

How responsive is body weight to transitory income changes? Evidence from rural Tanzania*

Niklas Bengtsson[†]

This draft, July 12, 2007.

Abstract

We use time-series of rainfall along with individual fixed effects to estimate the response of body weight to transitory changes in household income and expenditure. Our data consist of a longitudinal sample of subsistence farmers in rural Tanzania, representing one of the poorest populations in the world. We find that the response of body weight to transitory changes in household income is positive on average, but that the impact decreases with age and being male. For female children, a ten percent increase in household income implies an increase in body weight with about 0.4 kilo. The body weight of male adults is practically invariant to income changes.

JEL-classification: O12, D13, I12

Keywords: Income variability, Consumption, Nutrition, sub-Saharan Africa

*I would like to thank Sosina Bezu, Per Engström, Maria Floro, Hans Grönqvist, Bertil Holmlund, Erik Jakobsson, Peter Nilsson, Johan Vikström, Joachim de Weerd and seminar participants at Uppsala University and at the Nordic Conference in Development Economics for valuable comments and suggestions. Financial support from the Swedish International Development Cooperation Agency (SIDA) is gratefully acknowledged.

[†]Department of Economics, Uppsala University. Box 513, SE-751 20, Uppsala, Sweden. Contact information: niklas.bengtsson@nek.uu.se. Phone: +46 18 471 15 62

1 Introduction

Throughout rural Africa, household expenditure is very sensitive to weather-induced changes in rural production and income. In this paper, we ask to what extent such fluctuations are reflected in the physical capacity of the population. How sensitive is nutritional status to transitory changes in household income and expenditure? Are female children particularly vulnerable?

Although widespread malnutrition typically occurs after series of weather failure, it has been recognized for some time that malnutrition is not simply determined by the regional supply of food. In many sub-Saharan countries, endemic malnutrition is present also during normal periods, when the total amount of food production plus imports is enough to feed everyone. Amartya Sen was among the first to note that severe outbreaks of malnutrition not seldom coexist with regional prosperity, along with little or no decrease in the total level of regional food production (Sen, 1981). Rather than simply being a function of regional food availability, Sen argues, hunger catastrophes occur when the household's purchasing power with respect to food decreases.

Sen's entitlement approach has had a large influence on development policy, with clear policy guidance: if you want to reduce malnutrition, you need to reduce poverty. However, a second wave of literature has come to question whether weather poverty is the main, or even most important, cause of malnutrition. This literature builds on two empirical observations. A first one is that regional malnutrition seems to occur also when the total cost of the calories needed to assure "light physical capacity" – FAO's criterion for nutritional well-being – is so small that even the poorest households should be able to secure an adequate level of nutrient intake. Secondly, estimates of the demand for calories and other nutrients suggest that households seem reluctant to channel positive changes in income towards calories. Although the income elasticity of food is typically close to unity in developing countries, the income elasticity of calories seems to be closer to zero, even among households that would not be considered well-nourished by international standards. In fact, a handful of papers suggest that the demand for calories is completely invariant to income changes (Behrman and Deolalikar, 1987, 1990; Bouis and Haddad, 1994; a recent paper is Aromolaran, 2004). Behrman and Deolalikar (1989) propose an economic rationale for this finding, namely that households have a strong "taste for variety", even at the lowest levels of income.

This so-called "revisionist" literature has been contested (see Subramanian and Deaton, 1996), but the question whether the nutritional status of individuals in sub-Saharan Africa will improve as income increases is still relevant for development policy. If the demand for nutrients and calories is largely invariant to income changes, policymakers mostly concerned with improv-

ing nutrition may find “standard” economic reforms – trade liberalization, micro-credits, income tax policies, etc. – less relevant. On the other hand, if malnutrition is the result of binding income constraints, economic reforms will be aligned with more health-oriented policy initiatives. Today, the former attitude is probably the most prevalent in policy documents. In a paper in the World Bank Policy Review, Haddad et al. (2003) concluded that child malnutrition will not be eradicated by economic growth alone, if it is not accompanied by reforms aimed at improving nutrition directly.

In this paper, we estimate the impact of large transitory swings in household expenditure on nutritional status. Our data consist of a longitudinal sample of subsistence farmers in rural Tanzania, collected by the World Bank. The *Kagera Health and Development Survey* (KHDS) sample represents one of the poorest populations in the world and one whose nutritional status, defined as BMI and weight-for-age, is below international standards. Rather than focusing on nutrient intake, such as the quantity of calories consumed, we use body weight as a proxy for nutritional outcome, following the suggestion of e.g. Haddad et al. (2003) and Strauss and Thomas (1998). This approach allows us to compare the response of nutritional outcome to income changes across family members. In order to capture the transitory component in household expenditure, we use rainfall as an instrument along with individual fixed effects. As most of the households in our sample rely on crop yields as the main source of income, rainfall is likely to explain a non-trivial share of the intertemporal variation in both household income and expenditure.¹

An instrumental variable strategy should improve the estimates obtained by ordinary least squares for three main reasons. First, a typical problem in survey-based econometric exercises is the attenuation bias stemming from the classical case of measurement error; people are unlikely to be able to perfectly recall household expenditure or income for the last six months. Error in self-reported income is also known to be correlated to actual income, which can lead to biases in both directions. The second objective is to evade the simultaneity between nutritional status and income. The reciprocal relationship between nutrition and income has been a central feature in some important work in development microeconomics (following the work by Harvey Leibenstein in 1957). Third, rainfall fluctuations arguably captures

¹More broadly related papers have used weather variation in order to test hypothesis about the income-consumption nexus. This strand of research includes Wolpin (1982), Jacoby and Skoufias (1998), Paxson (1992) and Duflo and Udry (2004). Kochar (1998) studies whether weather-induced income shocks increases labor supply. Miguel et al. (2004) and Miguel (2005) use rainfall variation to estimate the causal impact of economic shocks on ethnic conflict and witch killings, respectively.

a transitory and exogenous component in household's income and expenditure, uncorrelated to life-cycle decisions, knowledge or other variables that may enter the households' preferences over nutrition.

Despite the large body of research on malnutrition – and the popularity of the “weather-instrument” – there is only a handful papers that tries to establish the causal impact of transitory income shocks on nutritional status.² This literature is mostly concerned with child nutrition. Hoddinott and Kinsey (2001) and Alderman et al (2005) study child growth in Zimbabwe and Ethiopia, respectively, and report that children belonging to “drought cohorts” experienced a slowdown in human growth. Rose (1999) studies the interaction between gender, “favorable weather shocks” during early childhood and the probability of survival, and finds that female child mortality is increased during periods of adverse weather shocks. Foster (1995) compares the impact of flooding on child weight in Bangladesh across land-owners and landless households, in order to test for liquidity constraints, and finds that the variation in child weight was more pronounced among landless households. Dercon and Krishnan (2000) use nutritional status for adult males and females in order to study risk-sharing within households in Ethiopia, and report asymmetric intrahousehold responses to various “shock”-variables.

We add to this literature, and contribute in some important methodological respects. First, we use a time-varying instrument along with individual fixed effects. By doing so we can disentangle the random element in year-to-year rainfall fluctuations, keeping the mean levels of rainfall and expenditure constant. Second, earlier estimates of the causal impact of income shocks have typically been based on reduced-form equations – i.e., studied the impact of the shock-variable in itself (the exception is Foster, 1995). IV-estimates are, however, more useful for comparisons across datasets and samples, and will provide an economically relevant measure of the severeness of income fluctuations. Such estimates are becoming more and more important as empirical contributions in development economics level a strong emphasis on internal validity (as opposed to producing externally commensurable estimates). Third, the completeness of our dataset allow us to control for incidences of malaria and other diseases that are known to be triggered by

²There is, of course, a large number of contributions that study correlations between development, interpreted broadly, and health outcomes, well summarized in Duncan and Thomas (1998). A relevant example is Alderman et al. (2006), who study the effects of program interventions and long-term growth using the same dataset as we do. They report a positive "association" between household consumption and child height (in some specifications, they use roof-quality and self-reported income as instruments for consumption, without using fixed effects, which unfortunately does not solve for simultaneity, omitted variables, or non-classical measurement error).

rainfall in order to study the stability of our IV-estimates. Finally, we know of no other paper that is able to compare the response of body weight to transitory changes in household income across *all* household members. Despite the theoretical and documented importance of intrahousehold allocation of calories and nutrients, direct evidence on asymmetric responses among family members is quite rare (but see Behrman and Deolalikar 1990, who study calorie intake in this fashion).

