Intrahousehold Resource Allocation in Côte d'Ivoire: Social Norms, Separate Accounts and Consumption Choices*

Esther Duflo† and Christopher Udry‡

July 25, 2001

1 Introduction

Anthropologists often insist on the lack of fungibility of income when describing the flow of money within households in traditional economies, particularly in Africa. First, each household member has specific claims on particular sources of income: he or she retains ownership or usufructuary rights on a plot of land and thus primary claim to the income from that plot, or he or she is entitled to the proceeds from particular crops. The obligation to share this income with other household members is limited. While household members cooperate in some productive activities and share their outcomes to some extent, they seem to be far from achieving perfect risk sharing.

> "Men control their own cash income, and the kinds of legitimate demands a wife can make can be quite limited. A Yoruba wife can expect her husband to provide the basic staples of the diet, housing, and other more irregular support depending on how

*We thank Mike Gough, Shawn Cole and Jonathan Robinson for excellent research assistance and the John D. and Catherine MacArthur foundation and the National Science Foundation (SES-0079115) for financial support. We thank members of the MacArthur Network on the Effects of Inequality on Economic Performance, Daron Acemoglu, Abhijit Banerjee, Bo Honoré, Emmanuel Saez, and seminar participants at UMass/Amherst for helpful comments.

†MIT and NBER

‡Yale University

much domestic work she devotes to him (...) Beti wives remain farmers throughout their lives. Before the recent expansion of food sales they used the depend on their husbands for all major cash expenses, but neither in theory nor in day-to day life is a wife's right to her own share of her husband cash income guaranteed (..) Family welfare and risk avoidance are probably improved by the family labor force having a variety of occupations which cater to different markets, but the need in bad times and the opportunity in good times for a woman to earn an independent income originate in a domestic organization with limited income sharing" Guyer (1987)

Furthermore, the source of income may also determine its legitimate uses, and the uses of money obtained through particular activities may be restricted. In Kenya, Shipton (1989) describes how money obtained form the sale of land, tobacco, or gold, is "bitter" and "... must be kept strictly apart from transactions involving permanent lineage wealth and welfare, notably from livestock or bridewealth transactions (p. 25-26)." In Cote d'Ivoire, the Gouro, studied by Meillassoux (1965) draw a sharp distinction between "appreciated products" (rice and yams), ordinary food products, products cultivated by women, and cash crops.

'Appreciated products' are always under the control of the household head *for redistribution to the entire household in the form of food*. In contrast, the control of cash crops and food products belongs to its producer. Cash crops and food crops, even when they are cultivated by the same individual, and even when food crops are sold on the market, are not put to the same use:

> "In the traditional community, as we have seen, most of the production comes back to producers in the form of food. The rest is incorporated into particular goods, which have a specific role at the time of marriage (...). These goods cannot be diverted to personal uses. Nor are they investment goods, used for the reproduction of material goods. Everything changes when the products of agriculture are cash crops, which can be put to other uses (...). A greater part of this income disappears into prestige expenditures, especially into investment into houses which are monuments to the glory of their owners." (p 335).

These descriptions are fundamentally at odds with conventional ways in which economists describe individual and household behavior. Standard models imply that there should be a unified budget constraint for the entire household. If there is more than one individual, the average share allocated to an individual's consumption may depend on her bargaining power (which may well be related to her average contribution to the household, and hence to her permanent individual income), but her consumption should not fluctuate on a day to day basis as a function of the realization of her income. But these descriptions suggest that resources generated from different activities within the household are used differently. Taken literally, these descriptions imply that households maintain a series of discrete 'accounts' into which different revenue flows are directed, out of which different expenditures are made, and between which transfers are not freely made. When an account gets a windfall, expenditures out of that 'account' increase more than others.[1] This has a parallel in the "mental accounting" described in the behavioral economics literature Thaler (1990): money placed by individuals in different "mental accounts" is not fungible.

This paper seeks to test the empirical relevance of these descriptions in the context of rural Cote d'Ivoire. Efforts to empirically validate the mental accounts framework in behavioral economics have mostly concentrated on comparing the propensity to consume out of income from various types of flows: the propensity to consume out of housing wealth is very low, and the propensity to consume out of current income is very high, for example.[2] In this paper, we do not focus on the marginal propensity to consume out of different sources of income. Instead, we explicitly recognize that a given increase in observed current income in a given account may be more or less permanent, depending on the type of accounts it falls in (and thus may affect consumption differently), and we test whether shocks to different types of income affect expenditure shares over and above their effects on overall expenditures.

[1] Many descriptions imply that the separation of these "accounts" is not limited to the uses of crop proceeds, but extend to income from nonfarm enterprises and to inputs: family or community work can be used for food crops without compensation other than a share in the common meal, while cash wages are paid to household workers who help with cash crops (Berger and White (1999); Ekejiuba (1995); Etienne (1980); Guyer (1984)).

[2] There are other examples: for example, most people who take a second mortgage on their house use the money to finance home improvements

The fact that the proceeds of different crops are generally used to buy different goods does not necessarily imply that the household really maintains separate accounts. If individuals in the household have ownership rights on specific income streams, those who earn more could have more bargaining power: their income will thus appear to be linked with different purchases. For example, using anthropological evidence from Cote d'Ivoire that attributes the proceeds from some crops to different genders Haddad and Hoddinott (1994) show that income from "male crops" tends to be put to different uses than income from "female crops". This is not consistent with a *unitary* model of the household (where all household members have the same utility function or a dictator makes decisions for everyone) but could be consistent with the more general collective model (proposed notably by Chiappori, see Browning and Chiappori (1998) for a survey), where individuals may bargain over the household allocation, but achieve Pareto efficiency. Thus, one response of most economists to descriptions such as those we quoted above would not necessarily be to deny the reality of the norms which underlie these descriptions, but to argue that households have sufficient flexibility on the margins to undue any binding constraints on expenditures that would otherwise result. On average, the norms will be respected, but at the margin money is fungible and it is possible to shift household expenditures in such a manner that the norm does not prevent the household from achieving an efficient allocation of resources.

In this paper, we present evidence that expenditure patterns in Cote d'Ivoire not only violate the restrictions implied by the collective household model, but that they do so in a way that corresponds closely to the descriptions that can be found in the literature on the norms of household provisioning in Cote d'Ivoire. The central observation underlying our empirical approach is that if the household is efficient, household members fully insure each other against short term variation in individual income. Therefore, non-persistent income shocks should not translate into differences in the allocation of resources within the household.

To identify short term income shocks, we use rainfall variation. While all household members are subject to the same rainfall, the same pattern of precipitation has different effects on the income produced by different crops. In particular, a particular rainfall pattern affects differently crops that tend to be produced by women and crops that tend to be produced by men. In a Pareto efficient household, conditional on total expenditure this should not translate into any difference in the allocation of that expenditure to different purposes within the household. The spirit of our test is thus to test whether two rainfall configurations that have the same effect on total expenditure have different effects on the types of goods consumed by the household. We present two tests which rely on different assumptions: an exclusion restriction test which treats total expenditure as exogenous (and ignores measurement error in expenditure) and an overidentification test, which assumes a specific functional form for the relationship between total expenditure and expenditure on different categories, but not the exogeneity of expenditure or the absence of measurement error. We examine broad expenditure aggregates and more detailed expenditures on food. The results from our two types of tests are entirely consistent: we reject the hypothesis of income pooling. Furthermore, the patterns of rejections we obtain are consistent with the anthropological descriptions of income flows in Ivoirian households. The evidence presented in this paper supports the validity and the empirical relevance of the descriptions of separate accounts within households in Cote d'Ivoire. This observation can have far-reaching consequences for our understanding of the behavior of households, both as consumers and as producers.

The remainder of the paper proceeds as follows: in section 2, we derive our empirical test. In section 3, we discuss the data and the context of agriculture in Cote d'Ivoire. In section 4, we discuss our results. Section 5 concludes.

2 Theoretical framework and derivation of the test

2.1 Theory

Our objective is to use rainfall as a source of exogenous variation in income from various sources to examine a testable restriction of the collective model, the assumption that income from all sources is pooled. To put the question more bluntly: does rainfall variation that affects farms cultivated by a wife change the pattern of expenditure within households differently than rainfall variation that affects farms cultivated by her husband? Or, does rainfall variation that affects cash crop production have the same effect on consumption as rainfall variation that affects staple foods?

We illustrate our empirical strategy first in the context of a simple one-period model of intra-household resource allocation in a risky environment and then move to the more general dynamic case. It will be seen that the lessons from the one-period model generalize in a straightforward manner.

To simplify the notation in this section, we consider the optimization problem of a household comprised of 2 individuals, each of whom produces only one type of crop. Of course, this generalizes in a straightforward way to a situation where each produces different types of crops. Each individual $i \in \{m, f\}$ consumes a private good c_i and a public good c_p. Individual i's preferences are summarized by the utility function $u_i(c_i, c_p)$. The results that follow are robust to significant generalizations of these preferences: an individual's utility may depend on the consumption or utility of his/her spouse, and c_i and c_p may be vectors. A more substantial assumption is that labor is supplied inelastically, or that preferences over leisure are separable from preferences over other consumption. This will be discussed below.

Each individual cultivates a farm using labor (L_i) that can be traded on a competitive market at wage w.[3] The production function on farm i is $f_i(L_i, r)$, where $r \equiv \begin{pmatrix} r_1 \\ r_2 \end{pmatrix}$ is a vector of two (it will be seen that this is simple to generalize) measures of rainfall that affect cultivation on plot i. For example, r_1 might be early season rainfall and r_2 might be rainfall late in the season.

Any *ex ante* efficient allocation of resources in the household can be characterized as the solution to the program

$$\max_{c_i, c_p, L_i} u_f(c_f, c_p) + \lambda u_m(c_m, c_p)$$

subject to

$$c_m + c_f + pc_p \leq f_f(L_f, r) + f_m(L_m, r) - w(L_f + L_m).$$

Note that the Pareto weight does not depend on r: in the efficient allocation risk is pooled. We do not investigate the process through which λ is set; it may depend on many observable or unobservable attributes of the household and its members. For example, long-run rainfall patterns that influence profitability differently on the husband's and wife's plots may influence the Pareto weight. Hence, in examining the relationship between rainfall variation and the allocation of resources within the household it is important to distinguish adequately between

[3] It is a trivial matter to extend the model to include a vector of inputs, which may be purchased, non-traded, or traded on imperfect markets.

the realizations of random variables (which in an efficient allocation are pooled, that is to say, they do not influence λ) and the distribution from which those realizations are drawn (which may well be a determinant of λ).

