

NBER WORKING PAPER SERIES

INTRAHOUSEHOLD RESOURCE ALLOCATION
IN CÔTE D'IVOIRE:
SOCIAL NORMS, SEPARATE ACCOUNTS
AND CONSUMPTION CHOICES

Esther Duflo
Christopher Udry

Working Paper 10498
<http://www.nber.org/papers/w10498>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2004

We thank Mike Gough, Shawn Cole and Jonathan Robinson for excellent research assistance and the John D. and Catherine MacArthur Foundation, the National Science Foundation (SES-0079115) and the National Institute of Health (R01-HD39922-01) for financial support. We thank members of the MacArthur Network on the Effects of Inequality on Economic Performance, Daron Acemoglu, Abhijit Banerjee, Michael Booser, Anne Case, Pierre-André Chiappori, Angus Deaton, Bo Honoré, Shakeeb Khan, Ethan Ligon, Emmanuel Saez, and numerous seminar participants for helpful comments. The views expressed herein are those of the author(s) and not necessarily those of the National Bureau of Economic Research.

©2004 by Esther Duflo and Christopher Udry. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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

Esther Duflo and Christopher Udry

NBER Working Paper No. 10498

May 2004

JEL No. O12, D13

ABSTRACT

We study resource allocation within households in Côte d'Ivoire. In Côte d'Ivoire, as in much of Africa, husbands and wives farm separate plots, and there is some specialization by gender in the crops that are grown. These different crops are differentially sensitive to particular kinds of rainfall shocks. We find that conditional on overall levels of expenditure, the composition of household expenditure is sensitive to the gender of the recipient of a rainfall shock. For example, rainfall shocks associated with high yields of women's crops shift expenditure towards food. Strong social norms constrain the use of profits from yam cultivation, which is carried out almost exclusively by men. In line with these norms, we find that rainfall-induced fluctuations in income from yams are transmitted to expenditures on education and food, not to expenditures on private goods (like alcohol and tobacco). We reject the hypothesis of complete insurance within households, even with respect to publicly observable weather shocks. Different sources of income are allocated to different uses depending upon both the identity of the income earner and upon the origin of the income.

Esther Duflo

Department of Economics

MIT, E52-252G

50 Memorial Drive

Cambridge, MA 02142

and NBER

eduflo@mit.edu

Christopher Udry

Yale University

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 much domestic work she devotes to him (...) Beti wives remain farmers throughout their lives. Before the recent expansion of food sales they used to 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’s 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 (1980), pp. 369-70)

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 from 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 Côte d’Ivoire, the Gouro, studied by Meillassoux (1965) draw a sharp distinction between “appreciated products” (e.g., 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 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.” (Meillassoux (1965), 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 season-to-season basis as a function of the realization of her income. However, 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 (1992)): Money placed by individuals in different “mental accounts” is not fully fungible.

This paper seeks to test the empirical relevance of these descriptions in the context of rural Côte d’Ivoire. Efforts to empirically validate the mental accounts framework in behavioral

¹Many descriptions imply that the separation of these “accounts” is not limited to the uses of crop proceeds, but extend to income from non-farm 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)).

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.

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 Côte 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 undo 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 Côte 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

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

of household provisioning in Côte 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 determine 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 examine broad expenditure aggregates and more detailed expenditures on types of food. 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. In particular, we find that rainfall shocks that increase the output of the “appreciated” crop, yams, 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 no effect on the purchases of food.

Our results do not seem to be explained by obvious alternative explanations, such as misspecification of the demand functions, lack of separability between labor and consumption, price effects, or lack of time separability of preferences. Moreover, because we are testing whether the household pools *observed* risk, these results do not have a straightforward explanation in the framework of simple models of imperfect information or moral hazard. They are consistent,

however, with models of informal insurance without commitment,³ where a household member who faces a favorable shock needs to be partially compensated in order to agree to remain part of the insurance arrangement, and therefore where insurance can only be partial. The fact that shocks to yam income are transmitted to expenditures on food and education, despite the fact that yam income is formally under the control of the male household head is consistent with the fact that there are social sanctions associated with “mis-use” of these proceeds, and therefore there is little temptation to deviate from the pooling of yam income.

The evidence presented in this paper supports the validity and the empirical relevance of the descriptions of separate accounts within households in Côte 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 Côte 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?

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

³For an application to the household, see Coate and Ravallion (1993); Kocherlakota (1996); Ligon, Thomas and Worrall (2002); and Ligon (2003).

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 cultivates a farm using labor (L_i) that can be traded on a competitive market at wage w .⁴ 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 measures (it will be seen that this is simple to generalize) 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.

After rainfall realization r , each individual $i \in \{m, f\}$ consumes a vector of private goods c_i . Individual i 's preferences are summarized by the expected utility function $Eu_i(c_i)$, where expectations are taken over potential realizations of rainfall. 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. 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.

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

$$\max_{c_i, L_i} Eu_f(c_f) + \lambda Eu_m(c_m) \tag{1}$$

subject to

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

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 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).

⁴It 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.

This problem is separable and equivalent to

$$\max_{c_i} Eu_f(c_f) + \lambda Eu_m(c_m) \quad (2)$$

subject to

$$p \cdot (c_m + c_f) \leq \pi_f^*(r) + \pi_m^*(r) \quad (3)$$

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 $x = p \cdot (c_m + c_f)$, we have:

$$c_i = c_i(\lambda, p, x) \quad (4)$$

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

Equation (4) has the 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.⁵ In other words, for any i in m, f and j in 1, 2 and any good k :

$$\frac{dc_i^k}{dr_j} = \frac{\partial c_i^k}{\partial x} * \frac{\partial x}{\partial r_j}. \quad (5)$$

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

$$\frac{\frac{dc_i^k}{dr_1}}{\frac{\partial x}{\partial r_1}} = \frac{\frac{dc_i^k}{dr_2}}{\frac{\partial x}{\partial r_2}}. \quad (6)$$

⁵To 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.

There is an analogous test for $c_m^k + c_f^k$ if only aggregate consumption of the private good is observed.⁶

The essential element of the restriction (6) 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.

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 (6). If labor markets are imperfect, 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 $w_t \equiv \{r_1, r_2, \dots, r_t\}$ individual i consumes a vector of goods c_{iw_t} . The expected utility of individual i is $E \sum_t \beta_i^t U_i(c_{iw_t})$. 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

$$p \cdot (c_{fw_t} + c_{mw_t}) + A_{w_t} \leq RA_{w_{t-1}} + \pi_{mw_t}^* + \pi_{fw_t}^*, \quad (7)$$

where A_{w_t} is the amount invested after history w_t by the household in a safe asset that earns a return R .⁷

⁶Data on rainfall and expenditures are required to estimate (4) or (6). 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.

⁷It 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

There is a budget constraint for each history of rainfall realization, 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:

$$\max_{\{c_{iw_t}\}} E \sum_t \beta_f^t U_f(c_{fw_t}) + \lambda E \sum_t \beta_m^t U_m(c_{mw_t}) \quad (8)$$

for some value of λ , subject to (7) 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

$$\max_{c_{iw_t}, A_{w_t}} EU_f(c_{fw_t}) + \lambda EU_m(c_{mw_t}) + V_{w_t}(A_{w_t}; \lambda) \quad (9)$$

subject to

$$p \cdot (c_{fw_t} + c_{mw_t}) + A_{w_t} \leq RA_{w_{t-1}} + \pi_{mw_t}^* + \pi_{fw_t}^*. \quad (10)$$

The function $V_{w_t}(\cdot)$ 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 $\{c_{iw_t}\}$ and A_{w_t} . Let $A_{w_t}^*$ be the efficient level of assets held after w_t . Then efficient consumption is

$$\{c_{iw_t}^*\} = \arg \max_{c_{iw_t}} EU_f(c_{fw_t}) + \lambda EU_m(c_{mw_t})$$

subject to

$$p \cdot (c_{fw_t} + c_{mw_t}) \leq RA_{w_{t-1}} + \pi_{mw_t}^* + \pi_{fw_t}^* - A_{w_t}^*.$$

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

$$c_{iw_t} = c_i(\lambda, p, x_{w_t}) \quad (11)$$

for $i \in \{f, m\}$. 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 (4) holds in the dynamic setting. It can now be

 allocation of current consumption only through the function $V_{w_t}(A_{w_t})$ in equation (9).

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 testable restriction that

$$\frac{\frac{\partial c_{iw_t}^k}{\partial r_1}}{\frac{\partial x_{w_t}}{\partial r_1}} = \frac{\frac{\partial c_{iw_t}^k}{\partial r_2}}{\frac{\partial x_{w_t}}{\partial r_2}} \quad (12)$$

must hold for any consumption good k .

Equation (12) 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. In the empirical work below, we present both this test and a particular linear combination of these restrictions that is more straightforward to interpret.

2.2 Empirical Implementation

In order to implement tests based on (12) it is necessary to make some assumptions regarding the form of the commodity demand function (11). In particular, we assume that the demand for a particular commodity c by household i in period t are of the form

$$\log(c_{it}) = \alpha \log(x_{it}) + f(\lambda_i) + X_{it}\delta + v_i + \nu_{it}, \quad (13)$$

where X_{it} are year and region interactions (for the four agro-climatic zones in Côte d'Ivoire), v_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 agro-climatic 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.