We find that the response of body weight to transitory changes in household expenditure and income is positive on average, but that the elasticity decreases with age. For male adults, the elasticity is not statistically different from zero. For female children, the response of body weight to transitory income changes is markedly higher: a ten percent decrease in household expenditure implies a body weight reduction of about 0.4 kilo. The results do not differ if we use self-reported income from agricultural profits or total expenditure as explanatory variable. A close-to-hand interpretation is that female child nutrition is considered a “luxury good” in this region, which receives attention after the adult household members have secured an adequate level of nutrition for themselves (we will show, however, that a non-negligible share of the adult sample would not be considered well-fed in terms of body mass index).

The paper is organized as follows. In Section 2, we present some summary statistics and background information for the study. In Section 3.1, we outline a basic framework for interpreting the response of rainfall-induced changes in household income. We put special focus on the measurement and interpretation of income and total household expenditure when rainfall is thought to enter the budget constraint via rural production. In Section 3.2, we address some econometric considerations regarding the definition of our instrument, lagged rainfall. The data is presented in Section 4. In Section 5 we present our results, and section 6 concludes.

2 Background and institutional setting

The Kagera region in northwestern Tanzania is situated on the western shore of Lake Victoria, just below the equator, bordering Uganda to the north and Burundi and Rwanda to the west. The population is a little more than 1.3 million (1988 Census). The area is about 30 000 square kilometers.

In Table 1, we present some summary statistics for expenditure, economic activity and nutritional status, based on KHDS sample. Rough calculation suggests that the median per capita expenditure in 1991 prices was about 180 \$ per capita (which is similar to the IMF figures on Tanzania’s per capita

income in 1991). The budget share of food expenditure is similar to the much-cited figures in Marshall's *Principles of Economics* for the lowest income group: about 60 percent of the total household budget is devoted to food expenditure (including the value of own-produced food). The median household produces a little less than half its total expenditure, and about three fourths of all food expenditure. There is an obvious discrepancy between self-reported income, household expenditure and cash savings. The value of consumption greatly exceeds the self-reported income, but there is still non-zero (albeit low) cash savings. The expenditure variable also has lower variance.

Household production is centered around the household's shambas ("gardens"), which in most cases are owned by the household or the household's family. Formal sharecropping arrangements are rare, but some farms are owned by the household's relatives or extended family. Own-farm work is typically unpaid and performed by members of the household. According to the KHDS community survey, members of rural households often perform market work in Bukoba town or elsewhere in Tanzania on a seasonal basis. Note, however, that the value of this labor is very low compared to total household expenditure (Table 1). The value of remittances from individuals and from organizations (e.g. the Evangelical Lutheran Church of Tanzania) is also low compared to total household expenditure.

The most important food crops are banana, maize, and beans, and the major cash crops include coffee, cotton and tea (grown at the subsistence level). Maize, sorghum and cassava are also produced, but are less popular in Kagera than elsewhere in Eastern Africa. Sugar cane is the only crop produced at a high commercial level by the Sugar Cane Company (TPC 1998). Both staple crops and cash crops are rain-fed. The rainfall pattern in Kagera is characterized by rather sharp seasonality, and follow a bimodal pattern, with two rain periods per year. The timing of the rain periods is however erratic, and when they come they vary in intensity. Drought is perceived as the major type of weather failure. Flooding is less of a problem due to the undulating terrain and economic emphasis of tree crops in the largest districts of Bukoba, Karagwe and Muleba.

A first glance at the anthropometric figures in the KHDS dataset suggests that the nutritional status of the people in Kagera is below international standards. Looking at Figures 1 and 2, and Table 1, about 20 % of the adult sample and about 50% of those younger than 18 years old would be considered malnourished to some degree by WHO-standards.³ The bulk of

³For adults, a body mass index below 18.5 is a common definition of underweight; 25 is the threshold for overweight. For children (below 18 years), an analogous indicator

the sample is below the standard reference value of normal weight, and not even those and the ninetieth percentile of the BMI distribution in this sample would be considered overweight. Further comparison reveals that girls are actually somewhat better nourished than boys.

Are these figures the result of poverty? The essence of the “revisionist” critique to the analysis of malnutrition is that the costs of improving nutrition, in terms of energy intake, are easily born by even the poorest households. FAO (2001) estimates that the daily energy requirement needed for an “active” man of 65 kilo is about 3000 kilocalories; the corresponding figures are 2 500 for females and about 1 500 for children.⁴ A simple (albeit rough) calculation reveals that if the household was to consume only cassava – a cheap and drought resistant root-vegetable, used as a staple crop in much of rural Africa – the cost of fulfilling the yearly energy requirement for a five-person family would be 9 125 Tanzanian schilling (Tnz) per household member. This implies that a household at the median income level in our sample could meet the energy requirements using one sixth of the household budget; for the lowest quartile of the income distribution, the corresponding figure is a little more than one fourth. Similar patterns in other developing countries has led Behrman and Dealolika (1989) to hypothesize that people demand luxury goods also at the very lowest levels of income.

The traditional way of testing this hypothesis has been to estimate the expenditure elasticity of calories, sometimes along with other health-improving nutrients. Data on calorie-intake have either been obtained by converting consumption data into calorie intake by using tables of conversion ratios, or by letting nutritionists observe nutrient intake directly.⁵ The former approach, which relies on the assumption that all available calories are consumed by the household members, has yielded a large number of estimates, typically within the range of 0.2 to 1 (see Deaton, 1997, and the references therein). The handful of papers that uses the latter approach has reported estimates that are closer to zero (Behrman and Dealolika, 1987, 1990; Bouis and Haddad, 1992).

Even with this latter approach, however, the variable calorie-intake is un-

is the normalized weight-for-age z -score. A child with a corresponding z -score below -2 is typically considered moderately malnourished, -1 indicates light malnutrition, and -3 severe malnutrition (see WHO, 1995).

⁴The FAO figures are among the most influential, but have been criticized for being too large (they are based on the assumption that labor is physically demanding in rural developing areas). For our purposes, they serve well as upper bound-estimates.

⁵Behrman and Dealolika (1990), for example, asked the most knowledgeable person to serve the typical amount of food given to each family member on different plates, and gauged the individual calorie intake on that basis.

likely free from measurement error. Recollections of past consumption are seldom perfect, and it is possible that the respondent's inability or unwillingness to recollect can be correlated to other household characteristics (see e.g. Strauss and Thomas 1998, for examples where this seems to be the case). A worse problem is measurement error in the explanatory variable, and income and consumption figures are notoriously unreliable in developing countries. Even household expenditure, which is considered less noisy than self-reported income, is likely to be measured with a non-trivial degree of error. In the classical zero-mean case of measurement error in the independent variable, elasticity estimates will typically be attenuated, i.e., biased towards zero.

In this paper, we will use body weight as a proxy for nutritional status. This approach is not only motivated by the fact that body weight is a more reliable statistic than nutrient-intake (anthropometrics are not self-reported but observed directly by the survey staff), but also because the success or failure of development policy often is evaluated with a reference to health *outcomes*, not health *inputs*. The by far most utilized measure of health outcome in the developing world is a person's anthropometric status, promoted by, among others, the *World Health Organization* (WHO) and the *Food and Agriculture Organization* (FAO) of the United Nations. Anthropometric status has been recommended for econometric practice as well (Strauss and Thomas, 1998; Haddad et al, 2003). If one think of body weight as the output of a health-production function, and calories and nutrients as inputs, anthropometrics can arguably be more informative for policy in many settings – just like GDP is often (but not always) more interesting than factor endowments. If nutritional outcome is our primary interest, the income or expenditure elasticity of body weight thus has the advantage of capturing the aggregated effect of all inputs, also those that are not readily observable, such as labor supply.