This problem is separable and equivalent to

$$\max_{c_i, c_p} u_f(c_f, c_p) + \lambda u_m(c_m, c_p) \tag{1}$$

subject to

$$c_m + c_f + pc_p \leq \pi_f^*(r) + \pi_m^*(r) \tag{2}$$

where $\pi_i^* \equiv \max_{L_i} f_i(r, L_i) - wL_i$. Note that rainfall enters the efficient allocation of resources only through its effect on cultivation and hence on the budget constraint, and thence on total expenditure.

Denoting $x = c_m + c_f + pc_p$, we have:

$$c_i = c_i(\lambda, p, x) \tag{3}$$

for $i \in \{m, f, p\}$. Conditional on expenditures, prices, and the preference and Pareto weight parameters, consumption of any particular good is independent of the rainfall realizations r_1, r_2.

Equation (3) has the additional implication that the effect of rainfall realizations on expenditure on any particular commodity depends only on the expenditure elasticity of demand for that commodity and on the effect of rainfall on overall expenditure.[4] In other words, for any i in m, f,and j in $1, 2$:

$$\frac{dc_i}{dr_j} = \frac{\partial c_i}{\partial x} * \frac{\partial x}{\partial r_j}. \tag{4}$$

The collective model, therefore implies the restriction that the ratio between the effect of rainfall in quarter j on consumption of good i and its effect on total expenditure should be equal across all rainfall realizations:

[4] To reduce notational clutter, we assume for the time being that the relative prices of consumption are not related to rainfall realizations ($\frac{\partial p}{\partial r_i} = 0$). This need not be the case, and it is addressed in our empirical specification.

$$\frac{\frac{dc_i}{dr_1}}{\frac{\partial x}{\partial r_1}} = \frac{\frac{dc_i}{dr_2}}{\frac{\partial x}{\partial r_2}}. \tag{5}$$

There is an analogous test for $c_m + c_f$ if only aggregate consumption of the private good is observed, and if c_f, x_m, or c_p are vectors (5) holds for any component of these vectors that is observed.[5]

It is useful at this point to comment that the essential element of the restrictions (3) and (5) is the assumption that rainfall variation dr_i affects the collective household's decision making only via its influence on the household's resource constraint. In a more general model in which rainfall entered preferences directly, these restrictions fail to hold. This caveat should be borne in mind when considering the results that follow. For example, if one component of c_p is the health status of children in the household, a rejection of these restrictions need not imply that the household allocation is inefficient if r_1 has a different direct effect on child health than r_2.

It is also useful at this point to note that the assumption of a perfect labor market is *not* essential to this analysis. Precisely the same test emerges in a model in which there are no inter-household labor flows; if supervision is required for non-household labor; for general forms of imperfect substitutability between household and non-household labor; or if inter-household labor flows through reciprocal or cooperative arrangements. In each of these instances, it remains the case that labor decisions affect consumption only through their influence on the household's resource constraint and the collective model continues to imply both (5) and (3). What *is* essential is the assumption that conditional on total expenditure, the consumption of leisure does not affect the marginal rates of substitution between the other components of consumption.

It is now a simple matter to generalize these observations to the more general dynamic case. Consider a collective household with a horizon of T periods. In period t after a history of rainfall realizations $w_t \equiv \{r_1, r_2, ..., r_t\}$ individual i consumes a private good c_{iw_t} and the household jointly consumes a public good c_{pw_t}. The expected utility of individual i is $E \sum_t \beta_i^t U_i(c_{iw_t}, c_{pw_t})$.

[5] Data on rainfall and expenditures is required to estimate (3) or (5). We do not observe π_i, nor is such data required for the test. Hence, we avoid the issues raised by Rosenzweig and Wolpin (2000), which provides a very useful discussion of the potential consequences of treating estimates of the relationship between rainfall variation and output variation as if it were the relationship between rainfall and profits.

The budget constraint facing the household in period t after history of rainfall realizations w_t (note that w_t includes the rainfall realization in the current period) is

$$c_{fw_t} + c_{mw_t} + c_{pw_t} + A_{w_t} \leq RA_{w_{t-1}} + \pi^*_{mw_t} + \pi^*_{fw_t}, \tag{6}$$

where A_{w_t} is the amount invested after history w_t by the household in a safe asset that earns a return R.[6]

There is a budget constraint for each history of rainfall realizations, so for example the budget constraint in period t following rainfall history $\hat{w}_t \equiv \{\hat{w}_{t-1}, r_t\}$ is not the same as that after history $\tilde{w}_t \equiv \{\tilde{w}_{t-1}, r_t\}$ if $\tilde{w}_{t-1} \neq \hat{w}_{t-1}$. For notational simplicity, we have not permitted any inter-household insurance, though this would leave the problem essentially unchanged.

Any efficient allocation of household resources can be characterized as the solution to:

$$\max_{\{c_{iw_t}, c_{pw_t}\}} E \sum_t \beta_f^t U_f(c_{fw_t}, c_{pw_t}) + \lambda E \sum_t \beta_m^t U_m(c_{mw_t}, c_{pw_t}) \tag{7}$$

for some value of λ, subject to (6) and a period T constraint on A_{w_T}. An efficient allocation must have efficient continuations after any history of rainfall w_t, so in period t an efficient allocation must be the solution of

$$\max_{c_{iw_t}, h_{w_t}, A_{w_t}} U_f(c_{fw_t}, c_{pw_t}) + \lambda U_m(c_{mw_t}, c_{pw_t}) + V_{w_t}(A_{w_t}; \lambda) \tag{8}$$

subject to

$$c_{fw_t} + c_{mw_t} + pc_{pw_t} + A_{w_t} \leq RA_{w_{t-1}} + \pi^*_{mw_t} + \pi^*_{fw_t}. \tag{9}$$

The function $V_{w_t}(.)$ is complex: it depends on the preference parameters (including λ) and on the information about the distribution of future profits that is incorporated in the history of rainfall through the current period. However, the maximand is separable between $\{c_{iw_t}, c_{pw_t}\}$ and A_{w_t}. Let $A^*_{w_t}$ be the efficient level of assets held after w_t. Then efficient consumption is

$$\{c^*_{iw_t}, c^*_{pw_t}\} = \arg \max_{c_{iw_t}, c_{pw_t}} U_f(c_{fw_t}, c_{pw_t}) + \lambda U_m(c_{mw_t}, c_{pw_t})$$

[6] It is trivial to generalize the investment process to make it so that people are investing (perhaps in their farms), that this return depends on rainfall, that it is uncertain, or that they allocate these savings across a portfolio of assets. The only change to the model will be the additional notation, because it will all affect the allocation of current consumption only through the function $V_{w_t}(A_{w_t})$ in equation (8).

subject to

$$c_{fw_t} + c_{mw_t} + pc_{pw_t} \leq RA_{w_{t-1}} + \pi^*_{mw_t} + \pi^*_{fw_t} - A^*_{w_t}$$

Since $x_{w_t} \equiv RA_{w_{t-1}} + \pi^*_{mw_t} + \pi^*_{fw_t} - A^*_{w_t}$, we have once again

$$c_{iw_t} = c_i(\lambda, p, x_{w_t}) \tag{10}$$

for $i \in \{f, m, p\}$. Conditional on expenditure, prices, and preference and Pareto weight parameters, consumption of particular goods is independent of rainfall realizations r_1, r_2. Thus an analogous form of the exclusion restriction (3) holds in the dynamic setting. It can now be seen that the general dynamic problem is akin to the static model discussed above. Again, the crucial restriction of the collective model is income pooling: realizations of rainfall influence the allocation of current consumption only through their affect on current expenditure. We have the additional testable restriction that

$$\frac{\frac{\partial c_{iw_t}}{\partial r_1}}{\frac{\partial x_{w_t}}{\partial r_1}} = \frac{\frac{\partial c_{iw_t}}{\partial r_2}}{\frac{\partial x_{w_t}}{\partial r_2}} \tag{11}$$

must hold for any observable consumption good.

Equation (11) describes a 'reduced form' test of the collective household model, and (10) is a test of the collective model based on exclusion restrictions. In the empirical work below, we present both these tests and particular linear combinations of these tests that are more straightforward to interpret.

2.2 Empirical implementation

Equations (10) and (11) form the bases of our tests that the household has a single budget constraint. Equation (10) is a exclusion restriction: once we control adequately for total expenditure and preferences in the household (including the sharing rule λ) the rainfall variables should not enter in the equation determining the demand for any good consumed in the household. Equation (11) is an overidentifying restriction. It says that a particular rainfall realization must influence the demand of a particular good only to the extent it influences expenditure. We will implement two types of test, based on these two restrictions.

2.2.1 Test based on exclusion restrictions

• Semi-parametric tests

We start by assuming a semi-parametric model for the demand function for any good c, $c(\lambda, x, p)$.

Let i denote the household and t the period. For any particular good c, we assume:

$$\log(c_{it}) = \phi(\log(x_{it})) + f(\lambda_i) + X_{it}\delta + \upsilon_i + \nu_{it}, \tag{12}$$

Where X_{it} are year and region interactions (for the four agroclimatic zones in Cote d'Ivoire), υ_i is a household fixed effect, and ν_{it} is an error term that potentially reflects the effect of changes in relative prices on demand as well as other shocks to the preferences of households. If we assume that markets are regionally integrated, so that at a given point in time relative prices are the same across a particular agroclimatic zone, the effect of rainfall on relative prices and hence patterns of demand is contained within the fixed region-year effect (X_{it}). The most controversial element of this assumption is that it does not permit the price faced by a household to differ from that prevailing in the region. If local markets are not well integrated into a regional system, then this assumption is violated. We discuss this possibility when we comment on our empirical results.

With these assumptions, the test suggested in equation (10) takes a very simple form: rainfall variables do not belong in equation (12). However, it is clear that this test should not be implemented using a single cross-section of data. This year's level of rain is related to the permanent level of rainfall in the area, which could in turn be related to households' demand for different type of goods, in particular through the bargaining power of each household member (and hence, λ_i and thus υ_i). With data from two periods ($t = 1, 2$) on each household, we can take the first difference of equation (12) to obtain:

$$\log(c_{i2}) - \log(c_{i1}) = \phi(\log(x_{i2})) - \phi(\log(x_{i1})) + (X_{i2} - X_{i1})\delta + (R_{i2} - R_{i1})\gamma + \nu_{i2} - \nu_{i1}. \tag{13}$$

To estimate the parameters γ and δ we re-write equation (13) as:

$$\log(c_{i2}) - \log(c_{i1}) = g(\log(x_{i2}), \log(x_{i1})) + (X_{i2} - X_{i1})\delta + (R_{ij} - R_{i1})\gamma + \nu_{i2} - \nu_{i1}. \tag{14}$$

We then follow Robinson (1988) and Hausman and Newey (1995) and estimate the coefficients of rainfall and the set of region × time period indicator variables in the partially linear specification (14). To simply notation, let y be the vector $\log(c_2) - \log(c_1)$, z_1 be the vector $\log(x_1)$, z_2 be the vector $\log(x_2)$, m be the matrix $[(X_{i2} - X_{i1}) \;\; (R_{i2} - R_{i1})]$ and β be the vector of parameters $[\delta, \gamma]$. The estimator of β is:

$$\hat{\beta} = \left[\sum_{i=1}^{N}(m_i - \hat{E}[m|z_{1i}, z_{2i}])(m_i - \hat{E}[m|z_{1i}, z_{2i}])'\right]^{-1} \left[\sum_{i=1}^{N}(m_i - \hat{E}[m|z_{1i}, z_{2i}])(y_i - \hat{E}[y|z_{1i}, z_{2i}])'\right]$$

In other words, the parameters of $\beta = [\delta, \gamma]$ are obtained by estimating (separately) the nonparametric relationships between y and m with (z_1, z_2), forming the residuals, and regressing the residuals of y on the residuals of m. We can then test whether $\gamma = 0$ with a regular F test. The non-parametric estimator we use to estimate the conditional expectations is the Fan (1992) locally weighted regression, with a quartic kernel.

