

# Modeling Welfare and Uncertainty within a Frisch Demand System

Elliott Collins

August 1, 2017

## Abstract

In 2007, BRAC initiated a nation-wide randomized evaluation of what has since become a widely replicated program offering microenterprise support to particularly poor households, known as “Targeting the Ultra-Poor” or TUP program (Bandiera et al. (2017)). The experiment concluded that the transfers lead to significant increases in household consumption after four years. Here we revisit these results to compare two alternative measures of material welfare that more thoroughly exploit information on the composition of households’ expenditures. Both measures rely on the parameters of a Frisch demand system in which consumption is a function of market prices and households’ marginal utility of consumption. The first alternative welfare measure simply takes this marginal utility parameter,  $\lambda_{it}$ , as a theoretically grounded welfare index. With this approach, we find that household welfare goes down two years after the program, but goes up significantly after four, with positive spillovers to richer households in all periods. Our estimates of program effects on aggregate consumption also finds significant increases after four years, but essentially no effect after two. (While we fail to replicate the original study’s point estimates, we find qualitatively similar results.)

The second measure builds on the parameters of this demand system to construct a measure of “vulnerability” based on that of Ligon and Schechter (2003), which accounts for the welfare implications of uncertainty in consumption. With this approach, we find that program participants are made better off not only through higher average household welfare, but also due to a fall in period-to-period variation, suggestive of a greater ability to smooth welfare over time. However, this accounts for only a small portion ( $\approx 10\%$ ) of the welfare gains realized, and are not generally shared by neighboring ineligible households.

## Introduction

In 2007, BRAC initiated a nation-wide randomized evaluation of what has since become a widely replicated program offering microenterprise support to particularly poor households, known as “Transfers to the Ultra-Poor” or the “graduation framework” (Bandiera et al. (2017)). This experiment randomly selected branch offices across the country to implement the program, which identified “ultra-poor” women and offered them a cow, along with two years of training and support in the use of livestock as a source of income. The results of the experiment were understandably complicated, but showed a notable increase in households’ average income and consumption. The cluster randomization also allowed for estimation of spillovers to ineligible households, which saw small positive effects on wages. Promising early results lead to the ambitious roll-out of this framework worldwide, as documented in Banerjee et al. (2015), Collins and Ligon (2017) and related work in microenterprise support like Blattman, Fiala, and Martinez (2013). Here, we use this experiment as an opportunity to explore alternative approaches to measuring the distribution of welfare among respondent households. We hope to provide further insight into the effect of the “graduation framework” for capital transfers on the economic lives of the poor, while exploring the merits of alternative methods of welfare measurement.

In assessing the progress of a group of households out of poverty, whether that progress is by way of general economic forces or a targeted anti-poverty program, we must first decide on an observable indicator of welfare. What it is to be rich or poor is itself a dynamic and multi-dimensional question, and a variety of methods have been proposed to distill and quantify this complexity. Economists’ typical approach has been to measure average income or consumption over a short period, as recalled in survey data. While fairly intuitive in its interpretation, this approach also faces both practical and conceptual shortcomings, a few of which we hope to address.

First, aggregate consumption is relatively difficult and expensive to elicit, requiring lengthy survey modules asking about a long list of consumption goods. This list is also likely to include items for which measurement error is high, or which conflate consumptive and productive expenses (e.g. transportation, airtime). Households will also generally have non-linear Engel curves, so that expenditure shares will covary with total expenditures. At the least, this implies that aggregation, by ignoring composition of expenditures, throws out useful information. It also means that when a survey excludes a set of goods whose elasticity is different from that of the basket being measured, then the degree of mismeasurement in aggregate consump-

tion will vary systematically with household welfare. For example, if inferior goods bought exclusively by the poor are not addressed, then aggregates for poor households will be systematically off relative to wealthier households. To address these concerns, our first task will be to reconsider the treatment effects estimated in Bandiera et al. (2017) using the consumption-based welfare metric proposed in Ligon (2016b), with application to estimating treatment effects in Collins and Ligon (2017). This method exploits the composition of expenditures to estimate a demand system and an index related to the marginal utility of consumption. This approach allows us to use a subset expenditure categories that excludes potentially problematic goods or services, and allows for non-linear Engel curves.

Another shortcoming of total consumption is that it does not account explicitly for the importance of risk in a household’s well-being. The task of smoothing consumption over time can itself be quite costly for poor households, as they work to insure against negative shocks and spend money (or forego investments) on commitment savings devices. We’ll set this aside for now and focus instead on risk as a constitutive part of economic welfare. Any household with concave preferences prefers less variation in consumption over time, *ceteris paribus*. This is especially true for those already living in poverty, for whom even a small tightening of the budget constraint might prove very difficult. Ligon and Schechter (2003) develop a welfare index that accounts for the welfare loss due to this sort of variation in consumption, which they term “vulnerability”. They begin with a CES preference structure and use intertemporal variation in aggregate household consumption to estimate consumption risk. Building on the demand system we’ve estimated, we will define an additively separable analog to this measure using disaggregated expenditure panels, again allowing for variable elasticities of substitution. This lets us relax the implicit homotheticity assumption built into their univariate model. This will ultimately allow us to speak to the TUP program’s effect not simply on average consumption, but on household exposure to risk as well.

## The Program

The TUP program was motivated largely by the insight that sufficiently poor individuals seem less able to benefit from the small uncollateralized loans BRAC markets to a large number of households worldwide. Instead offering these microloans, the TUP program sought to target these “ultra-poor” women using a participatory wealth ranking, then offered a direct transfer

of productive capital. Each woman was offered a menu of six livestock asset bundles, with 91% choosing a bundle that included a cow. The transfers were then followed by a classroom training session and a two-year period of training and monitoring by program staff. Of those identified as eligible at baseline, 86% were ultimately enrolled, with the other 14% either becoming ineligible later on or declining to participate<sup>1</sup>. In all, the livestock and the training are valued at around 560 USD PPP each. Finally, participants are given a subsistence food allowance for the first forty weeks to help compensate for potentially costly increases in household labor requirements. In all, the average per-person program cost measures as much as twice the baseline wealth of the average participating household.

## The Experiment

A census was started in 2006 in 1309 villages across Bangladesh. For each of 13 districts in which BRAC works, two subdistricts, served by one branch office each, were randomly selected. One of each pair of branch offices was then randomly selected to implement the TUP program. The census identified 99,775 households, of which 34% were identified as poor based on basic features of the household, asset stock, and employment activities. Of these, a participatory wealth ranking was conducted in each village which identified the so-called “ultra-poor”, which constituted half of poor households on average. Ultimately, 17% of households were deemed eligible for the TUP program. Nearly all of the poor households (both eligible and not) were sampled, along with a random 10% of non-poor households. This formed the initial sample of 23,029 households that formed the panel from 2007 to 2011. The treated households were enrolled shortly after the baseline survey, with 2009 and 2011 representing treatment effects two and four years after the fact, respectively.

## A Frisch Demand System

We start out by estimating the parameters of the flexible Frischian demand system laid out in Ligon (2016b). This system will be used to estimate an index of vulnerability which accounts not only for the average household welfare over time, but also the welfare implications of uncertainty. First however, we will consider treatment effects on households’ marginal utility in

---

<sup>1</sup>For more details about selection, see the original paper.

expenditures, denoted by the parameter  $\lambda_{it}$ . This outcome has the benefit of a clear theoretical interpretation, where it is the multiplier on a household's budget constraint. In exploiting variation in the composition of expenditures, rather than relying on simple summation, it also allows us to glean important information about household welfare with only a subset of the goods available for purchase. This is particularly valuable when we can ignore goods that are hard to measure or interpret (e.g. housing quality, gambling, financial services). We'll first present the model and resulting demand system.

