Coping with the Coffee Crisis in Central America:
The Role of Social Safety Nets in Honduras

David Coady*, Pedro Olinto*, and Natalia Caldes+

* International Food Policy Research Institute, Washington DC
+ Departamento de Economía Agraria - ETSI Agrónomos - UPM

1. Introduction

In spite of some recovery in 1994 and 1998, world coffee prices have declined dramatically since the mid 1980s. Real coffee prices are now at their lowest levels in more than 50 years. Such a substantial decrease in prices has obviously had enormous adverse implications for the incomes of many of the coffee producing countries in Central America and for coffee producers in particular. However, in spite of the widespread perception of rising poverty, little rigorous empirical evidence exists regarding the magnitude and nature of the poverty impact of the crisis. Even less evidence exists with regard to the potential role for social safety net programs both in protecting poor households from such shocks as well as facilitating more efficient responses.[1] This paper contributes to filling some of these knowledge gaps.

Honduras is one of the poorest countries in Latin America with around 70% of its population classified as poor. Unlike other countries in the region, poverty is predominantly a rural phenomenon, with 49% of the total population in 1998 (around 6.5 million) living in rural areas where poverty and, in particular, extreme poverty, are substantially higher (World Bank, 2001; Morris et al, 2002). Also, the extent of poverty in rural areas has been increasing in recent years and much of this is often attributed to the continued decline in coffee prices over this period.

The analysis in this paper uses household-level survey data collected as part of an evaluation of a recently introduced social protection program in Honduras (called PRAF). These data provide a very rich source of information for the purpose of addressing the above empirical issues. We intend to use these data to evaluate the impact of the ongoing crisis on some of the poorest rural regions and households in Honduras. In particular, we will evaluate the effectiveness of PRAF in

[1] See Morris et al (2002) for a discussion of the impact of Hurricane Mitch in late October 1998 on some of the poorest municipalities in Honduras, in particular the discussion on the extend and distribution of foreign aid.

enabling poor households to protect their welfare against such shocks. The analysis may also help to improve our understanding of the role of social safety net programs more generally in protecting against economic shocks as well as of the nature of risk coping strategies available to poor households.

In Section 2 we provide a brief overview of coffee in the context of the Honduran economy. In Section 3 we describe the program and survey design. Section 4 provides some descriptive analysis of the data, with special emphasis on the extent of involvement in coffee in the program area as well as differences in socio-economic outcomes across coffee/non-coffee and program/non-program households. Section 5 evaluates the impact of the program on household consumption. Section 6 provides some concluding remarks.

2. Coffee in Honduras

Coffee is extremely important to Honduras in terms of output, employment, and export earnings (McCarty and Sun, 2003; Partners, 2002; Varangis, 2003). In terms of output, in 2002/2003 the area under coffee cultivation reached 233,750 hectares, accounting for 65% of total permanent crop area.[2] A large amount of the coffee is also grown on small farms in isolated, high altitude areas; for 92% of growers annual coffee production is less than 100 quintales per farm, most of which is sold to the market through intermediaries. Coffee is also seen as a major source of employment, employing up to one third of the labor force in rural areas (Hearne et al, 2002). Coffee is the second most important export crop in Honduras in value terms and accounts for almost 25% of national GDP. In 2000 the value of coffee exports was $345.2 million, which represents over 26% of total export revenue for that year.

During 2000 and 2001, continued increases in world supply led to some of the lowest real coffee prices over the last few decades. Since mid 2001, both the nominal and real prices of Arabica coffee, both domestically and on international markets (Figure 1), have been declining. Yet, in spite of this, coffee production in Honduras (as elsewhere) has been increasing over this period.

[2] The main coffee variety produced in Honduras is the Arabica coffee variety. In recent years there has been concerns about the quality of the coffee being produced and its economic consequences, with Honduran coffee being subject to penalties of between 5-20 cents per pound on world markets over the past few years (Herne et al, 2002).

