

Gender, Social Norms, and Household Production in Burkina Faso

HAROUNAN KAZIANGA

Oklahoma State University

ZAKI WAHHAJ

Oxford University

I. Introduction

Empirical studies of intrahousehold allocation have revealed that, in many instances, gender is an important determinant in the allocation of resources within the household. For developing countries, in particular, the literature has produced evidence of systematic differences between genders in the allocation of resources relating to nutrition, health, and education (see, e.g., Haddad, Hoddinott, and Alderman [1998] and Duflo [2005] for reviews of this literature).

Furthermore, the literature has provided extensive evidence that the initial distribution of resources between men and women often affects household outcomes, thus leading to the discrediting of the “Unitary Model” of the household. Evidence on imperfect risk sharing within the household (Dercon and Krishnan 2000; Duflo and Udry 2003; Dubois and Ligon 2009) and on inefficiency in the allocation of productive resources within the household (Udry 1996; Goldstein and Udry 2008) has also led to the discrediting of the “Collective Model” of the household (Browning and Chiappori 1998), which posits that, regardless of the distribution of resources within the household, the allocation is efficient. In light of this evidence, the theoretical literature has recently turned its attention to possible constraints that can prevent an efficient outcome within the household and, in particular, to models of dynamic bargaining exhibiting lack of commitment (Ligon 2002; Lundberg and Pollak 2003; Basu 2006).

Yet, within the theoretical literature on intrahousehold allocation, why gender matters within the household remains an open question. The aforementioned

We are grateful to Chris Udry, Marcel Fafchamps, Sonia Bhalotra, and two anonymous referees for their advice and feedback on an earlier version of the article. The data used in this study were collected by a project funded by the Canadian Agency for International Development and conducted by the University of Ouagadougou (Burkina Faso) and Laval University (Quebec, Canada). We thank Taladidia Thiombiano and Frederic Martin for making the data available.

© 2013 by The University of Chicago. All rights reserved. 0013-0079/2013/6103-0002\$10.00

models treat the household as a collection of agents with possibly conflicting preferences, but it makes no distinction between the genders.

In this article we propose a simple model of intrahousehold allocation based on a particular social institution for the organization of agricultural production practiced among certain ethnic groups in West Africa. In the context of agricultural production by a household unit, the ethnographic literature makes a distinction between “common” farm plots, managed by the household head, to which all household members tend to contribute some labor, and “private” plots that are worked individually or by smaller groups within the household. The head of the household carries a particular obligation to provide for the entire household using the proceeds from the common plot.

We highlight how this institution, while resolving certain problems of commitment and informational asymmetry, can also lead to a gendered pattern in the allocation of productive resources and consumption within the household, one that is consistent with the existing empirical literature. Using a survey of agricultural households in Burkina Faso, drawn from eight villages around the country, we test the theoretical predictions of the model regarding production and consumption within the household.

Consistent with the theory, we show, using rainfall variations to capture exogenous weather shocks to agricultural output, that the head of the household has a higher propensity to spend on household public goods out of his farm income than do the other household members. This suggests that members of the household should be more willing to work on the plots farmed by the household head than on private plots, because household heads are able, thanks to a particular social institution, to commit to using the fruits of their toil on household public goods. Consistent with this reasoning, we show that common plots managed by the household head use family labor more intensively and achieve higher yields than private plots (after controlling for plot characteristics and the crops planted).

As in the previous literature (Udry 1996; Goldstein and Udry 2008), we find that farm plots managed by male household members, on average, use family labor more intensively and achieve higher yields than those managed by female household members in the same household. However, no such gender difference exists among the private plots: male and female private plots have similar yields, and each is farmed primarily with family labor of the same gender as the plot manager. Thus, the social institution that places a particular obligation on the head of the household—and the fact that the household head is usually a man—can account entirely for the gendered pattern in agricultural production documented in the existing literature.

Our empirical tests of productive efficiency within the household are closely related to two papers in the recent literature. First, Akresh (2008) finds, for a nationally representative sample of agricultural households in Burkina Faso, from a survey conducted in 1990, that households that experience negative rainfall shocks do satisfy Udry's test of Pareto efficiency in household production. Akresh argues that the efficient allocation of productive resources may require household members to engage in cooperative agreements—where, for instance, one household member promises to provide labor on the plot of another in exchange for some future compensation—and these may be costly to negotiate, monitor, and enforce. Therefore, members of a household may only find it worthwhile to negotiate such agreements when agricultural production is particularly hard-hit, as it would be during a season of very low rainfall.

In the case of the CEDRES/Laval survey used for the empirical analysis in this article, rainfall levels in the four provinces from which our sample is drawn were all above the 30-year average during 1994 (between 0.3 and 2.0 standard deviations) and no more than 0.4 standard deviations below the average in 1993. Therefore, our finding of higher yields on plots managed by the household head is consistent with the explanation provided by Akresh (although there is not sufficient variation in rainfall levels in our sample to test this formally). Our findings, in conjunction with the explanation provided by Akresh, would suggest that household members provide labor on farm plots managed by the household head, even when rainfall levels are favorable, because the “exchange” taking place in this instance does not require specific negotiations, monitoring, or enforcement: the social norms ensure that all parties concerned have the appropriate incentives.

Goldstein and Udry (2008) also find, for a sample of agricultural households in Southern Ghana, that women achieve significantly lower yields on their farm plots compared to men who belong to the same household, after controlling for observable plot characteristics and the types of crops planted. However, the entirety of the gender effect can be explained by the length of time that the plot was left fallow before planting. Goldstein and Udry provide evidence to show that fallowing improves soil fertility but that individuals with little political power within the community risk losing rights to their land if it is left fallow over extended periods.

It is conceivable that, in the context of the Burkinabe households considered in our analysis, the length of fallow also varies with status within the household or community and that this can account for some part of the variation in plot yields observed for our sample. In the absence of information

regarding the practice of fallowing land in the CEDRES/Laval survey, it is not possible to confirm or rule out the fallowing hypothesis. However, we would argue that a compelling hypothesis should also account for the dramatic variations in labor intensity and the response of household consumption to shocks to different types of household income that we observe in the data. The norm-related hypothesis proposed in this article, as well as the ethnographic accounts discussed in the next section, can account for these patterns, but the fallowing hypothesis, on its own, does not.

The remainder of this article is organized as follows. Section II discusses the ethnographic literature on the organization of agricultural production and the consumption among a number of ethnic groups in West Africa, based, in particular, around the Sahel region. Section III develops a simple model of intrahousehold allocation that draws on this ethnographic evidence. A brief description of the household survey used for the econometric analysis is given in Section IV. Section V presents some descriptive statistics, focusing on the differences between the household head's plots and plots farmed by other household members. Section VI discusses our empirical findings. Section VII explores alternative hypotheses, and Section VIII concludes.

II. Household Production and Consumption in Burkina Faso

In this section, we highlight two features of production and consumption within Burkinabe households by drawing on ethnographic studies. The first of these features is that the household typically carries out its farming activities on a number of distinct plots with well-defined boundaries, each under the control of a specific adult household member. Among these plots, there is a separation between common plots under the control of the head of the household and private plots controlled by other household members. The second feature is that according to existing social norms, the head of the household is expected to provide for the entire household using the proceeds of the common plot, but the other household members have more leeway in how they use the income from their own plots.

These two features will play an important part in the formal model introduced in the following section. Therefore, we discuss them in some detail here in the context of the ethnographic literature. The studies cited below are all based on the Mossi, an ethnic group that practices farming and that constitutes about 40% of the population of Burkina Faso (CIA World Factbook 2009). The Mossi are also the predominant ethnic group in two of the eight villages used for the CEDRES/Laval household survey that is used for the empirical analysis in this article.

The characterization of agricultural production within Mossi households by ethnographers consistently distinguishes between common plots, which are farmed communally and managed by the head of the household, and private plots, which are managed by other members of the household. Hammond (1966), with regard to the Mossi of the Yatenga region in the north of Burkina Faso, writes:

In the extended family, the father has final authority over the allocation of all farmland used by his wives and children. Certain fields are under the more specific proprietorship of each of his mature sons who may in turn delegate the right to use them to their wives and their own sons. In each instance such fields are worked principally by their respective proprietors who have the first right to their products. However, most of the fields farmed by the head of the extended family are worked with the help of his wife or his wives and children, by any younger brothers who might be attached to his household, and, occasionally, by a sister's son who has come to reside with his mother's brother. (76)

This distinction between fields managed by the head of the household and other family members is also evident in the characterization of the Mossi provided by Fiske (1991):

The [Mossi] typically work their fields communally. A small group of people led by a senior man jointly cultivates a collective field (*puugo*). A cultivating group typically consists of the senior man and his wives and children, often with his younger brothers and their wives and children. . . . The people who spend the bulk of the day cultivating the joint corporate field (*puugo*) of the extended family together also cultivate separate fields (*beolse*). . . . Such a separate field may be cultivated by a junior man with his wives and children (if any), or by a wife of the senior man together with her children (if any); sometimes older children have separate fields. . . . Sometimes a separate field is cultivated alone by an adolescent, by a young man not yet married, by a widower with no surviving children at home, by a young wife with no children yet, or by a wife with no surviving children. But often the separate field, like the corporate field, is cultivated collectively. (260–61)

The descriptions above also suggest that the pattern of family labor contribution to a plot may vary considerably between common and private plots. All household members are expected to contribute labor to the common plot. But private plots are farmed individually or with the assistance of a nuclear group within the extended family. However, there are also indications in the literature that rules of labor allocation are not strict. Fiske (1991) writes: "In these cultivating groups [farming the common plot], what really matters is participation, even token participation—if a member is making an effort, people do not assess the amount. Complete failure to participate in the collective

farming, however, produces tension and results in critical gossip, although the group in fact continues to feed a member who does not work" (260).

