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

Wyatt Brooks
University of Notre Dame

Kevin Donovan
Yale University

March 2020

Abstract

We measure the impact of increasing integration between rural villages and outside labor markets. Seasonal flash floods cause exogenous and unpredictable loss of market access. We study the impact of new bridges that eliminate this risk. Identification exploits variation in riverbank characteristics that preclude bridge construction in some villages, despite similar need. We collect detailed annual household surveys over three years, and weekly telephone followups to study contemporaneous effects of flooding. Floods decrease labor market income by 18 percent when no bridge is present. Bridges eliminate this effect. The indirect effects on labor market choice, farm investment, and savings are quantitatively important and consistent with the predictions of a general equilibrium model in which farm investment is risky, and households manage labor market risk and agricultural risk simultaneously. In the calibrated model, the increase in consumption-equivalent welfare is substantially larger than the increase in income due to the ability to mitigate risk.

JEL Classification Codes: O12, O13, O18, J43

Keywords: risk, farming, rural labor markets, infrastructure, smoothing

*Thanks to conference and seminar participants at Arizona State, Bristol, Edinburgh, Illinois, Notre Dame, Stony Brook, Toulouse, the Universitat Autònoma de Barcelona, Yale, BREAD, the Chicago Fed Development Workshop, the NBER Conference on "Understanding Productivity Growth in Agriculture," the SITE Development Workshop (Stanford), the Society for Economic Dynamics (Edinburgh), the Workshop on Macro-Development (London), and the Yale Agricultural Productivity and Development Conference, including Bill Evans, Manuel García-Santana, Andreas Hagemann, Lakshmi Iyer, Seema Jayachandran, Joe Kaboski, Tim Lee, Jeremy Magruder, Mushfiq Mobarak, Pau Pujolàs, Mark Rosenzweig, Nick Ryan, and especially Kelsey Jack for an insightful discussion of the paper. Thanks also to Avery Bang, Maria Gibbs, Brandon Johnson, and Abbie Noriega 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. This research was carried out under University of Notre Dame IRB protocol 16-09-3367. *Brooks:* 3060 Jenkins Nanovic Halls, Notre Dame, IN 46556, wbrooks@nd.edu; *Donovan:* Yale School of Management, 3475 Evans Hall, New Haven, CT 06511, kevin.donovan@yale.edu

1 Introduction

The majority of households in the developing world live in rural areas where labor markets are poorly integrated across space and productivity is particularly low (Gollin, Lagakos, and Waugh, 2014). Increased integration has potentially large benefits in rural areas where household income is derived from both farming and labor markets, a common feature of income-generating activities in the developing world (Foster and Rosenzweig, 2007).[1] Thus, understanding spillovers between between wage work and farm decisions are necessary to understand the full effect of labor market integration.

In this paper, we directly study the impact of integrating rural Nicaraguan villages with outside labor markets and show empirically that it has sizable effects on household wage earnings, farm investment decisions, and savings. We use seasonal flash floods as an unpredictable, exogenous, and observable source of variation in market access.[2] We then work with an NGO that builds footbridges that connect villages to markets, eliminating this uncertainty of market access. We conduct household-level surveys the year before the bridges are constructed and for two years after. In addition, we collect 64 weeks of data from a subset of households during the same period to understand the contemporaneous impact of flooding on household outcomes. As these rural households have many income streams with interrelated outcomes, we directly focus on how multiple margins are affected by improved outside market access, such as labor market outcomes and agricultural production choices.

Our identification strategy is based on the fact that many villages need bridges, but construction is infeasible for some villages due to the characteristics of the riverbeds that they aim to cross. Because these rivers are typically distant from the houses and farmland of the village (the average village household is 1.5 kilometers from the potential bridge site), the failure to pass the engineering assessment is plausibly orthogonal to any relevant household or village characteristics. We verify this by showing

[1] The direct effect is access to higher wages outside their village (Bryan, Chowdhury, and Mobarak, 2014; Bryan and Morten, 2018). However, to the extent that wage income allows farmers to relax credit constraints or better manage risk, it may simultaneously decrease farm-level distortions. Mobarak and Rosenzweig (2014) and Karlan, Osei, Osei-Akoto, and Udry (2014), among others, find benefits from formal rainfall insurance, while Jayachandran (2006) and Fink, Jack, and Masiye (2017) show how missing credit markets affect agricultural employment and production. Thus, these margins potentially play an important role.

[2] In addition to its benefits as a source of variation, flooding is a common phenomenon in the developing world and widely cited as a major development hurdle. This is true both of international policy organizations and citizens of Nicaragua (World Bank, 2008). More broadly, seasonal flooding or monsoons in the tropics have long been discussed as a contributor to poverty. See Kamarck (1973) for an early study on agriculture and health issues in the tropics.

that baseline characteristics are balanced across villages that do and do not fulfill the engineering requirements, which we detail in Section 2.

Our results imply that uncertain market access is an important constraint to both labor market access and agricultural productivity, and we find economically and statistically significant effects on both. In the absence of a bridge, floods depress contemporaneous weekly labor market earnings by 18 percent and increase the probability of reporting no income. When a bridge is constructed, both of these effects disappear. Floods therefore generate uncertain access to labor markets, and a bridge eliminates this uncertainty. We also find that labor market income increases in non-flood periods once a bridge is constructed. This is driven by the fact that men shift their time from relatively low paying jobs in the village to higher paying jobs outside the village, while new women enter the outside-village labor force. Moreover, those who stay in the village for work benefit from the general equilibrium increase in wage as village labor supply declines. This result is consistent with Mobarak and Rosenzweig (2014) and Akram, Chowdhury, and Mobarak (2016), who find changes in wages in response to increased agricultural investment and rural emigration respectively.

Finally, we find benefits on the farm as well. Farmers spend nearly 60 percent more on intermediate inputs (fertilizer and pesticide) in response to a bridge, while farm profit increases by 75 percent. One explanation is that a bridge makes it easier to purchase inputs or get crops to market for sale.[3] As we discuss in Section 2, the timing and duration of floods and ease of crossing during non-flood periods make this an unlikely source of the empirical results.

We therefore build a model to investigate a different potential mechanism linking high-frequency variation in labor market access to agricultural outcomes, taking seriously the various spillovers highlighted in our empirical results. Though adopted to the specifics of our setting, our model shares features with the literature focused on understanding firm investment decisions with missing markets.[4] The key distortion in the model is that farmers limit investment to hold a buffer stock of savings that guards

[3]This idea underlies standard theories of internal trade barriers between urban and rural areas, such as Adamopoulos (2011), Gollin and Rogerson (2014), Sotelo (2016), and Van Leemput (2016), along with most work studying the impact of new infrastructure explicitly (Donaldson, 2013; Asturias, Garcia-Santana, and Ramos, 2016; Alder, 2017, among others).

[4]See Bencivenga and Smith (1991), Acemoglu and Zilibotti (1997), and Angeletos (2007), along with Donovan (2018) for an agriculture-specific application. These theories also provide the theoretical basis in favor of formal agricultural insurance markets (e.g. Mobarak and Rosenzweig, 2014; Karlan, Osei, Osei-Akoto, and Udry, 2014).

against unforeseen shocks before investment pays off. Once a bridge is constructed, however, households have the ability to additionally smooth consumption via wage earnings in outside labor markets. Thus, resources previously held as a buffer stock are unlocked for productive investment. A result of this intervention is that savings declines as farmers redirect resources toward investment. We test this empirically and find support. Agricultural storage declines from 90 to 80 percent of harvest in the data in response to a bridge. Moreover, there is a strong negative correlation between changes in fertilizer expenditures and crop storage among treated households. That is, those increasing fertilizer expenditure are also the same households that are saving less in response to the bridge.

Since the model can theoretically match our empirical predictions, we summarize the total benefit of a bridge by calibrating the model and estimating the increase in consumption-equivalent welfare. We calibrate the model to match certain treatment moments, including the change in fertilizer expenditures. Moreover we find that the model matches untargeted treatment moments including the increase in village wages and decline in crop storage. Our results imply that the welfare benefits from bridges are large, as the introduction of a bridge increases consumption-equivalent welfare by 11 percent. We further show that the aggregate welfare change admits a simple accounting decomposition in which we can separately measure the benefit of higher average consumption, lower volatility of consumption, and changes along the transition between the two steady states. We find that higher average and less volatile consumption both play critical roles, which each accounting for half of the total impact. The transition path, correspondingly, plays little role.[5]

Furthermore, we note that the bridge changes both the first and second moment of the shock distribution by eliminating the tail of the market access shock. We use the model to study the relative importance of these two changes. Specifically, we introduce two counterfactual shock processes that change only the mean and only the variance of the outside earnings process. We find that the change to the mean play the quantitatively dominant role, but that the change in the variance is non-trivial

[5] Allen and Atkin (2016) also highlight the importance of second-moment variation when attempting to fully account for the impact of infrastructure, though they focus exclusively on the ability to more cheaply ship goods. Our intervention specifically targets the ability to move people more easily across space, while minimizing the direct effect on goods trade. Yet even in our context, we find economically important effects and provide another potential margin through which infrastructure benefits rural populations. Like their work, however, taking endogenous farmer responses into account is critical for properly capturing the gains from increased market access.

contributor to overall welfare gains.

Finally, we note that a major barrier to studying transportation infrastructure as an intervention is the high cost of construction, which typically limits the ability to identify the underlying mechanisms driving changes or the scope of outcomes considered. In our context, each bridge costs $40,000 because of high engineering standards required to survive powerful floods. Because of the high cost, our study includes household-level data from only 15 villages. Our ability to detect statistically significant effects is a function of the intensity of the treatment and low intra-cluster correlations, which average 0.06 among our outcomes. We correct our inference for the small number of clusters by using the wild bootstrap cluster-t procedure (Cameron, Gelbach, and Miller, 2008) throughout the paper and provide a number of robustness checks in the Appendix on both the inference procedure and regression specifications.[6]

2 Background

2.1 Flooding Risk

Over 40 percent of people affected by disasters worldwide since 2000 were affected by flooding. Of that, nearly all are due to river floods (EM-DAT, 2017). In Nicaragua, both policy makers and residents cite flooding and the resulting isolation as a critical development constraint (World Bank, 2008). The villages in our sample are located in mountainous areas that face annual seasonal flooding during the rainy season between May and November. This overlaps with the main cropping season as crops are planted in late May and harvested in November.

During the rainy season, floods cause stream and riverbeds that are usually passable on foot to rise rapidly and stay high for days or weeks. This flooding is unpredictable in its timing or intensity. Rainfall in the same location is not necessarily a good 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). During the baseline rainy season, the average village is flooded for at least one day in 45 percent of the two-week periods we observe it. The

[6] The sample size also implies our results are almost certainly an incomplete accounting of the aggregate impact of scaling such an intervention, as the villages are roughly 1 percent of the market to which they are connected. In principal, connecting all villages would have important effects in the receiving market. See Dinkelman (2011), Asher and Novosad (2016), and Shamdasani (2016) for results connecting infrastructure to structural transformation.

average flood lasts for 5 days, but ranges from less than one to 9 days (the ninetieth percentile). On average, this implies that a village is flooded for 2.25 days every two weeks.