Two important functional form assumptions are embedded in (13). First, we are assuming that commodity demands are multiplicatively separable between the Pareto weight λ and

household expenditure x . Second, we are assuming that commodity demands are log-linear in expenditure.

There are preferences that yield commodity demand functions of this form. For example, we obtain (13) if m and f have non-identical homothetic preferences.⁸ More important, it is possible to test these assumptions and we do so in section 4.4.1.

It is possible, in principle, to test the efficient households hypothesis by examining the exclusion restriction implied by (11): Conditional on total expenditure and a household fixed effect, demand for any particular good does not depend on the rainfall realization. Such tests, however, 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 or understated.⁹ 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 mis-specified, the coefficients of the rainfall variables will be inconsistently estimated as well, and misleading conclusions could be drawn.

These problems do not affect tests based on overidentifying restrictions suggested by equation (12). To implement these tests, in addition to the functional form of equation (13), we assume that rainfall in millimeters (disaggregated by season, and including indicator variables for extreme rainfall events) affects the logarithm of profits. In turn, assuming that the relationship between total expenditure and total profit is linear in logarithms, for a vector of rainfall R , we have the following relationship between rainfall and total outlay:

$$\log(x_t) = R_{it}\alpha + X_{it}\delta_x + \epsilon_{it}. \quad (14)$$

Combining equations (13) and (14), we obtain the following reduced form relationship be-

⁸For example, suppose $c_i = \{c_{i1}, c_{i2}\}'$, and that preferences for i are $Eu_i(c_i) = Ec_{i1}^{\alpha_i} c_{i2}^{1-\alpha_i}$. Then it can be shown that the Pareto efficient allocation described in (1) yields a demand for, say, commodity 1 ($c_1 = c_{f1} + c_{m1}$) which is $\ln(c_1) = \ln(\alpha_f g(\lambda) + \alpha_m(1 - g(\lambda))) + \ln x - \ln p_1$.

⁹Imagine that food expenditure is measured with error: Since it is an 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.

tween the demand for good i and rainfall precipitation:

$$\log(c_{it}) = R_{it}\pi + f(\lambda_i) + X_{it}\delta + v_i + \nu_{it}. \quad (15)$$

Taking first differences of equation (14) and (15), we obtain the reduced form system:

$$\log(x_{i2}) - \log(x_{i1}) = (R_{i2} - R_{i1})\alpha + (X_{i2} - X_{i1})\delta_x + (\epsilon_{i2} - \epsilon_{i1}), \quad (16)$$

$$\log(c_{i2}) - \log(c_{i1}) = (R_{i2} - R_{i1})\pi + (X_{i2} - X_{i1})\delta + (\nu_{i2} - \nu_{i1}), \quad (17)$$

The test suggested in equation (12) is a simple overidentification test: We want to test the hypothesis that

$$\pi = \kappa\alpha \quad (18)$$

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

One drawback of the test based on (18) is that it is not explicitly linked with variation in income from various origins, and thus cannot be directly linked to the anthropological evidence. To make this link explicit and increase power by testing a specific hypothesis, 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 separately for each group of crops (male crops, female crops, and yams) a linear regression 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:

$$\log(y_{is2}) - \log(y_{is1}) = (R_{i2} - R_{i1})\gamma_{ys} + (X_{i2} - X_{i1})\delta_{ys} + (\xi_{sw2} - \xi_{sw1}). \quad (19)$$

We then form: $DR_{is} = (R_{i2} - R_{i1})\hat{\gamma}_{ys}$. Note that equations (16) to (18) 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. We estimate:

$$\log(x_{i2}) - \log(x_{i1}) = \sum_{s=1}^S DR_{is}\alpha_s + (X_{i2} - X_{i1})\delta_x + (\epsilon_{i2} - \epsilon_{i1}) \quad (20)$$

and

$$\log(c_{i2}) - \log(c_{i1}) = \sum_{s=1}^S DR_{is}\pi_s + (X_{i2} - X_{i1})\delta + (\nu_{wi2} - \nu_{wi1}), \quad (21)$$

and we test whether $\frac{\pi_s}{\alpha_s} = \frac{\pi_{s'}}{\alpha_{s'}}$ for any $s, s' \in \{f, m, y\}$. Therefore, instead of looking generally at rainfall realizations we focus on three particular dimensions of rainfall variation that are related to different sources of income for the household and that therefore correspond to different aspects of ‘provisioning’ that are prominent in qualitative discussions of household economics in West Africa.¹⁰

3 Data and Context

3.1 Data

The data for this paper comes from the Côte d’Ivoire Living Standards Measurement Survey (CILSS). 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.¹¹

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

¹⁰ 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.

¹¹It is publicly available on the World Bank LSMS web site.

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 a lack of information on rainfall. Our final sample has a little over 800 households (each observed twice).

3.2 Gender, Ethnicity, and Agriculture in Côte D’Ivoire

Farmers in Côte d’Ivoire work in a variety of agro-climatic 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 rain fed. 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 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.¹² 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.

¹²Doss (1998), Doss (2001), Bassett (1985), Bassett (1988), Bigot (1979), Davison (1988), Dey (1993), Saito, Mekonnen and Spurling (n.d.), Gastellu (1987), Guyer (1980), 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).

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.¹³ 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.¹⁴ 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 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 (GLSS) 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

¹³Rice is a male crop in some groups, a female crop in others, and in others the gender pattern of rice cultivation is very complex.

¹⁴Meillassoux (1965), Weekes-Vagliani (1985), Weekes-Vagliani (1990), Bassett (1988), and 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.

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 equation (19) for male non-yam cash crops, yams, 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 four agro-climatic regions in Côte d’Ivoire) and their interactions. The normal pattern of rainfall in these seasons is very different in forest areas and in the savannah: 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 four 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 eight 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).¹⁵

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

¹⁵Our choice of a specification was driven by the agro-climate 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), and Sanders, Shapiro and Ramaswamy (1996)). A more parsimonious specification that includes no interactions between the rainfall variables and the savannah indicator produces results similar to those reported in Tables 4 to 7.

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 more strongly than they do income from 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 Unconstrained Tests: Rainfall and Demand

Table 3 presents the overidentification tests of Pareto efficiency, based on estimating equation (16) (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 use as dependent variables expenditures grouped in broad categories (columns (1) to (5)), and more detailed expenditures on particular goods (columns (6) to (11)). 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 of Pareto efficiency based on equation (18). The overidentification test never rejects equality. However, this test is likely to have very low power, and we turn to testing a more specific hypothesis in Table 4.

4.3 Constrained Tests: Predicted Crop Income and Demand

In order to explicitly link our specification to the anthropological literature we discussed above, we estimate equations (20) and (21). The results are presented in panel A of Table 4. In column 1, we present results from 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 at the one percent level and they are also significantly different from each other. The elasticity of total expenditure with respect to predicted income from the three sources varies from 0.1 (for male non-yam income) to 0.3 (for female crop income).

In the following columns, we present the coefficients of estimating equation (21). The final row presents the tests of the overidentification restrictions. The overidentification restrictions are rejected at the five percent level for prestige goods, adult goods, staples, and vegetables and at approximately the eleven percent level for education expenditures. It is also useful to note from the final two columns that the source of food consumed in the household is sensitive to income flows: The overidentification restrictions are rejected at the one percent level for purchases of food, and at the fifteen percent level for food consumed from family farms.

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. In addition, an examination of the coefficient estimates reveals that the deviations from efficiency correspond closely 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 (tobacco and alcohol) 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, just as strongly, for prestige goods (jewelry and adult clothing items such as “pagnes”). Income from yams, 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 prestige goods.

Expenditures on education are positively related to yam income, but inversely related to income from male non-yam crops and from female-controlled crops. In contrast, predicted increases in income from male non-yam income are associated with decreases in expenditure on purchased food, while increases in yam income are associated with decreases in expenditure on adult and prestige goods.

Not surprisingly, consumption of staples is much more strongly related to variations in yam income than to variations in male non-yam crop income or to female crop income. Consumption of vegetables is much more strongly related to female crop income than to either yam or non-yam male crop income. These results could be a consequence of local relative price movements where markets are not well-integrated, or to marketing costs of crops, which will lead households

in corner solutions to increase home consumption of home-produced commodities when they increase the production of this commodity. However, the results on other food items are not easily explained by this relative price effect: If all that was going on was that households substituted towards the goods that they produce in years when it is more abundant, one should see that both yam income and female income are less strongly associated with the consumption of other food items (in particular, food purchases) than male cash income. In practice, increased yam production is directly associated with increases in household consumption of all other foodstuffs, and not only staples. Moreover, variations in female crop income are much more strongly associated with *purchases* of staples than variations in yam or male non-yam income.¹⁶ 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.

All of these results regarding the relationship between yam income and expenditures on particular goods are consistent with the idea that income from yams is associated with household public goods and basic necessities. This corresponds to Meillassoux' description of yams as an "appreciated good" under the control of the household head for redistribution in the household. Moreover, these effects are large. A 10% increase in income from yams is associated with a 3% decline in expenditures on prestige goods, while a similar increase in female (male non-yam) income is associated with an 10% (7%) increase in expenditures on prestige goods. A 10% increase in yam income is associated with a 5% decline in expenditures on adult goods, while a similar increase in female (male non-yam) income is associated with a 15% (9%) increase in expenditure on adult goods. A 10% increase in yam income corresponds to a 3% increase in education expenditure, while a similar increase in female (male non-yam) income corresponds to a 1% (1%) decline in educational expenditure. Shifts in income from yam to either female-controlled or male non-yam crops are associated with strong declines in expenditure on education and staple food consumption, and strong increases in the consumption of adult and prestige goods.

