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

February 2017

Preliminary Draft – Comments Welcome

Abstract

We estimate the impact of new bridges in rural Nicaraguan villages facing seasonal floods that unpredictably eliminate access to outside markets. We collect detailed annual household surveys over three years and conduct weekly telephone followups with a subset of households for sixty-four weeks, both before and after construction. This information is collected in villages where bridges are built, and in comparable villages where no bridge was built for only engineering-related reasons. We find that bridges eliminate uncertainty in market access driven by floods: in absence of a bridge, household income falls significantly during floods. Bridges completely eliminate this effect. Compared to households in villages where no bridges are built, bridges cause substantial reallocation of activities between farming and wage work. There are also significant effects on agricultural choices: increased fertilizer spending, increased yields on farms, and lower crop storage. We develop a model of occupational choice and risky farm investment, and show that these results are a rational response to reduced market income risk from the bridges. We provide evidence that these results are not due to lower trade costs outside of flooding periods, nor due to looser collateral constraints.

*This paper has benefited from comments by Nilesh Fernando, Manuel García-Santana, Lakshmi Iyer, Joe Kaboski, Tim Lee, Jack Rossbach and participants at the Notre Dame Development Lunch and the 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). Understanding the technical and policy constraints that generate this productivity gap is critical to both increase the welfare of the poor and better sectoral productivity differences. In this paper we focus on a constraint common in the developing world: the uncertainty of seasonal rains or monsoons.[1] Combined with relatively poor infrastructure in these places, extreme weather has the ability to isolate villages even further by washing out roads or rivers, and acts as a physical barrier to market access.

We confront this issue directly in rural Nicaragua, where it is considered a major development hurdle by both policymakers and citizens (World Bank, 2008*a*). Every year between May and October, flooding occurs unpredictably that cuts off access to outside food, product and labor markets. Along with a partner NGO, we build footbridges in villages that face this seasonal flood risk. We furthermore 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, an important consideration for understanding the complete effect of the intervention, given the heterogeneity in income-generating activities in rural areas (Foster and Rosenzweig, 2007).

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.[2] This is true in our context as well, where

[1]The issue has long been thought to be a contributor to poverty. See Kamarck (1973) for an early study related to its relation to agriculture and health issues in the tropics.

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

bridge costs approximately $40,000. As such, our study includes 500 households from only 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 decreases 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 zero income. We also confirm these results in the annual surveys. 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. The result is driven by two effects. First, individuals who had wage earnings at baseline shift from within-village wages to outside-village earnings and increase earnings. Second, those who had no earnings at baseline enter the market. Relative to those in villages without a bridge, those who receive a bridge have 39 percent higher earnings. Decomposing these results, new entrants account for two-thirds of the average effect. Taken together, the results show that (1) floods generate uncertain access to labor markets and (2) bridges eliminate this uncertainty.

Motivated by these results, we build a model in which increased labor market access can generate changes in not only market earnings, but also agricultural decisions and occupational choice. Households are infinitely lived, and can choose each period to be a worker, farmer, or both and are subject to idiosyncratic shocks in both sectors. Farm production is risky, in that decisions are made before shock realizations are realized. We parameterize the model using the control group in our data, and show that the model makes a number of predictions about how changes in market access affect occupational choice, agricultural decisions, and savings. First, farm investment in fertilizer and farm yields increase. This is driven by the fact that decreased uncertainty in labor market access acts as a consumption smoothing technology for farmers. Because farm investment is risky, farmers can turn to labor market work if farm shocks are low, thus lower *ex post* consumption risk. Second, they decrease savings. Again, because the labor market is a smoothing technology, it acts as a substitute for high-cost storage technologies used by rural households, and thus decreases the need for precautionary savings. Lastly, the model predicts that households should reallocate both from working to farming *and* from farming to working. The intuition from this result follows a similar logic to work on financial frictions (e.g. Buera, Kaboski and Shin, 2011; Midrigan and Xu, 2012). Because households are less constrained by risk, they put more weight on sector-specific skills when deciding their sector of work.