During these periods, villages are cut off from access to outside markets. However, it is important to emphasize a number of features of this flooding risk that are relevant for interpreting our results. First, floods are intense torrents of water from the mountains, not simply villages situated next to rivers. Thus, crossing the river by swimming, or any other method, entails substantial risk of injury or death.[7] These floods usually generate prohibitively dangerous crossing conditions or a long journey on foot to reach the market by another route. For our purposes, we interpret a flood as a substantial increase in the cost of reaching outside markets.

Second, it is unlikely that the flooding has any direct effect on village farms, as the average household is nearly a mile (1.5 kilometers) from the river.[8] Moreover, given the relatively low spatial correlation between local rainfall and flooding, it is simultaneously unlikely to affect the urban labor market associated with our study villages.

Finally, these rivers are easily crossable when not flooded, and usually contain little to no standing water. Moreover, these villages are not located on deep ravines that make crossing difficult during dry times. This is important for the interpretation of our results, and contrasts this context from standard issues around transportation infrastructure that is used to generate a constant reduction in transportation costs, as in recent work by Adamopoulos (2011), Gollin and Rogerson (2014), and Sotelo (2016).[9]

2.2 Local Context

Crop Cultivation Our study takes place in the provinces of Estelí and Matagalpa in northern Nicaragua. The main cropping season coincides with the rainy season, with planting occurring at the beginning of the rainy season and harvesting happening after

[7] We are aware of at least two people (one on horseback) in our sample that died trying to cross flooded rivers during the last survey wave.

[8] These floods are torrents of water that rush through well-defined riverbeds. Thus, any household that located within it would likely be destroyed during a flood. As we will discuss below, the NGO we work with requires a well-defined high water mark to construct a bridge. Thus, this is in part of function of their selection procedure.

[9] We find no evidence of effects on the prices of goods, which confirms that those channels are inoperative in our context.

it ends. In relation to the discussion above, flooding is therefore unlikely to physically prohibit farmers from access to fertilizer or taking harvest to market.

At baseline, 51 percent of households farm some crop. Of those households, 47 percent grow beans and 41 percent grow maize. The next most prevalent crop is sorghum (8 percent). The key cash crops in the region are tobacco and coffee, as Northern Nicaragua climate and geography are well suited for both. However, tobacco and coffee are almost exclusively confined to large plantations. Only 3 percent of households in our sample grow coffee at baseline, while less than one percent grow tobacco. As we discuss below, coffee and tobacco jobs (picking, sorting, etc.) are an important source of off-farm wage work. The modal use of staple crop harvest is home consumption. Over 90 percent of maize and bean harvest is either consumed immediately or stored for future household consumption. The majority of those who sold crops either sell in the outside market (58 percent) or to middlemen who buy in the village and export to other markets (38 percent). Only 4 percent sell to local stores in the same village.

Fertilizer is used by 73 percent of all farming households. While for a developing country this is a relatively high prevalence of fertilizer, fertilizer expenditures are only 16 percent of total harvest value. This share is not quite as low as the poorest African countries, but substantially lower than developed countries (Restuccia, Yang, and Zhu, 2008).

The Labor Market We use bi-weekly data collected from households in our sample to show that nearly all households receive labor market income at some point (we discuss data collection in Section 3). Despite the fact that 51 percent of households farm at baseline, 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.[10] Households are almost never entirely specialized in farming, suggesting potential for a relationship between the labor market and on-farm outcomes, which we study in later sections.

[10]The online appendix provides a detailed distribution of labor market earnings across households. Furthermore, we note that this is a cell phone-based survey. Therefore, 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. We further 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. Therefore, the results are robust to even the most conservative possible assumptions on response rates.

Jobs held by village members are made up of those inside the villages (62 percent) and those employed in the outside markets (38 percent). The latter are at risk of being inaccessible during a flood. Connected markets have between 10,000 and 20,000 people, compared to 150 to 400 people in the small villages we study, so these villages make up only a small fraction of the labor supplied outside the village. Outside-village jobs also pay more on average. There is a 30 percent daily wage premium for men outside the village and an even larger 70 percent daily wage premium for women, though women are employed at a much lower rate.

In both cases, jobs are primarily on short term contracts. At baseline, 80 percent of primary jobs held were on short-term (less than one week) contracts. This differs somewhat depending on job location. In the village labor market, 90 percent of all jobs held are short-term, while outside the village 64 percent of jobs are short-term. In terms of occupations, farm labor makes up 61 and 41 percent of all wage employment inside and outside the village. Inside the village this work primarily consists of laboring on other farms, while outside the village this involves work on large coffee and tobacco plantations. Workers in outside markets cross the riverbed to reach the market town where trucks pick up workers to bring them to work. Workers are then dropped off at the same location at the end of the day. Thus, the market towns are important staging points for this work. Outside of farm work, village residents are employed as carpenters, teachers, maids, among other various occupations, at a substantially lower rate.[11]

3 Intervention, Data Collection, and Identification Strategy

3.1 Intervention

The bridges we build traverse potentially flooded riverbeds, thus allowing village members consistent access to outside markets. We partner with Bridges to Prosperity (B2P), a non-governmental organization that specializes in building bridges in rural communities around the world. B2P provides engineering design, construction materials, and skilled labor to the village. Bridges are designed by a lab of civil engineers in the United States in consultation with local field coordinators, who are also engineers.

[11] Details of occupations can be found in our included appendix files.

Bridges cannot be crossed by cars, but can support horses, livestock, and motorcycles. A bridge that can survive multiple rainy season requires durable, expensive materials and a sufficiently sophisticated design to overcome issues of rising water levels, soil erosion, and other risks that face infrastructure.

B2P 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. This assessment is made based on the number of people that live in the village, the likelihood that the bridge would be used, proximity to outside markets and available alternatives.

If the village passes the needs assessment, the country manager conducts an engineering assessment. The purpose of this assessment is to determine if a bridge can be built at the proposed site that would be capable of withstanding a flash flood. To be considered feasible, the required bridge cannot exceed a maximum span of 100 meters, and the crests of the riverbed on each side must be of similar height (a differential not exceeding 3 meters). Moreover, evidence of soil erosion is used to estimate water height during a flood. The estimated high water mark must be at least two meters below the proposed bridge deck.

We compare villages 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, the characteristics of the riverbed are unlikely to be correlated with any relevant village characteristics. We show that villages that do and do not receive bridges are balanced on their observable characteristics in Table 2.

Because a bridge costs \$40,000, the number of bridges that can be funded is limited.[12] 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.[13]

[12] We discuss cost-effectiveness in the online appendix. The internal rate of return to the bridge is 19 percent.

[13] The villages are far from one another, so there is no risk that the households in a control village could use the bridge in a treatment village.

Comparison to Other Nicaraguan Villages Our research design focuses on internally valid comparisons, but to think about external validity it is useful to see how these villages compare to the set of all rural Nicaraguan villages. We use data from the 2001 Nicaraguan Living Standard Measurement Survey (LSMS) and look for household characteristics that we can compare to our own data. These comparisons appear in Table 1. While we caution reading too much into characteristics derived from data sets 13 years apart (the 2001 LSMS is the most recent available), we compare the baseline values of several household characteristics from our sample to two categories: households from across rural Nicaragua, and households in the two departments where our study takes place. The regions of our study have lower mean earnings than in rural Nicaragua overall, according to the LSMS. Our households earn somewhat more than their regional counterparts in the 2001 LSMS.[14]

3.2 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 only to collect GPS coordinates from households and sign them up for the high frequency survey. The data used in our analysis comes from surveys conducted at the end of the main rainy season, in November 2014, November 2015, and November 2016. Bridges were constructed in early 2015. Therefore we have surveys from three years for all villages. For those that receive a bridge, we observe one survey without a bridge and two surveys with a bridge. We refer to these survey waves as $t = 0, 1, 2$.

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 all village households. The number of households identified in each village varied widely, from a maximum of 80 to a minimum of 24, with an average household size of 4.2. Participation in the first round of the survey was very high in general, with 97 percent 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

[14] Nicaraguan real GDP per capita grew by 35 percent from 2001 to 2014 in Nicarauga, which likely affects this comparison.

disclose our interest in the bridges because we suspected this would bias their answers, or may make them feel they are compelled to answer the survey when they would not otherwise choose to participate.

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

The second component of our data is biweekly follow-up surveys conducted by phone with a subset of households. Because floods are high frequency and short term events, this data shows the contemporaneous effect that flooding has on households. We carried out these surveys for 64 weeks, covering the rainy season before construction, along with the first dry and rainy seasons after construction. Each household was called every two weeks and asked questions about the previous two weeks, 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.

3.3 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. Identification requires that the features required to pass the feasibility test are independent of any relevant household or village-level statistics. To test that these villages are comparable, we run the regression

$$y_{iv} = \alpha + \beta B_v + \varepsilon_{iv}$$

on the baseline data, where $B_v = 1$ if village v gets a bridge between $t = 0$ and $t = 1$. We consider a number of different outcomes, and show that households show no observable differences across the two groups. Table 2 produces the results, and we find no difference across households in build and no-build villages.

3.4 High Frequency Sample Selection

Because the high frequency data was collected by phone, two issues are worth highlighting before turning to the results. First, the high frequency data is not representative of the villages under study as not every individual has a cell phone. In the online appendix, we show how high frequency respondents compare to the overall populations in the study. As one may suspect with a cell phone-based survey, household characteristics differ slightly between those who participate and those who do not, as respondents tend to be younger and slightly more educated. However, along dimensions such as as wage income and farming outcomes, both groups look similar. Importantly, within the high frequency sample, those in villages that receive a bridge and those that do not have similar characteristics.

4 Empirical Results on Labor Market Earnings

We begin by showing that labor market earnings respond positively to the introduction of a bridge.

4.1 Labor Market Earnings and Floods

We first estimate the relationship between floods and labor market earnings. In the high frequency data, we observe how realized labor earnings depend on contemporaneous flooding in villages. We interact an indicator variable for a bridge being present with flooding to estimate how the relationship between income and flooding changes once the bridge is built. We include household and time fixed effects to control for constant characteristics of households, and for seasonal variation in earnings. Our empirical specification in the high frequency data is:

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

The variable $B_{vt} = 1$ if village v has a bridge in week t, while $F_{vt} = 1$ if village v is flooded at week t. η_t and δ_i are week and individual fixed effects. P-values are computed using the wild bootstrap cluster-t where clustering occurs at the village level. We use two measures of income in regression (4.1): earnings in the past two weeks, and an indicator equal to one if no income was earned. Table 3 illustrates the effects of flooding on contemporaneous income realizations.

When bridges are absent, flooding has a strong effect on labor market outcomes. The decline in labor market earnings is $C\$143.5$ ($p = 0.030$), which is 18 percent of mean earnings.[15] Moreover, the propensity to earn no labor market income increases by 7 percentage points ($p = 0.042$) from a mean of 24.9 percent. However, when a bridge is built the effect on income disappears, and on net, a flood generates a economically insignificant change in income of $C\$5.1$. Similar results arise when we consider the likelihood of reporting no income. Figure 1 shows this is not an artifact of our specification, and plots the density of (raw) income realizations in villages without a bridge (left panel) and with a bridge (right panel) during periods of flooding and no flooding.

