Econometrica Supplementary Material

SUPPLEMENT TO "ELIMINATING UNCERTAINTY IN MARKET ACCESS: THE IMPACT OF NEW BRIDGES IN RURAL NICARAGUA"

(*Econometrica*, Vol. 88, No. 5, September 2020, 1965–1997)

WYATT BROOKS
Department of Economics, University of Notre Dame

KEVIN DONOVAN
School of Management, Yale University

THIS DOCUMENT contains further details in support of the main text. It includes robustness to various econometric specifications, more details on the high-frequency data used, and a calculation of standard return on investment. See the table of contents below for more details.

CONTENTS

APPENDIX SA: ROBUSTNESS OF MAIN RESULTS

Even though we had only 15 villages in the study, we obtained statistically significant effects. This is less surprising in the high-frequency results, as repeated measurement requires less sample size to detect effects. The low-frequency data do not benefit from the same design. Here, there are two reasons why we find statistically significant effects. First, the treatment effects are large. Second, the intra-cluster correlations are relatively low. In the main empirical results, the intra-cluster correlations range from 0.002 to 0.108 with both a mean and median of 0.057. This implies that for our median dependent variable, the minimal detectable effect is roughly 69 percent higher than if the randomization were done at the household level.[1] Combined with the large average treatment effects, we are able to detect statistically significant results. However, given the small number of clusters, it is instructive to show that our results are robust. First, we consider a different clustering procedure. Second, we vary the regression specification by including household fixed effects.

Wyatt Brooks: wbrooks@nd.edu
Kevin Donovan: kevin.donovan@yale.edu

[1]This calculation assumes all clusters have an average of 33.5 households per village, to simplify the exposition.

 https://doi.org/10.3982/ECTA15828

SA.1. *Using Randomized Inference*

We re-run the main regressions using Fisher "exact p-values" derived from randomized inference instead of the wild bootstrap cluster-t procedure in the main text. Roughly, while the bootstrap procedure fixes the treatment assignment and selects random samples of the data, randomized inference fixes the data sample but randomly varies the treatment. To compute these "exact p-values," we run our main results for each of the ${}_{15}C_6 = 5005$ possible treatment realizations across villages. Defining $\mathbf{T}_j \in \mathbf{T}$ as the vector of treatment assignments across villages for assignment $j \in \{1, \dots, 5005\}$, and $\widehat{\beta}_j(y) \in \{\widehat{\beta}_1(y), \dots, \widehat{\beta}_{5005}(y)\}$ as the estimated treatment effect for outcome y under assignment $\mathbf{T}_j$, we compute the exact p-value for outcome y as

$$p(y) = \frac{\sum_{j=1}^{5005} \mathbb{1}\big[|\widehat{\beta}_j(y)| \geq |\widehat{\beta}_{\text{obs}}(y)|\big]}{5005},$$

where $\widehat{\beta}_{\text{obs}}(y)$ is the estimated bridge effect for the actual treatment assignment. These are in Appendix SA, and we note that the results are quite similar to the results with the wild bootstrap.

In the main body of the paper, we prefer the wild bootstrap because, unlike exact p-values, it does not require that villages are i.i.d. between treatment and control. Since this is not the case here, the wild bootstrap cluster-t is the econometrically correct clustering procedure. However, given the use of permutation tests in other small-sample work, it is still instructive to show that our results are robust.

Table SI recomputes the main results using the randomized inference procedure. The p-values derived from this procedure are in brackets, while the wild bootstrap cluster-t p-values are included in parentheses for ease of comparison.

SA.2. *Using Household Fixed Effects*

As a robustness check to the regression specifications used here, we utilize the fact that we have three years of data, including observations before and after bridge construction, and estimate the main regressions with household fixed effects instead of village fixed effects. We find similar magnitudes and statistical significance for our estimates. Specifically, we compare the following two regression specifications:

$$y_{ivt} = \alpha + \beta B_{vt} + \eta_t + \delta_v + \varepsilon_{ivt},$$
$$y_{ivt} = \alpha + \beta B_{vt} + \eta_t + \delta_i + \varepsilon_{ivt}.$$

The first specification includes village fixed effects (δ_v) and is the main specification in the text. As a robustness test of the specification, we re-compute the main results using household fixed effects (δ_i) instead. The results are in Table SII. We also include the main estimates and p-values from the text for ease of comparison.