Let  $U$  be an individual's utility function,  $C$  a consumption vector of  $J$  goods  $(c_1, \dots, c_J)$ , and  $u_j(C)$  the marginal utility of consumption of good  $j$ . Starting with the standard first-order condition of a household's consumer's problem,  $u_j(c_j) = p_j \lambda$ , and taking logs, we have the additively separable form:

$$\log u_j(c_j) = \log p_j + \log \lambda \quad (1)$$

In each period, the consumer chooses consumption and the multiplier on the budget constraint  $\lambda_{it}$  such that marginal utility in consumption for each good is proportion to its price, so that marginal utility in expenditure is equal for each good. So we understand  $\lambda_{it}$  as the household's marginal utility in expenditures.

For the sake of flexibility, let  $z$  represent a vector of  $l$  distinct household characteristics that may shift demand. We'll assume that  $z_i$  enters into household  $i$ 's marginal utility with the form

$$u_j^*(c_j, z) = u(c_j) \exp(\gamma_j z) \quad (2)$$

so that household characteristics enter into (1) linearly. We can further parameterize the households vector of marginal utility functions as  $u_j^*(C) = (\alpha C)^{-\beta^{-1}}$ , where  $\beta$  is a  $J \times J$  matrix of consumption elasticities with respect to marginal utility, and  $\alpha$  a  $J$ -vector of demand shifters related to a given item's budget share.<sup>2</sup> This yields for each good

$$-\beta^{-1} (\log C + \log \alpha) = \log p + \log \lambda - \gamma z \quad (3)$$

Finally, moving from consumption to expenditures and allowing for classical measurement error  $\epsilon$ , let  $x_j = p_j c_j e^{\epsilon_j}$ . Then we can say that

$$\log X = \log \alpha + (I - \beta) \log p + \beta \gamma z - \beta \log \lambda + \epsilon \quad (4)$$

---

<sup>2</sup>Ligon (2016a) discusses some of the merits and implications of this sort of functional form in more detail.

## Estimating the Model

We can now move towards estimating the parameters of (??). Indexing a panel or cross-sectional dataset with households  $i$ , goods  $j$ , years  $t$ , and markets  $m$ , we can write the reduced form analog as

$$\log(x_{ijtm}) = \log \alpha_j + (1 - \beta_j) \log p_{jtm} + \delta_{it} z_{it} + ce_{ijt} \quad (5)$$

where  $ce_{ijt} = -\beta_j \log \lambda_{it} + \epsilon_{ijt}$

If we do not observe prices, we need only make the further assumption that households face the same prices within a given market unit. This assumption allows us to control for price variation via good-market-time fixed effects, which does not yield a direct estimate of price elasticities, but provides the same residuals in the service of deriving  $\log \lambda_{it}$ .

$$\log(x_{ijtm}) = \alpha_{jtm} + \delta_{it} z_{it} + ce_{ijt} \quad (6)$$

Ligon (2016b) lays out the strategy for estimating  $\beta_j$  and  $\log \lambda_{it}$  up to a common factor  $\phi$ , such that  $\beta_j(\log \lambda_{it})$  forms the least-squares approximation of the  $(N \times n)$  residual matrix,  $ce_{ijt}$ . In estimating average treatment effects on “neediness”, we can choose any satisfying normalization of  $\phi$  (Since the units of  $\log \lambda / \phi$  are invariant to affine transformations). For this purpose, we impose that  $\log \lambda_{it}$  has a standard deviation of one, since it offers some intuition as to the magnitude of effect sizes.

However, for the sake of estimating “Vulnerability” we will go on to fully characterize the preference structure and impose a utility function, which will require a fully identified system.

In our case, we do observe prices after a fashion, and use this variation to identify  $\phi$  and  $\beta_j$ . In lieu of official price statistics, we observe both expenditures and quantities consumed for a range of goods, from which we impute unit values. We acknowledge (without entirely addressing) the standard concern that this conflates price and quality variation by exploiting the relatively modest data requirements of the Frisch approach. Resting again on the assumption that quality-adjusted prices within markets should be roughly the same, we exclude goods with particularly high variance or interquartile ranges in prices at the household level within markets. We then use the median stated unit value of a good (expenditure/quantity) among all households within a geographic unit as  $p_{jtm}$ . For aggregate categories like Rice (which contains brown, white, and spiced rice), household-level unit values are taken to be the expenditure weighted average among component goods. We incorporate these imputed prices to estimate (5), which provides

estimates of  $(1+\beta_j)$  as well as  $\delta_{it}$ . Combined with the decomposition of  $ce_{ijt}$ , this serves to identify  $\phi$ , giving us unique estimates of  $\beta$ , and  $\log \lambda_{it}$ .

## Characterizing the Expenditure System

We start by presenting estimates of  $\beta_j$ , the Frisch demand elasticities (with respect to  $\log \lambda_{it}$ ) for each good. By far the most commonly consumed commodities are rice, oil, and onions, each having been consumed by more than 95% of sampled households in 2007. Onions in particular will act as our numeraire good. Fruit and fish (including dried fish) are the most elastic categories in our sample, though fruit is only consumed by 26% of households. We also present  $\log \alpha$ , which is a preference parameter which scales demand for each good equal. Our empirical strategy results in this being equal to mean log expenditures at baseline. These are estimated at the good-market level, though we present estimates assuming a single market. Unsurprisingly, rice, fish, and fruit make up a significant portion of household food budgets. We estimate  $\log \alpha$  for each good specifically among those households with non-zero levels of consumption, so it's indicative of expenditure weights among these households, and not over the sample overall.

Turning to the distribution of household welfare in Figure 2, we can see the average relative value of  $\log \lambda_{it}$  shifting from year to year, improving (i.e. going down) from 2007 to 2009, then getting worse in 2011. Given the non-random sample of households, this is hard to interpret here, though it illustrates the potential value of this method for tracking a population over a longer period of time. Figure 4 splits these distributions by household eligibility, showing that the eligibility criteria imposed on households does manage to target households that are marginally worse off (and which stay worse off over the course of the panel). For comparison,

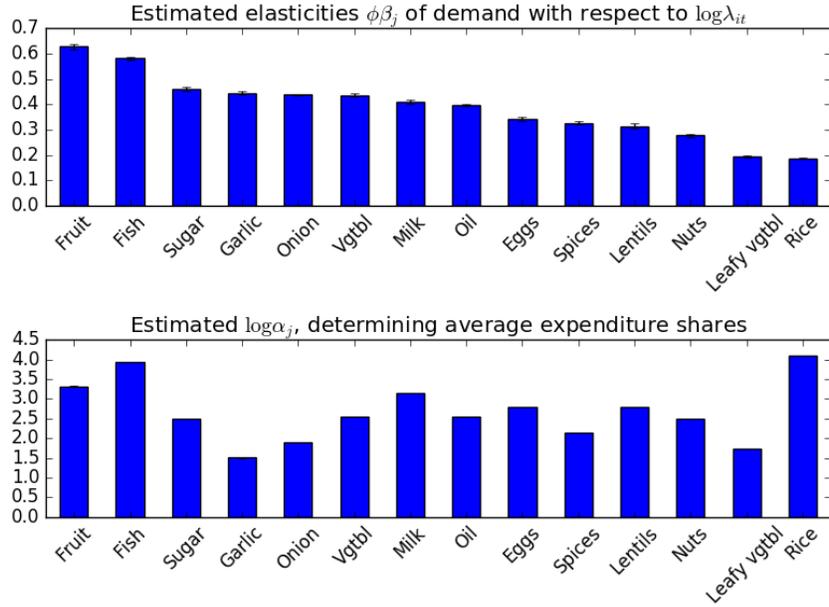


Figure 1:  $\beta_j$  for each good. Note that we're imposing here that elasticities do not vary by time or household.

