Theoretical Approaches to Measuring the Welfare Effects of Asset Transfers

by

Elliott Michael Collins

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Agricultural and Resource Economics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Ethan Ligon, Chair
Professor Jeremy Magruder
Professor Edward Miguel

Summer 2017

Theoretical Approaches to Measuring the Welfare Effects of Asset Transfers

Abstract

Theoretical Approaches to Measuring the Welfare Effects of Asset Transfers

by

Elliott Michael Collins

Doctor of Philosophy in Agricultural and Resource Economics

University of California, Berkeley

Professor Ethan Ligon, Chair

Millions of low-income households have received large one-time transfers of capital as part of private and public anti-poverty programs. Economists have studied the effect of such transfers in a wide variety of settings to understand their effect on household welfare and poverty rates. The standard metric used to study these effects is aggregate consumption, measured in price-adjusted dollars per capita per day, which in turn gets used to define the global poverty lines at \$2 and \$1.25. In this dissertation, I discuss some of the theoretical and practical issues that these measures face, and then suggest some alternatives pulled from a model of consumer demand. I demonstrate how they can be used to expand our understanding of two programs of direct capital transfers to poor households, with the broad goal of bringing insights from economic theory farther into the practice of impact evaluation.

This dissertation has as its chapters three papers, each written so that they stand alone for those interested in a particular part of the overall project. The first chapter develops a Frischian model of household expenditures, and uses this model to derive estimates of an individual household's marginal utility in expenditures at any given time, $\log \lambda_{it}$, which serves as the central welfare metric in each chapter. The application of this model to measure the welfare effects of asset transfers is illustrated using the early results of a pilot in South Sudan of the so-called "Targeting the Ultra-Poor" (TUP) program. It concludes that the large asset transfers improved welfare in the short-run.[1]

Chapter 2 uses the methods in Chapter 1 to reanalyze the first large experimental evaluation of the TUP framework in Bangladesh, previously studied in Bandiera et al. (2017). It goes on to estimate a flexible direct utility function which allows one to estimate welfare in a way that accounts for the immiserating effects of uncertainty faced by poor households. It concludes that this TUP program improved household welfare years after the transfers, and that this benefit was accompanied by a lower level of risk.

Finally, Chapter 3 goes on to study the experiment in South Sudan in more detail. It compares the TUP treatment to an experimental cash transfer, and showcases the practical value of the methods in Chapter 1 by continuing to cost-effectively monitor household welfare

[1]This chapter is referred to in later chapters as Collins and Ligon (2017).

after the long-form panel survey was completed. It then briefly speaks to how participation in the TUP program may have affected households' response to the outbreak of violent civil conflict in 2014. The results show that both cash and asset transfers improved household welfare only in the short run, but the asset transfers had a sustained positive impact on households' total stock of assets, which may have reduced the likelihood that those households were left without recourse upon the onset of violence.

Contents

List of Figures

List of Tables

Dedicated to my wife, Mora Collins, and my children, Acacia and Elias Collins.

Each of you, to the moon and back.

Acknowledgments

First and foremost, my thanks to Ethan Ligon, dissertation chair, my advisor, and mentor, who co-authored the Chapters 1 and 3, provided indispensable support on every chapter, and whose ideas motivate this thesis and much work to come.

Thanks to the members of my dissertation committee for their service.

Thanks also to BRAC, for implementing every transfer studied in this dissertation and for engaging in a rigorous research program to understand them.

For implementing the experiment and co-authoring the final chapter of this dissertation with me, I thanks Munshi Sulaiman and Reajul Chowdhury. Thanks also to Mushi Sulaiman and Narayan Das for generously sharing the data and code from Bandiera et al. (2017) in support of Chapter 2.

For invaluable comments and discussion throughout the writing of this thesis, my thanks to Eduardo Montoya, Angeli Kirk, Elisabeth Sadoulet, Alain de Janvry, Daniel Tregeagle, and many others.

For bringing me into the discipline and encouraging me to study development economics, I am grateful to Professor Jonathan Robinson.

Thanks also to Kiva.org, for providing the opportunity to continue this research and to apply it to important policy challenges, and for supporting me in the final stages of my work on this dissertation.

So many thanks to Mora Collins, whose love, patience, support, conversation, and intellect allowed me to complete this dissertation.

And finally, my thanks to Acacia and Elias Collins, who bring levity and inspiration to my work every day.

Chapter 1

Asset Transfers and Household Neediness in South Sudan

Abstract

What happens when you give an 'ultra-poor' household a productive asset, with training in how to use it? The answer depends on the ways in which markets are incomplete. Previous studies have found that in some settings this sort of program can have a significant impact on occupational choice and average income. Here we document the effects of such a program in South Sudan, but with a focus on the *welfare* of the household, using a measure related to the household's marginal utility of expenditures, or what Ligon (2016b) calls the households' neediness.

This construction allows us not only to see if the the program has a significant effect on household welfare, but also allows us to draw inferences regarding which households would benefit most from a hypothetical cash transfer. We use the fact that neediness is related not only to consumption expenditures, but also to key variables such as the marginal product of labor, investment, and participation in both market- and self-employment.

We report the results of an experiment which randomly assigns participation in such a program, and find large and significant effects on expenditures and a 0.2 standard deviation reduction in average neediness. These improvements in welfare are mirrored by increases in the number and value of assets held; increases in self-employment and skilled market employment, these last compensated for by a marked decrease in casual agricultural labor (and, less confidently, by an increase in leisure).

Introduction

We consider a program in South Sudan which provides training and productive assets to women in very poor households, which is intended to encouraging these women to create a productive enterprise. We have reasonably good data on the cost of this program to the NGO that has implemented it. Our question: how can we best measure the benefit?

We consider this question from the point of view of a given household. We think of this household as solving a dynamic program by simultaneously making decisions regarding consumption, investment, occupation, and production. All of these decisions are tied together by a quantity Ligon (2016b) calls 'neediness', which is simultaneously equal to the marginal benefit of additional consumption expenditures, time, investment, and inputs to production. We use data on disaggregate household expenditures and methods devised by Ligon (2016b) to measure changes in the logarithm of household neediness.

We find that the program results in a statistically significant 0.2 standard deviation reduction in the average log neediness of the treatment group relative to a control group, mirroring a 6.5 SSP (South Sudanese Pounds; about $1.62 USD) increase in the subset of daily household expenditures we observe. Other changes can be interpreted using our model, even when they're not necessarily predicted by that model. Those changes include some increase in both the number and value of productive assets held by the treated households, and a substitution away from casual agricultural labor into more skilled forms of market labor, self-employment, and perhaps leisure. Importantly, our estimates of these other changes can use estimated household neediness as an additional household covariate, which gives us a simple way to distinguish between 'wealth' and 'substitution' effects due to the treatment.

Background on 'TUP' and Related Interventions

Impoverished women in underdeveloped regions tend to be involved in low-return occupations, and frequently face both financial and human capital constraints (Banerjee and Duflo, 2007). A set of programs designed specifically to reduce poverty typically aim to alleviate these constraints simultaneously is the "ultra-poor graduation" framework, in which very poor individuals are offered both physical capital and some form of training or education to promote a particular kind of microenterprise activity. Broad outcomes for similar programs are described in (Banerjee, Duflo, et al., 2015). Bandiera et al. (2017) describes a large ultra-poor graduation initiative implemented in Bangladesh known as the "Transfers to the Ultra-Poor" (TUP) program. The program was implemented in 2007 by BRAC (cf., `www.brac.org`). Exploiting the randomized pattern of expansion, the study found persistent impacts on productivity, earnings, and participation in microenterprise.

Subsequently, BRAC decided to pilot a TUP program near the town of Yei in South Sudan. This paper uses randomized enrollment to evaluate the effects of this pilot program over the course of a year.[1] In late 2013, the program gave 249 women start-up capital at a marginal cost of around $240. Participants received some form of livestock, agricultural material, or retail inventory. They then participated in training specific to the assets provided and were given periodic food support valued at $110.

[1]A more complete description of the experiment and the program may be found in Chowdhury, Collins, et al. (2015)

The TUP program in South Sudan is a pilot program. As with other programs of its kind, it consists of four phases: targeting and selection, training and enterprise selection, asset transfers, and monitoring. Each of these phases is modeled after the process in the original program in Bangladesh, but has been modified in notable ways based on local conditions.

Targeting and Selection

The TUP program in Bangladesh targeted women based primarily on a participatory appraisal activity in which community members used subjective means to assign households to different wealth quantiles. By contrast, the TUP program in South Sudan relies primarily on a set of inclusion and exclusion criteria based on wealth correlates taken from a community-wide survey, de-emphasizing relative measures of poverty in favor of absolute criteria.

Targeting guidelines include characteristics correlated with poverty as both exclusion and inclusion criteria. Surveyed households are excluded on the basis of having a salaried worker in the household or participation in another NGO program. Participation is also limited to women with access to cultivable land, since this is necessary for some of the TUP enterprises. Of these women, BRAC identified as eligible 650 who fit at least three of the following criteria: (i) the household head works as a day laborer; (ii) the household has two or more children; (iii) at least one child is working; (iv) the household has fewer than three rooms; and (v) the household includes an adult female who has not completed secondary school.

Eligibility was established in a census conducted in April of 2013. A baseline survey was then conducted among eligible women in June and July, which provided stratification data for the random selection of 250 women into the TUP program, with 375 remaining as controls.

Training and Enterprise Selection

Of the eligible households, 250 were randomly selected to participate in the TUP program. After a general orientation to familiarize them with the program overall, each client was asked their preference over a menu of possible business types, which included selling dry fish at market, raising goats, raising ducks, and growing maize. BRAC set the number of participants in each group beforehand, ensuring that many but not all participants received their preferred asset type. Next, clients were enrolled in business skill training. Some of this training is program-wide, such as basic and financial literacy, though most of it is specific to the type of asset provided. Training occurred over four days at BRAC's own office or demonstration farm.

Asset Transfer and Monitoring

The standard program then provided clients with productive assets, with an effort to keep the market value of transferred goods constant across enterprises. In late 2013, each client in each enterprise group received assets valued at roughly $240.

After transfers were made, BRAC also provided weekly food transfers (bags of maize or maize flour) during group meetings. This was intended to ease clients' household budgets, compensate them for their time at trainings, and encourage them not to sell productive assets before their businesses got off the ground. These food transfers continued until about a month before the follow-up survey, and were valued at roughly \$110 per client, raising the value of physical transfers to \$350. BRAC estimates a marginal cost for an additional client equal to the value of transfers plus 10–20% of this in delivery and administrative costs. Initial intensive training sessions later gave way to monitoring and mentorship from local staff, as well as small support groups consisting of 8–12 clients, such as those found in BRAC's microfinance programs. These group meetings were ongoing when the final round of data was collected.

Data and Selection

Our data comes from three principal sources. First is a census of adult women proximate to BRAC's regional office in Yei, which was conducted in April of 2013. From this census a subset of 745 'eligible' women was identified, who were then selected to be surveyed in a second 'baseline' survey conducted in June and July of the same year. This baseline identified 649 of the eligible women, who were stratified by baseline asset holdings, participation in small trade and agriculture, and number of income earners, with 250 households being randomly selected. A third follow-up survey was subsequently conducted in July of 2014.

The first round of data collection consisted of a census of women in households within a six kilometer radius of the regional BRAC office. These women typically live on small plots of land with several small, mud, one-room buildings with thatched roofs. Eighty percent of surveyed women are between the ages of 20 and 40, with between one and three children.

The census survey was designed to establish program eligibility. BRAC's approach of selecting on a range of 'correlates' of poverty is designed to be less costly than the more intensive community-based ranking exercise used in the Bangladesh program, raising the question of targeting effectiveness. Do the eligibility requirements successfully separate out an especially poor group of women, and does it avoid excluding women who should be eligible? Of the 1,279 surveyed households, 58% met all of the eligibility requirements. A straightforward comparison of the sample averages between the selected and non-selected groups indicates that selected households are 17% less likely to have paid work, have fewer durable assets and less livestock, and are more likely to be eating sorghum, which is typically regarded as low-quality food. Most selected women work either as a housewife or in small-scale agriculture. Eighty percent lived in households with some agricultural output, 35% had some poultry or livestock, and roughly 36% were involved in small trade or retail. Average reported daily consumption expenditures amounted to roughly \$1.50 USD per person.

Summary statistics for surveyed eligible women are presented in Table 1.1, by treatment group. The table provides means of various outcome variables at baseline. The column "N" indicates the number of non-zero values across the entire sample; the column "Diff." gives the difference in means across these two groups, while p is related to a test of the hypotheses that "Diff." is equal to zero.

Table 1.1: Means of some analysis variables at baseline. Asterisks in the column labeled "Diff." are an indication of a significant difference between the means reported in the "CTL" and "TUP" columns.

	N	CTL	TUP	Diff.	p
Consumption					
Meat	378	4.19	3.64	-0.552	0.11
Fuel	456	0.74	0.72	-0.017	0.83
Clothesfootwear	595	0.68	0.64	-0.036	0.62
Soap	536	0.47	0.47	0.0	0.99
Fish	474	2.46	2.35	-0.107	0.61
Charities	134	0.03	0.02	-0.006	0.46
Cereals	605	9.27	8.24	-1.03	0.19
Transport	193	0.18	0.14	-0.033	0.30
Cosmetics	468	0.64	0.71	0.065	0.49
Sugar	604	1.66	1.64	-0.02	0.91
Egg	276	1.11	1.00	-0.103	0.47
Oil	613	1.32	1.23	-0.087	0.59
Ceremonies	152	0.14	0.14	-0.002	0.97
Beans	192	0.77	0.93	0.163	0.32
Fruit	272	0.69	0.60	-0.09	0.29
Textiles	376	0.17	0.15	-0.021	0.35
Utensils	442	0.25	0.24	-0.011	0.70
Dowry	126	1.28	1.23	-0.049	0.89
Furniture	368	0.21	0.18	-0.028	0.39
Salt	617	0.45	0.42	-0.028	0.39
Vegetables	471	1.49	1.38	-0.11	0.41
Asset Values					
Smallanimals	123	198.90	150.53	-48.368	0.36
Tv	42	36.28	45.94	9.659	0.54
Bicycle	171	105.58	96.52	-9.06	0.65
Shop	44	85.46	79.41	-6.043	0.89
Radio	260	53.39	52.48	-0.908	0.94
Motorcycle	93	450.07	534.69	84.621	0.48
Mosquito Nets...	423	19.24	19.83	0.592	0.77
...Some treated	181	8.18	9.04	0.854	0.56
Poultry	161	39.68	39.04	-0.642	0.94
Sewing	28	8.56	4.96	-3.597	0.42
Shed	9	1.85	0.02	-1.832**	0.03

Continued on next page

Table 1.1: Means of some analysis variables at baseline. Asterisks in the column labeled "Diff." are an indication of a significant difference between the means reported in the "CTL" and "TUP" columns.

	N	CTL	TUP	Diff.	p
Bed	521	251.30	249.26	-2.039	0.94
Chairtables	531	207.89	177.42	-30.476	0.31
Carts	17	2.31	3.48	1.173	0.45
Fan	16	3.56	1.84	-1.712	0.28
Homestead	274	4432.11	4738.73	306.621	0.77
Cows	35	222.79	112.70	-110.085	0.19
Mobile	414	96.25	110.16	13.912	0.14
Other Variables					
Daily Food	643	25.11	22.97	-2.136	0.15
Daily Exp	646	29.82	27.74	-2.079	0.22
No. Houses	543	2.87	2.86	-0.006	0.97
In Business	265	0.40	0.44	0.033	0.42
Cereals	605	9.27	8.24	-1.03	0.19
Asset Prod.	475	854.03	624.88	-229.151	0.18
# Child	594	3.30	3.38	0.085	0.61
Land Access (fedan)	542	2.50	2.05	-0.443*	0.07
Asset Tot.	603	1787.27	1712.26	-75.011	0.73
Cash Savings	431	216.07	265.42	49.352	0.42
HH size	648	7.32	7.06	-0.267	0.18
Cosmetics	468	0.64	0.71	0.065	0.49

Though the kind of information presented in Table 1.1 is more useful for thinking about magnitudes than it is for 'balance' between the two randomly assigned groups, it's nevertheless true that mean values for these groups are generally similar. Only one of the differences we compute is significant by the standard of a sequence of t-tests and 95% level of confidence, and this difference is instructive. It comes in the calculation of the average value of sheds, where the control group happens to have a total of 8 sheds, while treatment group has only four; further, though all of the households in the control group happen to report a that their sheds have a positive value, only one of the four shed-owning households in the treatment group does so. The probability of some kind of imbalance along these lines happening for *some* variable is quite high, and of course this is no kind of evidence against the quality of the random number generator used to manage the assignment. Nevertheless, the initial difference should be kept in mind, if only because (as we'll see in the results below) "Sheds" are one of the outcomes which seem to be affected by the TUP program.

Chowdhury and Morel (2014) employ a principal component index developed by the Consultative Group to Assist the Poor (CGAP) to evaluate the effectiveness of targeting in this experiment. They find that that roughly half of the selected individuals are in the bottom quartile, and nearly all are poorer than average for their community. Exclusion criteria based on NGO participation and lack of land ownership exclude a significant number of relatively poor women, suggesting that this targeting method has sacrificed some targeting effectiveness for the sake of program structure.

After the original census, two surveys (a "baseline" and "follow-up") were conducted in the summers of 2013 and 2014, respectively. These surveys contained modules on enterprise and income-generating activity, household composition, food security, and consumption of a range of food and non-food goods.

Among the 745 households identified as eligible in the census, enumerators were able to locate and interview 649 in the baseline survey in July 2013. It was using this baseline that households were stratified by potentially important characteristics and randomly selected for enrollment. Asset transfers and training began in December of 2013. In total, 554 of these were located and interviewed in the follow-up survey in July 2014.

Since BRAC had kept in much closer contact with the TUP participants in the intervening months, attrition is a source of concern.

A Modest Model

Bandiera et al. (2012) offer a simple static model of the behavior of an individual. The model itself is a version of an agricultural household model, of the sort discussed in Singh, Squire, and Strauss (1986), but with a focus on occupational choice, which Bandiera et al. identify as a critical feature in their study in Bangladesh.

Here we adopt the model of Bandiera et al. (2012) more or less wholesale, but extend it to allow for both time and uncertainty. The spirit of this extension is very similar to the "exogenously incomplete" model devised by Karaivanov and Townsend (2014). However, we interpret it as a model of *household*, rather than individual behavior, since most of the data we have to test this model is observed at the household level. This turns it into a dynamic model involving both asset accumulation and occupational choice, and we show how this extension allows us to nicely tie together the production, consumption, and investment decisions made by the household.

Our notation is adapted from Bandiera et al. (2012), with modest changes to generalize and allow for time and uncertainty. Households are indexed by $j \in \mathcal{J} = \{1, 2, \ldots, J\}$. In each period the economy is in a state $s \in \mathcal{S} = \{1, \ldots, S\}$; these states evolve according to finite-state Markov process with the probability of transitioning from state s to state r given by π_{sr}. Time is discrete, and in each period t the household derives utility from consumption of an n-vector of consumption goods C and from leisure R. Utility within a period can also depend on household characteristics θ. Bandiera et al. (2012) interpret this θ as skills, but we'd interpret it more broadly to include, e.g., household size and composition. Then momentary utility is given by $U(C, R, \theta)$, with this utility function increasing, concave, and continuously differentiable. The household makes plans over an infinite horizon, with utility in the next period discounted by a factor $\beta \in (0,1)$.

In each period the household allocates its time between leisure R, employment (by others) L, and self-employment S. All must be non-negative. We assume that no labor is hired in by the household (modifying the model to allow this would be straight-forward, but not empirically useful in our setting, as none of the households in our sample is observed to hire in labor). Earnings from employment depends on an individual and state-specific function $W_s^j(L,\theta)$. Income from self-employment involves a production process which depends not only on time allocated to this occupation, but also on the productive assets and a household-specific shock; household j's characteristics evolve according to a household-specific Markov process, so that $\theta_{t+1} = H_s^j(\theta_t)$ if the state at $t+1$ is s.

Asset accumulation depends on initial assets K, the state-specific idiosyncratic price for new assets q_s^j, and stochastic, household-specific returns to holding those assets $Q_s^j(K)$ (e.g., think of livestock fertility and mortality). Borrowing is limited, but these limits may depend on the state and vary across households, so that $K_{t+1} \leq B_s^j(K_t)$ if the $t+1$ state is s. The returns function Q_s^j is assumed to be weakly concave; both it and the borrowing limit functions B_s^j are also assumed to be increasing and continuously differentiable.

In any state s, given assets K, characteristics θ, and time spent in self-employment S, household j produces $F_{js}(K,S,\theta)$ units of the numÃľraire good, where we assume the F_{js} are increasing, weakly concave, and continuously differentiable. The cost of purchasing the consumption bundle C is taken to be $P_s^j(C)$ for household j in state s. In each period the cost of consumption plus net investment must not exceed income from employment and own production, so that household j faces the budget constraint

$$(1.1)\quad P_s^j(C) + q_s^j(K' - K) \leq F_s^j(K,S,\theta) + W_s^j(L,\theta),$$

where K' is a vector of the total assets invested for the next period.

Putting this altogether, we regard household j as solving the dynamic program

$$(1.2)\quad V_s^j(K,\theta) = \max_{C,S,L,K'} U(C, 1-L-S, \theta) + \beta \sum_{r\in\mathcal{S}} \pi_{sr} V_r^j\left(Q_r^j(K'), H_r^j(\theta)\right)$$

subject to the budget constraint (1.1) (with which we associate the Karush-Kuhn-Tucker multipliers λ_s^j); to non-negativity constraints on consumption goods $i = 1,\ldots,n$, (with associated multipliers ν_i^j); non-negativity constraints for time allocation $S \geq 0$ and $L \geq 0$ (with multipliers η_S^j and η_L^j, respectively); and subject finally to the borrowing constraint $K' \leq B_s^j(K)$ (with multipliers μ_s^j).

