{t=0}^{\infty}\beta^t u(c_t)\Big]$$

with discount factor $\beta \in (0, 1)$. We assume that u is increases, concave, and prudent.

A household can work in one of three "occupations:" only wage work (denoted w), only farming (f), or both (b). A working household uses its entire time endowment in wage work, while a farming household uses all its time farming. A household that both farms and works uses an exogenous fraction n of time in wage work and $1 - n$ in farming. Occupations can be changed each period depending on the optimal choice given a household's state.

7.1 Timing and Production

Time is comprised of two different types of periods. Every T periods, any households that farm enter the "harvesting and planting" period, t^H. This involves two things: harvesting the crops planted T periods ago at $t^H - T$, and also planting the crops that will be harvested T periods from today at $t^H + T$. Intermediate inputs (e.g. fertilizer and pesticide) have to be purchased at the planting period, with payoff coming in the form of harvest T periods later. The harvest amount is given by the production function

$$Y = ZX^{\alpha}N^{1-\alpha}$$

where Z is a random shock, X is total intermediate inputs invested, and N is labor. The labor input depends on occupational choice,

$$N = \begin{cases} 1 & \text{if occupation} = f \\ 1 - n & \text{if occupation} = b. \end{cases}$$

Working households (both occupations w and b) receive a stochastic wage each period, including the "inner periods" between planting and harvesting. We assume for simplicity that the shocks to wages and farms are independent and follow Markov processes, so that the joint distribution is given by $H(w_{t+1}, Z_{t+1}|w_t, Z_t) = G(w_{t+1}|w_t)F(Z_{t+1}|z_t)$.[13] For simplicity and ease of notation, define $s_t = (w_t, Z_t)$, and its associated cumulative distribution function $H(s_{t+1}|s_t)$.

7.2 Recursive Formulation

Given the timing described above, we can write the household problem recursively. First, the indirect utility of entering a period with assets A, wage draw w, and farm shock Z is

$$V(A, s) = \max \Big\{ V^A(A, s), V^W(A, s), V^B(A, s) \Big\}. \tag{7.1}$$

The value of becoming a worker to a household with savings S_0, wage realization w and farm shock Z is

$$\begin{aligned} V^W(A, s) \quad = \quad & \max_{\{c_t, A_{t+1}\}} \sum_{t=0}^{T-1} \beta^t \int u(c_t(s_t)) dH(s_t|s_{t-1}) + \beta^T \int V(A_T, s_T) dH(s_T|s_{T-1}) \quad (7.2) \\ & \text{subject to: } c_t(s_t) \leq w_t + RA_t(s_t) - A_{t+1}(s_t) \\ & \qquad A_0 = A \\ & \qquad A_{t+1}, c_t \geq 0 \end{aligned}$$

Note that once occupational choice is made, households must wait T periods to update. That is, a household cannot switch to farming if it had not planted anything. Turning to farming, the value of only farming is given by

$$\begin{aligned} V^A(A, s) \quad = \quad & \max_{\{\phi, c_t, A_{t+1}\}} \sum_{t=0}^{T-1} \beta^t \int u(c_t(s_t)) dH(s_t|s_{t-1}) + \beta^T \int V(A'(s_T), s_T) dH(s_T|s_{T-1}) \quad (7.3) \\ & \text{subject to: } c_t(s_t) \leq RA_t(s_t) - A_{t+1}(s_t) \\ & \qquad A'(s_T) = A_T(s_t) + Z_T(\phi A)^\alpha \\ & \qquad \phi \in [0, 1] \\ & \qquad A_{t+1}, c_t \geq 0 \end{aligned}$$

[13] This is not required for the results, but simplifies the analysis and proofs.