Figure 1: World and Honduras Domestic Coffee Prices and Production

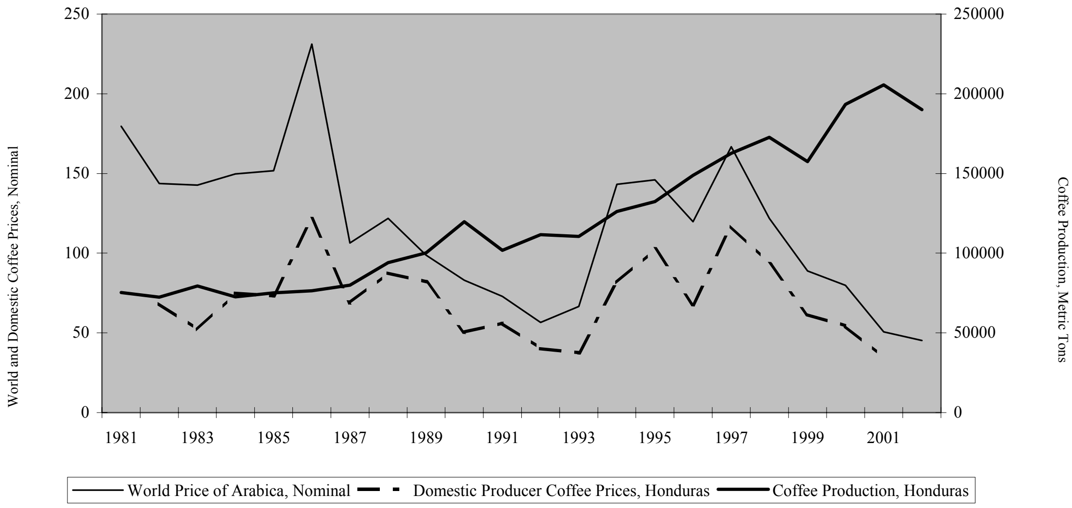

Source: International Coffee Organization, FAO

2. The Program and Survey Design

In late 1998, the Government of Honduras initiated the PRAF/IDB Phase II Project[3], which is one of the largest social investments in the history of Honduras. The program covers households in 70 (out of 297) of the most disadvantaged municipalities in 7 departments (Copan, Intibuca, Ocotepeque, F. Morazan, La Paz, Sta. Barbara and Lempira), these all being located in the western part of Honduras and covering many of the major coffee-growing departments. Program transfers are roughly 3% of total household consumption on average, but are substantially higher for the poorest households.

For the purpose of selecting participating households, the government first identified the poorest 70 municipalities based on the average height-for-age of children in first grade.[4] These municipalities were then categorized into five strata based on the same variable and, within each stratum, municipalities were randomly allocated to four program evaluation groups (i.e. demand- and supply-side transfers – 20 municipalities, demand-side transfers only – 20 municipalities,

[3] PRAF is an acronym for Programa de Asignacion Familiar (or Family Allowance Program).
[4] See the appendix for a more detailed discussion of the program and survey design.

supply-side transfers only – 10 municipalities, and 20 control municipalities). All households with children less than 13 years or with a pregnant woman were considered eligible for the program.

As part of the evaluation, a randomly selected sample of households was surveyed before and after the program intervention. This survey contains information on 5,600 households from the 70 municipalities covered by the program evaluation (i.e. 80 households per municipality). Households from 20 of these municipalities received cash transfers conditioned on attendance at school and health clinics, 10 received only increased investments in schools and health clinics, 20 received household transfers and supply-side investments, and 20 did not receive any component of the program and were used as a control for the purposes of program evaluation. The surveys were implemented between August-December 2000 (baseline) and June-September 2002 (follow-up). The surveys contain sixteen modules (including information on household conditions, household composition, remittances, education, expenditures, maternal/child health, anthropometrics, women/child time use, quality of schools/clinics, other programs, goods/animals, community survey) containing over 500 variables. In this preliminary version of the paper we will focus solely on one outcome variable, namely total household consumption. Other outcomes will be analyzed in a future revised version of the paper.

Two features of the data collection process have important implications for the way in which we analyze the impact of the program on consumption across coffee and non-coffee households. Firstly, the baseline and follow-up surveys were collected at different times of the year, with the 2000 baseline being collected over the months August to December and the 2002 follow-up survey collected over the months June to September. Secondly, for operational reasons, the baseline survey was implemented first in the treatment municipalities (i.e. in those where households were later to receive cash transfers). Both these timing discrepancies mean that great care needs to be taken when comparing differences in consumption (and other) outcomes across program treatment/control groups as well as over time. We will address such issues in more detail below.

3. Data Description

For the purpose of describing the data we start by classifying households into treatment/control as well as coffee/non-coffee households. In the dataset, we have 5,087 households that exist in both the baseline and follow-up surveys. The classification of households into treatment and control is determined by the sampling strategy described above. The survey includes a number of different variables we could use to classify households into coffee/non-coffee households. In the survey we have identified three sets of questions:

(i) A question that asks if a household has land and, if so, how much of this is devoted to coffee. This question is asked only in the 2002 follow-up survey and is in a module relating to landholdings and their allocation in 2001.

(ii) A question in each survey asking how each working individual's time was allocated between various activities over the seven days prior to the survey, including coffee-related activities mainly referring to coffee-harvesting activities.

(iii) A question on coffee stocks held as an asset in 2002.

We do not use the last two questions to classify households as coffee/non-coffee, although we will describe the responses below.

With respect to the second question relating to involvement in coffee activities, since the coffee harvesting season covers the months September-November, in the follow-up survey over June-September households are much less likely to report involvement in coffee activities since it is prior to the harvest. This is not the case in the baseline survey over August-December, which covers the harvest season. Therefore, one cannot gain any insights, other than possibly about seasonal patterns, from comparing the extent of households labor involvement in coffee-related activities. The different seasonal timing of the surveys means that one expects labor activity to be greater in the baseline survey. With respect to the third question relating to coffee stocks, very few households report such stocks so we do not analyze responses to this question in this paper.