## Estimating Vulnerability

Thus far, we have provided a practical and theoretically grounded analog to the standard aggregate consumption metrics. Spending enters into this simple consumption aggregate linearly regardless of what's being consumed or how much is already being spent. This makes it problematic as a welfare metric since it runs afoul of two fairly basic elements of demand theory, that welfare is concave in consumption and that elasticities differ across goods. Deriving  $\lambda_{it}$  models the fact that the price of additional utility depends on both income variation and different elasticities of substitution across goods. We carry this insight forward as we address another central insight of demand theory: the detrimental role of uncertainty and the frictions that prevent perfect welfare smoothing.

Like many of the risk-sensitive measures of welfare out there, ours rests on estimates of a household's expected utility. We draw on Ligon and Schechter's measure, which is essentially an estimate of the total welfare loss due to stochastic variation in consumption. In a standard dynamic

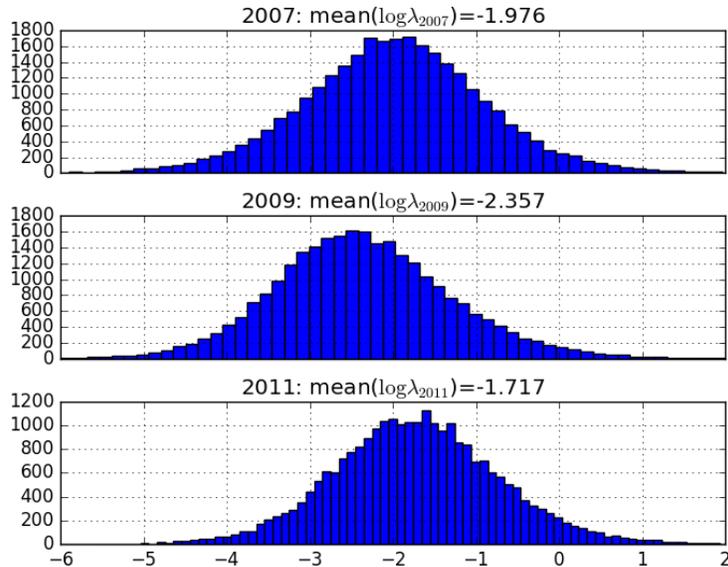


Figure 2: Histogram of  $\log \lambda_{it}$  by year.

model of household demand, consumption smoothing behavior is based on an objective of smoothing expected utility over time, typically expressed as some martingale condition. We lean on this insight as we modify the Ligon-Schechter formula to use the constant Frisch elasticity (CFE) demand system we’ve described, taking  $\log \lambda_{it}$  as the key parameter. Following the related literature, we call the measure “Vulnerability”.

We start by defining a momentary utility function, then estimating an upper bound on household Vulnerability (which will be biased upwards when the expenditure panel is measured with error). Then we decompose this variation by observable characteristics, providing a lower bound on welfare loss due to total vulnerability and some insight into of the specific role played by income variation.

## A Measure of Vulnerability

Maintaining our notation, we index households by  $i$ , goods by  $j$ , periods by  $t$ , and markets by  $m$ . We assume households within a given market share a preference structure represented by a utility function  $U$ . Ligon and Schechter (2003) and papers take  $U$  to be a univariate function of consumption, and

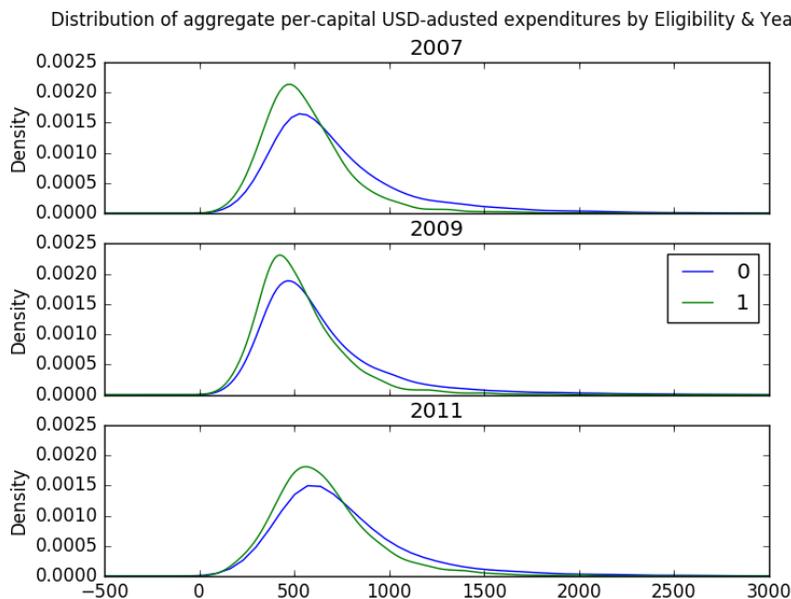


Figure 3: Distribution of  $\log \lambda_{it}$  by HH Eligibility. We can see that the eligible households have consistently higher marginal utilities in each year.

treat vulnerability as a shortfall of a household's expected utility from some certainty-equivalent level  $\bar{C}$ , yielding the formula

$$U(\bar{C}) - EU(C_i) \quad (7)$$

The logic of this approach is that, since  $U$  is concave and increasing in  $C_i$ , a household is taken as less vulnerable as the mean of  $C_i$  increases or as the variance decreases. We take the same approach with a Frischian indirect utility function of  $\lambda$  and prices,  $U(\lambda, p)$ . This approach is appealing in part since, in the Euler equation derived in a dynamic household model, the household's basic goal is to borrow, save, and invest such that  $\lambda_i$  is steady over time. Letting vulnerability be denoted  $\mathcal{V}(\lambda_i)$ , this yields

$$\mathcal{V}(\lambda_i, p_m) = U(\bar{\lambda}, p) - EU(\lambda_i, p) \quad (8)$$

The question arises what value is appropriate for  $\bar{\lambda}$ . Poverty lines frequently identify an absolute level of welfare or some standard basket of goods which social planners regard as significant, or which marks some real-world

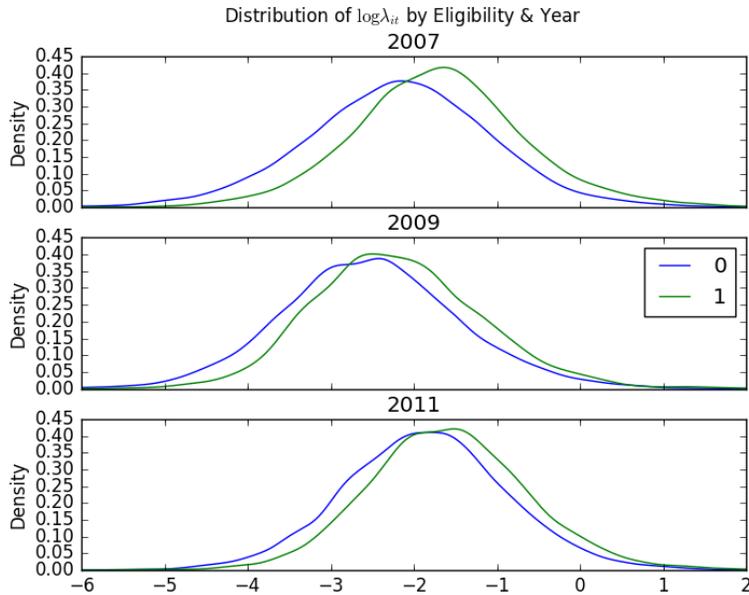


Figure 4: Distribution of  $\log \lambda_{it}$  by HH Eligibility. We can see that the eligible households have consistently higher marginal utilities in each year.

distinction between distinct classes. Lacking the sort of nuanced contextual information such a judgement requires, we follow the approach in Ligon and Schechter (2003), where they set  $\bar{C}$  as the average level of consumption in the population overall, by letting  $\bar{\lambda}$  be the mean level of  $\lambda$  across all household-year observations. This gives a nice interpretation to  $V(\lambda_i, p) = 0$ , which indicates that in each period, household  $i$  achieves the average level of welfare found in the panel, and reaches it without any uncertainty.