Along similar lines, Lallemand (1977), in a detailed study of two extended families in the Yatenga region, observes that "although the holders of the 'private' plots are meant to provide labor to the 'common' plot between eight in the morning and five in the evening, these time restrictions may be relaxed and flexible according to the vigilance or indulgence of the head worker" (46).¹

Since the rules regarding the allocation of one's labor across different farm plots cultivated by one's family are fluid and are not necessarily strictly enforced, a reasonable approach to adopt in the formal theory would be to model these allocations as voluntary contributions by the individuals concerned. However, this raises the question of why household members would make any contributions to the common plot at all beyond a nominal amount. One possible answer would be that such voluntary contributions are motivated by altruism. But the ethnographic literature points more strongly to an alternative explanation: the proceeds of the common plot are intended for the common good of the entire household. Hammond (1966) writes:

The produce from these fields (common plots) is stored in the granaries of the family head. There it is drawn upon principally to meet the individual needs of his household: to feed his wives and dependent children, to pay his tax, to be used as part of the bridewealth payment made by his sons when they marry, to offer hospitality to visitors, and to provide for the preparation of the millet gruel and millet beer used in sacrificial libations. But the millet in the granaries of the family head serves also as an emergency store to be used by any member of the family in need. (76)

Lallemand (1977, 100) provides a similar description of the use of the harvest from the common plot. It is likely that the head of the household is motivated, at least in part, by altruism in his decisions regarding the use of the proceeds from the common plot. However, we argue that a binding social rule is necessary to account for the pattern of usage described in the ethnographic literature, which is also borne out in our empirical analysis: since the extent of head's altruism, and the nature of this altruism, is likely to vary across households, altruism alone would lead to considerable heterogeneity across households in the usage of the proceeds from the 'common plot. A more regular pattern is more consistent with a social rule that is generally observed across households.

¹ Authors' translation of the following original text: "On sait que les détenteurs de ces champs personnels sont tenus de ne pas soustraire leurs efforts aux champs collectifs après huit heures du matin et avant cinq heures du soir; cependant, ces limites horaires sont plus souples et plus floues et dépendent surtout de la vigilance ou de l'indulgence du chef des travaux."

III. Theoretical Framework and Empirical Strategy

Imagine a household consisting of n adult members indexed by $i = 1, 2, \dots, n$.² There are m_0 possible consumption goods, of which m_1 are private goods and m_2 are household public goods ($m_0 = m_1 + m_2$). Let the vector $\mathbf{x}^i = (x_1, x_2, \dots, x_{m_1})$ denote private goods expenditures of household member i , and let $\mathbf{z} = (z_1, z_2, \dots, z_{m_2})$ denote a vector of public goods. Person i 's level of utility is given by $u^i(\mathbf{x}^i, \mathbf{z})$.

Each household member owns a plot of land with a vector of characteristics \mathbf{A}^i , which includes land area, soil quality, and so forth. There are m_k different crops that may be grown on a farm, indexed $k = 1, 2, \dots, m_k$. These crops may overlap with the private and public consumption goods. Each household member has access to the same production technology, described by

$$y_k = F_k(L_m, L_f, \mathbf{A}^i). \tag{1}$$

Here, y_k is the output level of crop k on person i 's farm plot, assuming that crop k has been planted on this farm (for simplicity, we assume that a single crop can be planted on a farm at any one time). Inputs L_m and L_f are, respectively, the levels of male and female labor used on person i 's farm plot.

Household member i has a labor endowment of L^i hours. Each person has the same labor productivity, except that male and female labor may vary in their relative productivity in different tasks. This last assumption is implicit in the description of agricultural production in (1). We assume that there is no labor market present, such that all farm work must be done using family labor (although this is a simplification, the data will show that there is only very limited use of hired labor in agriculture for our sample of households).

A. The Collective Model of the Household

First, we provide a characterization of the efficient allocation of resources within the household. The following maximization problem solves for such an allocation:

$$\max_{\substack{\mathbf{x}^1, \dots, \mathbf{x}^n, \mathbf{z} \\ (L_m^i, L_f^i, k_i)_{i=1}^n}} \sum_{i=1}^n \lambda^i u^i(\mathbf{x}^i, \mathbf{z}), \tag{2}$$

subject to

$$\mathbf{p}_z \mathbf{z} + \sum_{i=1}^n \mathbf{p}_x \mathbf{x}^i \leq \sum_{i=1}^n p_{k_i} F_{k_i}(L_m^i, L_f^i, \mathbf{A}^i), \tag{3}$$

² Hereafter, whenever we refer to a "household member," this will mean an adult member.

$$\sum_{i=1}^n L_m^i \leq \sum_{i \in \mathcal{I}_m} L^i, \tag{4}$$

$$\sum_{i=1}^n L_f^i \leq \sum_{i \in \mathcal{I}_f} L^i, \tag{5}$$

where \mathbf{p}_x , \mathbf{p}_z , and \mathbf{p}_k represent, respectively, a vector of prices of the private and household public goods, and the market price of crop k ; \mathcal{I}_s is the set of all household members belonging to gender s : $\mathcal{I}_s = \{i \in \mathcal{I} : i \text{ belongs to gender } s\}$; k_i is the choice of crop on the farm plot of household member i ; and λ^i is the Pareto weight assigned to individual i in the household welfare maximization problem.

It is well known that this problem satisfies the standard separation property of the agricultural household model (see, e.g., Bardhan and Udry 1999, chap. 1). In particular, Udry (1996) shows that if $F_k(\cdot)$ is increasing and concave in both types of labor, then the output levels are equal on plots that have the same characteristics and are planted with the same crop, that is, $F_{k_i}(L_m^i, L_f^i, \mathbf{A}^i) = F_{k_j}(L_m^j, L_f^j, \mathbf{A}^j)$ if $k_i = k_j$ and $\mathbf{A}^i = \mathbf{A}^j$. Then, adjusting the notation to allow for multiple households and time periods, the following specification enables us to test for efficiency in household production:

$$Q_{hct} = \mathbf{X}_{hct}\beta + \mathbf{G}_{hi}\gamma + \lambda_{hct} + \varepsilon_{hct}, \tag{6}$$

where Q_{hct} is the log of yield on plot i in year t , planted to crop c and belonging to household h ; \mathbf{X}_{ci} is a vector of physical characteristics of plot i ; λ_{hct} is a household-year-crop fixed effect; and \mathbf{G}_{hi} is a vector of individual characteristics of individual i in household h . Udry (1996) shows that if productive resources are allocated efficiently within the household, then $\gamma = 0$; that is, yield on a plot should not depend on the characteristics of the individual who controls it. Our data will enable us to replicate Udry’s test of productive efficiency within the household for our household sample.

Furthermore, efficient production implies that if $k_i = k_j$ and $\mathbf{A}^i = \mathbf{A}^j$, then $L_m^i = L_m^j$ and $L_f^i = L_f^j$; that is, if two farm plots with the same characteristics are planted to the same crop, then they should make use of the same amount of male labor and the same amount of female labor. Therefore, corresponding to (6), we can derive a specification for the level of each type of labor used on a particular plot:

$$l_{hct}^j = \mathbf{X}_{hct}\hat{\beta}^j + \mathbf{G}_{hi}\hat{\gamma}^j + \hat{\lambda}_{hct}^j + \hat{\varepsilon}_{hct}^j, \tag{7}$$

where l_{hci}^j is the amount of labor of type j applied to plot i per unit area, in year t , and plot i is planted to crop c and belongs to house h . Variable $\hat{\lambda}_{hct}^j$ is a household-year-crop fixed effect. Productive efficiency within the household implies that $\hat{\gamma}^j = 0$ for $j = m, f$.

B. A Model of Voluntary Contributions under a Social Norm

Next, we consider a noncooperative equilibrium for the model. We assume that each household member is able to allocate the proceeds of his farm plot and his labor endowment freely and that, because of lack of commitment, the members of the household cannot have a cooperative agreement that would enable them to implement the Pareto efficient outcome in household production. However, traditional institutions require that the proceeds of a common plot managed by the head of the household be spent entirely on certain household public goods, failing which the household head will be subject to social sanctions by the wider community.

For ease of exposition, we assume there are just three household members: the head of the household (represented by the letters h), who has a zero endowment of labor and manages only the common plot, and one adult male and another adult female member (represented by the letters m and f), who have labor endowments E^m and E^f , respectively, and manage their own private plots. There is a single crop, with market price of 1; a single private good x , with market price p_x ; and one household public good z , with market price p_z . We assume that there is no scope for saving. The social norm requires the household head to spend the proceeds from the common plot on the public good z .

Labor allocation and consumption decisions are made in two stages. In the first stage, household members m and f decide how to allocate their labor endowments across the three plots. Following these decisions, farm incomes are realized. Income from person i 's plot is given by $y_i = F(L_m^i, L_f^i, \mathbf{A}^i)$. In the second stage, m and f choose how to allocate their incomes across the different consumption goods. The household head h is constrained to spend all the income from the common plot on good z . Hence, we have $\bar{z}^h = (1/p_z)F(L_m^h, L_f^h, \mathbf{A}^h)$, where \bar{z}^h denotes the level of expenditures by the household head on the public good.

We assume that the household decisions constitute a subgame perfect equilibrium. Therefore, the equilibrium choices can be determined using backward induction. Given income $\mathbf{y} = (y^m, y^f, y^c)$ (representing incomes from the male, female, and common plots, respectively), the expenditure levels denoted by (\bar{x}^i, \bar{z}^i) , $i = m, f$, solve the following problems:

$$\begin{aligned}
 (\bar{x}^m, \bar{z}^m) &= \arg \max_{x^m, z^m} u^m(x^m, z^m + \bar{z}^f + \bar{z}^b) & (8) \\
 &\text{subject to } p_x x^m + p_z z^m \leq y^m,
 \end{aligned}$$

$$\begin{aligned}
 (\bar{z}^f, \bar{z}^f) &= \arg \max_{x^f, z^f} u^f(x^f, z^m + \bar{z}^f + \bar{z}^b) & (9) \\
 &\text{subject to } p_x x^f + p_z z^f \leq y^f.
 \end{aligned}$$

The interesting case occurs when both m and f make zero contribution to the public good z . In this case, we obtain $\bar{x}^m(\mathbf{y}) = y^m/p_x$, $\bar{x}^f(\mathbf{y}) = y^f/p_x$, and total expenditures on the public good, $\bar{z}(\mathbf{y}) = y^c/p_z$.