We then ask whether or not these model predictions are borne out in the data, both in terms of the sign of the effect and the quantitative magnitudes. We find support for both the mechanisms and magnitudes of the model results for all three predictions highlighted above. Of course, some of these results (in particular, higher fertilizer and yields) can be generated by many different models. We consider predictions from two other potential channels. The first is a standard trade theory, in which infrastructure development change agricultural decisions through changes in prices.[3] The short-term duration of the flood shocks make it unlikely that these price effects would occur here. While rivers flood quite often – during our high frequency survey, in one-third of weeks, at least one village was flooded – the floods are typically short. As long as crops are storable over a matter of weeks, price changes are unlikely. Indeed, we find no changes in prices in response to bridges. Second, we consider a theory of collateral constraints, in which households have to buy fertilizer before realizing harvest revenues, and thus may be constrained by insufficient income at planting. We show that this model has two predictions at odds with the data. First, savings should increase as households build a larger buffer stock. Second, the agricultural gains should only accrue to those who utilize the labor market to increase earnings. We test both and find no evidence for either.

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

[3]See Suri (2011) and Donaldson (2013), for example, along with the quantitative theories proposed in Adamopoulos (2011), Gollin and Rogerson (2014), and Van Leemput (2015).

trade costs and new construction (Alder, 2013; 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), and Van Leemput (2015). They find large gains from reducing the cost of movement across regions, as firms are able to more able to specialize in their comparative advantage. However, both the high costs and benefits of infrastructure make identification difficult, as infrastructure tends to be targeted toward areas with the largest benefits. Recently, a number of important papers have taken advantage of policy changes and natural experiments to identify the effects of infrastructure development, including Suri (2011) and Donaldson (2013). These papers primarily focus on the impact on prices, which are not present here. More closely related is Asher and Novosad (2016) who show that new roads in India generate movement out of agriculture using discontinuities in the policy design. 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.

We lastly find that the bridge increases on-farm productivity, but not through any direct changes to prices. Instead, the bridges allow for increased consumption smoothing through labor markets, which in turn endogenously promotes risk taking on farms. 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 lack of consistent labor market access limits agricultural productivity through this channel, which Restuccia, Yang and Zhu (2008) and Gollin, Lagakos and Waugh (2014) point out is substantially lower than nonagricultural productivity. 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 Village Selection

Our treatment directly confronts the theoretical predictions in the last section by building footbridges in rural Northern Nicaragua. 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 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. 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. As in the model, the time between planting and harvesting is also the time in which the flooding shocks occur.

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.

We therefore 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. Since 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.

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 are comparing 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.

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 Esteli and Matagalpa in northern Nicaragua.[4]

[4]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.

3 Data Collection and Design Validity

3.1 Data Collected

We conducted two types of data collection. 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 designed to give us an early indication of balance, and also to sign households up for the high frequency survey. 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 was primarily designed to (1) assess balance across groups and (2) register households for the phone surveys. 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) was having an effect on income generating activities. We therefore carried out these surveys for 64 weeks, covering the rainy season before construction, 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 [finish.]

Because the high frequency data was collected over the phone, two issues are worth highlighting before turning to the empirical results. First, this strategy implies that the high frequency data is not representative of the villages under study as not every individual has a cell phone. Second, it is an unbalanced panel of individuals as not everyone answered the phone each time.

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)

p-values in parentheses. Here, we do not cluster the standard errors as to give the model the greatest chance of finding a difference between the two groups.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

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 high frequency realizations of flooding and contemporaneous income realizations. In Section 4.3, 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. Before moving forward, two issues are worth highlighting. 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 of individuals as not everyone answered the phone each time. Figure 1 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 1: Number of Observations per Household

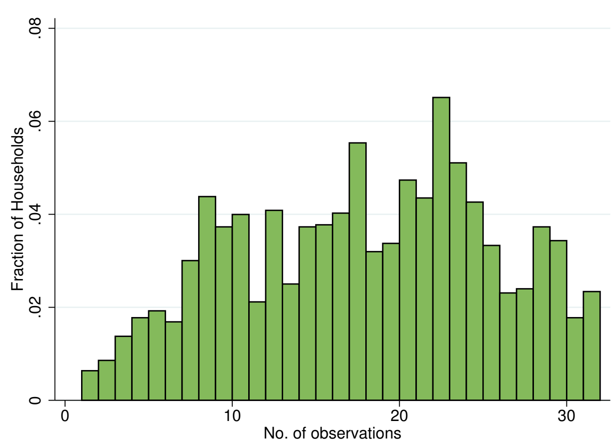

Despite these issues, the data still provides useful information on high frequency outcomes. To assess the impact of flooding on different outcomes, we run regressions of the form

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

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. First, labor income is ubiquitous even among farming households. Figure 2 is a histogram counting the share of weeks each household receives positive labor market income. Indeed, 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.[5]