3 Conceptual framework

3.1 How does rainfall affect rural income and expenditure?

In this section, we study under what assumptions an exogenous weather shock to rural production can be interpreted as a shift of the budget constraint. In a nutshell, such an interpretation is feasible if decisions regarding production are separable from those regarding consumption (see Rozensweig and Wolpin, 2000, for a discussion). If savings are negligible, it is furthermore possible to use total expenditure as a proxy for total income, which is preferable since

the latter is typically considered less noisy than the former.⁶

Consider a one-member household that extracts income from home production of crops and market labor in a one-period setting.⁷ The static framework implicitly assumes that the household is unwilling (or unable) to smooth consumption across time. This assumption can be given some credence by the low level of accumulated cash savings observed in the data (see Table 1). Whether households in rural Africa do not save and borrow across income shocks is of course an empirical question, but the received wisdom is that income and consumption covary to a large extent throughout the developing world (see the surveys by Townsend, 1995 and Morduch, 1995). As it turns out, rainfall has a decisive impact on expenditure in our sample (see Section 5.1).

The household faces the following problem:

$$\max_{c,l} U = u[\mathbf{c}, l, n(\mathbf{c}, l)], \quad (1)$$

subject to

$$I = \pi^*(w, r) + wT = \mathbf{p}\mathbf{c} + wl, \quad (2)$$

$$T = l + h^m + h^o, \quad (3)$$

$$L^d = h^h + h^o, \quad (4)$$

and where

$$\pi^* = \max_{L^d} [f(L^d, r) - wL^d]. \quad (5)$$

The vector \mathbf{c} represents the household's consumption bundle, and includes the good produced by the household; l is hours of leisure; n is the household member's "body weight function", which is determined by the consumption goods and leisure. The utility of the household member thus depends on both the consumption goods in themselves and the way these goods are transformed into nutritional well-being. It is assumed that the utility function (1) is increasing in all arguments. T is the household's time endowment,

⁶Even if our instrument takes care of the measurement-error bias, it is preferable that the dependent variable in the first stage has a low degree of error since it will increase the strength of the instrument.

⁷It is straightforward to extend the framework to a multi-member ("collective household") model with constant bargaining shares (see Duflo and Udry 2004 for an extension in that direction).

h^m is hours of supplied market work, and h^o is hours of supplied own-farm labor. Note that since only l enters utility, the household is indifferent between market work and own-farm work; only total labor supply ($h^o + h^m$) will cause disutility. The price of labor and the prices of the consumption goods are represented by w and \mathbf{p} , respectively.

The function $f(L^d, r)$ is the household's production function, which is determined by hours of effective labor, L^d , and the realized and exogenous rainfall shock r . The production function is increasing and concave in both arguments. As seen in Equation 4, L^d is defined as the sum of hired labor hours, h^h , plus the household's own-farm labor supply (in hours), h^o . Note that h^o and h^h are perfect substitutes in production. The essence of the model is that production is not constrained by the household's labor supply decisions as long as markets are complete (so that there is an unlimited supply of labor available at the going wage, w). In this case, decisions on household production are typically said to be separable from those of consumption. The household's demand for labor (L^d) is determined in the first step, when the household maximizes profits. In the second step, the household maximizes utility, taking profits π^* as given (see Bardhan and Udry, 1999, for a discussion of the so-called agricultural household model).⁸

As seen in the budget constraint (2), a shock to rural production via r will increase household profits π^* and induce a shift in household expenditure (\mathbf{pc}) plus the value of leisure (wl). It is attractive to measure rural income from the expenditure side. Estimates based on \mathbf{pc} are more readily comparable to the large body of "expenditure elasticity of calories"-estimates, and expenditure figures are, as mentioned, considered far less unreliable than profit-based income figures. The value of leisure is, however, not directly observable. To align our estimates with existing research, we will assume that leisure is invariant across time and control for it by using fixed effects, and thus use \mathbf{pc} (expenditure on market goods plus the value of own-produced goods) as a proxy for total income, for our baseline econometric models.

The assumption of invariant labor supply is, however, not a trivial one. As highlighted by Morduch (1995), seasonal market work or more intensive own-farm work can serve as an income insurance in liquidity constrained rural areas, and related empirical work has verified that such strategies exists (Kochar, 1999 and Beegle et al., 2006). If households engage in income smoothing activities, a shock to rural production may affect labor supply but

⁸Note that the model does not solve for the exact division between the household's supply of market work and supply of own-farm work, nor does it solve for how much hired labor is used in household production. However, *total* labor inputs in production (L^d) will be determined by the profit maximizing decision (5), just as *total* labor supply ($T - l$) will be determined by maximization of utility (1)

(in the extreme case) not observed expenditure. Given that physical activity is an important determinant of body mass, it is reasonable to consider an alternative measure of household expenditure that allows household to adjust their labor supply across time. To do this, we will use information on time-allocation and construct a proxy for the value of leisure. Note that

$$I = \pi^* + wT = \mathbf{p}\mathbf{c} + wl = \mathbf{p}\mathbf{c} + w(T - h^o - h^m). \quad (6)$$

Provided that wages can be controlled for, we can use fixed effects to control for T . Then it is possible to observe total income by subtracting market work income and the value of own-farm tasks from expenditure ($\mathbf{p}\mathbf{c}$).⁹

Using either the "raw" expenditure variable $\mathbf{p}\mathbf{c}$ or the leisure-adjusted one, the first-order conditions for utility optimization yield the household's demand for a particular good j :

$$c_j = d_j(I, p_j, \mathbf{p}), \quad (7)$$

where p_j is the good's own price and \mathbf{p} is a vector of the prices of other goods (including w). Using scalars to denote the consumption goods (including leisure), nutritional status is given by:

$$n(c_1, \dots, c_j) = n[d_1(I, p_1, \mathbf{p}), \dots, d_j(I, p_j, \mathbf{p})] = v(I, \mathbf{p}). \quad (8)$$

Provided that prices are kept constant, a rainfall shock will affect nutritional status only via the income argument (I) in Equation 8. Assuming a linear model, our baseline regression equation is:

$$\ln weight_{iht} = h_i + \beta_1 \ln(I_{ht}) + \mathbf{X}\boldsymbol{\beta}_2 + e_{iht} \quad (9)$$

where I_{ht} is household income and h_i is an individual-specific fixed effect.

The vector \mathbf{X} contains year-district interaction terms and seasonal dummies, which we include to control for district-level variation in prices. General equilibrium considerations suggest that relative prices will not remain unchanged if weather shocks determine rural output. If relative prices change, decisions on consumption may be influenced by the weather shock even in the absence of a shift of the budget constraint. However, assuming that local (district) markets are integrated implies that all households will face the same relative prices. By including \mathbf{X} , relative prices will thus be held constant.

The first stage equation is:

$$\ln(I_{ht}) = h_{hi} + \gamma_1 \ln(r_{ht-1}) + \mathbf{X}\boldsymbol{\gamma}_2 + u_{ht} \quad (10)$$

⁹The value of own-farm tasks is approximated using the going village-level agricultural wage, multiplied by the household's total number of own-farm labor hours.

In Equation 10, $\ln(r_{ht-1})$ is the natural logarithm of past rainfall, to be discussed shortly. With fixed-effects, the coefficient γ_1 is to be interpreted as the percentage deviation from normal rainfall for household h .

To recapitulate: for our baseline estimates, we will observe household income from the expenditure side. The assumption that rainfall does not affect nutritional status other than via observed expenditure is the exclusion restriction that makes our instrumental variable strategy feasible (i.e. that $cov(r_{ht-1}, e_{iht}) = 0$). As noted in the previous section, our identification strategy depends on the assumption that leisure is invariant across time, or, alternatively, that we can observe “the value of leisure”. Another issue is whether savings can be ignored, but this is less of a problem if saving does not affect nutritional status directly.¹⁰ We will evaluate several alternative measures of expenditure and income, as well as control for labor supply. But since we want to align our work with existing literature, we will depart from the “raw” expenditure variable (total market expenditure plus the value of own-produced goods).¹¹

3.2 Defining the rainfall instrument

The *Kagera Health and Development Survey* datasets include monthly rainfall from 1989, collected from weather stations in the five districts of Bukoba, Karagwe, Muleba, Biharamulu and Ngara. 1992 was a “failure”-year, characterized by a low level of rainfall, but not extremely so. Comparison with longer time-series of rainfall and reports from FAO indicates that the years 1992 and 1993 were plagued by food shortages, although not as severe as the Sahelian drought in the mid-eighties or earlier dry spells.