Given the functions $\bar{x}^m(\mathbf{y})$, $\bar{x}^f(\mathbf{y})$, and $\bar{z}(\mathbf{y})$, we can derive the equilibrium labor allocation decisions made by the household members. Let $\mathbf{L}^i = (L_m^i, L_f^i, L_c^i)$ represent the allocation of labor by household member i across the male, female, and common plots. Then, the equilibrium choices $(\bar{\mathbf{L}}^m, \bar{\mathbf{L}}^f)$ are given by

$$\begin{aligned}
 \bar{\mathbf{L}}^m &= \arg \max_{(L_m^m, L_f^m, L_c^m)} u^m(\bar{x}^m(\mathbf{y}), \bar{z}(\mathbf{y})) & (10) \\
 &\text{subject to } L_m^m + L_f^m + L_c^m \leq E^m,
 \end{aligned}$$

$$\begin{aligned}
 \bar{\mathbf{L}}^f &= \arg \max_{(L_f^f, L_c^f)} u^f(\bar{x}^m(\mathbf{y}), \bar{z}(\mathbf{y})) & (11) \\
 &\text{subject to } L_f^f + L_c^f \leq E^f,
 \end{aligned}$$

where $\mathbf{y} = (y^m, y^f, y^c)$, $y^i = F(L_m^i, L_f^i)$ for $i = m, f, c$.

Since all the proceeds from the private plots are used for private consumption, we obtain $L_f^m = L_m^f = 0$; that is, the household members do not contribute any labor to each other's private plots because doing so does not yield any benefits for themselves. Therefore, both m and f allocate all of his or her labor endowment between the common plot and his or her own private plot.

Thus, all household members are persuaded to provide some labor on the farm managed by the head of the household; given the social norm that proceeds from the common plot are to be used for household public goods only, both m and f can expect to reap some rewards from toiling on the common plot. However, they do not contribute any labor on each other's private farms because the managers of the private plots cannot commit to spend their in-

come on anything other than their own private goods. Thus, we obtain the result that the common plot is farmed with a mix of male and female labor, while the private plots are farmed almost entirely with labor of the same type as the gender of the plot manager.³

C. Predictions

These predictions can be tested by estimating (7). Let H_{bi} be a dummy variable that takes a value of one if individual i in household b is the household head, and zero otherwise; and let M_{bi} be a dummy variable that takes a value of one if individual i in household b is an adult male but not the household head, and zero otherwise. If $\mathbf{G}_{bi} = (H_{bi}, M_{bi})$, and correspondingly, $\hat{\gamma}^j = (\hat{\gamma}_H^j, \hat{\gamma}_M^j)$, then, under the noncooperative equilibrium analyzed above, we have $\hat{\gamma}_M^m > 0$, $\hat{\gamma}_M^f < 0$, and $\hat{\gamma}_H^m > 0$, $\hat{\gamma}_H^f > \hat{\gamma}_M^f$. In words, these conditions imply that (i) male plots use male labor more intensively than female plots with similar characteristics and planted to the same crops; (ii) female plots use female labor more intensively than male plots with similar characteristics and planted to the same crops; (iii) the common plot uses male labor more intensively than female plots and female labor more intensively than male plots, after controlling for the type of crop and the plot characteristics.

Note that the model also predicts that junior household members do not provide any labor on each other's private plots, but we acknowledge that these conditions may not hold in reality because of elements not captured in the model. For example, the junior household members may provide some labor on the private plots of family members as a result of side transactions. We discuss these possibilities after presenting the econometric results.

Furthermore, the model predicts that yields across different plots belonging to the same household can depend on the characteristics of the plot owner (after controlling for plot characteristics and the type of crop planted). In particular, if the junior household members value the household public good sufficiently, then the common plot will be farmed more intensively than either the male or the female plots. Therefore, we would expect yields to be higher on the common plots compared to the individual plots. In (6), if $\mathbf{G}_{bi} = (H_{bi}, M_{bi})$, and correspondingly, $\gamma = (\gamma_H, \gamma_M)$, then we have the prediction that $\gamma_H > 0$ and $\gamma_H > \gamma_M$. There is no clear prediction about the sign of γ_M .

If both m and f make voluntary contributions to the household public good, then equilibrium expenditures on consumption goods will be determined by total household income only, and these will be independent of the

³ Fafchamps (2001) discusses how land and labor transactions within the households are used to overcome the commitment problem.

income shares of the different household members (see Lundberg and Pollak 1993). In this situation, household members will allocate their labor across different plots so as to maximize total household income. Therefore, labor allocation within the household will be efficient, and the empirical predictions will correspond to those for the efficient model.

D. Consumption

The collective model of the household and the model of voluntary labor contributions introduced in Section III.B also provide testable predictions involving household consumption. First, in the case of the collective model, temporary shocks to household income should have no impact on expenditures on any particular good after total household expenditures have been controlled for. Otherwise, it would be possible to achieve better consumption smoothing by reallocating expenditures on that good across time periods. Duflo and Udry (2003) use this insight to devise a test for efficiency in household consumption using rainfall data to capture weather-related shocks. We replicate their test in this article using data on household public goods and income from the common and private plots. We briefly discuss here Duflo and Udry's test of efficiency in consumption.

We assume the following log-linear relationship between rainfall and household farm income:

$$\log(y_{hit}) = \mathbf{R}'_{vt} \lambda_{1i} + (\mathbf{P}'_{ht} \lambda_{2i}) \mathbf{R}_{vt} + \delta_i T_{94} + X_b + \xi_{hit}, \quad (12)$$

where y_{hit} represents period t income from plots belonging to household h managed by household members of type i (where i stands for either the household head, female household members, or junior male members); \mathbf{R}_{vt} is a vector of measures of the level of rainfall in province v (where household h is located) in period t ; \mathbf{P}_{ht} is a vector describing the characteristics of plots, such as soil type and topography, farmed by household h in period t ; T_{94} is a dummy that takes a value of one in the year 1994, and zero otherwise; X_b denotes household fixed effects to capture time-invariant household characteristics; and ξ_{hit} is the error term that captures other exogenous shocks that affect farm income for household h from plots managed by household members of type i in period t .⁴

The estimated coefficients are then used to compute a linear combination of the rainfall variables as follows: $\hat{y}_{hit} = \mathbf{R}'_{vt} \hat{\lambda}_{1i} + (\mathbf{P}'_{ht} \hat{\lambda}_{2i}) \mathbf{R}_{vt}$. These fitted

⁴ We do not include household plot characteristics, \mathbf{P}_{ht} , on their own in the specification, as the change in these characteristics from one year to the next is only slight, and therefore the effects are essentially subsumed in the household fixed effects.

values represent the component of household farm income that is explained by rainfall variations. If we assume that demand for each commodity can be expressed as a log-linear function of total expenditures, household Pareto weights, and other household and regional characteristics, then we can derive the following specifications relating household expenditures and income:

$$\log(e_{ht}) = \sum_{i=m,f,c} \pi_i \hat{y}_{hit} + \delta T_{94} + X_b + v_{ht}, \quad (13)$$

$$\log(x_{ht}) = \sum_{i=m,f,c} \pi_{xi} \hat{y}_{hit} + \delta_x T_{94} + X_b + v_{xb}, \quad (14)$$

where e_{ht} represents total expenditures in household h in period t and x_{ht} represents expenditures on some specific consumption good in household h in period t . The error terms v_b and v_{xb} capture other exogenous factors that may influence household expenditures. Under the collective model of the household, we have the following relationship:

$$\frac{\pi_{xi}}{\pi_i} = \frac{\pi_{xj}}{\pi_j} \text{ for each } i, j \in \mathcal{I}. \quad (15)$$

Following Duflo and Udry (2003), we test for (15) using a nonlinear Wald test.

Equation (14) also provides a test for the model of voluntary contributions under the social norm. Specifically, the model predicts that, if x is a household public good, then $\pi_{xm} = \pi_{xf} = 0$ and $\pi_{xc} = 1$; that is, public goods expenditures are responsive to exogenous shocks to income from the common plot (and, in particular, a one-on-one effect if the social norms require that all income from common plots be spent on household public goods), but they are not responsive to income shocks on private plots farmed by the household.

It should be noted that the estimates from (13) and (14) will be biased and the overidentification test in (15) will be invalid if our measure of rainfall does not satisfy the exclusion restriction. In particular, changes in prices of consumer goods may be correlated with rainfall, and rainfall may affect preferences between different household goods, including leisure. In these instances, the estimates of the coefficients π_i and π_{xi} would include a part of the direct effect of rainfall changes on household expenditures and demand for good x . We consider the validity of the identifying assumptions in more detail in Section VII.D after presenting our results.

IV. Description of the Household Survey

The household survey was conducted by the University of Laval, Quebec, and the University of Ouagadougou in Burkina Faso. It was carried out in four provinces in different parts of Burkina Faso: the Namentenga province in the Central Plateau, the Soum province in the North, the Kossi province in the West, and the Nahouri province in the Southeast. In each province, two villages were chosen, one to represent the wealthier villages and the other to represent the poorer ones in that region. A sample of 35 households was randomly selected in each village except in the North, where 40 households were chosen in anticipation of a higher dropout rate due to out-migration. Agriculture is the primary source of livelihood in all the villages, but livestock is relatively more important in the North, the region where population pressures and soil degradation have been most acute.