APPENDIX SB: HIGH-FREQUENCY DETAILS

This appendix covers additional results and details that are useful to understand the results of the paper. Sections SB.3 and SB.2 cover the high-frequency survey. The former discusses selection into the survey and balance, while the latter shows (1) response rates

TABLE SI

MAIN RESULTS WITH RANDOMIZED INFERENCE[a]

	Earnings			Farm Expenditures			Farm Outcomes					Storage	
	Total Earnings	Outside Earnings	Inside Earnings	Intermediates	Fertilizer	Pesticide	Maize Harvest	Maize Yield	Bean Harvest	Bean Yield	Farm Profit	Maize	Beans
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
Build	380.39	306.10	−27.70	659.97	383.31	166.52	1.81	11.90	1.02	2.19	1957.61	−0.085	−0.091
	(0.096)	(0.000)	(0.842)	(0.036)	(0.032)	(0.304)	(0.142)	(0.000)	(0.124)	(0.326)	(0.030)	(0.016)	(0.052)
	[0.042]	[0.001]	[0.822]	[0.057]	[0.085]	[0.252]	[0.163]	[0.006]	[0.246]	[0.322]	[0.063]	[0.032]	[0.010]††
Control Mean	1025.73	357.18	616.27	889.56	607.43	303.48	2.49	12.29	1.50	4.59	2559.20	0.942	0.928
Observations	1494	1493	1491	1492	1493	1492	1492	359	1499	356	1493	1507	1507
Time F.E.	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Village F.E.	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Intra-cluster correlation	0.073	0.073	0.050	0.061	0.040	0.081	0.054	0.082	0.122	0.032	0.088	0.036	0.048

[a]This table reports the main results using randomized inference to compute p-values. Those p-values are in brackets. For comparison, the original p-values using the wild bootstrap cluster-t are included as well, in parentheses. p-values for the wild bootstrap cluster-t are denoted $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$, while those using randomized inference are denoted †, ††, and † † †.

TABLE SII

MAIN RESULTS WITH HOUSEHOLD FIXED EFFECTS[a]

	Earnings			Farm Expenditures			Farm Outcomes					Storage	
	Total Earnings	Outside Earnings	Inside Earnings	Intermediates	Fertilizer	Pesticide	Maize Harvest	Maize Yield	Bean Harvest	Bean Yield	Farm Profit	Maize	Beans
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
Village FE	380.39	306.10	−27.70	659.97	383.31	166.52	1.81	11.90	1.02	2.19	1957.61	−0.085	−0.091
	(0.096)	(0.000)	(0.842)	(0.036)	(0.032)	(0.304)	(0.142)	(0.000)	(0.124)	(0.326)	(0.030)	(0.016)	(0.052)
Household FE	307.59	295.24	−41.76	646.48	437.81	152.94	1.65	14.76	1.16	3.15	1532.57	−0.085	−0.088
	(0.118)	(0.002)	(0.748)	(0.014)	(0.004)	(0.334)	(0.238)	(0.000)	(0.048)	(0.012)	(0.038)	(0.014)	(0.034)
Control Mean, $t = 0$	1025.73	357.18	616.27	889.56	607.43	303.48	2.49	12.29	1.50	4.59	2559.20	0.942	0.928
Observations	1494	1493	1491	1493	1492	1492	1492	359	1499	356	1478	1507	1507
Time F.E.	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Intra-cluster correlation	0.073	0.050	0.050	0.068	0.051	0.071	0.073	0.097	0.108	0.059	0.083	0.036	0.048

[a]This table reports the main regression specification using household and village fixed effects. Note that these are two separate regressions. p-values in parentheses are clustered using the wild cluster bootstrap-t with 1000 simulations. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

