1 Spatial and Dynamic Heterogeneity

> "You shared power calculations for effects between year 3 and year 1. Have you done power calculations for the coefficient on the time trend from year 1 to year 3 (or even year 4 if there were additional follow-up) and for the coefficient on regressions on the 'distance' trend (for households farther vs. closer from villages)? Getting reasonably well-powered estimates for these two is a crux for our case for this grant, so I wanted to follow-up on this piece."

To begin with distance effects of the treatment, we consider a MDE calculation that uses an estimating equation of the form:

$$Y_{ijt} = \alpha_j + \eta_t - \beta T_{jt} \times Distance_{ijt} + \gamma T_{jt} + \omega Distance_{ijt} + \varepsilon_{ijt}$$

Below is the MDE on the interaction term β, which measures the spatial heterogeneity in treatment effects (notice the minus sign in the estimating equation so that larger values denote steeper dropoff in effects) where distance is measured in kilometers:

OUTCOMES	**MDE, Average Effect**		**MDE, Distance Effect (per km)**	
	97 bridges	50 bridges	97 bridges	50 bridges
Market Food	275.7	316.12	19.51	22.38
Market Earnings	129.71	143.53	9.16	10.15
Use Fertilizer	0.0179	0.0206	0.0013	0.0014
Fertiilizer Spending	779.9	894.81	55.17	63.29
Work Outside	0.019	0.0218	0.0014	0.0015
Visit Clinic	0.0185	0.021	0.0013	0.0015
MUAC	0.08	0.1	0.0057	0.0071
Total Consumed	1649.05	1869.26	116.66	132.21
Total Income	1967.67	2177.27	139.17	153.98

This measure indicates the decrease in treatment effects for every kilometer from the bridge in the scenarios with both 50 and 97 bridges. A small estimate of β implies slower drop-off in treatment effects so that the treatment effects extend further out in space. One way to interpret the magnitude of this effect is: if the average treatment effect was only equal to its minimally detectable level, how far out (in km) would treatment effects have to extend for us to be unable to detect where the predicted treatment effect is equal to zero? We note that the MDEs for average treatment effects are 14-15 times greater than the MDEs for β. Therefore, a minimally detectable average treatment effect combined with a minimally detectable β would reach a predicted zero treatment effect at 14-15 km away from average household distance. The standard deviation of household distance is 6.26 km, so this is roughly 2.5 standard deviations. Moreover, this represents the minimum power we could have in any case in which both the average treatment effect and β are detectable. Since we already believe 15 km is quite a long distance in this context, we interpret this as very strong power to detect distance effects. This is due to the fact that our research design surveys 3 villages per site, which gives us substantial spatial heterogeneity to identify spatial effects.

As we discussed in the follow-up message, we don't expect to see a linear trend dynamic due to the readjustment we expect to occur after the bridge is built. So to think more flexibly about dynamics, I pooled Years 1 & 2 (the first two years after bridge construction) as a sort of "early effect" and I pooled Years 3 & 4 (the subsequent two years after bridge construction) as a "later effect" to see if we can detect changes in outcomes at that level:

OUTCOMES	**MDE, Years 1 & 2 vs. 3 & 4**	
	97 bridges	50 bridges
Market Food	519.25	641.43
Market Earnings	208.98	287.65
Use Fertilizer	0.0345	0.0419
Fertiilizer Spending	1474.62	1816.67
Work Outside	0.0359	0.0442
Visit Clinic	0.0328	0.0423
MUAC	0.15	0.19
Total Consumed	2923.75	3770.95
Total Income	3170.19	4363.55

We interpret this as informative about the value of adding a fourth year to the data collection. We can see that pooling two years together substantially improves precision over the case of only comparing years 1 and 3, and we can see that precision is better in the 97 bridge scenario than the 50 bridge scenario. Overall, the minimum detectable effects are somewhat large. In response to your question about additional data collection below, we make two proposals about avenues for data collection that, we believe, may offer more value for money.

2 Balance Checks

> "Have you done checks of baseline balance to confirm that there are no issues with the randomization? This is something we?re especially interested in, given the note in the pre-analysis plan (attached) that delays in bridge construction could compromise randomization."

Yes, the randomization order was not compromised by budget delays, which is the most critical thing for our identification strategy. This means that our treatments are orthogonal to observable and unobservable characteristics. Our balance checks look as we would expect with a large number of observations. In a few outcomes there are differences that are statistically significant (though the differences are not economically meaningful). The key is that, by design, treatment is orthogonal to characteristics. Following the literature, we will control for baseline characteristics in our estimating equations.

3 Spillovers and General Equilibrium Effects

> "How concerned are you about contamination of control areas (i.e., people in control areas using bridges in nearby treatment areas)? Are there ways we can account for or measure the extent of contamination/spillovers? Based on Figure 1 in the pre-analysis plan (attached), it looks like several of the control villages are fairly close to treatment villages."

This is possible, but unlikely due to B2P's selection process. Their needs assessment explicitly accounts for alternative means of crossing, including their own bridges when completed, and bridges are not built in places that have alternative options. Moreover, given the extremely rugged terrain in Rwanda, linear distance is a poor indicator of travel time, so two locations that look close on the map may actually be very difficult to travel between.