The survey was conducted during the 1993 and 1994 agricultural seasons and covered farm characteristics, production technologies, agricultural inputs and outputs, market activities, and household expenditures and consumption. Information relating to each farm plot was obtained from the individual in the household who had responsibility for it during that season, while information regarding the common plots was obtained from the head of the household. Information about expenditures, sales, livestock holding, and transfers was also gathered at the individual level, while information about housing and farming equipment was gathered at the household level, with the head of the household usually providing most of the information.

The household survey did not distinguish between the household head's private plots and common plots that he or she managed on behalf of the household. Therefore, for the purpose of the analysis in this article, we consider all plots farmed by the household head as common plots.

The questionnaire on household expenditures recorded information on the identity of the household member who made each purchase and, for the year 1994, also the identity of the person for whom the expense was incurred. We use the latter information to identify expenditures on household public goods. Specifically, we classify a particular good, say x , as a household public good if at least 75% of all reports of expenditures on x in the 1994 survey mentioned "all members of the household" or "other household members" as the intended beneficiary.⁵ This classification is then applied to all households in both survey years. The list of household public goods and private goods according to this classification is shown in the appendix in table A1.

⁵ This relatively high cut-off level was chosen since our goal is to identify a subset of expenditure items that we are confident are household public goods, not necessarily to capture all household public goods.

The household public goods include primarily types of food that require some amount of processing and preparation, typically done within the household, such as sorghum and millet flour/grain, raw meat, and fresh vegetables; as well as items used in food preparation, such as cooking oil, salt, spices and other condiments, and kitchen utensils. By contrast, the private goods include food items that may be consumed without any additional preparation, for example, doughnuts, cooked meals, local beer and other alcoholic beverages, tea, and coffee. Furthermore, private goods include durable goods, such as a radio, a bicycle, and a motorbike. All these items are, arguably, easier to exclude from other household members than unprocessed food items that must pass through the household kitchen. It should be evident from table A1 that adjusting the 75% rule leads to slightly different classifications, but this does not affect our qualitative results on household public goods and income shocks that will be reported in Section VI.D.⁶

V. Household Headship and Junior Household Members

As the theoretical framework used in this article makes the distinction, in the context of agricultural production, between gender and household headship, we explore in the data how closely these two variables are correlated. Only 22 of the random sample of 290 households interviewed for the CEDRES/Laval survey in 1993 were headed by women. Therefore, it is evident that household headship is predominantly a male phenomenon.

The female heads are invariably widows, often the sole adult in a household with young children or grandchildren, and in some instances living with unmarried sons or daughters. The distinctive structure of female-headed households raises the question whether the social norms described in Section II on the allocation of productive resources within the household should apply to them. Given the small number of households with female heads in our sample, we are not able to answer this question satisfactorily using the available data. For the main analysis, we do not distinguish between male-headed and female-headed households. Nevertheless, we return to this question after the main results have been presented.

To be able to test the model of voluntary labor contributions presented in Section III.B, we need to be able to compare farm plots managed by a male household head with those managed by other—both male and female—household members. Table 1 shows, for all plots farmed by the households in the CEDRES/Laval sample during the years 1993 and 1994, the rela-

⁶ Results based on alternative rules are not reported in this article but are available from us on request.

TABLE 1
RELATIONSHIP OF PLOT MANAGER TO HOUSEHOLD HEAD
 (% OF TOTAL SAMPLE)

	Male	Female
Self	65.3	3.6
Spouse	.1	17.5
Son/daughter	6.7	1.5
Son/daughter-in-law	.1	1.3
Grandchild	.3	.1
Sibling	1.5	.0
Spouse of sibling	.0	.9
Nephew/niece	.1	.2
Father/mother	.1	.5
Other relation	.1	.3

relationship of the plot manager to the household head. About 69% of these plots were farmed by the household head himself (or herself). For the purpose of our empirical analysis, these plot observations will be classified as “common” plots. In the remainder of the cases, the plot is managed (in order of decreasing frequency) by a spouse of the head, a son or daughter, a brother, a daughter-in-law, or a sister-in-law. These observations will serve as the “private” plots in this study. The managers of the private plots are henceforth referred to as “junior male” and “junior female” household members.

Table 2 shows that the common plots, as defined here, are, on average, about twice as large as the private plots farmed by the junior men and three times as large as the private plots farmed by the junior women. However, the common plots have the lowest average yields: about 15% lower than on junior female plots and barely half of that on junior male plots.

Table 2 also indicates the use of very different types of labor on the farm plots depending on the identity of the plot owner. The junior female plots are farmed primarily with female family labor, which very likely includes hours that the plot owner has herself spent working on the plot. The junior male plots use both male and female family labor, but the contribution of male labor is more significant (about twice as large) than that of female family labor. The common plots use substantial amounts of both male and female family labor. All three types of plots use some amount of child family labor.

The descriptive statistics show that the nature of production on common plots is, on average, very different from that on private plots, at least in terms of the scale of farming and the types of labor used. This differentiated pattern does not, in itself, imply that the allocation of labor across the household plots is inefficient. The pattern may be explained by the different labor requirements of crops grown by different household members, by differences in soil

TABLE 2
MEAN YIELD, AREA, AND LABOR INPUTS BY DIFFERENT TYPES OF CULTIVATOR

	Yield (FCFA/ Hectare)	Area (Hectare)	Child Labor (Labor Days/ Hectare)	Male Labor (Labor Days/ Hectare)	Female Labor (Labor Days/ Hectare)	Observations
Male heads	77,042 (305,190)	1.18 (1.26)	11.71 (20.9)	44.27 (53.8)	29.96 (45.3)	1,420
Female heads	177,185 (691,669)	.58 (.40)	6.19 (17.8)	21.17 (37.6)	46.29 (48.2)	67
All house- hold heads	81,579 (332,740)	1.15 (1.24)	11.43 (20.7)	43.05 (53.3)	30.81 (45.6)	1,487
Junior men	150,545 (589,342)	.60 (.57)	8.98 (20.1)	38.97 (51.0)	21.96 (42.4)	160
Junior women	95,392 (433,985)	.45 (.52)	7.35 (16.6)	9.17 (16.2)	40.9 (60.9)	333

Note. The table excludes, for each variable, the top percentile of observations. Due to missing values, labor averages are based on a total of 1,600 observations. The CFA franc (FCFA) refers to the currency of the monetary union Communauté Financière Africaine, which includes Burkina Faso. At the time of the survey, the approximate exchange rate was US\$1 = FCFA 800. Standard errors are in parentheses.

type, and so forth. Therefore, we estimate equations (6) and (7), which will enable us to see if these patterns persist after these factors have been controlled for.

VI. Results

A. Plot Yields

Table 3 shows the estimated coefficients for equation (6) with dummy variables for household headship and junior male status within the household (the default category being junior female status). The estimated coefficient for the household head dummy is large and significantly different from zero. Therefore, we are able to reject the hypothesis of Pareto efficiency in household production. The coefficient for the junior male dummy variable (which takes a value of one if the plot owner is an adult male who does not head the household and zero otherwise) is small and not statistically different from zero. Therefore, for plots owned by the same household that have similar characteristics and are planted to the same crops in the same year, there are no significant differences in yields between those controlled by junior women and those controlled by junior men.

However, the estimated coefficients are consistent with the model of voluntary contributions. Recall, from Section III.C that the model predicts that yields on common plots will be higher than those on private plots. The estimated coefficient for the household head dummy is greater than zero and

TABLE 3
PLOT YIELD

	Dependent Variable: Ln(Plot Yield in FCFA)	
	Coefficient	SE
Junior male dummy	.0127	.123
Male household head dummy	.7747	.105
Plot size (1st decile omitted):		
2nd decile	-.6671	.149
3rd decile	-.4843	.156
4th decile	-.9434	.162
5th decile	-1.208	.162
6th decile	-1.447	.175
7th decile	-1.254	.171
8th decile	-1.616	.181
9th decile	-1.854	.182
10th decile	-2.037	.178
Toposequence (mid-slope omitted):		
Near bottom	-.492	.268
Plain	-.461	.250
Soil type (clay omitted):		
Sand	.0326	.146
Laterite	-.3464	.187
Gravel	-.2647	.288
Location (village excluded):		
Outside village	.0225	.104

Note. Observations = 1,967. Household-crop-year fixed effects are included in each regression. The CFA franc (FCFA) refers to the currency of the monetary union *Communauté Financière Africaine*, which includes Burkina Faso. At the time of the survey, the approximate exchange rate was US\$1 = FCFA 800.

greater than the coefficient for the junior men, and the difference is statistically significant in both cases.

Note that these results are consistent with findings in the existing literature, which have shown, in similar contexts, that plot yields are higher for men than for women when comparing across farms owned by the same household that are planted to the same crops and with similar characteristics (see, e.g., Udry 1996; Goldstein and Udry 2008). However, our findings show that, for our sample of households, the higher yields achieved by men can be attributed entirely to household headship, since there is no statistically significant difference in plot yields between plots farmed by men and women who do not head the household in which they live.

The large, positive coefficient for the household head dummy obtained in the yield regression contrasts with the fact that common plots have lower yields on average than private plots, as shown in table 1. This is primarily due to the fact that yields decline with plot size (which is evident from the plot size coefficients in table 3) and that common plots are, on average, larger

than private plots. The estimated coefficients for equation (6) imply that, for a given plot size, yields are higher on common plots than on private plots.

B. Dispersion of Yields

Following a procedure adopted by Udry (1996), we compare the dispersion of plot yields within the village, the household, and across plots farmed by the same individual. Our primary aim in conducting this exercise is to explore how effective the household head is in allocating productive resources efficiently compared to junior household members and the household as a whole. The ethnographic literature discussed in Section II and the evidence in the preceding subsection suggests that the household head has better command over productive resources and, in particular, family labor, than do junior household members. If this is so, the head has better means to achieve efficient allocation of resources across his plots than would the village, the household as the whole, or junior household members.