Notice that there exists no credit market for intermediate inputs. That is, intermediates must be purchased out of existing assets through the term ϕ. Lastly, a household can engage in both farming and wage work. That household is required to spend (exogenously given) fraction $1-n$ of their time on the farm, and the rest in wage work. The value of doing both jobs is given by

$$V^B(A,s) = \max_{\{\phi,c_t,S_{t+1}\}} \sum_{t=0}^{T-1} \beta^t \int u(c_t(s_t))dH(s_t|s_{t-1}) + \beta^T \int V(A'(s_T),s_T)dH(s_T|s_{T-1}) \tag{7.4}$$

$$\begin{aligned} \text{subject to: } & c_t(s_t) \leq RS_t(s_t) - S_{t+1}(s_t) + nw_t(s_t) \\ & A'(s_T) = S_T(s_t) + z(s_T)(\phi A)^\alpha (1-n)^{1-\alpha} \\ & \phi \in [0,1] \\ & A_{t+1}, c_t \geq 0 \end{aligned}$$

7.3 Characterization and Analytical Results

Our assumption – which we verified empirically in Section 4 – is that a bridge allows for better market access, and in turn, better access to wage earning activities. In the context of the model this is both a decrease in the variance and an increase in the mean of the wage shock process. Our goal in this section is to show theoretically how this change affects agricultural production and occupational choice. We discuss them in turn.

7.3.1 Farming Decisions

We begin by discussing the returns to farming. First, we highlight the key distortions in the agricultural sector. Let η be the Lagrange multiplier on non-negativity of savings. Taking first order conditions in the farming problem in equation (7.3) and rearranging yields[14]

$$R^T + \frac{\sum_t \int \eta_t(s_t)dF(s_t)}{\int V_1(A'(s),w(s),Z(s))dF(s)} = \alpha(\phi A)^{\alpha-1} \int Z(s) \frac{V_1(A'(s),w(s),Z(s))}{\int V_1(A'(s),w(s),Z(s))dF(s)} dF(s). \tag{7.5}$$

[14]The implications for households that both farm and work in recursive program (7.4) are similar, with the inclusion of labor $n<1$.

Equation (7.5) shows both the credit and insurance distortions present in the agricultural sector. The first is driven by the lack of insurance combined with the fact that intermediates are chosen before the shock Z is realized. Without state-contingent assets, this production risk translates into consumption risk. The second is a credit market distortion: because farmers cannot borrow to purchase intermediate inputs, they may potentially be constrained from purchasing the profit-maximizing level of intermediates.

To see this more clearly, note that if households had access to intermediate input credit,[15] and had a complete set of state contingent assets, the first order condition would be

$$R^T = \alpha(\phi A)^{\alpha-1} \int Z(s) dF(s). \tag{7.6}$$

The left-hand side of (7.6) is the cost of fertilizer expenditures, as the opportunity cost is saving at return R. The right-hand side is the expected marginal return to fertilizer use. Comparing (7.5) and (7.6), there are two additional pieces in (7.5). The first is the inclusion of risk neutral productivities. That is, the probability that each shock occurs is weighted by its marginal utility, not only the likelihood that it occurs. This is driven by the fact that farmers do not have state-contingent assets, and therefore internalize how these decision affects future consumption. It immediately implies that farming households use less than the profit maximizing fertilizer amount. The second distortion is that farmers are subject to a non-negativity constraint on savings, the second term in equation (7.5). Some farmers have sufficiently low assets, which limits their ability to invest in fertilizer. This further shifts the choice of fertilizer away from its profit maximizing level.

Proposition 1 formalizes how these distortions change fertilizer expenditures in response to changes in the wage distribution.

Proposition 1. *(Characterization for households that farm) Suppose* $H(w', Z'|w, Z, s) = G(w'|w)F(Z'|Z)$*, and that the distribution* G *changes to a new distribution* $\bar{G}$*. Then within the problems in equations (7.3) and (7.4):*

[15]More formally, farmers would purchase intermediate inputs at a price of R and repay R^T after the harvest.