Figure 1: Density of Income Realizations

(a) No Bridge (b) Bridge

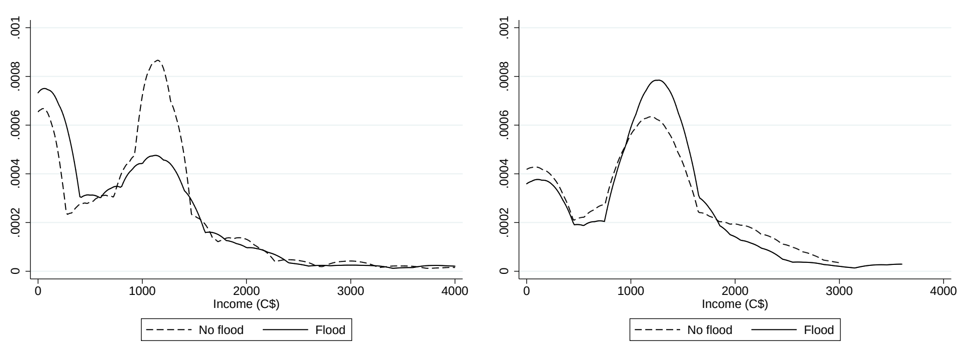

Figure notes: Figure 1a includes all village-weeks without a bridge, including those villages that eventually receive a bridge. Figure 1b includes all village-weeks post-construction.

Finally, it is notable that bridges increase income even in the absence of the flood. That is, during a non-flooded week, villagers with a bridge earn an average of C$160 ($p = 0.000$) more. We explore the cause of this finding in depth using the detailed annual surveys in Section 4.2 and find that a bridge causes workers to switch to jobs outside the village. The income gains, therefore, extend beyond just flooding periods. The bridges both smooth income during flood shocks and increases the average income level of households.

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

4.1.1 Do households substitute intertemporally?

If a household cannot access the labor market in a given week, they can potentially recoup their lost earnings by increasing earnings in the next (un-flooded) week. We test this by including lags in regression (4.1), with results in columns (2) and (4) of Table 3.

Column (2) shows that the results are inconsistent with control villages responding to floods by increasing future earnings. A flood two weeks in the past implies a statistically insignificant C\$17 decrease ($p = 0.698$), suggesting control households are not responding to past floods with increased current labor market earnings. Column (4) presents a similar result using an indicator for no income earned as the dependent variable. The returns among treatment villagers are consistent with the same theory. Households actually earn C\$126 less ($p = 0.150$) when they were flooded two weeks before, though it is not statistically significant. If anything, these results are consistent with the ability of the *treatment* villages to better adjust to shocks through utilization of the labor market.[16]

4.2 Earnings from Annual Surveys

In the previous sections, we showed that bridges eliminate labor market income risk during floods and also provide a benefit in non-flood periods. We next use our annual surveys to better understand these results. These surveys were conducted at the end of the rainy season from 2014 to 2016 ($t = 0, 1, 2$). Our baseline regression specification is

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

where $B_{vt} = 1$ if a bridge is built, η_t and δ_v are year and village fixed effects. Throughout, we use the wild bootstrap cluster-t at the village level.[17] The results are in Table 4, where we consider total earnings, and also break down the results by gender. Consistent with the previous results, labor market earnings increase by C\$380 ($p = 0.096$). This is almost entirely accounted for by the C\$306 increase in outside

[16] Theoretically, households need not intertemporally adjust this way. This would be true, for example, if on-farm productivity shocks are highly correlated with non-farm labor productivity shocks. In this case, the marginal product of on-farm labor would be high at exactly the time at which control households would wish to increase off-farm labor, thus dampening any effect. Anticipating the model, we allow endogenous responses of this sort.

[17] See the online appendix for further discussion of robustness. The results are robust to both the inclusion of household fixed effects and alternative inference procedures.

earnings ($p = 0.000$). Inside earnings decrease slightly (C\$27.70), but the change is statistically insignificant ($p = 0.842$). The same results hold when one distinguishes by gender. Columns 4 and 7 show that both men and women earn more, and these increases are entirely accounted for by earnings outside the village. For both genders, earnings inside the village decrease slightly, but both treatment effects are statistically insignificant.

We use the detailed employment information in the annual surveys to shed light on the mechanisms that generate these changes in earnings. Table 5 decomposes earnings by the number of household members, daily wages, and days worked. Men shift employment from inside to outside labor market work. In the average household, the number of males working outside increases by 0.19 ($p = 0.000$), compared to a 0.12 person decrease ($p = 0.130$) inside the village. Combined they generate a statistically insignificant net change in the number of males employed. Next, we find that male daily wages inside the village increase by C\$69 ($p = 0.102$), consistent with general equilibrium effects resulting from the decreased labor supply induced by the bridge. The male wages outside the village do not change (-C\$5.6, $p = 0.828$) because these villages account for a small fraction of labor market activity outside their village. The wage gap between inside-village and outside-village employment, therefore, converges for men.[18] Lastly, despite men moving to work outside the village, the number of in-village male days worked in the average household changes by an insignificant amount (-0.30, $p = 0.360$). Thus, those who remain in the village work more intensely at the higher wage. This implies an important spillover effect: even those who do not directly take advantage of the bridge still receive benefits in terms of higher in-village wages.

Panel B of Table 5 shows the results for women. The change in total household days worked mirror those for men. Days worked outside the village increase by 0.59 ($p = 0.002$) while number of days worked in the village do not change (-0.07, $p = 0.484$). However, the underlying mechanisms for this change are different. Instead of shifting job locations, we see a substantial increase in labor force participation. The average household increases the number of women employed for wages by 0.11 people ($p = 0.018$) over the baseline average of 0.17. This result is entirely due to entry in

[18] See the online appendix for the raw data showing the path of average daily wages over time in treatment and control villages.

the outside labor market. The number of employed women nearly doubles outside the village (from 0.12 to 0.22, $p = 0.000$) while there is no change in village employment. Consistent with this, we find no statistically significant changes in wages either inside or outside the village for women. Thus, while the bridge causes men to change where they work, it induces new women into labor market activity.

Both sets of results provide an explanation for the results of the high frequency surveys, namely, that bridges increase labor market earnings in non-flood weeks. As these more detailed results show, both men and women take up jobs outside the village. While the bridge increases access to the market during flood weeks, it also provides an opportunity to access jobs that pay more during non-flood weeks as well.

5 Impacts on Agricultural Outcomes and Savings

5.1 Agricultural Inputs and Outputs

Given that many households both earn wages in labor markets and operate their own farms, bridges may also affect agricultural choices. The results on agricultural outcomes using regression (4.2) are presented in Table 6. We first consider intermediate input (fertilizer plus pesticide) expenditures, and also the two components individually. These are columns 1-3 in Table 6. First, we see a very large increase in intermediate expenditures. Intermediate expenditures increase by C$659.97 ($p = 0.036$) on a baseline mean of C$890. The changes are primarily accounted for by fertilizer investment, which increases by C$383 ($p = 0.032$) compared to a statistically insignificant C$167 ($p = 0.304$) for pesticide.[19] Figure 2 plots the density of the natural log of intermediate expenditures in villages with and without a bridge. Not only does the mean increase, but variance across households falls from 1.33 to 1.21 among those using positive amounts of fertilizer and pesticide.

Columns 4–7 then consider how this increase in input use translates into yields on staple crops. We consider changes in yields for maize and beans, measured in total quintales (100 pounds) harvested.[20] Here, we find positive but mostly statistically

[19] These results are for the average household, so these effects combine extensive and intensive changes in intermediate usage.

[20] In additional results available online, we show that there is no shift into cash crops in response to a bridge, hence our focus on staple crops here. In Nicaragua, most coffee is grown on large plantations, so this type of shift is *a priori* unlikely. Moreover, newly planted coffee trees do not produce coffee for several years.

Figure 2: Density of Log Intermediate Expenditures (C$)

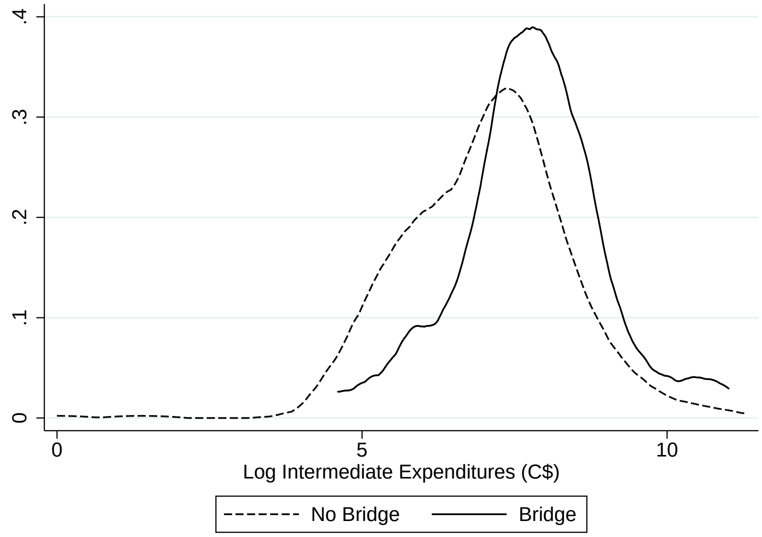

insignificant results, consistent with the fact that farm outcomes are subject to substantial shocks after investment is made. We do find that maize yield increases by 11.90 quintales per acre ($p = 0.000$).

Finally, we measure changes in farm profit.[21] We compute the value of harvest first using only maize and beans, as 90 percent of harvesting households harvest at least one of these two crops. No other crop is planted by more than 10 percent of households, so sale prices are limited outside these main staples. As a robustness check we also include the next two most prevalent crops, sorghum and coffee, as they are planted by 9 and 4 percent of farming households during this period. Profit rises substantially despite the increase in input costs (both higher wages and more fertilizer expenditure), suggesting farmers were initially subject to some distortion that decreases in response to a bridge.[22]

5.2 Savings Response

Bridges improve market access and increase earnings, which may result in greater savings. However, they also smooth the earnings process of workers that can now

[21]This is the value of output produced net of fertilizer and pesticide expenditures and payments to farm labor. Since not all crops are sold at market, we value harvest quantities at the median sale price during the year of production.

[22]Panel B decomposes all these results into differences between households that do and do not farm at baseline. The average effects are entirely driven by continuing farmers.

consistently reach those markets, which may reduce motives for precautionary savings. The key liquid savings vehicle in rural Nicaragua is storage of staple crops. Storage is defined as quantity harvested net of sales, debt payments, gifts, and land payments, measured as a share of total harvest quantity.[23] We measure this in terms of quantities for both maize and beans. Any household with no crop production is given a value of zero in this regression. Table 7 shows how bridges affect savings behavior. Regressions 1 and 3 show the average effect. Farmers save about 9 percentage points less of both their maize harvest ($p = 0.016$) and their bean harvest ($p = 0.052$). Columns (2) and (4) again show that the decrease in storage is concentrated among continuing farmers, the same subgroup as those who increase investment. Among continuing farmers, we find decreases of 13 percentage points for maize ($p = 0.012$) and 17 percentage points for beans ($p = 0.042$). Among those who did not farm at baseline, we see small and statistically insignificant changes in storage rates across build and no-build villages.