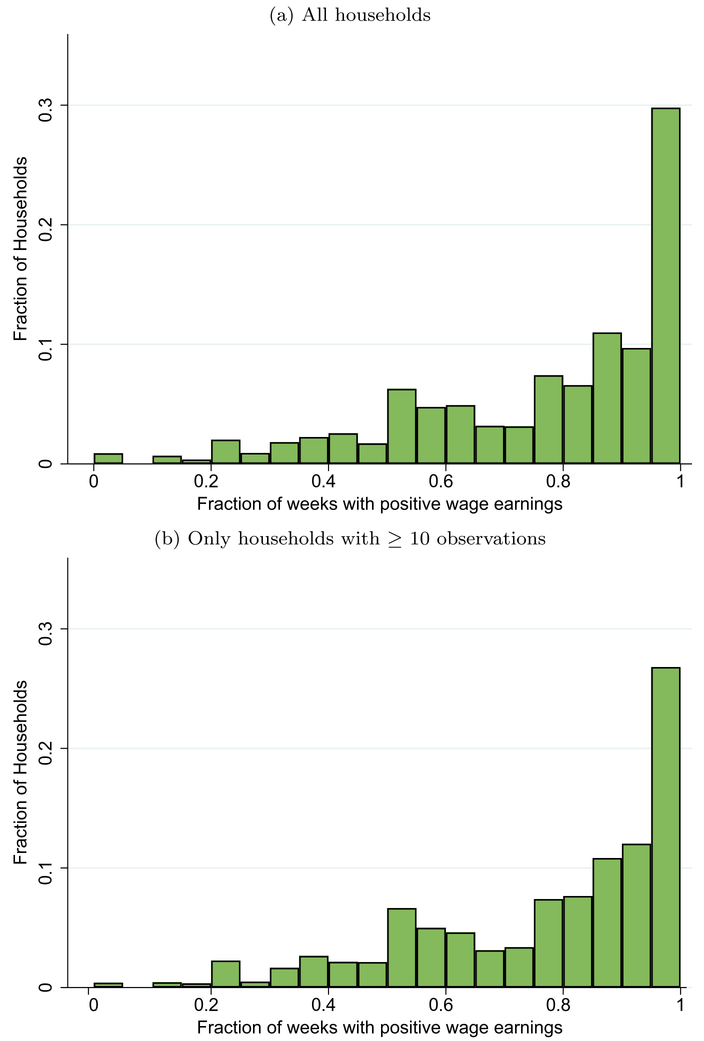

FIGURE S1.—Fraction of weeks with labor market income.

are uncorrelated with the likelihood of flooding and (2) even the most extreme assumption on missing values does not invalidate the fact that most individuals work in the labor market sometimes.

SB.1. *Importance of Labor Market Income*

Figure S1 plots the fraction of weeks we observe labor market income in households we contact during the high-frequency survey. One can see that almost all households receive at least some income from wage work.

SB.2. *How High-Frequency Survey Response Rates Change During Floods*

Figure S1 shows that almost all individuals in the high-frequency survey use the labor market to some degree. However, our survey is biased toward finding that result if floods

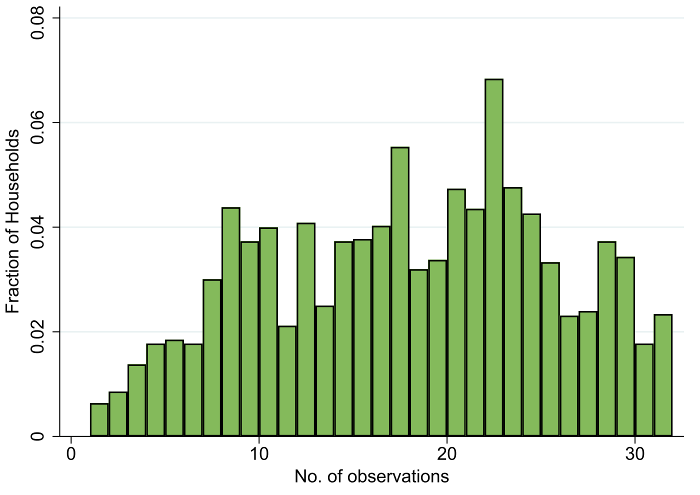

FIGURE S2.—Number of observations per household.

decrease the likelihood of answering the survey. Indeed, given that it is a cell phone-based survey, we do have an unbalanced panel. Figure S2 plots the histogram of the number of observations per household in the high-frequency data. The minimum is 1, the maximum is 32, and the average is 12. The maximum possible is also 32, as each village is surveyed bi-weekly.