We can break this expression into two conceptually distinct quantities, which we'll be calling "Poverty" and "Risk":

(9)

$$V(\lambda_i, p) = [U(\bar{\lambda}, p) - U(E[\lambda_i], p)] \text{ (Poverty)} \\ + [U(E[\lambda_i], p) - E[U(\lambda_i, p)]] \text{ (Risk)}$$

The first bracketed quantity represents the welfare shortfall (or surplus) experienced by household  $i$  in an *average* period, relative to  $\bar{\lambda}$ . We follow the

original authors and refer to this as “Poverty” (though this is not to imply that the final measure of vulnerability is not itself a worthwhile poverty metric). The second value, which we refer to as “Risk”, is the well-known quantity associated with Jensen’s inequality, and measures the welfare loss due to stochastic variation in  $\lambda_{it}$  due to frictions that prevent ideal smoothing (e.g. behavioral biases, credit constraints).

### An Indirect Frisch Utility Function

Before exploring this further, it will serve to specify the particular utility function we will be considering. The demand system described thus far can be represented by the separable direct utility function

$$U(C) = \sum_{j \in J} \alpha_j \frac{\beta_j}{\beta_j - 1} \left[ c_j^{\frac{\beta_j - 1}{\beta_j}} - 1 \right] \quad (10)$$

Here  $C$  is a vector of consumption of  $n$  goods, with  $\beta_j$  governing concavity for any given good, and  $\alpha_j$  governing its expenditure share. We like the flexibility of this function, which allows elasticities to vary across goods. However, our application calls instead for the Frischian *indirect* utility function. We simply take the corresponding demand functions for each good  $j$

$$\widehat{c}_{ijt}(\lambda_{it}, p_{ijt}) = \left( \frac{p_{ijt} \lambda_{it}}{\alpha_j} \right)^{-\beta_j} \quad (11)$$

and plug them into (10). Using estimated  $\lambda_{it}$  for all consumers,  $\beta_j$  and  $\alpha_j$  for all goods, and infer from stated unit values the disaggregated price panel  $p_{ijt}$ , we finally have the indirect Frisch utility function

$$U(\lambda_{it}, p_{ijt}) = \sum_{j \in J} \alpha_j \frac{\beta_j}{\beta_j - 1} \left[ \left( \frac{p_{ijt} \lambda_{it}}{\alpha_j} \right)^{1 - \beta_j} - 1 \right] \quad (12)$$

We then take  $E[\lambda_i]$  to be household  $i$ ’s average level of  $\lambda_{it}$  over time,  $(1/T) \sum_{t \in T} \widehat{\lambda}_{it}$ . We can define our relative poverty line accordingly as the average of  $E[\lambda_i]$  over all households,  $\bar{\lambda} = (1/N) \sum_{i \in I} E[\lambda_i]$ , the average marginal utility over both households and periods.

A final consideration is the role of price variation over time and across markets, which enters into utility in parallel with  $\lambda_{it}$ . By construction, households’ marginal utilities are orthogonal to prices, since they are included in the regression whose residuals are used to estimate  $\log \lambda$ . What’s more, This

vulnerability is designed to specifically capture the role of uncertainty in this measure of momentary household welfare. As such, we choose to leave intertemporal price variation to the side by holding the price of each good constant at the average level within each market over the course of the panel, which we'll call  $\overline{p_{jm}}$ . This lets us avoid conflating the importance of consumption smoothing and that of price variation (the welfare consequences of each being quite different), while also leaving out potentially significant amounts of measurement error in our price estimates.

So within a given market (where prices are shared), the distribution of vulnerability is wholly determined by variation in households' marginal utilities. Since this is an additively separable function, our vulnerability measure is similarly separable, with each sub-function having a different (and empirically estimated) curvature parameter. Our index thus preserves the essential feature that expected utility decreases as variance in the distributions of  $\lambda_{jt}$  or  $p_{jmt}$  increase, but unlike the Ligon-Schechter index, it does so in proportion with  $\beta_j$ , which is allowed to differ across goods and which is estimated empirically, rather than being imposed as an assumption of the model. Expected utility will also be unaffected by variation due to measurement error and substitution outside of their effect on our estimates of household neediness.

## Measuring Kinds of Risk

With this demand system and associated utility function in hand, we can consider how measured risk can be usefully decomposed to illuminate the association between variation in welfare around and various observable outcomes and characteristics. In particular, we will be interested to isolate the portion of risk associated with variation in income from different sources, as this is the primary channel by which program participation might influence household welfare.

We follow Ligon & Schechter in decomposing variation in welfare according to aggregate and idiosyncratic observables. Where  $k$  indexes individual-level observables, we denote aggregate economic variables by  $\overline{x_{mt}}$  and idiosyncratic variables as  $x_{itk}$ , respectively. We treat market areas as the natural unit over which aggregate shocks are distributed. An ordered sequence of  $K$  observables is then selected and  $\lambda_{jt}$  is then projected onto the first  $k$  variables in the list for  $k \in 1, \dots, K$ . In this case, we include farm income (including livestock), non-farm income, and value of the total reported asset stock.

Putting all of these parts together, we finally have the formula

(13)

$$\begin{aligned}
V(\lambda_{it}) &= U(\bar{\lambda}, \bar{p}_{jm}) - U(E_t(\lambda_i), \bar{p}_{jm}) \text{ (Poverty)} \\
&+ U(E_t(\lambda_i), \bar{p}_{jm}) - EU(E_t(\lambda_i | \bar{x}_{mt}), \bar{p}_{jm}) \text{ (Aggregate Risk)} \\
&+ EU(E_t(\lambda_i | \bar{x}_{mt}), \bar{p}_{jm}) - EU(E_t(\lambda_i | \bar{x}_{mt}, x_{it}^1), \bar{p}_{jm}) \text{ (}\downarrow\text{Idiosyncratic Risk)} \\
&\vdots \\
&+ EU(E_t(\lambda_i | \bar{x}_{mt}, x_{it}^1, \dots, x_{it}^{k-1})) - EU(E_t(\lambda_i | \bar{x}_{mt}, \dots, x_{it}^k)) \\
&+ EU(E_t(\lambda_i | \bar{x}_{mt}, \dots, x_{it}^k)) - EU(\lambda_i) \text{ (Unexplained Risk)}
\end{aligned}$$

Consider that if  $\lambda_{it}$  or  $p_{jmt}$  is measured with error, the term for unexplained risk will be biased upward, since variation due to measurement error will be conflated with the welfare-reducing variation due to uncertainty. As such, Equation (13) will represent an upper bound on welfare loss due to risk. Excluding the last line from the expression will thus represent “explained” risk, and will itself be a lower bound on the true value of  $V(\lambda_{it})$ .