One drawback of the test whether $\gamma = 0$ is that it is not explicitly linked with variation in income from various origins, and cannot thus not be directly linked to the anthropological evidence. To make this link explicit, we form linear combinations of the elements of $(R_2 - R_1)$ that reflect the variations in income from various sources. This is implemented by estimating a linear regression separately for each group of crops (male crops, female crops, and yams) of the difference over years of the logarithm of income from each crop in the group (output valued at market price minus inputs valued at market price) on the difference over the two years in rainfall realizations, and calculating the predicted values from these regressions.

Hence, with $s \in \{m, f, y\}$ defining a specific group of crops, we estimate:

$$\log(y_{is2}) - \log(y_{is1}) = (R_{i2} - R_{i1})\gamma_{ys} + (X_{i2} - X_{i1})\delta_{ys} + (\xi_{sw2} - \xi_{sw1}). \tag{15}$$

We then form: $DR_{is} = (R_{i2} - R_{i1})\hat{\gamma}_{ys}$, and we estimate the relationship:

$$\log(c_{i2}) - \log(c_{i1}) = g(\log(x_{i2}), \log(x_{i1})) + (X_{i2} - X_{i1})\delta + \sum_{s=1}^{S} DR_{is}\gamma_s + \nu_{i2} - \nu_{i1}. \tag{16}$$

We can then test whether any γ_s is individually different from 0, whether they are jointly significant, and whether they are significantly different from each other. The principle is identical to that in equation (14): conditional on $\{x, p\}$ and household characteristics, no indicator of rainfall realizations should influence the demand for any particular good. Here, though, instead of looking generally at rainfall realizations we focus on three particular dimensions of rainfall variation that are related to different sources of income to the household and that therefore correspond to different aspects of 'provisioning' that are prominent in qualitative discussions of household economics in West Africa.[7]

• **Parametric tests**

We are independently interested in the shape of ϕ. In particular, we want to assess whether it presents any strong non-linearity. To estimate it, we first obtain the estimate of $g(z_1, z_2)$ by partialling out the coefficient of m (see Robinson (1988), Hausman and Newey (1995)):

$$\hat{g}(z_1, z_2) = \hat{E}[y|z_1, z_2] - E[m|z_1, z_2]\hat{\beta}.$$

We then apply the partial means method suggested by Porter (1996) to recover the shape of $\phi(.)$ (up to an unidentified constant term) in equation 13:

$$\hat{\phi}(z) = 0.5 * \left(\frac{1}{N} \sum_{j=1}^{N} \hat{E}[y|(z, z_{2j})] \right) - 0.5 * \left(\frac{1}{N} \sum_{j=1}^{N} \hat{E}[y|(z_{1j}, z)] \right).$$

Pointwise confidence intervals for $\hat{\phi}(z)$ are constructed based on 50 bootstrap replications. As we shall see below, for almost all the goods we consider, $\phi(.)$ is very close to linear. We thus also estimate parametric versions of equations (14) and (16), where we impose $g(z_1, z_2) = \alpha(z_2 - z_1)$, e.g.:

$$\log(c_{i2}) - \log(c_{i1}) = \alpha(\log(x_{i2}) - \log(x_{i1})) + (X_{i2} - X_{i1})\delta + (R_{ij} - R_{i1})\gamma + \nu_{i2} - \nu_{i1}. \quad (17)$$

2.2.2 Test based on overidentifying restrictions

The tests based on exclusion restrictions (both the parametric and the non-parametric versions) present some potentially serious problems: in the presence of measurement errors in expenditure, the relationship between total expenditure and the expenditure in a particular good may be over

[7] DR_{is}, therefore, is the predicted relationship between rainfall realizations and the change in net output of crop group s for household i, where output is net of purchased inputs but not family labor.

or understated.[8] Moreover, shocks to total expenditure could be caused by events that also affect preferences (for example, a drop in expenditure could be due to sickness, and this could also lead to an increase in medical expenditure). If the model is misspecified, the coefficients of the rainfall variables (or of DR_{is}) will be inconsistently estimated as well and misleading conclusions could be drawn.

These problems do not affect the test based on overidentifying restrictions, suggested in equation (11). However, these tests require us to specify functional forms for the relationship between total expenditure and rainfall, and for the relationship between expenditure on each good and total expenditures. It seems reasonable to assume that rainfall in millimeters affects the logarithm of profits. In turn, assuming that the utility function of each household member exhibits constant relative risk aversion, the relationship between total expenditure and total profit is linear in logarithms, so that for a vector of rainfall R, we have the following relationship between rainfall and total outlay:

$$\log(x_t) = R_{it}\kappa + X_{it}\delta_x + \epsilon_{it}. \tag{18}$$

As described above, we have estimated non-parametric relationship between the logarithm of total expenditure and the log of expenditure for different types of goods, and we do not reject the hypotheses that these relationships are linear. Therefore we posit the following functional form:

$$\log(c_{it}) = \alpha \log(x_{it}) + f(\lambda_i) + X_{it}\delta + \upsilon_i + \nu_{it}. \tag{19}$$

Therefore, the reduced form relationship between the demand for good i and rainfall precipitation is

$$\log(c_{it}) = R_{it}\pi + f(\lambda_i) + X_{it}\delta + \nu_{it}. \tag{20}$$

Taking first differences of equation 18 and 20, we obtain the reduced form system:

[8] Imagine that food expenditure is measured with error: since it is a important part of total expenditure, the measurement error appears both on the left and on the right of the equation, leading us to overestimate the relationship between total expenditure and food expenditure. See Deaton (1997) and Bouis and Haddad (1992) for discussion.

$$\log(x_{i2}) - \log(x_{i1}) = (R_{i2} - R_{i1})\alpha + (X_{i2} - X_{i1})\delta_x + (\epsilon_{i2} - \epsilon_{i1}), \tag{21}$$

$$\log(c_{i2}) - \log(c_{i1}) = (R_{i2} - R_{i1})\pi + (X_{i2} - X_{i1})\delta + (\nu_{i2} - \nu_{i1}), \tag{22}$$

The test suggested in equation 11 is a simple overidentification test: we want to test the hypothesis that

$$\pi = \kappa\alpha \tag{23}$$

for some scalar κ. In the empirical work below, we use a non-linear Wald test to test this hypothesis.

As above, a difficulty with the test based on (23) is that it is not explicitly linked to any particular budgeting process within the household. However, note that equations (21) - (23) imply that if we consider *any* linear combination of the elements of $(R_{i2} - R_{i1})$ (say DR_{is}), we can construct an analogous test. Hence, we also construct a more specific version of the test using rainfall variation along the dimensions associated with changes in net output of particular types of goods instead of the entire vectors of rainfall variables. To do so, as above we estimate the regressions (15) to obtain $D\hat{R}_{is}$ and in turn estimate

$$\log(x_{i2}) - \log(x_{i1}) = \sum_{s=1}^{S} DR_{is}\alpha_s + (X_{i2} - X_{i1})\delta_x + (\epsilon_{i2} - \epsilon_{i1}) \tag{24}$$

and

$$\log(c_{i2}) - \log(c_{i1}) = \sum_{s=1}^{S} DR_{is}\pi_s + (X_{i2} - X_{i1})\delta + (\nu_{wi2} - \nu_{wi1}) \tag{25}$$

and we test whether $\frac{\pi_s}{\alpha_s} = \frac{\pi_{s'}}{\alpha_{s'}}$ for any $s, s' \in \{f, m, y\}$.

3 Data and context

3.1 Data

The data for this paper comes from the Cote d'Ivoire Living Standard Measurement Survey. The survey started in 1985, with 1,500 households. In 1986, half of these were re-surveyed, and 750 households were added to the survey. In 1987, the households newly introduced in 1986 were surveyed again and 750 new households were added. In 1988, a final wave of the survey was collected in the same fashion. For this study, we stack the 3 waves of the panel (1985/86, 1986/87 and 1987/88). The data set includes a wealth of information on the households, including information on their income from agriculture and other sources, health and education variables, ethnic affiliation, and a detailed expenditure survey.[9]

The data indicates separately the output of each crop cultivated by the household and the inputs spent on its cultivation. However, it does not record labor supply separately for each crop. It can also be merged with data from rainfall stations near the communities where the household is interviewed. Rainfall is recorded monthly for the past 14 years for most rainfall stations. We construct for each household aggregate rain recorded at the nearest rainfall station for each calendar quarter for the year that immediately preceded the most recent harvest (we label this as "current year") and for each quarter of the previous year.

We drop households that reside in Abidjan. We keep only households engaged in agriculture, where there is at least one man and one woman, and where households produce at least one crop defined as "male only" and one crop defined as "female only". In addition, some observations are dropped because of lack of information on rainfall. Our final sample has a little over 800 households (each observed twice).

3.2 Gender, ethnicity, and agriculture in Cote D'Ivoire

Farmers in Côte d'Ivoire work in a variety of agroclimatic conditions, from the rather dry savannah in the north to wet forest in the south. In no region is irrigation commonplace; almost all cultivation is rainfed. Rural households are heavily dependent upon crop income for their livelihoods: in rural areas of Côte d'Ivoire, farm income makes up 75% of total household income (Kozel (1990), Vijverberg (1988)).

An important characteristic of the organization of agriculture in Côte d'Ivoire as in other West African contexts is that much production takes place on plots that are managed by particular individuals within the household. Decision-making authority with respect to cultivation on these plots rests with that individual, cultivation expenses are paid by that individual and income from the plot is attributed to that individual. Household members commonly provide

[9] It is publicly available on the World Bank LSMS web site.

labor on each others' plots, at least partly as a consequence of a gender division of labor by task that cuts across the gender division of crops. Therefore, individuals in households rarely have absolute autonomy with respect to decision-making on their individual plots. However, a voluminous literature makes it clear that individuals have substantive control over decisions on their plots, and that nominal control over the output from a plot belongs to the cultivator.[10] One goal of this paper is an examination of the hypothesis that this nominal control over output from a plot influences the allocation of consumption within the household.