1. *G is a mean-preserving spread of $\bar{G}$* $\implies \phi(A,w,Z) < \bar{\phi}(A,w,Z)$

2. $G(w'|w) = \bar{G}(aw'|w)$ *for* $a > 1 \implies \phi(A,w,Z) < \bar{\phi}(A,w,Z)$

Proof. See Appendix A. ■

Proposition 1 links the wage distribution with farming outcomes. Intuitively, it makes use of the the risk neutral probabilities in the first order condition (7.5). A decrease in the variance of shocks shifts weight away from both positive and negative tail realizations. However, since decreasing marginal utility guarantees that households are more sensitive to negative shocks than positive ones, the weight shifts toward more positive realizations. This increases *ex ante* fertilizer investment. A similar result holds when the mean increases. Moreover, the result does not hold in the absence of these distortions, as can be seen from (7.6). Thus, if the bridge increases farming outcomes, this is evidence of agricultural distortions changing through infrastructure development.

7.3.2 Occupational Choice

The next set of results characterizes occupational choice. Here, we show that the agricultural distortions highlighted above imply that households move from farming to wage work, but also wage work to farming. First, Proposition 2 shows that the occupations follow cutoff rules.

Proposition 2. *(Characterization of occupation choices in the cross-section) Discrete occupation choices follow cutoff rules in A for every (w,z).*

1. *Wage cutoffs: For every z, there exists a $w_F(z)$ and $w_B(z)$ such that, for A large enough, if $w > w_B(z)$ then the household only works, if $w < w_F(z)$ then the household only farms, and otherwise the household both works and farms.*

2. *Asset cutoffs: For every (w,z) pair, (1) $w > w_B(z)$ implies that the household works at every asset level, (2) $w_F(z) < w < w_B(z)$ implies that the household works if and only if their assets are below a threshold $A_B(w,z)$ and both farms and works otherwise, and (3) $w < w_F(z)$ implies that the household only works if assets are below one threshold $A_B(w,z)$, both works and farms if instead assets are below $A_F(w,z)$ and otherwise only farms.*

Proof. See Appendix A. ■

Proposition 2 shows that the occupational choice is defined by a series of shock and asset holding cutoffs. Figure 6 shows occupational choice graphically for a given farm shock z, for both the model developed above, and the model in which there are no agricultural distortions. Figure 6 shows that agricultural distortions generate curvature in the decision cutoffs. when we consider the response of occupations to changes in the distribution of wage shocks, it is sufficient to consider how these cutoffs move. In Proposition 4, we show how the wage cutoffs affect occupational choice.

Figure 6: Occupational Choice for a Given Farm Shock z

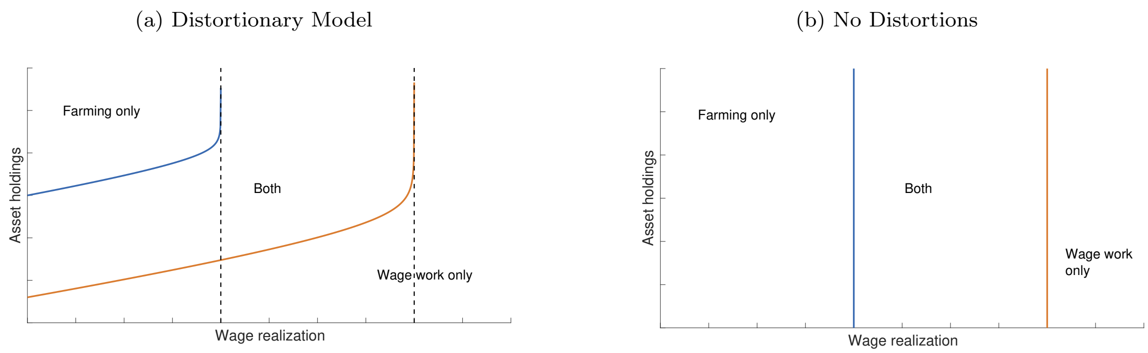

Proposition 3. *(Changes in occupational choice with changes in shock process) Suppose*

1. *G is a mean-preserving spread of $\bar{G}$, or*

2. *$G(w'|w) = \bar{G}(aw'|w)$ for $a > 1$*

Then when a household faces $\bar{G}$ instead of G: (1) $w_B(z)$ decreases and $w_F(z)$ decreases for every z

Proof. See Appendix A. ■

Notice that Proposition 4 implies that changes in the wage process generate a shift toward wage work. This is independent of the inclusion of distortions in the model: some households originally engaged in both farming and working switch exclusively to working, while other households engaged exclusively in farming begin to also work. Graphically, this is shown in Figure 7.