Figure 2: Fraction of weeks with labor market income

(a) All households (b) Only households with ≥ 10 observations

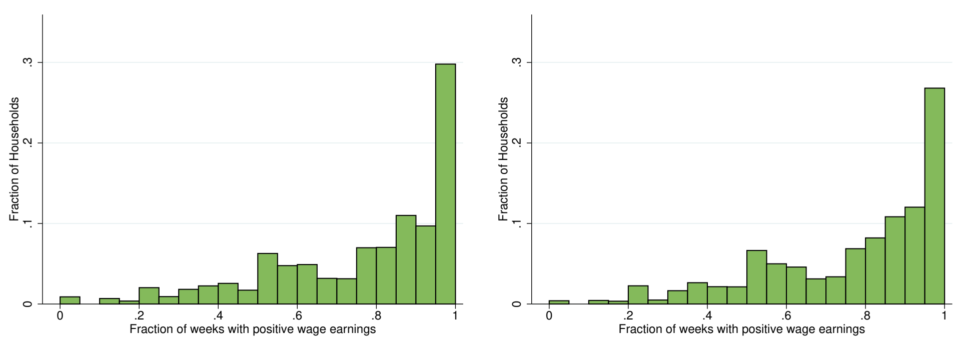

We therefore ask how income realizations change during flooding episodes, and how the bridge changes the results. We use two measures of income in regression (4.1): amount earned in the past two weeks and an indicator equal to one if no income was earned.

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

[5]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 A we show that there is no relationship between flooding and the likelihood of response to surveys.

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

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. This effect is not present in villages with a bridge. 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 3 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.

Figure 3: Density of Income Realizations

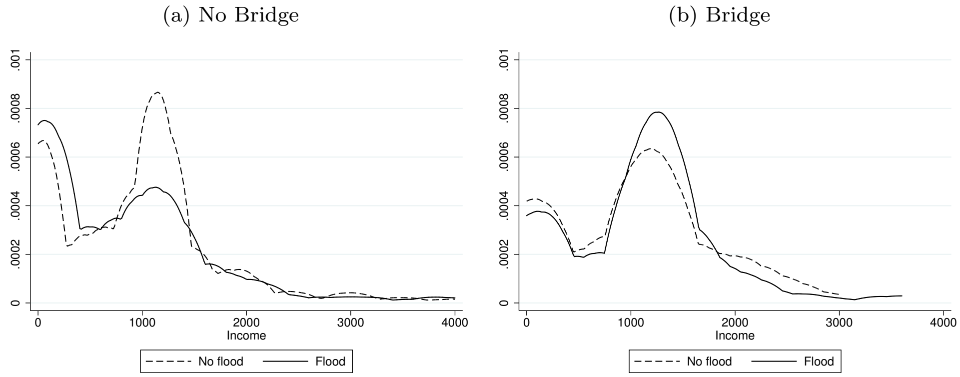

Figure notes: Figure 3a includes all village-weeks without a bridge, including those villages that eventually receive a bridge. Figure 3b includes all village-weeks post-construction.

4.2 Changes in Food Security

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 3 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 3: 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

4.3 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 $t = 0, 1, 2$ 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 individual fixed effects, and standard errors are clustered at the village level using a wild cluster bootstrap-t. Panel A of Table 4 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 4 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 4: Effects on Market Income, by Source

Panel A:	Total Earnings			Earnings Outside Village			Earnings Inside Village		
Earnings	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)
Baseline Average	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	Total Earnings			Earnings Outside Village			Earnings Inside Village		
Extensive Margins	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)
Baseline Average	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 boostrap-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. 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. We therefore build a model that highlights the potential relationship between labor market earnings, agricultural decisions, and occupational choice. We then highlight a number of predictions of the model, and use them to guide our analysis of further empirical predictions from the data.

5 Model

The model captures the key features of rural Nicaragua, including a mixture of farm and off-farm work and limited savings opportunities. We assume that there are two sources of income uncertainty in the model: wage risk (e.g. isolation through flooding) and farm risk (potentially also weather-related).

Time is discrete. There is a continuum of infinitely-lived households who maximize utility

$$\mathbb{E}_0\left[\sum_{t=0}^{\infty}\beta^t\frac{c^{1-\sigma}-1}{1-\sigma}\right]$$

with discount factor $\beta \in (0,1)$. Households do not have access to any state-contingent assets, but they can save b at exogenous gross rate R. They are endowed with one unit of time.