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

There are also some strong differences in expenditure patterns from transitory fluctuations in female and male non-yam income. A 10% increase in income from female-controlled crops is associated with a 4% increase in expenditure on purchased foods and a 5% increase in meat purchases, while a similar increase in income from male non-yam crops is associated with a .3% decline in purchases of food and no rise in meat consumption.

This pattern corresponds with discussions of the role of ‘chop money’ in the descriptive accounts of household resource allocation in West Africa. 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).¹⁷ 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 can be an important source of friction within the household and between the extended families (Goldstein (2000)). Women have access to this base contribution from their husbands to provide meals for the household, but “it is ultimately the woman’s responsibility to feed everyone, whatever the amount she receives from her husband”.¹⁸ When their own disposable incomes increase, some of this increase is used for purchased foods; the rest on goods that women privately consume.

We raised the possibility that changes in local relative prices might bias these estimates. We control for region and time interactions to deal with price effects if markets are regionally integrated. To confirm that markets seem to be regionally integrated, we use the information available on prices for a wide range of goods at the CILSS cluster level for 3 of the 4 years of

¹⁷Ekejiuba uses the term ‘hearth-hold’ to mean a mother and her children.

¹⁸This is a quote from Etienne (1980), describing the relationship between husband and wives in Côte 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.

the survey. We find in Table 5 almost no commodity for which there is a statistically significant relationship between rainfall and price, conditional on the region \times year effects. Only for palm oil and for plastic sandals are the predicted income variables jointly significant. Moreover, the fact that the overidentification tests are not rejected for total food consumption, but are for the consumption of other goods (adult goods, education, and prestige goods) whose prices are not likely to substantially vary with rainfall pattern suggests that this result is not entirely due to relative price effects. As we discussed above, some results may be explained by marketing costs for households whose marginal consumption of a good is not transacted on the market (the increase in the consumption of yams when the household produces more yams). However, most of these results are not consistent with such price effects, since positive shocks in male cash crop income are not associated with food purchases (which they would be, under this hypothesis), while shocks to yam and female income are both associated with increases in the consumption of other types of food. We conclude that there is no evidence that rainfall-induced variations in local prices can account for the association we observe between changes in consumption patterns and shocks in the flows of different categories of net crop income.

4.4 Robustness Checks

In this subsection, we examine several possible threats to our interpretation of the results: The assumptions of linearity and separability in commodity demand, which are central to the derivation of the overidentification test, the assumptions of time separability, and the assumption of separability between consumption and labor supply.

4.4.1 Testing for Separability and Linearity in Commodity Demand

Our test depends on the linearity and separability of the demand function. If the demand functions are not linear, then one expects different reactions of consumption to the same income shock, depending on the overall wealth of the household, which could spuriously translate into patterns similar to those present in the data. Moreover, if they are not linear, then in general the unobserved Pareto weight and overall expenditures will not enter additively in the log demand functions, which would preclude a simple strategy of flexibly controlling for total expenditure and checking for the exclusion of the rainfall variables in this equation.

Fortunately, under the null hypothesis that the household is Pareto efficient, it is possible to test our assumptions that commodity demand functions are separable and linear. Consider a more general form for the commodity demand function (13):

$$\log(c_{it}) = \Phi(\log(x_{it}), \lambda_i) + X_{it}\delta + v_i + \nu_{it}. \quad (22)$$

We will consider a group of households that share a particular characteristic (perhaps, ethnicity). For this group G , define

$$\Phi_G(\log(x)) \equiv E(\Phi(\log(x_{it}), \lambda_i) | x_{it} = x, i \in G).$$

The basis of testing our assumptions regarding separability and linearity is that $\Phi_G(\log(x))$ and $\Phi_H(\log(x))$ are identical (up to a constant) for all arbitrary groups G and H only if $\Phi(\log(x), \lambda) = \phi(\log(x)) + f(\lambda)$.

To see the idea of the test, consider an extreme example in which λ_i does not vary across households within groups, but does vary across groups. If $\Phi(\log(x), \lambda)$ is not separable – say, $\Phi(\log(x), \lambda) = \lambda \log(x)$, then $\Phi_G(\log(x))$ has a different slope than $\Phi_H(\log(x))$.

So for households in G , we consider estimating $\Phi_G(x)$ in

$$\log(c_{it}) = \Phi_G(\log(x_{it})) + X_{it}\delta + v_i + \nu_{it} + \eta_{it} \quad (23)$$

where $\eta_{it} = \Phi(\log(x_{it}), \lambda_i) - \Phi_G(\log(x_{it}))$. With data from two periods ($t = 1, 2$) on each household, we can take the first difference of equation (23) to obtain:

$$\log(c_{i2}) - \log(c_{i1}) = \Phi_G(\log(x_{i2})) - \Phi_G(\log(x_{i1})) + (X_{i2} - X_{i1})\delta + \nu_{i2} - \nu_{i1} + \eta_{i2} - \eta_{i1}, \quad (24)$$

which we re-write as

$$\log(c_{i2}) - \log(c_{i1}) = g(\log(x_{i2}), \log(x_{i1})) + (X_{i2} - X_{i1})\delta + \nu_{i2} - \nu_{i1} + \eta_{i2} - \eta_{i1}. \quad (25)$$

We then follow Robinson (1988) and Hausman and Newey (1995) to estimate δ . To simplify 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 $[(X_{i2} - X_{i1})]$. The estimator of δ conditional on membership in group G is:

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

In other words, $\hat{\delta}_G$ is obtained by estimating (separately) the nonparametric relationships between y and (z_1, z_2) and m and (z_1, z_2) , forming the residuals, and regressing the residuals of y on the residuals of m . The non-parametric estimator we use to estimate the conditional expectations is the Fan (1992) locally weighted regression, with a quartic kernel.

To calculate $\hat{\Phi}_G(x)$, we first obtain the estimate of $g(z_1, z_2)$ by partialling out the coefficient of m (see Robinson (1988) and Hausman and Newey (1995)):

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

We then apply the partial means method suggested by Porter (1996) to recover the shape of $\Phi_G(\cdot)$ (up to an unidentified constant term) in equation (23):

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

Pointwise confidence intervals for $\hat{\Phi}_G(z)$ are constructed based on 50 bootstrap replications.

We estimate $\hat{\Phi}_G(x)$ and $\hat{\Phi}_H(x)$ for groups defined by observables that are *a priori* likely to be associated with different λ . We define a first pair of groups by ethnicity, and a second pair by estimates of the relative variance of the incomes of husbands and wives.

Figures 1a–k present estimates of Φ_{Akan} and $\Phi_{Non-Akan}$, where group “Akan” households are members of matrilineal Akan ethnic groups, and “Non-Akan” households are members of other ethnic groups. Inheritance systems differ in important ways across these groups, and this could translate into systematic differences in λ_i across these groups. However, for all but vegetable (and perhaps meat) consumption, the estimated $\hat{\Phi}_{Akan}(x)$ and $\hat{\Phi}_{Non-Akan}(x)$ are difficult to distinguish across these broad ethnic classes. Moreover, in almost all cases the estimated demand function is very close to linear.

Figures 2a–k present similar estimates for a pair of groups defined by the relative variance of incomes of husbands and wives. The estimate of female (male) income variance is simply the squared change in income attributable to crops controlled by the female (male) across the two years the household is observed. Households are assigned to the “High Female Variance” (“High Male Variance”) group if the estimate of the variance of female (male) income is higher than that of male (female) income. If λ_i is related to the riskiness of the productive activities of the individuals in household i , then it may differ systematically across these groups. However,

we again see that for almost all goods, the demand functions are very similar across these groups (here, the exceptions are education expenditures and vegetable consumption). Again, the estimated demand functions are very close to linear.

Therefore, we see no evidence of important non-separabilities between the Pareto weight λ_i and total expenditure in commodity demands. In addition, log demand is estimated to be approximately linear in log expenditure.

4.4.2 Labor Supply and Commodity Demand

Another important assumption we made in the derivation of the overidentification test is the separability between labor supply and the consumption choices. If labor markets are well functioning, this assumption will be satisfied. In practice, however, it may not be satisfied in several cases. 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. If different crops have different effort requirements, or if the labor market works better for some crops than others (which, according to Meillassoux (1965), is true – there is a spot labor market for cash crops, but not for yams), then the demand for calories will vary with rainfall even conditional on total expenditure. Another possibility is that some of the laborers are paid in food, rather than in cash (again, a possibility raised by Meillassoux (1965)), which may again show up in differential consumption.