We drop the term G_{hi} from (6) and reestimate the equation after replacing λ_{hc} with village-crop-year fixed effects. Figure 1 shows the kernel estimate of the density of the error term, which represents the deviation of yields, for each plot in our sample, from the mean yield of the corresponding village-crop-year group. The figure gives an indication of the variation in output within each village across apparently identical plots, due to risk as well as the misallocation of resources within the village. The same figure shows the de-

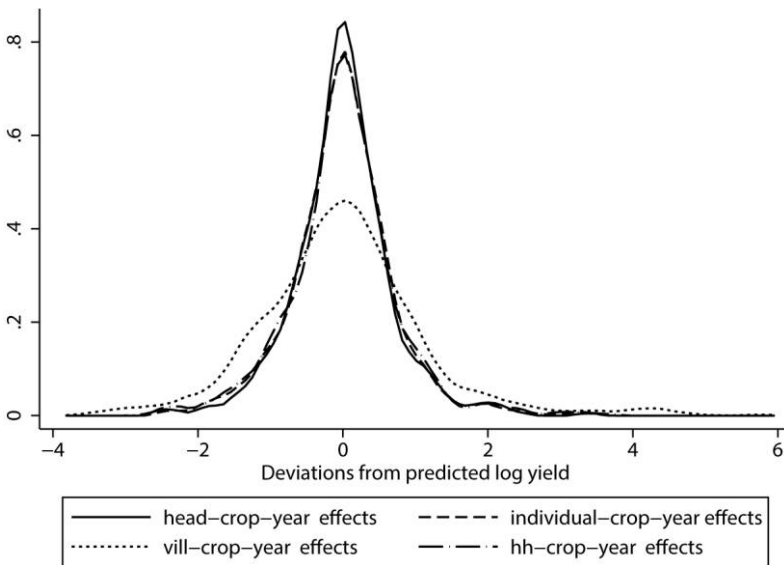


Figure 1. Kernel estimates of density of yield function error terms

viation of yields from the mean yield of the corresponding household-crop-year-group and from the mean yield of the corresponding individual-crop-year-group. The dispersion around the household mean is noticeably smaller than the dispersion around the village mean. It is evident from the figure that, as with the ICRISAT data set used in Udry's (1996) study, the household serves to allocate productive resources more efficiently across farm plots. The deviations around the individual mean are virtually indistinguishable from the deviations around the household mean. Thus, the contrast between the individual and the household as vehicles for the efficient allocation of productive resources is much smaller than the contrast between the household and the village.⁷

To compare the effectiveness of the household head in allocating productive resources efficiently, with the individual and the household, we restrict the sample to plots managed by the household head and recompute the deviation of yields from the individual-crop-year group. If the household head were similar to other members of the household in his or her ability to distribute productive resources efficiently, then this last density function would mirror those for the individual. Visually, we find that the dispersion for the household head is noticeably smaller.⁸ Thus, the figure provides suggestive evidence that the head has better control over the allocation of productive resources across his or her farm plots than does the junior household member, and, in particular, that he or she is better able to call upon family members to work on his farm plots. Next, we trace the difference in plot yields noted in Section VI.A to variations in labor use across plots.

C. Labor Allocation

Table 4, column 1, shows the estimated coefficients for equation (7). The results show that plots controlled by the household head use all types of family labor—adult male, adult female, and child—more intensively than other plots owned by the same household, after controlling for plot characteristics, the crop planted, and the year of planting (with the difference being statistically significant in each case). This pattern is inconsistent with the predictions of the

⁷ Using the two-sample Kolmogorov-Smirnov test for the equality of distributions, we can reject the null hypothesis that the village-level and household-level distributions are identical (p -value = .00), but we cannot reject the hypothesis that the household-level and individual-level distributions are identical (p -value = .873).

⁸ However, it should be noted that, using a two-sample Kolmogorov-Smirnov test of the equality of distributions, we cannot reject the null hypothesis that the individual-level distribution for junior household members and the corresponding distribution for the household head are identical (p -value = .664).

TABLE 4
LABOR INTENSITY

	Male Labor (1)	Female Labor (2)	Child Labor (3)
Junior male dummy	1.211 (.137)	-1.163 (.126)	.2114 (.131)
Male household head dummy	2.162 (.114)	.5196 (.104)	.8146 (.108)
Plot size (1st decile omitted):			
2nd decile	-.6086 (.178)	-.5258 (.164)	-.7544 (.170)
3rd decile	-.9657 (.188)	-.8277 (.173)	-.7673 (.180)
4th decile	-1.312 (.195)	-1.389 (.179)	-.8364 (.186)
5th decile	-1.474 (.188)	-1.317 (.173)	-1.046 (.179)
6th decile	-1.723 (.198)	-1.71 (.182)	-1.133 (.189)
7th decile	-1.757 (.201)	-1.814 (.185)	-1.083 (.192)
8th decile	-1.643 (.203)	-1.881 (.187)	-1.012 (.194)
9th decile	-1.872 (.210)	-2.267 (.193)	-1.035 (.201)
10th decile	-2.204 (.217)	-2.577 (.199)	-1.21 (.207)
Toposequence (near bottom omitted):			
Mid-slope	.3143 (.304)	.3994 (.279)	.2027 (.290)
Plain	-.1746 (.192)	-.0093 (.177)	-.3498 (.184)
Soil type (clay omitted):			
Sand	-.2028 (.152)	-.0766 (.140)	.2597 (.146)
Laterite	.0724 (.177)	-.1066 (.163)	.4236 (.169)
Gravel	.3381 (.308)	.0471 (.284)	.8797 (.295)
Location (village excluded):			
Outside village	.2246 (.112)	-.0013 (.103)	-.0065 (.107)

Note. Dependent variable = $\ln(\text{labor days per unit area on plot of land})$. Observations = 1,979. Household-crop-year fixed effects are included in each regression. Standard errors are in parentheses.

Pareto efficient (collective) model. As discussed in Section III.A, the model predicts that the allocation of labor across plots can depend on the plot characteristics and the crop planted but that it should be independent of the characteristics of the plot owner.

While plots farmed by junior men use male family labor more intensively than those farmed by junior women, the opposite is true for female family labor, by about the same order of magnitude. There are no significant dif-

ferences in the use of child labor between the junior men and junior women. The pattern of labor allocation across agricultural plots appears consistent with the variation in plot yields discussed earlier, with the common plots achieving the highest yields and also making use of family labor and outside labor most intensively.

This pattern is also, at least in part, consistent with the predictions of the model of voluntary labor contributions discussed in Section III.D. Recall that one of the model's predictions is that the common plot uses male family labor more intensively than private plots farmed by women and female family labor more intensively than private plots farmed by junior men. If we interpret the plots farmed by the household head as common plots, then the estimated coefficients correspond to these predictions.

However, the model of voluntary labor contributions also generated the prediction that household members do not provide any labor on each others' private plots, and the descriptive statistics in table 1 show that they clearly do (since family labor of the opposite gender is used on private plots). We speculate that such labor contributions, which are much smaller than contributions to the common plots, may be due to side transactions between the junior male members and the women within the same household, a dynamic flow of goods and services in which labor is just one element. Hammond (1966) describes some exchanges along these lines: a junior male, with a bicycle, may help out a female household member to carry some of her produce to a distant market and, in exchange, she may spend some time working on his plot during the farming season. Our theoretical model did not allow for such exchanges.⁹

The pattern of labor allocation indicates that the household head is able to induce family members to provide labor on common plots much more effectively than the junior men and women. It also shows that while these junior household members have the assistance of some family labor of the opposite gender on their private plots, they are relying primarily on labor of their own gender. This labor of their own gender is potentially their own

⁹ An alternative hypothesis is suggested by the ethnographic studies on the Mossi discussed in Sec. II. A junior household member, if married, usually cultivates the private plots assigned to him with the assistance of his own nuclear family, living under the same household head (Hammond 1966; Fiske 1991) and may have similar rights and obligations in relation to them as the household head does in relation to all household members. This would imply positive levels of female family labor on junior male plots but no corresponding use of male family labor on female plots. However, the latter prediction does not fit the data well: table 4 shows similar intensities in the use of male family labor on female plots and the use of female family labor on male plots (after controlling for plot characteristics and the crop being grown), and therefore this hypothesis, on its own, does not explain the labor patterns in the data adequately.

labor, although there is no way of verifying this on the basis of the available data.

Why is the head of the household so much more successful at commanding family labor? The model presented in Section III.B posits that it is because the household head is committed to spending the proceeds from the common plot on household public goods. We analyze consumption data next to see if this hypothesis has any support in the data.

D. Expenditures on Household Public Goods and Rainfall Shocks

Using data on seasonal rainfall for each province in Burkina Faso, we estimate equation (12) for three different sources of household income: junior female, junior male, and household head farm income. The specification includes monthly rainfall during the months June, July, August, and September in the current year and annual rainfall during the previous year, interacted with plot characteristics. The monthly measures for rainfall in the current year were chosen as the June to September period corresponds with the planting season when adequate rainfall is critical. The level of rainfall during the previous year can affect the moisture content in the soil (see, e.g., Xie et al. 2002), which can be an important determinant of soil fertility (Marten and Vityakon 1986), but the timing of the rainfall is, arguably, unimportant in this regard. For a given level of rainfall, the amount of moisture can also depend on topography and type of soil in the farm plot; hence, we interact previous year's total rainfall with plot characteristics. However, the optimal timing for rainfall during the planting reason does not depend on plot characteristics per se but on the crops being planted; therefore, we do not interact current year's rainfall with plot characteristics or with the crop choices, as these are endogenous variables.

The coefficients from the first-stage regressions are shown in table 5. The *F*-tests indicate that the coefficients for the rainfall variables and the plot characteristics and rainfall interactions are jointly significant in all three regressions. In the presence of multiple endogenous regressors, however, the first-stage *F*-test is not sufficient to assess weak identification. For such cases, Stock, Wright, and Yogo (2002) and Stock and Yogo (2005) recommend the Cragg-Donald weak instrument test statistic (Cragg and Donald 1993) as an alternative to the traditional *F*-test. We will return to the Cragg-Donald test when we discuss the second-stage results.