In the case of one explanatory variable (namely income), we can reduce the expression to a more concise form:

(14)

$$\begin{aligned}
V(\lambda_{it}) &= U(\bar{\lambda}, \bar{p}_{jm}) - U(E_t(\lambda_i), \bar{p}_{jm}) \text{ (Poverty)} \\
&+ U(E_t(\lambda_i), \bar{p}_{jm}) - EU(E_t(\lambda_i | \bar{x}_{mt}), \bar{p}_{jm}) \text{ (Aggregate Risk)} \\
&+ EU(E_t(\lambda_i | \bar{x}_{mt}), \bar{p}_{jm}) - EU(E_t(\lambda_i | \bar{x}_{mt}, income_{it}^1), \bar{p}_{jm}) \text{ (Income Risk)}
\end{aligned}$$

### Estimating Vulnerability

We can briefly specify the empirical analogues to the conditional expectations specified above. The unconditional expectations  $\bar{\lambda}$  and  $E_t(\lambda_i)$  are simply sample averages. Meanwhile,  $E_t(\lambda_i | \bar{x}_{mt})$  is calculated as the average level of  $\lambda_{it}$  for each good in a household’s market area  $m$  and period  $t$ . Finally, the conditional expectations, conditioning on the first  $k$  observables,  $E_t(\lambda_i | \bar{x}_{mt}, x_{it}^1, \dots, x_{it}^k)$  will be calculated as fitted values of a linear model relating marginal utility to household-time fixed effects and those  $k$  variables. Thus for the  $k$ ’th row of Equation (13) related to idiosyncratic risk, we estimate the parameters of

$$\lambda_{ijt}^k = \eta_t + \delta_i + \beta X_{it}^k + e_{ijt} \quad (15)$$

where  $X_{it}$  is a matrix of the first  $k$  observable variables. The fitted values are then used as the conditional expectations in (13). When we exclude “Unexplained risk”, we call this “Explained Vulnerability”, which in our case is the sum of *Poverty* and *Income Risk*.

### Distribution of vulnerability

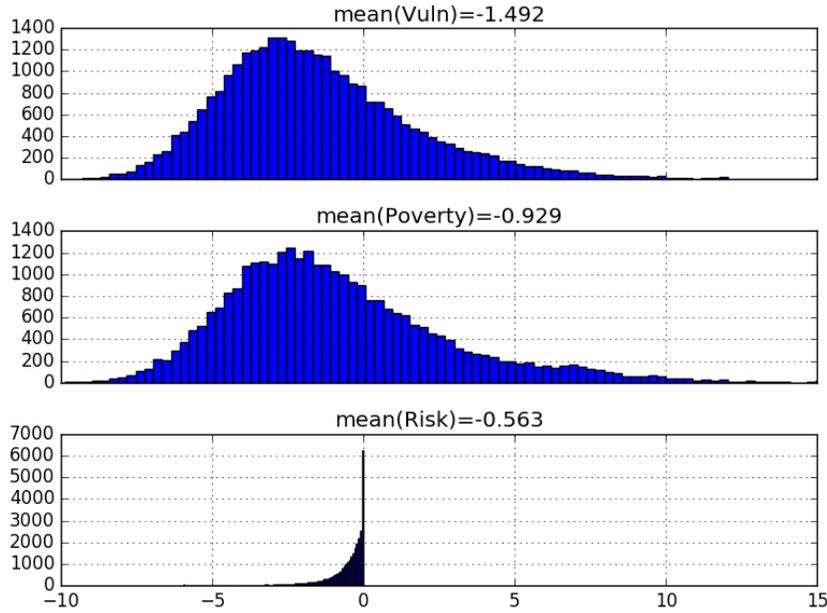


Figure 5: Histograms of Vulnerability, Poverty, and Risk for the entire Bangladesh Sample

Figure 5 shows us the distribution of total vulnerability based on (13) using estimated  $\lambda_{it}$ 's. As we'd expect, the distribution is log-normal, with the risk component strictly negative. We can see that Risk makes up a relatively small portion of total Vulnerability. This indicates that, according to our model, nearly all of a household's deviation from the mean level of welfare can be attributed to the “average” level over time, with the welfare loss due to year-to-year uncertainty accounting for much less. The last histogram shows us the distribution of “Income Risk”, which is what the distribution

of risk looks like if we consider only the variation in welfare explained by variation in come.

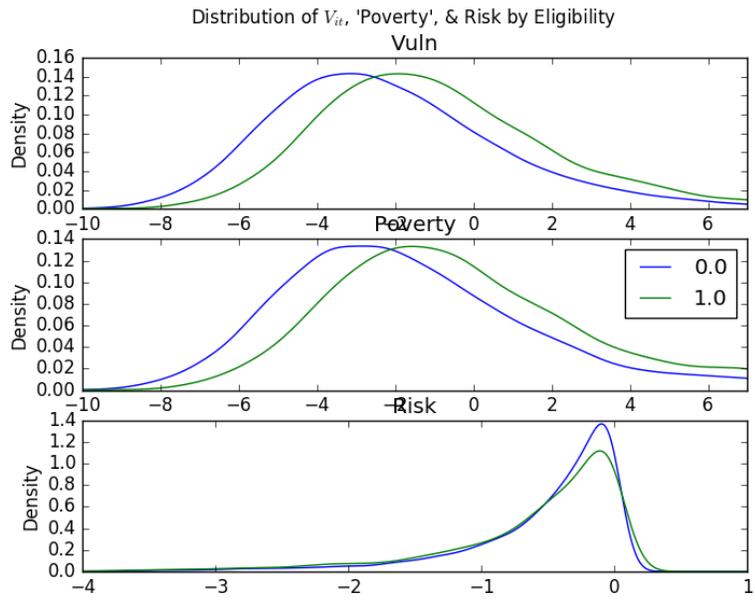


Figure 6: Comparing Vulnerability, Poverty, and Risk by Eligibility

To better understand how the distribution of vulnerability relates to the characteristics of the sample, we also look at cross-sectional regressions of each component of vulnerability on a range of covariates measured at baseline. Table 2 estimates the joint linear model with all of these covariates to show us how each relates to vulnerability, holding several important factors constant. Table 3 reports the coefficients and standard errors for each univariate regression to show us simple unconditional associations.

As expected, most of the variation in vulnerability can be explained by variation in poverty, with coefficients on risk proving relatively small in each case. Unsurprisingly, we find that wealthier households rank lower on this poverty index. Larger households also appear more vulnerable, especially those with many adults. However, those same households appear to face less risk overall. Female and older household heads tend to be poorer as well, but we cannot say that they face different levels of risk. Cash Savings, borrowing activity, and productive assets (including cows) are all associated with lower poverty overall. Interestingly, households with cows also appear to face more risk, which will prove consistent with the TUP program’s results in the next section.

## Treatment Effects on Welfare

Our final question is to what extent various measures of household welfare can be seen to vary between treatment and control groups. We start with aggregate consumption, food consumption, and marginal utilities.

### Consumption and Marginal Utility

#### Empirical Specification

Treatment effects are estimated using the random assignment of households to treatment or control villages. Following Bandiera et al. (2017), we include subdistrict fixed effects (since this is the level of stratification), with standard errors clustered at the branch office level (since this is the largest geographic unit which might plausibly have economic spillovers from the program). This yields the specification

$$Outcome_{itm} = \alpha + \sum_{t \in (1,2)} \beta_t (Y_t \times T_{im}) + \gamma T_{im} + \sum_{t \in (1,2)} \delta_t Y_t + \eta_m + \epsilon_{itm} \quad (16)$$

where  $T_{im}$  is the treatment status of individual  $i$  in market or subdistrict  $m$  and  $Y_t$  is an indicator as to whether it is period or year  $t$ . Periods  $t=0$ ,

1, and 2 refer to 2007, 2009, and 2011, respectively. Here, the interaction terms provide average intent-to-treat effect estimates for 2009 and 2011, representing the effect of the program after two and four years, respectively.