Using lower case letters to indicate partial derivatives, the first order conditions then can be written

$$
(1.3)\quad
\begin{aligned}
C_i : u_i(C,R,\theta) - \nu_i^j &= p_{si}^j \lambda_s^j \qquad \text{for all } i = 1,\dots,n\\
L : u_R(C,R,\theta) - \eta_L^j &= w_s^j \lambda_s^j\\
S : u_R(C,R,\theta) - \eta_S^j &= f_S^j \lambda_s^j\\
K' : \beta \sum_r \pi_{sr} v_r^j q_r^j + \mu_s^j &= q_s^j \lambda_s^j.
\end{aligned}
$$

Here u_i denotes the marginal utility of consumption good i, and u_R is the marginal utility of leisure. Similarly f_s^j is the marginal product of S in production for household j, while $v_r^j = \frac{\partial V_r^j}{\partial K}(Q_r^j(K'), H_r^j(\theta))$ is household j's marginal valuation of an additional unit of realized capital in state r, and q_r^j is j's marginal return to investment in state r. In addition to these optimality conditions we have the envelope condition with respect to K,

$$
(1.4)\quad v_s^j(K,\theta) - \mu_s^j b_s^j(K) = \lambda_s^j \left(q_s^j + f_{sK}(K,S,\theta) \right).
$$

Now, the key variable which ties together all of these is the multiplier on the budget constraint, which measures the marginal benefit of having additional resources. Since this marginal value depends in turn on not only the state s but also the current values of (K,θ), we use (1.3) and (1.4) to implicitly write it as a function $\lambda_s^j(K,\theta)$. We have

$$
(1.5)\quad \lambda_s^j(K,\theta) = \frac{u_i - \nu_i^j}{p_{si}^j} = \frac{u_R - \eta_L^j}{w_s^j} = \frac{u_R - \eta_S^j}{f_{sS}^j} = \frac{\beta \sum_r \pi_{sr} v_r^j q_r^j + \mu_s^j b_s^j}{q_s^j} = \frac{v_s^j - \mu_s^j b_s^j}{q_s^j + f_{sK}^j}.
$$

In words, the household is allocating its resources to equate returns measured in terms of *utility* across different margins; none of these are returns in physical quantities that we can directly measure. "Utility return" would be an accurate way of describing these quantities: Taking each equality in (1.5) one at a time, λ_s^j is equal to household j's utility return of consuming an additional unit of good i (this holds for every $i = 1,\dots,n$, of course); is equal to the utility return to taking an hour off from employment; is equal to the utility return to taking an hour off from self-employment; and is equal to the utility return to an additional unit of investment, which finally is equal to the utility value of having additional assets in the current state s.

But while "utility return to an additional unit of investment" may be accurate, we think the English language already has a suitable word: the variables λ_s^j measure the *neediness* of household j. When λ^j is high relative to those of other households, so is $(u_i^j - \nu_i^j)/p_i^j$, and household j stands in greater need of food; similarly when $\lambda_s^j > \mathrm{E}\lambda^j$ the household is particularly in need of labor; of investment; of consumption; of leisure.

The neediness variables λ_s^j have other interpretations as well. If we were to consider the static consumer's problem being equal to the partial derivative of the household's indirect utility function with respect to total consumption expenditures in state s.

Notice that the different expressions for neediness in (1.5) involve three different kinds of objects. First, there are some prices which may be directly observable in the data (e.g., prices of consumption goods; individuals' wages; purchase prices of assets such as livestock). Second, there are shadow prices that will *not* be directly observable; these include the key λ_s^j as well as multipliers on the non-negativity constraints and the multiplier on the borrowing constraint. Third, there are unknown functions, including the marginal utility functions (u_i, u_R) and the marginal productivities of assets and labor in the self-employment technology (f_{sS}^j, f_{sK}^j).

Modeling our experiment

We want to think now about how our experiment can be thought of in terms of the model of the households we've developed—only by putting the experiment "into" the model can we think coherently about how a household might react to the experimental treatments we introduce. Or as Rubin (1974) might put it, we think of putting the experiment into the model as the construction of a logical argument establishing circumstances under which only some particular variables should be expected to have a causal effect on particular dependent variables.

Accordingly, consider partitioning the space $\mathcal{S} = \mathcal{C} \cup \mathcal{E}$. Then for any state $s \in \mathcal{E}$ we begin our experiment (we can always specify $\mathcal{S}$ and choose the transition probabilities π_{sr} to ensure that we only start the experiment once). Further, let $T_0(s)$ and $T_1(s)$ be subsets of the index set of households, so that for $\hat{s} \in \mathcal{E}$ if $j \in T_0(\hat{s})$ then household j is assigned to a 'control group' in our experiment, while if $j \in T_1(\hat{s})$ then household j is assigned to a 'treatment group' which receives assets, training, and so on. Assignment is random if, for any pair of households $(j_0, j_1) \in T_0 \cup T_1$ each had an equal probability of being assigned to T_1.

In partitioning $\mathcal{S}$ into states where the experiment is conducted and states where it is not, we think of $\mathcal{C}$ as the set of 'counterfactual' states. Thus, for an 'experiment' state $\hat{s} \in \mathcal{E}$ there exists another 'counterfactual' state $\tilde{s} \in \mathcal{C}$ such that for any household $j \in T_1(\hat{s})$, the 'treatment' consists of an $\hat{K}$, a $\hat{\theta}$, and a $\hat{C}$ such that

$$(1.6)\quad Q_{\hat{s}}^j(K') = Q_{\tilde{s}}^j(K') + \hat{K}; \quad H_{\hat{s}}^j(\theta) = H_{\tilde{s}}^j(\theta) + \hat{\theta}; \quad \text{and } P_{\hat{s}}^j(C) = P_{\tilde{s}}^j(C - \hat{C})$$

for all K', θ, and C. Note that we are *not* assuming that consumption or investment will be unchanged by the treatment; it would be surprising if they were not. The content of the assumption is that the technology producing returns to investment or the cost of a consumption bundle only be affected by the experiment in an additive way.

Further, we assume that for any household in the *control* group outcomes are the same in both the experimental state $\hat{s} \in \mathcal{E}$ and the counterfactual state $\tilde{s} \in \mathcal{C}$, or, for any $j \in T_0(\hat{s})$ that we have

$$(1.7)\quad Q_{\hat{s}}^j(K') = Q_{\tilde{s}}^j(K'); \qquad H_{\hat{s}}^j(\theta) = H_{\tilde{s}}^j(\theta); \qquad \text{and } P_{\hat{s}}^j(C) = P_{\tilde{s}}^j(C),$$

also for all K', θ, and C. Together, these two conditions just assert that our experiment only affects the treated, and give the effect of the treatment on treated households. Left unstated is a third assumption, that the treatment's effects on treated households are channeled solely through the transfers of $(\hat{K}, \hat{\theta}, \hat{C})$.

This notation may seem unnecessary, if our goal is simply to discuss what it means to have experimental treatments and random assignments. But now we ask—within the context of the model—what effects we'd expect from the experimental treatment. There turns out to be a very simple way to measure these. Equation (1.5) implies that changes in any aspect of the household's economic behavior (consumption, labor supply, production, credit constraints) will be reflected in the neediness λ^j_s, so one way of thinking about what we want to measure experimentally is the ratio $\lambda^j_{\hat{s}}/\lambda^j_{\tilde{s}}$ for $j \in T_1(\hat{s})$. This ratio would tell us the proportional difference in utility returns for a treated household due to the experiment.

Viewed through this lens, the expected "average treatment effect" on (the log of) neediness can be written as

$$\text{ATE} = \text{E}\left(\frac{1}{\#T_1(\hat{s})}\sum_{j\in T_1(\hat{s})}\log\lambda^j_{\hat{s}}\right) - \text{E}\left(\frac{1}{\#T_1(\hat{s})}\sum_{j\in T_1(\hat{s})}\log\lambda^j_{\tilde{s}}\right).$$

The problem, of course, is that we can't observe the λ^js in the counterfactual state $\tilde{s}$. But using the assumption (1.7) and the assumption of random assignment, it follows that

$$\text{E}\left(\frac{1}{\#T_1(\hat{s})}\sum_{j\in T_1(\hat{s})}\log\lambda^j_{\tilde{s}}\right) = \text{E}\left(\frac{1}{\#T_0(\hat{s})}\sum_{j\in T_0(\hat{s})}\log\lambda^j_{\hat{s}}\right),$$

so that we have the average treatment effect on the logarithm of neediness given by

$$\text{ATE} = \text{E}\left(\frac{1}{\#T_1(\hat{s})}\sum_{j\in T_1(\hat{s})}\log\lambda^j_{\hat{s}}\right) - \text{E}\left(\frac{1}{\#T_0(\hat{s})}\sum_{j\in T_0(\hat{s})}\log\lambda^j_{\hat{s}}\right). \tag{1.8}$$

This now only involves needing to observe outcomes in realized states.

Empirical Strategy

Notice that the utility returns in (1.5) involve three different kinds of objects. First, there are some prices which may be directly observable in the data (e.g., prices of consumption goods; individuals' wages; purchase prices of assets such as livestock). Second, there are shadow prices that will *not* be directly observable; these include the key λ^j_s as well as multipliers on the non-negativity constraints and the multiplier on the borrowing constraint. Third, there are unknown functions, including the marginal utility functions (u_i, u_R) and the marginal productivities of assets and labor in the self-employment technology (f^j_{sS}, f^j_{sK}).

These last unknown functions depend on variables which we may be able to observe. Consider in particular a good i of which household j consumes a positive quantity. This gives us the equality $u_i(C^j_s, R^j_s, \theta^j_s)/p^j_{si}(C^j_s) = \lambda^j_s$. This equation holds for all states and for every good $i = 1, \ldots, n$ with positive consumption, so it must hold in any realized state. To celebrate this fact we simplify notation, letting t indicate the $u_i(C^j_t, R^j_t, \theta^j_t)/p^j_{ti}(C^j_t) = \lambda^j_t$ to simplified notation we also introduce some additional assumptions: first, that utility from leisure is additively separable from utility from consumption, or that $u_{iR} = 0$. Second, we partition the index set of households into sets of households that reside within m distinct areas; i.e., we take $\mathcal{J} = \mathcal{J}_1 \cup \mathcal{J}_2 \cdots \cup \mathcal{J}_m$. Then we assume that within each of these m areas households all face the same prices for consumption goods, or that $p^j_{ti}(C^j) = p_{ti}$.

Now, with this we return to the equation defining the expected average treatment effect (1.8). Using the fact that for goods consumed in positive amounts we now have $\lambda^j_t = u_i(C^j_t)/p_{ti}$, we substitute into (1.8), obtaining

$$\log u_i(C^j_t, \theta^j_t) = \log p_{ti} + \sum_g \mathbb{1}(j \in T_g)\overline{\log \lambda_t}^{T_g} + \epsilon^j_{ti},$$

where $\overline{\log \lambda_t}^{T_g}$ is the average value of the $\log \lambda$s for treatment group T_g, ϵ^j_{ti} is a residual which, by (1.8) will be equal to $\lambda^j_t - \overline{\log \lambda_t}^{T_g}$ if household j is a member of treatment group g.

Estimating Marginal Utilities

If we observed prices p_{ti} and happened to know the values of $\log u_i(C^j_t, \theta^j_t)$ we could go ahead and straight-forwardly estimate the average treatment effect we're interested in. Of course we do not know the latter. However, we do observe expenditures on multiple kinds of food and other non-durable consumption. If we re-arrange the first equality in (1.5) and use our assumption that leisure is separable in utility then we can write the vector of marginal utilities of consumption as

$$u(C, \theta) = p\lambda.$$

Next, following the long line of work following Heckman and MaCurdy (1980) and MaCurdy (1983), we parameterize the log of marginal utilities, assuming $\log u(C, \theta) = \Gamma \log C + \zeta\theta$, where Γ is an $n \times n$ matrix of parameters having full rank, and where ζ is an $n \times l$ matrix.[2]

With this parameterization, we can write

$$\Gamma \log C + \zeta\theta = \log p + \log \lambda.$$

This is getting close to something we can estimate, but we have data on the value, not quantity, of food consumption. Let $X_i = p_i C_i e^{\epsilon_i}$, where ϵ_i is some measurement error, be the value of expenditures on consumption good i. Then rearranging, we have the system of equations

$$\log X = (I + \Gamma^{-1}) \log p - \Gamma^{-1}\zeta\theta + \Gamma^{-1} \log \lambda + \epsilon. \tag{1.9}$$

[2]The development and justification of this particular preference structure is discussed in Ligon (2015).

This system is what we might call a Frischian expenditure system (Browning, Deaton, and Irish, 1985). Ligon (2016b) provides methods for estimating this system; showing that with data on at least some expenditures and household characteristics one can obtain not only estimates of the parameters but also of the neediness measures $\log\lambda$ (up to a normalization).

Differences in the mean of the inferred neediness $\log\lambda$ between treatment and control group will be equal to the average treatment effect that most interests us, but we can also obtain estimates of this effect directly from (1.9). Consider the following standard ANCOVA specification of the sort championed by McKenzie (2012). Key features of the standard specification include a set of fixed effects for time and place; linear covariates as controls; baseline values of the outcomes as additional controls; and finally a collection of average treatment effects, which are ordinarily the object of interest. We adopt just such a specification, letting X_{ti}^{jga} be expenditures on good i in period t for a household j in area a and in treatment group g. Then we can write

$$\log X_{ti}^{jga} = \alpha_{ti}^{a} + \tau_i^g + \delta_i(\theta_t^j - \bar{\theta}_t^g) + \gamma_i \log X_{t-1,i}^{jga} + u_{ti}^j. \tag{1.10}$$

Now, in the standard interpretation of this regression $\tau_i^1 - \tau_i^0$ will be the average treatment effect on expenditures on good i, while the terms involving the θ and the lagged outcomes improve power by accounting for covariance between household characteristics and expenditures (and perhaps accounting for unbalanced outcomes in the baseline). Because the latent variables α_{ti}^a capture differences in means across areas as well as goods and periods, it is the variation that is within an area that is being exploited here to estimate the τ_i^g.

This ANCOVA specification has an intimate relationship with the Frischian expenditure system (1.9) which allows us to give a structural interpretation of the reduced-form ANCOVA. In particular, the good-area-time effects α_{ti}^a estimate the effects of changes in prices on expenditures, the vector $(I + \Gamma^{-1})\log p_t$ in (1.9). The terms involving the idiosyncratic covariate characteristics $\delta_i\theta_t^j$ match up with the effects of characteristics on expenditure demand $\Gamma^{-1}\zeta\theta_t$, while the average treatment effect estimates $\tau_i^g = \beta_i(\overline{\log\lambda_t}^{T_1} + \zeta_i\bar{\theta}_t^g)$, where the β_i are equal to the row sums of the matrix Γ^{-1}.

So, the average treatment effect in these ANCOVA regressions with log consumption expenditures as outcomes can be interpreted as the product of a demand elasticity and neediness. Further, these can be decomposed, giving us both parameters useful for understanding demand systems and Engel curves and measures of neediness useful for measuring welfare. Even better, these neediness parameters are key to understanding the connections between consumption, investment, production, and occupational choice, and allow us to measure the extent to which an intervention operates via its effects on wealth versus effects it may have on production or occupational choice.

What assumptions have we had to make in order to give this 'structural' interpretation to our average treatment effects? There are only really four 'structural' assumptions we need to make. All pertain to the household's utility function, and seem fairly unobjectionable, or at least conventional in applied empirical work. The first two are that the household's utility function is intertemporally separable and von Neumann-Morgenstern; these allow us to think of the household as solving a 'two-stage' intertemporal budgeting problem (Gorman, 1959). The third is that the utility function is separable in consumption and leisure; the last that Frischian consumption expenditure elasticities are constant. This is much less restrictive than what is usually assumed in parametric Engel curve estimation.

Results

We offer results in three parts. First we discuss the average treatment effect on consumption expenditures, and use estimates of this effect across different consumption goods to estimate the average treatment effect on neediness, as well as the distribution of neediness in both treatment and control groups. Second, we consider outcomes related to both the number and value of assets held by the household. The estimates of household neediness previously obtained can be used to control for the effects of treatment on wealth. The link between these holdings and the model is considerably looser than in the case of consumption, but certainly both the average number and value of assets we observe is positively affected by TUP. The distribution of resources *across* different assets is less easy to predict, but we see large average treatment effects on the value of livestock owned, consistent with the focus of TUP on increasing livestock ownership for treated households which choose this. We finally examine self-employment and occupation. There are quite large effects on participation in self-employment, broadly consistent with what one would expect from a purely administrative analysis (BRAC gave animals to so many treated households, of which a certain known number already had significant livestock holdings). Finally, we turn to a broader and more detailed notion of occupation: here we see members of treated households leaving housework and casual agricultural employment. Some of these people seem to enter non-agricultural day labor, but it's less clear what they're doing instead. However, one possibility consistent with both the evidence and the model is that people in the average treated household move out of low-skill market employment, instead increasing labor in more skilled market employment, and possibly increasing both household leisure and participation in home production.

Consumption Expenditures and Neediness

Our principal results may be found in Table 1.2. As suggested above, these are 'ANCOVA' estimates of the effects of being in either of the two groups "CTL" (Control) and "TUP" (targeted ultra-poor), the latter of which received assets, training, and food subsidies. Other household characteristics included as controls are the number of people in the household as well as the number of children. Baseline values of expenditure were included as an additional control, with a complete set of village/area fixed effects (constrained to sum to zero). Where recorded values of consumption expenditure are equal to zero, these are regarded as missing and dropped from the analysis. There are two motivations for this treatment of zeros: first, at an entirely practical level, our dependent variable is the logarithm of expenditures, which is undefined at zero. But second, if a household is at a corner when it chooses a particular consumption item, then the first order condition in (1.3) for that consumption good won't be correct (we'd be missing a multiplier related to non-negativity). By simply dropping observations for goods where consumption is zero we are effectively dropping observations where expenditures do not correctly reveal household neediness. In any event, treating zero consumptions as missing results in our 'panel' of goods by households being unbalanced, so we estimate the ANCOVA equations as a single system.

We see in the first instance that the average treatment effect for TUP participants on the value of these consumption goods are almost uniformly positive, and significantly positive (three stars indicates a 99% level of confidence, two stars 95%, and one star 90%) for 8 of 14 different goods. The exceptions are informative. The estimated sign for the difference in the value of salt consumption is negative, but very small and insignificant, consistent with the view that the income elasticity of salt is very small for this population. The other negative difference is for transportation expenditures. We've included transportation in this table with the idea that transportation services enter the utility function. But another view is that transportation is an expense associated with employment or production. One of the principal findings of Bandiera et al. (2012) is that TUP participants in Bangladesh switched from wage employment to self-employment, which one presumes may have reduced the demand for transportation, and it's very possible that something similar is happening here.

The differences in average treatment effects are also highly jointly significant: a test of the hypothesis that these are all zero yields a χ^2_{14} statistic of 75.43, with an associated p-value less than 10^{-9}.

Now recall that according to our model each treatment effect is equal to the product of an elasticity parameter β_i and average log neediness for the group. By redefining the 'treatment groups' so that there are 554 of them, each group consisting of exactly one household, we can obtain estimates of *individual* effects on the value of goods consumed, or $\beta_i \log \lambda^j_t$. We adopt the normalization that $\text{var}(\lambda^j_t) = 1$, and scale the elasticities β_i so that their sum weighted by expenditure shares is equal to one. Scaled in this way these Frisch elasticities would be equal to Marshallian income elasticities provided each household had a coefficient of relative risk aversion of one. These Frisch elasticities are reported in the final column of Table 1.2. As the differences in estimated average treatment effects would suggest, all but salt appear to be normal goods. Because the scale is only identified by an arbitrary normalization, we can't say based on this evidence what goods are necessities or a luxuries. But we can say that fuel, transport, soap, and cosmetics (all the non-food items) appear to be the four most income elastic goods, followed by vegetables, sugar, cooking oil, and cereals. And whatever the scale, the least income elastic good seems to be salt, with an elasticity orders of magnitude smaller than that of the most income elastic goods.

We now turn our attention to the relationship between (log) neediness and treatment. The final row of Table 1.2 reports the mean values of neediness for both CTL and TUP groups. As with the individual goods, there's a highly significant difference between these means. Because the standard deviation of the pooled $\log \lambda^j_t$ is equal to one (because of our normalization), we can interpret the difference between these means as evidence that neediness for the treatment group fell by a highly statistically significant 0.2 standard deviations relative to the control.