3.1 The Extent of Coffee Involvement

We start by examining the pattern of coffee involvement across treatment and control areas for the 5,087 households that can be matched over both the baseline and follow-up surveys. For the purposes of impact evaluation households were allocated into four groups according to whether they received: demand-side transfers only (demand), supply-side transfers only (supply), both demand- and supply-side transfers (dem+sup), nothing (control). The distribution of surveyed households over these groups was: demand 29%, supply 14%, dem+sup 28%, and control 28%. For the purposes of discussion, and given our current focus on the role of cash transfers, we will refer to those areas where households receive demand-side cash transfers simply as "treatment" areas and those not receiving such transfers as controls. According to this definition, then, 57% of households are located in treatment areas and the remaining 43% in control areas.

As already indicated, for operational reasons (i.e. the desire to get transfers out to beneficiaries as soon as possible) the baseline survey was applied first to the treatment areas (i.e. from August 1st, 2000, to October 15th, 2000) and then to the control areas (i.e. from October 16th, 2000, to November 30th, 2000). This systematic variation in the timing of the survey means that one needs to be extremely wary when comparing patterns across these groups. For example, one expects coffee-related labor activities to be higher during the coffee harvesting period (i.e. September to November) and therefore higher in the control groups.

Table 1 presents survey information regarding coffee-related activities across program groups. We can see that the areas covered by the program are very coffee-intensive areas. Overall, 76% of households report having some land and 32% report having some land devoted to coffee production. In 2000, around 17% of households report having at least one household member involved in coffee activities over the seven days preceding the survey. Reflecting the timing differences in the survey implementation across treatment and control groups in 2000, we also observe a greater amount of coffee-related labor activity in the control group. The smaller extent of coffee-related activities in the 2002 follow-up survey also reflects differences in timing of survey implementation since this survey was implemented before the coffee harvesting season.

Table 1: Coffee intensity by program group

Program group	Have land	Active in coffee 2000	Active in coffee 2002	Land in coffee	In coffee cluster
Demand	0.77	0.13	0.07	0.30	0.68
Demand+suply	0.77	0.14	0.08	0.32	0.76
Supply	0.75	0.23	0.10	0.38	0.86
Control	0.75	0.20	0.06	0.31	0.80
Total	0.76	0.17	0.08	0.32	0.76

In the empirical analysis below we use two definitions of a "coffee household". In the first, we simply use a binary variable that indicates whether or not a household has some land devoted to coffee. Around 32% report having some land devoted to coffee. In the second, we use a binary variable that indicates whether or not a household lives in a coffee cluster, i.e. in a cluster where at least one household reports having some land devoted to coffee production. This is obviously a much broader definition of a coffee household and 76% of households are classified as coffee households under this definition. To the extent that labor hiring for coffee activities is local, one expects the additional households incorporated under the second definition to include coffee laborers. Below we will refer to these distinct groups "coffee land", "other coffee", and "non-coffee" households.

3.2 Consumption Patterns

Table 2 presents data on the pattern of total household per capita consumption across coffee and non-coffee households and over time.[5] Overall, we observe a 2% decline in consumption levels over the two years. It is also noteworthy that households with coffee land experience a substantially larger 5% decline compared to the other two groups, which both experience a 1% decline. The fact that "other coffee" households have a much smaller decline also suggests that households not directly involved in coffee production can mitigate the welfare impacts of a worsening coffee environment, e.g. by switching to other non-coffee activities. Note also that approximately as many households experience increases as decreases highlighting the diversity of impacts across households.

Table 2: Consumption by coffee grouping

	Per capita consum 00	Per capita consum 02	Change in pc consum	Percentage change pccon	% with decrease in consumption
Coffee Land	22.19	21.53	-1.08	-0.05	0.53
Other Coffee	19.27	19.46	0.10	-0.01	0.51
Non Coffee	17.92	18.43	0.55	-0.01	0.49
Total	19.87	19.87	-0.17	-0.02	0.51

[5] Total household per capita consumption is in nominal terms (i.e. not deflated by a consumer price index). Deflating would lead to the decrease over 2000-2002 being larger. Note that although one cannot sensibly compare changes across program groups because of the different survey timing, this is not the case for comparisons across coffee/non-coffee groups since these are distributed uniformly across survey periods. However, the timing differences between the survey rounds (as well as within the baseline) means that the time pattern contains an element of seasonality. For this reason we focus more on relative changes across groups.