Taken together, these results suggest that farmers are selling a greater share of their harvests and using the proceeds to invest more in their farms. We test this directly by correlating changes from baseline intermediate expenditures with changes for baseline storage among treatment households. The correlation is -0.28 when using corn storage and -0.34 for bean storage. Both are statistically significant at the one percent level. This demonstrates that those who are increasing fertilizer the most are also those decreasing their savings the most.

5.3 Heterogeneous Effects

We lastly investigate two sources of heterogeneity in effects across households. First, we consider physical distance. While a bridge reduces the cost of crossing a flooded river, our intervention does not affect the cost of traveling from home to the bridge site. The bridge is most easily used by households located in close proximity to it, and therefore distance to the bridge is a measure of the intensity with which households are treated. Households vary substantially in their distance from the bridge. The average household at baseline is 1.5 kilometers from the bridge site, with a ninetieth and tenth percentile of 2.9 and 0.2 kilometers. We use household and bridge GPS

[23] In our supplementary material, we present the results when we define storage as the amount of each crop currently held in the household, which we ask directly. The results are similar. However, "amount currently stored" is net of any already-consumed harvest and is therefore not the total measure of harvest stored. For this reason, we prefer the in-text measure of storage.

locations to construct the distance in kilometers to the bridge site for each household, normalized by the distance of the median village household.[24] The second dimension of heterogeneity that matters here is initial consumption. With standard curvature in the utility function, the benefits to the poor should be larger than those for the rich. We therefore interact the treatment with baseline consumption expenditures. Using these two sources of heterogeneity, we run the regression

$$X_{ivt} = \alpha + \beta B_{vt} + \gamma D_{iv0} + \delta C_{iv0} + \zeta(B_{vt} \times D_{iv0}) + \theta(B_{vt} \times C_{iv0}) + \eta_t + \delta_v + \varepsilon_{ivt}. \quad (5.1)$$

where X, B, D, and C are intermediate expenditures, bridge, distance, and baseline consumption respectively. Intermediates and consumption are measured as inverse hyperbolic sines (to allow for zeros in intermediate expenditures).[25] The results are in Table 8. The interaction terms, ζ and θ, are both negative. This is consistent with our interpretation of distance as a measure of the intensity of treatment. Moreover, increasing baseline consumption by one percent implies a -0.83 percent ($p = 0.002$) decrease in treatment impact. Our interpretation is that richer individuals are already sufficiently unconstrained that they gain less from the consumption smoothing provided by the bridge. We formalize this argument in our quantitative model.

5.4 Brief Discussion of Robustness and Additional Results

Given the small sample, one may be concerned about the robustness of our results. We take this up in the online appendix, where we re-run our results using randomized inference instead of the wild bootstrap, and also include household fixed effects. Neither affects the results. Furthermore, we provide a number of additional results, including the impact on land rentals and purchases, price changes, and household size, and furthermore separate the treatment effect by year.

[24] In interpreting these results, it is worth noting that we did not find any households that relocated within their village at any point in the survey period. Nicaragua has weak land title rights, and most households report that they have lived in the same place since the Sandinista land reforms of the 1980s. As such, the location of households is unlikely to change in response to a bridge.

[25] The inverse hyperbolic sine function is $sinh^{-1}(x) := \ln\left(x + (1+x^2)^{\frac{1}{2}}\right)$. This function approximates the natural log of x while retaining zero-valued observations.

6 Model and Welfare Gains from the Bridge

Sections 4 and 5 highlight the fact that a bridge affects rural life along multiple margins. In this section, we build a model to help rationalize our results for two reasons. First, we use the model to derive a credible a link from the bridge-induced change in the earnings process to other changes we observe empirically.[26] Second, we ultimately use the model to compute the welfare gains derived from a bridge. In an economy with poor households and changing volatility, it is likely that changes in income understate the true changes in welfare. Thus, we use the model to decompose the importance of these various changes induced by a bridge.

6.1 Model Details

Basics and Timing We model a village as a small open economy. Each household owns a farming technology and is endowed with one unit of labor (we use household and farmer interchangeably in what follows). Households use a fraction of their labor endowment to access labor markets either inside their village, outside their village or both. The remainder of the labor endowment is employed in these labor markets to earn wages.

Time is discrete and consists of an infinite number of years. A year is T sub-periods long. Sub-periods $1, \ldots, T_w - 1$ make up the "dry season" while $T_w, \ldots, T$ is the the "wet season." Throughout the year, households make consumption and savings decisions. Income comes from two sources: farming, in which inputs are chosen at the beginning of the year and pay off T sub-periods later at the end of the year, and wage work. During the dry season, there is no flood risk and village agricultural productivity is zero. Flood risk appears during the wet season. As we discuss next, this flood risk limits the benefits of wage work outside the village.[27] This formalizes the temporal notion of agricultural decisions.

[26] As discussed in Section 2, we do not believe our results are driven by the ability to more easily trade goods across space.

[27] Our assumptions imply a period of riskless earnings at the higher outside labor market wage. This dampens the need for buffer stocks, and reduces the role of risk in this model. As we discussed in Section 2, however, work opportunities are limited during this period. Thus, we model the dry season in the most conservative way possible. Nonetheless, as we show later, risk plays an important role in our welfare analysis.

Labor Market Earnings Wage income can be derived from earnings in both the inside and outside labor market. Access to each market comes at a cost. Working in the village labor market requires f_i units of labor. If paid, the household earns the market clearing wage w_i per unit of labor in each sub-period. Alternatively the household can work in the outside labor market at the exogenous wage w^o. The ability to work outside costs f_o labor units. Moreover, access to this labor market is subject to an aggregate shock τ, which reduces earnings in the outside labor market. τ is the formalization of the flood shock cost during the wet season.

Villagers choose which labor market to work in at the beginning of the dry season and again at the beginning of the wet season, and can re-optimize that decision every season.[28] Specifically, they can work outside only at a cost of f_o, inside only at a cost of f_i, or both at a cost of $f_i + f_o$. If a household chooses both, it optimally decides how to split time between these two markets.

On-Farm Production A farm produces output using labor and an intermediate input (e.g., fertilizer or pesticide). At the beginning of each year, each household makes an irreversible intermediate input investment in their farm and a choice of labor input.[29] Labor is employed only during the wet season. Output is harvested at the end of the year. The farm technology is given by the constant elasticity of substitution function

$$y' = z^{1-\mu} \left(\alpha^{\frac{1}{\eta}} x^{\frac{\eta-1}{\eta}} + (1-\alpha)^{\frac{1}{\eta}} n^{\frac{\eta-1}{\eta}} \right)^{\frac{\mu\eta}{\eta-1}}, \tag{6.1}$$

where z is an idiosyncratic shock to farm productivity, x is the intermediate input, n is the labor choice and y' is the harvest. The shock z is realized at the end of the year and is i.i.d. across years. We assume that z is log-normally distributed with mean z_m and standard deviation z_s.

The output from the farm technology can be used to purchase inputs in the next season or can be stored by the household. We assume that all resources stored must

[28] Our choice on how to model labor market choices is motivated by two facts. First, we find that bridges increase inside-village wages even during non-flood periods. Second, workers shift from village wage work to outside wage work, even during non-flood periods. This is inconsistent with purely spot labor markets, as the returns to working outside have not changed during non-flood weeks. Instead, to rationalize the empirical result, the model requires some "sluggishness" in the response of labor across space. This guarantees that workers do not immediately move back into the village labor market when a flood ends. We make a particularly stark assumption – that the choice is fixed for the season – for analytical and expositional clarity.

[29] Alternatively, we could assume that intermediate decisions are made at the beginning of the wet season. Given that there is no uncertainty during the dry season, since there is no flooding nor a dry season harvest, these timing assumptions are equivalent.

be fully consumed by the next harvest.[30]

6.2 Recursive Formulation

The timing described above can be represented recursively. At the beginning of a year, a household enters into the period with y resources available. It chooses the faction of time to devote to village labor during the wet season (ϕ), and how to split resources between farming inputs (x, n) and savings (a_1) to carry into the sub-periods. Recursively, that process can be written as

$$\begin{aligned} V(y) \quad &= \quad \max_{\{a_1,n,x,\phi\}\geq 0} \quad W(a_1,\phi) + \beta \mathbb{E}_z \left[V(y'(z)) \right] \\ s.t. \quad & \quad y = a_1 + x + T_w w_i n \\ & \quad y'(z) = z^{1-\mu} \left(\alpha^{\frac{1}{\eta}} x^{\frac{\eta-1}{\eta}} + (1-\alpha)^{\frac{1}{\eta}} n^{\frac{\eta-1}{\eta}} \right)^{\frac{\mu\eta}{\eta-1}} \\ & \quad \phi \in [0,1] \,, \end{aligned} \tag{6.2}$$

where $W(a_1, \phi)$ is the value of entering the first sub-period of the season with a_1 savings, and assigning ϕ fraction of working time to the inside labor market during the wet season (the last T_w periods of the year).[31] The first constraint is the budget constraint: resources must be either used to purchase intermediates, labor, or held as savings to eventually consume during the year. The second defines how those input choices translate into harvest resources at the end of the year. The last is the simply a constraint to guarantee that the amount of time spent working inside the village does not exceed the the time endowment of the household.

Labeling the history of flooding shocks up to time t as s_t and $\pi(s_t)$ as the probability that s_t is realized, the value function of entering the year with a_1 savings and ϕ time

[30] This is again in the interest of notational simplicity and clarity, as it simplifies the state space of the household.

[31] Note that all households work outside during the dry season, as there is no value to labor on the firm during this time.

spent in the village labor market is given by

$$
\begin{aligned}
W(a_1, \phi) \quad &= \quad \max_{a_{t+1} \geq 0} \quad \sum_{t=1}^{T} \sum_{s_t} \pi(s_t) u(c_t(s_t)) \qquad (6.3) \\
s.t. \qquad & a_{t+1}(s_t) = [w_i \phi + (1 - \tau(s_t)) w_o (1 - \phi)] \, L(\phi) - c_t(s_t) + a_t(s_{t-1}), \quad \forall t > T - T_w, s_t \\
& a_{t+1}(s_t) = w_o(1 - f_o) - c_t(s_t) + a_t(s_{t-1}), \quad \forall t \leq T - T_w, s_t, \\
& L(\phi) = \begin{cases} 1 - f_i & \text{if } \phi = 1, \\ 1 - f_o & \text{if } \phi = 0, \\ 1 - f_i - f_o & \text{otherwise.} \end{cases} \\
& a_1(s_0) = a_1 \quad \text{for all } s_0
\end{aligned}
$$

(6.4)

We assume that households have the constant relative risk aversion utility function $u(c) = c^{1-\sigma}/(1-\sigma)$. The first budget constraint is during the wet season, while the second covers the dry season (note the lack of τ in this equation). The third constraint is the market choice, while the final says that the household begins the the year with whatever resources they decided not to invest in farm inputs.

The tradeoffs we discussed in the previous section are evident in the recursive formulation here. If a household decides to work in the outside market at higher expected wages, they are subject to fluctuations in earnings via floods. Given the non-negativity condition on storage, this creates a need to maintain a buffer stock of savings (e.g. Aiyagari, 1994). In addition, households make a farm investment decision. Thus, this need to hold a buffer stock is further complicated by the importance of devoting resources to farm investment. This tradeoff pulls resources away from investment, thus lower farm productivity. These tradeoffs links together seasonal farm decisions with high frequency changes in labor market access.