5.1 Occupational Choice

A household can choose to be a worker or a farmer. The ability of a household in each occupation is given by the vector $\mathbf{z} = (z_a, z_w)$, which varies stochastically according to the transition function $Q(\mathbf{z}_{t+1}, \mathbf{z}_t)$. Every household owns an agricultural technology that they can operate, and has the option to work in the labor market for a wage.

Their problem is represented recursively by the function:

$$v(\mathbf{z_{-1}}, b) = \max_{\phi \in \{0,1\}} \ \phi \mathbb{E}_{z|z_{-1}} \Big[v^a(\mathbf{z}, b) \Big] + (1 - \phi) \mathbb{E}_{z|z_{-1}} \Big[v^w(\mathbf{z}, b) \Big] \tag{5.1}$$

where $\phi \in \{0, 1\}$ is the choice to operate the agricultural technology or not. The value function v^a is the value of operating the technology, which requires payment of a fixed cost, and v^w is the value of not operating the technology and not paying the cost.[6]

5.2 Wage Work

A household that chooses to not operate the agricultural technology uses its entire time endowment on market work. Their earnings are therefore equal to z_w. In recursive form, this implies that the value of being a worker with shock $\mathbf{z}$ and savings b is:

$$\begin{aligned} v^w(\mathbf{z}, b) \quad &= \quad \max_{c, b'} \ u(c) + \beta v(\mathbf{z}, b') \\ s.t. \quad & \quad c + b' - Rb = z_w w \\ & \quad b' \geq 0 \end{aligned}$$

Note the continuation value v implies that the household can re-optimize its occupational choice each period given shock realizations and savings decisions.

5.3 Farming

If a household decides to farm, they are required to pay a fixed cost ψ. Before the shock is realized, farmers must choose the quantity of intermediate inputs (fertilizer and pesticide) on their farm, consistent with the importance of risk for this decision.[7] Once they choose intermediates, the shock is realized. After the shock, farmers can choose how to delineate their time between working on the farm and working in the market for wages, along with their consumption and savings decisions.[8]

This timing implies a two-step problem for farmers. After the shock is realized,

[6] We think of the fixed cost to operate in agriculture as capturing the sunk costs made at the beginning of the growing season, such as land rental. We model variable costs, such as fertilizer and seed, separately as discussed below.

[7] See the growing literature on rainfall insurance, such as Mobarak and Rosenzweig (2012) and Karlan et al. (2014), among many others.

[8] We abstract from hired labor. In our sample, 93 percent of farmers hires no labor.

the value of being a farmer with shock $\mathbf{z}$, savings b, and an intermediate choice x is given by

$$\begin{aligned}
\widetilde{v}^a(\mathbf{z}, b, x) &= \max_{c,n,b'} \; u(c) + \beta v(\mathbf{z}, b') \\
s.t. \quad & c + b' - Rb = z_a x^\theta n^\eta - p_x x + z_w w(1-n) - \psi \\
& n \in [0,1] \\
& b' \geq 0.
\end{aligned}$$

This defines decision rules as a function of x, $\widetilde{b}'(\mathbf{z}, b, x)$ and $\widetilde{n}(\mathbf{z}, b, x)$. Moving to before the shock, the value of entering the period as a farmer is given by

$$v(\mathbf{z}_{-1}, b) = \max_{x \geq 0} \; \mathbb{E}_{z|z_{-1}} \Big[\widetilde{v}^a(\mathbf{z}, b, x) \Big]$$

which defines the decision rule $x(\mathbf{z}_{-1}, b)$, and thus implicitly the decision rule $b'(\mathbf{z}_{-1}, \mathbf{z}, b) := \widetilde{b}'(\mathbf{z}, b, x(\mathbf{z}_{-1}, x))$.

5.4 Discussion

The discrete occupational choice makes an analytical characterization of the model difficult, but is required to guarantee that some households specialize in a single occupation. Therefore, we instead quantify the model and test a number of predictions against the data. Before doing so, we briefly digress to highlight the important channels that map the model to data. We view the exogenous weather that generates flooding as part of the governing force behind the wage shock w. That is, flooding risk increases the likelihood of a low wage realization. A bridge eliminates this risk, which acts as both an increase in the mean and decrease in the variance of the wage shock, which we can use the model to decompose. At the same time, rainfall may impact agricultural productivity. While this feature of the economy is not directly subject to change in response to a bridge, we will show how it responds to a change in the bridge.