3.3 Consumption Groups

Table 3 presents the pattern of consumption changes and coffee involvement across consumption quintiles. The first column of numbers presents the ratio of quintile mean income to that for the lowest quintile. The steeply increasing ratio tells us that even in these poor municipalities there is substantial inequality with, for example, per capita consumption for the middle quintile being 2.3 times that of the lowest quintile. The pattern of consumption changes is also very different across consumption groups with the lowest quintiles experiencing large increases and the highest experiencing large decreases. This indicates that inequality in consumption would appear to be decreasing substantially over time, albeit reflecting households becoming more equally poor! The extent of coffee involvement, as captured by being classified as a "coffee land" household, also differs substantially across consumption quintiles with only 25% of households classified as coffee households in the lowest quintile and this increasing substantially to 71% in the highest quintile. Therefore, one expects that the relatively large decreases in consumption in the highest quintiles partly reflect their greater exposure to some coffee-related shock against which it is difficult to protect, e.g. an inability to switch coffee land quickly out of coffee in response to trend price decreases. The broader definition of coffee involvement exhibits a more similar involvement across quintiles consistent with the involvement of the poorest quintiles in coffee reflecting their role as coffee laborers.

Table 3: Consumption and coffee by consumption quintiles

Quintiles	Consumption ratio	% change in in pc consum.	% with decrease in consumption	Coffee land	Coffee cluster
Low	1.00	0.35	0.25	0.25	0.74
2	1.62	0.10	0.45	0.34	0.76
3	2.28	-0.05	0.55	0.29	0.76
4	3.34	-0.18	0.61	0.33	0.75
High	8.43	-0.33	0.71	0.38	0.80

3.4 Consumption Regressions

Above we presented simple mean comparisons across coffee and non-coffee groups. However, comparing means and establishing whether any differences are statistically significant requires that we adjust for the sampling design of the survey. Therefore, in this section we present the results of regression analyses, which make such corrections.[6] When we regress the change in the log of per capita consumption simply on a dummy for having land devoted to coffee we find a coefficient of -0.041, significant at the 5% level (t=1.98). Table 4 presents the results from a similar regression of the change in the log of per capita consumption on dummy variables identifying whether or not a household is involved in coffee. However, now all the regressions also include controls for household characteristics (i.e. land size, and the education, age and gender of the head of household; all at their initial 2000 values) as well as departmental dummies (coefficients not reported). Controlling for such variables means that the coefficient on coffee is more likely to capture a "coffee impact" as opposed to capturing the impact of other important characteristics correlated with coffee involvement.

The first column of results presents the estimates using a "coffee land" dummy; the coefficient shows the percentage point difference in the decrease in per capita consumption for coffee-land households. We see that the decrease in consumption is 3.8 percentage points greater for households with coffee land, and this difference is significant at the 5% level. Focusing on the second column, we do not find any significant difference between coffee and non-coffee clusters. The results in column three directly mirror the comparisons presented earlier in Table 2. But although the magnitudes and their pattern across the groups are similar, the differences are not statistically significant.

[6] We correct for the sampling design using the svyreg command in stata specifying municipalities as the primary sampling units and stratification based on five nutrition municipality clusters.

Table 4. Consumption Differences Across Coffee and Non-Coffee Households

	Dependent variable: *Difference in the log of the per capita nominal expenditures (Lmps)*		
Constant	-0.057 (0.82)	-0.041 (0.55)	-0.051 (0.69)
Coffee Land Dummy	-0.038 (2.08)*		-0.042 (1.28)
Coffee Cluster Dummy		-0.018 (0.62)	
Other Coffee Dummy			-0.006 (0.21)
Land size	0.002 (1.38)	0.002 (1.23)	0.002 (1.38)
Land size squared	-0.000 (1.04)	-0.000 (0.88)	-0.000 (1.04)
Sex head	-0.071 (2.69)**	-0.074 (2.83)**	-0.071 (2.68)**
Age head	0.010 (3.38)**	0.010 (3.25)**	0.010 (3.37)**
Age head squared	-0.000 (2.75)**	-0.000 (2.65)*	-0.000 (2.74)**
Education head	-0.001 (0.04)	-0.000 (0.03)	-0.001 (0.04)
Education head squared	0.001 (0.42)	0.001 (0.41)	0.001 (0.42)
Observations	4989	4989	4989
R-squared	0.06	0.06	0.06

Note: Absolute value of t statistics in parentheses: * significant at 5%; ** significant at 1%. All regressions include municipality dummies (not reported).

4. Program Impact

In this section we evaluate the impact of the PRAF transfer program on household consumption as well as trying to identify any differential impact across coffee and non-coffee households. For this purpose we take two approaches. First, because of the systematic differences in survey timing across program and non-program communities, we use only the follow-up survey in 2002. In other words, we focus on the "difference estimator". Second, we estimate a "double difference" impact based on both rounds of data. This approach enables us to address the issue of the channels through which the program has any impact. For example, if the cash transfer has an effect through relaxing liquidity constraints then it can be expected to increase the returns to existing assets (e.g. the returns to land). In the absence of such effects the transfer has only a pure income effect.