To show that this variation is not driven by rainfall, we run the regression

$$\mathbb{1}[\text{answer}]_{ivt} = \alpha + \beta \text{Flood}_{vt} + \eta_t + \delta_i + \varepsilon_{ivt},$$

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

To further emphasize this point, Figure S3 reproduces Figure S1 with one key difference. Here, we assume that every period a household does not answer the survey, they received zero income that period. That is, we replace all missing values with zeros. This extreme assumption generates the lowest possible bound on the results driven by the unbalanced nature of the panel.

TABLE SIII

EFFECT OF FLOODING ON SURVEY RESPONSE[a]

	(1)	(2)
Flood	0.026	-0.031
	(0.128)	(0.060)
Mean Answer Rate	0.376	0.376
Observations	18,079	18,079
Individual F.E.	Y	Y
Week F.E.	Y	N

[a]"Mean answer rate" calculates share of weeks answered for each household, then takes the mean. p-values in parentheses are clustered using the wild cluster bootstrap-t with 1000 simulations. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

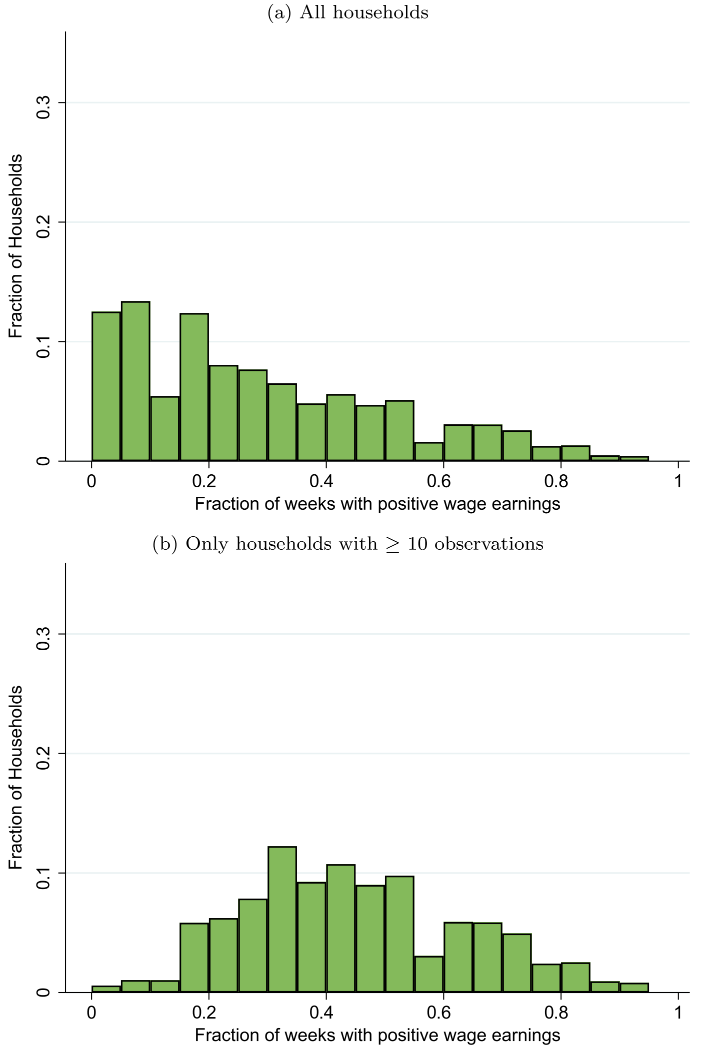

FIGURE S3.—Fraction of weeks with labor market income.

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

SB.3. *High-Frequency Data Balance Checks*

Table SIV shows the results from the regression

$$y_{iv} = \alpha + \beta \text{Bridge}_v + \gamma HFNum_{iv} + \eta(\text{Bridge}_{iv} \times HFNum_{iv}) + \varepsilon_{iv}.$$

TABLE SIV

PRE-BRIDGE DIFFERENCES HIGH-FREQUENCY DATA[a]