### Results

Finally, we find that when comparing eligible households, treatment appears to have led to a reduction in welfare in the short-term, but improved welfare in 2011 4 years after enrollment (Table 4). Treatment is associated with a statistically significant *increase* of .14 standard deviations in 2009, followed by a fall in  $\log \lambda_{it}$  in 2011 of 0.22 standard deviations. The first column of Figure 8 presents this shift graphically, showing the fall in  $\lambda_{it}$  in the final year. We find similar results when looking at aggregate consumption. A small (insignificant) fall in consumption in 2009, followed by a considerable rise in 2011. Aggregate consumption is calculated as described in Table 5 of Bandiera et al. (2017), but unfortunately, we are not able to replicate their results exactly. In 2009, they find the treatment group has higher (though still insignificant) consumption, and their 2011 point estimate is around 40% higher than ours. We also find that aggregate food consumption is improved, again only in the 4-year results from 2011.

Comparing ineligible households with and without neighbors enrolled in the TUP program (Table 5), we find that they were actually made better off in both years according to the marginal utility results. However, this does not match the per-capita expenditure results, which only statistically significant results for food consumption in 2011, where they are negative.

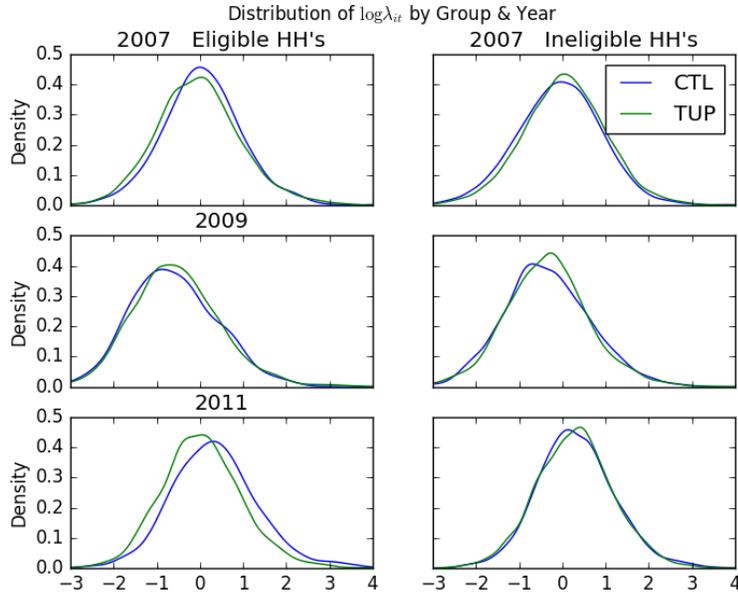


Figure 7: Density of  $\log \lambda_{it}$  by group and year. We can see that the means are close in 2007 (at baseline), and there is a marked downward shift in neediness for the treatment group in 2009 & 2011. (Plot currently censored at 1st and 99th percentiles to exclude extreme outliers)

## Vulnerability, Risk, and Poverty

### Empirical Specification

In the case of our Vulnerability estimates (and the related components Poverty and Risk), our specification must be different from that of the published results since we are relying on time-series variation. Instead, we're using the much simpler cross-sectional model

$$Outcome_{im} = \eta_m + \beta T_{im} + \epsilon_{im} \quad (17)$$

It's worth reflecting again on the definitions of these metrics and how we should interpret the related treatment effect parameters. Poverty, for example, is defined as

$$[U(z, p) - U(E[\lambda_i], p)]$$

which we can see as the difference between the market-wide average and

households’ level of welfare in an “average” year. So a reduction in this value suggests an improvement in welfare overall across the duration of our panel.

Risk, on the other hand, which we define as

$$[U(E[\lambda_i], p) - E[U(\lambda_i, p)]]$$

provides a measure of welfare loss due to period-to-period variation in  $\lambda_{it}$  in the face of concave preferences. This is a loss *relative* to the household’s own certainty-equivalent level of welfare, and so should be strictly negative. So a reduction in this value indicates that a household’s welfare has been made less variable over time due to the treatment. The magnitude of this effect will naturally be more pronounced for the poor, who live on a more concave region of the utility function. It may at first feel counter-intuitive that a sudden exogenous increase in welfare will appear harmful along this axis. The intuition stems from the fact that this is an unequivocally beneficial shock, which is captured in our measure of *Poverty*. The increase in “Risk” stems from the way this benefit is diminished by the fact that it can’t be smoothed out retroactively. This may be because it was unforeseen, its future benefits were too uncertain to borrow against, or because the household is subject to some binding financial constraint. Either way, the benefit of such a transfer is less than one that can be saved and borrowed against in a complete capital market. In this sense, our risk index proxies how nice it would have been to be able to get a smaller benefit on a more consistent basis.

This somewhat mechanical feature of our model of household risk does not imply that the TUP program must increase risk overall, though. Thinking back to the motivating idea behind the TUP framework, having a large productive asset may well make household welfare less variable overall by providing a more consistent source of income or by improving access to credit. On the other hand, substituting away from relatively predictable farm labor towards household production is by no means guaranteed to reduce year-to-year variation in incomes. The treatment thus relates to our Risk measure in two conceptually distinct ways, first that the benefit of the transfer is diminished by its unpredictability, and second that the transfer may affect the underlying welfare uncertainty faced by a given household.

The aggregate impact of both of these factors on eligible households is reported in Table 6, with spillovers reported in Table 7. However, we can also try to disentangle the results by restricting our estimation to the two post-treatment years. This excludes from our analysis the initial exogenous shock to program households, allowing us to focus only on variation in welfare

due to factors shared across treatment and control villages. Thus treatment effects on this post-treatment variation may be seen as more directly checking whether the treatment reduced the level of welfare uncertainty faced by households. These results are listed in Table 8, with spillovers reported in Table 9.

### **Results**

Finally moving to the results, we find that overall Vulnerability is significantly lower for the treatment households. This corresponds to an even larger fall in Poverty, since we find that Risk is actually higher (in line with our discussion above). However, looking specifically at the risk that can be explained by household income, we find that households are slightly less exposed to our measure of income risk, though as seen in the histograms above, this ends up being a fairly small effect relative to the program's effect on poverty. When we restrict our sample to the two post-treatment years of the panel as discussed above, we find qualitatively similar results on all measures. The Risk measures are of particular interest in Table 8. Here we see that risk is even higher than in the full sample, suggesting that welfare uncertainty from 2007 to 2009 was notably lower than that from 2009 to 2011, even accounting for the precipitous shift upwards caused by the BRAC's programs.

Overall, there do not appear to be spillover effects on vulnerability among ineligible households in treated areas. While the point estimates are not stable between the two-year and full-sample estimates, neither finds statistically significant results. The notable exception is in our measure of income risk. However, the point estimates found thus far appear somewhat implausible, perhaps pointing to a flaw in how we use income to derive conditional expectations in welfare.

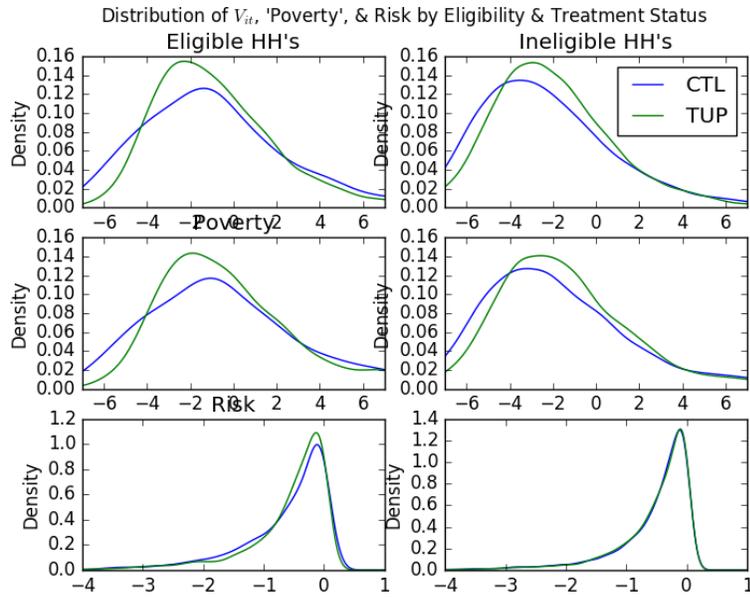


Figure 8: Density of Vulnerability and its components, Risk and Poverty. “1” indicates that a household lived in a treatment area. The columns show the distributions among eligible and ineligible households, respectively. We can see that the distributions appear broadly similar among ineligible households. Among eligible households, being in a treatment area is associated with lower vulnerability and poverty, and slightly higher ( $p < .05$ ) but broadly similar levels of Risk exposure.