We examine this possibility in Table 6, where we show: First, that agricultural labor demand is not strongly associated with rainfall realizations (neither household labor nor amount of money spent in paid labor strongly changes with income increases predicted by rainfall). Second, that to the extent labor demand moves with rainfall realizations, it does so in a way that is not consistent with a nutrition-productivity explanation of our results. In particular, we see that rainfall that is associated with higher output of female crops and male cash crops is associated with a *higher* demand for labor. If nutrition productivity links were the cause of the rejection of Pareto efficiency, these rainfall realizations would be associated with higher demand for calories, not higher demand for adult or prestige goods. Finally, the last line of Table 6 show overidentification

tests similar to those performed for consumption expenditures. With this test, we are testing whether each type of labor is more strongly associated with rainfall favorable to one crop than with rainfall favorable to another crop, over and above their effect on total expenditure. If this were the case, and these different types of labor were also associated with different consumption choices, then our assumption would not be satisfied. In fact, none of the overidentification tests are significant. This provides some confidence in our interpretation that the expenditure results are not the consequence of differential increases in the demand for labor that then translate into different consumption requirements.

So far, we have treated the household as a well-defined entity, without taking into account the fact that household composition may be endogenous to rainfall realization. A possible reason for some of the effects we found in section 4.3 is that household composition is actually responding to rainfall: For example, the household could take on more prime age adults when the yam crop is going to be plentiful, both for providing extra labor supply, and as a way to provide insurance to other villagers or relatives. For the same reason, the household could also take on more children. If this phenomenon were driven by total expenditure only, this would not translate into differential consumption of particular goods conditional on total expenditure. However, the increase in the production of some crops may lead to larger changes in household composition. It can, for example, happen that nephews move in with their aunt when she is in good financial standing (Akresh (2003)).

To the extent that taking on relatives is, to some extent, private consumption, the finding that this is actually happening would not necessarily invalidate our results, but rather provide an interpretation for them. In practice, however, as we show in Table 7, household size and composition do not seem to respond differentially to different types of rainfall. Household composition *is* indeed sensitive to increases in predicted income due to rainfall: Family size increases with predicted income, and this seems mostly driven by 5 to 14 year-old children (presumably due to fostering) and to some extent prime age adults (both males and female). However, these increases are proportional to the effect of the predicted income on total expenditure, and none of the overidentification tests are significant: While there is evidence that household size and composition is sensitive to household income, there seems to be no evidence that it reacts differentially to increases in different types of predicted income.

4.4.3 Time Separability of Preferences

If, contrary to the model in section 2.1, the marginal utility of the consumption of $\{c_{iwt}\}$ is affected by consumption in period $t - 1$, then the recursivity that produces equation (9) fails, as do the testable restrictions. It is important to note that this is not a matter of imperfect financial markets: equation (9) 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. However, in panel B of Table 4 we examine this possibility and find no evidence of non-separabilities over time in demand.

The specification in panel B is identical to that of panel A, with the exception that rainfall realizations are lagged 8 seasons (2 years) relative to those in panel A.¹⁹ “Predicted lagged male non-yam income”, for example, is defined as 2-year lagged rainfall changes times the coefficients estimated in columns (5)-(6) of Table A1. As anticipated in (9), column (1) provides evidence that lagged changes in income influence current expenditure. There are indeed non-trivial dynamics in the expenditure process: A particular rainfall realization has implications for the distribution of future profits over multiple years, and therefore for saving and expenditure choices over time. Also as one would expect, these effects are strongest for rainfall realizations that influence male non-yam income, which is largely generated from cocoa and other tree crops.

As a consequence, we see effects of lagged changes in predicted income on the demands for particular consumption goods in columns (2)-(11). However, all of these effects are closely proportional to those on overall consumption: In no instance is the overidentification test even close to rejected. Therefore, we find no evidence that lagged income realizations are associated with changes in the composition of current demand, and thus no evidence of non-separabilities over time in consumption.

¹⁹We construct $DR_{it} = (R_{i,t-2} - R_{i,t-3})\hat{\gamma}_{ys}$ for $s \in \{m, y, f\}$ where time periods are years, and $\hat{\gamma}_{ys}$ is estimated in (19).

4.5 Interpretation: Can the Results Be Explained by Information Asymmetries?

The full information collective model with which we began may seem unrealistic. A plausible alternative is a household that achieves a Pareto efficient allocation of consumption subject to a series of information constraints. Husbands and wives operate separate enterprises. In neighboring Ghana, there is direct evidence from surveys that spouses have poor information about profits from each others' enterprises, and they also know little about each others' private consumption (Goldstein (2000)). Could our pattern of empirical results emerge from households that achieve an efficient allocation of resources subject to these information constraints?

In this context, if the researcher observed individual incomes, one would expect to see a correlation between consumption (either private or publicly observed) and individual income: Insurance should only be partial, as shown by the literature on insurance in the presence of imperfect information (see, for example, Ligon (1998)). Note, however, that the exercise we are carrying out here is different: Since we are only using the part of income variation that is *predicted* from rainfall realization, and the rainfall realization is observed by everybody in the household, we are testing whether the household members are fully insuring themselves against the *observable* part of income shocks. *Prima facie*, one would expect that the household members can insure themselves against the observed part of income shocks even in an environment where there is imperfect information (unless they do not know the model that links rainfall to their partner's income, which seems unlikely). That said, there may be circumstances where the variance of the private income of a household member is correlated with the realization of the observed shocks, which would lead the household to compensate her with more consumption.

The point is made most clearly with a very simple model. Income from individual i 's enterprise depends on publicly-observed rainfall and on a shock that is private information: $f_i(r) + \varepsilon_i$. $f_i(r)$ is defined so that $E_i(\varepsilon_i|r) = 0$. The distribution of ε_i is defined by the density $h_i(\varepsilon_i|r)$.

Individual consumption (for simplicity, of a single good) is unobserved, but there can be an observed net transfer of income t from the husband to the wife. An efficient allocation within

the household satisfies

$$\max_{c_f, t} Eu_f(c_f) + \lambda Eu_m(c_m)$$

subject to the household resource constraint

$$c_f + c_m \leq f_f(r) + \varepsilon_f + f_m(r) + \varepsilon_m \quad (26)$$

In addition to the standard non-negativity constraints, there are now two new incentive compatibility constraints:

$$c_f \in \arg \max_{c_f} Eu_f(c_f) \quad (27)$$

$$s.t. c_f \leq t + \varepsilon_f$$

$$c_m \in \arg \max_{c_m} Eu_m(c_m) \quad (28)$$

$$s.t. c_m \leq f_m(r) + \varepsilon_m + f_f(r) - t$$

The accounting is arbitrary but without loss of generality for what follows. The female chooses her consumption to maximize her utility subject to not spending more than her observed income (t) plus her private income (ε_f), and the male does the same (his observed income is $(f_m(r) + f_f(r) - t)$ and his private income (ε_m)). The constraints in (27) and (28) bind in any efficient allocation. Together they imply (26), so substituting we see that the efficient allocation satisfies

$$\max_{t(r)} Eu_f(\varepsilon_f + t(r)) + \lambda Eu_m(f_m(r) + f_f(r) + \varepsilon_m - t(r)). \quad (29)$$

An efficient household chooses an allocation of observed income between the wife and husband when rainfall r is realized to equalize the expected marginal utility of consumption between the two:

$$\int \frac{\partial u_f(\varepsilon_f + t(r))}{\partial c} h_f(\varepsilon_f; r) d\varepsilon_f = \lambda \int \frac{\partial u_m(f_m(r) + f_r(r) + \varepsilon_m - t(r))}{\partial c} h_m(\varepsilon_m; r) d\varepsilon_m. \quad (30)$$

Consider the analogue to our earlier result. Conditional on total observed household expenditure ($= f_m(r) + f_r(r)$), does the observed expenditure on, say, the female private good ($= t(r)$) depend on the rainfall realization?

In this model, the observed composition of consumption does *not* depend on rainfall realizations conditional on observed total expenditure, unless the distribution of unobserved income

depends on rainfall. For example, if two distinct realizations of rainfall are associated with the same observed expenditure, but the second involves higher variance of ε_f than the first (but the same variance of ε_m), then the net transfer from the husband to the wife will be higher in the second.²⁰

If there is a particular relationship between mean (observed) output for individual i and the variance of i 's private output across rainfall realizations, then this effect could underlie some of the empirical regularities we observe. If “better” rainfall for i (in the sense that $f_i(r)$ is higher) is associated with higher variation in ε_i , then in an efficient allocation, higher observed income to i is associated with higher observable expenditures by i conditional on total observed expenditure.

While this cannot be tested directly (since we do not observe private income), it seems unlikely that our patterns of results can be explained by this fact. We find that increases in female and male cash crop predicted income are associated with an expenditure towards adult and prestige goods, and that increases in predicted yam income are associated with no increases in this expenditure. To explain this pattern uniquely by an increase in the variance of individual expenditures, the argument would require that the variance of ε_y decreases with rainfall that increases $f_y(r)$, while the opposite pattern holds for female crops and other male crops. While this pattern is possible, it seems less than likely.