Figure 7: Wage-cutoff shifts in response to changing shock process for given z

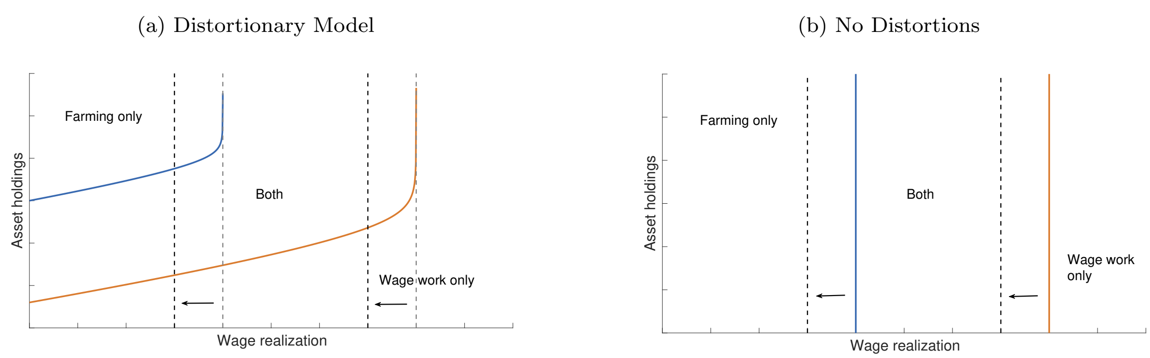

However, there is an additional effect when distortions are present.

Proposition 4. *(Changes in occupational choice with changes in shock process) If the Lagrange multiplier η is sufficiently high on a household engaged in only wage work, that household will switch to farming in response to the shift from G to $\bar{G}$.*

Proof. See Appendix A. ■

Intuitively, this result makes use of the fact that the labor market allows households to better overcome its agricultural distortions. Therefore, if a household is sufficiently skilled in farming, but too poor, the labor market allows for the household to overcome this. Critically, this is not an outcome of a model with no agricultural distortions.

8 Conclusion

We consider the impact of new footbridges in rural Northern Nicaragua. The villages that we study are subject to sporadic seasonal flooding that cuts off households from local markets. Working with an NGO partner, we construct footbridges to link these villages back to markets, and use the small but critical engineering requirements to identify the effect. Despite the fact that we construct only 6 bridges in 15 villages, we identify a number of important changes among households. First, the bridge eliminates any change in income realizations during floods. When we consider longer run outcomes, the bridges induce substantial changes in economic activity, as the persistence of both farming and wage work decrease. Moreover, farmers increase fertilizer and pesticide investment, while storing less. This is consistent with the bridge as a income smoothing technology.

The bridge induces a reduction in both extensive and intensive margin misallocation. Finding evidence of these multiple channels is important for policy, given the variety of income-generating activities in rural areas (World Bank, 2008*b*). Given the relatively small sample, however, we have little to say about general equilibrium effects here. This is an important component of understanding the full effect of such interventions, and we leave this to future research.

References

Adamopoulos, Tasso. 2011. "Transportation Costs, Agricultural Productivity, and Cross-Country Income Differences." *International Economic Review*, 52(2): 489–521.

Adamopoulos, Tasso, Loren Brandt, Jessica Leight, and Diego Restuccia. 2017. "Misallocation, Selection and Productivity: A Quantitative Analysis with Panel Data from China." NBER Working Paper 23039.

Alder, Simon. 2017. "Chinese Roads in India: The Effect of Transport Infrastructure on Economic Development." Working Paper, University of North Carolina.

Asher, Sam, and Paul Novosad. 2016. "Market Access and Structural Transformation: Evidence from Rural Roads in India." World Bank Working Paper.

Asturias, Jose, Manuel Garcia-Santana, and Roberto Ramos. 2016. "Competition and the Welfare Gains from Transportation Infrastructure: Evidence from the Golden Quadrilateral in India ." Working Paper, Universitat of Pompeu Fabra.

