

Households and the Social Organization of Consumption in Southern Ghana

Christopher Udry
Department of Economics
Yale University
christopher.udry@yale.edu

Hyungi Woo
Department of Economics
Columbia University
hw2198@columbia.edu

November 8, 2006

1 Introduction

In her chapter “Balances: Household Budgets in a Ghanaian Study,” Jane Guyer rediscovers a remarkable fact, and draws from it important, wide ranging, and surprising implications. The fact is this: in predominately Akan areas of Ghana households of a wide variety of types (headed by women and men, early and late in their life cycle, migrant and non-migrant, even those consisting of only a single person) have astonishingly consistent patterns of expenditure. About 40% of all expenditures by households of all different types is spent on food, and this share declines only moderately with income. This finding provides the foundation for Guyer’s conclusion that “... the poor person lives a diminished, but otherwise similarly balanced, version of the good life to the better off” (p. 146). This is a fascinating conclusion, because we are all familiar with the evident complexity of living arrangements in Akan households, and with the “multiplicity of criteria of worth” (p. 148) that are expressed in interminable negotiations in multiple realms of social interaction. Thus we are perhaps surprised to find a social gradient in Ghana; a cardinal and one-dimensional scale of differentiation that underpinned the social organization of consumption at least at the historical moment of the Ghana Living Standards

Surveys that provide her key raw material in this chapter.

In this note, we examine Guyer's finding in some more detail and in some other contexts, and we provide some suggestions for further interpretation.

2 Consumption Patterns: Ghana and Abroad

Guyer finds that the pattern of consumption expenditure is very consistent across many different types of households, over a wide range of levels of overall expenditure. This is surprising in that it contrasts so strikingly with some of the early studies of consumption in England in the 1930s. In those early studies, Engel's law was found to be decisive: the proportion of total expenditure that is devoted to 'necessities' such as food declines as income rises. In figure 1, we replicate Guyer's fundamental result, using the same subsample of the GLSS that she examined. We use a slightly different statistical technique than she uses, to facilitate comparisons across different samples of households, and to be somewhat more flexible. The horizontal axis measures total household expenditure per person (in logarithms, so a move from '11' to '12' corresponds to somewhat more than a doubling of total *per capita* expenditure). The vertical axis measures the share of household expenditure that is devoted to food. The figure clearly shows Guyer's main point: huge increases in household expenditure are associated with very minor changes in the share of expenditure that is devoted to food, at least until one reaches the very upper ranges of the distribution of expenditures. The lower half of figure 1 portrays the distribution of per capita expenditure, to give the reader a better idea of where in the distribution of expenditures

most households fall.^{1,2}

In table 1, we quantify the relationship between total expenditure and share devoted to food. On average, a ten percent rise in expenditure within the household is associated with a decline in the food share of less than half a percentage point. This number gets a bit larger when we take into account the other household characteristics examined by Guyer. In columns 2 and 3, we show that a ten percent increase in expenditure is associated with a decline of just less than 1 percentage point in the share devoted to food. This implies that poor households tend, on average, to have other characteristics that depress the proportion of expenditure that they devote to food (poorer households tend to be larger, for example). Figure 2 shows this result graphically; the steeper curve depicts the relationship between total expenditure and the food share taking into account the effect of other variables. The shallower curve represents the simple correlation between per-capita expenditure and the food share and is identical to the same curve in figure 1.³ Even this steeper slope, however, remains quite shallow. A doubling of overall expenditure is associated with only a small decline in the share of expenditure devoted to food. Guyer’s key empirical finding is robust to all of these investigations.

¹The details: The top half of figure 1 is the Lowess locally weighted regression of the food share on the log of per capita expenditure. I use the quartic kernel with a bandwidth of .5. The estimates of the pointwise standard errors are obtained from 50 bootstrap replications. The lower half of figure 1 is the kernel estimate of the density of log per-capita expenditures, with an Epanechnikov kernel.