A similar example can be constructed if moral hazard is the source of the imperfect information (the household members cannot adequately monitor each other's labor). Here again, *prima facie*, the household members should still be able to insure each other against observable shock. Differential individual consumption would arise only if higher output translated into

²⁰Let \hat{r} and \tilde{r} be such that $f_m(\tilde{r}) + f_f(\tilde{r}) = f_m(\hat{r}) + f_f(\hat{r})$ so that aggregate public resources are identical. However, let $h_f(\varepsilon_f|\hat{r})$ be a mean preserving spread of $h_f(\varepsilon_f|\tilde{r})$, while $h_m(\varepsilon_m|\tilde{r}) = h_m(\varepsilon_m|\hat{r})$. Then if marginal utility is convex and denoting the efficient net transfers with rainfall \tilde{r} and \hat{r} as $t(\tilde{r})$ and $t(\hat{r})$:

$$\lambda \int \frac{\partial u_m(f_m(\tilde{r}) + f_r(\tilde{r}) + \varepsilon_m - t(\tilde{r}))}{\partial c} h_m(\varepsilon_m|\tilde{r}) = \lambda \int \frac{\partial u_m(f_m(\hat{r}) + f_r(\hat{r}) + \varepsilon_m - t(\hat{r}))}{\partial c} h_m(\varepsilon_m|\hat{r}) =$$

$$\int \frac{\partial u_f(\varepsilon_f + t(\tilde{r}))}{\partial c} h_f(\varepsilon_f; \tilde{r}) d\varepsilon_f > \int \frac{\partial u_f(\varepsilon_f + t(\hat{r}))}{\partial c} h_f(\varepsilon_f; \tilde{r}) d\varepsilon_f.$$

So

$$\lambda \int \frac{\partial u_m(f_m(\tilde{r}) + f_r(\tilde{r}) + \varepsilon_m - t(\tilde{r}))}{\partial c} h_m(\varepsilon_m|\tilde{r}) > \int \frac{\partial u_f(\varepsilon_f + t(\hat{r}))}{\partial c} h_f(\varepsilon_f; \tilde{r}) d\varepsilon_f.$$

Therefore, $t(\tilde{r}) > t(\hat{r})$.

higher variance of the required effort. So to explain our results, it should be the case that while higher cash crop predicted income translates into a higher variance of individual effort by males, whereas higher yam predicted income translates into a lower variance of individual effort by males. Again, it does not seem very likely.

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 no effect on the purchases of food. This result does not seem to be explained by changes in market prices or prices faced by the household, and they are robust to several specification checks: The assumption of linearity of the demand functions seems verified in the data, the results do not seem to be explained by non-separability between labor supply and consumption demands, nor do they arise from differential changes in household composition as a response to these shocks.

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. Finally, because the variation in this paper comes from observable (and common) rainfall shocks, these results are not easy to reconcile with simple models of imperfect information (such as imperfect observation of the output, consumption, or the inputs that went into production). A more radical departure from

the conventional model is required.

One model that is consistent with our results on the differential impact of male cash crop and female income is the model of informal insurance with limited commitment. In this model (Coate and Ravallion (1993), Kocherlakota (1996), and Ligon et al. (2002)), individuals cannot commit to remain part of the informal insurance arrangement. In any single period, an individual who received a high realization of income compares the short-term loss of remaining in the insurance arrangement (the payment he must make to the common pool) and the long-term insurance gain. As a result, perfect insurance is often not achievable, and individuals who receive high incomes in a specific period consume more than others. This model has been recently proposed as a model of insurance in the household by Ligon (2003). Note that in this setting, the household cannot fully insure even against fully observable income shocks: It would thus explain why males and females consume more of the goods they prefer when their own predicted income is bigger.

It is less direct to reconcile the model of limited commitment with our finding that expenditures on food and education increase with yam predicted income. While the notional property rights over yam income are attributed to men, yam, as an “appreciated product”, comes with strings attached. Deviations from accepted use of this income can provoke strong punishment from the community. Correspondingly, our results seem to imply that yam income is put into a separate account, not fungible with the rest of male income, and spent on different goods. 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. We hypothesize that this institution could have arisen as an endogenous response to the limited commitment problems faced by the households. Faced with the consequences of these commitment problems, society has constructed a new type of property right, such that the basic needs of the “household” can be met with income flows from the appreciated products. The social sanction associated with deviation from the accepted use limits the enforcement problem for these income streams.

This suggests 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 intra-household negotiations and/or by constrained fungibility of resources across uses. For example,

inter-temporal decisions such as the allocation of household resources into the human capital of children could be affected by the labelling of income if husbands and wives face different opportunities in financial markets. More generally, our results suggest 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 intra-household 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.

References

- Akresh, Richard (2003) 'Risk, network quality and family structure: Child fostering decisions in Burkina Faso.' Mimeo, Yale University
- 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., University of California
- (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, Institute des Savanes
- 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
- Coate, Stephen, and Martin Ravallion (1993) 'Reciprocity without commitment: Characterization and performance of informal insurance arrangements.' *Journal of Development Economics* 40, 1–24
- 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 (1998) 'Do households fully share risk? Evidence from Ghana.' Mimeo, Williams College
- (2001) 'Men's crops? women's crops? The gender pattern of cropping in Ghana.' Mimeo, University
- 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.' Ph.d, University of California, Berkeley
- Guyer, Jane (1980) 'Food, cocoa, and the division of labour by sex in two West African societies.' *Comparative Studies in Society and History* 22(3), 355–373
- (1984) 'Family and farm in southern Cameroon.' Technical Report 15, African Studies Center, Boston University
- (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, and P. E. Peters (1987) 'Conceptualizing the household - Issues of theory and policy in Africa - Introduction.' *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* 18(2), 197–214
- 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 Mook (Boulder: Westview Press)
- Kocherlakota, Narayana (1996) 'Implications of efficient risk sharing without commitment.' *Review of Economic Studies* 63, 595–610

- Kozel, Valerie (1990) ‘The composition and distribution of income in Cote d’Ivoire.’ Technical Report, World Bank LSMS Working Paper No. 68
- Ligon, Ethan (1998) ‘Risk sharing and information in village economics.’ *The Review of Economic Studies* 65(4), 847–864
- (2003) ‘Dynamic bargaining in households (with an application to Bangladesh).’ MIME, UC Berkeley
- Ligon, Ethan, J.P. Thomas, and T. Worrall (2002) ‘Informal insurance arrangements with limited commitment: Theory and evidence from village economies.’ Forthcoming, *The Review of Economic Studies*
- 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 ‘Raising the productivity of women farmers in sub-saharan Africa.’ Technical Report, World Bank Discussion Papers, Africa Technical Department Series, No. 230
- Sanders, John, Barry Shapiro, and Sunder Ramaswamy (1996) *The Economics of Agricultural Technology in Semiarid Sub-Saharan Africa* (Baltimore: Johns Hopkins University)
- Shipton (1989) *Bitter Money: The Classification of Forbidden Commodities among the Luo of Kenya* (Washington, DC: American Anthropological association)

- Thaler, Richard (1992) *The Winner's Curse* (Princeton, New Jersey: Princeton University Press)
- Vijverberg, Wim (1988) 'Nonagricultural family enterprises in Cote d'Ivoire.' Technical Report, World Bank LSMS Working Paper No. 46
- 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'ajustement en Cote d'Ivoire.' Documents Techniques no. 9, Centre de Developpement de l'OCDE, January

TABLE 1: Descriptive statistics

	Mean, year 1 1000's FCFA	Difference in logs Year2-Year1	Observations
	Standard errors in parentheses		
	(1)	(2)	(3)
Income from male crops	553.78 (26.22)	-0.07 (.06)	1025
Income from female crops	103.46 (9.28)	0.00 (0.10)	1025
Total "male income"	593.89 (26.17)	-0.08 (0.06)	1025
Total "female income"	143.58 (10.21)	0.06 (0.09)	1025
Unattributed income	144.85 (6.24)	0.17 (0.12)	1025
Total expenditure	1111.45 (28.41)	-0.10 (0.02)	1008
Food consumption	639.37 (12.68)	-0.06 (0.02)	973
Adult goods	45.67 (2.61)	-0.32 (0.12)	1025
Clothing	263.92 (9.41)	-0.17 (0.06)	1025
Prestige goods	218.11 (8.06)	-0.17 (0.07)	1025
Staples	442.58 (12.32)	-0.03 (0.03)	1025
Meat	142.72 (7.62)	-0.12 (0.04)	1025
Vegetables	51.21 (2.67)	-0.07 (0.08)	1025
Processed foods	41.15 (1.74)	-0.25 (0.04)	1025
All purchased foods	309.86 (9.09)	-0.17 (0.03)	1020
All food consumed at home	424.49 (13.78)	-0.06 (0.05)	1025

Table 2: First stage summary statistics

	Dependent variables		
	Current		
	Male cash crop	Yam income	Female Income
	(1)	(2)	(3)
F statistics (p value)			
All rainfall variables are significant	1.99 (0.014)	3.50 (0.000)	2.53 (0.000)
Current year rainfall variables significant	1.18 (0.315)	3.38 (0.000)	2.43 (0.005)
Past year rainfall variables significant	2.79 (0.005)	4.64 (0.000)	2.64 (0.001)
Rainfall variables significantly different from:			
Male cash crop	NA		
Yam income	2.10 (0.010)	NA	
Female income	2.10 (0.009)	2.38 (0.002)	NA

Note

(1) The full results are presented in Appendix, table 1

(2) The specification include year dummies, region dummies,
and their interactions

Table 3: unconstrained overidentification tests

	Dependent variable: Change in log(item consumption)										
	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)
F test (p value)	3.16	2.21	3.22	3.25	1.56	3.04	1.90	1.92	3.33	5.95	2.04
Rainfall variables jointly significant	(0.000)	(0.000)	(0.000)	(0.000)	(0.039)	(0.000)	(0.003)	(0.003)	(0.000)	(0.000)	(0.001)
F test (p value)	0.59	0.85	1.02	1.01	0.59	0.75	0.40	0.74	0.94	1.29	0.94
Test of overidentification restrictions	(0.955)	(0.681)	(0.443)	(0.453)	(0.942)	(0.823)	(0.997)	(0.833)	(0.551)	(0.150)	(0.549)

Note: The second row presents the F statistics for a non-linear Wald Test of overidentification restriction that all year to year differences in rainfall variables affect year to year differences in expenditures only through their impact on year to year difference in total expenditure (see text, equation 19). The p value is shown in parentheses. The regressions include year dummies, region dummies, and their interactions.