5.5 Model Characterization

This model generates rich predictions for how households choose between farming and market labor, and how changes in the wage process that they face will affect their choices.

5.5.1 Allocation of Labor across Occupations

In the model there are no credit markets and no insurance. Therefore, since fixed costs and intermediate inputs must be chosen before agricultural productivity is known, the wealth of the household partly determines both whether or not they engage in agriculture, and how much intermediate input is used. A household with very low wealth does not engage in agriculture because the possibility of low agricultural productivity means that, lacking the ability to borrow, they would realize very low consumption in some states of the world. Therefore, that household would be more inclined to engage in market labor, which requires no sunk costs.

The fact that household asset levels and occupational choice have a relationship is inefficient in the sense that it reduces total output relative to an economy without borrowing constraints. That is, some very poor households would be more productive as farmers, but are not due to their borrowing constraint.

5.5.2 Reduction in Risk

Motivated by the results we discussed before, we interpret the bridge in the model as a reduction in the risk of the wage process for households. This has three effects in the model, which are explained in detail below: 1) it mitigates the relationship between asset levels and occupation choice, 2) farmers use more intermediate inputs, and 3) households save less.

Even when paying the fixed cost to operate the agricultural technology, households have the option of working in outside labor markets. Therefore, when the riskiness of wages falls, it provides better insurance against negative realizations of the agricultural shock. This causes the process for total income after having paid the fixed cost less risky than it was before. Because risk was the cause of the inefficient relationship between assets and farming status, this reduction in risk necessarily mitigates it. The model therefore predicts that some households should move from market work to farming because of the bridge.

By a similar argument, the model predicts that spending on intermediates rises when the bridge is built. Intermediate spending is a risky form of investment because it must be chosen before the agricultural shock is realized. When the income process becomes less risky, the household changes their portfolio of saving from the risk-free savings technology to the risky farming technology. Therefore, spending on agricultural intermediates should rise when the bridge is built.

Finally, households in this model primarily save to self-insure income risk because there are no state-contingent assets. Therefore, when there is less risk, they choose to save less.

We next test for these predictions in our data by looking for changes in these behaviors after bridges are built.

6 Testing Model Predictions Against Empirical Results

Our model generates predictions for how the presence of a bridge should affects other choices. In particular, it has important effects on savings. Because households face less risk, they put less of their income into risk-free savings and invest more in risky agriculture. Moreover, because the inherent risk in agriculture is less salient when labor markets are an effective mode of responding to risk, the relationship between household wealth and farming is weaker with the bridge than without it.

6.1 Occupational Choice and Misallocation

Table 4 shows that the impact on earnings is in large part driven by new entrants into the labor market. 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. While agricultural production and wage work are persistent over time, they are not perfectly so. Thirty percent of households with wage income at baseline have no wage income in periods one and two, and the same number holds true for farming. Moreover, twenty-five percent of households with no wage employment at baseline have some wage income at periods one or two, and similarly 29 percent engage in farming in periods one or two despite no farming at baseline. We therefore ask how this persistence changes after a bridge is constructed. In particular, we run the regressions

$$\begin{aligned}
\mathbb{1}[Wage]_{ivt} &= \alpha + \beta\Big(B_{vt} \times \mathbb{1}[Wage]_{iv0}\Big) + \gamma\Big(B_{vt} \times (1 - \mathbb{1}[Wage]_{iv0})\Big) + \eta_t + \delta_i + \varepsilon_{ivt} \\
\mathbb{1}[Agr]_{ivt} &= \alpha + \beta\Big(B_{vt} \times \mathbb{1}[Agr]_{iv0}\Big) + \gamma\Big(B_{vt} \times (1 - \mathbb{1}[Agr]_{iv0})\Big) + \eta_t + \delta_i + \varepsilon_{ivt}
\end{aligned}$$

where $\mathbb{1}[Wage]_{ivt} = 1$ if a household is engaged in labor market activities at time t, while $\mathbb{1}[Agr]_{ivt} = 1$ if the household is engaged in agricultural activities.[9] We wish to see whether baseline engagement in the sector predicts post-treatment engagement in the sector. The results are presented in Table 5. Interestingly, we find that the bridge makes households significantly less likely to engage in both agricultural and labor market activities. Among households not engaged in farming at baseline, those who receive a bridge are 18 percentage points more likely to begin farming ($p = 0.04$), and similarly, those not engaged on wage work are 29 percentage points more likely to being doing so ($p = 0.00$). Correspondingly, those engaged in both farming and wage work at baseline are less likely to remain so after the bridge is built.