²There is one puzzling difference between these estimates and those presented by Guyer. We estimate that the food share is typically about 60 percent of total expenditure, which corresponds to the summary statistics reported by Glewwe and Twum-Baah (1991). Guyer reports a lower food share (about 40 percent). We use a version of the GLSS that has been updated since Guyer’s work, which may be one source of the difference. Another possibility is that Guyer’s food consumption figures may not include the value of food produced on the farm and consumed at home. Neither of these changes has *any* substantive impact on Guyer’s conclusions regarding the social gradient.

³Technically, the new curve is the semiparametric portion of the partial linear relationship

$$\omega_i = f(x_i) + Z_i\beta + \varepsilon_i$$

where ω_i is the share of total expenditure devoted to food by household i , x_i is the log of per-capita expenditure, Z_i is a vector of characteristics of household i including its size, the gender of its head, and an indicator that is 1 if the household has only a single member. The estimate of the function $f()$ is labeled in figure 2 as the curve with “Controls for HH Size, etc...”. The other curve in figure 2 is the same as that portrayed in figure 1 (though on a slightly different scale).

This robust average pattern, of course, conceals a great deal of variation. Figure 3 depicts some of this variation. It is simply the distribution of the food share across the GLSS sample. While the median household devotes about 60% of its total expenditure to food, 10% of households devote more than 75% of their expenditure to food, and another 10% devote less than 40% of their expenditure to food. While some of this variation is undoubtedly a consequence of measurement error, it is important to always keep in mind that, especially in matters of consumption, there is typically enormous variation around sample means. There may be one social gradient; but there seems to be a good deal of room for individual innovation within this pattern.

Perhaps even more interesting, however, is that the fact that Guyer uncovers for Akan households in the GLSS is not at all unique to these households. We examined consumption patterns as revealed in the Côte d'Ivoire Living Standards Survey and, more adventurously, in the Kagera, Tanzania Health and Development Survey. In both Côte d'Ivoire and in Kagera, we find a remarkably stable pattern of food consumption that is quite similar to Guyer's finding for Akan households. Figure 4 summarizes this finding for the Côte d'Ivoire. Over a range of income similar to that we examined among Akan households, we find that the share of expenditure devoted to food falls from about 57% to about 48%. There is a hint that the share of expenditure devoted to food actually *increases* with total expenditure at very low levels of expenditure, but this is not statistically significant. As in Ghana, the rate at which the food share declines with increases in expenditure is higher among richer households.

Table 2, column 2 presents regression results for Côte d'Ivoire analogous to those for the southern Ghanaian households, which we repeat for clarity in column 1 of Table 2. As can be seen from the coefficients on the logarithm of per capita expenditure, increases in expenditure are associated with almost precisely the same declines in the proportion of expenditure devoted to food in Ghana and in Côte d'Ivoire.

Ghana and Côte d'Ivoire share many features, and a significant proportion of the households in the CILSS have Akan ethnicity. Therefore, it is perhaps not entirely surprising that consumption patterns are relatively similar between southern Ghana and Côte d'Ivoire. It is more remarkable that yet another similar pattern is found in the Kagera region of Tanzania. Figure 5 looks familiar: it describes the relationship between per capita expenditure and the share of ex-

penditure devoted to food found in the Kagera Health and Development Survey. Once again, we find a very minor fall in the food share as expenditure rises. Over a range of expenditure similar to that we examined among the Ghanaian Akan households, the food share in Kagera falls from 67% to 64%. There is, as in Ghana and Côte d’Ivoire, a hint that this rate of decline is steeper for relatively wealthier households, but the overall pattern is absolutely apparent: households at all levels of expenditure typically have very similar average patterns of consumption.

Column 3 of Table 2 presents regression results for Kagera analogous to those for Ghana and Côte d’Ivoire. The similarity in the relationship between overall expenditure and the share devoted to food across these three very different samples of households is particularly apparent in this table.