Our basic test of the efficiency of the pooling of income and risk within households in Côte d'Ivoire does not rely on any particular mapping between the gender of the cultivator and the crops he or she cultivates. However, directed by the descriptive literature, we refine the test by constructing three linear combinations of rainfall realizations, one for the cash crops cultivated by men, one for yams (which are cultivated by men), and the other for crops cultivated by women. We follow the method of Haddad and Hoddinott (1994) by drawing on the ethnographic literature to carry out the assignment.

We treat separately yams, the main "appreciated product", and the only major food crop controlled by men throughout the country. [11]. The other crops assigned to men are cocoa, coffee, wood, pineapple and kola nuts. Coconut, plantain, oil palm, taro, sweet potato, vegetables, banana, fruit trees and some minor crops are assigned to women.[12] For cassava, maize, tobacco, and sugar cane the evidence is not sufficiently strong that the crops are substantially more likely to be grown on the plots of one gender or the other, so they are not assigned. In addition, there is some ethnographic evidence that cotton, rice, millet, sorghum and fonio can be assigned to particular genders in some ethnic groups, but we do not consider them. Approximately 80

[10] Doss (1998), Doss (2001), Bassett (1985), Bassett (1988), Bigot (1979), Davison (1988), Dey (1993), Saito, Mekonnen and Spurling (1994), Gastellu (1987), Guyer (1987), Guyer and Peters (1987), Jones (1986), Meillassoux (1975), Berry (1993), von Braun and Webb (1989), Carney and Watts (1991), Goldstein (2000), Weekes-Vagliani (1985), Weekes-Vagliani (1990).

[11] Rice is a male crop in some groups, a female crop in others, and in others the gender pattern of rice cultivation is very complex.

[12] Meillassoux (1965), Weekes-Vagliani (1985), Weekes-Vagliani (1990), Bassett (1988), Gastellu (1987) are the primary sources for the assignment. The sources used by Haddad and Hoddinott (1994) are a subset of this group. Our assignments differ from theirs only in that ours are somewhat more conservative; some crops that they assign to a gender we leave unassigned.

percent of the value of agricultural output can be attributed in this manner.

It is important to note that no crop is exclusively cultivated by farmers of only one gender. Reporting from neighboring Ghana, Doss (2001) relates, "...I spoke with a woman who emphatically explained that yams were a man's crop and then invited me to see her yam farm." The 1991-92 round of the Ghana Living Standards survey provides information on the crops cultivated on particular plots and responses to the question "Who keeps the revenue from the sale of the produce?" Unfortunately, data on plot-specific crop production is not collected, but it is possible to examine the frequency with which farmers of different genders engage in the cultivation of particular crops. Doss (2001) carries out this exercise and shows that substantial numbers of both male and females are engaged in the cultivation of each of the 31 crops specified in the GLSS data. For no crop are women a majority of the cultivators. However, it is the case that there are systematic differences across crops in the likelihood that they are cultivated by women relative to men. For example, plantain farmers are approximately 50 percent more likely to be female than are cocoa farmers.

4 Results

4.1 Effects of rainfall on income from crops

Columns 1, 2 and 3 in table 2 present F statistics obtained after the estimation of equations 15 for yams, other male crops, and female crops. The estimated equations are presented in table A1. We include as male (or female) crops only those crops that are cultivated by males (or females) in all ethnic groups. In all equations, we include year and region effects (for the 4 agroclimatic regions in Cote d'Ivoire) and their interactions. The normal pattern of rainfall in these seasons is very different in forest areas and in the Savanna: in the forest, there are two rainy seasons (March to June and September to November) and two dry seasons, while there is only one rainy reason in the savannah. We partition the year into 4 seasons (December to February, March to June, July and August, and September to November), and we allow for different coefficients in the Savannah and in the Forest. We include rainfall for the 8 seasons prior to the most recent harvest. We use two types of rainfall variables: rainfall precipitation in millimeters, and a variable that indicates a particularly severe 'shock' when the rainfall precipitation was more than one standard deviation above or below its 14 year mean in this station. Therefore, we estimate 32 coefficients for each equation (except for the male cash crops which are cultivated only in the forest).[13]

As the F tests in table 2 indicate, rainfall variables are jointly significant in all regressions, and the coefficients are significantly different in each of them. Specifically, past year rainfall matters more than this year rainfall for the male cash crops (mostly tree crops), while both past and current year rainfall realizations matter for female crops and yams. The coefficients in the appendix reveal that in the savannah, rainfall shocks influence yam income in every quarter where they occurred, while the savannah interactions are not significant for women's crops. In the forest, shocks in the most recent long dry and long rainy season negatively affect both yam and female crops.

Thus there are strong differences across crop groups in the relationship of rainfall realizations to net income. This suggests that a test of income pooling based on evaluating whether rainfall patterns that affect different crop groups influence expenditure shares of different goods over and above their effects on total expenditures could have some power.

4.2 Exclusion restriction tests

We begin by presenting the results of the exclusion restriction tests: controlling for variation in total expenditure, does the specific pattern of rainfall affect expenditures on particular types of goods? In particular, do rainfall realizations that affect income from cash crops cultivated by men, crops controlled by women, and yam income respectively have different effects on the allocation of expenditures across goods?

[13] Our choice of a specification was driven by the agroclimate of Côte d'Ivoire, because there is clear evidence that: (1) both current and lagged rainfall influence yields; (2) the effect of rainfall on yields is often nonlinear, with exceptional events having a role; and (3) rainfall patterns and their effects on yield are very different in forest and savannah regions (see Amanor (1994), Hopkins (1973), Nicholson (1980), Sanders, Shapiro and Ramaswamy (1996)). A more parsimonious specification that includes no interactions between the rainfall variables and the savanah indicator produces results similar to those reported in tables 4 to 7.

4.2.1 Unconstrained tests: rainfall and demand

In table 3, we present in panel A an F-test of the joint significance of all of the rainfall variables in equation (14), where we control non-parametrically for expenditures. In panel B we present the F-test for the significance of the rainfall variables in its parametric counterpart, equation (17).

We use as dependent variables expenditures grouped in broad categories (columns (1) to (5)), and more detailed expenditures on particular goods (columns (6) to (11)).The results obtained in the non-parametric and the parametric specifications are extremely similar: this is not surprising, since the relationship between expenditures on each type of goods and total expenditures seem close to being log linear: in figures 1 to 3, we show the semi-parametric estimates of these relationships, along with a bootstrapped confidence interval.

In all of these regressions except total food consumption and education, rainfall variables are jointly significantly different from zero: rainfall affects the allocation of expenditures across goods, over and above its effect on total expenditures. Hence, conditional on our maintained auxiliary assumptions, we reject the hypothesis that households in Côte d'Ivoire achieve a Pareto efficient allocation of resources.

These auxiliary assumptions are neither innocuous nor self-evident. There are three that are of particular concern. First, these regressions condition on region $\times$ year interactions, but not for relative prices faced by each specific household. The fact that rainfall variation does not affect total food consumption, but does affect the consumption of other goods (adult goods, clothing, and prestige goods) whose prices are not very likely to substantially vary with rainfall pattern suggests that this result is not entirely due to relative price effects. However, in tables 5 and 6 below, we examine further the possibility that within regions, local relative prices vary with rainfall. We find little evidence to suggest that this is so, or that our results are substantially affected by local variation in price.

Second, we maintain the assumption that either labor (and other input) markets operate smoothly or that leisure is separable from other consumption in preferences. If labor demand moves with rain and labor markets are imperfect, and leisure is *not* separable from other consumption, then in an efficient allocation rainfall could influence commodity demand even conditional on total expenditure. For example, if there is an important nutrition-productivity effect, then the demand for (say) calories might vary with rainfall, conditional on total expenditure.

Third, we maintain the assumption that preferences are time-separable. If, contrary to the model in section 2.1, the marginal utility of the consumption of $\{c_{iw_t}, c_{pw_t}\}$ is affected by consumption in period $t-1$, then the recursivity that produces equation (8) fails, as do the testable restrictions. It is important to note that this is not a matter of imperfect financial markets: (8) remains valid with imperfect or absent credit markets. Instead, it arises when consumption is spread out over multiple years, as in a model with habit formation or with durable consumption when rental markets or markets for used products are absent.

We do not examine directly these later two auxiliary assumptions. Instead, in the following sections we examine a sequence of more detailed hypotheses drawn from the descriptive literature on household allocation processes in West Africa. As we do so, we evaluate as well the possible alternative hypothesis that the patterns we see are generated by violations of these auxiliary assumptions.

4.2.2 Constrained tests: predicted crop income and demand

To make this discussion more concrete and to explicitly link our specification to the anthropological literature we discussed above, we estimate equation (16). Table 4 presents the estimates of the relationships between expenditures on particular goods and predicted net income from male cash crops, yams, and from female crops, conditional on total expenditures. Panel A contains the results of the partial linear model (16), while panel B reports the results of the linear formulation (17). In each equation the individual coefficients are of interest, as is their joint significance.

For many goods, one or the other forms of predicted income is significantly related to consumption, conditional on total expenditure. Moreover, not only do the effects of predicted male and female income differ, but although men typically farm yams, the effect of predicted yam income often differs radically from that of income from the other male crops. Moreover, these differences accord with those described in the anthropological literature.

Predicted income from yam is, not surprisingly, associated with an increase in the consumption of staples (mostly yams), and a decrease in the consumption of vegetables. This could be a consequence of local relative price movements where markets are not well-integrated (but this interpretation is not supported in tables 5-6 below). More interestingly, yam income is negatively associated the purchase of adult goods (tobacco and alcohol), and especially with "prestige goods" (jewelry and adult clothing items such as "pagnes"). However, it is not the case that an increase in yam income is systematically associated with a decrease in expenditures on all goods except yams (conditional on total expenditure), which would be the case if yams were never sold. There is a positive relationship between predicted income from yams and education expenditures (schools fees, books, etc..), which are cash expenditures, while the coefficients in the purchased food and children's clothing equations are close to 0.

All of these results regarding the relationship between yam income and expenditures on particular goods are consistent with the idea that income from yam is associated with household public goods and basic necessities. This corresponds to Meillasoux' description of yams as an "appreciated good" under the control of the household head for redistribution in the household. Moreover, these effects are large. Conditional on total expenditure, a 10% increase in income from yam is associated with a 5% decline in expenditures on prestige goods, a 2% rise in expenditures on education, and an 8% decline in expenditures on tobacco and alcohol.