Using the fitted values for the different types of farm income and expenditures, we estimate equation (13) and also (14) for household public goods. The coefficients from the second-stage regression provide a measure of the responsiveness of household expenditures to changes in household income due to

TABLE 5
HOUSEHOLD EXPENDITURES AND INCOME: FIRST-STAGE REGRESSION RESULTS

	Household Head	Junior Male	Female
Current-year rainfall, amount below historical average in:			
June	-.07745*** (.01928)	-.00132 (.01833)	.05651*** (.02106)
July	-.62151*** (.09562)	-.24146*** (.09090)	-.47169*** (.10448)
August	.16894 (.04257)***	.00503 (.04047)	-.17768*** (.04652)
September	-.76340 (.12168)	-.32114*** (.11567)	-.41689*** (.13296)
Previous year's rainfall interacted with household plot characteristics:			
Junior male plots:			
Plain	.00010** (.00004)	-.00004 (.00004)	-.00003 (.00005)
Laterite	.00007 (.00011)	.00018* (.00010)	.00000 (.00012)
Junior female plots:			
Plain	-.00110*** (.00040)	.00260*** (.00038)	.00055 (.00044)
Laterite	-.00340** (.00132)	.00658*** (.00125)	-.00082 (.00144)
Household head plots:			
Plain	-.00162** (.00067)	.00032 (.00063)	.00700*** (.00073)
Laterite	.00149 (.00105)	-.00129 (.00100)	-.00261** (.00115)
1994 year dummy	91.383*** (15.19390)	45.042*** (14.44245)	92.803*** (16.60152)
Constant	146.947*** (23.14951)	58.022*** (22.00459)	137.588*** (25.29416)
Test for all coefficients being jointly significant <i>F</i> -statistics (<i>p</i> -values)	5.747	14.90	18.10

Note. Dependent variable = $\ln(\text{income from plots managed by } X)$. Observations = 550. Standard errors are in parentheses.

* Significant at the 10% level.

** Significant at the 5% level.

*** Significant at 1% level.

exogenous rainfall shocks. The second-stage estimation results, with clustered standard errors for each household across the two time periods, are shown in table 6. We report estimates of two-stage least squares (2SLS) for household public goods in column 3. The coefficients indicate that a 10% increase in the household head's income due to variations in rainfall leads to a 4.03% increase in expenditures on household public goods. This effect is statistically significant at the 1% level. By contrast, the corresponding values for income from junior male private plots is 2.20% (significant at the 1% level), and for income from junior female private plots it is statistically insignificant. Using

TABLE 6
HOUSEHOLD EXPENDITURES AND INCOME: SECOND-STAGE REGRESSION RESULTS

Dependant Variable (Fitted Values from First Stage)	Ln(Total Expenditures)		Ln(Public Goods Expenditures)	
	2SLS (1)	LIML (2)	2SLS (3)	LIML (4)
Ln(household head income)	.421*** (.137)	.463*** (.148)	.549*** (.071)	.629*** (.077)
Ln(junior male income)	.220*** (.079)	.247*** (.086)	.249*** (.082)	.295*** (.105)
Ln(junior female income)	.072 (.055)	.068 (.063)	.100 (.064)	.092 (.080)
1994 year dummy	.449* (.268)	.467 (.298)	.227 (.368)	.269 (.419)
χ^2 tests:				
χ^2 : coefficient for household head = junior male	3.173	2.843	8.154	6.97
Prob > χ^2	.07	.09	.00	.01
χ^2 : coefficient for household head = female	6.448	6.568	16.14	15.16
Prob > χ^2	.01	.01	.00	.00
χ^2 : coefficient for junior male = female	2.086	2.26	1.428	1.563
Prob > χ^2	.15	.13	.23	.21
Hansen J-test for overidentification:				
$\chi^2(7)$	4.276	4.003	5.008	4.645
Prob > χ^2	.75	.78	.66	.70
Cragg-Donald Wald F-statistic	6.568	6.568	6.568	6.568

Note. 2SLS = two-stage least squares; LIML = limited information maximum likelihood. In the top portion of the table, standard errors, shown in parentheses, are clustered at the household level across the two time periods.

* Significant at the 10% level.

*** Significant at 1% level.

chi-square tests, we can reject the null hypothesis of equality between the household head and the junior male coefficients and between the household head and junior female coefficients at the 1% level. The null hypothesis of equality between the coefficients for the junior male and junior female household members cannot be rejected.

We report the Cragg-Donald Wald test statistics at the bottom of table 6. We can reject the null that the 2SLS bias may exceed 20% of the OLS bias (a statistic of 6.57 vs. a critical value of 5.83) but not that it exceeds 10% of the OLS bias.¹⁰ Therefore, weak instrument bias may be a concern with our 2SLS estimates. To address this concern, we estimate the model using limited information maximum likelihood (LIML), a procedure that is known to be more robust to weak instruments than 2SLS (Stock et al. 2002; Stock

¹⁰ For the case of three endogenous variables and 10 excluded instruments, the critical values for the null that the IVE bias may exceed 20%, 10%, or 5% of the OLS bias are 5.83, 9.64, and 16.80, respectively. For tabulation of the critical values, see Stock and Yogo (2005, 100).

and Yogo 2005). The LIML estimates shown in column 4 of table 6 are consistent with the 2SLS estimates shown in column 1. Moreover, we have exactly the pattern of rejection from the chi-square tests.

The Cragg-Donald test statistic for LIML, reported in the last row of table 6, column 4, is 6.57. Stock and Yogo do not provide a critical value for LIML with more than one endogenous variable. However, Bernal and Keane (2011) argue that the critical value for the null that the bias for the LIML test statistics is no greater than 10% of the OLS bias is roughly 5.4 in a model with three endogenous variables and 78 instruments. Given that the critical value increases with the number of instruments, it is apparent that in our case—three endogenous variables with 10 instruments—the critical value for rejecting the null, that the LIML statistic is no greater than 10% of the OLS bias, is smaller than 5.4. This value is easily exceeded by our Cragg-Donald statistic of 6.57. Thus, there is no evidence of weak instrument bias with the LIML estimates.

Because our equations are overidentified, we can test whether the excluded instruments are appropriately independent of the error terms. Therefore, we present the Hansen J -test for overidentification. In both cases, we cannot reject the joint null hypothesis that the instruments are valid at the 10% level. This supports the plausibility that our instruments are “truly” exogenous and can be excluded from the second-stage regressions.

Overall, it is apparent that the household head has a higher marginal propensity to spend on household public goods out of income from farms managed by himself or herself than do the junior household members. In this case, the household members would have a stronger incentive to provide labor on the fields managed by the household head than on the private plots of the other household members, and this can explain the pattern of allocation family labor across household plots described in the preceding section.

The results are also consistent with the reasoning underlying the model of voluntary contributions under a social norm presented in Section III.B. However the estimated coefficients do not match exactly with the predictions of the model, and therefore we discuss briefly why. First, the model predicts that, for household public goods, the coefficient π_{xc} in (14) should equal one. However, if we allow for intertemporal savings within the model, then an argument akin to the Permanent Income Hypothesis should imply that a temporary income shock, as rainfall variations would produce, should not lead to a one-for-one change in expenditures, even if the household head is socially obliged to spend the proceeds from his (common) farm plots on household public goods. This would explain why $\hat{\pi}_{xc} < 1$ in the second-stage regression. Second, the model predicts that $\pi_{xm} = \pi_{xf} = 0$ (with x denoting

expenditures on household public goods). However, if there are multiple household public goods, and the corner solution in consumption expenditures assumed in Section III.B applies to only a subset of them, this would explain why we obtain $\pi_{xm} > 0$ and $\pi_{xf} > 0$.

We report the estimated coefficients for total household expenditures (eq. 13) in columns 1 and 2 in table 6. We obtain similar values for the Cragg-Donald weak instruments test and the Hansen J -test for overidentification as with equation (14) and therefore conclude that there is no evidence for weak instrument bias or invalidity of the instruments. The estimated coefficients for total expenditures are somewhat smaller than the corresponding estimates for household public goods, although the differences are not statistically significant in every case. Thus, there is some evidence that household members are able to smooth expenditures on private goods better than for household public goods. We test for efficiency in consumption expenditures using equation (15) for the category household public goods as defined above. The null hypothesis of efficiency is strongly rejected, with a Wald statistic of 18.32 and a p -value of .00038.

We should note that, according to the estimates in table 4, the common plot, on average, uses all types of family labor more intensively than private plots managed by household members of either gender. This suggests that the marginal product of family labor is higher on a private plot than on a household common plot. But if so, can the promise of household public goods be sufficient inducement for an individual, with his or her own parcel of land, to work on the common plot? Could he (she) not provide the public good more cheaply by allocating more of his (her) labor to his (her) private plot at the expense of the common plot and using the proceeds to pay for the public good?

We offer two possible answers to this question. First, the pattern of allocation of family labor suggests that the model of voluntary labor contributions is not capturing all the tools at the disposal of the household head to induce family members to work on the common plot; that is, the patriarch can threaten or coax them in ways not captured in the model.

A second explanation is that labor activities on the common plot follow a more specialized division of family labor than the private plots. If the family members are given specific and well-defined tasks on the common plot, than any withdrawal of labor could have a significant impact on production. If so, the marginal product of each person's labor may be higher on the common plot than on his or her private plot, even if the common plot is being farmed more intensively. This may happen even if the labor inputs of different family members are close substitutes in terms of the production possibility set. Un-

fortunately, in the absence of more detailed information on labor activities on farm plots, it is not possible to verify the second hypothesis.

E. Female Headship

To determine whether the pattern of production and consumption previously discussed also extends to female-headed households, we reestimate equations (7) and (12) with separate dummies for male and female household heads. The plot-level data includes 67 observations where the plot manager is a female household head. The estimated coefficients for the main variables of interest are shown in table 7. In the case of plot yields, the estimated coefficient for the female head dummy is smaller than that of the male head and larger than that of a junior male household member. However, because of large standard errors, the difference is not statistically significant in either case.