In related papers (e.g. Hoddinott and Kinsey, 2001, and Paxson, 1992), the authors have taken some care in identifying the correct sowing and harvest seasons in order to strengthen the link between rainfall and income. Our first idea of instrumental variable was a monthly average of the last six months, hoping to extract as much variation in expenditure as possible due to the bimodal cropping regime. There are, however, several problems attached to this approach. First, since the timing of waves was not perfectly

¹⁰If there is perfect consumption smoothing (which is conceptually different from income smoothing; see Morduch 1995), expenditure will not be affected by rainfall, but neither will body weight. If there is partial consumption smoothing, we can not interpret transitory changes in total household expenditure as changes in household income. But as long as body weight remains invariant to saving *per se*, the issue is (simply) one of interpretation, and not identification.

¹¹As it turns out, the first stage relationship is surprisingly insensitive to various definitions of expenditure and income (see Section 5).

semi-annual but sometimes spanned more than nine months, the variance in household expenditure induced by rainfall might be confounded by a seasonal component. Second, cultivation techniques are somewhat dispersed, and not all households in our sample exploits the bimodal cropping regime (the traditional tree crops, like bananas, are usually harvested only once a year). This implies that the impact of a six-month instrument could be different in different periods (maybe even have the opposite sign), also for the very same household. Third, a typical observation is that consecutive weather failure is particularly severe for rural households, and longer time-series of rainfall statistics should be a stronger instrument. We chose to collapse the seasonal variation in perturbation levels, and use a twelve-month moving average as instrument.

The intertemporal variation in rainfall amounts to four points in time. The within-wave variation in rainfall has two sources. The first is the interregional variation across the five weather stations from which the rainfall data was collected. The second source of variation is cluster-based (a “cluster” is typically a village; see World Bank, 2004). For logistic reasons, the timing of the interviews often differed across clusters, sometimes with as much as six months. Under the assumption that the timing of the survey staff visit is uncorrelated to rainfall, after controlling for fixed effects, year and season, we can exploit this “accidental” variation within waves. For our baseline estimates, we associated the rainfall shock with the *month immediately preceding the date of the interview*. The reason not to lag the rainfall shock further and associate it with some time before the recollection period is basically that I_{ht} is theoretically determined by “current production”. Crop yields from tree plants (including the economically important banana tree) are in practice determined by rainfall until the very month of harvest. If this date coincides with the recollection period, our instrument might not capture more recent consumption, especially if our assumption of binding liquidity constraints is true. This can be quite important, since current body weight is likely to be influenced more by recent consumption than past consumption – in fact, the last week of consumption may practically determine most of the intertemporal variation in body weight. Our approach is thus to see recollected expenditure as proxy for current income, and use rainfall up to the date of the interview to correct for any measurement error the usage of such a proxy entails.

The intertemporal relationship between own-mean deviations in household expenditure and past rainfall realizations is depicted in Figure 5, 6 and 7.¹² The time-series of rainfall in these figures represents a regional average

¹²In longitudinal datasets, the households selected in the first wave typically grow in

across the five different weather stations, and the expenditure and income figures represent the mean values for the households surveyed in the relevant month. For some months, only a handful of households was surveyed, which can explain part of the volatility in household expenditure and income. Comparing these three figures is informative since it reveals that a) total expenditure, food expenditure and income covary with rainfall in much the same fashion, implying a low degree of consumption smoothing, and b) the intertemporal variation in income seems to be much higher than the variation in expenditure. In Figure 3 and 4, we study the same relationship across clusters (villages). There is a clear positive pattern, especially between food expenditure and rainfall, but there seems to be much unexplained variance in village-level expenditure left.

Since the instrumented variation in rainfall is based on survey-timing, the year and season effects are based on the actual date of the interview, and not on the specific wave. There were, however, some wave-specific changes to the questionnaire. In particular, in the first wave, the most knowledgeable person in the household was asked to recollect consumption for the last twelve months; in the consecutive waves, the recollection period was six months. We standardized the twelve month figure by dividing it by two. In Section 5.2, we explore the issue further by controlling for wave-specific traits using questionnaire dummies.

As for standard errors, the month of the interview was typically cluster-specific, implying that our instrument is measured on the cluster-level, and, consequently, that the instrumented variation in expenditure is likely to be correlated within clusters. All reported standard errors are therefore corrected for arbitrary correlation within clusters (the “cluster” is also the main stratum of selection in the KHDS dataset, see Section 4 and World Bank 2004). The asymptotic properties of clustered sampling have recently been subject to some interesting research (see Wooldridge 2003, 2006; Donald and Lang, 2006; Bertrand et al. 2004), and simulation studies suggests that if the number of clusters is “large”, cluster-adjusted standard errors perform well in fixed-effects analyses when an explanatory variable is a clustered variable. In our case, there are 52 clusters, which, according to Bertrand et al. (2004) is an acceptable number of groups.

size across time. Since we cannot control for year and seasonal effects in a two-dimensional figure, we choose to use per capita expenditure and per capita income for Figures 3 to 7, dividing e.g. expenditure with the number of current household members.

4 Sample description

The Kagera Health and Development Survey was conducted in four rounds between 1991 and 1994.¹³ The sample is fully longitudinal in the sense that all households that were interviewed the first period were interviewed in subsequent rounds (it was not a rotating panel, common in similar datasets in Africa). The survey staff visited each household four times between 1991 and 1994, in intervals of between six and nine months. In total, around 820 households were surveyed. 4895 individuals was observed at least twice, and the full individual sample size used for the panel analysis (N*T) was 16 640. The sample is stratified, as it was based on a variable probability sampling procedure in order to capture incidences of illness. We will evaluate if and to what extent this procedure reduces the scope of inference in Section 5.4.

There was moderate attrition. About 77% of the population was available for anthropometric records in at least two periods, and only a little more than half the sample was available for all periods. As is the case in all survey-based analyses of health and mortality, it is plausible that this attrition is associated with our dependent variable. For example, insufficient calorie intake may induce illness or even death, and therefore absence. If this is the case, $\hat{\beta}_1$ will be biased towards zero. It is fairly easy to come up with other potential correlations between absence and body weight (due to migration or work) that will distort our estimates. In Section 5.4, we will address these concerns by restricting the sample to individuals that was present at all survey round. If attrition is unimportant, the restricted and unrestricted estimates will not differ.

Finally, there is the issue of outliers in our dependent variable. The measurement of body weight differed somewhat across cohorts and waves. Infants under the age of two were weighted using hanging Salter scales; for older cohorts, standard scales were used. In the third wave, the adult scale was replaced by a digital one, which, according to the KHDS technical report, reduced the variance of young children’s weight. With individual fixed effects, our dependent variable is to be understood as the deviation from the individual’s mean body weight. About 90 individuals had a very high relative standard deviation in body weight across time, and about a handful had extremely high. It is reasonable to ask whether such outliers, be they the result of measurement error or not, are driving our results. To answer this question, we will compare our estimates across various sample restrictions. We will, however, use the sample “as is” for our baseline estimates, throwing

¹³A fifth round was conducted ten years afterwards, in 2004. In this paper, we will use the first four rounds.

away nothing but missing values.

5 Results

5.1 First stage estimation

In Table 2, we study the relationship between rainfall and household income and expenditure.¹⁴ The first model represents a regression using our raw expenditure variable, which corresponds to total expenditure on market goods plus the value of consumption goods produced at home. This specification can be interpreted as a shift of the budget constraint (2) under the assumption that there is no saving and that labor is supplied inelastically. In the second model, we relax the assumption of invariant labor supply, and subtract the value of market labor and own-farm labor from the expenditure figure. Model 3 to 6 consider alternative specifications.