Moreover, it appears that this pattern may not be unique to Africa. Subramanian and Deaton (1996) provide a careful examination of food demand patterns in rural Maharashtra, India. They show that the share of expenditure devoted to food declines from 73% for the bottom 10% of the expenditure to about 67% at the sample mean, with a somewhat more rapid decline at the upper end of the expenditure distribution.⁴

The very rapid decline in the food share as expenditure rose in England in the 1930s observed by Prais and Houthakker appears to be the exceptional case, not the relative stability that Guyer uncovered for Akan households in Ghana. It is worth considering what it was about Europe in the 19th and 20th centuries that made Engel’s law seem so plausible and easy to replicate. Guyer proposes that the rapidly declining share of food expenditure in that context might “reflect a rising implication of household expenditure in frameworks for capital management in national financial institutions” (p. 150). Perhaps so, but if the growth of formal financial intermediation is to account for Engel’s law being so strong in that context, we believe that the route may be quite complex and subtle. Virtually all saving in Ghana takes a physical form like livestock, or stored grain (Aryeetey, 2004). Typically, such savings are at least partially included in measured total expenditures in household surveys. In contrast, saving in financial institutions are never

⁴These figures *overstate* the actual rate of decline of the food share as expenditure increases in Maharashtra. Subramanian and Deaton provide the average proportion of expenditure on food (that’s the 67% number). Since the relationship between per-capita expenditure and the food share is concave, the average food share understates the food share of those households whose expenditure is equal to the average. A better estimate of the food share at the sample mean would be something like 68 or 69%.

included in measured total expenditure. Therefore, the overall gradient of the food share of expenditure with respect to *measured* total expenditure is a complex function of savings rates, portfolio choice, and the rate at which households enter the formal financial system as their total income increases.

Turning this hypothesis on its head, the relative stability of consumption patterns may reflect the high degree of poverty among households in Ghana, Côte d'Ivoire and Kagera. Polly Hill dismissed this possibility, because “Ghana is not the Far East and such a preposterous suggestion must be at once disposed of” (1958, 16, cited by Guyer, 134). While chronic caloric deficiencies are not widespread in rural Ghana, it is true that a large fraction (about 40%) of Ghanaians live below the poverty line as defined by Ghana. Households in poverty (roughly the households in the left half of the graphs in Figure 1) have little room for expenditure on luxuries, so their consumption patterns are relatively stable. It is indeed these households for whom the gradient is least steep in our graphs. Despite Hill’s dismissal, this hypothesis should be considered more carefully.

Thus we have two complementary directions to consider for further work. First, it is imperative to understand the relationship between formal financial intermediation and the shift in the composition of expenditure as total expenditure increases. Second, we need to consider the relationship between ‘poverty’, as defined by a political process in Ghana and the shape of expenditure decisions within households.

3 ‘Households’, and Consumption

Any more careful consideration of consumption patterns, however, has to address a concern raised by Guyer in the opening section of ‘Balances’. That is the interpretation of the responses that people made to the questions on the GLSS. We are certainly not about to argue that nothing can be learned from people’s responses to standardized surveys; they are our primary research tool. Along with Guyer, We believe that “people say *something* true about themselves ...” in response to such questions (p. 131). One of our tasks is to think carefully about what that *something* may be.

Guyer’s (1981) caution regarding the dangers of working with households as units has had an

enormous impact on the methods used by economists studying consumption and other aspects of social interaction, particularly in West Africa. The GLSS was designed with some of these concerns in mind; in particular, the Ghana Statistical Service understands that different individuals in a household neither pool their income nor make joint consumption decisions. In contrast to budget surveys in many parts of the world, the GLSS expended considerable effort to obtain information on consumption from all of the household members who “are mainly responsible for making the household purchases” (GLSS 3 Household Questionnaire, Section 6, Question 6). Nevertheless, as a practical matter, most of the questions regarding purchases of particular types of product (food, consumer durables, non-food expenses) were answered by a single individual. If more than one individual responded, this typically occurred during an interview at which all were present. Diaries were also provided to households with literate members; it is not clear how that individual obtained expenditure information from other household members. As a consequence, expenditure information on particular items in the GLSS tends to come from the individual most responsible for that type of purchase. This reliance has important consequences for the reliability of data on consumption in Ghanaian (and most West African) households.