In contrast, controlling for total expenditure, predicted income from cash crops controlled by men has a negative effect on the purchase of food (even though this is a cash expenditure) and a positive effect on expenditure on prestige goods. It has a negative effect on education expenditures. Fluctuations in income from male cash crops are unrelated to changes in the consumption of particular food goods. Again, this corresponds to the descriptive evidence: the critical point is not so much that the income from these crops is in cash and thus readily available to buy things without additional transaction cost, but rather that shocks to this income are not positively associated with changes in food purchases. In much of West Africa, the male head of household is responsible for a "statutory contribution" to his wife to prepare meals, but after that generally fixed obligation is met, he "acts on his own account .... He contributes to, but is never solely responsible for, the total expenditure of the component hearth-hold(s)" (Ekejiuba (1995), pp. 52-53).[14] In neighboring Ghana, the 'chop money' provided by a husband to a wife for the preparation of meals is a regular, fixed amount that can be changed only after negotiations that often involve extended family members; when a husband does not meet this obligation it

[14] Ekejiuba uses the term 'hearth-hold' to mean a mother and her children.

can be an important source of friction within the household and between the extended families Goldstein (2000)).

Income derived from crops controlled by women, however, is positively associated with all types of food consumption with the exception of staples. In particular, when the predicted income from women's crops increases, the consumption of meat and (perhaps) processed foods increases. In contrast with income from yams, income from women's crops is negatively associated with expenditures on education, and positively associated with the consumption of adult goods. Like the income from male cash crops, it is strongly associated with expenditures on prestige goods. Again, the effects can be large: conditional on total expenditure, a 10% increase in income from women's crops is associated with a 5% increase in expenditures on prestige goods, a 2% increase in expenditures on meat and perhaps a 9% increase in consumption of alcohol and tobacco. This pattern, again, corresponds to the anthropological evidence discussed above: women control the production on their own farms, and "it is ultimately the woman's responsibility to feed everyone, whatever the amount she receives from her husband". [15] Women have access to a base contribution from their husbands for these purposes, but when their own disposable incomes increase, expenditure on these goods and on goods that they consume privately increases.

We raised the possibility that changes in local relative prices might bias these estimates. There is information available on prices for a wide range of goods at the CILSS cluster level for 3 of the 4 years of the survey. At the cost of a 1/3 reduction in sample size, a vector of prices can be incorporated into the model. Accordingly, in table 5 we report the results of estimating

$$\log(c_{i2})-\log(c_{i1}) = \alpha(\log(x_{i2}-\log(x_{i1}))+(X_{i2}-X_{i1})\delta+(P_{i2}-P_{i1})\delta_p+\sum_{s=1}^{S} DR_s\gamma_s+\nu_{i2}-\nu_{i1} \quad (26)$$

[15] This is a quote from Etienne (1980), describing the relationship between husband and wives in Cote d'Ivoire. Etienne (1980) describes how among Baule households in Côte d'Ivoire, "in the case of some essential subsistence products, production was entirely the responsibility of one or the other sex and the producer was the "owner" of the product or, in other words, controlled its distribution. In the case of other products, both sexes contributed to production, each being in charge of specific tasks or phases of the production process; the sex that was considered to have initiated the process and taken responsibility for it "owned" the product or controlled its distribution" (p. 219-220). In addition, see Guyer (1995) who describes how Senufo women in Côte d'Ivoire are responsible for the production of certain crops, and that they have control over the incomes from those crops

where P_{it} is a vector containing the prices of the 20 commodities in the cluster of household i (and 20 indicator variables for the 20% of cluster-years in which a price is missing).

Moving from table 4 to table 5, one first notes that standard errors have almost uniformly grown as a consequence of the fact that the sample size falls from approximately 850 to approximately 500. Overall, there is little change in the pattern of coefficient estimates. In no case does a coefficient estimate that is significantly different from zero at the five percent level in table 4 change signs in table 5. In several cases, the price variables are jointly significantly different from zero conditional on the region $\times$ year indicator variables, providing some evidence that local markets are not fully integrated into regional marketing systems. Turning to the determinants of prices themselves, we find in table 6 almost no commodity for which there is a statistically significant relationship between rainfall and price, conditional on the region $\times$ year effects. Only for peanut butter and for plastic sandals are the predicted income variables jointly significant: higher yam income is associated with a higher peanut butter price, while higher output of female goods is associated with a higher price for plastic sandals. We conclude that there is no evidence that rainfall-induced variations in local prices are driving the association we observe between changes in consumption patterns and shocks in the flows of different catagories of net crop income.

4.3 Overidentification restriction tests

4.3.1 Unconstrained reduced form tests

Panel C in table 3 presents the reduced form tests of Pareto efficiency, based on estimating equations 21 and 13 (with the 32 rainfall variables) and then testing jointly the hypothesis that all the ratios between the corresponding coefficients in the two reduced form equations are equal. We present in the first row a test that the rainfall coefficients are jointly significant. The rainfall variables are jointly significant in all regressions.

The second row presents the overidentification tests, which are the reduced form test of Pareto efficiency described earlier in equation 11. The overidentification test never rejects equality. However, this test is likely to have very low power, and we turn to restricted version in table 7.

4.3.2 Constrained reduced form tests

In column 1 of table 7, we present the result of regressing differences in the logarithm of total expenditures on predicted changes in the logarithm of income from yams, male cash crops, and female crops. The coefficients are all significant and not significantly different from each other, although the coefficient on predicted male non-yam income is smaller. The elasticity of total expenditure with respect to yam and female income are very similar (0.34 and 0.32 respectively), and close to the elasticities of expenditures to total income estimated by Townsend (1994) for India and Deaton (1997) for Cote d'Ivoire (using the same data set).

In the following columns, we present the coefficients of estimating equation 25 for the same outcomes reported in table 4. The final row of each panel presents the test of the overidentification restrictions. The overidentification restrictions are rejected at the five percent level in five cases: prestige goods, adult goods, staples, vegetables and purchased foods. Apparently, households do not allocate consumption expenditures in a manner consistent with Pareto efficiency. Moreover, an examination of the coefficient estimates reveals that the deviations from efficiency are consistent with the patterns shown in table 4 and with the anthropological accounts discussed above.

Variations in income from male non-yam crops and from female-controlled crops are much more strongly associated with the consumption of adult goods than are variations in yam income (for this comparison, as for all those that follow, it is understood that these statements are relative to the effects of these income flows on total expenditure). Precisely the same pattern is observed, but more strongly, for prestige goods. Income from yam, it seems, is associated with household public goods and basic necessities while income from the individually-controlled female and male cash crops is associated with expenditures on alcohol, tobacco, and the prestige goods. Consumption of staples and expenditures on education are much more strongly related to variations in yam income than to variations in male non-yam crop income or to female crop income. Interestingly, while increased yam production is directly associated with increases in household consumption of staples, variations in female crop income are much more strongly associated with purchases of staples than variations in yam or male non-yam income.[16] Consumption of vegetables is much more strongly related to female crop income than to either

[16] These results and those regarding vegetable purchases below are not shown, but are available from the authors.

yam or non-yam male crop income, but this is rather unsurprising because vegetable income is a component of female crop income. More interestingly, it is also the case that both overall consumption and purchases of vegetables are much more strongly related to income from female non-vegetable income than to yam or male non-yam income. Overall food purchases (and consumption of processed foods, albeit at a low level of statistical significance) are much more strongly associated with variations in income from female-controlled crops than income from yam or male non-yam crops. Again, this pattern corresponds to that shown in table 4 and with discussions of the role of 'chop money' in the descriptive accounts of household resource allocation in West Africa.

5 Conclusion

We have shown that expenditure patterns in households in Côte d'Ivoire are not consistent with a Pareto efficient allocation of household resources. Moreover, the deviations from Pareto efficiency that we document correspond closely to the descriptions of provisioning norms available in the literature. In particular, we find that rainfall shocks that increase the output of the "appreciated" crop, yam, are associated with strong shifts in the composition of expenditures towards education, staples, and overall food consumption and away from adult goods and "prestige" goods such as jewelry. In contrast, rainfall shocks that increase the output of crops cultivated individually by either men or women are associated with strong expenditure shifts toward adult and prestige goods. Shocks that increase the output of crops predominantly cultivated by women shift expenditures toward all types of food consumption (except staples), while similar shocks affecting cash crops cultivated by men have a negative effect on the purchases of food. The results are similar across two different types of test: a test based on exclusion restrictions that permits arbitrary nonlinearities in the expenditure elasticities of demand for different goods; and an overidentification test that is robust to measurement error in expenditures. In no case does it seem that these results are strongly influenced by changes in local relative prices.

There is a rich literature that describes norms of behavior that guide resource allocation in West African households. The pattern of actual household expenditures in Côte d'Ivoire corresponds to these norms. The important lesson from this finding is not (for example) that income from cash crops controlled by men is directed towards prestige goods. Rather it is the more general, albeit modest, claim that even when investigating such core economical topics as demand analysis, economists may have much to learn from the detailed observations available from neighboring disciplines. This is particularly so in a case such as that of intrahousehold resource allocation in West Africa, where the broad contours of the descriptions are at once so similar across many studies in a large number of local settings and so strongly inconsistent with the routine models available to applied economists.

The norms that are so prominent in the discussions of household provisioning in West Africa appear to have real consequences for the allocation of resources in Côte d'Ivoire. The immediate implication of these results is that the conventional unitary household model employed, for example, in the permanent income hypothesis is insufficiently rich to capture important aspects of demand behavior. Nor does the more general collective model provide an adequate framework for the interpretation of these results. A more radical departure from the conventional model is required. The results of this paper suggest two potentially complementary approaches. First, the interactions between household members might involve a bargaining situation characterized by important information asymmetries and/or significant enforcement problems. Second, resources flows might occur through discrete 'accounts' with imperfect fungability. The literature contains a great deal of descriptive material supportive of each of these interpretations.

Both of these considerations suggest that a wide range of household outcomes could respond to changes in the economic environment in ways that do not correspond to the predictions of simple collective models. Decisions regarding investment in children's human capital, production decisions, and the allocation of land and other productive assets could all be affected by inefficient intrahousehold negotiations and/or by constrained fungability of resources across uses. For example, the allocation of household labor to different uses could be influenced by and individuals' desire to guide income into her own hands, or into the appropriate 'account' for a needed expenditure, violating the separation of consumption and production decisions that occurs in the efficient household model. Similarly, intertemporal decisions such as the allocation of household resources into the human capital of children could be affected by the labeling of income if husbands and wives face different opportunities in financial markets.

References

Amanor, Kojo Sebastian (1994) *The New Frontier: Farmers' Response to Land Degradation, A West African Study* (London: Zed Books)

Bassett, Thomas J (1985) 'Food, peasantry and the state in northern ivory coast, 1898-1982.' Ph.D.