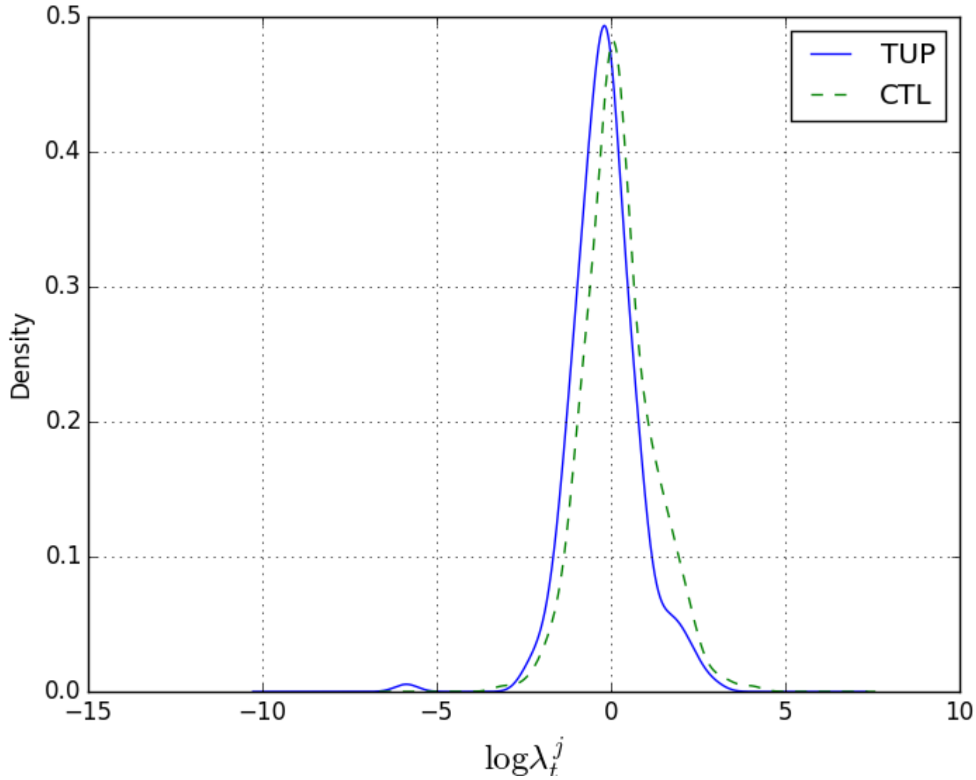

Figure 1.1: Distribution of Neediness, by Treatment.

Of course, knowing just that the mean neediness is less in the TUP group tells us little about how changes in welfare are distributed across households. Giving assets to would-be entrepreneurs might have very disparate effects on welfare, as many standard models of entrepreneurship predict (Banerjee and Newman, 1993; Paulson, Townsend, and Karaivanov, 2006; Karaivanov and Townsend, 2014) and as a number of recent experiments tend to confirm (De Mel, McKenzie, and Woodruff, 2008; McKenzie and Woodruff, 2008; Fafchamps et al., 2011). Perhaps some fortunate or skilled few benefit hugely while others experience little benefit.

To understand the distribution of benefits in our setting, consider Figure 1.1, which presents kernel density estimates of the distribution of $\log \lambda_t^j$ across households conditional on whether they are members of either the treatment or control group. Two things are visually evident from the figure. The first is that average neediness for the TUP group is smaller than it is for the control group. Related, the second is that the distribution of welfare gains for the TUP group may first-order stochastically dominate the distribution for the control group: it's not just that mean neediness falls, it's that mean neediness appears to fall for *everyone*, save for the least needy (consistent with the idea that the utility function U is concave).

Other Testable Predictions of the Model

The model presented here is written so as to be quite general in some dimensions, and we lack the data to construct structural estimates of the full model. However, with only fairly modest maintained assumptions we can estimate parts of this model, and test others. For example, the previous section has outlined methods for estimating a parametric utility function and corresponding demands for non-durable consumption, which we exploit below. We have also described an approach to measuring the effects of the TUP program on average household welfare.

With what we've been able to estimate, we'd like to be able to use the model to ask two counterfactual questions about the TUP program. The first: what size of cash transfer would yield the same welfare benefits as what we observe from the experiment? We'll call this the "welfare-equivalent cash transfer." The second: in what ways is the behavior induced by the TUP program different from what we'd expect from the welfare-equivalent cash transfer?

Results

We offer results in three parts. First we discuss the average treatment effect on consumption expenditures, and use estimates of this effect across different consumption goods to estimate the average treatment effect on neediness, as well as the distribution of neediness in both treatment and control groups. Second, we consider outcomes related to both the number and value of assets held by the household. The estimates of household neediness previously obtained can be used to control for the effects of treatment on wealth. The link between these holdings and the model is considerably looser than in the case of consumption, but certainly both the average number and value of assets we observe is positively affected by TUP. The distribution of resources *across* different assets is less easy to predict, but we see large average treatment effects on the value of livestock owned, consistent with the focus of TUP on increasing livestock ownership for treated households which choose this. We finally examine self-employment and occupation. There are quite large effects on participation in self-employment, broadly consistent with what one would expect from a purely administrative analysis (BRAC gave animals to so many treated households, of which a certain known number already had significant livestock holdings). Finally, we turn to a broader and more detailed notion of occupation: here we see members of treated households leaving housework and casual agricultural employment. Some of these people seem to enter non-agricultural day labor, but it's less clear what they're doing instead. However, one possibility consistent with both the evidence and the model is that people in the average treated household move out of low-skill market employment, instead increasing labor in more skilled market employment, and possibly increasing both household leisure and participation in home production.

While Banerjee, Duflo, et al. (2015) do not estimate neediness according to our procedure, they both find comparable average treatment effects on the sum of consumption for the expenditure categories they measure. The consistency of average treatment effects across the distribution also mirrors the distributional results in Banerjee, Duflo, et al. (2015). Bandiera et al. (2017) find increases in consumption among the ultra-poor, but do not report short-term estimates like the ones we present here.

Assets

We have seen that the TUP treatment has a positive and significant effect on consumption expenditures and leads to a significant and sizable reduction in neediness. From (1.3), we might expect this reduction in neediness to also show up in investment and assets. Of course, since the TUP program revolves around actually giving assets to treated households, it may appear obvious that assets should increase. But in fact this is not at all a foregone conclusion. From (1.3) we have an indication that a decrease in neediness (such as the one we measured above) may decrease the marginal value of assets (consistent with an increase in the holdings of those assets). But the assets may be valued simply because they can be sold to finance increased consumption or leisure—a pure wealth effect, which would be reflected in a reduction in neediness λ^j_s. This use certainly improves welfare, and may help extend the benefits of the TUP program to future periods, but this is a role that would be played equally well by a (simpler) financial transfer. For the asset transfers to play an important role in *production*, we should look for the effects they may have on the production function, where a transfer of particular assets may either directly enter the production function, or may help to relax a borrowing constraint (perhaps by serving as a security), allowing the household to finance the purchase of other inputs to production should it wish.

Here we explore the effect of the TUP program on physical assets by estimating the average treatment effect on both the number (Table 1.3) and value (Table 1.4) of different sorts of assets.

Both sets of regressions are estimated just as the average treatment effects for consumption was, with the sole difference that reports of "zero" assets (whether count or value) were not treated as missing data. In particular, we include a complete set of village fixed effects, constrained to sum to zero; baseline (2013) values were included as controls, along with the number of people and number of children in the household.

Results for the *number* of assets are reported in Table 1.3. Consider first the column labeled "Diff. (no $\log\lambda$)", which excludes estimated neediness from the list of controls. In contrast to the case of consumption goods, few of these individual items are significant: at a 90% level of confidence the TUP program results in significant increases only in the number of chairs and tables, mobile telephones, poultry, and the number of sheds. The hypothesis that none of these differences is significant is rejected; it yields a χ^2_{15} statistic of 142.7 with an associated p-value less than 10^{-9}. The finding that treatment results in more poultry[3] and sheds is unsurprising, as some of the enterprises selected in the TUP program explicitly involved duck acquisition and shed construction. The finding that furniture or mobile phone

[3]Some care should be taken in interpreting the magnitudes of the effects on poultry, as the elicitation of both the number and value of poultry was handled slightly differently in the 2013 baseline and the 2014 follow-up survey.

purchases are significant is less expected, but it is perfectly possible that the operation of a small businesses might benefit from having a mobile or a table, of course. Sewing machines have obvious productive uses, but none directly related to the enterprises the TUP program was designed to encourage. Other surprises are that some other outcomes do *not* have significant treatment effects. In particular there is no significant effect of the TUP on the number of small animals owned—this is surprising as 35 of the treated women chose to rear goats (these out of a total of 246 treated households that received some kind of asset).

A deeper insight into the mechanisms behind asset acquisition can be gained by re-estimating the ANCOVA regression behind Table 1.3, but this time controlling for neediness. The coefficients on $\log\lambda$ are reported in the final column of the table; this allows us to see that less needy households are more likely to have more assets, as without exception the estimated coefficient on neediness is negative. Of these, 11 of 15 are significant at a 90% level of confidence. But perhaps more importantly, we can now re-interpret the effects of the TUP program on the number of assets held *controlling* for a measure of wealth; the relevant estimates are reported in the column labeled "Diff. (with $\log\lambda$)". When we control for neediness, we see that the increase in chairs and tables or mobile phones appears to be due only to the wealth effect of the TUP program ($\log\lambda$ is significant in these regressions, but the estimated average treatment effect is no longer significantly different from zero). The coefficients on poultry and small animals remain significant, as we'd expect. The coefficients on mosquito nets we do not understand: they suggest that less needy households are more likely to own mosquito nets, but that the TUP treatment resulted in fewer mosquito nets.

Referring to Table 1.4 may help to resolve the puzzle of the missing small animals; treatment is associated with a significant increase in the *value* of both poultry and small animals. No other differences are individually significant at the 95% confidence level, but we easily reject the hypothesis that *none* of these differences in value is significant. To summarize: the average treatment effect on the value of different assets is significant for poultry and small animals. Perhaps it would be surprising if this was *not* the case, since the treatment involves giving ducks and goats to more than half of the treated households, but the fact that those ducks and goats haven't been eaten or sold six months after the asset transfers provides suggests that the asset transfers affect production as intended, and serve as more than just a store of wealth. Both Banerjee, Duflo, et al. (2015) and Bandiera et al. (2017) find significant effects on total asset holdings in the short term, as well, with a similar emphasis on livestock and other productive assets.

Employment and Occupation

The model we've described above requires an explicit decision from the household about the allocation of time between leisure, production, and employment. Our results related to consumption expenditures and neediness tell us that TUP households' neediness falls, and our model suggests that we should expect this decrease in neediness to be related not only to consumption but also time allocation. Equation (1.3) describes the relation, with

$$\log \lambda_s = \log(u_R - \eta_L) - \log w_s = \log(u_R - \eta_S) - \log f_{sS}.$$

Suppose that labor and consumption are separable, as assumed in our calculation of neediness. There are four cases to consider.

First, it might be the case that the household supplies no labor at all, so that $\eta_L^j \eta_S^j > 0$. In this case a small decrease in neediness caused by an increase in K will not affect the marginal utility of leisure, u_R, and cannot affect the 'wage' w_s^j, so that the entire decrease in neediness will (from the point of view of employment) be reflected in the shadow cost of not being able to take *more* leisure. Taking the appropriate derivatives in this case yields $d\eta_L/d\lambda = w$; note that only 'shadow' quantities are changed in this case. For the second equality, the household in this case is still assumed to be at a corner in leisure, so u_R will remain unchanged, and changes in λ will be reflected in changes in η_S, but also in increases in the marginal product of labor f_{sS}. Under reasonable specifications of the production function F this makes perfect sense: the provision of greater capital inputs to home production are apt to yield exactly this sort of response.

Second, consider the case where the household is at a corner in L, because the wage w faced by the household is less than the marginal product of labor from own production f_S given assets K. In this case for the second equality we may expect to see increases in the marginal product of labor f_S. The effect on leisure is indeterminate, depending on the curvature of F.

Third, consider the case where the household is at a corner in S, because the wage w faced by the household exceeds the marginal product of labor from own production f_S. In this case there may be an increase in own production, but if wages are fixed there will certainly be an increase in leisure.

Finally we come to the fourth case, in which the household supplies labor both in the market and in own-production. As in the third case, if wages are fixed leisure must increase, resulting in a decrease in u_R. But the reduced labor previously supplied to the market can be divided between leisure and additional self-employment, though whether and how much time in self-employment increases will depend on the curvature of the production function.

Thus, the model at this level of generality leaves us with only some weak predictions about outcomes. The main prediction we have is that a small decrease in neediness will (weakly) increase leisure, *unless* the household is initially only self-employed, in which case the change in leisure is ambiguous. Note also that even this weak prediction hinges on the marginal product of labor in the market being taken as given—this may not be the case for, e.g., piecework labor, where decreasing marginal products may be the rule.

So, compared to the case of consumption the model gives us much less guidance regarding what to expect in terms of employment, and we can turn to the results without being surprised.

Table 1.5 provides the unsurprising results regarding the effects of the program (which we've established results in a decrease in neediness) on self-employment. The respondent is said to be "in business" if they claim to have been involved in any non-farm self-employment in the past year ("non-farm" here explicitly means agriculture, livestock, and poultry). They are "cultivating" if they report actively cultivating any land, whether owned, rented, or common. Being occupied in rearing livestock is taken here to mean that the respondent reports owning total livestock valued at more than 50 South Sudanese Pounds (roughly 12 USD at the time of the survey).

So how do the less-needy households of the TUP program change their self-employment? Noting that these questions elicit *participation* in different forms of self-employment, we see significant increases in participation both in business and livestock. In terms of the model, these changes have more to do with the multiplier η_S than they do with hours spent in self-employment, but the average treatment effects seem robust; we can think of this as being leading to roughly 20 per cent (5% plus 17% less some who leave cultivation) more of TUP households moving into self-employment than out of as a consequence of their participation in the program. The TUP program had a total of 216 participants, of which 116 chose to receive livestock, so our estimated participation rate here is slightly less than half of what the administrative data would tell us *provided* that none of these households had any livestock previously, but from our baseline data we can see that in fact 67 TUP households did previously have livestock, suggesting a reasonable match between the changes in participation expected from administrative data and the estimated average treatment effect. This may seem of little note, but evidence that people who were given ducks have chosen to raise them instead of eating or selling them is important to have.

In a separate part of the survey we elicit occupational information for all members of the household, where "occupation" can include not only various forms of self-employment, but also many other possible uses of one's time. Though it is possible to report several occupations for each person, in only three instances was more than one occupation reported for a single person. Of the 4304 individuals in the sample households in 2014, occupations are reported for 3886. The rather old or the very young are heavily over-represented among those with no reported occupation.

We add up the total number of people in each of 35 occupations in each household. Table 1.6 reports, in its first column, the total number of people in each occupation (for occupations with more than 30 people). This evidence on occupation paints either a disturbing picture of the economic environment in South Sudan or an encouraging testament to BRAC's ability to identify and target the 'ultra-poor': of people in the top twelve occupations listed, less than 22 per cent are engaged in what we might think of as remunerated productive work (students and housewives work, but aren't remunerated; beggars are remunerated but aren't productive). Of this 22 per cent, 61 percent are engaged in cultivating household land, either for home consumption or for sale. An additional 12 per cent are reported to work in their own small business, while the balance (26 per cent) sell their labor to others.

Because these occupations are reported for all household members they include children, and in the second column of Table 1.6 we show the number of children (under the age of 17) in different occupations. It's no surprise that most (three quarters) of students are children, and it's reassuring to see that 70 percent of those who are unemployed and not seeking work are young. Less reassuring is the fact that two thirds of the beggars in this sample are children. Twenty-two per cent of the unemployed looking for work are also under the age of 17.

Using data on reported occupations for all individuals, we estimate an ANCOVA regression, on the model we've used in earlier tables; the only important difference is that though we control for household size and numbers of children (and for neediness for figures reported in the last two columns), we do *not* control for baseline occupational counts (data on occupation in the baseline was elicited in a way not directly comparable to data in the later survey). Thus, this is an entirely "Post-treatment" comparison, and we are unable to control for any pre-treatment differences in occupation across groups.

The joint hypothesis that all differences are zero is soundly rejected, with a p-value less than 0.02. The principal finding from Table 1.6 is that the households in the TUP group are significantly less likely to be engaged in unpaid housework or as agricultural laborers (on someone else's land). This last finding seems to echo a result of Bandiera et al. (2017), who find that a TUP treatment seems to play an important role in causing women to shift from wage- to self-employment. Banerjee, Duflo, et al. (2015) similarly finds a significant increase in time spent working, without an increase in hours of wage labor. Those experiments find the increase in labor hours driven in part by increases in agriculture and livestock activities related to the program itself, which we will see is not the case in our setting. However, the other significant effect is an *increase* in employment as a non-agricultural laborers. We have no compelling explanation for this second finding, though we note that the total numbers of such workers is quite small. But though only the only individual occupations that demonstrate a significant treatment effect are casual agricultural and non-agricultural labor, *overall* there seems to be a quite significant effect of TUP on occupation (the hypothesis that all of the coefficients in either of the "Diff." columns of Table 1.6 are equal to zero is rejected with a very high degree of confidence).

One might have thought that we would see people reporting occupations related to the TUP program: increases in household land cultivation or vegetable farming; increases in poultry or livestock husbandry; or increases in the operation of a small business; all of these were explicit offered as possible occupations. But we do not see significant effects for any of these. Further, of the 83 women who were given ducks, and of the 35 women given goats, exactly one of each reports their occupation as "poultry husbandry" or "livestock husbandry". Possibly the participants in these programs regard the corresponding activity as something less than or different from an "occupation".

One clear prediction of our model was that treated households which were not initially active in both self- and market-employment would respond, in part, by increasing leisure. Table 1.6 offers some tantalizing but not conclusive evidence on this. Looking just at the point estimates in the first "Diff." column, it appears that the average treated household moves out of casual agricultural employment (this is significant), but also moves out of *all other* listed occupations save for casual *non*-agricultural labor (significant), skilled labor (not significant) and unemployment (both seeking employment and idle). Thus, a consistent account one can give to explain Table 1.6 is that the average treated households moves away from casual agricultural labor and perhaps some other unskilled occupations. The time freed is allocated to more skilled market employment, and perhaps to increased leisure.

To summarize: the introduction of the TUP program does induce a significant occupational response, with particular identifiable responses including a decrease in casual agricultural labor, an increase in casual non-agricultural labor, and an increase in unemployment. We do not, however, see direct evidence of *particular* TUP enterprises changing occupation. One tempting general conclusion is that role played by TUP in determining occupation may depend more on its general loosening of budget and borrowing constraints than it does on changing the relative returns of wage- and self-employment. However, given the lack of reliable baseline data on occupation this is largely speculative.

Conclusion

We finish by making a general observation. We have arranged an randomized control trial of a complicated intervention in a low income country. This is the kind of endeavor where theory and 'structural' approaches to estimation may not seem to have very much to contribute: the number of outcomes affected by the complicated intervention may be large and uncertain, and the demands made by a 'structural' model to explain all these outcomes may seem absurd. But combined with randomization, sometimes a little structure can go a long way. With only quite modest assumptions on household preferences we've developed a rather general model of household behavior. This model is not very structural in the sense that we've adopted very few assumptions about precise functional forms or laws of motion.

Our main approach to estimation is both conventional and modest: we identify a number of "outcomes", and use ANCOVA regression to estimate average treatment effects. The main methodological contribution of the paper is the recognition that when the outcomes include the logarithms of different consumption expenditures, then the average treatment effects can be interpreted within our modest model as the product of a price elasticity and the average value of the log of the multiplier on the household's budget constraint. With this recognition one sees that these average treatment effects can be easily decomposed, recovering estimates of those elasticities and of the welfare measure we're calling 'neediness.' These quantities are useful to know for a wide variety of purposes, as knowing these may allow one to conduct any of a number of interesting counterfactual exercises.

Table 1.2: Average treatment effects for value of consumption of different goods from ANCOVA regression, along with estimated Frisch elasticities β_i (proportional to income elasticities). Controls include baseline values of expenditures, household size, and numbers of children. Asterisks indicate statistical significance at 90, 95, or 99% level of confidence.

Goods	N	CTL	TUP	Diff.	β_i
Beans	464.000	-0.034^{**}	0.033^{**}	0.067^{***}	0.236^{***}
		(0.017)	(0.015)	(0.022)	(0.021)
Bread	311.000	-0.015	0.014	0.029	0.252^{***}
		(0.024)	(0.021)	(0.032)	(0.033)
Cosmetics	397.000	-0.079^{***}	0.080^{***}	0.160^{***}	0.514^{***}
		(0.028)	(0.027)	(0.039)	(0.032)
Egg	91.000	-0.048	0.050^{**}	0.098^{**}	0.186^{***}
		(0.044)	(0.020)	(0.049)	(0.046)
Fish	420.000	-0.034^{*}	0.036^{**}	0.070^{***}	0.224^{***}
		(0.020)	(0.016)	(0.026)	(0.026)
Fruit	114.000	-0.028	0.028	0.056	0.234^{***}
		(0.046)	(0.042)	(0.062)	(0.059)
Fuel	521.000	-0.032	0.030	0.062	0.627^{***}
		(0.034)	(0.029)	(0.045)	(0.036)
Maize	308.000	-0.063^{*}	0.063^{**}	0.125^{***}	0.233^{***}
		(0.037)	(0.030)	(0.047)	(0.051)
Meat	169.000	-0.053	0.055	0.109^{*}	0.210^{***}
		(0.042)	(0.040)	(0.058)	(0.051)
Millet	59.000	-0.044	0.101	0.144^{*}	-3.172^{***}
		(0.048)	(0.070)	(0.085)	(0.268)
Oil	514.000	-0.024	0.022	0.045^{*}	0.322^{***}
		(0.021)	(0.016)	(0.026)	(0.024)
Rice	415.000	-0.016	0.016	0.032	0.252^{***}
		(0.021)	(0.018)	(0.027)	(0.026)
Salt	535.000	0.002	-0.002	-0.004	-0.002
		(0.006)	(0.004)	(0.007)	(0.007)
Soap	543.000	-0.077^{***}	0.080^{***}	0.157^{***}	0.635^{***}
		(0.028)	(0.025)	(0.038)	(0.026)
Sorghum	211.000	-0.028	0.023	0.051	0.171^{***}
		(0.031)	(0.027)	(0.041)	(0.039)
Sugar	513.000	-0.023	0.020	0.043	0.370^{***}
		(0.021)	(0.016)	(0.027)	(0.023)
Sweetpotato	57.000	0.021	-0.036	-0.057	0.280^{***}
		(0.035)	(0.060)	(0.069)	(0.091)
Transport	116.000	0.009	-0.026	-0.035	0.704^{***}
		(0.061)	(0.060)	(0.086)	(0.068)
Vegetables	512.000	-0.054^{**}	0.052^{***}	0.106^{***}	0.362^{***}
		(0.023)	(0.018)	(0.029)	(0.026)
$\overline{\log \lambda}^g$	554	0.138^{***}	-0.060^{***}	-0.198^{***}	—

Table 1.3: Average treatment effects for number of assets of different types from ANCOVA regression; controls include baseline values of dependent variable, household size, number of children, and log neediness. Asterisks indicate statistical significance at the 90, 95, or 99% level of confidence. Estimates in columns labeled "CTL" and "TUP" do not control for $\log \lambda$.