To assess this in more detail, we decompose the occupational space into the four mutually exclusive groups – only agricultural production, only labor 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 6, and a number of results emerge. First, the bridge induces households to engage in market economic activity. For those currently engaged

[9] Note that these are not mutually exclusive, as households can be engaged in both or neither.

Table 5: Effects on Persistence of Activities

	Agriculture	Labor Market
	(1)	(2)
Build × Engaged	-0.315***	-0.113*
	(0.000)	(0.059)
Build × Not engaged	0.178**	0.286***
	(0.038)	(0.001)
Control mean	0.501	0.538
Observations	1,347	1,347
Time F.E.	Y	Y
Household F.E.	Y	Y

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

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 5, 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 6. For any current economic activity, a bridge makes it significantly less likely that a household is engaged in that same activity post-treatment. Again, this is consistent with misallocation across households.

6.2 On-Farm Impact

We first consider intermediate input use on farms, with results presented in Table 7. 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

Table 6: 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 mean	0.329	0.366	0.172	0.133
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: Engaged = 1 if the household is engaged in the relevant activity at baseline, and *Not engaged* is the converse. p-values in parentheses are clustered using the wild cluster boostrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

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 ($p = 0.01$) on a baseline of C\$934, and its components fertilizer and pesticide increase by C\$438 ($p = 0.01$) and C\$153 ($p = 0.29$) respectively. Interestingly, however, when we decompose the results based on baseline farming, we find that there are similar changes on both continuing and new farmers. Intermediate expenditures increase by C\$674 ($p = 0.13$) among continuing farmers and C\$614 ($p = 0.04$) among new farmers, relative to those in villages without a bridge. Similar results are found when considering fertilizer. Continuing farmers increase fertilizer expenditure by C\$465 ($p = 0.03$) and new farmers by C\$407 ($p = 0.03$). We do see some difference in changes in pesticide spending, however, where changes across treatment and control is primarily driven by new farmers. Note that these results are in contrast to changes labor market earnings in Table 4, which were primarily driven by new entrants to wage work. Thus, in addition to allowing individuals to better sort into their preferred occupation, it also allows farmers who do not switch to invest more on their farms.

Table 7: 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)
Baseline Average	934.25***	934.14***	598.10***	597.99***	344.31***	344.67***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
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 boostrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Therefore, we next consider changes in harvest for maize and beans, measured in total quintales (100 kilograms) harvested.[10] The results are in Table 8. 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). In Appendix A.3, we consider farm yields of the same crop, measured in quintales per manzana (approximately 1.73 acres). This eliminates a substantial fraction of observations, as those with no harvest have no land in production, and thus the notion of yield is undefined. Nevertheless, we find positive significant increases in yield in that regression.

6.3 Savings Decisions

Consistent with the model developed in Section 5, households have both higher off-farm income and invest more on their farms. The last prediction of the model is that households decrease the amount of crop stored in the household. This is the main form of savings in these rural villages, yet is a high cost savings technology. In the baseline

[10] In Appendix A.2, we show that the bride has no effect on crop selection by farmers, hence our focus directly on yields here.

Table 8: Harvest of Staple Crops

	Maize		Beans	
	(1)	(2)	(3)	(4)
Build	1.63		1.06	
	(0.234)		(0.123)	
Build × Farming		2.11		0.96
		(0.380)		(0.363)
Build × No farming		1.05		1.18
		(0.164)		(0.206)
Baseline Average	2.89***	2.88***	1.64***	1.64
	(0.000)	(0.000)	(0.000)	(0.000)
Observations	1,601	1,601	1,601	1,601
Time F.E.	Y	Y	Y	Y
Household F.E.	Y	N	Y	N

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 boostrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

pre-bridge data, 38 percent of households store their crops in plastic bags, 36 percent use plastic barrels, and 23 percent use small personal silos.[11] These technologies are prone to substantial spoilage and infestations (Grolleaud, 2002; Hodges et al., 2010), thus making storage a high cost method to move goods across time. The model predicts that access to the labor market should substitute for this savings technology, as households can more easily insure consumption through labor market earnings. We test this here.