—— (1988) 'Breaking up the bottlenecks in food-crop and cotton cultivation in northern cote d'ivoire.' *Africa* 58(2), 147–173

Berger, Iris, and E. Frances White (1999) *Women in Sub-Saharan Africa* (Bloomington, IN: Indiana University)

Berry, Sara (1993) *No Condition is Permanent: The Social Dynamics of Agrarian Change in Sub-Saharan Africa* (Madison: University of Wisconsin Press)

Bigot, Yves (1979) 'Analyse techno-economique des systemes de production coton-cereale dans le nord-ouest de la cote d'ivoire.' Technical Report

Bouis, Howarth, and Lawrence Haddad (1992) 'Are estimates of calorie-income elasticities too high? a recalibration of the plausible range.' *Journal of Development Economics* 39(2), 333–364

Browning, Martin, and Pierre-Andre Chiappori (1998) 'Efficient intra-household allocations: A general characterization and empirical tests.' *Econometrica* 66(6), 1241–1278

Carney, J., and M. Watts (1991) 'Disciplining women: Rice, mechanization and the evolution of mandinka gender relations in senegambia.' *Signs* 64(4), 651–681

Davison, Jean (1988) 'Land and women's agricultural production: The context.' In *Agriculture, Women and Land: The African Experience,* ed. Jean Davison (Boulder, CO: Westview)

Deaton, Angus (1997) *The Analysis of Household Surveys* (World Bank, International Bank for Reconstruction and Development)

Dey, Jennie (1993) 'Gender asymmetries in intrahousehold resource allocation in sub-saharan africa: Some policy implications for land and labor productivity.' In *Intrahousehold resource allocation in developing countries : models, methods, and policy,* ed. Lawrence Haddad, John Hoddinott, and Harold Alderman (Baltimore: Johns Hopkins University Press)

Doss, Cheryl (2001) 'Men's crops? women's crops? the gender pattern of cropping in ghana'

Doss, Cheryl R. (1998) 'Is risk fully pooled within the household? Evidence from Ghana.' Mimeo, Williams College

Ekejiuba, Felicia (1995) 'Down to fundamentals: Women-centered hearth-holds in rural west africa.' In *Women Wielding the Hoe: Lessons from Rural Africa for Feminist Theory and Development Practice,* ed. Deborah Fahy Bryceson (Oxford: Berg Publishers) pp. 47–62

Etienne, Mona (1980) 'Women and men, cloth and colonization: The transformation of production-distribution relations among the baule (ivory coast).' In *Women and Colonization: Anthropological Perspectives,* ed. Mona Etienne and Eleanor Leacock (New York: Praeger) pp. 214–238

Fan, Jianqing (1992) 'Design-adaptive nonparametric regression.' *Journal of the American Statistical Association* 87, 998–1004

Gastellu, Jean-Marc (1987) 'Matrilineages, economic groups and differentiation in west africa: A note.' *Development and Change* 18, 271–280

Goldstein, Markus Paul (2000) 'Intrahousehold allocation and farming in southern ghana.' PhD dissertation

Guyer, J. I., and P. E. Peters (1987) 'Conceptualizing the household - issues of theory and policy in Africa - introduction.' *Development and Change* 18(2), 197–214

Guyer, Jane (1984) 'Family and farm in southern cameroon.' 15

—— (1995) 'Women in the rural economy: Contemporary variations.' In *African Women South of the Sahara,* ed. Margaret Jean Hay and Sharon Stichter, ii ed. (New York: John Wiley and Sons) pp. 23–43

Guyer, Jane I. (1987) 'Food, cocoa, and the division of labour by sex in two West African societies.' *Development and Change* 18(2), 197–214

Haddad, Lawrence, and John Hoddinott (1994) 'Women's income and boy-girl anthropometric status in the Cote d'Ivoire.' *World Development* 22, 543–553

Hausman, Jerry A., and Whitney Newey (1995) 'Nonparametric estimation of exact consumers surplus and deadweight loss.' *Econometrica* 63(6), 1445–1476

Hopkins, A.G. (1973) *An Economic History of West Africa* (London: Longman)

Jones, Christine (1986) 'Intra-household bargaining in response to the introduction of new crops: A case study from north Cameroon.' In *Understanding Africa's Rural Households and Farming Systems,* ed. Joyce Lewinger Moock (Boulder: Westview Press)

Kozel, Valerie (1990) 'The composition and distribution of income in Cote d'Ivoire.' Technical Report

Meillassoux, Claude (1965) *Anthropologie Economique des Gouros de Cote d'Ivoire* (Paris: F. Maspero)

——— (1975) *Femmes, greniers et capitaux* (Paris: F. Maspero)

Nicholson, Sharon E. (1980) 'The nature of rainfall fluctuations in subtropical west africa.' *Monthly Weather Review* 108, 473–487

Porter, Jack (1996) 'Essays in econometrics.' PhD dissertation, MIT

Robinson, Peter (1988) 'Root-n consistent semiparametric regressions.' *Econometrica* 56, 931–954

Rosenzweig, Mark R., and Kenneth I. Wolpin (2000) 'Natural natural experiments in economics.' *Journal of Economic Literature* 38(4), 827–74

Saito, Katrine, Hailu Mekonnen, and Daphne Spurling (1994) 'Raising the productivity of women farmers in sub-saharan africa'

Sanders, John, Barry Shapiro, and Sunder Ramaswamy (1996) *The Economics of Agricultural Technology in Semiarid Sub-Saharan Africa* (Baltimore: Johns Hopkins University)

Shipton, Parker (1989) *Bitter Money: The Classification of Forbidden Commodities among the Luo of Kenya* (Washington D.C.: American Anthropological Association)

Thaler, Richard (1990) *The Winner's Curse*

Townsend, Robert (1994) 'Risk and insurance in village India.' *Econometrica* 62(4), 539–591

Vijverberg, Wim (1988) 'Nonagricultural family enterprises in cote d'ivoire.' Technical Report

von Braun, Joachim, and Patrick J. R. Webb (1989) 'The impact of new crop technology on the agricultural division of labor in a west african setting.' *Economic Development and Cultural Change* 37(3), 513–34

Weekes-Vagliani, Winifred (1985) 'Actors and institutions in the food chain: The case of the ivory coast.' Technical Report, OECD Development Centre

——— (1990) 'Analyse des variables socioculturelles et de l'adjustement en cote d'ivoire.' Documents Techniques no. 9, Centre de Developpement de l'OCDE, January