Table 4: Restricted overidentification tests

	Dependent variable: Change in log (item consumption)											
	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)
PANEL A												
OLS coefficients:												
Predicted change in male non-yam income	0.126 (0.049)	0.062 (0.054)	0.870 (0.425)	-0.164 (0.334)	0.683 (0.209)	-0.101 (0.128)	0.113 (0.072)	0.002 (0.126)	0.345 (0.210)	0.004 (0.139)	-0.029 (0.078)	0.098 (0.119)
Predicted change in yam income	0.207 (0.037)	0.227 (0.041)	-0.473 (0.320)	0.296 (0.252)	-0.272 (0.158)	0.320 (0.108)	0.345 (0.054)	0.135 (0.096)	0.023 (0.159)	0.122 (0.105)	0.087 (0.059)	0.444 (0.090)
Predicted change in female income	0.309 (0.056)	0.235 (0.061)	1.537 (0.490)	0.535 (0.382)	0.993 (0.239)	-0.098 (0.159)	0.193 (0.082)	0.492 (0.144)	0.995 (0.239)	0.474 (0.159)	0.412 (0.089)	0.313 (0.136)
F tests (p value) :		0.934	5.064	0.514	7.595	2.260	5.870	1.824	3.277	1.397	4.777	1.912
Overidentification		(0.393)	(0.007)	(0.598)	(0.001)	(0.106)	(0.003)	(0.162)	(0.038)	(0.248)	(0.009)	(0.148)
Restriction test												
PANEL B: LAGGED RAINFALL												
OLS coefficients:												
Predicted change in lagged male non-yam income	0.073 (0.020)	0.039 (0.022)	0.350 (0.169)	0.044 (0.133)	0.047 (0.082)	0.091 (0.056)	0.038 (0.029)	0.150 (0.050)	0.039 (0.083)	0.115 (0.055)	0.155 (0.031)	-0.007 (0.047)
Predicted change in lagged yam income	-0.003 (0.009)	0.004 (0.009)	0.008 (0.073)	-0.125 (0.059)	-0.076 (0.036)	-0.031 (0.029)	-0.021 (0.013)	0.015 (0.022)	0.011 (0.036)	0.027 (0.024)	0.024 (0.013)	-0.018 (0.021)
Predicted change in lagged female income	-0.001 (0.026)	0.018 (0.028)	-0.024 (0.220)	-0.251 (0.173)	-0.289 (0.107)	0.093 (0.079)	0.044 (0.038)	0.023 (0.064)	-0.054 (0.107)	-0.010 (0.071)	0.062 (0.040)	-0.035 (0.061)
F tests (p value) :		0.105	0.128	0.254	0.043	0.016	0.049	0.052	0.024	0.058	0.054	0.057
Overidentification		(0.900)	(0.880)	(0.776)	(0.958)	(0.984)	(0.952)	(0.949)	(0.976)	(0.943)	(0.948)	(0.945)
Restriction test												

Note: The table presents the OLS coefficient of the difference in log consumption of each item on the difference in predicted log income (obtained from the equation presented in table A1). Standard errors are shown in parentheses. The regressions include year dummies, region dummies, and their interactions.

The overidentification test is a non-linear wald test for the hypothesis that the coefficients in each regression are proportional to their coefficients in column (1)

Table 5 : Relationship between predicted income shocks and local prices

	Dependent variable: change in 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 change in male non-yam income	0.09 (0.126)	-0.19 (0.134)	-0.04 (0.157)	0.32 (0.189)	-0.23 (0.195)	-0.07 (0.064)	0.23 (0.288)	0.44 (0.246)	0.36 (0.214)	-0.01 (0.103)
predicted change in yam income	0.00 (0.047)	0.01 (0.050)	-0.07 (0.058)	0.04 (0.070)	-0.05 (0.073)	-0.01 (0.024)	0.01 (0.107)	-0.14 (0.091)	0.00 (0.079)	0.09 (0.038)
predicted change in female income	0.13 (0.158)	0.12 (0.168)	0.04 (0.197)	0.02 (0.237)	-0.23 (0.245)	0.10 (0.080)	-0.50 (0.362)	0.53 (0.308)	-0.06 (0.268)	-0.01 (0.129)
F statistics: predicted income variables jointly significant (p-value)	0.33 (0.80)	1.04 (0.38)	0.46 (0.71)	1.04 (0.39)	0.71 (0.55)	1.23 (0.31)	1.05 (0.38)	2.80 (0.05)	1.08 (0.36)	1.76 (0.17)
	cassava	yams	plantain	oil palm nuts	peanuts	eggs	cloth	fish	sandals	enamel bowl
	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)
predicted change male non-yam income	-0.27 (0.250)	0.40 (0.225)	0.26 (0.320)	-0.14 (0.213)	0.34 (0.192)	-0.15 (0.220)	0.34 (0.212)	0.05 (0.147)	-0.04 (0.141)	0.18 (0.222)
predicted change in yam income	-0.03 (0.093)	0.04 (0.084)	-0.07 (0.119)	0.01 (0.079)	-0.03 (0.071)	0.03 (0.082)	0.08 (0.079)	-0.04 (0.055)	-0.06 (0.052)	0.07 (0.082)
predicted change in female income	-0.10 (0.314)	-0.05 (0.283)	-0.31 (0.402)	0.09 (0.267)	-0.12 (0.240)	-0.46 (0.276)	0.23 (0.266)	-0.19 (0.185)	0.40 (0.176)	0.19 (0.278)
F statistics: predicted income variables jointly significant (p-value)	0.41 (0.75)	1.18 (0.33)	0.65 (0.58)	0.23 (0.87)	1.40 (0.26)	1.06 (0.38)	1.20 (0.32)	0.58 (0.63)	2.53 (0.07)	0.50 (0.69)

Note: item prices are obtained in the market for each enumeration area. The regressions include year dummies, region dummies, and their interactions. Standard errors in parentheses

Table 6: Restricted overidentification tests, labor supply

	Dependent variable: Change in log (labor supply)							
	Total expenditure (1)	hours worked, men	hours worked, women (1)	hours worked in agriculture, men (2)	hours worked in agriculture, women (3)	hours work, non agriculture, men (4)	hours work, non agriculture, women (5)	amount spent on paid labor (6)
Predicted male non-yam income	0.126 (0.049)	0.187 (0.140)	0.070 (0.142)	0.160 (0.177)	0.085 (0.186)	0.315 (0.207)	0.177 (0.204)	-0.927 (0.586)
Predicted yam income	0.207 (0.037)	0.167 (0.106)	0.163 (0.108)	0.076 (0.134)	0.181 (0.140)	-0.119 (0.157)	-0.089 (0.154)	-0.341 (0.444)
Predicted female income	0.309 (0.056)	0.200 (0.160)	0.239 (0.162)	0.413 (0.202)	0.280 (0.212)	0.307 (0.236)	-0.031 (0.233)	-1.205 (0.669)
F tests (p value):		0.228	0.020	0.486	0.013	1.422	0.483	0.467
Overidentification		(0.797)	(0.980)	(0.615)	(0.988)	(0.242)	(0.617)	(0.627)
Restriction test								

Note: The table presents the OLS coefficient of difference in log(hours) for each type of labor supply on difference in predicted log income (obtained from the equation presented in table A1). Standard errors are shown in parentheses. The regressions include year dummies, regions dummies, and their interactions. The overidentification test is a non-linear wald test for the hypothesis that the coefficients in each regression are proportional to their coefficient in column (1)

Table 7: Restricted overidentification tests: family composition

	Dependent variable: Change in the numbers of household members in each category											
	total	Family	Infants 0-4		Children 5-14		Teenagers 15-19		Prime age 20-60		Older adults >60	
	expenditure	size	male	female	male	female	male	female	male	female	male	female
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(9)	(10)	(11)
male non yams	0.126 (0.049)	0.311 (0.243)	0.174 (0.125)	-0.192 (0.131)	0.129 (0.171)	0.316 (0.217)	-0.144 (0.156)	-0.050 (0.138)	0.070 (0.083)	0.118 (0.089)	0.034 (0.095)	0.080 (0.116)
Yam	0.207 (0.037)	0.504 (0.184)	0.190 (0.094)	0.058 (0.100)	-0.010 (0.134)	0.212 (0.170)	0.106 (0.117)	-0.099 (0.115)	0.210 (0.063)	0.110 (0.067)	-0.004 (0.067)	-0.152 (0.088)
female	0.309 (0.056)	0.843 (0.278)	0.106 (0.144)	-0.005 (0.144)	0.468 (0.203)	0.674 (0.241)	0.192 (0.187)	0.168 (0.180)	0.136 (0.096)	0.190 (0.100)	0.093 (0.106)	-0.066 (0.141)
F tests (p value):		0.028	0.950	0.880	0.942	0.173	0.899	0.668	1.139	0.110	0.257	1.448
Overidentification		(0.973)	(0.387)	(0.416)	(0.391)	(0.841)	(0.408)	(0.514)	(0.321)	(0.896)	(0.774)	(0.237)
Restriction test												