The general wisdom conveyed in Table 2 is that rainfall seems to affect expenditure much in the same way as it affects rural income. The rainfall-expenditure elasticity is practically unchanged when we allow for intertemporal variation in labor supply (see Footnote 9 for a definition). At least for this sample, own-farm labor seems to be quite invariant to weather shocks (this is verified in Model 8, Table 2). The point estimate in Model 3, in which our dependent variable is *net income from agriculture*, is somewhat higher than the model using expenditure, but also less precise, and we cannot reject that it is the same as our baseline estimate. This confirms the tight relationship between fluctuations in income and expenditure, depicted in Figure 5 and 7. It is however noteworthy that by using food expenditure as dependent variable, we obtain a higher point estimate which is statistically different from our baseline point estimate. This suggests that “food”, defined in a broad sense, is indeed an Engel good in this area (provided that prices are kept constant by our controls).

In Table 3, we evaluate the relevance of alternative instruments and some additional household-level covariates. *Total household members* has a positive impact on household expenditure, but the relevant rainfall-expenditure elasticity is unchanged. In Column 3, we include wave-specific dummies in the regression. With year-district effects, the wave dummies captures any remaining variation in expenditure induced by the questionnaire-specific traits

¹⁴All regressions in Table 2 are made on the household level. Neither point estimates nor standard errors differ when we scale up the variables to the cluster level or scale down the variables to the individual level (which is expected since we use cluster-robust standard errors on a sufficient number of clusters). These results are available on request.

mentioned in Section 3.2 and 4. We find that controlling for wave has little impact on the coefficient of interest, and the individual wave effects enter in an statistically insignificant fashion (not shown in Table 3 is the result that the wave dummies drives all other year and seasonal effects into insignificance). Since the inclusion of the wave dummies resulted in a rather severe case of multicollinearity, we choose to exclude them from our baseline estimates.

We find no support for a second order relationship between rainfall and expenditure (Column 4). Interacting our weather instrument with agronomic zone suggest that households residing in “urban” areas are less weather dependent (note, however, that “urban” in the KHDS dataset is a very loose definition, as garden production is still practiced by most households). Lagging the rainfall variable further reduces the point estimate somewhat, but still suggests that there is some persistence in the weather shock effect. Finally, we interact the weather instrument with the mean share the household’s food expenditure (the “Engel ratio”), and find a somewhat stronger effect. The intuition here is that expenditure among households with a high share is more sensitive to rural productivity shocks.

5.2 IV-estimates

In Table 4, estimates of the response of body weight to transitory changes in household expenditure are presented, along with the reduced-form estimates. Our preferred IV-model is that of Column 4. The contrast between our IV estimates and the OLS counterparts is stark. The general impression one would get from regressing body weight on expenditure using OLS is that nutritional outcome is invariant to the purchasing power of the household, since the elasticity is statistically insignificant and very close to zero. We argue that the discrepancy between the IV and OLS estimates is most likely driven by measurement error in the main explanatory variable. As pointed out in Section 2 and elsewhere, expenditure figures based on recollection is likely to be associated with a non-trivial degree of error. In the classical case of zero-mean error in the explanatory variable, the variable will be attenuated towards zero. It is our belief that our instrumental variable strategy evades this bias.

In terms of magnitude, the IV-estimates suggest that 10-percent increase in household expenditure will lead, on average, to a 0.7 percent increase in body weight. For a person that weights 50 kilo, this implies an increase in about 0.35 kilo, which is arguably a non-trivial effect. This estimate is informative insofar it clearly challenges the OLS-result of a statistically and economically insignificant relationship between body weight and expenditure.

However, it is reasonable to ask whether the elasticity is sensitive to age and gender. In Section 5.5, we will provide more meaningful point estimates by disaggregating the sample by age and gender.

Before doing so, however, we will focus on the robustness of the IV estimates reported in Table 4 by addressing some issues that might be thought to violate the exclusion restriction (i.e. the assumption that $cov(r_{ht-1}, e_{iht}) = 0$). If the exclusion restriction does not hold, our IV estimates will be inconsistent. It is therefore crucial that we can maintain the assumption that rainfall affects body weight only via the income channel. A first concern is that body weight is not constant across time. Autocorrelation in our dependent variable would be no problem if our time-series of rainfall would be completely stationary, but as seen in Figure 5, 6 and 7, there seems to be autocorrelation in rainfall as well (by construction, this is so because we use a moving average). Note first, however, that if our panel was completely balanced and every individual in the same district had the same growth rate, our year-district interaction dummies would root out general growth trends. However, human growth is not constant throughout the life cycle, and age-specific autocorrelation in body weight will be magnified if there is attrition, since household members that are observed less frequently would have had more time to naturally increase or reduce weight. We correct for this by including the *time since last survey*-variable, interacted with an adult dummy (Initial age > 15). The variation in *time since last survey* arises from (individual-specific) attrition and the fact that the spacing between interviews differed between waves and across households. As seen in Table 4, controlling for an age-specific growth trend adds much to precision, but does not drastically change the point estimates.

An alternative and perhaps more straightforward way to control for human growth is simply to include *height* as an explanatory variable. However, since child growth can be retarded by malnutrition (see Hoddinott and Kinsey, 2001; Alderman et al. 2005), *height* can be viewed as a proxy for our dependent variable and therefore unsuitable as an explanatory variable. However, as seen in Model 5 in Table 4, *height* does not seem to have a large impact on expenditure coefficient, nor on its standard error. The variable *height* is however statistically significant, even when controlling for an age-specific trend.

Another concern is that climactic factors may induce the spread of certain diseases, which, in turn, may reduce body weight. Although well-documented in epidemiological literature, this fact has been given less attention in related economic exercises (but see the discussions in Thomas and Strauss 1998 and Foster 1995). In the region of Kagera, both malaria and cholera epidemics have been triggered by rainfall in this fashion, so these concerns clearly have

some merit in this context. Fortunately, the KHDS dataset is sufficiently rich for us to control for incidences of malaria and other illnesses (the survey included questions of both self-reported and diagnosed health status). Under the assumption that illness represents a random shocks to body weight, the IV estimates should be robust to the inclusion of such variables. Since many household members reported being sick but unable or unwilling to seek professional care, we based our illness variables on the respondent’s own diagnosis. The symptoms of *malaria* are typically well-known, whereas for cholera, we used *diarrhea* as a proxy.

As seen in Table 4, incidences of diarrhea and malaria are associated with a reduction in body weight, and the effects are statistically significant. Our expenditure coefficient is only marginally changed. It is important to note that the regressions including the illness variables are by no means to be seen as preferable to the baseline estimate. The problem is again that body weight is a proxy for “health” in a broad sense, and including additional health covariates will bias the elasticity downwards (note, however, that health expenditures seem to be invariant to the weather shock; see Table 2). The reason that the illness-variables enter in a statistically significant way in our regression without affecting the point estimate implies that losing weight is not regarded as an illness *per se* among the respondents in our sample.¹⁵ The results conveyed in Table 2 suggest that most of the variation in reported illness is uncorrelated to the error term, given our other covariates.

In Table 4, we also provide some supplementary tests of the robustness of our IV-estimates by including additional covariates that might be thought to confound the empirical strategy. For example, *schooling* and *own-farm work* can be regarded as time allocation decisions that may be incompletely captured in household expenditure but may nevertheless affect body weight. We also control for *pregnancy* as a way of addressing the seasonal component in fertility decisions, as well as family composition (*total household members* as recorded on the wave-roster). Finally, we include controls for questionnaire design (i.e. differences in how the survey was implemented across waves). None of these additional covariates has any dramatic impact on expenditure-body weight relationship (although most of them enter in a statistically significant way). The wave dummies has some impact on the *F*-test of the excluded instrument in the first stage regression, which is most likely a result of the covariance between the wave dummies and the district-year and seasonal dummies. The importance of these wave controls should

¹⁵Indeed, when those who reported being ill was asked for a diagnosis, only a few reported "malnutrition".

thus not be over-dramatized, and they have only a slight impact on point estimate and standard error of our income effect. The effect is still statistically significant at the 5%-level.