To define storage, we asked first about the amount harvested of each crop. We then asked what part was sold, used to pay debt, gifted, or given as land payment. Storage is defined as harvest net of sales, debt payments, gifts, and land payments.[12] Any household with no crop production is given a value of zero in this regression. Table 9 shows how bridges affects savings behavior.

Consistent with the model, farmers store a significantly smaller proportion of their harvest. In build villages, farmers save 7 percentage points less of corn harvest than

[11] The remaining three percentage points are split between doing nothing and more complicated storage technologies, such as crop cellars.

[12] In the Appendix we present the results when we define storage as the amount of each crop currently held in the household. The results are quite similar. However, "amount currently stored" is net of any already-consumed harvest and thus is not the total measure of harvest stored. For this reason, we prefer the in-text measure of storage.

Table 9: Farm Savings Choices

	Maize		Beans	
	(1)	(2)	(3)	(4)
Build	-0.071**		-0.083**	
	(0.032)		(0.026)	
Build × Farming		-0.072		-0.101
		(0.138)		(0.121)
Build × No farming		-0.069**		-0.061**
		(0.038)		(0.027)
Baseline Average	0.942***	0.942***	0.928***	0.928***
	(0.000)	(0.000)	(0.000)	(0.000)
Observations	1,601	1,601	1,601	1,601
Time F.E.	Y	Y	Y	Y
Household F.E.	Y	N	Y	N

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 boostrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

those in non-build villages ($p = 0.03$), and 8 percentage points less ($p = 0.03$) of bean harvest. This affect is found in both continuing and new farmers, though the higher variance among continuing farmers implies that only new farmers have a significant effect from the bridge. Among new farmers, the bridge induces a 7 percentage point decrease in maize storage ($p = 0.04$) and a 6 percentage point decrease in bean storage ($p = 0.03$). Among continuing farmers, we find similar decreases of -0.07 ($p = 0.14$) and -0.10 ($p = 0.12$), and though we slightly miss statistical significance at the 10 percent level.

We also consider other household financial transactions. For example, households could substitute storage for debt if there was any associated change in (formal or informal) financial markets. We therefore consider debt positions in the household, including to banks, businesses or any other institutions or individuals. Table 10 includes evidence on household debt levels. Regressions 1-4 are outstanding debt levels in córdoba, while regressions 5-8 are indicators for any outstanding debt. Consistent with the increased wealth level of households that receive a bridge, we find a decrease in outstanding bank debt, both on the intensive and extensive margins. Bank debt decreases by C$573 ($p = 0.03$) among households in villages that receive a bridge, and they are 7.6 percentage points less likely to have debt owed to a formal financial institution ($p = 0.09$). Along other dimensions, including to local businesses or other debt, we find no change in either likelihood of outstanding debt or the amount.

Table 10: Household Debt

	Outstanding Debt				Any outstanding?			
	Total	Bank	Business	Other	Total	Bank	Business	Other
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Build	-788.18***	-573.08**	-41.65	-0.98	-0.090	-0.076*	-0.036	-0.029
	(0.005)	(0.025)	(0.538)	(0.916)	(0.292)	(0.088)	(0.614)	(0.743)
Baseline Average	1291.00***	936.63	285.91***	29.43***	0.413***	0.169***	0.297***	0.299***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	1347	1347	1347	1347	1347	1347	1347	1347
Time F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Household F.E.	Y	Y	Y	Y	Y	Y	Y	Y

Table notes: Total is the sum of bank, business, and other debt. p-values in parentheses are clustered using the wild cluster boostrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

7 Conclusion

We consider the impact of new footbridges in rural 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.

Alder, Simon. 2013. "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.

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.

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.

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.

Grolleaud, Michael. 2002. "Post-harvest losses: Discovering the full story." Food and Agriculture Organization of the United Nations, Food and Agriculture Organization of the United Nations.

Hodges, R.J., M. Bernard, H. Knipschild, and F. Rembold. 2010. "African Postharvest Losses Information System - a network for the estimation of cereal weight losses." In *Proceedings of the 10th International Working Conference On Stored Products Protection.* , ed. M.O. Carvalho, 958–964.

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.

Midrigan, Virgiliu, and Daniel Xu. 2012. "Finance and Misallocation: Evidence from Plant-Level Data." Working Paper.

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

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.

Suri, Tavneet. 2011. "Selection and comparative advantage in technology adoption." *Econometrica*, 79(1): 159–209.