Banerjee, Abhijit V., and Andrew F. Newman. 1993. "Occupational Choice and the Process of Development." *Journal of Political Economy*, 101(2): 274–298.

Bryan, Gharad, and Melanie Morten. 2015. "Economic Development and the Spatial Allocation of Labor: Evidence from Indonesia." London School of Economics Working Paper.

Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak. 2014. "Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Banladesh." *Econometrica*, 82(5): 1671–1748.

Buera, Francisco J., Joseph P. Kaboski, and Yongseok Shin. 2011. "Finance and Development: A Tale of Two Sectors." *American Economic Review*, 101(5): 1964–2002.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *Review of Economics and Statistics*, 90(3): 414–427.

Casaburi, Lorenzo, Rachel Glennerster, and Tavneet Suri. 2013. "Rural Roads and Intermediated Trade: Regression Discontinuity Evidence from Sierra Leone." Working Paper, MIT.

Dinkelman, Taryn. 2011. "The Effects of Rural Electrification on Employment: New Evidence from South Africa." *American Economic Review*, 101(7): 3078–3108.

Donaldson, Dave. 2013. "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure." *American Economic Review*. forthcoming.

Donovan, Kevin. 2016. "Agricultural Risk, Intermediate Inputs, and Cross-Country Productivity Differences." Working Paper, University of Notre Dame.

EM-DAT. 2017. "The CRED/OFDA International Disaster Database." Université Catholique de Louvain, Brussels, Belgium.

Foster, Andrew D., and Mark R. Rosenzweig. 2007. "Economic Development and the Decline of Agricultural Employment." In *Handbook of Development Economics*. , ed. T. Paul Schultz and John A. Strauss, 3051–3083.

Gollin, Douglas, and Richard Rogerson. 2014. "Productivity, Transport Costs and Subsistence Agriculture." *Journal of Development Economics*, 107(1): 38–48.

Gollin, Douglas, David Lagakos, and Michael E. Waugh. 2014. "The Agricultural Productivity Gap in Developing Countries." *Quarterly Journal of Economics*, 129(2): 939–993.

Guiteras, Raymond, Amir Jina, and A. Mushfiq Mobarak. 2015. "Satellites, Self-reports, and Submersion: Exposure to Floods in Bangladesh." *American Economic Review Papers and Proceedings*, 105(2): 232–236.

Kamarck, Andrew K. 1973. *The Tropics and Economic Development: A Provocative Inquiry into the Poverty of Nations.* Baltimore:The Johns Hopkins University Press.

Karlan, Dean, Robert Osei, Isaac Osei-Akoto, and Christopher Udry. 2014. "Agricultural Decisions After Relaxing Credit and Risk Constraints." *Quarterly Journal of Economics*, 129(2): 597–652.

Kochar, Anjini. 1999. "Smoothing Consumption by Smoothing Income: Hours of Work Responses to Idiosyncratic Agricultural Shocks in Rural India." *Review of Economics and Statistics*, 81: 50–61.

Midrigan, Virgiliu, and Daniel Yi Xu. 2014. "Finance and Misallocation: Evidence from Plant-Level Data." *American Economic Review*, 104(2): 422–458.

Mobarak, A. Mushfiq, and Mark Rosenzweig. 2012. "Selling Formal Insurance to the Informally Insured." Yale Economics Growth Center Discussion Paper No. 1007.

Restuccia, Diego, and Raul Santaeulalia-Llopis. 2017. "Land Misallocation and Productivity." University of Toronto Working Paper 575.

Restuccia, Diego, and Richard Rogerson. 2016. "The Causes and Costs of Misallocation." *Journal of Economic Perpsectives.* in preparation.

Restuccia, Diego, Dennis Tao Yang, and Xiaodong Zhu. 2008. "Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysis." *Journal of Monetary Economics*, 55(2): 234–250.

Rosenzweig, Mark, and Christopher Udry. 2016. "External Validity in a Stochastic World." NBER Working Paper 22449.

Sotelo, Sebastian. 2016. "Domestic Trade Frictions and Agriculture." Working Paper, University of Michigan.