Asset	CTL	TUP	Diff. (no $\log \lambda$)	Diff. (with $\log \lambda$)	$\log \lambda$
Bed	-0.30	0.64	0.93	0.68	-1.28^{*}
	(0.37)	(0.61)	(0.72)	(0.72)	(0.66)
Bicycle	0.00	0.01	0.01	-0.01	-0.06^{***}
	(0.02)	(0.01)	(0.02)	(0.02)	(0.02)
Chairs & tables	0.04	0.24^{***}	0.20^{*}	0.09	-0.56^{***}
	(0.07)	(0.07)	(0.10)	(0.10)	(0.10)
Cows	0.07	-0.05	-0.12	-0.17	-0.26
	(0.17)	(0.05)	(0.17)	(0.17)	(0.16)
Fan	-0.01	0.01	0.02	0.01	-0.05^{***}
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Mobile	-0.02	0.11^{**}	0.13^{**}	0.06	-0.33^{***}
	(0.04)	(0.04)	(0.06)	(0.06)	(0.05)
Motorcycle	0.01	-0.00	-0.01	-0.02	-0.03^{*}
	(0.01)	(0.01)	(0.02)	(0.02)	(0.02)
Mosquito Net	0.14^{***}	0.05	-0.09	-0.14^{**}	-0.24^{***}
	(0.04)	(0.04)	(0.06)	(0.06)	(0.06)
Poultry	-1.13^{***}	1.40^{***}	2.53^{***}	2.33^{***}	-1.00^{***}
	(0.11)	(0.20)	(0.22)	(0.22)	(0.22)
Radio	0.02	0.02	0.00	-0.01	-0.08^{***}
	(0.02)	(0.01)	(0.02)	(0.02)	(0.02)
Sewing	-0.02	0.04	0.06^{*}	0.06	-0.04
	(0.03)	(0.03)	(0.04)	(0.04)	(0.04)
Shed	-0.02^{**}	0.03^{***}	0.06^{***}	0.04^{***}	-0.07^{***}
	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)
Shop	-0.00	0.01	0.01	-0.00	-0.06^{***}
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Small animals	0.09	-0.02	-0.11	-0.22	-0.59^{*}
	(0.33)	(0.08)	(0.34)	(0.34)	(0.32)
Tv	0.01	-0.00	-0.01	-0.01	-0.02^{**}
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Total	-1.21^{*}	2.46^{***}	3.67^{***}	2.73^{***}	-4.71^{***}
	(0.66)	(0.71)	(0.97)	(0.95)	(0.90)

Table 1.4: Average treatment effects for value of assets of different types from ANCOVA regression. Asterisks indicate statistical significance at the 90, 95, and 99% level of confidence. Estimates in columns labeled "CTL" and "TUP" do not control for $\log\lambda$.

Asset	CTL	TUP	Diff. (no $\log\lambda$)	Diff. (with $\log\lambda$)	$\log\lambda$
Bed	2.82	18.57**	15.75	0.78	−75.56***
	(9.36)	(9.23)	(13.14)	(12.80)	(12.27)
Bicycle	1.47	3.23	1.76	−1.57	−16.84**
	(5.69)	(4.78)	(7.43)	(7.40)	(7.26)
Chairtables	−0.29	13.53***	13.83**	7.53	−31.92***
	(4.67)	(4.72)	(6.64)	(6.52)	(6.31)
Cows	−12.54	18.22	30.76	14.12	−84.55***
	(16.69)	(18.63)	(25.01)	(24.77)	(24.68)
Fan	−0.07	0.66	0.73	0.46	−1.37
	(0.96)	(0.81)	(1.26)	(1.26)	(1.26)
Mobile	1.92	6.79**	4.87	−1.46	−32.05***
	(3.85)	(3.23)	(5.02)	(4.85)	(4.73)
Motorcycle	25.43	−11.31	−36.73	−54.23	−88.49***
	(29.04)	(18.70)	(34.54)	(34.32)	(33.87)
Net	1.13**	0.34	−0.79	−1.36*	−2.94***
	(0.54)	(0.46)	(0.71)	(0.70)	(0.69)
Poultry	−37.10***	46.50***	83.61***	76.89***	−33.97***
	(4.07)	(6.72)	(7.86)	(7.79)	(8.07)
Radio	1.59	1.84	0.24	−1.99	−11.30***
	(2.39)	(2.08)	(3.17)	(3.13)	(3.06)
Sewing	3.26	−1.99	−5.25	−6.32	−5.36
	(3.91)	(2.29)	(4.53)	(4.52)	(4.51)
Shed	−2.54	3.99*	6.53**	4.32	−11.28***
	(1.89)	(2.04)	(2.78)	(2.74)	(2.75)
Shop	2.41	0.02	−2.38	−9.77	−37.38***
	(7.16)	(5.19)	(8.84)	(8.69)	(8.67)
Small animals	−23.26**	32.79***	56.04***	46.61***	−48.66***
	(10.05)	(12.45)	(16.00)	(15.89)	(15.67)
Tv	2.73	−1.62	−4.35	−6.47*	−10.69***
	(3.25)	(2.17)	(3.91)	(3.88)	(3.88)
Total	−37.57	131.75***	169.32***	71.40	−494.94***
	(49.76)	(42.89)	(65.69)	(62.84)	(59.53)

Table 1.5: Average treatment effects for nature of self-employment, from ANCOVA regression. Asterisks indicate statistical significance at 95% level of confidence. Estimates in columns labeled "CTL" and "TUP" do not control for $\log\lambda$.

Self-employment	N	CTL	TUP	Diff. (no $\log\lambda$)	Diff. (with $\log\lambda$)	$\log\lambda$
In business	229	-0.02	0.03^{**}	0.05^{***}	0.05^{***}	0.01
		(0.01)	(0.01)	(0.02)	(0.02)	(0.02)
Cultivating	452	0.03^{***}	0.01	-0.02	-0.02	-0.01
		(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
Livestock business	229	-0.05^{***}	0.12^{***}	0.17^{***}	0.16^{***}	-0.07^{***}
		(0.01)	(0.01)	(0.02)	(0.02)	(0.02)

Table 1.6: Occupations of individuals in surveyed households, along with average effects for control and TUP groups. Occupations with fewer than 30 people are excluded. Estimates in columns labeled "CTL" and "TUP" do not control for $\log \lambda$.

Occupation	N	<17	CTL	TUP	Diff. (no $\log \lambda$)	Diff. (with $\log \lambda$)	$\log \lambda$
Student	1932	1484	0.16^{***}	0.09^{*}	-0.07	-0.10	-0.14^{*}
			(0.06)	(0.05)	(0.07)	(0.07)	(0.08)
Cultivation	357	34	0.04	0.00	-0.03	-0.04	-0.04
			(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
Idle	308	212	-0.01	0.04	0.05	0.02	-0.14^{***}
			(0.03)	(0.03)	(0.04)	(0.04)	(0.04)
Beggar	278	184	0.05	-0.01	-0.05	0.03	0.41^{***}
			(0.04)	(0.03)	(0.05)	(0.05)	(0.05)
Housewife	193	8	0.02	0.00	-0.02	-0.04^{**}	-0.11^{***}
			(0.02)	(0.01)	(0.02)	(0.02)	(0.02)
Seeking employment	134	29	0.00	0.01	0.01	0.00	-0.03
			(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
Vegetable farming	126	0	0.03	-0.01	-0.03	-0.01	0.13^{***}
			(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
Small business	98	1	0.01	0.00	-0.00	0.00	0.02
			(0.01)	(0.01)	(0.02)	(0.02)	(0.02)
Ag. Laborer	78	4	0.03^{**}	-0.02^{***}	-0.05^{***}	-0.03^{**}	0.10^{***}
			(0.01)	(0.01)	(0.02)	(0.02)	(0.02)
Skilled labor	56	0	-0.00	0.01	0.01	0.01	-0.03^{**}
			(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Driver	41	1	0.01	-0.01	-0.02	-0.02	-0.00
			(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Non-ag Laborer	31	1	-0.01^{***}	0.02^{**}	0.03^{***}	0.03^{***}	-0.02^{*}
			(0.01)	(0.01)	(0.01)	(0.01)	(0.01)

Chapter 2

Modeling Welfare and Uncertainty within a Frisch Demand System

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 led 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, λ_{it}, as a theoretically grounded welfare index. With this approach, we find that household welfare does not change 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 a small statistically insignificant improvements after two years, followed by a larger and precisely estimated effect after four years. This matches reasonably with the point estimates reported in Bandiera et al. (2017).

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 negative welfare implications of uncertainty in welfare. With this approach, we find that program participants are made better off not only through higher average household welfare, but through a fall in period-to-period variation, suggestive of a greater ability to smooth welfare over time. In fact, the welfare effects of this reduction in risk exposure are found to be similar in magnitude to the welfare gains from improvements in the year-to-year average.

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 "Targeting 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, Duflo, 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 co-vary 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 consumption will vary systematically with household welfare. For example, if inferior goods bought exclusively by the poor are not included, 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 of 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 and states can itself be quite costly for poor households, as they work to insure against negative shocks, spend money (or forego investments) on commitment savings devices, and pay for credit services. 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 a flexible analog to this measure using disaggregated expenditure panels and 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 of offering them 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[1]. 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

[1]For more details about selection, see the original paper.

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 negative implications of uncertainty. First however, we will consider treatment effects on households' marginal utility in expenditures, denoted by the parameter λ_{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 $(c_1, \ldots, c_J)$, p a J-vector of prices, and $u_j(C)$ the marginal utility of consumption of good j. Starting with the standard first-order condition of a consumer's problem and taking logs, we have the additively separable form:

$$(2.1)\quad \log u_j\left(c\right) = \log p_j + \log \lambda$$

In each period, the consumer chooses consumption and the multiplier on the budget constraint λ_{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 λ_{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

$$(2.2)\quad u_j^*(c, z) = u(c)\exp(\gamma_j z)$$

so that household characteristics enter into (2.1) linearly. We can further parameterize the households vector of marginal utility functions with a $J \times J$ matrix of consumption elasticities with respect to marginal utility, β, and a J-vector of demand shifters related to a given item's budget share, α.[2] Letting x be the consumer's maximum total expenditure, this yields for each good

$$(2.3)\quad -\beta^{-1}(\log C + \log \alpha) = \log p + \log \lambda(x,p) - \gamma z$$

where $\lambda(x,p)$ is implicitly defined by

$$(2.4)\quad \sum_{j \in J} e^{\alpha_j} p_j^{1-\beta_j} \lambda(x,p)^{-\beta_j} = x$$

Finally, moving from consumption to expenditures and allowing for classical measurement error ϵ, let $x_j = p_j c_j e^{\epsilon_j}$. Then we can characterize a consumer's system of demands as

$$(2.5)\quad \log X = \log \alpha + (I - \beta)\log p + \beta\gamma z - \beta \log \lambda(x,p) + \epsilon$$

Estimating the Model

We can now move towards estimating the parameters of (2.5). 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

$$(2.6)\quad \log(x_{ijtm}) = \log \alpha_j + (1 - \beta_j)\log p_{jtm} + \delta_{it} z_{it} + ce_{ijt}$$

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 $\log \lambda_{it}$. These fixed effects simplify (2.6) to

$$(2.7)\quad \log(x_{ijtm}) = \alpha_{jtm} + \delta_{it} z_{it} + ce_{ijt}$$

Ligon (2016b) lays out the strategy for estimating β_j and $\log \lambda_{it}$ up to a common factor ϕ, such that $\beta_j(\log \lambda_{it})$ forms the least-squares approximation of the $(N \times n)$ residual matrix,

[2] Ligon (2016a) discusses some of the merits and implications of this sort of functional form in more detail.

ce_{ijt}. In estimating average treatment effects on "neediness", we can choose any satisfying normalization of ϕ (Since the units of logλ/ϕ 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 ϕ and β_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 (2.6), which provides estimates of $(1 + \beta_j)$ as well as δ_{it}. Combined with the decomposition of ce_{ijt}, this serves to identify ϕ, giving us unique estimates of β, and $\log \lambda_{it}$.

Characterizing the Expenditure System

We start by presenting estimates of β_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 estimates of $\log \alpha$, the preference parameter which scales demand for each good. 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. It may be interesting that there does not seem to be a clear correlation between $\phi\beta$ and α. As demand elasticities, we can think of β as a vector of weights determining how important expenditure in each good is to our final estimate of $\log \lambda_{it}$. The vector α by contrast relates to each good's *share* of total expenditures and so suggests how heavily each plays into measures of aggregate expenditures. In this light, the lack of a clear correlation between them goes to show that, while $\log \lambda_{it}$ is closely related to the total expenditures of household i, it is clearly not simply a different way of measuring the same thing.

Table 2.1: Estimated elasticities of demand, $\phi\beta_j$ with respect to λ_{it}. (For goods with %Zero<95)

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

Turning to the distribution of household welfare in Figure 2.2, 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 2.3 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 2.4 presents the analogous results for aggregate household

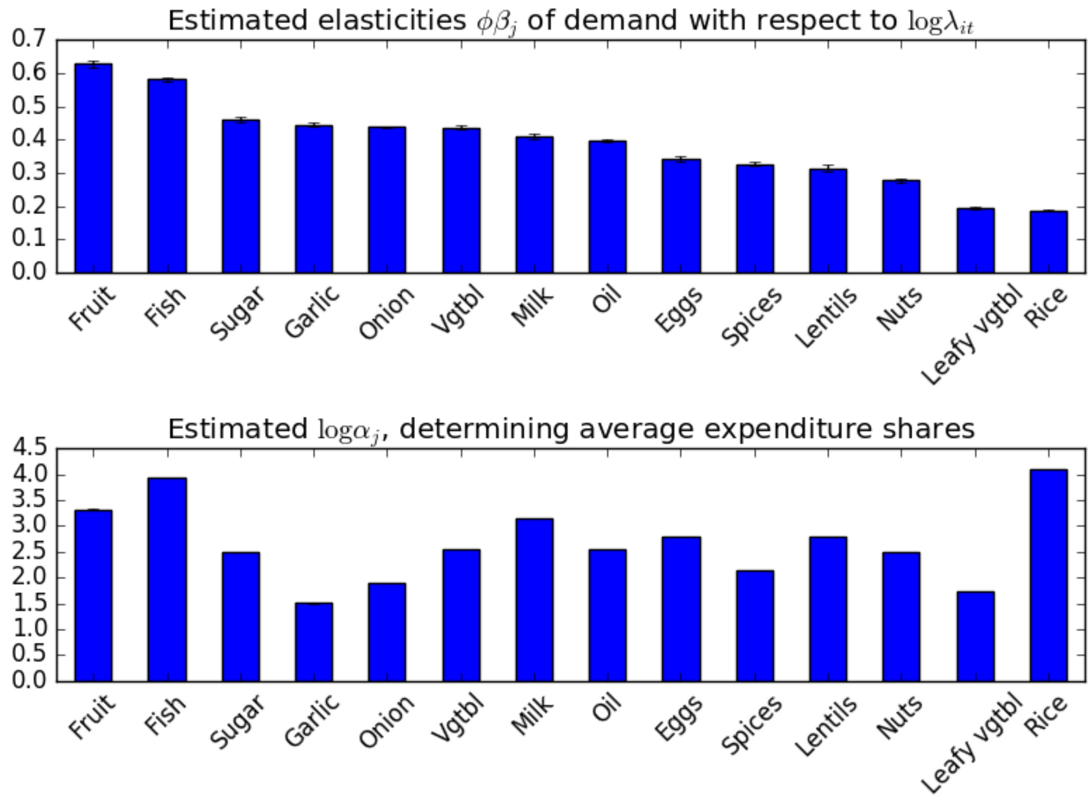

Figure 2.1: β_j for each good. Note that we're imposing here that elasticities do not vary by time or household.

expenditures by eligibility for each year. We can see again that households deemed eligible for participation in the ultra-poor program do in fact have lower expenditures on average in each year. However, both Figure 2.3 and Figure 2.4 reveal that while households targeted by the Targeting the Ultra-Poor households are indeed *more* poor, the distributions are far from disjoint, and seem drawn from much the same support. It is by no means certain that a given household in the TUP program will be worse off than a given household excluded by the targeting mechanism.

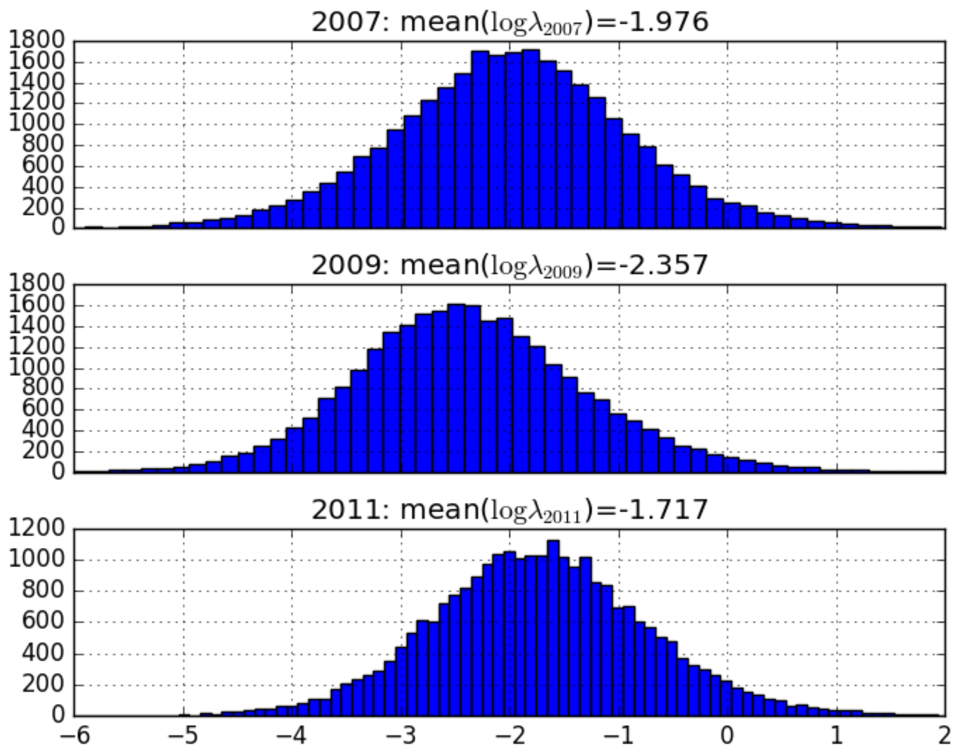

Figure 2.2: Histogram of $\log \lambda_{it}$ by year.

Estimating Vulnerability

Thus far, we have provided a practical and theoretically grounded analog to the standard aggregate consumption metrics. Spending enters into this standard 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 λ_{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 model of household demand, consumption smoothing behavior is based on an objective of smoothing marginal utility over time, typically expressed as some martingale condition. We lean on this insight as we modify the Ligon-Schechter formula to

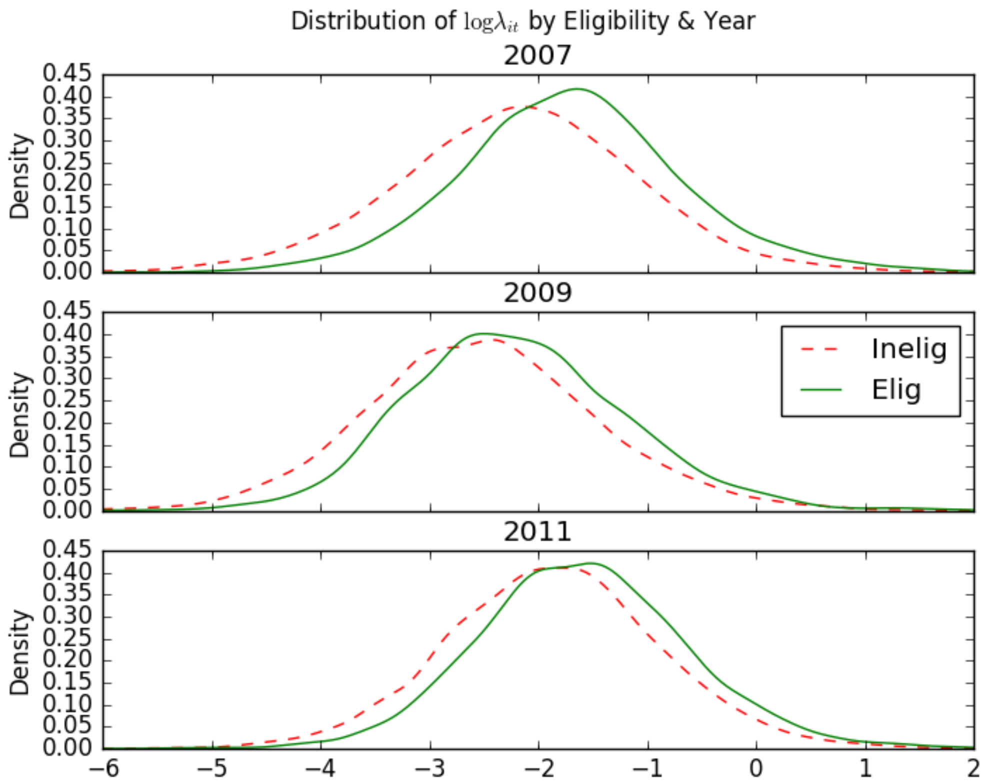

Figure 2.3: Distribution of $\log \lambda_{it}$ by HH Eligibility. We can see that the eligible households have consistently higher marginal utilities in each year.