In the case of the allocation of family labor, we find that the male and female household heads have similar levels of labor intensity in the case of both adult male and adult female labor. Interestingly, the female head uses child labor much more intensively than the male head, with the difference being statistically significant. This may be due to the fact that some of the female-headed households consist of widowed grandmothers living with, possibly orphaned, grandchildren. The female household head also uses all types of family labor more intensively than junior male and female members of the same household, but the differences are not statistically significant in every case). Therefore, there is some suggestive evidence that female-headed households exhibit, broadly, similar patterns in the allocation of productive resources as the male-headed households, at least along the dimensions measured. But given the small sample size, these results are, at best, tentative.

TABLE 7
ESTIMATED COEFFICIENTS FOR FEMALE-HEADED HOUSEHOLDS

	Dependent Variable: Ln(Plot Yield)	Dependent Variable: Ln(Labor Days per Unit Area on Plot of Land)		
		Male Labor	Female Labor	Child Labor
Junior male dummy	.021 (.12)	1.216 (.14)	-1.164 (.13)	.206 (.13)
Male head dummy	.778 (.11)	2.185 (.12)	.517 (.11)	.790 (.11)
Female head dummy	.377 (.75)	1.658 (.48)	.582 (.44)	1.348 (.46)

Note. Standard errors are in parentheses.

VII. Testing for Alternative Hypotheses

A. Unobserved Plot Characteristics

In this section we consider a number of alternative explanations for the differences in agricultural yields between common plots and private plots. First, we explore the possibility that the yield differentials across different types of plot managers shown in table 3 are being driven by unobserved plot characteristics. Although the regression controls for available plot characteristics, it should be noted that the information on soil quality and toposequence available in the CEDRES/Laval data set, which we use, is very limited. When considering the same question, Udry (1996) showed that when controls for plot characteristics are removed, the gender differential in yields disappear, implying that female plots are of better quality than male plots, at least along the dimensions we are able to observe. In the CEDRES/Laval data set, we obtain a similar pattern, as shown in table 8. If soil quality, toposequence, and plot location are removed from the plot yield equation but plot size is retained, the coefficient for household headship is similar in magnitude to that in the original specification and statistically significant at the 5% level. If all plot characteristics are removed, then the corresponding coefficient is only one-fifth as large as in the original specification and no longer statistically significant. The differential between junior male and junior female plots remains small and statistically insignificant in both these alternative specifications. These results imply that the farm plots managed by the junior household members are, in fact, of better quality than (in terms of plot size) or of similar quality to (in terms of soil quality, toposequence, and location), those managed by the head of the household along the dimensions observed. This evidence goes against the hypothesis that the yield advantage of household heads may be due to unobserved differences in plot quality.

B. Varying Distance to Farm Plots

Second, we note that some farm plots are at a considerable distance from the homestead of the household responsible for its cultivation. In these circumstances, the cost of traveling to the plot may be nonnegligible and therefore, as Udry (1996) shows formally, applying labor inputs with varying levels of intensity on plots with similar characteristics planted with the same crops may be consistent with efficiency in household production. For example, to save on traveling time, it may be efficient for one household member to apply all his labor to a farm plot that lies 2 kilometers east of the homestead, while another applies all her labor to a farm plot which lies 2 kilometers west, even if their labor inputs are not perfect substitutes.

TABLE 8
ALTERNATIVE SPECIFICATIONS FOR PLOT YIELD EQUATION

	Household-Crop- Year Fixed Effects		Household-Crop- Year Fixed Effects and Plot Size		Original Specification	
	Coefficient	SE	Coefficient	SE	Coefficient	SE
Junior male dummy	-.0619	.122	-.0419		.0127	.123
Male household head dummy	.1502	.090	.6948		.7747	.105
Plot size (1st decile omitted):						
2nd decile			-.6984	.143	-.6671	.149
3rd decile			-.4548	.149	-.4843	.156
4th decile			-.8705	.155	-.9434	.162
5th decile			-1.1463	.153	-1.208	.162
6th decile			-1.3222	.165	-1.447	.175
7th decile			-1.2102	.163	-1.254	.171
8th decile			-1.5405	.170	-1.616	.181
9th decile			-1.6982	.167	-1.854	.182
10th decile			-1.9068	.169	-2.037	.178
Toposequence (mid-slope omitted):						
Near bottom					-.492	.268
Plain					-.461	.250
Soil type (clay omitted):						
Sand					.0326	.146
Laterite					-.3464	.187
Gravel					-.2647	.288
Location (village excluded):						
Outside village					.0225	.104
Observations		2,068		2,068		1,967

Note. Dependent variable = $\ln(\text{plot yield in CFA francs})$. Household-crop-year fixed effects are included in each regression.

However, even if the cost of traveling to distant farm plots were a significant factor, the above reasoning would not apply to those located close to the homestead. Efficiency would still imply similar yields across plots with similar characteristics, planted to the same crops, if they are near the homestead. This appears not to be the case in the data. Table 9 shows that the difference in yields between plots farmed by the household head and those farmed by other household members continue to hold when we restrict the sample to plots less than 4 kilometers, 3 kilometers, 2 kilometers, or 1 kilometer from the homestead. This pattern is very similar to that obtained by Udry (1996).

C. Misspecification of the Relationship between Yield and Plot Size

Next, we consider the possibility that the estimated yield differential across plots managed by different household members, shown in table 3, may be driven in part by a misspecification of the relationship between yield and plot size. In particular, if the relationship between yield and plot size varies across common and private plots, then, following Udry's reasoning, this may lead to

TABLE 9
ESTIMATED YIELD DIFFERENTIAL ACROSS DIFFERENT PLOT MANAGERS BY DISTANCE TO HOMESTEAD

	Distance to Plot from Homestead					
	All Plots (1)	<4 Kilometers (2)	<3 Kilometers (3)	<2 Kilometers (4)	<1 Kilometer (5)	Inside Village (6)
Junior male dummy	.0127 (.123)	-.0182 (.133)	.0186 (.140)	.0727 (.152)	.0184 (.166)	.0200 (.183)
Male head dummy	.7747 (.105)	.7553 (.113)	.7623 (.116)	.7518 (.123)	.7161 (.140)	.6949 (.156)
N	1,967	1,782	1,697	1,595	1,452	1,215

Note. Standard errors are in parentheses.

an overestimate of the slope of the yield-versus-plot-size relationship and an upward bias in the estimated yield differential between the household head’s plots and those managed by other household members. Following Udry (1996), we perform a nonparametric regression of plot yields on plot size, controlling for household-crop-year fixed effects for three categories of household members: household heads, junior males, and junior females.

The results are shown in figure 2. We find that for plots larger than 0.4 hectares, the yields on the plots managed by the junior household members are consistently below those managed by the household head. For plots smaller than 0.4 hectares, the yields across the three types of plots are similar.



Figure 2. Nonparametric regression of plot yield on plot area with household-year-crop fixed effects

finding is consistent with the idea that the smaller plots managed by the household head are actually private plots, which we cannot distinguish from the common plots in our data. If so, the observed pattern remains consistent with the predictions of the model introduced in Section III.B but inconsistent with efficiency.

D. Direct Impact of Rainfall on Household Consumption

As noted in Section III.D, if rainfall affects household consumption patterns directly, independently of its effects on household income, then this will lead to a bias in the estimated coefficients for equations (13) and (14) and invalidate the test of overidentification in (15). We discuss here the validity of these assumptions.

First, we noted that if price changes are correlated with rainfall, then we may be capturing the effect of price changes in our estimates of the effects of household agricultural income on consumption expenditures. In the absence of good information on local prices, we cannot test whether they are indeed correlated with rainfall patterns. However, we can argue, in theoretical terms, that our results are unlikely to be driven by bias arising from such correlation. If the correlation relates to prices of inputs or outputs of farm production, then they are accounted for in our measures of farm plot incomes. Changes in these prices affect consumption via farm income and, therefore, such correlation would not violate the exclusion restriction. However, correlation between rainfall patterns and the prices of consumption goods can directly affect consumption expenditures. But, to the extent that rainfall affects farm output from each type of plot in broadly similar ways, such correlation on its own would lead to similar estimated effects for all three sources of income on household public goods, which is contrary to what we find.

Another possibility highlighted in Section III.D was that the household's demand for certain consumption goods is sensitive to weather variations (after controlling for household income) and that these effects are included in our estimates of (13) and (14). These effects could lead to a reject of the test of overidentification in (15) even if there is efficiency in household consumption; but, once again, it would not explain why the coefficient for income from the household head's plots is larger than the corresponding coefficient for the junior members' plots.

Another potential issue is that labor demand is sensitive to variations in rainfall. This would not affect our estimates if there are perfect labor markets or if leisure is separable from other consumption goods in the utility function. But if both of these assumptions fail, then rainfall patterns would also directly influence household demand for consumption goods. Therefore, we test

TABLE 10
LABOR SUPPLY AND RAINFALL

Fitted Values from First Stage	Male Family Labor	Female Family Labor
Ln(household head income)	-.095** (.041)	-.150** (.064)
Ln(junior male income)	-.052* (.028)	-.078** (.037)
Ln(junior female Income)	-.042** (.018)	-.029 (.023)
1994 year dummy	-14.308** (5.774)	-19.525** (8.422)
Test for all coefficients being jointly significant <i>F</i> -statistic	2.47	1.89

Note. Observations = 473. Dependent variable = $\ln(\text{total labor days used by household})$. Standard errors, shown in parentheses, are clustered at the household level across the two time periods.

* Significant at the 10% level.

** Significant at the 5% level.

whether labor demand responds to rainfall patterns, using data on the household's use of family labor. The results are shown in table 10.¹¹ There is some evidence that a change in rainfall that positively affects income from farm plots managed by the household head or by junior male or female household members may cause a decline in the household's use of male and female family labor. An *F*-test for the null hypothesis that the coefficients for all the rainfall variables equal zero has a *p*-value of .0621 in the case of male family labor and a *p*-value of .1311 in the case of female family labor.