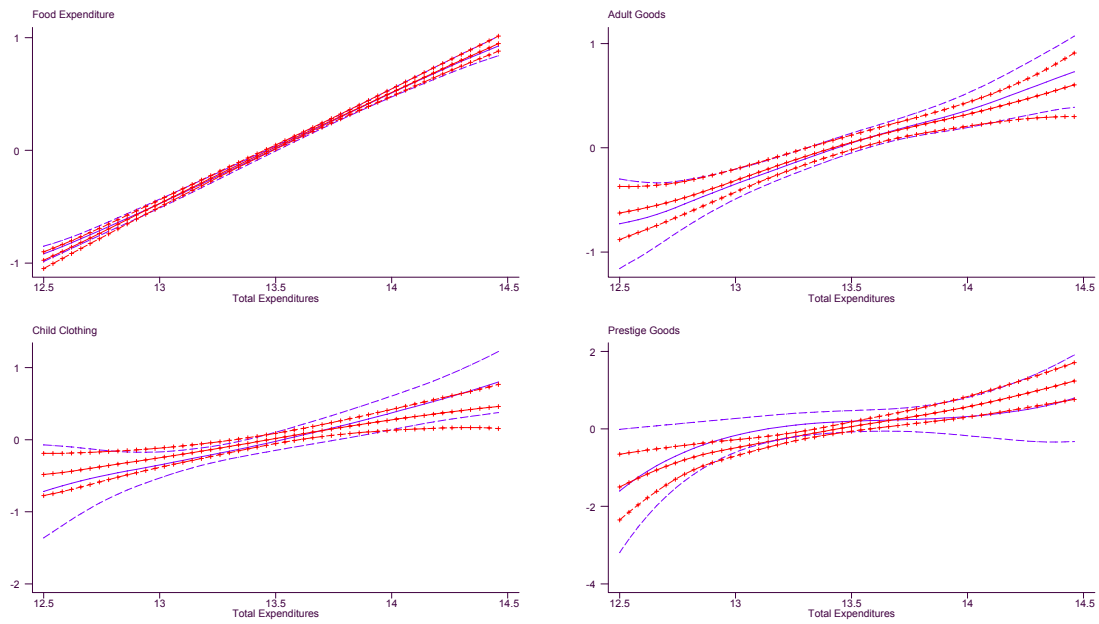
Note: The table present the OLS coefficient of difference the number of household members on difference in predicted log income (obtained from the equation presented in table A1). Standard errors are shown in parenthesis. The regression include year dummies, region dummies, and their interactions. The overidentification test is a non-linear wald test for the hypothesis that the coefficients in each regression are proportional to their coefficient in column (1)

Appendix Table A1: First stage regression results

	Dependent variables					
	Female crop income		Male cash crop		Yam income	
	Forest coefficients (1)	Savannah interaction (2)	Forest coefficients (3)	Savannah interaction (4)	Forest coefficients (5)	Savannah interaction (6)
Difference (year 2 - year 1) in:						
Aggregate rainfall current year, season 1	-0.0015175 (0.001)	0.0040317 (0.003)	0.0004811 (0.002)		-0.003153 (0.002)	-0.010761 (0.006)
Aggregate rainfall current year, season 2	0.0007268 (0.000)	0.0013814 (0.002)	-0.001099 (0.001)		0.0015603 (0.001)	0.0015827 (0.004)
Aggregate rainfall current year, season 3	-0.0006134 (0.001)	0.0038313 (0.001)	0.0001552 (0.001)		-0.002321 (0.001)	-0.003099 (0.002)
Aggregate rainfall current year, season 4	0.0007069 (0.001)	-0.0042 (0.005)	-0.000169 (0.001)		0.0005378 (0.001)	-0.006442 (0.006)
Aggregate rainfall past year, season 1	-0.0003565 (0.002)	0.0068233 (0.008)	-0.004016 (0.003)		-0.00618 (0.003)	-0.010605 (0.011)
Aggregate rainfall past year, season 2	0.0000808 (0.000)	-0.006707 (0.005)	0.0008669 (0.001)		0.0023795 (0.001)	-0.000265 (0.006)
Aggregate rainfall past year, season 3	-0.00138 (0.001)	0.0033809 (0.001)	-9.57E-05 (0.001)		-0.00226 (0.001)	0.0027378 (0.001)
Aggregate rainfall past year, season 4	-0.0007686 (0.001)	-0.003408 (0.005)	0.0014161 (0.001)		0.0007269 (0.001)	0.0053683 (0.006)
Dummy for shock, current year, season 1	-0.476418 (0.233)		-0.093278 (0.364)		-0.894238 (0.439)	
Dummy for shock, current year, season 2	0.4592265 (0.193)	0.3583756 (0.485)	0.127267 (0.300)		0.4623188 (0.283)	-2.75326 (0.828)
Dummy for shock, current year, season 3						
Dummy for shock, current year, season 4	-0.4114966 (0.378)		0.6722299 (0.654)		-2.134331 (0.520)	
Dummy for shock, past year, season 1	0.2208379 (0.208)	-0.023197 (0.531)	0.1107528 (0.362)		0.2016122 (0.262)	3.537023 1.107312
Dummy for shock, past year, season 2	-0.0744996 (0.119)	0.1403303 (0.429)	-0.037784 (0.183)		-0.133787 (0.204)	-2.962664 0.9110861
Dummy for shock, past year, season 3	-0.3152398 (0.245)	0.5705816 (1.027)	-1.324416 (0.384)		-0.124188 (0.386)	-3.387585 1.388615
Dummy for shock, past year, season 4	-0.7206122 (0.267)	0.4587139 (1.366)	0.7792504 (0.437)		-1.748257 (0.408)	1.238107 1.274639
Number of observations	976		614		607	

Note: the specifications also include year dummies, region dummies, and their interactions
Standard errors in parentheses

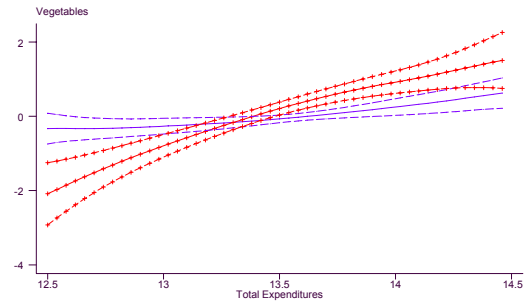
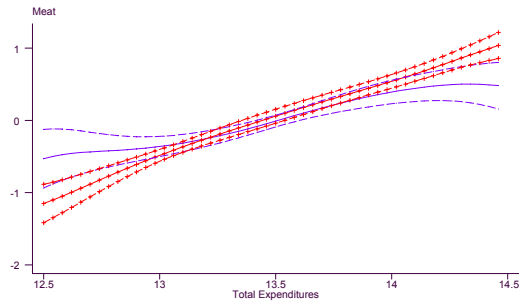
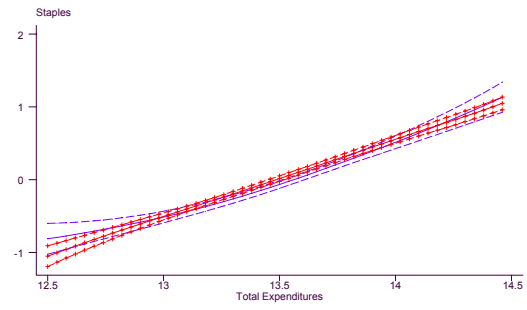
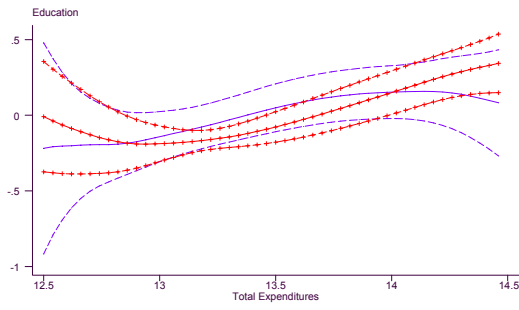
Akan Households --+--+--+ Non-Akan Households



+/- 2 std. dev. pointwise confidence intervals
Figs. 1a-d: Partial Linear Expenditure Functions