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.

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) take U to be an indirect utility function and treat vulnerability as a shortfall of a household's expected utility from some certainty-equivalent level $\overline{x}$, yielding the formula

(2.8) $U(\overline{x}) - EU(x_i)$

The logic of this approach is that, since U is concave and increasing in x_i, a household

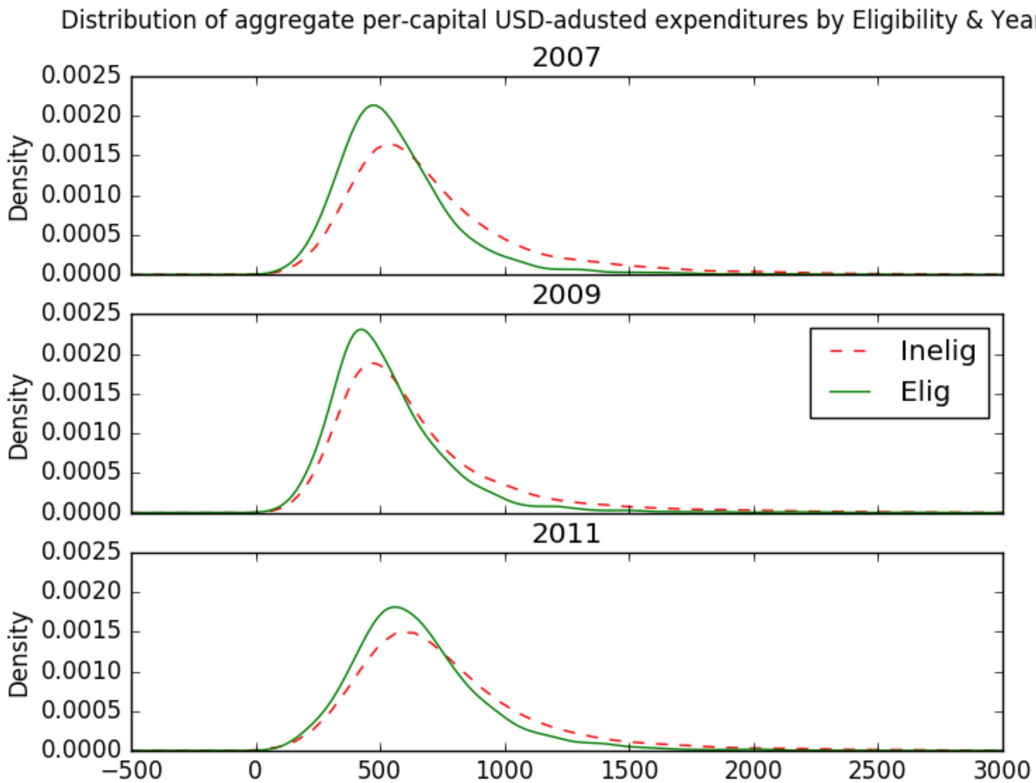

Figure 2.4: Distribution of aggregate expenditures by HH Eligibility. We can see that the eligible households have consistently lower consumption in each year.

is taken as less vulnerable as the mean of x_i increases or as the variance decreases. They decompose this expression into two conceptually distinct quantities, which they call "Poverty" and "Risk".

$$\begin{aligned} &U(\overline{x}) - U(E[x_i]) \qquad \text{(Poverty)} \\ &+ U(E[x_i]) - E[U(x_i) \qquad \text{(Risk)} \end{aligned} \tag{2.9}$$

The first quantity represents the consumption shortfall (or surplus) experienced by household i in an *average* period, relative to some set poverty line, $\overline{x}$. In their case, they take $\overline{x}$ to be the average consumption for all households in all periods. 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 x_{it} due to frictions that prevent ideal consumption smoothing (e.g. behavioral biases, credit constraints).

We take a similar approach with an indirect utility function of x and prices, but instead of treating x simply as the sum of observed expenditures, we calculate the level of total expenditures implied by our estimates of λ_{it}. This approach is appealing in part since, where

Ligon & Schechter must impose some reasonable univariate utility function, we are able to use parameters that have been empirically estimated for this particular sample and time, namely α and β.

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

$$(2.10)\quad U(C) = \sum_{j\in J} \alpha_j \frac{\beta_j}{\beta_j - 1}\left[c_j^{\frac{\beta_j - 1}{\beta_j}} - 1\right]$$

Here C is a vector of consumption of n goods, with β_j governing concavity for any given good, and α_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 an *indirect* utility function. Fortunately, we have the Frischian demand functions for each good j

$$(2.11)\quad \widehat{c_{ijt}}\left(\lambda_{it}, p_{ijt}\right) = \left(\frac{p_{ijt}\lambda_{it}}{\alpha_j}\right)^{-\beta_j}$$

and a corresponding Frischian expenditure function,

$$(2.12)\quad x_{it}\left(\lambda_{it}, p_{ijt}\right) = \sum_{j\in J} p_{ijt}^{1-\beta_j}\left(\frac{\lambda_{it}}{\alpha_j}\right)^{-\beta_j}$$

Plugging these demands into (2.10) and using estimated λ_{it} for all households, β_j and α_j for all goods, and inferring from stated unit values the disaggregated price panel p_{ijt}, we have the indirect Frisch utility function

$$(2.13)\quad U(\lambda_{it}, p_{ijt}) = \sum_{j\in J} \alpha_j \frac{\beta_j}{\beta_j - 1}\left[\left(\frac{p_{ijt}\lambda_{it}}{\alpha_j}\right)^{1-\beta_j} - 1\right]$$

This is an interesting and useful function in its own right. However for our purposes, since we are interested in measuring welfare loss due to uncertainty, it's essential that we be dealing with a concave function, and the expression in (2.13) is clearly convex in λ_{it}. Instead, recall that $\log\lambda$ is implicitly defined in (2.4) as a function of prices and aggregate expenditures, x. We use this fact to take our flexible parameterization and disaggregated preference structure back into a concave Marshallian indirect utility function.

When β_j is allowed to vary across goods, we cannot in general invert (2.12) to obtain an analytical expression for $\log \lambda_{it}(x_{it}, p_{it})$, but it can be readily calculated numerically (Ligon, 2017). This allows us to back up and treat (2.13) as a function of x, leaving us with $U(\lambda_{it}(x_{it}, p_{it}), p_{it}) = U(x_{it}, p_{it})$, a familiar concave and increasing Marshallian indirect utility function like the one used in Ligon and Schechter (2003), but with a flexible, empirically estimated parameterization.

A Frischian Measure of Vulnerability

With this Marshallian utility function in hand, we can define a Frischian analog to the Marshallian definition of vulnerability in (2.8). Letting vulnerability be denoted $V(x_i)$, this yields

$$V(x_i, p_m) = U(\overline{x}, p) - EU(x_i, p_m) \tag{2.14}$$

The question arises what value is appropriate for $\overline{x}$. 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 distinction between distinct classes. Lacking the sort of nuanced contextual information such a judgment requires, we let $\overline{x}$ be the mean level of λ across all household-year observations. This gives a nice interpretation to $V(x_i, p) = 0$, which indicates that in each period, household i achieves the average level of welfare found in the panel, and reaches it without any uncertainty.

As in (2.9), we can break this expression into two conceptually distinct quantities, which we'll be calling "Poverty" and "Risk":

$$\begin{aligned} V(x_i, p) &= [U(\overline{x}, p) - U(E[x_i], p)] && \text{(Poverty)} \\ &+ [U(E[x_i], p) - E[U(x_i, p)] && \text{(Risk)} \end{aligned} \tag{2.15}$$

We then take $E[x_i]$ to be household i's average level of x_{it} over time, $(1/T)\sum_{t \in T} \hat{x_{it}}$. We can define our relative poverty line accordingly as the average of $E[x_i]$ over all households, $\overline{x} = (1/N)\sum_{i \in I} E[x_i]$, the average of x_{it} 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 λ_{it}. 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 baseline 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 (through their effect on

x_{it}). Our index preserves the essential feature that expected utility decreases as variance in the distributions of x_{it} or p_{jmt} increase, but unlike the Ligon-Schechter index, it does so in proportion to β_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 decompose variation in welfare by aggregate and idiosyncratic observables. Where k indexes individual-level observables, we denote aggregate economic variables by $\overline{w_{mt}}$ and idiosyncratic variables as w_{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 x_{it} is then projected onto the first k variables in the list for $k \in 1, ..., K$. 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

$$
\begin{aligned}
V\left(x_{it}\right) = {} & U\left(\overline{x}\right)) - U\left(E\left(x_i\right)\right) \qquad \text{(Poverty)} \\
& + U\left(E\left(x_i\right)\right) - EU\left(E\left(x_i \mid \overline{w}_{mt}\right)\right) \qquad \text{(Aggregate Risk)} \\
& + EU\left(E\left(x_i \mid \overline{w}_{mt}\right)\right) - EU\left(E\left(x_i \mid \overline{w}_{mt}, w_{it}^1\right)\right) \qquad (\downarrow\text{Idiosyncratic Risk}) \\
& \vdots \\
& + EU\left(E\left(x_i \mid \overline{w}_{mt}, w_{it}^1, \ldots, w_{it}^{k-1}\right)\right) - EU\left(E\left(x_i \mid \overline{w}_{mt}, \ldots, w_{it}^k\right)\right) \\
& + EU\left(E\left(x_i \mid \overline{w}_{mt}, \ldots, w_{it}^k\right)\right) - EU\left(x_i\right) \qquad \text{(Unexplained Risk)}
\end{aligned}
\tag{2.16}
$$

Consider that if x_{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 (2.16) 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 $V\left(x_{it}\right)$.

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

$$
\begin{aligned}
V(x_{it}) &= U(\overline{x}) - U(E(x_i)) \qquad \text{(Poverty)} \\
&\quad + U(E(x_i)) - EU(E(x_i \mid \overline{w}_{mt})) \qquad \text{(Aggregate Risk)} \\
&\quad + EU(E(x_i \mid \overline{w}_{mt})) - EU\left(E\left(x_i \mid \overline{w}_{mt}, \text{income}^1_{it}\right)\right) \qquad \text{(Income Risk)}
\end{aligned}
\tag{2.17}
$$

Estimating Vulnerability

We can briefly specify the empirical analogs to the conditional expectations specified above. The unconditional expectations $\overline{x}$ and $E(x_i)$ are simply sample averages. Meanwhile, $E(x_i \mid \overline{w}_{mt})$ is calculated as the average level of x_{it} for each good in a household's market area m and period t. Finally, the conditional expectations, conditioning on the first k observables, $E\left(x_i \mid \overline{w}_{mt}, w^1_{it}, \ldots, w^k_{it}\right)$ 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 (2.16) related to idiosyncratic risk, we estimate the parameters of

$$
\log x^k_{ijt} = \eta_t + \delta_i + \beta W^k_{it} + e_{ijt} \tag{2.18}
$$

where W_{it} is a matrix of the first k observable variables. The fitted values are then used as the conditional expectations in (2.16). 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 2.5 shows us the distribution of total vulnerability based on (2.16) using estimated x_{it}'s. 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. As expected, risk is strictly positive.

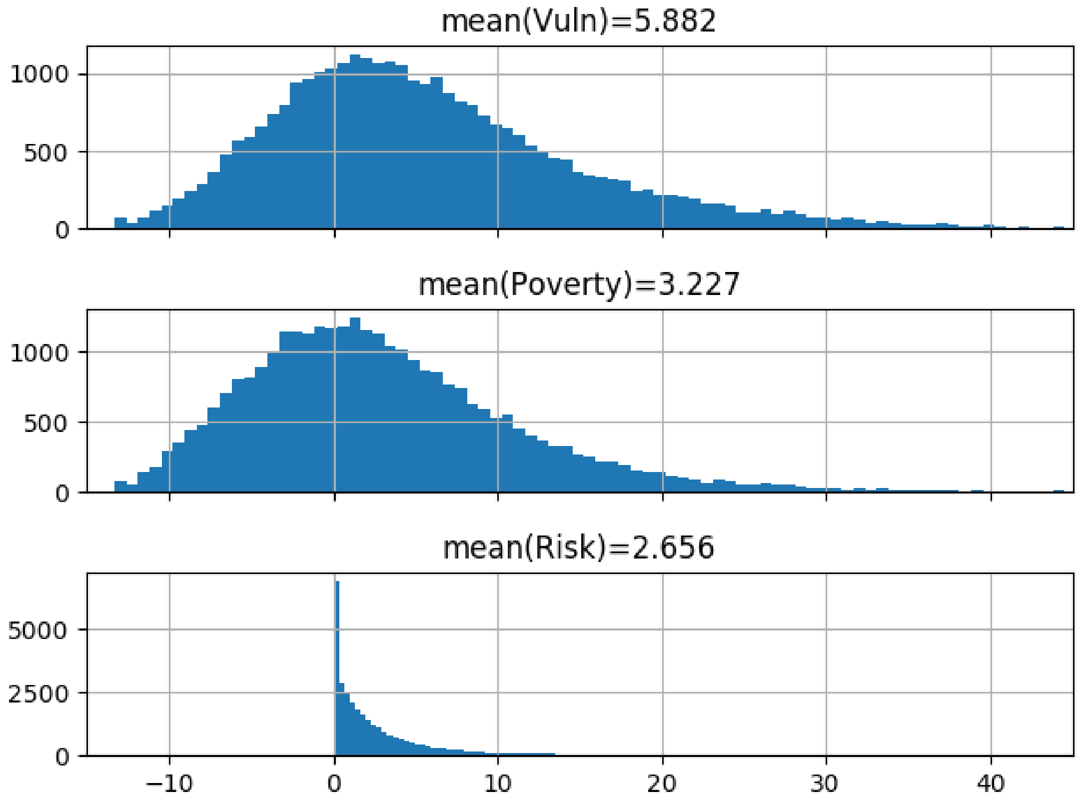

Figure 2.5: Histograms of Vulnerability, Poverty, and Risk for the entire Bangladesh Sample

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 2.2 estimates the joint linear model with all of these covariates to how us how each relates to vulnerability, holding several important factors constant. Table 2.3 reports the coefficients and standard errors for each univariate regression to show us simple unconditional associations.

As Figure 2.5 would lead you to expect, most of the variation in vulnerability can be explained by variation in poverty, with coefficients on risk proving relatively small in each case. Larger households also appear more vulnerable, especially those with many adults. Female and older household heads tend to be poorer as well, and both factors (but especially gender) are associated with more risk. Cash savings, land ownership, 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. Among the selected variables, the second strongest correlate of vulnerability after land ownership is our simple measure of "food security", namely whether households report having no trouble affording two meals per day. These two binary variables are the only two with a stronger association than gender. Gender and food security

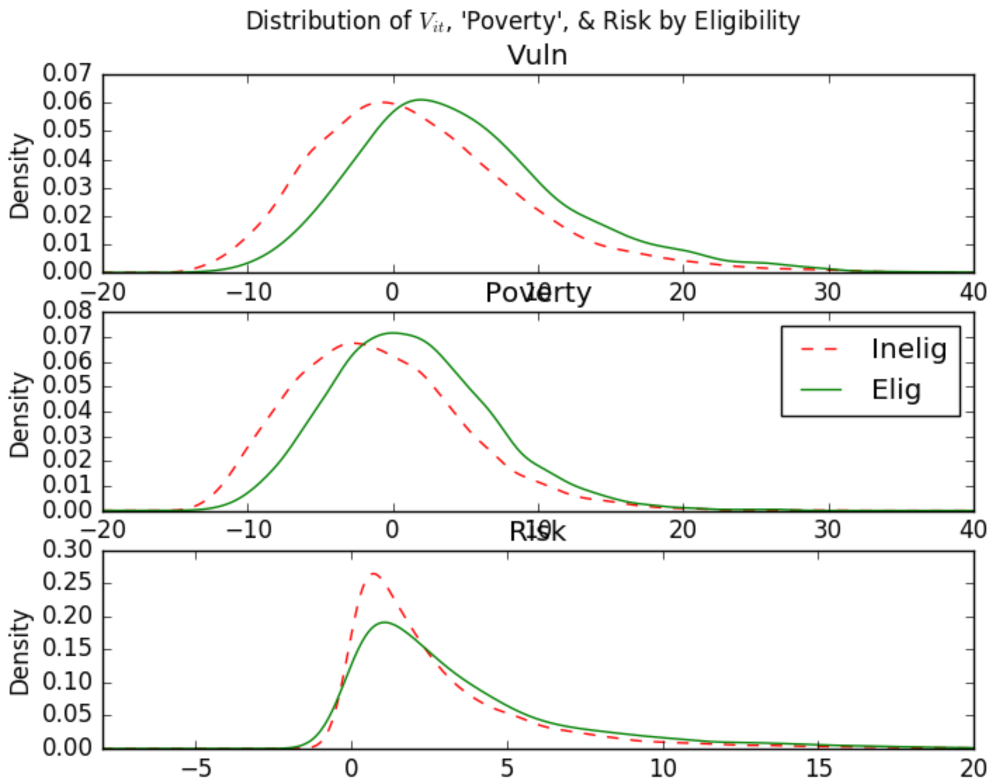

Figure 2.6: Comparing Vulnerability, Poverty, and Risk by Eligibility

are also by far the variables most closely related to risk in this model.

Unsurprisingly, we find that wealthier households rank lower on this poverty index. The wealth ranking we are examining was used to determine perhaps the most interesting baseline characteristic for our context, eligibility in the TUP program. The distributions of Vulnerability, Poverty, and Risk broken down by eligibility status in Figure 2.6. The first two are highly reminiscent of Figure 2.3 in that eligible households are clearly worse off on average, but selected households by no means appear to come from a wholly separate from the general population. This could indicate that the targeting mechanism used was weak in some way. However, discussion of the local economic context suggests that the ultra-poor are in some sense readily identifiable in rural Bangladesh. So perhaps the more likely explanation is that regional variation is more pronounced than within-region variation, and Figure 2.6 pools dozens of regions together.

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. Since aggregate consumption and $\log \lambda_{it}$ form a panel, while vulnerability and its constituents are cross-sectional, we lay out the empirical strategy and results separately for each of these kinds of outcomes. 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

$$Outcome_{itm} = \alpha + \sum_{t\in(1,2)} \beta_t \left(Y_t \times T_{im}\right) + \gamma T_{im} + \sum_{t\in(1,2)} \delta_t Y_t + \eta_m + \epsilon_{itm} \tag{2.19}$$

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 no measurable effect on welfare in the short-term, but improved welfare in 2011, 4 years after enrollment (Table 2.4). Treatment is associated with a statistically significant fall in $\log \lambda_{it}$ in 2011 of 0.104 standard deviations. The first column of Figure 2.7 presents this shift graphically, showing the fall in λ_{it} in the final year. The graph of eligible households in 2009 confirms that the distributions remained very similar for the bulk of the distribution in that year.

We find similar results when looking at aggregate consumption. Treatment is associated with a small and statistically insignificant increase in consumption in 2009 of $20 USD per year (adjusting for PPP and inflation). This matches the comparison for eligible households in 2009 in Figure 2.8, where we see almost identical distributions. Aggregate consumption is calculated as described in Table 5 of Bandiera et al. (2017), and comparing our results to theirs, the point estimates for aggregate expenditures in price-adjusted expenditures per year are similar to our own. They found a modest increase in 2009 of $30 per year, and their point estimates were also statistically insignificant.

We find that this is followed by a more considerable rise in 2011 of \$70.8 per year. Unfortunately, we are not able to replicate their point estimate (\$62.62/year) exactly, but this is arguably not far off. Regardless, our results match theirs qualitatively, touting a small, noisy result in 2009 followed by a larger, precise result in 2011. We also find that aggregate food expenditures increased by \$8 per year, again only in the four-year results from 2011.

Comparing ineligible households with and without neighbors enrolled in the TUP program (Table 2.5), we find that they were actually made better off in both years according to the marginal utility results, by 0.15 and 0.11 standard deviations, respectively. This matches the per capita expenditure results, which find statistically significant results for food consumption in both years of \$33 and \$39, respectively. This did not translate into a change in food consumption for the ineligible households. Despite these linear regression results, we do not see evidence of a precipitous drop in the distribution of $\log \lambda_{it}$ or aggregate expenditures in Figure 2.7 or Figure 2.8.