Van Leemput, Eva. 2016. "A Passage to India: Quantifying Internal and External Barriers to Trade." International Finance Discussion Papers 1185, Board of Governors of the Federal Reserve System.

World Bank. 2008*a*. "Nicaragua Poverty Assessment." World Bank Report, Washington D.C.

World Bank. 2008*b*. *World Development Report 2008: Agriculture for Development.* Washington, DC:The World Bank Group.

A Proofs

A.1 Proof of Proposition 1

Proof. Let η be the Lagrange multiplier on non-negativity of savings. Taking first order conditions in the problem defined in equation (7.4) and rearranging yields:

$$R^T \int V_1(A'(s), w(s), Z(s))dF(s) + \sum_t \int \eta_t(s_t)dF(s_t) = \alpha(\phi A)^{\alpha-1} \int Z(s)V_1(A'(s), w(s), Z(s))dF(s) \tag{A.1}$$

For now, suppose all values of $\eta_t = 0$. Then this condition implies:

$$R^T = \alpha(\phi A)^{\alpha-1} \int Z(s) \frac{V_1(A'(s), w(s), Z(s))}{\int V_1(A'(s), w(s), Z(s))dF(s)} dF(s) \tag{A.2}$$

Suppose ϕ is held fixed at its optimal value under the distribution G. Note that the case of independent wage and agricultural productivity shocks implies that $S(s_T)$ is independent of $Z(s_T)$. Under the distribution $\bar{G}$, the variance of $S(s_T)$ across states is lower than under G with ϕ held fixed. Because V_1 is convex and decreasing, this reduction in variance implies that $E(V_1|Z(s_T))$ falls by more for small values of $Z(s_T)$ than for large values of $Z(s_T)$. Note that the integral term on the right hand side of equation (A.2) is the risk neutral-weighted expectation of the shock Z. Then the previous argument implies that probability mass under $\bar{G}$ is shifted more toward higher values of Z than under G, which implies that this integral increases. It is trivial to show that the right hand side is strictly decreasing in ϕ. Therefore, ϕ is higher under $\bar{G}$ than under G. Notice that because the η_t terms are unambiguously lower under $\bar{G}$ than under G, then this implies the first result.

The proof for the second result is very similar. A uniform increase in w implies that for every value of $Z(s_T)$, $S(s_T)$ is greater (again, keeping ϕ fixed). Because V_1 is convex and decreasing, the average marginal value of assets falls by more for low values of $Z(s_T)$ than for high values. Therefore, by the same argument as above, the right hand side of equation (A.2) increases if ϕ is fixed. Therefore, ϕ must increase. As before, this argument is only reinforced by the reduction in the η_t terms. ■

A.2 Proof of Proposition 2

Proof. It is useful to show that:

$$\frac{\partial V^W(A,w,z)}{\partial z} < \frac{\partial V^B(A,w,z)}{\partial z} < \frac{\partial V^A(A,w,z)}{\partial z} \tag{A.3}$$

$$\frac{\partial V^W(A,w,z)}{\partial w} > \frac{\partial V^B(A,w,z)}{\partial w} > \frac{\partial V^A(A,w,z)}{\partial w} \tag{A.4}$$

$$\frac{\partial V^W(A,w,z)}{\partial A} < \frac{\partial V^B(A,w,z)}{\partial A} < \frac{\partial V^A(A,w,z)}{\partial A} \tag{A.5}$$

The first two inequalities follow immediately from the envelope theorem. The third is due to two effects. First, when the household does more farming and assets are the same, the marginal utility of consumption is higher. Second, the marginal product of on-farm investment is higher the greater is the intensity of farming. Both these effects imply the third inequality.