6.3 Stationary Equilibrium

In our quantitative exercise, we will assume that the baseline data we collect is from the stationary equilibrium of the economy. A stationary equilibrium of this economy is defined by an invariant cumulative distribution function $M^*(y)$, value functions V and W, decision rules $x(y)$, $n(y)$, $a_1(y)$, $\phi(y)$, $c_{it}(a_1)$, $c_{ot}(a_1, s_t)$, and a wage w_i such that

1. The value functions solve the household's problem given by (6.2), and (6.3) with associated decision rules
2. The village labor market clears: $\int_y \phi(y)L(\phi(y))dM^*(y) = \int_y n(y)dM^*(y)$
3. The law of motion for M, denoted $\Lambda(M)$ is such that

$$\Lambda(M(\hat{y})) = \int_y Prob\Big[\hat{y} \geq z^{1-\mu}\left(\alpha^{\frac{1}{\eta}}x(y)^{\frac{\eta-1}{\eta}} + (1-\alpha)^{\frac{1}{\eta}}n(y)^{\frac{\eta-1}{\eta}}\right)^{\frac{\mu\eta}{\eta-1}}\Big]dM(y)$$

   and $\Lambda(M^*(\hat{y})) = M^*(\hat{y})$ for all $\hat{y}$.

6.4 Discussion: Nature of the Exercise

Before characterizing the model, it is useful to highlight how our model and analysis map to the data. Our goal is to compare two different infrastructure regimes: one without a bridge and another that mimics the introduction of a bridge. Specifically, we assume the τ process takes two values, τ_f ("flood") and τ_d ("dry"). We interpret the bridge as a reduction in τ_f. That is, while the bridge does not change the likelihood that a flood occurs, $Pr[\tau_f]$, it does directly lower the cost of market access during a flood.

Note, then, that this interpretation implies that a bridge changes both the mean and the variance of the shock process. In addition to computing welfare, we will use the model to decompose the relative importance of the change in mean versus the change in variance.

6.5 Quantitative Exercise

We next solve and parameterize the model to study the welfare implications of the introduction of a bridge. We start the economy from the stationary distribution without a bridge, with shock realizations (τ_d, τ_f^{NB}).[32] We then shock the economy by shifting the cost of a flood from τ_f^{NB} to τ_f^{B} and trace out the full transition to the new steady state.

[32] The market access shock during a non-flood period is assumed equal across infrastructure regimes, and looking ahead, will be set at $\tau_d = 0$.

This allows us to do two things. First, we can use the model to measure the welfare impact of a bridge taking into account the value of reduced earnings risk and a full transition. Second, the model allows us to separately consider the quantitative importance of a bridge changing both the mean and variance of outside earnings.

6.5.1 Parameterization

Given that our goal is to measure the welfare effects from the outcomes that we observe, we parameterize the model to match both baseline characteristics of the villages and average treatment effects. To begin, we interpret each sub-period to be two weeks, as in the high frequency data. Therefore, we set $T = 26$ and, with the wet season lasting half of the year, we set $T_w = 13$. Throughout, we normalize the outside wage to $w_o = 1$. This leaves us with 13 parameters. Four of these are either set exogenously or match one-to-one with a given moment (4). The remaining 9 are jointly estimated to match 9 moments.

First, we set the household discount factor to $\beta = 0.98$ and the returns to scale of the farm technology to be $\mu = 0.8$. In terms of flooding risk, we can normalize the cost of outside market access during a non-flood to $\tau_d = 0$. The likelihood that a flood occurs is matched to the empirical flood realization rate, implying $Pr(\tau_f) = 0.41$.

This leaves 9 parameters that are jointly determined to match 9 moments. The parameters are household risk aversion (σ), the weight on intermediates (α), elasticity of substitution between intermediates and labor (η), the mean and standard deviation of the farm shock (z_m, z_s), the cost of reaching the outside market during a flood with and without a bridge (τ_f^{NB}, τ_f^{B}), and the cost to work in the inside and outside market (f_i, f_o). We choose these parameters to jointly match the pre-bridge outside wage premium, the change in the wage premium induced by the bridge, the fraction of wages earned in the outside market, the fraction of households earning wages in both markets, mean expenditure on agricultural expenditures, the change in agricultural expenditures induced by the bridge, the standard deviation of log-harvest values, and the change in high-frequency earnings in response to a flood, both pre- and post-bridge. These parameters are identified from the initial steady state and, as necessary, only the first two periods of the transition path so that it remains consistent with our data collection period. Appendix A provides a detailed description of how these moments help identify the parameters of interest.

The list of all parameter values is in Table 9, along with the relevant comparison of model and data moments.

6.5.2 Non-Targeted Outcomes

Before turning to the welfare gains, we ask whether the model can deliver some of the other responses we observed in the empirical data. These results are in Panel B of Table 10, which also includes the welfare calculations we discuss in the next subsection.

We begin with asking whether the model predicts a decline in savings, which we do not target in the calibration. Consistent with our empirical analysis, we consider the share of resources devoted to savings during the first two periods of the transition path (as our data is for two years following the intervention). We find that the bridge caused a meaningful decrease in storage (8.25 percent), as it is no longer needed to mitigate consumption risk. This is in the same order of magnitude as the empirical results, where storage rates declined between 9 and 10 percent, despite not being targeted directly.

The only post-bridge change in labor markets that our parameterization targets is the outside wage premium. We therefore consider two related moments in terms changes in labor market outcomes. First, we check whether the change in outside earnings implied by the model is similar to the data. We find a 97 percent increase in outside earnings, similar in magnitude to the 86 percent found in the data. Likewise, we can measure the total increase in earnings during the wet season induced by the bridge in the model and data. The "Bridge" coefficient estimated in the high frequency earnings regression is equal to 20 percent of mean earnings in the data, which is again similar to the corresponding change in the model of 24 percent.

Overall, the model matches both targeted and untargeted moments well, and we therefore turn to estimating the welfare gains of a bridge.

6.6 Welfare Gains from Bridges

With the model parameterized, we now compute welfare gains from the bridges. To reiterate, the exercise works as follows. We solve the stationary equilibrium under the parameterized model and shock realizations $(\tau_d, \tau_f^{NB}) = (0, 0.545)$. We then shock the model by changing τ_f^{NB} to the lower value $\tau_f^B = 0$ and compute the full transition path.

The consumption-equivalent welfare change is the proportion by which consumption would have to be increased in every state of the no-bridge economy in order for average utility across households to be equal to that in the economy with a bridge. With our assumed utility function, this proportion γ solves:

$$\int_y \gamma^{1-\sigma} V^{NB}(y) dM^{NB}(y) = \int_y V_1^B(y) dM^{NB}(y) \tag{6.5}$$

where V^{NB} and M^{NB} are the value function and distribution in the stationary equilibrium of the economy with no bridge, and V_1^B is the value function in the first period of the transition when the bridge has been introduced. The righthand side of (6.5) is the value of introducing a bridge in the stationary equilibrium of a no-bridge economy before any new decisions are made by households (hence the relevant distribution is M^{NB}). Thus, this γ includes the transition path to the steady state of the economy with a bridge.

Then our welfare measure γ can be solved for as

$$\log(\gamma) = \frac{\log\left(\frac{\int_y V_1^B(y) dM^{NB}(y)}{\int_y V^{NB}(y) dM^{NB}(y)}\right)}{1-\sigma}. \tag{6.6}$$

The main result of this exercise is in the first row of Table 10, where we shock the stationary equilibrium of a no-bridge economy with a bridge. Welfare increases by 11.42 percent. A bridge therefore increases welfare substantially.[33]

A key benefit of measuring welfare is that it takes into account economic features that would not necessarily be captured by measuring the change in average income or consumption. Here, there are two such additional pieces. The first is any transition cost or benefit associated with the movement between the two steady states. Second, in addition to changes in average consumption between the two steady states, consumption volatility declines as well. With curvature in the utility function, households

[33]For some perspective on the magnitude of this number, we compare to the cost of a bridge. A bridge costs 1,100,000 C$ to service an average of 33.5 families, which is equal to 30.7 weeks of mean earnings. Within the model, we calculate that financing this cost with a perpetual bond with an interest rate implied by the household's discount factor $(1/\beta - 1)$, this cost is equal to 4.68 percent of household consumption. Therefore, the total welfare effect more than justifies the cost of a bridge. In the online appendix, we also compute a standard return on investment measure, and come to a similar conclusion.

value this response.

How important are these changes relative to the change in average consumption between the steady states? It turns out that our model implies an exact decomposition of these various forces. Specifically, the γ from equation (6.6) can be written as

$$\log(\gamma) = \log(\Omega^T) + \log(\bar{c}^B/\bar{c}^{NB}) + \log(\Phi^B/\Phi^{NB}), \tag{6.7}$$

where Ω^T is the benefits derived by households during the transition path between the two stationary equilibria, $\bar{c}^j$ is average consumption in the stationary equilibrium of economy $j \in \{B, NB\}$ and Φ^j is a measure of the welfare loss from deviations from that average consumption in the stationary equilibrium of economy j.[34] We leave the derivation and details of these various terms to Appendix B.

The results of the decomposition are available in Table 10, in rows 2 − 4. In accounting for increase in welfare, the transition matters little, generating only 1.19 percentage points of the 11.42 percent change in total welfare. The remainder of the welfare gains is almost evenly split between the importance of the higher mean and and lower variance of consumption in the new steady state, implying a substantial role for the second moment of consumption here. This importance of volatility here are derived from two margins. The first is that the bridge directly eliminates the risk on outside labor earnings. Second is that access to outside labor market induces a change in the composition of annual income away from risky harvests.[35]

The decomposition results highlight an important point when studying welfare in a development context: changes in average income or consumption are almost sure to understate true welfare gains when the intervention helps to eliminate risk. These effects can be large and economically relevant among individuals near subsistence, as our results demonstrate.

6.6.1 Counterfactuals: Changing Mean vs Changing Variance

Finally, we conduct two counterfactual exercises. As mentioned above, the bridge-induced change in τ_f affects both the mean and variance of the shock distribution.

[34] Specifically, $\Phi = 1$ if consumption is always equal to $\bar{c}$ and Φ declines as the true consumption path deviates from $\bar{c}$. Thus, if a bridge decreases consumption volatility, the third term will be positive.

[35] One may be concerned that the importance of volatility is driven by an outsized value of the CRRA utility parameter σ. Interestingly, however, our estimation implies a value of $\sigma = 4.4$, similar to estimated values of this parameter in other contexts.

We therefore construct two counterfactuals in which the mean and variance are varied independently, holding the other moment fixed, to study the relative importance of these two channels.