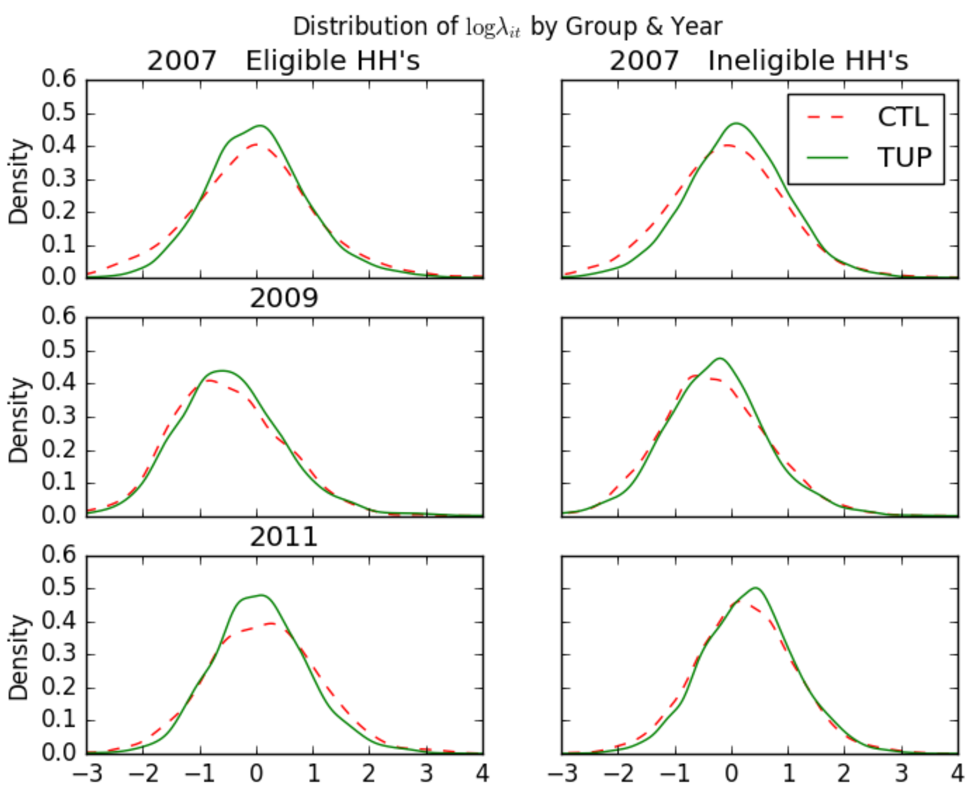

Figure 2.7: Density of $\log \lambda_{it}$ by group and year. We can see that the means are close in 2007 (at baseline), and there is a marked downward shift in neediness for the treatment group in 2011. (Plot currently censored at 1st and 99th percentiles)

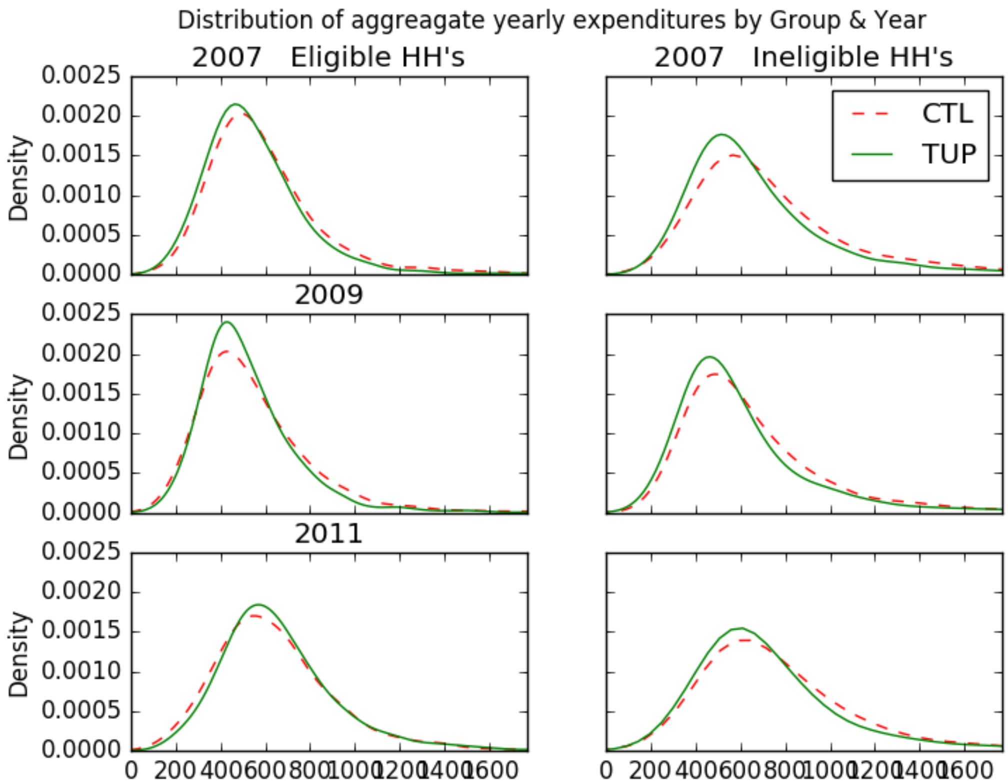

Figure 2.8: Density of HH Expenditures in \$US/year by group and year. We can see that the means are close in 2007 (at baseline) and 2009, and there is a slight upward shift in consumption for the treatment group in 2011. (Plot currently censored at 1st and 99th percentiles to exclude extreme outliers)

Vulnerability, Risk, and Poverty

Empirical Specification

We move now to the cross-sectional analysis of how the TUP program may have affected vulnerability, risk, and poverty. Our specification must be different from that of Bandiera et al. (2017), since we are relying on time-series variation to estimate uncertainty. Instead, we're using the much simpler cross-sectional model

$$Outcome_{im} = \eta_m + \beta T_{im} + \epsilon_{im} \tag{2.20}$$

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

$$[U(\overline{x}) - U(E[x_i])]$$

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

$$[U(E[x_i]) - E[U(x_i)]$$

provides a measure of welfare loss due to period-to-period variation in λ_{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 2.6, with spillovers reported in Table 2.7. 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 2.8, with spillovers reported in Table 2.9.

Results

Finally moving to the results, we find that overall Vulnerability is significantly lower for the treatment households, due to reductions in both Poverty and Risk. These measures all have *utils* as their units, and so may seem hard to interpret. For some economic context, recall the interpretation of $V(x) = 0$. The "poverty line" $\overline{x}$ is defined so that a household will have zero vulnerability if they achieve the panel-wide average level of x in each period, and do so with certainty. To get a sense of relative magnitudes, we report the average level of each measure as well.

The fall in overall vulnerability represents a 22% fall relative to the mean value. While not statistically significant, 38% of this difference is accounted for by a fall in Poverty. The rest is accounted for by an apparent fall in risk facing households. However, looking specifically at the risk that can be explained by household income, we find that households are slightly more exposed to our measure of income risk (8%).

Looking to Table 2.8, when we restrict our sample to the two post-treatment years of the panel as discussed above, we find less vulnerability overall. The average level of household vulnerability is 54% of what it was in the full-panel specification. We still find that vulnerability is reduced by treatment (15%), but find notably different results on its constituents. Poverty has gone down more (30%) and the effect is statistically significant. The treatment effect on poverty is also larger than on vulnerability because the treatment effect on risk appears to be positive. Looking only as post-treatment rounds might suggest that households face more welfare risk, despite being better off overall. Income risk in this specification is lower, but only very slightly relative to the sample average.

Overall, there do not appear to be spillover effects on vulnerability among ineligible households in treated areas, and 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.

Conclusion

In conclusion, we have developed and explored two 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 households' marginal utility has remained stable, but after four years, marginal utility is significantly lower. Our estimates of aggregate consumption and total food consumption have the same signs and similar p-values to the marginal utility estimates, finding no statistically significant change after two years, and a precipitous increase after four. Our results qualitatively replicate the point estimates found in Bandiera et al. (2017), though not precisely. Looking at spillovers in household welfare, we find statistically significant improvements in expenditures and significantly lower marginal utilities (by at least 0.1 standard deviations) in both years.

Turning to treatment effects on vulnerability, risk, and poverty, we find that treatment households are less vulnerable overall due both to decreased poverty (i.e. a smaller shortfall in a household's average welfare), and decreased risk. Looking to the relative magnitudes, there is evidence that, based on the cardinalization we estimate, risk mitigation is actually the more significant part of the TUP's effect on overall welfare.

Future developments of this investigation will hopefully look more closely at quantile effects, especially in light of Figure 2.4 and Figure 2.3. 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.

Table 2.2: Joint linear models associating a range of baseline observables with each component of Vulnerability.

	Vuln	Poverty	Risk	IncRisk
WealthRnk	-0.600^{***}	-0.604^{***}	0.004	-0.124^{***}
	(0.053)	(0.048)	(0.020)	(0.012)
girls	0.885^{***}	0.844^{***}	0.042^{**}	-0.009
	(0.057)	(0.051)	(0.021)	(0.013)
boys	0.792^{***}	0.777^{***}	0.015	-0.051^{***}
	(0.056)	(0.051)	(0.021)	(0.013)
men	0.329^{***}	0.264^{***}	0.064^{*}	-0.115^{***}
	(0.092)	(0.083)	(0.034)	(0.021)
women	0.394^{***}	0.356^{***}	0.039	-0.070^{***}
	(0.100)	(0.090)	(0.037)	(0.022)
age	0.052^{***}	0.041^{***}	0.011^{***}	-0.002^{*}
	(0.004)	(0.004)	(0.002)	(0.001)
gender	1.437^{***}	1.319^{***}	0.118^{*}	0.147^{***}
	(0.172)	(0.155)	(0.064)	(0.041)
pension	0.111	0.075	0.036	0.182^{***}
	(0.140)	(0.126)	(0.052)	(0.033)
OwnLand	-1.693^{***}	-1.608^{***}	-0.085	0.097^{***}
	(0.155)	(0.139)	(0.058)	(0.035)
LandSize	-0.006^{***}	-0.005^{***}	-0.001^{***}	0.000
	(0.001)	(0.000)	(0.000)	(0.000)
Save($100)	-0.120^{***}	-0.101^{***}	-0.018^{***}	-0.008^{**}
	(0.018)	(0.016)	(0.007)	(0.004)
AnySave	-0.207^{*}	-0.140	-0.067	-0.060^{**}
	(0.118)	(0.106)	(0.044)	(0.027)
Debt($100)	-0.084^{***}	-0.076^{***}	-0.009^{**}	0.009^{***}
	(0.011)	(0.010)	(0.004)	(0.002)
AnyDebt	-0.351^{***}	-0.376^{***}	0.025	0.035
	(0.128)	(0.116)	(0.048)	(0.029)
Food Sec	-1.505^{***}	-1.314^{***}	-0.191^{***}	0.142^{***}
	(0.122)	(0.110)	(0.045)	(0.029)
# Cows	-0.378^{***}	-0.321^{***}	-0.057^{***}	0.003
	(0.045)	(0.040)	(0.017)	(0.009)
Assets($100)	-0.014^{***}	-0.013^{***}	-0.002	0.000
	(0.003)	(0.003)	(0.001)	(0.001)
N	28780.000	28780.000	28780.000	19928.000

Table 2.3: 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	-1.917^{***}	-1.783^{***}	-0.134^{***}	-0.112^{***}	30530.000
	(0.039)	(0.035)	(0.014)	(0.008)	
age	0.024^{***}	0.014^{***}	0.010^{***}	-0.002^{***}	28836.000
	(0.004)	(0.004)	(0.001)	(0.001)	
gender	3.005^{***}	2.710^{***}	0.295^{***}	0.294^{***}	28837.000
	(0.138)	(0.125)	(0.048)	(0.031)	
pension	1.640^{***}	1.459^{***}	0.181^{***}	0.254^{***}	30577.000
	(0.143)	(0.130)	(0.049)	(0.032)	
OwnLand	-5.335^{***}	-4.942^{***}	-0.393^{***}	-0.139^{***}	30577.000
	(0.119)	(0.108)	(0.042)	(0.025)	
Save($100)	-0.451^{***}	-0.405^{***}	-0.046^{***}	-0.018^{***}	30577.000
	(0.017)	(0.015)	(0.006)	(0.003)	
Debt($100)	-0.247^{***}	-0.227^{***}	-0.020^{***}	0.001	30577.000
	(0.010)	(0.009)	(0.004)	(0.002)	
Food Sec	-4.041^{***}	-3.651^{***}	-0.390^{***}	-0.029	30577.000
	(0.113)	(0.103)	(0.040)	(0.026)	
# Cows	-1.352^{***}	-1.226^{***}	-0.126^{***}	-0.045^{***}	30577.000
	(0.035)	(0.031)	(0.012)	(0.007)	
Assets($100)	-0.083^{***}	-0.075^{***}	-0.008^{***}	-0.003^{***}	30595.000
	(0.002)	(0.002)	(0.001)	(0.000)	

Table 2.4: Year interactions with T show treatment effects.

Treatment	lambdas	ExpUSD	Food
T	0.001	356.885^{***}	43.082^{***}
	(0.012)	(5.439)	(0.687)
T*Y09	0.015	20.784	0.757
	(0.038)	(16.838)	(2.127)
T*Y11	-0.104^{***}	70.825^{***}	8.030^{***}
	(0.039)	(17.138)	(2.165)
Y09	-0.485^{***}	295.256^{***}	35.747^{***}
	(0.029)	(12.780)	(1.615)
Y11	0.162^{***}	391.529^{***}	46.082^{***}
	(0.029)	(13.028)	(1.646)
const	Mkt-FE	Mkt-FE	Mkt-FE
$Mean_{2007}$	0.0	744.11	84.79
	(1.0)	(514.75)	(57.02)
N	13756.000	13756.000	13756.000

Table 2.5: Year interactions with T show treatment effects.

Spillovers	lambdas	ExpUSD	Food
T	0.000	387.141***	44.128***
	(0.008)	(5.501)	(0.570)
T*Y09	−0.154***	33.488**	1.419
	(0.021)	(15.267)	(1.581)
T*Y11	−0.108***	39.178**	1.601
	(0.021)	(15.496)	(1.605)
Y09	−0.218***	454.925***	52.056***
	(0.013)	(9.464)	(0.980)
Y11	0.330***	577.718***	65.337***
	(0.013)	(9.618)	(0.996)
const	Mkt-FE	Mkt-FE	Mkt-FE
$Mean_{2007}$	0.0	574.62	69.35
	(1.0)	(300.44)	(40.12)
N	40178.000	40178.000	40178.000

Table 2.6: Treatment effects for enrolled households on Vulnerability, Poverty, Risk, and Income Risk

Treatment	Vuln	Poverty	Risk	IncRisk
T	−0.898***	−0.348	−0.550***	0.336***
	(0.324)	(0.244)	(0.174)	(0.070)
const	5.832***	1.878***	3.954***	−4.242***
	(0.284)	(0.214)	(0.153)	(0.062)
$Mean$	4.03	1.52	2.51	−4.15
	(9.07)	(8.19)	(3.08)	(1.73)
N	2512.000	2512.000	2512.000	2512.000

Table 2.7: Spillover effects for ineligible households on Vulnerability, Poverty, Risk, and Income Risk

Spillovers	Vuln	Poverty	Risk	IncRisk
T	−0.358**	−0.340***	−0.018	0.440***
	(0.151)	(0.121)	(0.068)	(0.033)
const	2.349***	−0.497***	2.845***	−4.495***
	(0.112)	(0.089)	(0.050)	(0.025)
$Mean_{2007}$	7.97	5.04	2.93	−3.9
	(9.54)	(8.69)	(3.51)	(1.66)
N	8625.000	8625.000	8625.000	8625.000

Table 2.8: Treatment effects excluding initial baseline data for enrolled households on Vulnerability, Poverty, Risk, and Income Risk

ATE (no BL)	Vuln2Yr	Poverty2Yr	Risk2Yr	IncRisk2Yr
T	-0.346^{**}	-0.529^{***}	0.183^{***}	-0.043^{***}
	(0.142)	(0.159)	(0.034)	(0.016)
const	-1.774^{***}	-1.113^{***}	-0.661^{***}	0.234^{***}
	(0.124)	(0.139)	(0.030)	(0.014)
$Mean$	−2.18	−1.76	−0.41	3.91
	(3.5)	(3.74)	(0.63)	(44.09)
N	2512.000	2512.000	2512.000	2512.000

Table 2.9: Spillover effects excluding initial baseline data for ineligible households on Vulnerability, Poverty, Risk, and Income Risk

Spillovers (no BL)	Vuln2Yr	Poverty2Yr	Risk2Yr	IncRisk2Yr
T	-0.278^{***}	-0.296^{***}	0.019	1.550^{**}
	(0.065)	(0.072)	(0.014)	(0.765)
const	-2.668^{***}	-2.189^{***}	-0.478^{***}	3.569^{***}
	(0.048)	(0.053)	(0.011)	(0.566)
$Mean_{2007}$	−1.09	−0.66	−0.43	−0.97
	(3.65)	(3.9)	(0.66)	(19.75)
N	8625.000	8625.000	8625.000	8625.000

Chapter 3

Comparing Cash and Asset Transfers to Low-Income Households in South Sudan

Abstract

Several previous studies have found that the "graduation" or "Transfers to the Ultra-Poor" (TUP) framework is an effective approach to alleviating the constraints that prevent extremely poor households from increasing their productivity. The framework consists of a sizable transfer of productive physical capital, coupled with training and continuous support over the course of one or two years. A second and related literature suggests that unconditional cash transfers (UCT's) may have a comparable effect on productivity and welfare (with fewer fixed costs). This field experiment, examining the first two years of BRAC's TUP pilot in South Sudan, offers a direct comparison of these very different approaches to alleviating capital constraints. We consider the effect of each on consumption, income, asset holdings, and a number of intangible outcomes. We also consider the TUP program's effect on households' response to the outbreak of violence in 2014, which led to a level of instability in which the Graduation framework has not been previously tested. We find evidence of positive consumption effects from both treatments, but only in the short-run, and a persistent wealth effect only from the TUP. We also elicit suggestive evidence that BRAC's support may have helped TUP beneficiaries cope with the short-term economic effects of the outbreak of violence in 2014. We tentatively conclude that targeted asset transfers can play a constructive role in helping poor, self-employed households when they face economic uncertainty. And while cash increases household consumption, the goal of improving income or wealth is aided by the additional services that the ultra-poor graduation framework offer.

Introduction

Poor rural households typically earn money from low-return activities like small-scale cultivation or casual day labor, and face both financial and human capital constraints, keeping them from investing and expanding into more lucrative activities. Experience and research over many years has led many to believe that households facing particularly acute poverty are unable to solve this problem through the small, high-interest loans typically marketed to them. It was these considerations that led to the development of the initial "Targeting the Ultra-Poor" (TUP) program in Bangladesh as a supplement or precursor to credit services. First implemented by BRAC in 2007, the program aims to simultaneously alleviate physical and human capital constraints by providing households with a significant transfer of food and productive assets, followed by two years of training and support by extension officers. (Bandiera et al. (2017)) The general framework[1] has since expanded to a wide range of countries, with a general pattern of success in increasing aggregate investment, labor supply, and aggregate consumption. (Banerjee, Duflo, et al., 2015)

A second, older literature has gained new interest in parallel with this literature which examines the effect of offering direct unconditional cash transfers (UCT's) to poor households. (Haushofer and Shapiro, 2013) (Blattman, Jamison, et al., 2014) (Blattman, Fiala, and Martinez, 2013) While this and the TUP framework are both direct capital transfer interventions, they are very different in their approach, with TUP programs guiding and constraining the use of capital towards productive investment while UCT's allow households to invest and consume as they see fit. The natural question that arises is how these additional features and constraints in the TUP framework change how households use their capital transfers.

Here, we examine the experimental evaluation of BRAC's pilot TUP program in South Sudan and compare it to a round of unconditional cash transfers. Our results contribute to the general literature in two important ways. First, South Sudan's political and economic institutions have overwhelmingly politically unstable since this study's inception, which may affect the value of the program for households in important ways. Second, a randomly selected group of households received cash transfers equal in market value to the assets provided to the TUP households. While an experimental literature has been established studying the graduation framework in isolation, this is among the first experiments attempting to directly compare it to a obvious alternative investment.

The Program

The pilot program itself was similar to the other TUP programs completed by BRAC. It consisted of four phases: targeting and selection, training and enterprise selection, asset transfers, and monitoring.

[1]Known as the "graduation framework" pointing to the original ambition to move households into an activity where they are able to finance further income growth without costly transfers.

Targeting, Selection, & Training

The fist phase of the program was to complete a census of households in the area around BRAC's office in the town of Yei in Western Equatoria. This census contained questions to assess eligibility for the program. First, households were excluded if they had a salaried worker in the household, were participating in another NGO program, or had no access to cultivable land (which was in some cases necessary for the program's model). Households were then deemed eligible if they fit at least three criteria in a list of five poverty indicators.[2] The census was completed in April of 2013 and 745 were identified as eligible. Of these, 649 were identified in a baseline survey. These households were stratified on employment, asset ownership, and size and selected into treatment groups. 250 were enrolled in the TUP program, 125 in the UCT group, and the final 274 in a pure control group.

Asset Transfers & Monitoring

The second phase of the program was training and enterprise selection. Unlike most programs of this type, the number of households given each kind of asset was set in advance, with 75 enrolled in agricultural activities (vegetable cultivation), 85 in duck rearing, 45 in goat rearing, and the rest in small trade businesses. While the staff tried to map households' asset types to their respective preferences and skills, a disproportionate number stated a preferences for goats and small trade. Households then attended training sessions. The first of these were for general business skills around literacy, numeracy, and financial management. The next were sector specific and focused on how to properly raise livestock or gardens.