4.1 Difference Estimates

We first use the 2002 follow-up survey to get a "difference" estimate of program impact. The model we use is as follows:

$$Y = \alpha + \beta.P + \gamma.(PxC) + \Gamma.X + \epsilon$$

where Y is the log of total per capita household consumption, P is a binary variable taking the value unity if the household is "in the program" (zero otherwise), C is a binary variable taking the value unity if the household is classified as a "coffee household" (zero otherwise), X is a vector of other household characteristics (including location characteristics), and ϵ is a random disturbance term with mean zero. The coefficient β then gives an estimate of the impact of the program on non-coffee households (the base category) and $\beta+\gamma$ gives an estimate of the impact on coffee households. Therefore, γ is an estimate of the differential impact of the program on coffee households. We also try alternative specifications, for example, by dropping the interaction term (PxC), in which case β is the common impact on all program households, or by using different combinations of X variables.

We also estimate three different "program impacts". Firstly, we estimate the impact of the program on households in program municipalities, i.e. both eligible and non-eligible households. In this case we use the sample of all households and the program dummy represents whether a household resides in a program municipality. In the sample, 57% of households reside in program municipalities. Second, since only households with children aged less than 13 years or with a pregnant woman are eligible, we restrict our sample to such households. The program dummy represents whether or not an eligible household resides in a treatment municipality. This program impact is commonly referred to as the impact of the "offer to treat". In the sample, 88% of households are eligible for the program and 58% of these are in treatment municipalities. Third, since some households may decide not to take-up their transfers, e.g. because they do not want to meet the program conditions, or some households may not receive transfers for operational reasons, we also estimate the impact of actually receiving transfers. In this case the sample is also made up of only eligible households and the program dummy represents whether or not a household received a transfer. This is commonly referred to as the "treatment effect". In the sample, 58% of eligible households report receiving a transfer. Since program take-up is potentially endogenous, we instrument the program dummy using the dummy representing whether a household is eligible. We expect this third estimate to be higher than the second, which is expected to be higher than the first. In other words, we expect the differences to be highest for those who receive transfers.

Table 5 presents our three sets of estimates of program impact on per capita household consumption. Although we tried numerous specifications, the only ones that generated statistically significant program impacts were those which included only the program dummy, its interaction with the coffee dummy, and departmental dummies. All other specifications generated coefficients of similar magnitude but all were statistically insignificant. We therefore only present the former sets of regression estimates. For each of the three program impacts identified above, we present two different estimates, one that assumes a common program impact across all households in the sample, and a second which tests for differential impacts across coffee and non-coffee households.

Table 5. Program Impact on Coffee and Non-Coffee Households

	Dependent variable: *Log of the per capita consumption in 2002*					
	Program community		Household eligible		Received transfer	
Constant	2.666 (48.24)**	2.625 (44.60)**	2.614 (47.86)**	2.575 (42.24)**	2.575 (42.24)**	2.575 (42.02)**
Program Dummy	0.075 (1.85)	0.095 (2.04)*	0.072 (1.74)	0.087 (1.76)	0.093 (1.74)	0.114 (1.76)
Coffee Land Dummy		0.125 (2.86)**		0.119 (2.57)*		0.119 (2.56)*
Program*Coffee Interaction		-0.055 (0.89)		-0.040 (0.62)		-0.054 (-0.66)
Observations	5087	5087	4464	4464	4464	4464
R-squared	0.01	0.02	0.01	0.02	0.02	0.002

Note: Absolute value of t statistics in parentheses; * significant at 5%; ** significant at 1%. All regressions include department dummies (not reported).

The first two columns of numbers present the program impact on households in program communities. The estimate, which is significant at the 5% level, indicates that the program increases household consumption by 7.5 percent. However, the results in column two indicate that the program impact does not differ across coffee and non-coffee households. So coffee and non-coffee households in program municipalities appear to benefit equally from the program. But coffee households have 12.5% higher consumption levels compared to non-coffee households.

The second pair of columns presents our estimates of the impact of the program on eligible households. The estimates in column three indicate that the program increases the consumption of eligible households by 7.2 percent. But, again, column four indicates that the program impact appear to be equal across coffee and non-coffee households. Similar to above, coffee households have a 11.9% higher consumption level relative to non-coffee households.

Finally, the final two columns present the impact of receipt of transfers. The results in column five indicate that the receipt of transfers increases household consumption by 9.3%. But, again, this impact does not appear to be significantly different across coffee and non-coffee households. As above, coffee households have on average an 11.9% higher consumption level than non-coffee households.

4.2 Panel Estimates

In this section we take a more structural approach and investigate the impact of the program and the coffee crisis. We do so by studying the channels via which the drop in coffee prices may have affected household expenditures, and identifying the impact of transfers in mitigating such shock. Throughout we focus on the effects of receiving transfers as opposed to effect of eligibility.