Writing p_f as the probability of a flood, the first two moments of outside earnings in economy $j \in \{NB, B\}$ are

$$\begin{aligned} \mathbb{E}^j[w] &= p_f w_o(1-\tau_f^j) + (1-p_F)w_o \\ Var^j(w) &= p_f(w_o(1-\tau_f^j) - \mathbb{E}^j[w])^2 + (1-p_F)(w_o - \mathbb{E}^j[w])^2 \end{aligned}$$

Thus, in the economy with a bridge, $\mathbb{E}^B[w] = w_o$ and $Var^B[w] = 0$ due to the fact that $\tau_f^B = 0$. To isolate the change in the mean, we hold τ_f fixed at its no-bridge level $\tau_f^{NB} = 0.54$ but increase w_o until we reach $\mathbb{E}^B[w]$. Conversely, to isolate the change in variance, we allow τ_f to adjust to its calibrated τ_f^B level, but then decrease w_o until the economy faces an expected wage of $\mathbb{E}^{NB}[w]$. These cases are referred to as "mean only" and "variance only" respectively.

Table 10 reports the results from these cases in the last two columns. Increasing mean outside earnings plays a larger role are higher than those of reducing the variance, and that the sum of the two (9.84 percent) is smaller than the effect in the baseline scenario (11.42 percent) because lowering the variance and raising the mean have complementary effects. Not surprisingly, in the "variance only" case, a large part of the welfare benefits come from reducing dispersion in marginal utility, which lower volatility has almost no effect in the "mean only" case. Moreover, Table 10 also shows that both margins play an important role in allowing the model to match the untargeted moments discussed above.

7 Conclusion

We study the impact of integrating rural villages with more urban markets. Our intervention is footbridges that eliminate the risk of unpredictable seasonal flooding. These bridges have a substantial impact on the rural economy. Bridges eliminate the decrease in contemporaneous income realizations during floods, while allowing individuals to move into better jobs. This increases income during non-flood periods as well via general equilibrium effects. Second, agricultural investment in fertilizer and yields on staple crops both increase. These results imply that (1) lack of consistent outside market access can have a substantial impact on long-term agricultural decisions in rural economies and (2) the benefits of infrastructure extend beyond the ability to move goods more easily across space.

We then build a model that links these results together, in which bridges unlock resources for investment via more consistent labor market access, and show that it is consistent with the data. We use the model to synthesize these various channels into a single measure of consumption-equivalent welfare. Welfare increases by 11 percent. When we decompose the channels through which a bridge increases welfare, both the increase in mean consumption and decrease in consumption volatility play equally critical roles. Uncertainty about the ability to access outside markets affects *ex ante* decisions and therefore impacts welfare in a quantitative important way. This possibility has received little attention in the context of developing countries, where this issue is likely to be the most salient (with a prominent exception being Allen and Atkin (2016) in the context of physical goods trade).

Lastly, we note that these bridges are cost-effective in both consumption-equivalent terms (see Section 6.6) or by standard return on investment measures (see the online appendix). Critical is understanding the full set of channels affected by a bridge.

References

Acemoglu, D., and F. Zilibotti (1997): "Was Prometheus Unbound by Chance? Risk, Diversification, and Growth," *Journal of Political Economy*, 105(4), 709–751.

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

Aiyagari, S. R. (1994): "Uninsured idiosyncratic risk and aggregate saving," *The Quarterly Journal of Economics*, 109(3), 659–684.

Akram, A., S. Chowdhury, and A. Mobarak (2016): "Effects of Emigration on Rural Labor Markets," Working Paper, Yale University.

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

Allen, T., and D. Atkin (2016): "Volatility and the Gains from Trade," NBER Working Paper 22276.

Angeletos, G. (2007): "Uninsured idiosyncratic investment risk and aggregate savings," *Review of Economic Dynamics*, 10(1), 1–30.

Asher, S., and P. Novosad (2016): "Market Access and Structural Transformation: Evidence from Rural Roads in India," World Bank Working Paper.

Asturias, J., M. Garcia-Santana, and R. Ramos (2016): "Competition and the Welfare Gains from Transportation Infrastructure: Evidence from the Golden Quadrilateral in India ," Working Paper, Universitat of Pompeu Fabra.

Bencivenga, V., and B. Smith (1991): "Financial Intermediation and Endogenous Growth," *Review of Economic Studies*, 58(2), 195–209.

Bryan, G., S. Chowdhury, and A. Mobarak (2014): "Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh," *Econometrica*, 82(5), 1671–1748.

Bryan, G., and M. Morten (2018): "The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia," *Journal of Political Economy*, forthcoming.

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

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

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

Donovan, K. (2018): "Agricultural Risk, Intermediate Inputs, and Cross-Country Productivity Differences," Working Paper, Yale University.

EM-DAT (2017): *The CRED/OFDA International Disaster Database*Université Catholique de Louvain, Brussels, Belgium.

Fink, G., B. Jack, and F. Masiye (2017): "Seasonal Liquidity, Rural Labor Markets and Agricultural Production: Evidence from Zambia," Working Paper.

Foster, A., and M. Rosenzweig (2007): "Economic Development and the Decline of Agricultural Employment," in *Handbook of Development Economics*, ed. by T. Schultz, and J. Strauss, pp. 3051–3083.

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

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

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

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

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

Karlan, D., R. Osei, I. Osei-Akoto, and C. Udry (2014): "Agricultural Decisions After Relaxing Credit and Risk Constraints," *Quarterly Journal of Economics*, 129(2), 597–652.

Mobarak, A., and M. Rosenzweig (2014): "Risk, Insurance and Wages in General Equilibrium," Yale Working Paper.

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

Shamdasani, Y. (2016): "Rural Road Infrastructure and Agricultural Production: Evidence from India," Working Paper, University of Pittsburgh.

Sotelo, S. (2016): "Domestic Trade Frictions and Agriculture," Working Paper, University of Michigan.

Van Leemput, E. (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): "Nicaragua Poverty Assessment," World Bank Report, Washington D.C.

Main Tables for Text

Table 1: Comparison to LSMS Data

	LSMS: Rural Nicaragua	LSMS: Matagalpa and Estelí	Sample
Adult Males	1.68	1.69	1.45
Adult Females	1.55	1.57	1.43
Total HH Wage Earnings	411.31	352.37	452.25
Avg Male/Avg Female Earnings	3.25	4.15	4.76
Fertilizer + Pesticide over Harvest	0.08	0.10	0.13

Table notes: Matagalpa and Estelí are the departments where this study took place.

Table 2: Pre-Bridge Differences

	Constant	Bridge
Distance to Nearest Town		
Km from bridge site to town	5.44**	2.80
	(0.01)	(0.35)
Flooding Intensity		
Days flooded	2.40***	-0.45
	(0.00)	(0.46)
Flood likelihood	0.47***	-0.06
	(0.00)	(0.54)
Flood length (days)	5.10***	-0.36
	(0.00)	(0.84)
Household Characteristics		
Distance to bridge site (km)	1.52***	-0.09
	(0.00)	(0.33)
HH head age	45.05***	-0.08
	(0.00)	(0.95)
HH head yrs. of education	3.43***	0.26
	(0.00)	(0.36)
No. of children	1.28***	0.04
	(0.00)	(0.68)
HH size	4.15***	0.07
	(0.00)	(0.62)
Occupational Choice		
Agricultural production	0.49***	0.05
	(0.00)	(0.26)
Off-farm work	0.57***	-0.03
	(0.00)	(0.47)
Total wage earnings (C$)	1063.80***	1.11
	(0.00)	(1.00)
Male wage within the village (C$)	138.98***	-7.63
	(0.00)	(0.74)
Male wage in outside labor markets (C$)	182.02***	-10.41
	(0.00)	(0.77)
Farming		
Corn harvest	2.49***	1.00
	(0.00)	(0.21)
Bean harvest	1.50***	0.26
	(0.00)	(0.57)
Sale price of corn (C$)	189.33***	-22.38
	(0.00)	(0.69)
Sale price of beans (C$)	871.43***	43.39
	(0.00)	(0.86)
Plant staples (maize or beans)?	0.34***	0.03
	(0.00)	(0.45)
Fertilizer + pesticide expenditures	899.56***	99.50
	(0.00)	(0.59)
Joint F-test (linear), p-value	0.332	
Chi-squared test (probit), p-value	0.268	

Table notes: Distance is at the village level. Thus, it includes 15 observations and should be interpreted with caution. Flood intensity measures are from high frequency data and refer to the previous two weeks in the pre-construction rainy season. An observation in these three regressions is a community-week. The remaining variables are at the household level. The F and Chi-squared tests are conducted excluding distance from town, the flood intensity measures, wages and prices (since they are not defined for all households). p-values in parentheses. We do not cluster the standard errors here to give the regression the greatest chance of finding a statistically significant difference between the two groups.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Effects of Flooding on Income

	Household Income	Household Income	No Income Earned	No Income Earned
	(1)	(2)	(3)	(4)
Bridge_t	159.555***	84.224	0.061	0.055
	(0.000)	(0.338)	(0.104)	(0.308)
Flood_t	-143.512**	-116.397*	0.069**	0.044
	(0.030)	(0.084)	(0.042)	(0.196)
$\text{Bridge}_t \times \text{Flood}_t$	148.680**	153.102***	-0.107***	-0.125***
	(0.042)	(0.000)	(0.004)	(0.002)
Bridge_{t-2}		77.452		0.018
		(0.320)		(0.786)
Flood_{t-2}		-17.111		0.007
		(0.698)		(0.782)
$\text{Bridge}_{t-2} \times \text{Flood}_{t-2}$		-125.924		0.032
		(0.150)		(0.406)
Control mean	783.563	783.563	0.249	0.249
Observations	6,443	4,394	6,756	4,589
Individual F.E.	Y	Y	Y	Y
Week F.E.	Y	Y	Y	Y
Intra-cluster correlation	0.080	0.080	0.027	0.027

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

Table 4: Effects on Market Earnings

	All			Men			Women		
	Total	Outside	Inside	Total	Outside	Inside	Total	Outside	Inside
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Bridge	380.39*	306.10***	-27.70	267.09*	189.34***	-64.37	80.65*	79.21***	-7.53
	(0.096)	(0.000)	(0.842)	(0.072)	(0.004)	(0.326)	(0.092)	(0.000)	(0.778)
Control Mean, $t = 0$	1063.80	357.18	616.27	473.54	210.19	170.43	113.51	62.60	18.23
Observations	1,494	1,493	1,491	1,494	1,492	1,491	1,494	1,491	1,494
Time F.E.	Y	Y	Y	Y	Y	Y	Y	Y	Y
Village F.E.	Y	Y	Y	Y	Y	Y	Y	Y	Y
Intra-cluster correlation	0.073	0.050	0.050	0.049	0.023	0.015	0.027	0.018	0.005

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

Table 5: Decomposing Earnings Changes

Panel A: Men	No. of HH Members			Daily Wage		Days	
	Total	Outside	Inside	Outside	Inside	Outside	Inside
	(1)	(2)	(3)	(5)	(6)	(8)	(9)
Bridge	0.048	0.192***	-0.120	-5.63	68.57	0.866***	-0.303
	(0.464)	(0.000)	(0.130)	(0.828)	(0.102)	(0.000)	(0.360)
Control Mean, $t = 0$	0.543	0.294	0.251	182.025	138.980	1.401	1.299
Observations	1,507	1,507	1,507	306	349	1,494	1,497
Time F.E.	Y	Y	Y	Y	Y	Y	Y
Village F.E.	Y	Y	Y	Y	Y	Y	Y
Intra-cluster correlation	0.048	0.041	0.020	0.105	0.000	0.042	0.032
Panel B: Women	No. of HH Members			Daily Wage		Days	
	Total	Outside	Inside	Outside	Inside	Outside	Inside
	(1)	(2)	(3)	(5)	(6)	(8)	(9)
Bridge	0.109**	0.107***	0.013	44.99	4.45	0.589***	-0.072
	(0.018)	(0.000)	(0.568)	(0.382)	(1.00)	(0.002)	(0.484)
Control Mean, $t = 0$	0.171	0.118	0.055	206.754	121.894	0.538	0.183
Observations	1,507	1,507	1,507	147	107	1,493	1,498
Time F.E.	Y	Y	Y	Y	Y	Y	Y
Village F.E.	Y	Y	Y	Y	Y	Y	Y
Intra-cluster correlation	0.043	0.021	0.019	0.035	0.061	0.019	0.003