Finally, we also report the reduced-form estimates in Table 4. As illustrated in Section 2.2, the IV estimates identify an income effect under the tacit assumption that decisions regarding household production are separable from those of consumption. We have shown some evidence that labor supply seems to be invariant to the rainfall shock, but another concern may be that malnutrition reduces household productivity (e.g., if “calories” are production factors). The reduced-form estimates relies on the weaker assumption that a rainfall shock enters the production function (and not preferences), and represent the net effect of this productivity shock. Now, the exact point estimates of the “rainfall elasticity of body weight” may convey little in terms of external relevance, but can be informative when comparing how the effects differ internally, e.g. across household members. We will return to this issue in Section 5.5, where we also provide a strong case against the possibility of a pronounced simultaneity between household productivity and nutrition.

5.3 Alternative instruments and explanatory variables

In Table 5, we study three alternative forms of weather-shocks as instruments across four alternative definitions of income and expenditure. As can be seen in the three wide columns in Table 5, different rainfall definitions produces slightly different point estimates (the typical differences is about one standard error). The effects are still statistically significant. One can further note that the rainfall-Engel ratio instrument is stronger in terms of marginal contribution to the explanatory power of the first stage regression. Those who were concerned that an F -statistic around 15 was too close to the lower bound suggested by e.g. Staiger and Stock (1997) should be relieved to learn that the point estimates seem robust to a stronger, but perhaps less intuitive, instrument.

Using *net income from agriculture* as an explanatory variable produces lower point estimates than using household expenditure. The difference is within one standard error of the baseline point estimate, and it would probably be a little drastic to conclude that this discrepancy is evidence of (partial) consumption smoothing. Rather, the difference is likely driven by the fact that our instrument is not as strong for agricultural profits as it is for expenditure, plausibly due to the measurement issues surrounding this variable. Using *household expenditure plus the value of leisure* also produces lower estimates, although well within one standard error from the baseline estimates. These results once again verify that although household income and

household expenditure are conceptually different in both theory and according to survey-based stylized facts, they seem to play very similar roles in our fixed-effects instrumental variable models.

5.4 Further robustness tests

In Table 6, we address three sample-oriented concerns raised in Section 4, namely that the sample was stratified with respect to plausibly endogenous variables, that there was non-negligible attrition, and that there were some outliers with respect to intertemporal body weight variation.

The stratification was based on “mortality-risk” at both the community and household level (the survey objective was to study the impact of HIV/AIDS on local economies in Africa). In the first step, clusters of households were randomly selected from predetermined PSUs corresponding to different agronomic zones. The probability that the randomly selected clusters were kept was proportional to the level of mortality reported in the 1988 Tanzanian Census. The second stage kept individual households in a similar fashion. Preferably, we would use sample weights to evaluate whether the external validity of our results goes beyond the sample at hand. Unfortunately, the household weights included in the KHDS dataset are incomplete and inconsistent with the alleged sampling methodology (World Bank, 2004). Of course, even with perfect weights, it is also questionable whether they would have added anything to consistency and inference when the objective is to estimate a behavioral model. In the econometric model (Equation 9), the population parameter β is assumed to be identical in each population stratum. If this assumption holds, then the unweighted IV estimator will be consistent, and since 2SLS is most efficient, it is preferable.¹⁶ Still, it is possible to address the sampling issue in a more rudimentary way, namely by disaggregating the sample across indicators of village-level adult mortality – the main selection variable – and study whether $\hat{\beta}$ differs across these subsamples. The results from such an exercise are presented in the first and second column in Table 6. The results suggest that any heterogeneity between selected and non-selected clusters is controlled for by our fixed-effects and other controls.

In the next three columns in Table 6 (3-5), we test whether our results are sensitive to the exclusion of individuals who experienced an extremely large variation in body weight across time. To do this, we stepwise excluded those

¹⁶For the “econometric argument” against sample weights, see the discussion in Deaton (1999) and Duncan and DuMouchel (1994). Furthermore, if there is first-stage heterogeneity, our TSLS estimates are informative only for those whose expenditure is affected by rainfall (see e.g. Angrist and Imbens 1995), and the scope of inference is further limited.

whose standard deviation in body weight was higher than a given threshold (20, 15 and 10 percent of their own mean body weight, respectively). As seen, excluding these outliers has little impact on the main result. In the final column, we also impose the restriction that each individual had to be present all survey rounds (once again excluding outliers), in order to study whether the results are robust to non-response. We find no evidence of selection into survey participation.

5.5 Heterogeneity (subsample results)

The “income elasticity of body weight” is, on average, about 0.05 to 0.07. The relevance of this figure is best appreciated in light of the weak performance of the OLS estimates, but is perhaps of limited usefulness for policy as long as we do not provide further estimates across age and gender.¹⁷ This section aims to do just that.

Ideally, since age and sex are the two most important determinants of body mass and stature, we would like to disaggregate our sample in narrow age categories – perhaps even “age in months” for the youngest children – much like how the World Health Organization and UN’s Food and Agricultural Organization have obtained their estimates of human energy requirements (FAO, 2001). Our sample is unfortunately too small for such an approach. Instead, we use age windows of two to four years, and capture the heterogeneity by using rainfall-age interaction terms and disaggregating the sample across gender. We also interacted initial age with the linear trend variable; however, due to lacking observations among some subgroups, we had to assume that the year-district and seasonal effects are the same across age categories.

Another issue is the question of first-stage heterogeneity. If the weather instrument is weaker (in terms of efficiency) for “adult households”, then the second step estimates for these households will be less precise (and a potential bias will be blown up). In order to evaluate this concern, we present the reduced form estimates and first stage estimates in Table 7. Two important messages can be read from that table. First, the reduced-form estimates suggest that children between age 2 and 15 are more sensitive to weather shocks, and that the body weight of females, in general, is more sensitive. The body weight of adult males is completely invariant to changes in rainfall. This seems to be the case for infants (Age<2) as well. For both male and

¹⁷Behrman and Deolalikar (1990) report that female nutrition is more sensitive to changes in food prices – although they find no income effect for any family member. Regarding excess female mortality in related papers, see Rose (1999) and Rosenzweig and Schultz (1982).

female infants, the coefficient is within one standard error from zero. This suggests that the nutritional status of infants is invariant to transitory income shocks. However, we have already mentioned that the scale was replaced at age two, and there is also some attrition in this group (because some of them were not born the first wave). Interpreting the estimates for the infant groups should be done with caution.

The second message is that the reduced-form heterogeneity is not mirrored by first stage heterogeneity. This is perhaps not so shocking given that expenditure is measured on the household level. However, there could have been selection into weather-dependent households; for example, households with many children in the household might depend on rain-fed crop yields to a greater extent than others. As seen in columns 4 to 6, however, we cannot reject that the age-rainfall interaction terms are statistically the same in the first stage regressions. This suggests that the reduced-form heterogeneity can be interpreted as differences in how individual nutritional status responds to a reduction in total household expenditure (as opposed to differences in how expenditure responds to weather shocks across households with different family composition).

The IV-estimates reported in Table 8 verify this interpretation, and gives a quantitative dimension to it. In particular, they reveal that the effects are quite substantial for female children (2-9). These results do not seem to be driven by spurts in growth or incidences of illness. In Column 4 and 8, we exclude outliers in intertemporal body weight variation – i.e. persons that experienced extremely volatile fluctuations in weight. As seen, this had some impact on the estimates on the group of young teenagers (10-14 years), but the general picture is quite robust.

For female children (with a mean weight of 20 kilo), a 10-percent increase in household expenditure induces a change in body weight with 0.4 kilo; for boys, the magnitude is 0.2 kilo. Again, we argue that it is reasonable to interpret these effects as income effects. If we take WHO's weight-for-age z-scores seriously, a negative income shock of that magnitude is enough to increase the fraction of moderately malnourished female children in the KHDS sample from about 19 percent (the first wave) to 25 percent. Conversely, the body weight of adults males seem to be practically invariant to changes in household income. This finding adds further credence to our assumption that decisions on consumption are separable from those of production, since the feedback from child malnutrition to household income should be marginal compared to how adult malnutrition may affect productivity (marginal, but not negligible; see Beegle et al, 2006).