Therefore, it is possible that the increase in expenditures on household public goods noted earlier is not caused by an increase in income per se but, at least in part, by increased leisure enjoyed by adult household members when the level of rainfall is beneficial for farm output. However, we shall argue that, in spite of this caveat, the estimated coefficients in table 10 remain broadly consistent with the hypothesis discussed in Section III.B. According to the estimates in table 10, a rainfall shock that leads to a 10% increase in income from the household head's plot would cause male family labor to decline by 0.9% and female family labor to decline by 1.5%. The corresponding effects for income from plots managed by junior female and junior male members are roughly half as large. Thus, compared to other household farm plots, a positive rainfall shock on the household head's plot leads to a greater enjoyment of leisure by adult household members and also (according to the results in table 6) greater enjoyment of household public goods. In this sense, the picture remains that the gains on the head's plots are more widely distributed

¹¹ We do not provide corresponding results for child labor and nonfamily labor because of the high frequency of observations for which these inputs equal zero.

within the household, as suggested by the ethnographic literature. And if this is so, the argument that household members have a stronger incentive to work on the common plots remains consistent with the evidence. It is also the case that the impact of rainfall on labor use is roughly half as large as the corresponding effect on household public goods. Therefore, it seems unlikely that the pattern of expenditures on household public goods observed in table 6 can be accounted for entirely by changes in leisure, with no explanatory role for increased income.

VIII. Conclusion

Empirical studies of households in developing countries have repeatedly found that the allocation of resources within the household is correlated with gender. By contrast, theories of intrahousehold allocation, for the most part, have remained gender-blind. The household is treated as a collection of agents with distinct preferences but with no *a priori* difference between a man and a woman.

In this article, we suggest that some of the gender-related patterns observed in the empirical investigation of households may be explained by the fact that the head of the household, who is bestowed with authority over as well as responsibilities toward the other members of the household by existing social norms, is in most instances male. We suggest that gender-related patterns in the allocation of resources within the household can be explained without resorting to any assumption of innate differences in preferences or power between men and women but rather by differences in their positions created by social norms. The key piece of evidence we provide for this argument is that, for agricultural households in Burkina Faso, yields on plots owned by men who are not household heads are similar to those achieved on plots owned by women in the same household (controlling for plot characteristics and the crops planted), while yields achieved by household heads (who are, in most instances, male) are significantly higher.

The variations in plot yields can be accounted for by the pattern of labor allocation within the household: both male and female household members provide significant amounts of labor to the farm plots managed by the household head but substantially less labor on plots managed by junior household members of the opposite gender. This pattern of labor allocation does not appear to be Pareto efficient. However, the pattern is consistent with a scenario where family labor is offered voluntarily and to the farm plots where one calculates the greatest benefit for oneself. If the household head is obliged, by social norms, to provide for household public goods using the proceeds from

his farm plots, then the junior household members would have much stronger incentives to work for the household head than on some other household plot over which he or she has no claim. This explanation finds substantial support in the ethnographic literature, which describes a distinction between common and private plots belonging to the same household and the different rules that apply in each case. Finally, in support of this explanation, we show that the head of the household has, indeed, a higher propensity to spend on household public goods out of his agricultural income than do the other household members.

The hypothesis and evidence provided in this article naturally beg the question whether female-headed households exhibit the same pattern of behavior as male-headed households. This is also related to the following, broader question: if the pattern of allocation of productive resources and consumption expenditures within a household is heavily influenced by patriarchal norms of responsibilities and obligations, do these norms extend to settings where there is no longer a patriarch, as traditionally defined? The small number of female-headed households in our sample means that we are not able to answer this question satisfactorily with the available data. Therefore, it remains an open question to be explored in future research.

Appendix

TABLE A1
HOUSEHOLD PUBLIC AND PRIVATE GOODS

Expenditure Item	Proportion of Reported Purchases Intended for All/ Other Household Members	Expenditure Item	Proportion of Reported Purchases Intended for All/ Other Household Members
Fresh vegetables	1.000	Lantern	.695
Raw meat	1.000	Cooked rice	.577
Cooking oil	1.000	Millet doughnut	.547
Lantern	1.000	Clothing and shoes	.521
Social events	.988	Millet porridge	.512
Cereals	.987	Red sorghum product	.500
Sweet potatoes	.982	Fresh fruits	.500
Other nut grains	.972	Cooked meat	.500
Sesame-grain	.970	Coffee	.462
Millet-grain	.968	Sugar	.450
White sorghum-grain	.968	Corn porridge	.333
Bread	.960	Peanuts	.333
Salt	.956	Honey	.324
Red sorghum	.940	Misc. durable goods	.319
Millet flour	.935	Motorbike	.308
White sorghum-flour	.932	Local beer	.298
Cassava	.911	Tea	.244
Rice-grain	.896	Transportation costs	.230
Other condiments and spices	.882	Eggs	.222
Kitchen utensils	.878	Cooked bean products	.184
Bicycle	.738	Other alcoholic beverages	.182
Cooked millet couscous	.709	Soft beverages	.000

Note. The table shows the proportion of all reported purchases in each expenditure category that were intended for "all" or "other" household members, in the 1994 household expenditures survey. The cut-off used for identifying household public goods is .750.

References

- Akresh, Richard. 2008. "(In)Efficiency in Intrahousehold Allocations." IZA Discussion Paper no. 1858, Institute for the Study of Labor, Bonn.
- Bardhan, Pranab, and Christopher Udry. 1999. *Development Microeconomics*. Oxford: Clarendon.
- Basu, Kaushik. 2006. "Gender and Say: A Model of Household Behavior with Endogenously Determined Balance of Power." *Economic Journal* 15 (April): 558–80.
- Bernal, Raquel, and Michael P. Keane. 2011. "Child Care Choices and Children's Cognitive Achievement: The Case of Single Mothers." *Journal of Labor Economics* 29 (July): 459–512.
- Browning, Martin, and Pierre-Andre Chiappori. 1998. "Efficient Intra-household Allocations: A General Characterization and Empirical Tests." *Econometrica* 66, no. 6:1241–78.

- Cragg, John G., and Stephen G. Donald. 1993. "Testing Identifiability and Specification in Instrumental Variable Models." *Econometric Theory* 9 (June): 222–40.
- Dercon, Stefan, and Pramila Krishnan. 2000. "In Sickness and in Health: Risk-Sharing within Households in Rural Ethiopia." *Journal of Political Economy* 108, no. 4:688–727.
- Dubois, Pierre, and Ethan Ligon. 2009. "Nutrition and Risk Sharing within the Household." Working paper, Institut d'Economie Industrielle, Toulouse.
- Duflo, Esther. 2005. "Gender Equality in Development." BREAD Policy Paper no. 011, Bureau for Research in Economic Analysis of Development, Duke University.
- Duflo, Esther, and Christopher Udry. 2003. "Intrahousehold Resource Allocation in Cote d'Ivoire: Social Norms, Separate Accounts and Consumption Choices." BREAD Working Paper no. 016, Bureau for Research in Economic Analysis of Development, Duke University.
- Fafchamps, Marcel. 2001. "Intrahousehold Access to Land and Sources of Inefficiency: Theory and Concepts." In *Access to Land, Rural Poverty, and Public Action*, ed. Alain de Janvry. Oxford: Oxford University Press.
- Fiske, Alan P. 1991. *Structures of Social Life: The Four Elementary Forms of Human Relations*. New York: Free Press.
- Goldstein, Markus, and Christopher Udry. 2008. "The Profits of Power: Land Rights and Agricultural Investment in Ghana." *Journal of Political Economy* 116, no. 6: 981–1022.
- Haddad, Lawrence, John Hoddinott, and Harold Alderman. 1998. *Intra-household Resource Allocation in Developing Countries: Models, Methods and Policies*. Baltimore: Johns Hopkins University Press.
- Hammond, Peter B. 1966. *Yatenga: Technology in the Culture of a West African Kingdom*. New York: Free Press.
- Lallement, Suzanne. 1977. *Une famille Mossi*. Recherches Voltaïques 17. Paris: CNRS-CVRS.
- Ligon, Ethan. 2002. "Dynamic Intra-household Bargaining (with an Application to Bangladesh)." Working paper, Giannini Foundation, University of California, Berkeley.
- Lundberg, Shelly, and Robert Pollak. 1993. "Separate Spheres Bargaining and the Marriage Market." *Journal of Political Economy* 101, no. 6:988–1010.
- . 2003. "Efficiency in Marriage." *Review of Economics of the Household* 1: 153–67.
- Marten, Gerald, and Patma Vityakon. 1986. "Soil Management in Traditional Agriculture." In *Traditional Agriculture in Southeast Asia: A Human Ecology Perspective*, ed. Gerald G. Marten. Boulder, CO: Westview.
- Stock, James H., Jonathan H. Wright, and Motohiro Yogo. 2002. "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments." *Journal of Business Economics and Statistics* 20, no. 4:518–29.
- Stock, James H., and Motohiro Yogo. 2005. "Asymptotic Distributions of Instrumental Variables Statistics with Many Instruments." In *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, ed. Donald W. K. Andrews and James H. Stock. Cambridge: Cambridge University Press.

- Udry, Christopher. 1996. "Gender, Agricultural Production, and the Theory of the Household." *Journal of Political Economy* 104 (October): 1010–46.
- Xie, Yun, Erik van den Elsen, Baoyuan Liu, Kim Trouwborst, Coen Ritsema, Yonqui Wu, and W. B. Zhang. 2002. "Soil Moisture and Distribution and Runoff Generation in a Gully System." In *Sustainable Utilization of Global Soil and Water Resources: Proceedings of the 12th International Soil Conservation Organization Conference, Beijing, China*. Beijing: Tsinghua University Press.