Akan Households

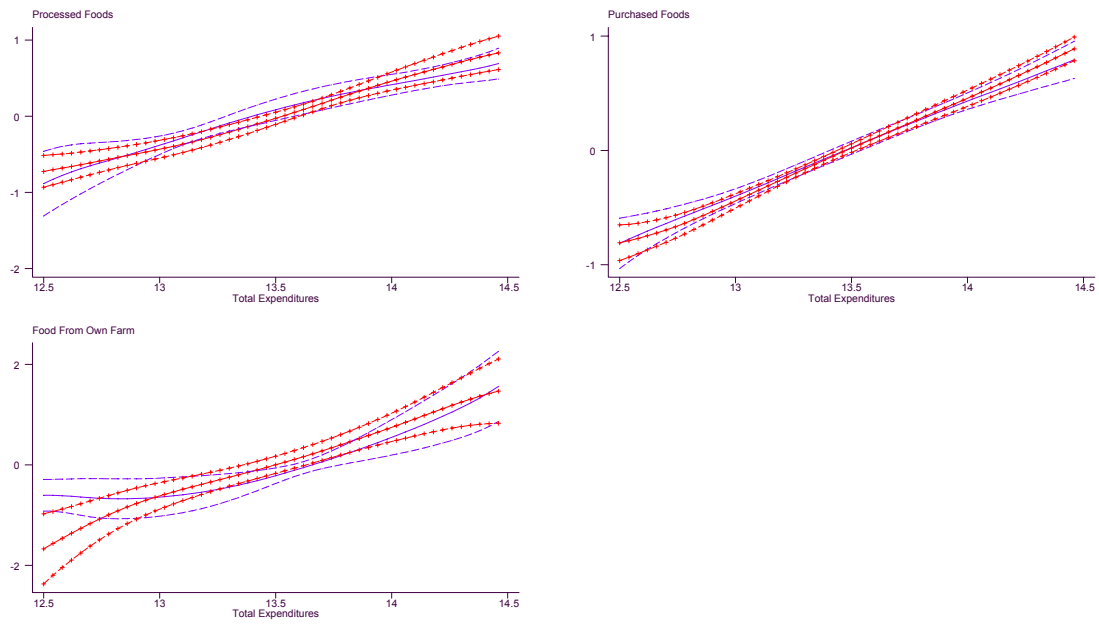
--+--+--+-- Non-Akan Households



+/- 2 std. dev. pointwise confidence intervals

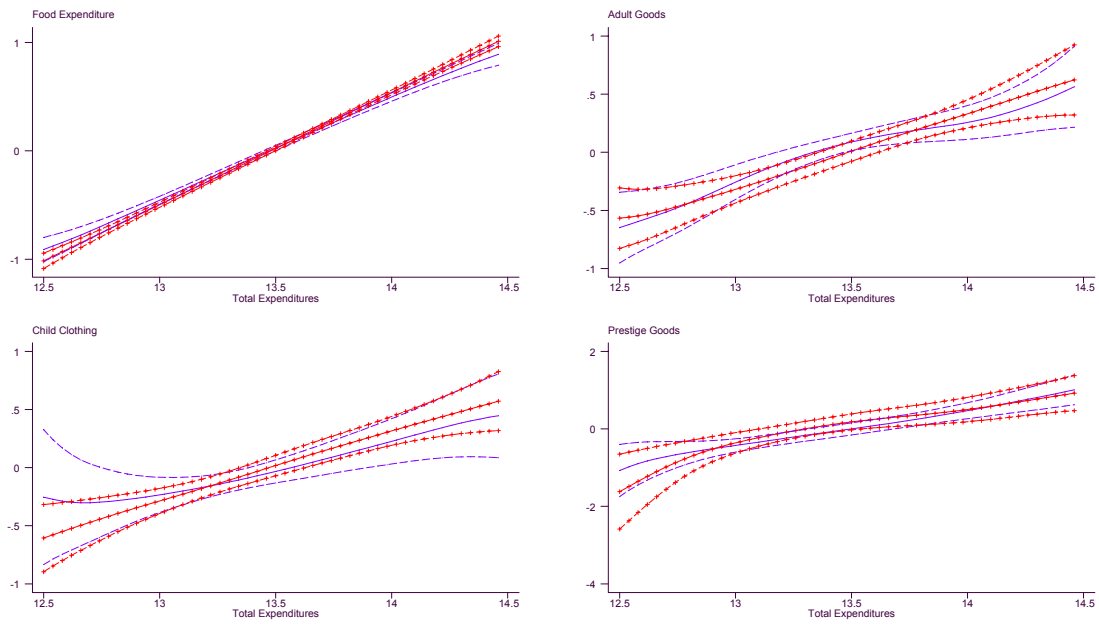
Figs. 1e-h: Partial Linear Expenditure Functions

Akan Households --+--+--+ Non-Akan Households



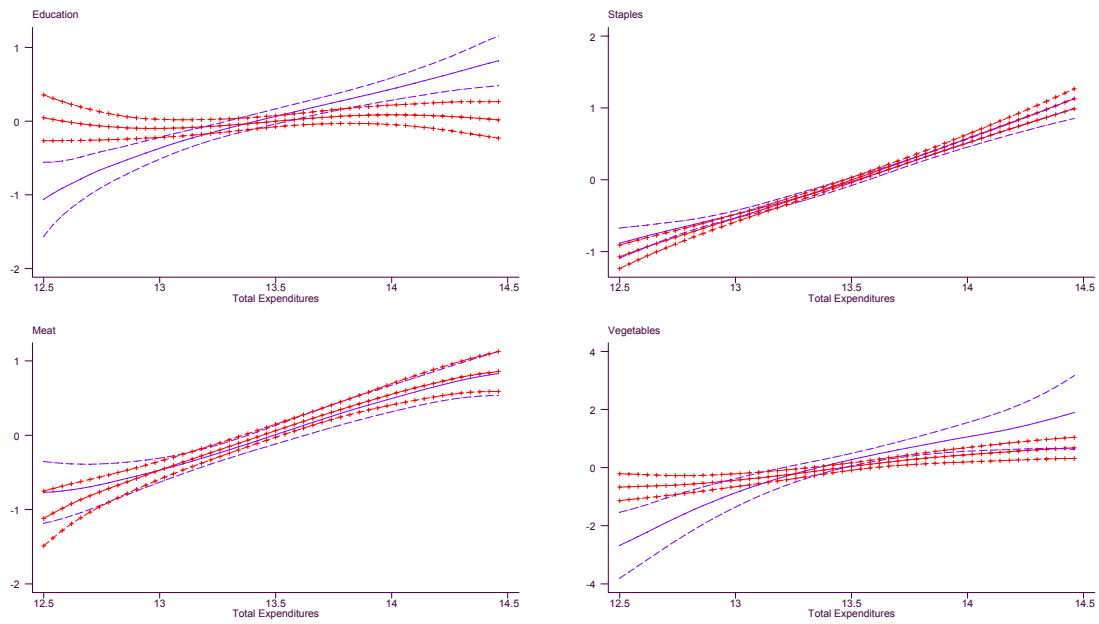
+/- 2 std. dev. pointwise confidence intervals
Figs. 1i-k: Partial Linear Expenditure Functions

High Female Variance --+--+--+ High Male Variance



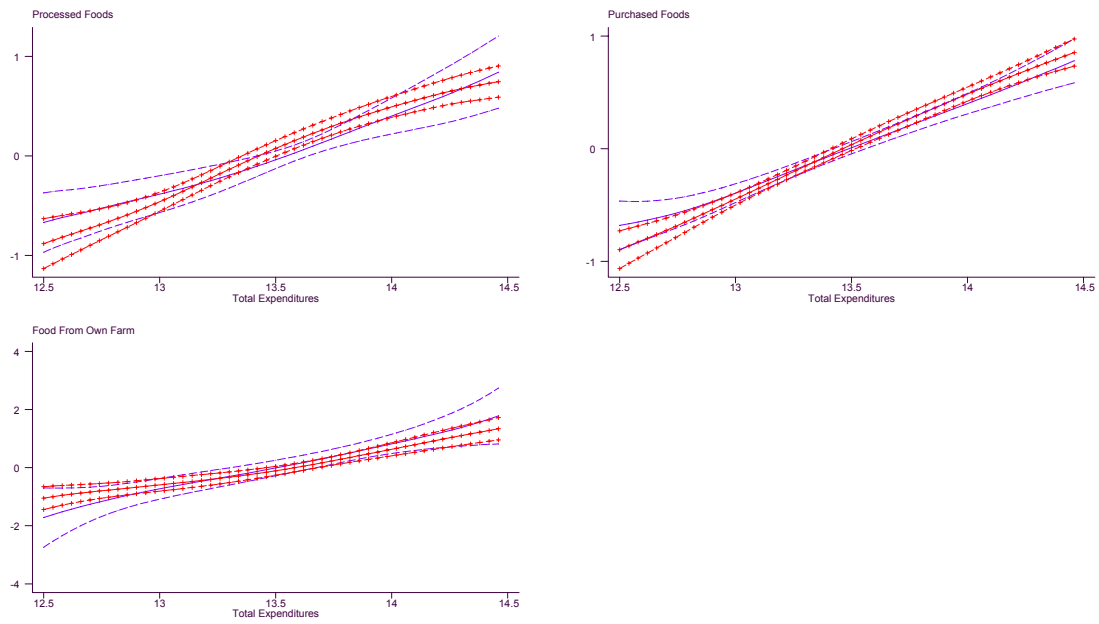
+/- 2 std. dev. pointwise confidence intervals
Figs. 2a-d: Partial Linear Expenditure Functions

High Female Variance --+--+--+ High Male Variance



+/- 2 std. dev. pointwise confidence intervals
Figs. 2e-h: Partial Linear Expenditure Functions

High Female Variance --+--+--+ High Male Variance



+/- 2 std. dev. pointwise confidence intervals
Figs. 2i-k: Partial Linear Expenditure Functions