	Constant	Bridge	High-Frequency Responses	Interaction
Household Composition				
Distance to bridge site (km)	1.67	−0.15	−0.01	0.00
	(0.00)	(0.29)	(0.02)	(0.83)
HH head age	47.04	1.26	−0.19	−0.21
	(0.00)	(0.53)	(0.04)	(0.18)
HH head yrs. of education	3.02	0.30	0.04	0.00
	(0.00)	(0.48)	(0.03)	(0.96)
No. of children	1.16	−0.02	0.01	0.01
	(0.00)	(0.91)	(0.08)	(0.36)
HH size	3.81	−0.01	0.03	0.02
	(0.00)	(0.97)	(0.00)	(0.26)
Occupational Choice				
Agricultural production	0.44	0.10	0.004	−0.005
	(0.00)	(0.09)	(0.01)	(0.27)
Off-farm work	0.53	−0.02	0.00	−0.00
	(0.00)	(0.77)	(0.26)	(0.91)
Total wage earnings (C$)	1037.68	182.72	2.45	−21.10
	(0.00)	(0.49)	(0.84)	(0.31)
Farming				
Maize harvest	2.52	0.23	−0.00	0.09
	(0.00)	(0.85)	(0.95)	(0.31)
Bean harvest	1.34	0.72	0.01	−0.05
	(0.00)	(0.27)	(0.63)	(0.32)
Plant staples (maize or beans)?	0.34	0.06	0.00	−0.00
	(0.00)	(0.33)	(0.89)	(0.51)
Fertilizer + pesticide expenditures	934.26	187.54	−3.22	−11.45
	(0.00)	(0.49)	(0.80)	(0.59)

[a]Flood intensity measures as measured from high-frequency data and refer to the previous two weeks during rainy season only. p-values in parentheses. We do no clustering procedure here as to give the regression the greatest chance of finding a statistically significant difference between the two groups. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Here, y_{iv} is some outcome at baseline for household i in village v, $\text{Bridge}_v = 1$ if village v will receive a bridge, while $HFNum_{iv}$ is the number of responses for household i in the high-frequency survey.

We re-do the same exercise except with an indicator for whether a household takes part in the high-frequency survey:

$$y_{iv} = \alpha + \beta \text{Bridge}_v + \gamma HF_{iv} + \eta(\text{Bridge}_{iv} \times HF_{iv}) + \varepsilon_{iv}.$$

Here, y_{iv} is some outcome at baseline for household i in village v, $\text{Bridge}_v = 1$ if village v will receive a bridge, while $HF_{iv} = 1$ if household i participates in the high-frequency survey. The results are in Table SV.

SB.4. *Does Local Rainfall Predict Contemporaneous Floods?*

We use daily rainfall from the Climate Hazards Group InfraRed Precipitation with Station data (CHIRPS) which covers 1981–2016. CHIRPS provides rainfall estimates at the 0.05 degree resolution. We combine GPS coordinates for potential bridge sites in our

TABLE SV

PRE-BRIDGE DIFFERENCES HIGH-FREQUENCY DATA[a]

	Constant	Bridge	High-Frequency	Interaction
Household Composition				
Distance to bridge site (km)	1.78	−0.34	−0.34	0.32
	(0.00)	(0.06)	(0.01)	(0.13)
HH head age	51.37	−1.21	−8.31	0.83
	(0.00)	(0.64)	(0.00)	(0.78)
HH head yrs. of education	2.48	0.76	1.25	−0.65
	(0.00)	(0.18)	(0.00)	(0.32)
No. of children	0.96	0.11	0.42	−0.05
	(0.00)	(0.56)	(0.00)	(0.80)
HH size	3.66	0.00	0.64	0.16
	(0.00)	(0.99)	(0.00)	(0.62)
Occupational Choice				
Agricultural production	0.45	0.08	0.06	−0.04
	(0.00)	(0.31)	(0.37)	(0.64)
Off-farm work	0.53	−0.03	0.05	0.00
	(0.00)	(0.70)	(0.45)	(0.96)
Total wage earnings (C$)	906.02	297.20	207.95	−407.73
	(0.00)	(0.39)	(0.44)	(0.32)
Farming				
Maize harvest	2.78	−1.28	−0.39	3.26
	(0.00)	(0.39)	(0.74)	(0.07)
Bean harvest	1.49	0.78	0.01	−0.76
	(0.01)	(0.36)	(0.99)	(0.45)
Plant staples (maize or beans)?	0.35	0.05	−0.01	−0.03
	(0.00)	(0.49)	(0.91)	(0.71)
Fertilizer + pesticide expenditures	962.20	124.82	−82.31	−44.42
	(0.00)	(0.73)	(0.77)	(0.92)