A typical framework for thinking about these consequences is provided by models of *classical measurement error*. In this model, it is assumed that an item reported in a survey (say, the household’s purchases of onions over the past week) equals the true amount of such purchases, plus a random error that is uncorrelated with anything else. This is a very useful framework in certain circumstances, but not here.

In the late 1990s, one of us and a colleague (Markus Goldstein) organized a detailed survey that collected data over a two year period from members of a sample of approximately 240 households in 4 clusters of villages near Akwapim. Information on the survey, and the data itself is available at <http://www.econ.yale.edu/~udry/ghanadata.html>.

One of the features of this data is that we collected cross-reports from household members about each other’s expenditures, in order to investigate information flows within households. In an innovative paper, Boozer and Goldstein (2003) investigate the characteristics of these cross reports. Figures 6 and 7 summarize one of the key conclusions of their paper. Figure 6 shows the reports of both the husband and the wife regarding the husband’s expenditure on types of purchases of food that Boozer and Goldstein categorize as ‘public’, which includes most of the

basic foodstuffs that are purchased for consumption in meals at home (for example, purchases of maize). The horizontal axis reports the man's report of his own purchases; the vertical axis records his wife's report of the same purchases. If both individuals reported accurately the same purchases, this graph would be a collection of points all lying on the 45° line. Classical measurement error results in a symmetrical spread of points around that line, reasonably close to what we actually observe. In contrast, Figure 7 shows similar cross-reports on 'private' food expenditure – most importantly, meals outside the house and drinks. There is little evidence of classical measurement error in this picture. Women report a preponderance of zeros for their husband's private expenditures. Aggregate statistics based on women's responses would seriously underestimate total 'private' expenditures, if in fact men report their own spending relatively reliably. Boozer and Goldstein interpret this pattern as an indication of lack of information; these women simply do not know what their husbands are spending on these private purchases and they are unwilling to guess. When they do hazard a guess, their accuracy is strikingly low. Intensive informal conversations with respondents in our survey, and several focus group discussions support this interpretation.

In our Akwapim survey, the data are revelatory regarding information flow about consumption expenditure within households because these cross reports permit explicit comparisons of the responses of different individuals. One can literally see the reluctance of people to guess about expenditures are private; this is apparent in Figure 7, and even more so in the cross reports of more fundamentally private information such as gifts to family members, or income from nonfarm enterprises. However, cross reporting is not an element of conventional expenditure surveys like that embedded in the GLSS. As a consequence, we have to be more careful about our interpretation of the recorded numbers. When a woman responds to the question "How much was spent on fufu and soup since my last visit?", does she report only those expenses she actually knows about? Or, since she does not know how much her husband spent, does she report an estimate of how much one would have expected a household like hers to have spent on this item?

4 Conclusion

In section 2, we raised the possibilities that the social gradient proposed by Guyer in ‘Balances’ is a consequence of the high degree of stress faced by Ghanaian households (households in poverty have little flexibility in spending patterns) or that it reflects the relative unimportance of a formal financial intermediation for these households. The analysis here leads us to a third hypothesis concerning the social gradient. Guyer hints at this hypothesis: “the enormous variety of descriptive criteria is reduced to two principles – one fixed by chronology and the other volatile according to power – to produce a momentary performance of a single scale of value” (p. 149). This ‘single scale of value’ may be particularly apparent in certain responses in household surveys. When information about consumption patterns in the household is very imperfect, then surveys that pose impossible questions could well lead to responses that correspond to norms that are at least partially shared. If so, then at least some of the systematic pattern of proportional rises in food and total expenditure may reflect a normative social gradient. The answers may tell us less about actual food expenditures than about respondents’ beliefs about what expenditures *should be*.