Once training is completed, asset transfers began in late 2013 and continued through the first few months of 2014. The productive assets related to each enterprise were valued at around \$240 per household, with a random subset receiving an additional \$60 in assets later in 2014. Shortly thereafter, households started to attend weekly or semi-weekly meetings with other nearby participants to discuss with each other and a BRAC extension officer the details of their businesses. These meetings also included food transfers for a while, which were designed to help get households to the point of receiving revenue from their assets without having to sell them.

In all, the market value of these food transfers were valued at \$110, bringing the total value of all transfers to \$350-\$410. The 125 households in the UCT group were randomly divided in half to receive cash in these amounts. Unfortunately, political instability disrupted NGO operations throughout South Sudan, preventing the simultaneous disbursal of the cash and asset transfers. Instead, a second survey was conducted in June of 2014, with the cash transfers being disbursed immediately thereafter. This resulted in a timing difference of 3 to 6 months between the two.

[2]These criteria were that the household had a head working as a day laborer (generally an occupation with poverty wages), two or more children, at least one child working, fewer than three rooms, or a woman who has not completed secondary school.

The Data

The census was conducted in April of 2013 in the area around BRAC's offices in Yei County to identify women eligible for participation. A baseline survey was conducted that Summer, which successfully interviewed 649 of these women and randomly selected them into the TUP, UCT, and control groups. Half of each beneficiary group was randomly selected to receive additional "top-up" transfers with market value of $60 (around 20% of the original transfers).

In response to the outbreak of violence in late 2013 and subsequent closing of the offices in Yei, a midline survey was conducted in June 2014 to try to separate pre- and post-conflict changes in outcomes. For lack of a valid comparison group, we will in Yei, though we will report estimates of treatment effects on the severity or likelihood of having been effected exposure to the conflict. Some of the original asset transfers were done before the office closure, which may affect estimates of the difference between programs if rates of return changed in the few intervening months. Finally, an endline survey was conducted in mid-2015 to estimate the effect of program participation on households' financial situation and overall welfare. The key here is that the survey conducted in mid-2014 provides us with *short-term* treatment effects of the TUP program within 6 months of the asset transfers, while providing a second baseline for the Cash transfers. Likewise, the 2015 survey allows us to estimate treatment effects one year after the cash transfers, and 15-18 months after the asset transfers.

This unfortunately left us without data past one year for the cash transfer effects. To get some point estimates on household welfare for this group in the slightly longer term, we conducted a series of five short surveys on a monthly basis from November of 2015 to March of 2016. These collected only a subset of the full consumption modules and a few questions tracking major transactions and shocks. The short length of the survey allowed them to be administered via the mobile network, reducing cost and improving response rate.

Empirical Strategy

For the main panel, we estimate a single model using interactions between time effects and group assignment, as well as baseline values of the outcome variable where available.

$$Y_{it} = \sum_{t=2014}^{2015} \delta_t + \beta_t^{Cash} I_t * Cash_{it} + \beta_t^{TUP} I_t * TUP_{it} + \gamma Y_{i,2013} + \epsilon_i$$

where δ_t are time fixed effects and I_t is an indicator if the year t, and Y_{it} is an outcome of interest for household i in year t. We take the interactions of TUP assignment with 2014 and 2015 indicators as the treatment effects at 6-8 and 15-17 months respectively. The analogous interactions with the Cash group offer a second baseline and a 12-month treatment effect, respectively. Since those transfers happened after the midline survey, its interaction with *2014* acts as a placebo; there is no *ex ante* reason to expect that they were different from the rest of the control group at that point. Given the slight difference in timing, we report a t-test of the hypothesis $\beta_{TUP,t} - \beta_{Cash,2015} = 0$ for both $t \in 2014, 2015$. Since the difference in timing is much smaller, we consider $\beta_{TUP,2015} - \beta_{Cash,2015} = 0$ to be the central hypothesis of interest.

For the supplementary analysis of the high-frequency panel, we estimate a separate model, since the underlying data is so different. A constant parameter takes the place of the fixed effects. We include 2013 levels as a covariate where possible. Since we collect expenditures on only ten consumption items, we report not only the total value of spending on those goods, but also a more theoretically grounded measure described in Collins and Ligon (2017), which uses the composition of expenditures to derive the marginal utility of expenditures for each household. We chose ten relatively demand-elastic items specifically for this purpose, as those will tend to be the most responsive to changes in welfare.

Results

Randomization Check

A crucial assumption is that the treatment and control groups were selected appropriately. We check this by presenting summary statistics by group on a range of factors related to consumption, asset holdings, and household characteristics. We check for balance on observables in Table 3.1.

Table 3.1: Means of some analysis variables at baseline. Asterisks indicate p<.1, .05, and .01 respectively

Consumption	CTL	Δ TUP	Δ CSH	N
Meat	4.21	-0.568	-0.052	378
Fuel	0.76	-0.039	-0.072	456
Clothesfootwear	0.67	-0.026	0.033	595
Soap	0.48	-0.008	-0.026	536
Fish	2.50	-0.154	-0.156	474
Charities	0.03	-0.006	0.0	134
Cereals	9.19	-0.947	0.27	605
Transport	0.18	-0.033	0.002	193
Cosmetics	0.68	0.027	-0.125	468
Sugar	1.71	-0.078	-0.189	604
Egg	1.10	-0.091	0.038	276
Oil	1.36	-0.13	-0.141	613
Ceremonies	0.13	0.006	0.026	152
Beans	0.70	0.232	0.226	192
Fruit	0.69	-0.089	0.001	272
Textiles	0.16	-0.004	0.056*	376
Utensils	0.25	-0.009	0.008	442

Continued on next page

Table 3.1: Means of some analysis variables at baseline. Asterisks indicate p<.1, .05, and .01 respectively

Consumption	CTL	Δ TUP	Δ CSH	*N*
Dowry	1.27	-0.041	0.028	126
Furniture	0.20	-0.014	0.045	368
Salt	0.45	-0.026	0.007	617
Vegetables	1.54	-0.165	-0.18	471
Assets	CTL	Δ TUP	Δ CSH	*N*
Smallanimals	236.60	-86.068	-123.133	123
Bicycle	109.08	-12.555	-11.414	171
Radio	58.45	-5.968	-16.529	260
Motorcycle	341.74	192.956	353.836**	93
Net	19.16	0.668	0.247	423
Poultry	42.40	-3.365	-8.894	161
Bed	241.27	7.992	32.762	521
Chairtables	206.79	-29.368	3.617	531
Mobile	97.54	12.627	-4.198	414
Netitn	7.82	1.215	1.178	181
Cosmetics	0.68	0.027	-0.125	468
Household	CTL	Δ TUP	Δ CSH	*N*
Daily Food	25.18	-2.215	-0.261	643
Daily Exp	29.90	-2.167	-0.288	646
No. Houses	2.83	0.031	0.118	543
In Business	0.40	0.038	0.017	265
Cereals	9.19	-0.947	0.27	605
# Child	3.26	0.118	0.108	594
Asset Tot.	1757.05	-44.791	98.654	603
Cash Savings	236.90	28.52	-66.812	431
HH size	7.23	-0.175	0.3	648

This is simply suggestive evidence that the treatment and control groups were similar in observables at baseline, with the exception that the cash group has atypically more motorcycles and clothing. But it does suggests that our stratified randomization was not too far from creating comparable groups.

Attrition

Another crucial question is to what extent attrition in 2014 and 2015 was small or balanced. Table 3.2 reports the total number of households identified in each treatment arm and year for the whole sample. Table 3.3 reports the same numbers restricting ourselves to households with baseline surveys. In the TUP group, we were unable to find 21 participants in 2014

(8% attrition), but found 5 not identified in the baseline survey. We found 8 additional TUP households with baseline surveys again in 2015 for a final attrition rate of 5%. Of those in the Cash group, 12 were lost (9.6%), then two more in 2015 (11%). The control group saw very high attrition in 2014 (22%), but also found a large number of households not found at baseline, yielding a comparison group only 6% smaller. The high attrition was due largely to the fact that these households did not enjoy the same consistent contact that BRAC had with the TUP group, and the local area lacked infrastructure to easily locate people. This was exacerbated by the uncertain political situation and early harvest. Attrition in 2015 was 6.7%, with a full 85 more households identified who were not in the baseline survey. In order to take advantage of the households not included in the baseline, the main specification below includes an indicator for whether the household was in the baseline.

Table 3.2: Total number of households in sample by group and round

Full Sample	2013	2014	2015
Cash	124.0	113.0	111.0
Control	281.0	265.0	347.0
TUP	244.0	228.0	236.0
All	649.0	606.0	694.0

Table 3.3: Number of households in sample with baseline survey by group and round

Balanced Sample	2013	2014	2015
Cash	124.000	112.000	110.000
Control	281.000	219.000	262.000
TUP	244.000	223.000	231.000
All	649.000	554.000	603.000

Next we ask how those who did not turn up in subsequent rounds differed by a range of baseline characteristics. Table 3.4 reports the average level of various characteristics in 2013. Then we report the coefficient from a linear regression on indicators for whether they were in the midline or endline surveys. Here we see that overall, households found in the midline survey were larger with more children and larger reported asset stocks. Households found in 2015 seemed to have, at baseline, significantly smaller asset stocks and less consumption.

Table 3.4: Means of household baseline characteristics and regression coefficients for whether they were ultimately found at baseline or endline. (Note that this does not consider households found only in 2014 or 2015).

HH Features	$Mean_{Bsln}$	β_{Mid}	β_{End}
HH size	7.223	0.595**	0.428
# Child	3.328	0.656***	0.423
Asset Prod.	512.822	126.360	-369.190
Asset Tot.	1494.324	361.889	-689.174*
Daily Exp	25.212	1.257	-4.150
Daily Food	24.300	0.299	-4.790*
In Business	0.415	0.038	0.007
Land Access	2.324	0.014	0.305
No. Houses	2.863	0.305	0.367
Cash Savings	178.662	46.322	54.295
Assets			
Bed	250.534	12.649	-51.133
Bicycle	102.174	11.179	4.212

Continued on next page

Table 3.4: Means of household baseline characteristics and regression coefficients for whether they were ultimately found at baseline or endline. (Note that this does not consider households found only in 2014 or 2015).

HH Features	$Mean_{Bsln}$	β_{Mid}	β_{End}
Mobile	101.482	6.336	-13.028
Motorcycle	481.885	213.002	-241.819
Carts	2.751	1.929	2.962
Cows	181.402	67.862	-89.273
Smallanimals	180.716	18.966	-79.014
Consumption			
Cereals	8.882	-0.084	-3.714**
Beans	0.826	0.269	-0.382
Ceremonies	0.141	-0.020	-0.038
Charities	0.027	0.007	-0.001
Clothesfootwear	0.663	0.180*	-0.206
Cosmetics	0.668	0.005	0.229
Dowry	1.263	0.755	-0.399
Egg	1.069	-0.005	0.106
Fish	2.417	-0.132	0.036
Fruit	0.656	0.009	-0.151
Fuel	0.733	0.105	-0.049
Meat	3.981	0.254	0.300
Other	0.0	0.000	0.000
Poultry	39.437	23.634*	-2.243
Salt	0.438	-0.140***	-0.043
Soap	0.475	-0.181*	0.047
Sugar	1.647	-0.285	-0.020
Textiles	0.165	0.010	0.011
Transport	0.163	0.004	0.018
Tv	39.915	-16.377	0.845
Utensils	0.247	0.062	-0.023
Vegetables	1.446	0.096	-0.151

Consumption

The first measure of welfare we consider is household consumption, defined as the market value of goods or services used by the household. A sizable basket of goods were included in the survey module. These are separated into three categories: Food items (with a 3-day recall window), non-durables (a 30-day recall window), and durables and large expenditures (a one-year recall window). This is perhaps the most appropriate measure of the welfare or poverty of a household in our survey.

The results for several important consumption measures are presented in Table 3.5. Importantly, we do not know about prices for each good in this time, though we can say that inflation was as high as 100% between 2014 and 2015. Nonetheless, we take the sum of all consumption and expenditure questions together as a measure of welfare. In light of the fact that we have data on an incomplete basket, we also follow Collins and Ligon (2015), which details a method for deriving treatment effects on a model-based estimate of households' marginal utility, which we include here as $\log \lambda_{it}$.

The main result is that TUP participants had higher consumption consumption in 2014, several months into the primary monitoring phase after the asset transfers. Similarly, the cash group has higher consumption in 2015, measured just over a year after disbursal. Food transfers had ceased weeks before the 2014 survey was conducted, and the assets had been transferred 6-8 months prior. The TUP group sees no notable effect in 2015. The short-term consumption effects of either program are economically significant, representing a roughly 16% increase in average total consumption for both TUP and Cash. These results are consistent with a story in which either sort of transfer has a short-term consumption effect. Importantly, we do not reject the null hypothesis that the two effects are equal to one another. In either group, the increase in total consumption appears to be driven mainly by increased food consumption, with smaller effects on non-food consumption goods and durables. As such, there is no evidence that the share of food consumed falls, as might be predicted by Engel's law.

Table 3.5: Average treatment effects by Group-Year, controlling for baseline levels.

	Tot	$\log \lambda_{it}$	Food	logTot
CTL mean	115.404	0.159	38.468	4.509^{***}
	(78.750)	(0.967)	(26.250)	(0.756)
CSH*2014	-2.745	0.127	-0.915	0.007
	(8.008)	(0.110)	(2.669)	(0.079)
CSH*2015	18.023^{**}	-0.145	6.008^{**}	0.160^{**}
	(7.831)	(0.108)	(2.610)	(0.077)
TUP*2014	18.590^{***}	-0.365^{***}	6.197^{***}	0.212^{***}
	(6.426)	(0.089)	(2.142)	(0.063)
TUP*2015	4.179	-0.055	1.393	0.045
	(6.130)	(0.084)	(2.043)	(0.060)
2014	76.831^{***}	0.214^{***}	25.610^{***}	3.931^{***}
	(5.318)	(0.062)	(1.773)	(0.113)
2015	105.702^{***}	0.188^{***}	35.234^{***}	4.175^{***}
	(5.001)	(0.057)	(1.667)	(0.111)
Bsln2013	0.081^{**}	0.022	0.081^{**}	0.073^{***}
	(0.038)	(0.029)	(0.038)	(0.026)
Bsln NA	20.521^{***}	-0.119	6.840^{***}	0.447^{***}
	(6.964)	(0.088)	(2.321)	(0.121)
$\beta^{TUP}_{2014} - \beta^{CSH}$	0.566	-0.220	0.189	0.052
	(9.994)	(0.137)	(3.331)	(0.098)
$\beta^{TUP}_{2015} - \beta^{CSH}$	-13.844^{*}	0.090	-4.615^{*}	-0.115
	(8.125)	(0.111)	(2.708)	(0.080)
F-stat	10.142	4.169	10.142	8.131
N	1291.000	1296.000	1291.000	1291.000

Table 3.6: Average treatment effects using mobile data collection (results are robust to controlling for baseline levels)

	$\log \lambda_{it}$	Tot	logTot
CTL mean	-0.018	30.851	3.158^{***}
	(1.001)	(27.768)	(0.734)
TUP	0.023	-0.624	-0.011
	(0.041)	(1.152)	(0.030)
CSH	0.056	0.776	0.028
	(0.052)	(1.459)	(0.038)
const	-0.018	30.851^{***}	3.158^{***}
	(0.027)	(0.753)	(0.020)
$\beta^{TUP} - \beta^{CSH}$	-0.033	-1.399	-0.039
	(0.055)	(1.524)	(0.040)
F-stat	0.584	0.434	0.475
N	2877.000	2878.000	2878.000

This result leaves open the question of whether the cash treatment had a persistent effect on consumption, or whether the short-term effect found in 2015 is similarly temporary. It was this question that motivated the collection of an additional five rounds of data over a 6-month period in late 2015 and early 2016, in which we asked about ten items, five food and five non-food. We consider the average treatment effect on households sampled for these phone interviews, both for $\log \lambda_{it}$ and for "Total consumption", which in this case we take a simple sum over the goods discussed. We find that, consistent with the TUP program's results in 2015, all evidence of an effect seem to be gone by 18th months after the transfer

Food Security

Generally speaking, observed changes in total consumption don't translate into an increase in reported food security. In each year, we ask how often in a given week the respondent has had experiences indicative of food insecurity. Included are (from left to right) going a whole day without eating, going to sleep hungry, being without any food in the house, eating fewer meals than normal at mealtimes, and limiting portions. We report the percentage of people who report experiencing each in a typical week, as well as a standardized composite z-score using all of these questions. There is little evidence of a significant treatment effect at endline in 2015.

Table 3.7: Percentage of respondents reporting a food security problem occurs at least once a week.

	Z-score	Whole Day	Hungry	No Food	Fewmeals	Portions
CTL mean	-0.01	0.21	0.21	0.28	0.32	0.36

Continued on next page

Table 3.7: Percentage of respondents reporting a food security problem occurs at least once a week.

	Z-score	Whole Day	Hungry	No Food	Fewmeals	Portions
	(1.00)	(0.41)	(0.40)	(0.45)	(0.47)	(0.48)
TUP*2014	−0.10	−0.02	−0.05	−0.03	0.01	0.01
	(0.09)	(0.03)	(0.03)	(0.03)	(0.04)	(0.04)
TUP*2015	−0.02	0.03	−0.01	−0.03	0.05	−0.02
	(0.09)	(0.03)	(0.03)	(0.03)	(0.04)	(0.04)
CSH*2014	−0.05	−0.00	−0.04	−0.01	−0.03	−0.00
	(0.11)	(0.04)	(0.04)	(0.04)	(0.05)	(0.05)
CSH*2015	0.03	0.06	0.03	−0.01	−0.00	−0.04
	(0.11)	(0.04)	(0.04)	(0.04)	(0.05)	(0.05)
Bsln2013	0.07**	−0.00	0.02	0.03	0.06**	−0.02
	(0.03)	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)
2014	0.07	0.09***	0.10***	0.09***	0.17***	0.22***
	(0.06)	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)
2015	0.03	0.22***	0.21***	0.26***	0.30***	0.39***
	(0.06)	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)
Bsln NA	−0.17*	−0.02	−0.03	0.03	−0.02	−0.08*
	(0.09)	(0.03)	(0.03)	(0.03)	(0.04)	(0.04)
F-stat	1.45	9.34	8.36	10.84	6.70	5.91
N	1299.00	1282.00	1297.00	1293.00	1297.00	1292.00
$\beta^{TUP}_{2014} - \beta^{CSH}$	−0.13	−0.08	−0.08*	−0.01	0.01	0.05
	(0.14)	(0.05)	(0.05)	(0.05)	(0.06)	(0.06)
$\beta^{TUP}_{2015} - \beta^{CSH}$	−0.06	−0.03	−0.04	−0.02	0.06	0.02
	(0.12)	(0.04)	(0.04)	(0.04)	(0.05)	(0.05)

Assets

Turning now to asset holdings for the households, we estimate treatment effects for total value of assets owned, total value of potentially productive assets, as well as land and financial assets.

Total Asset Holdings

The cash group does not appear to have seen an increase in the value of assets measured, with negative and imprecise point estimates. The most important result is that the TUP group has significantly more asset wealth than the cash or control groups in both 2014 and 2015, 18 months after receipt of transfers. The TUP group has a change of 536 SSP on average (43% increase over controls, $p<.01$). So-called "Productive" assets include anything that could plausibly be used in productive activity.[3] Here we see the TUP group has 320 SSP (95%) more in this area over the control group, with a similar magnitude at midline.

Importantly, this is not due to a precipitous increase in assets reported over this time. Note also that the effect on total assets is higher in absolute value than the effect on productive asset value, suggesting that the increased wealth cannot be explained purely by households holding onto asset transfers for the length of the program's monitoring phase. Instead, the TUP group is the only one for whom total measured asset holdings did not fall on average over these two years, which saw (including the savings effect below) is the only feature of households' financial situation on which we we see a persistent effect.

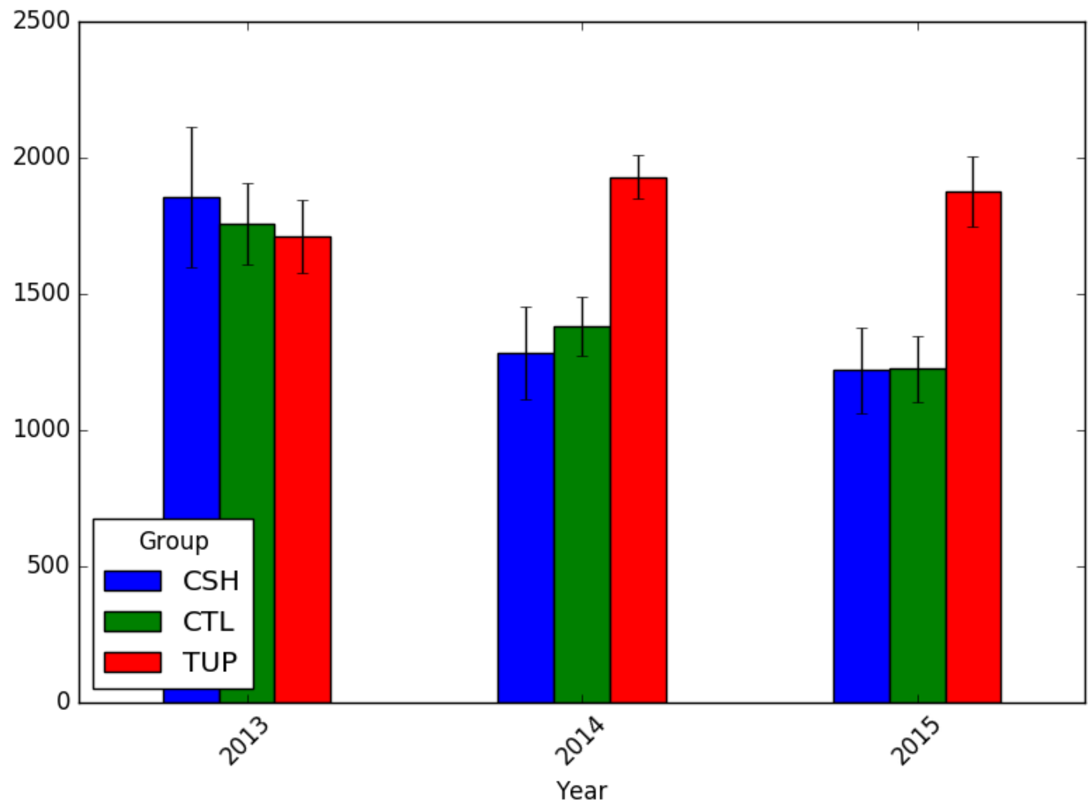

Figure 3.1: Measured asset wealth by group-year

[3]For now, we include in this list: small and large livestock, farm equipment, mobiles, carts, sewing equipment, sheds, and shop premises.