The first set of results follows immediately from the second inequality coupled with the fact that clearly $V^W(A,0,z) < V^B(A,0,z) < V^A(A,0,z)$. The second set of results follow from how the wage cutoffs were defined, and from the third inequality. ■

A.3 Proof of Proposition 4

Proof. The first results are obvious given that households are risk-averse, so a less risky wage process is more attractive to them. $A_F(w,z)$ rises because the wage process does not enter V^A except through the continuation term, which also appears in V^B, and does increase V^B. Households on the cutoff between working and both farming and working are more inclined to work, because the wage process is less risky, but are more inclined to also farm, because the reason they did not farm was the need for a buffer stock of assets. As the need for a buffer stock is diminished they would be more inclined to farm. Which effect dominates depends on parameters, such as n. ■

B More Results and Robustness

B.1 How high frequency survey response rates change during floods

Figure 3 in the text shows that almost all individuals in the high frequency survey use the labor market to some degree. However, our survey is biased toward finding that result if floods decrease the likelihood of answering the survey. To show that this is not the case, we run the regression

$$\mathbb{1}[answer]_{ivt} = \alpha + \beta Flood_{vt} + \eta_t + \delta_i + \varepsilon_{ivt}.$$

where $\mathbb{1}[answer]_{ivt} = 1$ if an individual answers the survey in week t, and is zero otherwise. The results are in Table 8. We find no statistically different effect of flood on the response rate, and the point estimate is small. If we remove time fixed effects we are able to generate a negative response to flooding, but again, the point estimate is quite small.

Table 8: Effect of flooding on survey response

	(1)	(2)
Flood	0.026	-0.025**
	(0.151)	(0.035)
Constant	0.580***	0.498***
	(0.000)	(0.002)
Observations	13,705	13,705
Individual F.E.	Y	Y
Week F.E.	Y	N

Table notes: p-values in parentheses are clustered using the wild cluster bootstrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

To further emphasize this point, Figure 8 reproduces Figure 3 in the main text with one key difference. Here, we assume that every period a household does not answer the survey, they received zero income that period. That is, we replace all missing values with zeros. This extreme assumption generates the lowest possible bound on the results driven by the unbalanced nature of the panel.

Naturally, this shifts the distribution toward zero. However, even when considering

Figure 8: Fraction of weeks with labor market income

(a) All households

(b) Only households with ≥ 10 observations

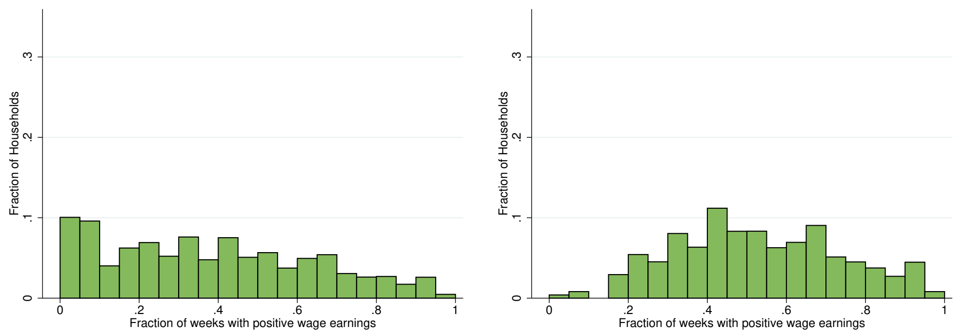

all households, the fifth percentile household still receives labor market income in 3 percent of its observations. The median household receives labor market income in 36 percent of weeks. Thus, individuals are still utilizing the labor market to varying degrees of intensity. When we condition on households that have at least ten observations, the numbers look quite similar to the text. The fifth percentile household receives labor market income in 21 percent of weeks. Thus, even under the most extreme assumptions about non-response, the labor market is still an important part of most households income strategy.

B.2 Changes in Food Security from High Frequency Surveys

While flooding increases income uncertainty by eliminating access to labor markets, it also directly contributes to consumption uncertainty. In particular, households are prohibited from purchasing as much food from local stores as they desire. This occurs when vendors are worried about running out of stock during flood episodes. Indeed, Table 9 shows that this behavior is concentrated in flooding episodes. Once a village has a bridge, this rationing ends abruptly. For instance, there are no observations of households being rationed in corn purchases in any village that has a bridge. Moreover, the frequency of bean and rice rationing is also greatly mitigated by a bridge. This is consistent with less worry about access to food during flooding times, because households and food vendors are able to access outside food supplies when they are necessary. Regression three shows total food spending in the previous month. Food spending actually increases by C$ 32 ($p = 0.000$) during a flood among households with a bridge, while there is a small, statistically insignificant change of C$ 0.50 ($p = 0.959$) among those without a bridge. During non-flooding weeks, food spending cannot be distinguished between the build and no-build villages. These results together show the impact on food consumption uncertainty. They imply that a bridge allows households to better control consumption during floods.