Consider a model in which household expenditure is given by:

(1) $$y_{it} = X_{it}\beta + Z_i\theta + P_t\delta + \tau_t + \nu_i + \varepsilon_{it}$$

where y_{it} is expenditure (per capta), X_{it} are observable time-variant assets and other endowments from which households generate income, Z_i are observable time invariant (or predetermined as of 2000) assets and endowments, P_t is a vector of observable time specific, but household invariant, factors (as prices, weather, etc.), τ_t is a vector of unobservable time-specific effects, ν_i is an unobservable time-invariant (or predetermined as of 2000) household specific effect, and ε_{it} is a household- and time-specific random disturbance, with zero mean and assumed orthogonal to all other variables in the right hand side of (1).

The level of assets X_{it} employed to generate income y_{it} is likely to be correlated to P_t and τ_t since households adjust their endowments to changing exogenous incentives. But since P_t is invariant across households we cannot identify δ. Nevertheless, the effects of changes in both P_t and τ_t may be captured by the de inclusion of a dummy variable d_t that equals 1 if the year of the survey is 2002 and 0 otherwise.

Note that since β measures the marginal benefit of each asset in X_{it}, it is conceivable that β changes over time in response to changes in P_t and τ_t. Therefore we allow β to vary over time by interacting X_{it} with the dummy d_t. Then (1) becomes:

(2) $$y_{it} = X_{it}\beta_{00} + d_t X_{it}\gamma + Z_i\theta + d_t\eta + \nu_i + \varepsilon_{it}$$

where η captures the effects of changing prices and other time varying macro factors. The marginal benefit of the assets X_{it} in 2000 are therefore given by β_{00}, while the marginal benefits in 2002 are given by $\beta_{02}=(\beta_{00}+\gamma)$.

To estimate the parameters in (2) consistently, we specify a model in which households adjust their asset holding in response to changes in exogenous macro factors P_t and τ_t, and idiosyncratic (observed and unobserved) household characteristics Z_i and ν_I respectively. That is:

(3) $$X_{it} = Z_i\pi_1 + d_t\pi_2 + \phi\nu_i + u_{it}\,.$$

where the effects of exogenous macro factors are capture by π_2, and the effects of household observed characteristics are capture by π_1.

Note that unobserved household characteristics v_i appear as an error component of both (2) and (3). If $\phi \neq 0$, the X_{it}'s are endogenous, which implies that the OLS estimator of the parameters in (2) will be biased and inconsistent. While we could employ fixed-effects panel data methods to consistently estimate some of the parameters in (2), we opt for an *Instrumental Variable* (IV) approach to allow us to estimate the coefficients of the time-invariant variables in (2).[7] For some of the time-variant variables, the IV approach employed is equivalent to the fixed-effect estimator because we use the within transformation of X_{it} as instrumental variables. Table 6 below lists all the explanatory variables included in (2), the instruments employed for each of them, and the IV regression results. Note that the dummy variable for whether or not the household received a transfer is zero for all households in 2000, and takes the value of one only for those households that actually received transfers. Therefore, this variable is endogenous since some of the households might have self-selected out of the program due to unobserved heterogeneity.

Table 6 presents the results of the IV estimation of the parameters in model (2). The only difference between the models in columns (A) and (B) is that the latter includes an interaction term between the dummy for program participation (transfer receipt) and the ratio of coffee land owned to total land owned.

As it can be seen, owning land for non-coffee uses in 2000 and 2002 had positive and statistically significant effects on household expenditures. In 2000, one manzana of land for non-coffee use seemed to have pushed median expenditures upward by approximately 0.44%(P=0.043). In 2002 it seemed to have impacted it by 0.56%(P=0.015). At the median daily per capita expenditure of approximately 14Lps, this represents a low increase of about 22Lps per-capita per year (or 125Lps per household per year) in 2000, and 28Lps per-capita per year (or 160Lps per household per year). These results reflect the low returns to non-coffee land--used mostly for pasture and low productivity crops—in western Honduras.

[7] In addition, as we discuss below, a standard fixed-effects approach will not allow us to identify the effect of receiving a CCT on per capita expenditures.