Guyer’s foray into the statistics of Ghanaian household budgets turns out to have been quite fruitful. A very stable proportion of expenditure is devoted to food, across a wide range of total expenditure, at least until one reaches the upper quarter of the distribution of expenditure. We reconfirm her finding in Ghana, and extend it to other areas of Africa. In contrast, Engel’s law (that the food share declines dramatically with total expenditure) holds strongly in Europe and the United States. We propose that three hypotheses that would reconcile these results be explored. First, Ghanaian households, by and large, do not have the flexibility within their budgets to make dramatic changes to the food share.⁵ Second, a more rapid decline in the food share as income increases in the U.S. and Europe may reflect the relatively greater importance of the formal financial sector in these countries. This structural difference transforms savings and portfolio choices, and could lead households to alter expenditure patterns more rapidly as

⁵Put more precisely, it may be that the expenditure elasticity of food demand falls with total expenditure. That is, the food share remains fairly constant as expenditure increases for very poor households, and then declines more rapidly as expenditure increases for better-off households. This basic pattern is observed in budget data from the US and Europe.

income rose. Third, some of the apparent consistency of the food share across expenditure classes may reflect a normative pattern that is revealed in responses to questions when the respondent has limited information about actual expenditures.

5 References

Aryeetey, Ernest. 2004. "Household Asset Choice Among the Rural Poor in Ghana," paper presented at the workshop on *Understanding Poverty in Ghana*, Accra, January.

Boozer, Michael and Markus Goldstein. 2003. "Poverty Measurement and Dynamic." Manuscript: Yale University.

Guyer, Jane. 1981. "Household and Community in African Studies," *African Studies Review* 24, 2/3: 87-137.

Guyer, Jane. 2004. *Marginal Gains: Monetary Transactions in Atlantic Africa*. Chicago: University of Chicago Press.

Subramanian, Shankar and Angus Deaton. 1996. "The Demand for Food and Calories," *Journal of Political Economy*, 104: 133-62.