## Conclusion

In conclusion, we have developed and explored several theoretically motivated methods of looking at household welfare. These relax the stringent data constraints of standard aggregate consumption measures by allowing for estimation of welfare with only a subset of goods, and without the observation of prices for particularly problematic or difficult items. In exchange, we require disaggregate consumption or expenditure data for estimating  $\log \lambda_{it}$ , and a panel of such goods with at least two prices to study vulnerability and risk.

Taking these measures to the nation-wide evaluation of the original TUP program, we find that after two years, treatment household marginal util-

ity has actually increased, indicating reduced household welfare. After four years, however, marginal utility is significantly lower. Our estimates of aggregate consumption and total food consumption both find no statistically significant change after two years, and a precipitous increase after four. We have not yet successfully replicated the point estimates found in Bandiera et al. (2017), which see higher consumption among treatment households in both 2009 and 2011, but our conclusions are broadly similar since they do not find statistically significant two-year effects either. Looking at spillovers in household welfare, we only find statistically significant effects for food consumption, and only in 2011. However, we find that marginal utilities are significantly lower (by at least .1 standard deviations) in both years.

Turning to treatment effects on vulnerability, risk, and poverty, we find that treatment households are much less vulnerable overall due to decreased poverty (i.e. a smaller shortfall in average welfare), and face slightly lower income risk specifically, but slightly more risk overall. Ineligible households may see less income risk as well, though the estimated magnitudes may seem implausibly large and are not borne out in any of our estimates of unconditional vulnerability or risk.

Future developments of this investigation will hopefully look more closely at quantile effects, especially in light of Figure 3 and Figure 4. These figures suggest that the ultra-poor targeting exercise used by BRAC did successfully identify poorer-than-average households. But it also suggests that the welfare of the two groups is drawn from generally similar ranges, raising the familiar question of how valuable the means testing element of the program really was. One wonders how much greater could the total treatment effect have been under retrospectively optimal targeting, or if the trouble of means testing had been replaced by simple lottery, without any attempt to target the ultra-poor.

## References

- Bandiera, Oriana et al. (2017). “Labor markets and poverty in village economies”. In: *The Quarterly Journal of Economics* 132.2, pp. 811–870.
- Banerjee, Abhijit et al. (2015). “A multifaceted program causes lasting progress for the very poor: Evidence from six countries”. In: *Science* 348.6236, p. 1260799.

- Blattman, Christopher, Nathan Fiala, and Sebastian Martinez (2013). “Generating skilled self-employment in developing countries: Experimental evidence from Uganda”. In: *The Quarterly Journal of Economics* 129.2, pp. 697–752.
- Collins, Elliott and Ethan Ligon (2017). “Asset transfers and household neediness”. CUDARE Working Paper. URL: <http://escholarship.org/uc/item/24t720fk>.
- Ligon, Ethan (2016a). “All  $\lambda$ -separable Frisch demands and corresponding cardinal utility functions”. CUDARE Working Paper. URL: <http://escholarship.org/uc/item/1w13q2f1>.
- (2016b). “Estimating Household Neediness From Disaggregate Expenditures”. Unpublished manuscript.
- Ligon, Ethan and Laura Schechter (2003). “Measuring Vulnerability”. In: *Economic Journal* 113.486, pp. C95–C102. URL: <http://dx.doi.org/10.1111/1468-0297.00117>.

Table 1: Estimated elasticities of demand,  $(\phi\beta_j)$  with respect to  $\lambda_{it}$ . (For goods with %Zero<95)

	$\phi\beta_i$	%Zero	logalpha	Girls	Boys	Men	Women	log HHSIZE
Fruit	0.63*** (0.01)	74.90	3.32*** (0.01)	-0.11*** (0.01)	-0.08*** (0.01)	-0.02* (0.01)	-0.01 (0.01)	0.72*** (0.03)
Fish	0.58*** (0.01)	28.80	3.94*** (0.00)	-0.06*** (0.01)	-0.04*** (0.01)	0.07*** (0.00)	0.04*** (0.00)	0.50*** (0.02)
Sugar	0.46*** (0.01)	79.60	2.49*** (0.01)	-0.06*** (0.01)	-0.07*** (0.01)	0.01 (0.01)	0.03*** (0.01)	0.38*** (0.03)
Garlic	0.45*** (0.01)	52.10	1.51*** (0.00)	-0.02*** (0.01)	-0.01** (0.01)	0.05*** (0.00)	0.06*** (0.00)	0.25*** (0.02)
Onion	0.44*** (0.00)	4.90	1.90*** (0.00)	-0.00 (0.00)	0.01*** (0.00)	0.09*** (0.00)	0.08*** (0.00)	0.27*** (0.01)
Vgtbl	0.44*** (0.01)	17.90	2.55*** (0.00)	-0.02*** (0.00)	0.00 (0.01)	0.06*** (0.00)	0.04*** (0.00)	0.46*** (0.01)
Milk	0.41*** (0.01)	79.00	3.15*** (0.01)	-0.10*** (0.01)	-0.09*** (0.01)	-0.02*** (0.01)	-0.01* (0.01)	0.72*** (0.03)
Oil	0.40*** (0.00)	1.80	2.56*** (0.00)	-0.01*** (0.00)	0.01*** (0.00)	0.10*** (0.00)	0.08*** (0.00)	0.32*** (0.01)
Eggs	0.34*** (0.01)	74.40	2.80*** (0.01)	-0.06*** (0.01)	-0.04*** (0.01)	-0.00 (0.01)	-0.01 (0.01)	0.49*** (0.02)
Spices	0.33*** (0.01)	65.40	2.14*** (0.00)	-0.04*** (0.01)	-0.04*** (0.01)	0.00 (0.01)	0.03*** (0.01)	0.42*** (0.02)
Lentils	0.31*** (0.01)	79.60	2.79*** (0.01)	-0.04*** (0.01)	-0.03*** (0.01)	0.01** (0.01)	0.03*** (0.01)	0.39*** (0.02)
Nuts	0.28*** (0.01)	41.80	2.50*** (0.00)	-0.02*** (0.00)	-0.03*** (0.01)	0.05*** (0.00)	0.05*** (0.00)	0.20*** (0.01)
Leafy	0.19*** (0.00)	28.20	1.74*** (0.00)	-0.00 (0.00)	0.00 (0.00)	0.03*** (0.00)	0.03*** (0.00)	0.31*** (0.01)
Rice	0.19*** (0.00)	0.10	4.10*** (0.00)	-0.06*** (0.00)	-0.04*** (0.00)	0.01*** (0.00)	-0.01*** (0.00)	0.86*** (0.01)

Table 2: Joint linear models associating a range of baseline observables with each component of Vulnerability.