That said, we can check and account for this possibility since we collect data on travel times and frequently visited locations outside the village.

For example, we expect that bridges will cause intensive margin changes in travel times to schools, banks, markets and clinics, and extensive margin changes in the locations that are visited. We can check to see if nearby villages in control locations experience changes in these outcomes because of construction of nearby bridges.

Moreover, even if new bridges in treatment locations are not directly used by control villages, we would expect spill-overs from general equilibrium effects. This relates to your next question.

> "Maybe related to above ? Will it be possible to measure general equilibrium effects on non-targeted communities (e.g., exploiting differences in outcomes between control villages with fewer bridges close by vs. control villages with more bridges close by)? I?m thinking about the GiveDirectly general equilibrium effects paper (Egger et al.) and couldn?t remember if you were planning to do something similar."

Studying general equilibrium effects and contamination are naturally related. Our idea is the following: control and treatment villages share the same local markets. Therefore, understanding GE effects requires understanding how the markets evolve: how the supply of workers affect wages, where crop sales occur, prices of locally produced crops, and availability of agricultural intermediates. Currently, we have market surveys in 9 districts. Adding scale to this exercise by extending to our full study area would be extremely useful for studying the general equilibrium effects of markets. The methodology you suggest (exploiting differences in bridge density) is exactly what we would propose. These types of exercises naturally demand more power (they're at the market level instead of the household level). So adding more market surveys would be first order here.

Moreover, taking general equilibrium effects into account is crucial for getting a complete picture of how many people are affected by bridges. As discussed above, it is very interesting to think about how far bridge use extends beyond the first village being connected. But taking into account how adding all these new villages to outsides markets, general equilibrium effects on wages and prices in those market start to impact the entire region. Again, this is crucial for informing a complete picture of how bridges impact all the people in the areas where they are built.

4 Ideas for Additional Data Collection

> "We've talked about potentially providing $650k for an additional year of surveys. That seems valuable, but I wanted to check if there are alternative uses of that funding that could answer questions of interest for GiveWell. Two specific ideas that I'd be interested in your reaction to:
>
> 1. You mentioned that the number of household surveys remains constant even when you add more bridges. Would it be possible to improve the power for the effects on income and consumption by paying for more household surveys, or have you already hit diminishing returns?
> 2. Another key parameter we're interested in is the number of households who see benefits from bridges and how many more distant households see some consumption gains. Are there ways to meaningfully improve the ability to detect those effects (e.g., adding more surveys in distant villages)?"

On the first question, yes, our costs scale in direct proportion to the number of household surveys, and the diminishing returns involved make additional household surveys an unfavorable use of money. In a cluster design, it would be a little more efficient to add more survey locations, but of course that is limited by our identification strategy (and even then, the gains are not large). We believe instead of surveying more intensively, it would be better to concentrate on other forms of data either currently uncollected, or collected on a small scale.

These are the two really valuable opportunities that we think would have high returns for additional funding.

First, as previously mentioned we are currently gathering market data in a subset of districts so that we have precise information about how local prices evolve when villages are connected. Again, this is crucial for understanding how connecting villages to markets affect prices and therefore indirectly affect all other villages in the area. Moreover, market price data is useful in constructing our measure of household consumption. In particular, the aggregate consumption of a household is the value of all products consumed by the household. For market purchases, we measure value as expenditure, which is well-defined. However, this is more difficult for crops produced on the farm and consumed within the household, since it is not transacted at a price. If the household also sells the crop, we apply the sale price to the quantity consumed within the household when computing consumption. However, it is quite common that there are crops entirely consumed within the household. Instead, it would be useful to have data from markets on prices farmers receive for crops so that this can be applied to home consumption in the local villages. We expect this to be much more precise that computing unit values (total value divided by quantity) or applying inferred prices from neighboring households or even neighboring villages. This increase in precision will reduce mismeasurement in the key consumption outcome and therefore increase power.

Second, we are very interested in the spatial extent of effects, and are very interested in leveraging remote sensing techniques as a way of measuring agricultural impacts of bridges. From satellite imagery, it would be relatively cheap and scalable to measure changes in crop mix and agricultural productivity in the area around new bridges. The crucial input for this is accurate measurement of ground truth data that can be used to train a machine learning algorithm, which can then accurately classify images of fields across Rwanda. Recent improvements in techniques for these procedures makes this feasible, but requires context-specific data that we would need to collect on the ground.

5 Pre-specification and Data Availability

> "Do you plan to pre-specify how primary outcomes in the pre-analysis plan will be measured? I think the outcomes we?re most interested in are income and consumption effects, and I wanted to see if your team would be pre-specifying how those variables are constructed."
>
> "Do you plan to make the data and code publicly available?"
>
> "Our understanding is that the pre-analysis plan will be published. Is that correct?"

The answers to all of these questions is "Yes." Our pre-analysis plan is in the publication process at this moment. Also, data and codes will be publicly available in a de-identified form at the conclusion of the project in fulfillment of commitments made to other donors, and according to the requirements of our affiliations with J-PAL. This will be available on a public server, such as through Yale's library system, or the J-PAL dataverse.