Figure 1a: Food Share of Expenditure in Ghana

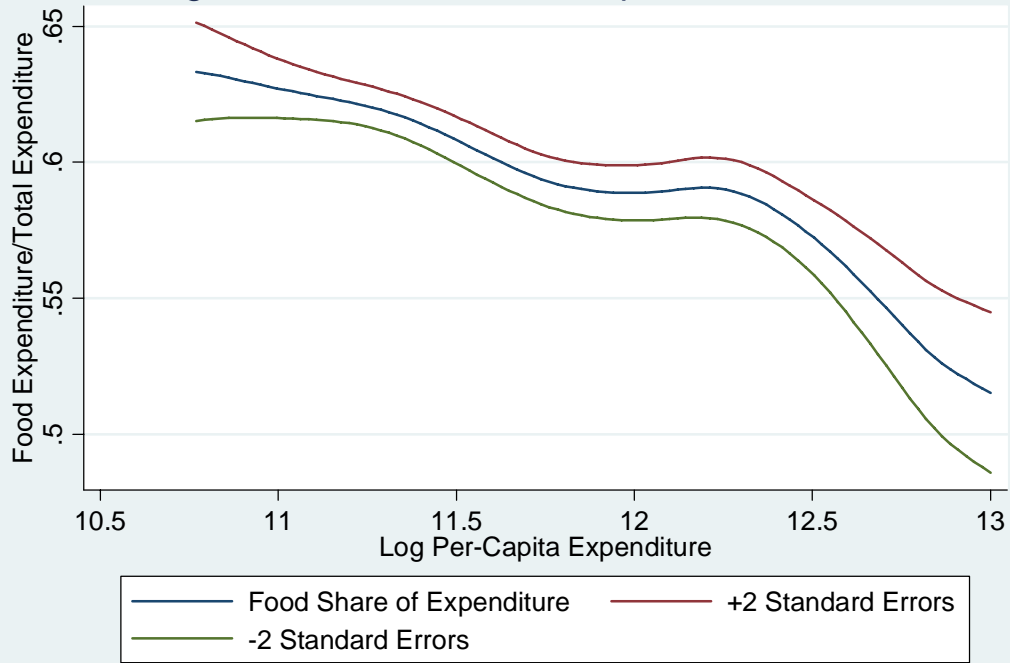



Figure 1b: The Distribution of Expenditure in Ghana



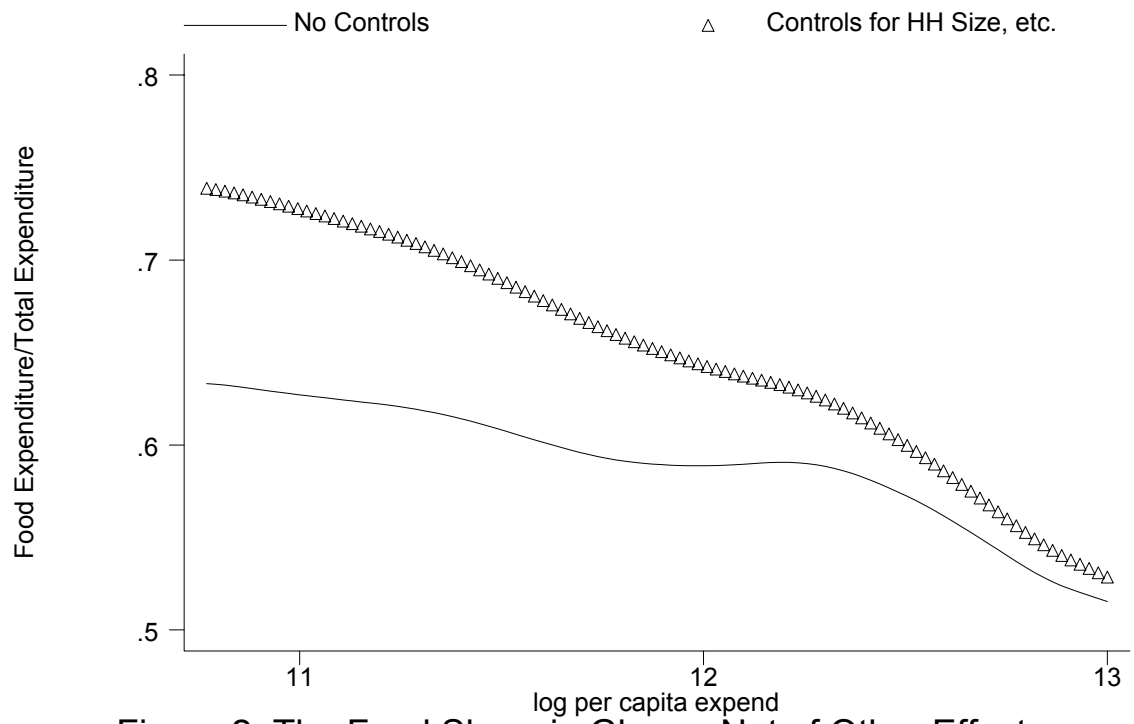


Figure 2: The Food Share in Ghana, Net of Other Effects

STATA™

Figure 3: Distribution of the Share of Expenditure on Food

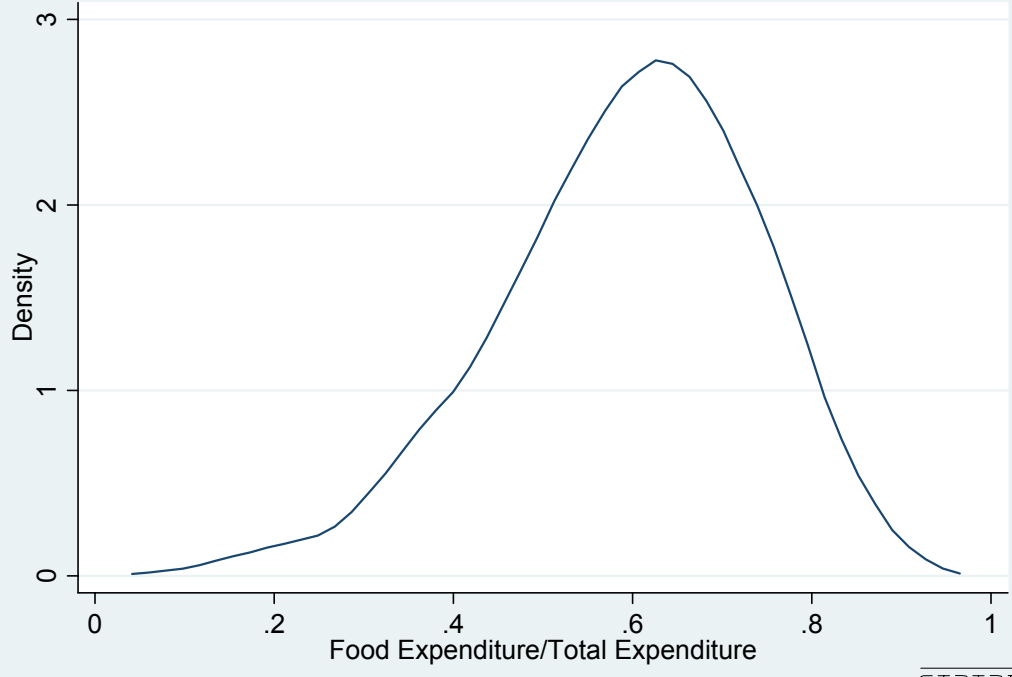


Figure 4a: Food Share in Cote d'Ivoire




Figure 4b: The Distribution of Expenditure in Cote d'Ivoire

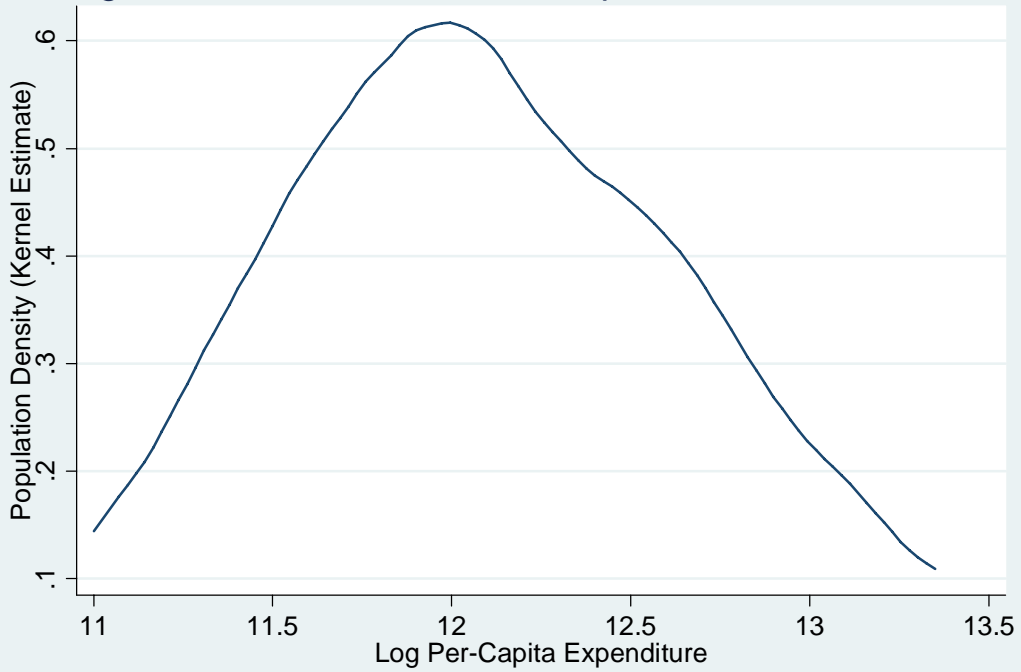


Figure 5a: Food Share in Kagera