Table 6. Instrumental variable estimates of the household income equation using panel households observed in 2000 and 2002

Dependent variable = Log of daily per-capta expenditures in Lps.	(A)	(B)
Dummy = 1 if received CCT[1]	0.139*	0.134+
	(0.056)	(0.079)
Total land area in Manzanas[2]	0.004*	0.004*
	(0.002)	(0.002)
Share of land under coffee[2]	0.117+	0.117+
	(0.062)	(0.061)
Log of household population[2]	-0.549**	-0.549**
	(0.044)	(0.044)
Share of the population ages >=11 and < 19[2]	0.081	0.080
	(0.096)	(0.096)
Share of the population ages >=19 and < 49[2]	0.389**	0.389**
	(0.102)	(0.101)
Share of the population ages >=49 and < 65[2]	0.321*	0.321*
	(0.135)	(0.135)
Dummy for year 2002[2]	-0.054	-0.053
	(0.059)	(0.063)
Dummy 2002 X Total land in Manzanas[2]	0.001**	0.001**
	(0.000)	(0.000)
Dummy 2002 X Share of coffee land[2]	-0.076**	-0.084
	(0.027)	(0.090)
Dummy 2002 X Log of household population[2]	-0.025	-0.024
	(0.029)	(0.029)
Dummy 2002 X Share of the population ages >=11 and < 19[2]	0.154*	0.153*
	(0.069)	(0.071)
Dummy 2002 X Share of the population ages >=19 and < 49[2]	0.005	0.005
	(0.070)	(0.070)
Dummy 2002 X Share of the population ages >=49 and < 65[2]	0.041	0.042
	(0.071)	(0.071)
Age of the household head[3]	0.003**	0.003**
	(0.001)	(0.001)
Age of the household head squared[3]	-2.907e-07**	-2.906e-07**
	(9.360e-08)	(9.387e-08)
Dummy = 1 if head of household is a male[3]	0.010	0.010
	(0.032)	(0.032)
Education of the household head[3]	0.071**	0.071**
	(0.022)	(0.022)
Education of the household head squared[3]	0.010	0.010
	(0.006)	(0.006)
Dummy = 1 if received CCT X Share of land under coffee[1]		0.020
		(0.254)
Constant	2.986**	2.986**
	(0.094)	(0.095)
Observations	9180	9180
Number of Households	4590	4590
R-squared	0.24	0.24

Standard errors in parentheses
+ significant at 10%; * significant at 5%; ** significant at 1%
[1]Municipal level availability of the program used as instrument
[2]Within transformation used as instrument
[3]Exogenous variable

Table 6 shows that, in 2000, ownership of coffee-land yielded much higher returns than land used for staple crops and pasture. In 2000, a manzana converted to coffee would have increased household per-capita expenditure by approximately 3% (P=0.038)[8], which represents an approximate increase of 150Lps per-capita (or 860Lps per household) in expenditure per year.[9]

For 2002, the data shows that there was a substantial, and statistically significant, deterioration on the returns to coffee land, which is consistent with the reported drastic drop in international coffee prices. As the results show, returns to coffee land were knocked down by approximately 1.5(P=0.009) percentage points, to 1.4%(P=0.271), and seized to be statistically significant. This represented a drop in expenditures of about 430Lps per year for the median household owning coffee land.

What has been the impact of the transfers on mitigating the negative effects of the drop in international coffee prices? As the results in Table 6 show, the program appears to have more than offset the drop in household expenditure of coffee growers caused by the lower returns to coffee land. Participation in the program seemed to have boosted expenditure per-capita by about 14% (or about 715Lps per-capita per year, and 4000Lps per household per year).

In addition to compensating coffee growers for lower returns to their assets, the cash flow provided by transfers may also help growers obtain higher returns on their coffee area by providing much needed liquidity in times of crisis. That is, since part of the drop in returns to coffee land may be caused by lack of resources to fund the maintenance of coffee trees, it is conceivable that program participants were better able to preserve their returns because of their greater access to liquidity. To test this hypothesis we interact the dummy variable indicating

[8] Under the specification on Table 6, the marginal impact of an extra unit of coffee land is given by the following expression:

$$\frac{\partial \ln(y)}{\partial C} = \beta_1 + \beta_2\left(\frac{T - C}{T^2}\right)$$

where y is per-capita expenditure, C is the mean area of coffee land owned by growers, T is the mean area of total land owned by growers, and the β's are parameters to be estimated.

[9] Calculated at median per capita expenditure of 14Lps per day.

participation in the program with the share of land owned. As can be seen in column (B) of Table 6, however, access to transfers does not seem to have had any impact on returns to coffee land.

5. Concluding Remarks

In this paper we have described the nature of coffee involvement in some of the poorest rural municipalities in western Honduras. Our analysis indicates that the region is very coffee intensive with a large proportion of land devoted to coffee and many poor rural laborers relying on employment on coffee farms. We have also examined the change in total household consumption over the period 2000 to 2002, a period when households were hit with two economic shocks, namely, a drought and a continuing decline in international coffee prices. We find that, although there is evidence of a decline in consumption, the pattern varies substantially across consumption quintiles, with the lowest quintile experiencing a substantial increase in consumption levels and the highest experiencing substantial decreases. We also find that the highest quintiles are substantially more likely to have land devoted to coffee.