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

Table 6: On-Farm Impact

Panel A: Average Farm Outcomes	Input Expenditures Farm Profit			Maize		Beans		Farm Profit	
	Intermediates	Fertilizer	Pesticide	Harvest Quantity	Yield	Harvest Quantity	Yield		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Bridge	659.97**	383.31**	166.52	1.81	11.90***	1.02	2.19	2223.43***	1957.61**
	(0.036)	(0.032)	(0.304)	(0.142)	(0.000)	(0.124)	(0.326)	(0.004)	(0.030)

Panel B: Intensive and Extensive Margins	Input Expenditures			Maize		Beans		Farm Profit	
	Intermediates	Fertilizer	Pesticide	Harvest Quantity	Yield	Harvest Quantity	Yield		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Bridge $\times$ Farm at $t = 0$	1231.48**	702.07**	315.48	4.13*	12.84***	1.58	2.23	3211.46***	2990.37**
	(0.028)	(0.016)	(0.250)	(0.072)	(0.006)	(0.152)	(0.360)	(0.000)	(0.000)
Bridge $\times$ No farm at $t = 0$	-7.16	11.60	-7.96	-0.94	9.22**	0.35	2.07	1086.28	768.53
	(1.000)	(0.914)	(0.994)	(0.274)	(0.012)	(0.550)	(0.356)	(0.124)	(0.346)
Control mean, $t = 0$	889.56	607.43	303.48	2.49	12.29	1.50	4.59	2351.69	2559.20
Observations	1,492	1,493	1,492	1,492	359	1,499	356	1,478	1,478
Time F.E.	Y	Y	Y	Y	Y	Y	Y	Y	Y
Village F.E.	Y	Y	Y	Y	Y	Y	Y	Y	Y
Intra-cluster correlation	0.068	0.051	0.071	0.073	0.097	0.108	0.059	0.082	0.083

Table notes: Farm at $t = 0 = 1$ if the household is engaged in any crop production at baseline ($t = 0$), where *No Farm = 1 - Farm.* Farm profit in (8) uses only maize and beans (two main crops), while (9) includes sorghum and coffee as well. p-values in parentheses are clustered at the village level using the wild cluster bootstrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Farm Savings Choices

	Maize		Beans	
	(1)	(2)	(3)	(4)
Bridge	-0.085**		-0.091*	
	(0.016)		(0.052)	
Bridge × Farm at $t = 0$		-0.129		-0.172
		(0.012)**		(0.042)**
Bridge × No farm at $t = 0$		-0.032		0.007
		(0.280)		(0.842)
Control mean	0.942	0.942	0.928	0.928
Observations	1,507	1,507	1,507	1,507
Time F.E.	Y	Y	Y	Y
Village F.E.	Y	Y	Y	Y
Intra-cluster correlation	0.036	0.036	0.048	0.048

Table notes: *Farm* = 1 if the household is engaged in any crop production at baseline, where *No Farm = 1 - Farm*. p-values in parentheses are clustered at the village level using the wild cluster bootstrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Heterogeneous Impact on Farm Expenditures

	Intermediates	Intermediates	Intermediates	Intermediates	Intermediates
	(1)	(2)	(3)	(4)	(5)
Bridge	1.045**	9.616***	1.767***	10.845***	7.96**
	(0.038)	(0.000)	(0.008)	(0.000)	(0.012)
Consumption		0.652**		0.657**	0.646**
		(0.036)		(0.016)	(0.034)
Bridge × Consumption		-0.831***		-0.877***	-0.901***
		(0.002)		(0.002)	(0.006)
Distance			0.144	0.166	
			(0.598)	(0.552)	
Bridge × Distance			-0.586*	-0.614**	
			(0.086)	(0.048)	
Log(Distance)					-0.010
					(0.996)
Bridge × Log(Distance)					-0.522
					(0.262)
Control mean, $t = 0$	3.458	3.458	3.458	3.458	3.458
Observations	1,507	1,507	1,483	1,483	1,483
Time F.E.	Y	Y	Y	Y	Y
Village F.E.	Y	Y	Y	Y	Y
Intra-cluster correlation	0.075	0.075	0.075	0.075	0.075

Table notes: The dependent variable is the inverse hyperbolic sine of intermediate input expenditures. *Consumption* is the inverse hyperbolic sine (IHS) of baseline consumption expenditures. Regressions (3)-(4) measure *Distance* as kilometers from house to bridge site normalized by median distance in the village. Regression (5) uses log of raw distance measured in kilometers. p-values in parentheses are clustered at the village level using the wild cluster bootstrap-t with 1000 simulations. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 9: Parameter Values and Moments

Panel A: Parameters		
Parameter	Value	Method/Moment matched
Utility function		
Discount Factor, β	0.960	Set exogenously
Risk Aversion, σ	4.397	Jointly estimated
Agricultural production		
Returns to Scale in Agriculture, μ	0.800	Set exogenously
Scale Parameter on Intermediates, α	7.381 $\times 10^{-4}$	Jointly estimated
Elasticity of Substitution, η	3.191	Jointly estimated
Mean of Log(z), z_m	2.759	Jointly estimated
St. Dev. of Log(z), z_s	0.915	Jointly estimated
Flood risk		
Outside market access cost (non-flood), τ_d	0	Normalized
Outside market access cost (flood), pre-bridge, τ_f^{NB}	0.545	Jointly estimated
Outside market access cost (flood), post-bridge, τ_f^{B}	0	Jointly estimated
Probability of flood, $Pr(\tau_f)$	0.41	Empirical flood realizations
Market access costs		
Cost of Inside Market, f_i	0.007	Jointly estimated
Cost of Outside Market, f_o	0.018	Jointly estimated
Panel B: Joint estimation		
Target moments for joint estimation	Model moment	Data moment
Outside wage premium, pre-bridge (as %)	37.2	37.2
Outside wage premium, post-bridge (as %)	1.1	1.1
Fraction of labor earnings from outside (as %)	33.4	33.4
Households working both inside and outside (as %)	22.4	22.4
St. dev. of log(harvests)	1.31	1.31
Mean intermediate expenditure, baseline	0.88	0.88
Increase intermediate purchases (as %)	70.9	70.2
Earnings decline during flood, pre-bridge (as %)	18.0	18.0
Earning decline during flood, post-bridge (as %)	0.0	0.0
Root Mean-Squared Error (RMSE)	0.004	

Table notes: There are 13 parameters in the model. Four are either set exogenously or match one-to-one with a given moment. The remaining 9 are jointly estimated. These are listed in Panel A. Panel B lists the 9 moments for the joint estimation, and the corresponding model and data moments. Mean Intermediate Expenditure is expressed as a fraction of mean household earnings. RMSE takes the ratio of each model and data moment, subtracts one, raises to the second power, takes the mean across all moments, then takes the square root.

Table 10: Simulation Results

	Effect of Bridge	Mean Only	Variance Only
Panel A: Welfare Decomposition			
Consumption Equivalent Welfare, log(γ)	11.42	8.66	1.18
Effect of Transition, log(Ω)	1.19	0.46	0.10
Mean Consumption Gain, $\log(\bar{c}^B/\bar{c}^{NB})$	4.82	8.10	0.45
Gain from Lower Volatility, $\log(\Phi^B/\Phi^{NB})$	5.42	0.09	0.62
Panel B: Non-targeted Moments			
Change in Storage	-8.25	4.67	0.60
Change in Outside Earnings	97.06	126.01	-7.12
Bridge Effect on Wet Season Mean Income	24.01	47.25	-3.69

Table notes: All results are measured as percentages. The non-targeted model moments are computed as the first two periods after the shock, consistent with the data we collected.

Appendix

A Identification of Model Parameters

As discussed in the text, the model has 13 parameters once we set T, T_w and normalize $w_o = 1$. Four of those 13 are set exogenously, leaving 9 parameters to be jointly matched to 9 moments. This appendix discusses how these 9 moments help identify each parameter. For simplicity, we denote $P[\tau_f] := \pi$ throughout this appendix.
The remaining parameters are:

1. σ, the coefficient of relative risk aversion in the utility function
2. α, the weight on intermediates relative to labor in the production function
3. η, the elasticity of substitution between intermediates and labor in the production function
4. z_m, the mean of the lognormal farm productivity shock
5. z_s, the standard deviation of the log normal farm productivity shock
6. f_i, the required payment to work in the inside labor market
7. f_o the required payment to work in the outside labor market
8. τ_f^{NB}, the income loss from flooding without a bridge
9. τ_f^{B}, the income loss from flooding with a bridge

The moments are moments are:

1. The fraction of income derived from the outside labor market (34 percent),
2. The fraction of households deriving income from both inside and outside labor (22 percent),
3. The pre-bridge outside wage premium (37 percent),
4. The post-bridge outside wage premium (1 percent),
5. Mean intermediate spending divided by average household labor earnings (0.88),

6. The post-bridge percentage increase in intermediate spending (73 percent),
7. The standard deviation of log-harvest value (1.31)
8. The decrease in income during flooding as a percentage of the mean without a bridge (18 percent),
9. The decrease in income during flooding as a percentage of the mean with a bridge (0 percent),

Below, we detail how these parameters match to various moments. Throughout, we will denote post-bridge variables with $'$'s. We begin by deriving some useful relationships between prices, parameters and moments that inform the identification.

Identifying the cost of outside market access during flood, τ Note that we can compute the loss in income due to flooding as a fraction of mean income. For this, we make use of the fact that we will match the fraction of income coming from the outside labor market, which is measured in a non-flood week. Then

$$\text{Loss from Flood} = \tau_f^j \times \text{Outside Income Share} \implies \tau_f^j = \frac{\text{Loss from Flood}}{\text{Outside Income Share}} \tag{A.1}$$

where $j \in \{NB, B\}$. This relationship follows from the assumption that households cannot adjust their market choice in response to a flood in-season. With this, we can pin down τ_f^B and τ_f^{NB}. Since empirically there are no losses from a flood one a bridge is built, this immediately implies that $\tau_f^B = 0$. This further implies that there is no variance of in earnings after a bridge is built, which will be useful below. This leaves only the baseline cost τ_f^{NB}. Taking the ratio of these moments at baseline gives us $\tau_f^{NB} = 0.545$.