Van Leemput, Eva. 2015. "A Passage to India: Quantifying Internal and External Barriers to Trade." Working Paper, Federal Reserve Board of Governors.

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 More Results and Robustness

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

Figure 2 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 11. 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 11: Effect of flooding on survey response

	(1)	(2)
Flood	0.026	-0.025**
	(0.151)	(0.035)
Baseline Average	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 boostrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

To further emphasize this point, Figure 4 reproduces Figure 2 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.

Figure 4: Fraction of weeks with labor market income

(a) All households (b) Only households with ≥ 10 observations

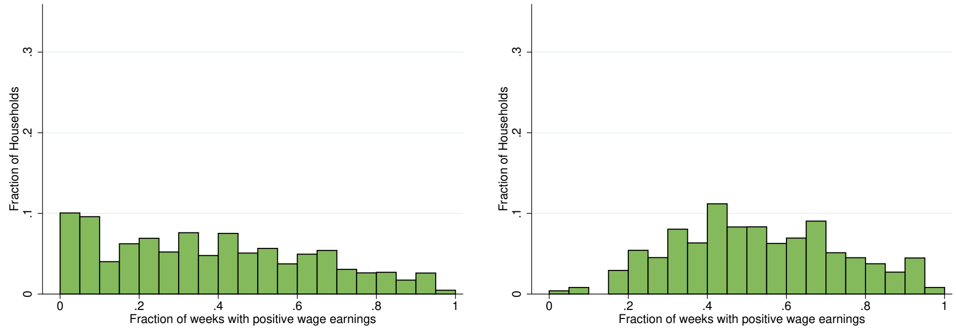

Naturally, this shifts the distribution toward zero. However, even when considering all households, the fifth percentile household still receives labor market income in 3 percent of its observations. The median household receives labor market income in 36 percent of weeks. Thus, individuals are still 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.

A.2 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.[13] The outcome variable here is an indicator equal to one if the crop is planted (not necessarily harvested), and the results are in Table 12.

Table 12: 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)
Baseline Average	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: Farming = 1 if the household is engaged in any crop production at baseline. p-values in parentheses are clustered using the wild cluster boostrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 13: Harvest of Staple Crops

	Maize		Beans	
	(1)	(2)	(3)	(4)
Build	13.47**		3.26***	
	(0.019)		(0.003)	
Build × Farming		15.07**		2.89**
		(0.018)		(0.022)
Build × No farming		6.27		6.48
		(0.112)		(0.181)
Baseline Average	9.10***	9.28***	4.36***	4.22***
	(0.000)	(0.000)	(0.000)	(0.000)
Observations	313	313	324	324
Time F.E.	Y	Y	Y	Y
Household F.E.	Y	N	Y	N

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 boostrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

A.3 Farm Yields in Response to Bridge

A.4 Using "current storage" as a direct measure of stored crops

Table 14 shows storage levels using a direct measure of storage. The measure of storage used here is

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

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 14: 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)
Baseline Average	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.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

B A model of collateral constraints

In this section, we show that the empirical results are inconsistent with a model of collateral constraints. The model is exactly the same as that provided in the text,

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

except for the farming decision problem which is detailed below.

If a household decides to farm, they are required to pay a fixed cost ψ. Unlike the model in the text, farmers know their shock before choosing intermediate inputs. However, they are subject to a collateral constraint $x \leq \lambda b$. The idea behind this constraint is that households purchase fertilizer and other intermediate inputs at the planting stage, which revenues unrealized until harvest. Thus, if they are not holding enough savings, they cannot purchase enough. As in the text, they can also choose how to delineate their time between working on the farm and for wages.

The value of being a farmer with shock $\mathbf{z}$, savings b is therefore given by

$$\begin{aligned}
v^a(\mathbf{z}, b) \quad &= \quad \max_{c,n,x,b'} \; u(c) + \beta v(\mathbf{z}, b') \\
s.t. \qquad & c + b' - Rb = z_a x^\theta n^\eta - p_x x + z_w w(1-n) - \psi \\
& n \in [0, 1] \\
& b' \geq 0 \\
& x \leq \lambda b
\end{aligned}$$

The rest of the model is identical to the model in the text. We provide two results from this model that distinguish it from the model in the text. First, increases in fertilizer use should only accrue to those who utilize the labor market, and second, savings should increase.[14]

[14]In the model, average yields increase as well only for those who utilize the labor market. However, there are no aggregate shocks in the model, so this result is inherently less convincing a test compared to fertilizer spending.