Table 3.8: Average treatment effects by group-year on total value (in SSP) of all assets measured and of productive assets measured

	Total	Productive
CTL mean	1225.61	337.60
	(1502.46)	(605.57)
TUP*2014	535.79***	361.80***
	(154.02)	(74.19)
TUP*2015	624.79***	320.74***
	(146.01)	(68.68)
CSH*2014	−125.86	18.50
	(191.31)	(95.80)
CSH*2015	−49.99	−5.00
	(187.32)	(88.40)
Bsln2013	0.08***	0.00
	(0.02)	(0.01)
2014	1259.75***	465.53***
	(112.68)	(55.96)
2015	1124.61***	392.97***
	(103.46)	(50.21)
Bsln NA	21.30	−131.14**
	(146.51)	(51.35)
N	1305.00	1247.00
F-stat	8.53	10.19
$\beta^{TUP}_{2014} - \beta^{CSH}$	585.78**	366.79***
	(239.76)	(114.58)
$\beta^{TUP}_{2015} - \beta^{CSH}$	674.78***	325.74***
	(194.72)	(92.26)

Savings

Both treatment arms had significant impact on the average value of cash savings within households in 2015. The TUP households are strongly encouraged to pay into a savings account maintained by BRAC each time they meet. Anecdotally, this has discouraged some women from attending the meetings, but it results in TUP participants being 44% (20 pp) more likely to report having any savings at all. It's worth noting though that since the TUP households also regard their savings behavior as much more transparent to BRAC (and have received pressure to save from them) than the other groups, these households may simply be more likely to reveal that they are saving when asked. Among those who have savings, TUP households report having roughly 43% (81 SSP) more in value.

Cash households appear no more likely than the control households to report having cash savings (around 45% in each group), but households that report having any savings at all report having 47% (91.4 SSP) more in value. This is significantly less than was given to these households, but combined with the short-term consumption results, goes some distance in explaining the lack of effect on physical asset wealth. The cash seems to have gone primarily to consumption and savings.

It is common in this community (and most in the region) to store non-perishable food like maize, cassava, or millet as a form of savings. This would seem particularly reasonable in a high-inflation context, where the price of grain had doubled in the previous year. At least as many households report saving in food (53%) as in cash (46%), with an average market value of 106 SSP. However, we find no evidence that either treatment group increased food savings.[4]

Neither do we find evidence that either treatment increased the size or likelihood of giving or receiving interhousehold transfers, either in cash or in kind. These results are omitted since only 35 and 60 households reported giving and receiving transfers respectively, with no difference in group means.

Table 3.9: Average treatment effects by group-year on percentage of households reporting any savings or land access

% > 0	Savings	Food Sav	LandCult	LandOwn
CTL mean	0.45	0.82	0.82	0.90
CSH*2014	−0.06	0.00	−0.04	−0.01
	(0.06)	(0.04)	(0.04)	(0.04)
CSH*2015	0.03	0.02	0.05	0.02
	(0.05)	(0.04)	(0.04)	(0.04)
TUP*2014	0.22***	−0.02	−0.03	−0.00
	(0.04)	(0.03)	(0.03)	(0.03)
TUP*2015	0.21***	−0.03	0.01	−0.01
	(0.04)	(0.03)	(0.03)	(0.03)
2014	0.43***	1.00***	0.83***	0.82***
	(0.04)	(0.02)	(0.06)	(0.05)
2015	0.39***	0.82***	0.77***	0.84***
	(0.04)	(0.02)	(0.05)	(0.05)
Bsln2013	0.05		0.05	0.07
	(0.04)		(0.05)	(0.04)
Bsln NA	0.08*		0.05	0.05
	(0.04)		(0.06)	(0.05)
$\beta_{2014}^{TUP} - \beta^{CSH}$	0.19	−0.04	−0.07	−0.02
$\beta_{2015}^{TUP} - \beta^{CSH}$	0.18	−0.05	−0.03	−0.03
F-stat	8.83	15.60	0.79	0.76
N	1259.00	870.00	1231.00	1251.00

[4]Note that food savings was not measured at baseline, so these controls are omitted.

Table 3.10: Average treatment effects by group-year on total value (in SSP) of all cash and food savings and area (in fedan) of land being cultivated by the household (including rented or temporary-use) and owned by the household.

Amt.	Savings	Food Sav	LandCult	LandOwn
CTL mean	191.19	114.78	61.88	46.00
CSH*2014	28.74	0.22	10.18	10.50
	(42.93)	(15.38)	(15.07)	(12.57)
CSH*2015	91.40**	−14.34	−39.18***	−32.37***
	(40.89)	(14.98)	(14.90)	(11.95)
TUP*2014	−27.09	17.16	−4.76	−3.02
	(29.76)	(12.33)	(11.94)	(10.04)
TUP*2015	81.33***	1.13	−17.38	−12.56
	(29.32)	(12.26)	(11.65)	(9.41)
2014	106.72***	62.03***	11.37	17.31**
	(24.85)	(8.36)	(9.94)	(8.56)
2015	163.04***	114.78***	61.52***	51.89***
	(24.13)	(7.60)	(9.54)	(7.88)
Bsln2013	0.05**		0.94	−2.43
	(0.02)		(3.07)	(1.95)
Bsln NA	40.07*		−1.60	−6.02
	(21.24)		(9.92)	(8.29)
$\beta^{TUP}_{2014} - \beta^{CSH}$	−118.49	31.50	34.42	29.35
$\beta^{TUP}_{2015} - \beta^{CSH}$	−10.07	15.47	21.79	19.80
F-stat	7.41	7.14	4.91	3.72
N	671.00	777.00	1042.00	1114.00

Land Holdings

We also examine land ownership and cultivation in each year. We find no evidence that either group is more or less likely to report owning or cultivating at least some land, though this may be in part because land ownership and cultivation is already very common. Anecdotally, divesting from land ownership entirely could be seen as a relatively drastic decision. However, members of the cash group who are involved in agriculture are found to be cultivating significantly less land after the fact, which reports cultivating 65% less and owning 70% less land than the control group. This raises the interesting question of whether the cash group was likely to switch occupations from farming to non-farm self-employment. It could also raise questions around the underlying logic of the more agrarian transfer in the TUP program, if unconstrained transfers prompt households to divest participants primarily stated a preference for small retail training and transfers over small animal husbandry or vegetable gardening.

Income

Income was reliably measured only in 2015, and so our estimates do not control for baseline values. The control group in 2015 has a measured income of roughly 4325 SSP per year, or roughly \$540 US (assuming an exchange rate of around 8). The TUP group sees a 327 SSP (\$41 US, 7%) increase in annual average income, but with a fairly skewed distribution and high standard errors. The related figure shows that total income is not particularly different among groups. Perhaps the main lesson is that the TUP group has measurably more reported livestock-related income, and less farm income, indicating a shift away from farming. The cash group may exhibit some substitution away from farm and livestock, but as is evident graphically, we do not observe sizable changes in total income for either treatment group.

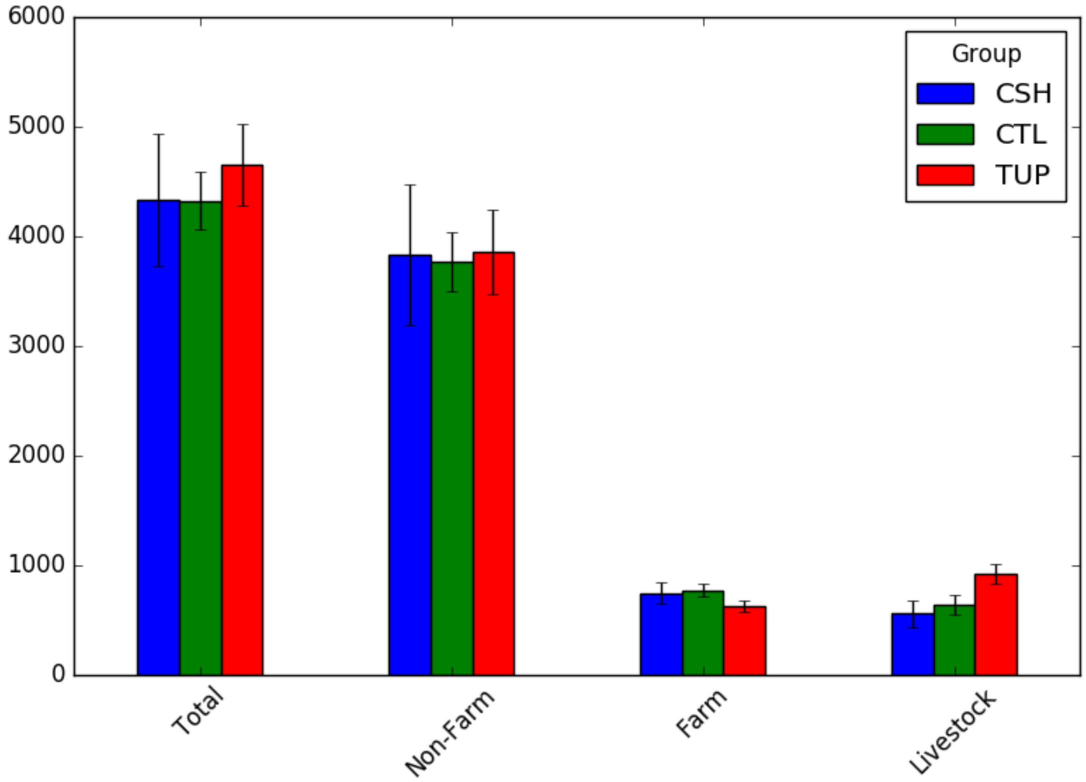

Figure 3.2: Distribution of total observed income by group

Table 3.11: Average treatment effects by group-year on total value (in SSP) of income reported in 2015 by sector.

	Farm	Livestock	Non-Farm	Total
CTL mean	773.05	640.33	3774.49	4325.54
TUP	-142.20^{*}	281.12^{**}	86.24	327.83
	(77.21)	(126.30)	(469.48)	(455.95)
CSH	-26.15	-83.81	61.80	7.92
	(100.82)	(177.25)	(620.53)	(600.43)
N	531.00	380.00	606.00	671.00
F-stat	1.75	3.48	0.02	0.28
$\beta^{TUP} - \beta^{CSH}$	-116.05	364.94^{**}	24.44	319.91
	(105.79)	(174.74)	(651.27)	(629.93)

Exposure to Conflict

In 2014, households were surveyed shortly after the NGO's offices had re-opened in the wake of the outbreak of widespread armed conflict. Respondents were asked a short set of questions about whether they were directly affected, and if so, in what way. There has only been a few incidents of violence near Yei town at that point, and the most directly involved ethnic groups made up a small portion of the local population. There is no clear comparison group to which we might compare our sample, and the economic climate changed over this same period in several ways that were probably not directly caused by the violence. As such, we have no clear means of identifying the effect of the conflict itself on household welfare. Nonetheless, it is interesting to consider correlates with self-reported exposure to the conflict, and to see if program assignment had any effect on households' exposure or response.

Our main outcomes of interest are whether individuals say they were "worried" or "directly affected" by the violence, unable to invest in a farm or business as a result, migrated as a cautionary measure, or did something else to protect the lives of family members. A final question among those who took no cautionary measures was whether this because they did not have the means (i.e. "NoMeans"). TUP participants are 24% (13 pp.) less likely to report having been "affected" by the conflict, and 38% (6 pp.) less likely to report that they were affected specifically by being unable to plant crops or invest in their business. This was the second most common way in which households reported being affected behind "needed to relocate or migrate", where respondents are not clearly different. Nonetheless, this raises the possibility that having received a significant asset transfer and the expectation of NGO support around the outbreak of conflict may have helped mitigate the conflict's negative effect on investment and protect households from being affected overall.

Table 3.12: Average treatment effects by group-year on the probability of having been affected in a significant way by the outbreak of violence in late 2013

	Affected	Migrated	NoInvest	NoMeans	ProtectLives	Worried
CTL mean	0.53***	0.33***	0.16***	0.33***	0.38***	0.93***
	(0.03)	(0.02)	(0.02)	(0.02)	(0.03)	(0.01)
TUP	−0.13***	0.04	−0.06**	−0.06	0.02	−0.02
	(0.04)	(0.04)	(0.03)	(0.04)	(0.05)	(0.02)
F-stat	9.20	0.96	3.95	2.55	0.19	0.49
N	601.00	655.00	655.00	655.00	585.00	603.00

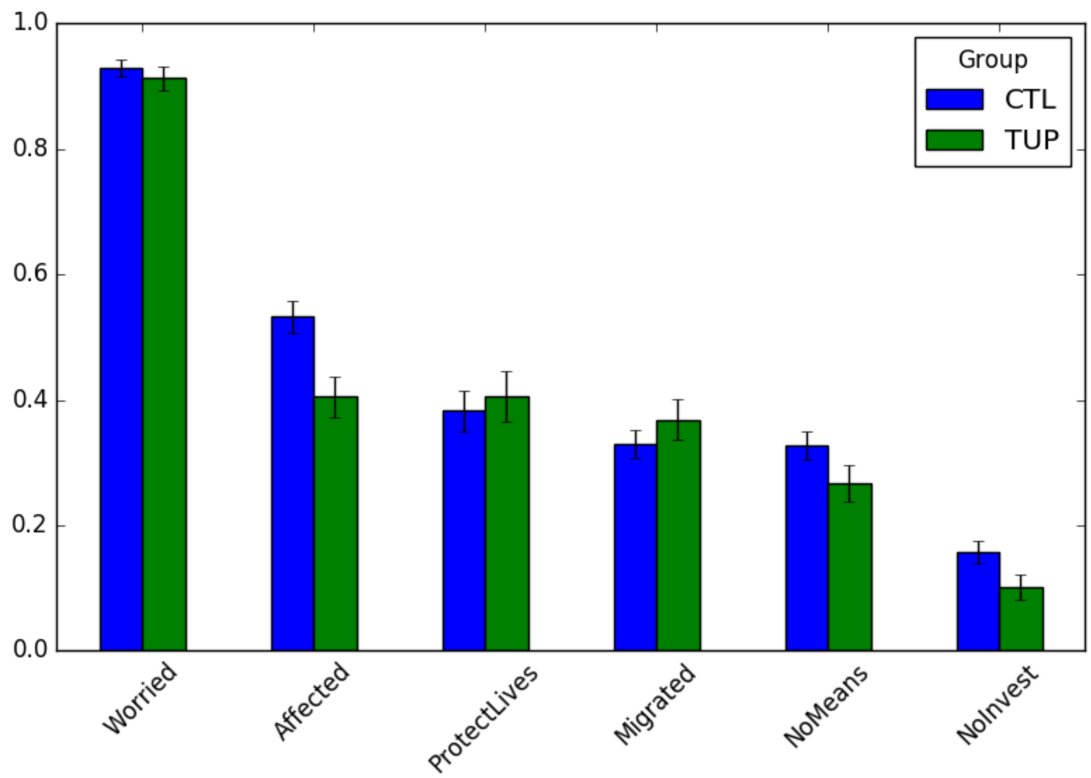

Figure 3.3: % of Sample reporting exposure to conflict by group.

Concluding Remarks

BRAC's South Sudan pilot of the TUP program represents the only such test of the ultra-poor graduation framework conducted in an area of significant political and economic instability.

It also represents among the only direct comparisons of this model to a similarly expensive unconditional cash transfer, arguably its most sensible benchmark for success. As such, it provides suggestive evidence as to the best way of transferring wealth in order to help poor and vulnerable households.

Cash transfers appear to increase consumption and possibly shift investment from agriculture to non-farm activities, without a related increase in wealth or income. Conversely, the TUP program increased wealth and directly shifted work from agriculture to livestock, with increased consumption in the short run. We also find that having received asset transfers dampened the negative investment effects following the outbreak of violence.[5] We tentatively conclude that targeted asset transfers can play a constructive role in helping poor, self-employed households when they face economic uncertainty. And while cash increases household consumption, the goal of improving income or wealth is aided by the additional services that the ultra-poor graduation framework offer.

[5]Whether a cash transfer would have had a similar mitigating effect is hard to say.

Bibliography

Bandiera, Oriana et al. (2012). "Can Basic Entrepreneurship transform the economic lives of the poor?" Unpublished Ms.
— (2017). "Labor markets and poverty in village economies". In: *The Quarterly Journal of Economics* 132.2, pp. 811–870.
Banerjee, Abhijit and Esther Duflo (2007). "The Economic Lives of the Poor". In: *Journal of Economic Perspectives* 21.1, pp. 141–168. DOI: 10.1257/jep.21.1.141.
Banerjee, Abhijit, Esther Duflo, et al. (2015). "A multifaceted program causes lasting progress for the very poor: Evidence from six countries". In: *Science* 348.6236, p. 1260799.
Banerjee, Abhijit and Andrew F. Newman (1993). "Occupational Choices and the Process of Development". In: *Journal of Political Economy* 101.2, pp. 274–298.
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.
Blattman, Christopher, Julian Jamison, et al. (2014). "The Returns to Cash and Microenterprise Support Among the Ultra-Poor: A Field Experiment". In: *Available at SSRN*. URL: http://ssrn.com/abstract=2439488.
Browning, Martin, Angus Deaton, and Margaret Irish (1985). "A Profitable Approach to Labor Supply and Commodity Demands over the Life Cycle". In: *Econometrica* 53, pp. 503–543.
Chowdhury, Reajul, Elliott Collins, et al. (2015). "Effects of Transfers to the Ultra-Poor on Consumption, Productivity, Assets and Employment in South Sudan". Unpublished manuscript.
Chowdhury, Reajul and Ricardo Morel (2014). "Measuring targeting effectiveness of BRAC's 'Targeting the Ultra-Poor' program in a different context: Evidence from South Sudan". Unpublished manuscript.
Collins, Elliott and Ethan Ligon (2017). "Asset transfers and household neediness". CUDARE Working Paper. URL: http://escholarship.org/uc/item/24t720fk.
De Mel, Suresh, David McKenzie, and Christopher Woodruff (2008). "Returns to capital in microenterprises: evidence from a field experiment". In: *The Quarterly Journal of Economics*, pp. 1329–1372.
Fafchamps, Marcel et al. (2011). *When is capital enough to get female microenterprises growing? Evidence from a randomized experiment in Ghana*. Tech. rep. National Bureau of Economic Research.

Gorman, W. M. (1959). "Separable Utility and Aggregation". In: *Econometrica* 27, pp. 469–81.

Haushofer, Johannes and Jeremy Shapiro (2013). "Policy Brief: Impacts of Unconditional Cash Transfers". Unpublished manuscript. Massachusetts Institute of Technology.

Heckman, James J and Thomas E MaCurdy (1980). "A life cycle model of female labour supply". In: *The Review of Economic Studies*, pp. 47–74.

Karaivanov, Alexander and Robert M. Townsend (2014). "Dynamic Financial Constraints: Distinguishing Mechanism Design from Exogenously Incomplete Regimes". In: *Econometrica* 82.3, pp. 887–959.

Ligon, Ethan (2015). "Tracking Household Neediness Over Time, With an Application to Uganda". Unpublished manuscript.

— (2016a). "All λ-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.

— (2017). *`CFEDemands`: Tools for estimating and computing Constant Frisch Elasticity (CFE) demands.* URL: `https://pypi.python.org/pypi/CFEDemands/0.1.4.2dev`.

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`.

MaCurdy, Thomas E. (1983). "A Simple Scheme for Estimating an Intertemporal Model of Labor Supply and Consumption in the Presence of Taxes and Uncertainty". In: *International Economic Review* 24.2, pp. 265–289.

McKenzie, David (2012). "Beyond baseline and follow-up: The case for more T in experiments". In: *Journal of Development Economics* 99.2, pp. 210–221. URL: `http://ideas.repec.org/a/eee/deveco/v99y2012i2p210-221.html`.

McKenzie, David and Christopher Woodruff (2008). "Experimental evidence on returns to capital and access to finance in Mexico". In: *The World Bank Economic Review* 22.3, pp. 457–482.

Paulson, Anna L., Robert M. Townsend, and Alexander Karaivanov (2006). "Distinguishing limited liability from moral hazard in a model of entrepreneurship". In: *Journal of political Economy* 114.1, pp. 100–144.

Rubin, Donald B (1974). "Estimating causal effects of treatments in randomized and nonrandomized studies." In: *Journal of Educational Psychology* 66.5, p. 688.

Singh, I., L. Squire, and J. Strauss (1986). *Agricultural Household Models: Extensions, Applications and Policy.* Baltimore: John Hopkins University Press.