Figure 5b: The Distribution of Expenditure in Kagera






Fig 6: Cross reports of Male Public Spending

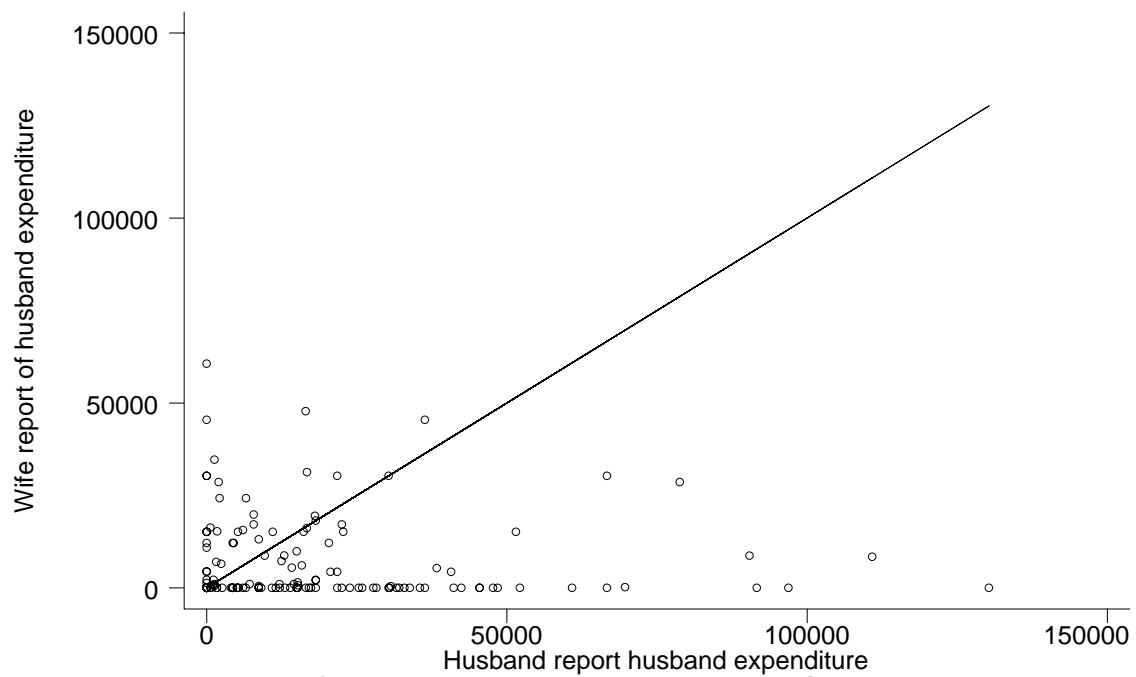


Fig 7: Cross reports of Male Private Spending

Table 1: Determinants of Food Share of Expenditure in Ghana

dependent variable	1		2		3	
	OLS		OLS		OLS	
	Food Share		Food Share		Food Share	
	estimate	t-ratio	estimate	std error	estimate	std error
Constant	1.20	23.54	1.80	28.50	1.76	27.80
Log Per-Capita Expenditure	-0.05	-11.94	-0.10	-18.81	-0.09	-18.51
Log Household Size			-0.07	-14.92	-0.08	-11.53
Female Headed Household					0.02	3.78
Single Person Household					0.03	2.30
Observations	2329		2329		2329	
R-Squared	0.06		0.14		0.15	

**Table 2: Determinants of the Food Share
in Ghana, Cote d'Ivoire and Kagera, Tanzania**

dependent variable	1 OLS		2 OLS		3 OLS	
	Ghana Food Share		Cote d'Ivoire Food Share		Kagera Food Share	
	estimate	t-ratio	estimate	std error	estimate	std error
Constant	1.80	28.50	1.88	33.89	1.46	21.64
Log Per-Capita Expenditure	-0.10	-18.81	-0.10	-23.62	-0.07	-12.1
Log Household Size	-0.07	-14.92	-0.07	-15.45	-0.03	-4.34
Observations	2329		2349		915	
R-Squared	0.14		0.21		0.14	