6 Concluding remarks: How are these estimates helpful for policy?

We find that quite small fluctuations in household income may induce severe malnutrition among children – and, in particular, among female children. This suggests that fluctuations in household expenditure (i.e. an inability or unwillingness to insure against income fluctuations) can have quite severe secondary effects on child health. Being a normative measure of the severeness of consumption fluctuations, our results further highlight the need for proper insurance mechanisms in rural developing areas.

But it is important to note that a high income elasticity of body weight is not only bad news. If there is a pronounced relationship between nutritional status and current income, exogenous *increases* in household income will have a *positive* impact on the nutritional status of children. In other words: income-augmenting policies that manage to reach the households might be as effective in reducing child malnutrition as those that address nutrition more directly. This implication is seldom highlighted in related literature on rural income variability, possibly since most previous findings are based on reduced-form estimates that study the effects of a decidedly “adverse” or “unfavorable” shock in itself. To the extent that our estimates can be interpreted as income-effects – and we argue that they can – one must acknowledge that the elasticities work in both directions.

A third question that our estimates shed light on is the nature of the health-income nexus. Since the nutritional status of adults remains invariant to income changes, it seems unlikely that negative income shocks will induce so-called poverty traps, via feedback from nutrition to productivity, in the short run. Rather, such effects will be visible first in the long run when the children affected by malnutrition have grown up. A relevant question is to what extent illness during childhood induces permanent health and welfare effects. There is an emerging body of literature that assesses the long-term effect of transitory health shocks, which hopefully will receive more contributions in the future.

A Appendix: Tables and Figures

Table 1: Descriptive statistics of the KHDS sample (first wave)

	Mean	Median	First quartile	Third quartile
Household expenditure (Tnz)	387 331	274 028	159 221	448 438
Household income (Tnz)	255 790	175 506	96 189	322 302
Per capita expenditure (Tnz)	82 447	54 180	34 054	91 423
Per capita expenditure (Dollar)	276	181	306	114
Household members	5.6	5	3	7
Food exp. / Total exp.	.6355	.646	.542	.756
Cash savings / Total exp.	.0418	.00782	.00131	.0237
Cons. of own-prod. goods / Total exp.	.449	.465	.321	.595
Cons. of own-prod. food / Total food cons.	.663	.727	.850	.545
Market labor income / Total exp.	.0698	0	0	.046
Ass. from organizations / Total exp. ¹	.012	.0046	.0023	.011
Ass. from oth. individuals / Total exp. ¹	.0580	.0249	.0105	.0597
<i>Nutritional statistics</i>				
Body mass index (Age>18)				
Females	21.46	21.06	19.17	22.97
Males	20.43	20.3	18.75	21.82
Weight-for-age z-score				
Females				
Age 0-5	-1.111	-1.28	-1.96	-0.343
Age 6-10	-1.04	-1.04	-1.72	-0.440
Age 11-18	-1.20	-1.28	-1.82	-0.666
Males				
Age 0-5	-1.32	-1.39	-2.10	-0.545
Age 6-10	-1.11	-1.15	-1.70	-0.562
Age 11-18	-1.77	-1.83	-2.39	-1.226

All expenditure variables are annual values, based on recollection. "Household members" is the number of current household members recorded on the roster. Dollar exchange rate: 298.5Tnz/Dollar, based on the average exchange rate between the first quarter of 1990 and the first quarter of 1992. The weight-for-age z-score is calculated using a sex and age specific reference weight. A z-score of 0 means normal nutrition status, -1 light malnutrition, -2 moderate malnutrition, and -3 severe malnutrition. See WHO (1986).

¹ Mean across all waves.

Table 2: The impact of rainfall on expenditure, income, time-allocation and prices in Kagera, 1991-1994. Household sample.

Dependent variable	Coefficient	Std. error	Obs.	R^2
1. Household expenditure	0.463***	0.132	2974	0.055
2. HH expenditure minus value of leisure	0.486***	0.160	2442	0.047
3. Net income from agriculture	0.614***	0.204	2865	0.046
4. Expenditure divided by nr of hh members	0.527***	0.139	2974	0.075
5. Expenditure divided by price index	0.456***	0.134	2974	0.101
6. Food expenditure	0.686***	0.154	2968	0.071
7. Any health expenditures?	0.0206	0.0677	2974	0.014
8. Any hh member performed own-farm work?	.0454	0.0685	2974	0.018
9. Local (cluster-level) price index	0.00707	0.0287	2974	0.938

Cluster-adjusted standard errors in parenthesis. All models include year-district dummies and seasonal controls. Models 1-6 are log-log estimations. Explanatory variable: average monthly rainfall the year preceding the interview (log).

Table 3: The impact of rainfall on total household expenditure. Household sample.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Rainfall (log)	0.463*** (0.132)	0.492*** (0.132)	0.478*** (0.184)	0.761 (2.141)			
Total hh members		0.0690*** (0.0130)	0.0685*** (0.0129)				
Wave 1			-0.130 (0.267)				
Wave 2			-0.0809 (0.185)				
Wave 3			-0.0588 (0.0930)				
Rainfall squared				-0.0319 (0.232)			
Rainfall x Tree crop zone					0.529* (0.306)		
Rainfall x Cereal zone					0.552** (0.261)		
Rainfall x Cotton zone					0.583*** (0.191)		
Rainfall x Urban zone					0.188 (0.203)		
Rainfall (lagged 6 months)						0.398*** (0.118)	
Rainfall x Engel ratio							0.744*** (0.199)
Observations	2974	2974	2974	2974	2974	2974	2974
R^2	0.055	0.075	0.075	0.055	0.056	0.055	0.056

All models include year-district interaction terms, seasonal (quarterly) dummies, and fixed effects at the individual level. Cluster-adjusted standard errors in parenthesis. "Engel ratio" is the household's mean food expenditure divided by mean total expenditure (a time-invariant variable). Dependent variable: total household expenditure (log). Rainfall is average monthly rainfall the year preceding the interview. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: How responsive is body weight to transitory changes in expenditure? Comparison between OLS and IV. Full sample.

	(1) OLS	(2) OLS	(3) IV	(4) IV	(5) IV	(6) IV	(7) IV	(8) IV
Expenditure (log)	0.00189 (0.00182)	0.00177 (0.00191)	0.0619** (0.0272)	0.0742*** (0.0253)	0.0772*** (0.0262)	0.0766*** (0.0266)	0.0781*** (0.0261)	0.0774** (0.0323)
Time since last survey	0.00796*** (0.00117)	0.00488*** (0.00152)	0.00780*** (0.00158)	0.00502** (0.00206)	0.00502** (0.00206)	0.00498** (0.00206)	0.00537*** (0.00199)	0.00873** (0.00433)
Time since last survey x Initial age > 15	-0.00957*** (0.000217)	-0.00682*** (0.000861)	-0.00946*** (0.000227)	-0.00653*** (0.000855)	-0.00653*** (0.000855)	-0.00654*** (0.000851)	-0.00673*** (0.000867)	-0.00672*** (0.000859)
Height in Centimeters		0.00662*** (0.00197)	0.00662*** (0.00197)	0.00662*** (0.00197)	0.00662*** (0.00197)	0.00662*** (0.00197)	0.00654*** (0.00195)	0.00656*** (0.00196)
Malaria last 4 weeks		-0.00882*** (0.00198)	-0.00881*** (0.00196)	-0.00881*** (0.00196)	-0.00881*** (0.00196)	-0.00881*** (0.00196)	-0.00862*** (0.00195)	-0.00856*** (0.00199)
Diarrhea last 4 weeks		-0.0157 (0.00975)	-0.0215** (0.0103)	-0.0215** (0.0103)	-0.0215** (0.0103)	-0.0215** (0.0103)	-0.0209** (0.0104)	-0.0208** (0.0102)
Currently Pregnant		0.0467*** (0.00450)	0.0467*** (0.00450)	0.0467*** (0.00450)	0.0467*** (0.00450)	0.0467*** (0.00450)	0.0443*** (0.00469)	0.0445*** (0.00483)
Attended school last week		-0.0204*** (0.00352)	-0.0204*** (0.00352)	-0.0204*** (0.00352)	-0.0204*** (0.00352)	-0.0204*** (0.00352)	-0.0199*** (0.00335)	-0.0198*** (0.00333)
Own-farm work last week		-0.00327** (0.00161)	-0.00327** (0.00161)	-0.00327** (0.00161)	-0.00327** (0.00161)	-0.00327** (0.00161)	-0.00277 (0.00202)	-0.00260 (0.00200)
Total household members		-0.000978 (0.000767)	-0.000978 (0.000767)	-0.000978 (0.000767)	-0.000978 (0.000767)	-0.000978 (0.000767)	-0.00587*** (0.00224)	-0.00572** (0.00269)
Wave dummies	No	No	No	No	No	No	No	Yes
Rainfall (<i>reduced form</i> ; <i>separate regression</i>)			0.0283** (0.0112)	0.0363*** (0.00907)	0.0379*** (0.00917)	0.0375*** (0.00939)	0.0393*** (0.00961)	0.0339*** (0.0104)
Observations	16638	16617	16638	16638	16617	16617	16617	16617
R ²	0.425	0.474	0.213	0.341	0.375	0.377	0.382	0.384
Craig-Donald F			122.4	131.6	132.1	131.0	141.7	47.02
Partial F-test of excl. instr.			13.64	14.87	14.91	14.64	15.27	8.23