Table 9: Food Rationing and Spending During Floods

	Maize Rationed	Beans Rationed	Food Spending
	(1)	(2)	(3)
Flood × No Bridge	0.110**	0.147***	0.497
	(0.010)	(0.001)	(0.959)
Flood × Bridge	0.014	0.027	32.002***
	(0.749)	(0.505)	(0.007)
Bridge	-0.052	-0.051*	17.324
	(0.139)	(0.058)	(0.321)
Control mean	0.140	0.159	476.164
Observations	6750	6750	6750
Individual F.E.	Y	Y	Y
Week F.E.	Y	Y	Y

Table notes: p-values in parentheses computed using the wild cluster bootstrap-t with 1000 simulations, clustered at the village level. Control mean is average dependent variable over entire time horizon for households in villages that never receive a bridge. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

B.3 Crop Planting Decisions

We look at planting decisions, where we consider the two key staple crops maize and beans along with the main cash crop in Northern Nicaragua, coffee.[16] The outcome variable here is an indicator equal to one if the crop is planted (not necessarily harvested), and the results are in Table 10.

Table 10: Planting Decisions

	Maize		Beans		Coffee	
	(1)	(2)	(3)	(4)	(5)	(6)
Build	0.003		0.080		0.004	
	(0.606)		(0.178)		(0.766)	
Build × Farming		-0.034		0.045		-0.003
		(0.679)		(0.598)		(0.863)
Build × No farming		0.047		0.123^{**}		0.127
		(0.159)		(0.012)		(0.523)
Constant	0.217^{***}	0.218^{***}	0.272^{***}	0.272^{***}	0.018^{***}	0.018^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)	(0.001)
Observations	1,601	1,601	1,601	1,601	1,601	1,601
Time F.E.	Y	Y	Y	Y	Y	Y
Household F.E.	Y	Y	Y	Y	Y	Y

Table notes: p-values in parentheses are clustered using the wild cluster bootstrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 11: Land Use and Farm Size

	Total Land Owned	Total Land Cropped	Rent out any land?
	(1)	(2)	(3)
Build	-0.343	-0.088	-0.015
	(0.539)	(0.543)	(0.509)
Control mean	2.636	1.074	0.067
Observations	1,601	1,601	1,601
Time F.E.	Y	Y	Y
Household F.E.	Y	Y	Y

Table notes: Regressions one and two are measured in manzanas (1.73 acres), while regression three is an indicator for whether or not you rent land to someone else, including formal and informal arrangements. p-values in parentheses are clustered using the wild cluster bootstrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

B.4 Land Use

B.5 Using "current storage" as a direct measure of stored crops

Table 12 shows storage levels using a direct measure of storage. The measure of storage used here is

$$\frac{\text{Current Quantity Stored in Household}}{\text{Total Quantity Harvested}}.$$

[16]We considered other cash crops as well, and find similar results to coffee.

This measure does not measure the total amount of harvest stored, as some was consumed prior to the survey wave. Nevertheless, the results are similar to those in the main text.

Table 12: Farm Savings Choices

	Fraction Corn Saved		Fraction Beans Saved	
	(1)	(2)	(3)	(4)
Build	-0.10*		-0.10*	
	(0.08)		(0.06)	
Build × Near		-0.12*		-0.12**
		(0.10)		(0.04)
Build × Far		-0.08		-0.08
		(0.38)		(0.24)
Far		-0.01		0.00
		(0.94)		(0.90)
Constant	0.85***	0.85***	0.90***	0.90***
	(0.00)	(0.00)	(0.00)	(0.00)
Observations	926	926	926	926

p-values in parentheses computed using the wild cluster bootstrap-t with 1000 simulations, clustered at the village level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$