[a]Flood intensity measures as measured from high-frequency data and refer to the previous two weeks during rainy season only. p-values in parentheses. We do no clustering procedure here as to give the regression the greatest chance of finding a statistically significant difference between the two groups. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

study, our high-frequency flood realizations, and CHIRPS to assess whether contemporaneous rainfall predicts flooding. To do so, we run regressions of the form

$$f_{vt} = \alpha + \beta \text{Rain}_{vt} + \gamma_v + \varepsilon_{vt},$$

where v is a village and t is a two-week time period. Table SVI shows the results. We consider three different rainfall measures. The first is the millimeters of rainfall that fall in the given two-week period. The second is the z-score of rainfall, where the mean and standard deviation are taken over that same two-week period since 1981. Finally, we consider an indicator if the realized rainfall is more than 2 standard deviations above the mean realization.

As one would expect, there is a positive relationship between rainfall and flooding. Using both millimeters of rainfall and the z-score, we find positive and statistically significant relationships in at least some specifications. More importantly, however, the R^2s are low. This suggests that rainfall has relatively little predictive power. As a final test, in columns (3) and (6), we ask whether extreme rainfall events are correlated with floods, which one may suspect ex ante, and find that they are not.

TABLE SVI
RAINFALL AND FLOODING[a]

	Flood Indicator			Days Flooded		
	(1)	(2)	(3)	(4)	(5)	(6)
mm Rain	0.004			0.020		
	(0.001)			(0.004)		
Rain z-score		0.038			0.145	
		(0.017)			(0.099)	
Rain $\geq \mu + 2\sigma$			−0.089			−0.377
			(0.113)			(0.650)
Total R^2	0.084	0.012	0.003	0.055	0.007	0.002
Between R^2	0.075	0.074	0.201	0.013	0.121	0.113
Within R^2	0.103	0.011	0.001	0.065	0.005	0.001
Observations	465	465	465	465	465	465
Village F.E.	Y	Y	Y	Y	Y	Y

[a]Standard errors are in parentheses. Mean and standard deviation are always taken over the identical week from 1981 to 2014.

APPENDIX SC: RETURN ON INVESTMENT OF THE BRIDGE

For the sake of completeness, we compute the cost-effectiveness of a bridge. Each bridge costs approximately 40,000 USD, or C$1,100,000 at an exchange rate of 0.036 USD = 1 córdoba. We first compute the annualized benefit in terms of increased labor market earnings per household, which is derived from our high-frequency data using changes in flooding and average time flooded. In particular, it is computed as

$$\begin{aligned}
&\text{Annual Effect on Earnings} \\
&\quad = 26 \times (\%\text{ with flood } \times \text{Wage effect in flood weeks} \\
&\qquad + \%\text{ with no flood } \times \text{ Wage effect in no flood}) \\
&\quad = 26 \times (0.095 \times 308.12 + 0.905 \times 159.42) \\
&\quad = \text{C\$}4512.21.
\end{aligned}$$

The annual effect on the farm is derived directly from the treatment effect on farm profit, C$1957.61, and together imply a total annual effect of C$6489.82. On average, there are 33.5 households per village, which implies a total village benefit of C$216,739. The internal rate of return can be computed as the solution to

$$1{,}100{,}000 = \sum_{t=1}^{T} \frac{216{,}739}{(1+r)^t},$$

where T is the useful life of the bridge in years. Bridges to Prosperity designs bridges to last 40 years.[2] This implies that the internal rate of return is 19.69 percent. If one were to compute the same rate of return using only the impact on wage earnings, the return would be 13.66 percent. Therefore, not considering the spillovers from labor market access to farm profit would underestimate the annual return of labor market access for rural communities by 31 percent.

[2]This estimate is based on internal tower corrosion rates of 25 microns per year. After 40 years, this is 1 millimeter, which no longer satisfies the design criteria for safety.

Note, however, that there are potentially other effects we are missing. For example, 18 percent of households in the control group mentioned clinic and maternal health care access as an important service cut off by flooding. Unfortunately, our sample size limits our ability to say much more than this when considering low-probability outcomes such as births. Such channels are potentially important additions when measuring the full impact of infrastructure, but require a substantially larger-scale study. We think this is an interesting margin for future work.

Co-editor Dave Donaldson handled this manuscript.

Manuscript received 1 November, 2017; final version accepted 9 March, 2020; available online 30 March, 2020.