	Vuln	Poverty	Risk	IncRisk
WealthRnk	-0.222*** (0.018)	-0.220*** (0.020)	-0.001 (0.004)	0.500*** (0.161)
girls	0.087*** (0.019)	0.086*** (0.022)	0.001 (0.004)	-0.455*** (0.172)
boys	0.037* (0.019)	0.028 (0.021)	0.009* (0.004)	-0.371** (0.171)
men	0.246*** (0.031)	0.265*** (0.035)	-0.019*** (0.007)	4.020*** (0.278)
women	0.367*** (0.034)	0.384*** (0.038)	-0.018** (0.008)	0.666** (0.302)
age	0.008*** (0.001)	0.010*** (0.002)	-0.002*** (0.000)	-0.020 (0.012)
gender	0.231*** (0.059)	0.248*** (0.065)	-0.017 (0.013)	2.786*** (0.518)
pension	0.012 (0.048)	0.022 (0.053)	-0.010 (0.011)	0.640 (0.421)
OwnLand	-0.613*** (0.053)	-0.625*** (0.058)	0.012 (0.012)	-1.882*** (0.467)
LandSize	-0.002*** (0.000)	-0.002*** (0.000)	0.000*** (0.000)	0.034*** (0.002)
Save(\$100)	-0.044*** (0.006)	-0.047*** (0.007)	0.003** (0.001)	0.273*** (0.054)
AnySave	-0.062 (0.040)	-0.073 (0.044)	0.012 (0.009)	-0.157 (0.355)
Debt(\$100)	-0.032*** (0.004)	-0.033*** (0.004)	0.001 (0.001)	0.055 (0.034)
AnyDebt	-0.118*** (0.044)	-0.104** (0.049)	-0.014 (0.010)	-0.912** (0.388)
Food Sec	-0.516*** (0.042)	-0.548*** (0.046)	0.033*** (0.010)	-0.252 (0.368)
# Cows	-0.132*** (0.015)	-0.144*** (0.017)	0.012*** (0.003)	-0.087 (0.135)
Assets(\$100)	-0.005*** (0.001)	-0.006*** (0.001)	0.000 (0.000)	0.122*** (0.009)
N	28780.000	28780.000	28780.000	28780.000

Table 3: The coefficients for each univariate linear model, showing the simple sample-wide association between Vulnerability and each household features at baseline.

	Vuln	Poverty	Risk	IncRisk	N
WealthRnk	-0.650*** (0.013)	-0.670*** (0.014)	0.020*** (0.003)	3.413*** (0.119)	30530.000
age	0.004*** (0.001)	0.006*** (0.002)	-0.002*** (0.000)	0.103*** (0.012)	28836.000
gender	0.843*** (0.047)	0.888*** (0.052)	-0.045*** (0.010)	-3.356*** (0.406)	28837.000
pension	0.483*** (0.049)	0.517*** (0.054)	-0.034*** (0.010)	-1.875*** (0.432)	30577.000
OwnLand	-1.841*** (0.041)	-1.899*** (0.045)	0.059*** (0.009)	8.079*** (0.367)	30577.000
Save(\$100)	-0.156*** (0.006)	-0.164*** (0.006)	0.007*** (0.001)	1.096*** (0.050)	30577.000
Debt(\$100)	-0.088*** (0.003)	-0.091*** (0.004)	0.002*** (0.001)	0.386*** (0.031)	30577.000
Food Sec	-1.330*** (0.039)	-1.391*** (0.042)	0.061*** (0.008)	4.266*** (0.347)	30577.000
# Cows	-0.464*** (0.012)	-0.485*** (0.013)	0.021*** (0.003)	2.835*** (0.106)	30577.000
Assets(\$100)	-0.028*** (0.001)	-0.029*** (0.001)	0.001*** (0.000)	0.272*** (0.007)	30595.000

Table 4: Year interactions with  $T$  show treatment effects.

Treatment	$\lambda_{it}$	ExpUSD	Food
T*Y11	-0.216*** (0.050)	46.469** (18.837)	31.675*** (7.989)
T*Y09	0.141*** (0.049)	-12.323 (18.651)	0.248 (7.910)
T	0.000 (0.014)	388.999*** (5.315)	130.854*** (2.254)
Y09	-0.661*** (0.042)	242.092*** (15.739)	116.065*** (6.675)
Y11	0.276*** (0.042)	329.599*** (15.860)	173.124*** (6.726)
const	Mkt-FE	Mkt-FE	Mkt-FE
Mean <sub>2007</sub>	0.0 (1.0)	553.83 (284.63)	186.46 (107.28)
N	11709.000	11709.000	11709.000

Table 5: Year interactions with  $T$  show treatment effects.

Spillover	$\lambda_{it}$	ExpUSD	Food
T*Y11	-0.138*** (0.025)	-5.529 (16.085)	-20.861*** (7.523)
T*Y09	-0.110*** (0.024)	3.833 (15.894)	-10.697 (7.434)
T	0.001 (0.008)	402.502*** (5.427)	174.480*** (2.538)
Y09	-0.299*** (0.017)	404.559*** (11.096)	238.138*** (5.190)
Y11	0.373*** (0.017)	542.389*** (11.208)	342.705*** (5.242)
const	Mkt-FE	Mkt-FE	Mkt-FE
Mean <sub>2007</sub>	0.0 (1.0)	709.82 (444.4)	307.56 (207.52)
N	33590.000	33590.000	33590.000

Table 6: Treatment effects for enrolled households on Vulnerability, Poverty, Risk, and Income Risk

<b>ATE</b>	Vuln	Poverty	Risk	IncRisk
T	-0.526*** (0.105)	-0.614*** (0.119)	0.088*** (0.028)	-0.032*** (0.010)
const	-0.302*** (0.091)	0.452*** (0.103)	-0.754*** (0.024)	0.199*** (0.009)
Mean	-1.98 (2.94)	-1.44 (3.24)	-0.54 (0.65)	3.66 (28.83)
N	3583.000	3583.000	3583.000	3583.000

Table 7: Spillover effects for ineligible households on Vulnerability, Poverty, Risk, and Income Risk

<b>Spillovers</b>	Vuln	Poverty	Risk	IncRisk
T	0.074 (0.055)	0.079 (0.062)	-0.005 (0.013)	-1.130** (0.539)
const	-1.956*** (0.041)	-1.362*** (0.046)	-0.594*** (0.010)	4.179*** (0.404)
Mean	-0.69 (2.96)	-0.07 (3.29)	-0.62 (0.72)	0.2 (0.26)
N	10437.000	10437.000	10437.000	10437.000

Table 8: Treatment effects excluding initial baseline data for enrolled households on Vulnerability, Poverty, Risk, and Income Risk

<b>ATE (no BL)</b>	Vuln	Poverty	Risk	IncRisk
T	-0.701*** (0.133)	-0.860*** (0.143)	0.160*** (0.027)	0.037 (0.156)
const	-0.279** (0.115)	0.348*** (0.124)	-0.627*** (0.023)	0.226* (0.135)
Mean	-2.03 (3.42)	-1.61 (3.66)	-0.43 (0.63)	3.6 (32.21)
N	3583.000	3583.000	3583.000	3583.000

Table 9: Spillover effects excluding initial baseline data for ineligible households on Vulnerability, Poverty, Risk, and Income Risk

<b>Spillovers (no BL)</b>	Vuln	Poverty	Risk	IncRisk
T	-0.078 (0.066)	-0.095 (0.071)	0.017 (0.013)	0.937* (0.559)
const	-1.873*** (0.050)	-1.397*** (0.053)	-0.476*** (0.010)	2.722*** (0.418)
Mean	-0.94 (3.56)	-0.49 (3.82)	-0.45 (0.68)	0.26 (3.83)
N	10437.000	10437.000	10437.000	10437.000