In the second half of the paper we evaluate the impact of a transfer program implemented in the area over the period. Our results indicate that while the program has a substantial positive impact on the consumption of households, there is no differential impact across coffee and non-coffee households. We also find evidence that while the relative returns to coffee production are high, this return has decreased over the period consistent with a negative price shock. We also find that the program does not affect this return, as might be expected if it was substantially alleviating liquidity constraints.

References

FAO (Food and Agriculture Organization of the United Nations). At <http://apps.fao.org/subscriber>

Hearne, R., B. Barbier and J.M. Gonzalez (2002): "Development of a minimum cost, incentive-based plan for the implementation of a technology standard for coffee processing in Honduras". Paper presented at the 2002 AAEA Annual Meetings, Long Beach, CA.

ICO (International Coffee Organization). At <http://www.ico.org>

Morris, S., O. Neidecker-Gonzales, C. Carletto, M. Munguia, J.M. Medina and Q. Wodon (2002): "Hurricane Mitch and the livelihoods of the rural poor in Honduras", World Development, 30(1), pp49-60.

Partners (Partners of the Americas, Washington, DC) (2002): "Small scale coffee production in Honduras". At <www.partners.net>

Varangis, P., P. Siegel, D. Giovannucci and B. Lewin (2003): "Dealing with the coffee crisis in Central America: Impacts and strategies", Policy Research Working Paper, No. 2993, World Bank.

World Bank (2001): Honduras: Poverty Diagnostic 2000. Report No. 20531-HO, World Bank, Washington, DC.

Appendix: Program and Survey Design

The evaluation experiment was conducted in 70 municipalities in the west of Honduras with a total population of 660,000 in 2001 [Instituto Nacional de Estadística, 2002]. The municipalities were selected because they had the highest prevalence of malnutrition in the country according to a height census of first-grade primary school children conducted in 1997 [Government of Honduras, 1997]. For some benefits (see below), eligibility was restricted to households whose residence in a particular municipality had been recorded in a special census of the area conducted in mid-2000.

Two packages of conditional cash transfer interventions were planned and implemented in this area by the program. The first, which we term the *health and nutrition package*, was a cash transfer paid to pregnant women and women looking after children less than three years of age in households of established residence. Each eligible household received up to a maximum of two freely exchangeable vouchers worth 55 Lempiras per beneficiary per month. The second package, the *educational package*, was targeted to households with children 6 to 12 years of age (inclusive) enrolled in primary school in grades one through four received up to three vouchers worth 80 Lempiras per beneficiary per month for ten months of the year.

The vouchers were distributed on three occasions between the baseline and post-intervention surveys reported here: in November of 2000, and in May/June and October/November of 2001. A fourth round of voucher distribution coincided with the post-intervention survey in 2002. The payments were conditional in that all beneficiaries were informed that their payments would be suspended if they did not keep up to date with routine ante-natal care and preventive health care for children under three, and the school enrolment and minimum attendance requirement (85% of classes).

The total number of trial municipalities was determined by the budget available to the program. Because the entire population of each municipality could be receiving program benefits, it was necessary to restrict data collection activities to a representative sample of households in each municipality. Sample size calculations took into account the cluster-randomization, and were based on an *ex-post* comparison of 20 intervention and 20 control municipalities, with 80%

power to detect a significant difference (P=0.05, two-sided). We used the formula presented by Murray [1998; 368-9] for group-randomized trials with repeat observations of groups. The final size of the evaluation cohort was eighty households per municipality.

The representativeness of the evaluation cohort at the municipality level at baseline was ensured by: (i) randomly sampling eight census enumeration areas in each municipality with probability proportional to size; (ii) mapping all the dwellings in the enumeration area and numbering them consecutively; (ii) choosing a random start-point, and conducting interviews in ten consecutive inhabited dwellings following the direction of the numbering on the map. The same households were interviewed in the post-intervention survey. Women and young children in these households who had moved in the intervening period were followed up to their new homes (referred to as 'derived' households), provided these were located in one of the seventy trial municipalities or an adjacent one.

Each of the seventy municipalities was randomly assigned to one of four groups: (a) the household-level demand-side package alone, (b) the supply-side package, (c) both packages, and (d) control group. Before randomization, the municipalities were stratified into five groups of fourteen on the basis of the prevalence of stunting reported in the 1997 school height census [Government of Honduras, 1997]. Within each stratum, random allocation was achieved by drawing colored balls from a box after the name of each municipality was read out. Four balls represented the household-level package (alone), two the service-level package (alone), four the combination of both packages, and four the control group. Thus, the randomization was blocked as well as stratified.

The randomization was carried out by children in the presence of legal authorities and representatives of Honduras' agency for administrative probity. The aperture of the box was sufficiently small that once the child had inserted his/her arm, it was impossible for him/her to see the colored balls. From the day of the randomization onwards, there was no attempt to conceal the allocation. Because of the 2000 residence requirement, no household could become eligible for the cash transfers by moving house after randomization.