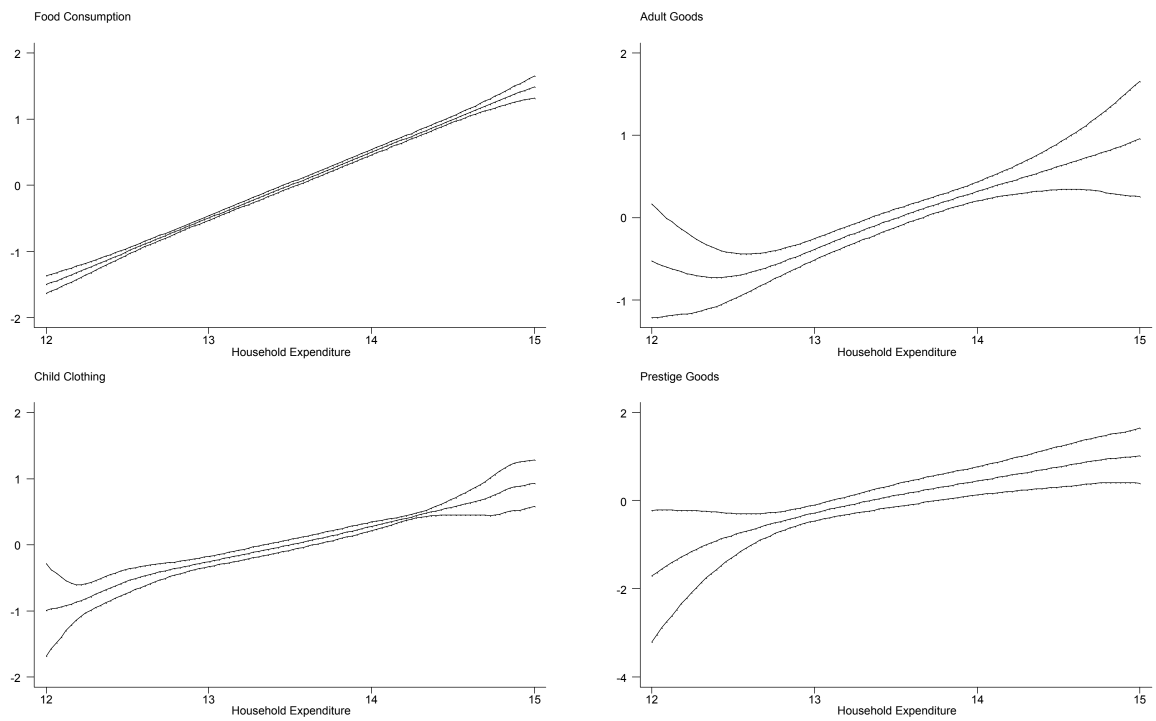

+/- 2 std. dev. pointwise confidence bands

Fig. 1: Partial Linear Expenditure Functions

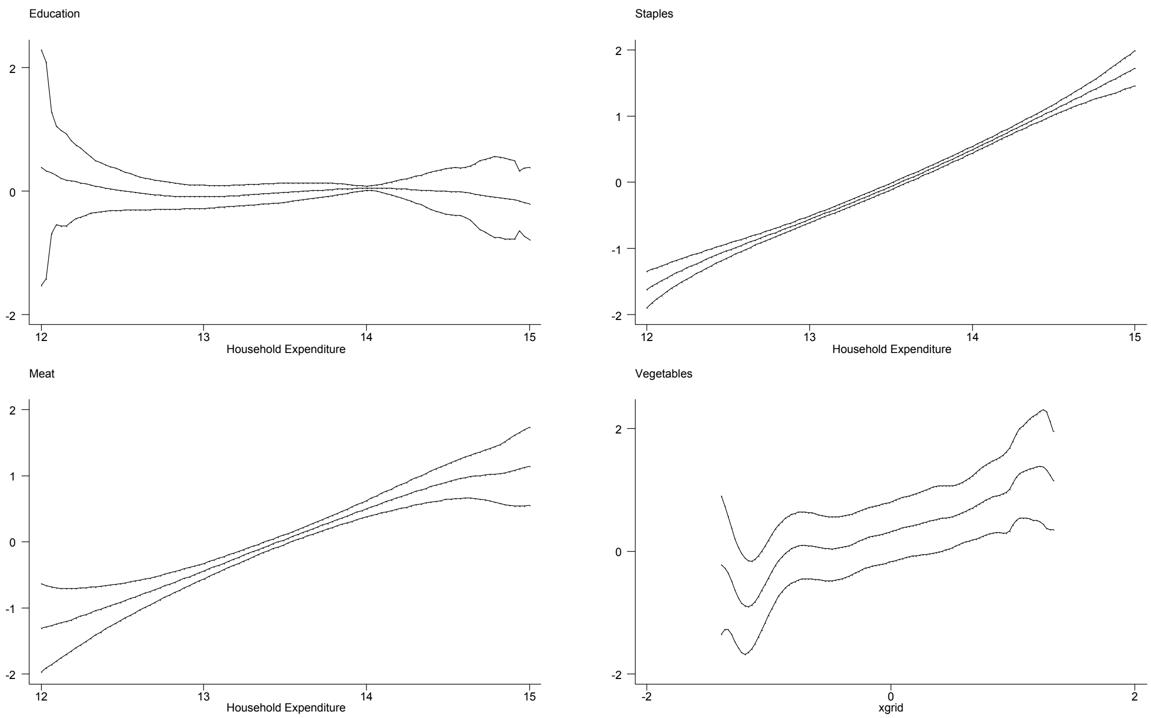

+/- 2 std. dev. pointwise confidence bands

Fig. 2: Partial Linear Expenditure Functions

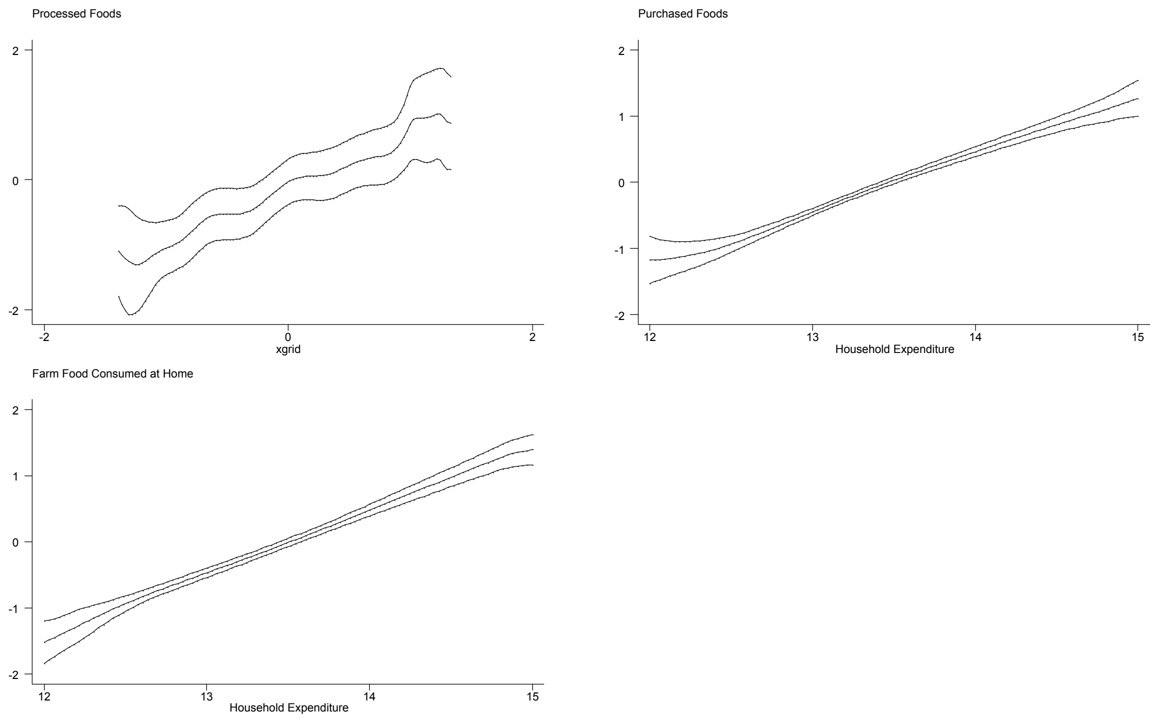

+/- 2 std. dev. pointwise confidence bands

Fig. 3: Partial Linear Expenditure Functions

TABLE 1: Descriptive statistics

	Mean, year 1 1000's FCFA	Difference in logs Year2-Year1	Observations
	Standard errors in parantheses		
	(1)	(2)	(3)
Income from male crops	583.87	-0.03	924
	(28.64)	(.04)	
Income from female crops	111.09	-0.12	924
	(10.25)	(0.04)	
Total "male income"	624.75	-0.04	924
	(28.55)	(0.04)	
Total "female income"	151.97	-0.05	924
	(11.16)	(0.04)	
Unattributed income	146.75	-0.02	924
	(6.57)	(0.05)	
Total expenditure	1146.77	-0.09	911
	(29.97)	(0.02)	
Food consumption	599.56	-0.06	841
	(11.41)	(0.02)	
Adult goods	46.01	-0.30	853
	(2.44)	(0.14)	
clothing	18.55	0.09	866
	(0.75)	(0.10)	
prestige goods	60.17	-0.12	924
	(1.94)	(0.06)	
Staples	430.06	-0.02	924
	(12.09)	(0.02)	
Meat	120.73	-0.13	919
	(4.61)	(0.03)	
Vegetables	48.60	0.00	899
	(1.72)	(0.03)	
Processed foods	56.74	-0.23	917
	(56.40)	(0.04)	
All purchased foods	267.54	-0.17	918
	(8.67)	(0.03)	
All food consumed at home	349.77	0.03	838
	(6.94)	(0.03)	

Table 2: First stage summary statistics

	Dependent variables		
	Female crop income	Male cash crop	Yam income
	(1)	(2)	(3)
F statistics (p value)			
All rainfall variables are significant	2.38 (0.000)	1.99 (0.014)	3.45 (0.000)
Current year rainfall variables significant	2.26 (0.008)	1.18 (0.315)	3.38 (0.000)
Past year rainfall variables significant	2.53 (0.001)	2.79 (0.005)	4.23 (0.000)
Rainfall variables significantly different from:			
Female crop income:	NA		
Male cash crop income	2.02 (0.012)	NA	
Yam income	2.36 (0.003)	1.96 (0.019)	NA

Table 3: unconstrained overid and exlusion restrictions tests

	Food consumption	Adult goods	Clothing	Prestige goods	Education	Staples	Meat	Vegetables	Processed foods	Purchased foods	Food consumed at home
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
PANEL A: EXCLUSION RESTRICTION TEST: SEMI PARAMETRIC FORMULATION											
F test (pvalue)	0.83	1.89	1.45	2.30	1.18	1.82	1.65	2.28	2.06	3.88	1.95
Rainfall variables jointly significant	(0.713)	(0.004)	(0.065)	(0.000)	(0.245)	(0.006)	(0.019)	(0.000)	(0.001)	(0.000)	(0.003)
PANEL B: EXCLUSION RESTRICTION TEST: LINEAR FORMULATION											
F test (pvalue)	0.75	1.88	1.45	2.17	1.31	1.94	1.55	2.27	1.91	3.68	1.94
Rainfall variables jointly significant	(0.822)	(0.004)	(0.065)	(0.000)	(0.142)	(0.003)	(0.036)	(0.000)	(0.003)	(0.000)	(0.003)
PANEL C: OVERIDENTIFICATION RESTRICTIONS											
F test (p value)	3.57	1.60	2.63	2.15	1.26	3.96	2.17	3.17	2.98	5.05	3.02
Rainfall variables jointly significant	(0.000)	(0.038)	(0.000)	(0.001)	(0.189)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
F test (pvalue)	0.39	0.69	1.11	0.89	0.40	0.65	0.42	0.82	0.88	0.92	0.80
Test of overidentification restrictions	(0.997)	(0.876)	(0.317)	(0.618)	(0.996)	(0.915)	(0.996)	(0.726)	(0.639)	(0.589)	(0.746)

Table 4: Restricted exlusion restriction tests

	Food consumption	Adult goods	Clothing	Prestige goods	Education	Staples	Meat	Vegetables	Processed foods	Purchased foods	Food consumed at home
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
PANEL A: RESTRICTED EXCLUSION RESTRICTION TEST: SEMI PARAMETRIC FORMULATION											
Prediced male non-yam	-0.037	0.178	0.112	0.550	-0.139	0.015	-0.053	-0.054	0.004	-0.176	0.068
income	(0.029)	(0.464)	(0.267)	(0.233)	(0.099)	(0.077)	(0.090)	(0.142)	(0.131)	(0.090)	(0.133)
Predicted yam	0.047	-0.705	0.094	-0.491	0.212	0.142	-0.093	-0.167	-0.005	-0.018	0.100
income	(0.032)	(0.588)	(0.282)	(0.155)	(0.136)	(0.061)	(0.073)	(0.097)	(0.110)	(0.071)	(0.073)
Predicted female	-0.006	0.845	0.214	0.534	-0.210	-0.117	0.195	0.574	0.266	0.127	-0.013
income	(0.034)	(0.623)	(0.370)	(0.192)	(0.130)	(0.080)	(0.103)	(0.144)	(0.164)	(0.135)	(0.104)
F tests (pvalue) :	1.147	0.837	0.203	7.057	1.895	2.696	2.880	5.640	1.014	1.803	0.952
Predicted income variables jointly significant	(0.339)	(0.479)	(0.894)	(0.000)	(0.143)	(0.055)	(0.044)	(0.002)	(0.393)	(0.157)	(0.422)
Predicted income variables	1.711	1.252	0.041	10.584	2.635	3.871	4.280	8.229	1.055	1.790	0.630
significantly different	(0.190)	(0.294)	(0.960)	(0.000)	(0.082)	(0.027)	(0.019)	(0.001)	(0.355)	(0.177)	(0.537)
Coefficient of female crops	1.268	2.501	0.054	17.596	4.059	7.066	8.440	15.467	2.092	1.180	0.996
and yam income equal.	(0.265)	(0.120)	(0.818)	(0.000)	(0.049)	(0.010)	(0.005)	(0.000)	(0.154)	(0.282)	(0.323)
PANEL B: RESTRICTED EXCLUSION RESTRICTION TEST: LINEAR FORMULATION											
Prediced male non-yam	-0.045	0.250	-0.008	0.509	-0.177	-0.007	-0.020	-0.031	0.007	-0.160	0.032
income	(0.029)	(0.445)	(0.256)	(0.204)	(0.096)	(0.074)	(0.101)	(0.147)	(0.126)	(0.086)	(0.134)
Predicted yam	0.051	-0.878	0.035	-0.458	0.235	0.151	-0.054	-0.194	-0.059	-0.046	0.111
income	(0.032)	(0.508)	(0.273)	(0.157)	(0.122)	(0.058)	(0.067)	(0.085)	(0.110)	(0.071)	(0.069)
Predicted female	-0.005	0.971	0.196	0.504	-0.232	-0.114	0.167	0.613	0.246	0.127	-0.018
income	(0.034)	(0.607)	(0.376)	(0.178)	(0.137)	(0.078)	(0.106)	(0.146)	(0.155)	(0.127)	(0.109)
F tests (pvalue) :	1.472	1.427	0.150	6.456	2.505	3.618	1.784	7.197	1.081	2.562	1.218
Predicted income variables jointly significant	(0.232)	(0.245)	(0.929)	(0.001)	(0.070)	(0.019)	(0.161)	(0.000)	(0.365)	(0.064)	(0.312)
Predicted income variables	2.207	2.130	0.189	9.415	3.484	5.194	2.676	10.792	1.466	1.990	0.720
significantly different	(0.120)	(0.128)	(0.828)	(0.000)	(0.038)	(0.009)	(0.078)	(0.000)	(0.240)	(0.146)	(0.490)
Coefficient of female crops	1.329	4.215	0.095	15.316	5.602	8.988	5.161	21.465	2.932	1.857	1.416
and yam income equal.	(0.254)	(0.045)	(0.759)	(0.000)	(0.022)	(0.004)	(0.027)	(0.000)	(0.092)	(0.178)	(0.239)