Post-bridge wage premium pins down f_i as function of f_o In any period, the normalization $w_0 = 1$ implies that the wage premium is matched when the equilibrium value of w_i satisfies

$$w_i = \frac{1}{\text{Outside Wage Premium}}. \tag{A.2}$$

This equation holds both pre- and post-construction. Thus, matching the pre-bridge wage premium (37 percent) and post-bridge wage premium (1 percent) amount to choosing parameters that guarantee the inside labor market clears at this price.

Next since there is no earnings variation after the bridge (via the $\tau_f^B = 0$ result above), labor market clearing requires households be indifferent between working in the two markets,

$$(1 - f_i)w_i' = (1 - f_o)w_o. \tag{A.3}$$

Combining (A.2) and (A.3) implies that the post-bridge wage premium is matched if and only if

$$f_i = 1 - (1 - f_o) \times \text{Outside Wage Premium, post} \tag{A.4}$$

This means that given a value of f_o, the post-bridge wage premium pins down f_i exactly.

Identifying τ as a function of outside income share Next, we can compute the loss in income due to flooding as a fraction of mean income. For this, we make use of the fact that we will match the fraction of income coming from the outside labor market, which is measured in a non-flood week. Then:

$$\text{Loss from Flood} = \tau_f^j \times \text{Outside Income Share} \implies \tau_f^j = \frac{\text{Loss from Flood}}{\text{Outside Income Share}} \tag{A.5}$$

where $j \in \{NB, B\}$. Hence, the cost of a flood τ_f is identified conditional on matching the Outside Income Share from the data. Note that since empirically there are no losses from a flood one a bridge is built, this immediately implies that $\tau_f^B = 0$.

Production Function Parameters Next we derive some useful relationships for the production function parameters. The first order conditions on household input choices imply that

$$\frac{x}{n} = \frac{\alpha}{1 - \alpha} \left(T_w w_i\right)^\eta. \tag{A.6}$$

First we can derive an equation for η as

$$\eta = \frac{\log(x'/x) - \log(n'/n)}{\log(w_i'/w_i)}. \tag{A.7}$$

Given the conditions above that pin w_i and w_i' to data moments, the observed change in intermediate spending (which equals x'/x) would identify ηnet of the fact that n'/n is still unknown. However, for any given value of n'/n, this equation pins down η.

Given a value of η, we can also solve for α using expenditures on intermediates from the data. The moment is expressed as expenditure divided by the mean labor income of households. Therefore

$$\text{Intermediate Expenditure} = \frac{\frac{\alpha}{1-\alpha}(T_w w_i)^\eta E[n]}{(1-f_i)\phi_i w_i + (1-f_o)\phi_o w_o + (1-f_i-f_o)\phi_b \tilde{w_b}} \tag{A.8}$$

where ϕ_i, ϕ_o, ϕ_b are the measures of household choosing to work inside, outside and both, and $\tilde{w_b}$ is the mean earnings of households working in both. Solving this equation for α gives a closed form for α as a function of labor market choices (which also imply η), f_o (which also implies f_i), and moments from the data.

Identifying the cost of working outside, f_o The parameter f_o is identified by the measure of households that choose to work in both inside and outside labor markets. Recall that the *difference* between f_i and f_o is pinned down by the post-bridge outside wage premium in equation (A.4). Then the *level* of f_o is given by the number of households willing to pay both fixed costs to get the ability to mix between labor markets. First consider the case where $f_o = 0$ so that all households work in both markets and choose an optimal mixture of the two income processes. Households with low assets (and low savings) prefer to work more in the inside labor market due to its lower risk profile while households with a greater buffer stock of savings prefer the outside labor market with higher average earnings. As we increase f_o – with f_i correspondingly increasing by equation (A.4) – very high asset households that work mostly in outside markets prefer to not pay f_i and low asset households that mostly work in inside markets prefer to not pay f_o. In the extreme, as fixed costs get very high (e.g., if $f_o + f_i > 1$) then no one works in both. Hence, there are guaranteed to be parameter values of f_o that rationalize any percentage of households choosing both from 0% to 100%. Notice that the choice of labor markets depends on the earnings processes (w_i and τ), risk aversion σ, savings decisions given assets, and the distribution of households over asset levels.

Risk aversion, σ The risk aversion parameter σ is identified by the fraction of total earnings coming from outside markets. The choice of whether or not to work in outside labor markets is a tradeoff between the higher average returns in outside markets against the certainty of the inside market income process. If households were risk neutral, 100% of earnings would come from outside, as expected inside earnings per period is the reciprocal of the Outside Wage Premium (0.71), and the expected outside earnings is $((1-\tau_f)\pi+1-\pi)w_o$ (0.78). These values follow from the earlier discussion. As households become infinitely risk averse, they care only about earnings in the lowest state, which is still 0.71 in the inside market and 0.55 in the outside market. Thus, no household would choose the outside market. Therefore, conditional on savings rates, the distribution of households over asset values, fixed cost parameters, and income processes, there exists a value of σ that can match the fraction of total earnings from outside from the data.

Recap So Far Up to this point, we have solved the cost of a flood both with (τ_f^B) and without (τ_f^{NB}) a bridge. The next step is that we have shown that risk aversion σ and the fixed costs (f_i, f_o) are identified by the share of income coming from outside, the post-bridge wage premium and the percentage of households deriving income from both outside and inside. This is true conditional on a wage w_i, a distribution over assets, and savings rules. Those fixed costs and labor choices (together with the optimal labor choices post-bridge, given by $E[n']$) are sufficient to compute η and α from the equations above. This is 7 of the 9 parameters. The last two deal with the farm shock.

Mean (z_m) and Standard Deviation (z_s) of the Farm Shock The only remaining parameters we have to identify are z_m and z_s, the parameters of the farm shock distribution. First, we choose z_m so that the value of w_i that matches the pre-bridge wage premium is an equilibrium. To do this, we set the wage w_i to the reciprocal of the pre-bridge wage premium, then compute labor supply (given values of σ, f_i and f_o found before). Then the parameter z_m, which controls the mean productivity in agriculture, changes inside labor market demand. That is, as z_m increases then conditional on w_i (= 1/Outside Wage Premium, pre), labor demand increases. This allows us to find a value of z_m such that the inside labor market clears at the target wage premium.

Finally we are left with one parameter and one moment. The variation in harvests is directly affected by z_s, since $Var(Log(\text{Harvest})) = Var(Log(Z_A)) + Var(Log(F(x,n)) = z_s^2 + Var(Log(F(x,n))$. Obviously there is interplay between the value of z_s and the choices of agricultural inputs, but this relationship is sufficient to identify z_s.

A.1 Simulating changes to moments to test identification strategy

While we have shown heuristically how to link parameters to moments, they naturally interact with each other in equilibrium. To demonstrate the close link between the individual parameters and individual moments, we simulate changes in target moments when each of the five parameters is changed individually. Table 11 shows the change in each model moment when each parameter is changed by 1 percent. We arranged the table such that – if the arguments we have made above are correct – the largest changes (in absolute value) should appear along the main diagonal of this table. Specifically, the row numbers of the parameters are equal to the column numbers of the moments that identify them, according to our arguments above. As we can see from the table, this is the case. For clarity, the largest change in each column appears in bold.

Table 11: Effects on Moments

Parameter	Labor Marketing Clearing	SD(log harvest)	Intermediate Spending	Fraction of Earning Outside	Percentage Earning Both
z_m	**11.27**	0.46	10.51	-3.62	0.00
z_s	6.08	**1.21**	2.28	-3.55	2.80
α	9.89	0.32	**13.38**	-2.40	-0.56
σ	3.71	0.10	1.07	**4.04**	-16.26
f_o	-3.09	0.02	0.40	1.86	**-20.92**

Table notes: This table shows the change in moments when each parameter is changed by 1 percent. All changes are listed as percentages. As an example, changing z_s by 1 percent induces a 1.21 percent change in the standard deviation of log harvests. The largest entry in each column is in bold, to show heuristically how each parameter matches to the moment discussed above.

Recall that $(\tau_f^{NB}, \tau_f^{B})$ are directly pinned down by moments, while η and f_i are implied directly by the choices made on the 5 parameters in Table 11, and thus do not need to be included here.

B Details of Welfare Decomposition

Note that average consumption can be written as

$$\bar{c}^{NB} = \int_y \phi^{NB}(y) T c_{it}^{NB}(a_1^{NB}(y)) + (1-\phi^{NB}(y)) \sum_{t=1}^{T} \sum_{s_t} \pi(s_t) c_{ot}^{NB}(a_1^{NB}(y)) dM^{NB}(y) \tag{B.1}$$

In the stationary equilibrium, this is equal to average income net of agricultural inputs.[36] Therefore, $\bar{c}^{NB}$ is a measure of value added. We can define $\bar{c}^{B}$ analogously for the stationary equilibrium with a bridge.

Within a stationary equilibrium we can also measure the welfare loss from variation in marginal utility. Define:

$$\Phi^{NB} = \left[(1-\sigma)(1-\beta) \int_y \frac{V^{NB}(y)}{(\bar{c}^{NB}/T)^{1-\sigma}} dM^{NB}(y)\right]^{\frac{1}{1-\sigma}} \tag{B.2}$$

$$= \left[\int_y \frac{\phi^{NB}(y) \sum_{t=1}^{T} c_{it}^{NB}(a_1^{NB}(y))^{1-\sigma} + (1-\phi^{NB}(y)) \sum_{t=1}^{T} \sum_{s_t} \pi(s_t) c_{it}^{NB}(a_1^{NB}(y))^{1-\sigma}}{(\bar{c}^{NB}/T)^{1-\sigma}} dM^{NB}(y)\right]^{\frac{1}{1-\sigma}}$$

This term Φ^{NB} is a measure of the welfare loss from variation in marginal utility. If consumption was equalized across all states, then consumption would be equal to $\bar{c}^{NB}/T$ in each sub-period and state, and $\Phi^{NB} = 1$. Since utility is concave, Jensen's Inequality implies that any other distribution of consumption across states with the same average consumption gives a smaller value of Φ^{NB}, so that it is a measure of the welfare effect of consumption dispersion across states. Again, we can define Φ^{B} in the same way in the stationary equilibrium of the economy with a bridge.

Finally, we can measure the effect of the transition from the initial stationary equilibrium with no bridge to the new stationary distribution with a bridge. We

[36] This follows because consumption in any season is equal to the sum of all earnings and storage, and harvest income is equal to the sum of storage and agricultural input spending. Hence, total consumption is equal to total earnings plus total harvest income less total agricultural input spending. Note, however, that this argument does not hold along the transition path.

measure this simply as:

$$\Omega^T = \left(\frac{\int_y V_1^B(y) dM^{NB}(y)}{\int_y V^B(y) dM^B(y)} \right)^{\frac{1}{1-\sigma}} \tag{B.3}$$

where $V^B(y)$ is the value function in the stationary equilibrium with a bridge. It is possible that Ω could be less than one or greater than one, depending on whether or not transitions involve costs that must be paid, and how $M^{NB}(y)$ compares to $M^B(y)$. Moreover, the magnitude of Ω depends on how many periods it takes to transition from one stationary equilibrium to another.

Taking logs and adding those pieces together, then doing some algebra, implies the decomposition. It follows that

$$\log(\gamma) = \log(\Omega^T) + \log(\bar{c}^B/\bar{c}^{NB}) + \log(\Phi^B/\Phi^{NB}). \tag{B.4}$$