Table 5: Restricted exlusion restrictions tests with Prices: Linear Formulation

	Food consumption	Adult goods	Clothing	Prestige goods	Education	Staples	Meat	Vegetables	Processed foods	Purchased foods	Food consumed at home
	1.000	2.000	3.000	4.000	5.000	6.000	7.000	8.000	9.000	10.000	11.000
Prediced male non-yam	-0.12	0.12	0.42	0.16	0.34	0.26	-0.11	-0.13	0.19	-0.01	-0.05
income	(0.077)	(0.406)	(0.276)	(0.554)	(0.650)	(0.142)	(0.252)	(0.281)	(0.259)	(0.160)	(0.162)
Predicted yam	0.13	-0.37	-0.35	-0.99	0.00	0.25	-0.27	-0.40	-0.40	-0.24	0.42
income	(0.067)	(0.365)	(0.243)	(0.454)	(0.580)	(0.116)	(0.206)	(0.235)	(0.213)	(0.132)	(0.140)
Predicted female	-0.02	0.17	-0.04	0.65	-0.29	-0.18	-0.08	0.46	-0.28	-0.21	-0.15
income	(0.116)	(0.590)	(0.426)	(0.765)	(0.911)	(0.196)	(0.347)	(0.385)	(0.358)	(0.221)	(0.240)
F tests (pvalue) :											
Predicted income variables	2.07	0.39	1.41	1.77	0.11	2.76	0.69	1.37	1.67	1.55	3.19
jointly significant	(0.103)	(0.759)	(0.240)	(0.153)	(0.956)	(0.042)	(0.556)	(0.251)	(0.172)	(0.200)	(0.024)
Predicted income variables	3.08	0.45	2.11	2.17	0.15	1.91	0.18	1.67	1.65	0.67	3.63
significantly different	(0.047)	(0.637)	(0.123)	(0.115)	(0.857)	(0.149)	(0.836)	(0.190)	(0.193)	(0.511)	(0.028)
Coefficient of female crops	1.24	0.44	0.39	3.17	0.06	3.30	0.22	3.32	0.08	0.02	4.32
and yam income equal.	(0.267)	(0.509)	(0.535)	(0.076)	(0.813)	(0.070)	(0.637)	(0.069)	(0.779)	(0.889)	(0.038)
Price variables											
jointly significant (p)	0.000	0.130	0.040	0.000	0.700	0.060	0.087	0.260	0.000	0.000	0.000

Tables 6: Relationship between predicted income shocks and local prices

	Dependent variable: log(item price)									
	beef	imported rice	local rice	onion	salt	tomato paste	peanut butter	palm oil	local maize	local millet
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
predicted male non-yam income	0.15	-0.17	-0.30	0.36	-0.20	-0.04	-0.04	0.39	0.20	-0.03
	(0.130)	(0.130)	(0.130)	(0.130)	(0.130)	(0.130)	(0.130)	(0.130)	(0.130)	(0.130)
predicted yam income	0.08	-0.03	-0.12	0.19	-0.03	0.04	0.42	0.34	-0.04	0.13
	(0.130)	(0.130)	(0.130)	(0.130)	(0.130)	(0.130)	(0.130)	(0.130)	(0.130)	(0.130)
predicted female income	0.15	-0.17	-0.30	0.36	-0.20	-0.04	-0.04	0.39	0.20	-0.03
	(0.130)	(0.130)	(0.130)	(0.130)	(0.130)	(0.130)	(0.130)	(0.130)	(0.130)	(0.130)
F statistics: predicted income variables jointly significant	0.46	0.48	1.88	1.83	0.66	1.78	3.83	1.7	1.08	0.86

	cassava	yams	plantain	oil palm nuts	peanuts	eggs	cloth	fish	sandals	enamel bowl
	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)
predicted male non-yam income	-0.23	0.28	-0.06	-0.12	0.38	-0.14	0.56	0.17	0.05	0.28
	(0.253)	(0.184)	(0.436)	(0.230)	(0.198)	(0.192)	(0.274)	(0.146)	(0.107)	(0.324)
predicted yam income	-0.18	0.12	-0.30	-0.08	0.22	0.01	-0.13	-0.02	0.16	0.43
	(0.185)	(0.176)	(0.337)	(0.181)	(0.185)	(0.133)	(0.207)	(0.126)	(0.130)	(0.344)
predicted female income	-0.33	-0.03	-0.58	0.15	-0.02	0.39	0.25	-0.11	0.61	0.29
	(0.335)	(0.406)	(0.463)	(0.197)	(0.222)	(0.248)	(0.273)	(0.163)	(0.158)	(0.314)
F statistics: predicted income variables jointly significant	0.82	0.81	1.05	0.36	1.24	0.84	1.64	1.1	5.33	0.63

Table 7: Restricted overid tests

	Total expenditure	Food consumption	Adult goods	Clothing	Prestige goods	Education	Staples	Meat	Vegetables	Processed foods	Purchased foods	Food consumed at home
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Predicted male non-yam	0.125	0.053	0.428	0.127	0.607	-0.131	0.124	0.101	0.043	0.110	-0.047	0.158
income	(0.052)	(0.057)	(0.460)	(0.330)	(0.195)	(0.134)	(0.077)	(0.105)	(0.111)	(0.120)	(0.083)	(0.093)
Predicted yam	0.335	0.351	-0.328	0.430	-0.196	0.359	0.501	0.262	0.003	0.208	0.255	0.435
income	(0.043)	(0.048)	(0.387)	(0.276)	(0.163)	(0.127)	(0.064)	(0.088)	(0.091)	(0.100)	(0.070)	(0.078)
Predicted female	0.316	0.251	1.399	0.551	0.751	-0.052	0.215	0.464	0.801	0.498	0.410	0.257
income	(0.054)	(0.060)	(0.488)	(0.348)	(0.204)	(0.154)	(0.080)	(0.110)	(0.114)	(0.125)	(0.087)	(0.097)
F tests (pvalue) :		1.647	3.664	0.088	6.961	2.501	5.483	1.576	10.471	1.900	3.311	0.657
Overidentification		(0.193)	(0.026)	(0.915)	(0.001)	(0.083)	(0.004)	(0.207)	(0.000)	(0.150)	(0.037)	(0.519)
Restriction test												

Appendix Table A1: First stage regression results

	Dependent variables					
	Female crop income		Male cash crop		Yam income	
	Forest coefficients	Savanah interaction	Forest coefficients	Savanah interaction	Forest coefficients	Savanah interaction
	(1)	(2)	(3)	(4)	(5)	(6)
Difference (year 2 - year 1) in:						
Aggregate rainfall current year, season 1	-0.0014814	0.0042065	0.0004811		-0.003153	-0.010761
	(0.001)	(0.003)	(0.002)		(0.002)	(0.006)
Aggregate rainfall current year, season 2	0.0007278	0.0012071	-0.001099		0.0015603	0.0015827
	(0.000)	(0.002)	(0.001)		(0.001)	(0.004)
Aggregate rainfall current year, season 3	-0.0006083	0.0039253	0.0001552		-0.002321	-0.003099
	(0.001)	(0.001)	(0.001)		(0.001)	(0.002)
Aggregate rainfall current year, season 4	0.000697	-0.004462	-0.000169		0.0005378	-0.006442
	(0.001)	(0.005)	(0.001)		(0.001)	(0.006)
Aggregate rainfall past year, season 1	-0.0003562	0.007357	-0.004016		-0.00618	-0.010605
	(0.002)	(0.008)	(0.003)		(0.003)	(0.011)
Aggregate rainfall past year, season 2	0.0000884	-0.007043	0.0008669		0.0023795	-0.000265
	(0.000)	(0.005)	(0.001)		(0.001)	(0.006)
Aggregate rainfall past year, season 3	-0.001364	0.0033784	-9.57E-05		-0.00226	0.0027378
	(0.001)	(0.001)	(0.001)		(0.001)	(0.001)
Aggregate rainfall past year, season 4	-0.0007839	-0.003545	0.0014161		0.0007269	0.0053683
	(0.001)	(0.005)	(0.001)		(0.001)	(0.006)
Dummy for shock, current year, season 1	-0.4749124		-0.093278		-0.894238	0
	(0.233)		(0.364)		(0.439)	(0.000)
Dummy for shock, current year, season 2	0.4609491	0.3619415	0.127267		0.4623188	-2.75326
	(0.193)	(0.486)	(0.300)		(0.283)	(0.828)
Dummy for shock, current year, season 3			0		0	0
			(0.000)		(0.000)	(0.000)
Dummy for shock, current year, season 4	-0.4170426		0.6722299		-2.134331	0
	(0.378)		(0.654)		(0.520)	(0.000)
Dummy for shock, past year, season 1	0.2248324	-0.017546	0.1107528		0.2016122	3.537023
	(0.208)	(0.531)	(0.362)		(0.262)	(1.107)
Dummy for shock, past year, season 2	-0.0740645	0.1357927	-0.037784		-0.133787	-2.962664
	(0.119)	(0.429)	(0.183)		(0.204)	(0.911)
Dummy for shock, past year, season 3	-0.310151	0.5810194	-1.324416		-0.124188	-3.387585
	(0.245)	(1.027)	(0.384)		(0.386)	(1.389)
Dummy for shock, past year, season 4	-0.7240219	0.4450545	0.7792504		-1.748257	1.238107
	(0.267)	(1.367)	(0.437)		(0.408)	(1.275)
Number of observations	976		614		607	

Note: the specifications also include year dummies, region dummies, and their interactions