All models include year-district interaction terms, seasonal (quarterly) dummies, and fixed effects at the individual level. Cluster-adjusted standard errors in parenthesis. Dependent variable: log of body weight in kilos. Excluded instrument is mean monthly rainfall the year preceding the interview (log). Preferred estimates in bold.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: How responsive is body weight to transitory income shocks? Comparison across alternative instruments and explanatory variables.

	Instrument: Rainfall (log)			Instrument: Rainfall (log) x Engel ratio			Instrument: Rainfall					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Expenditure (log)	0.0742*** (0.0253)				0.0613*** (0.0207)				0.0881*** (0.0317)			
Expenditure plus the value of leisure (log)		0.0675*** (0.0231)				0.0557*** (0.0193)				0.0777*** (0.0262)		
Net income from agriculture (log)			0.0576*** (0.0200)				0.0438*** (0.0150)				0.0748*** (0.0252)	
Food expenditure (log)				0.0543*** (0.0170)				0.0477*** (0.0152)				0.0612*** (0.0192)
Observations	16638	14313	15967	16606	16638	14313	15967	16606	16638	14313	15967	16606
R ²	0.341	0.317	0.315	0.357	0.369	0.353	0.362	0.372	0.306	0.279	0.238	0.338
Craig-Donald F-test	131.6	103.6	98.25	167.5	147.2	114.5	124.9	163.8	120.1	102.5	76.73	170.0
Partial F-test of excl. instr.	14.87	18.20	11.61	16.99	19.66	25.39	15.55	21.24	11.50	15.15	11.12	16.14

All models include year-district interaction terms, seasonal (quarterly) dummies, an age-specific linear trend and fixed effects at the individual level. Cluster-adjusted standard errors in parenthesis. Dependent variable: log of body weight in kilos. "Engel ratio" is the household's mean food expenditure divided by mean total expenditure (a time-invariant variable). "Rainfall" is mean monthly rainfall the year preceding the interview.
 * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: How responsive is body weight to transitory changes in expenditure? Sensitivity to stratification, outliers and non-response.

	(1) High mortality rate	(2) Low mortality rate	(3) Relative weight variation <0.2	(4) Relative weight variation <0.15	(5) Relative weight variation <0.1	(6) Relative weight variation <0.15 - No attrition
Expenditure (log)	0.0778*** (0.0260)	0.0753* (0.0448)	0.0469*** (0.0171)	0.0508*** (0.0160)	0.0573*** (0.0197)	0.0615*** (0.0210)
Observations	8349	8289	16046	15514	13197	10916
R^2	0.341	0.333	0.531	0.507	0.353	0.496
Craig-Donald F	77.85	66.76	123.7	124.5	88.24	81.55

All models include year-district interaction terms, seasonal (quarterly) dummies, an age-specific linear trend and fixed effects at the individual level. Cluster-adjusted standard errors in parenthesis. Dependent variable: log of body weight in kilos. Excluded instrument: mean of monthly rainfall the year preceding the interview.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: How responsive is body weight to transitory income shocks? Age-gender decomposition.

	Reduced-form estimates			First stage estimates		
	Dep. var: log of body weight			Dep. var: log of hh expenditure		
	All (1)	Males (2)	Females (3)	All (4)	Males (5)	Females (6)
Rainfall (log) x Initial age <2	0.00316 (0.0460)	0.0392 (0.0601)	-0.0385 (0.0529)	0.514*** (0.152)	0.600*** (0.190)	0.417** (0.180)
Rainfall (log) x Initial age 2-4	0.0728*** (0.0162)	0.0453*** (0.0148)	0.102*** (0.0273)	0.469*** (0.143)	0.482*** (0.164)	0.456*** (0.156)
Rainfall (log) x Initial age 5-9	0.0625*** (0.0114)	0.0497*** (0.0149)	0.0765*** (0.0147)	0.423*** (0.129)	0.516*** (0.132)	0.363** (0.146)
Rainfall (log) x Initial age 10-14	0.0228* (0.0138)	0.00120 (0.0139)	0.0455** (0.0226)	0.545*** (0.132)	0.502*** (0.151)	0.593*** (0.138)
Rainfall (log) x Initial age 15-19	0.00978 (0.00938)	-0.00269 (0.0121)	0.0185 (0.0155)	0.608*** (0.147)	0.647*** (0.155)	0.588*** (0.186)
Rainfall (log) x Initial age 19<	0.00990 (0.00796)	0.00410 (0.00961)	0.0177* (0.0104)	0.452*** (0.131)	0.395*** (0.130)	0.502*** (0.140)
Observations	16642	7982	8653	16640	7982	8651
R^2	0.513	0.514	0.520	0.074	0.072	0.083

All models include year-district interaction terms, seasonal (quarterly) dummies, age-specific time trends and fixed effects at the individual level. Cluster-adjusted standard errors in parenthesis. "Rainfall" is average monthly rainfall the year preceding the interview.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: How responsive is body weight to transitory changes in expenditure? IV estimates across gender and age categories. Log-log estimation.

	Males			Females				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Expenditure (log)	0.0615	0.0450	0.0595	0.0326	-0.114	-0.128	-0.102	-0.103
x Initial age <2	(0.101)	(0.0948)	(0.102)	(0.0416)	(0.156)	(0.152)	(0.153)	(0.103)
Expenditure (log)	0.0901**	0.109**	0.0893**	0.114***	0.229***	0.237***	0.232***	0.208***
x Initial age 2-4	(0.0400)	(0.0483)	(0.0400)	(0.0390)	(0.0874)	(0.0884)	(0.0889)	(0.0625)
Expenditure (log)	0.0886***	0.0977***	0.0875***	0.0867***	0.216**	0.224**	0.218**	0.207**
x Initial age 5-9	(0.0316)	(0.0304)	(0.0315)	(0.0261)	(0.0877)	(0.0909)	(0.0896)	(0.0813)
Expenditure (log)	-0.00456	-0.00972	-0.00590	0.0393*	0.0749*	0.0834**	0.0746*	0.0470**
x Initial age 10-14	(0.0299)	(0.0321)	(0.0300)	(0.0217)	(0.0394)	(0.0406)	(0.0397)	(0.0234)
Expenditure (log)	-0.00858	-0.00133	-0.00943	-0.0197	0.0300	0.0380	0.0294	0.0175
x Initial age 15-19	(0.0189)	(0.0190)	(0.0189)	(0.0177)	(0.0258)	(0.0256)	(0.0262)	(0.0242)
Expenditure (log)	0.00350	0.00729	0.00143	-0.00183	0.0314	0.0326	0.0328	0.0290
x Initial age >19	(0.0258)	(0.0271)	(0.0262)	(0.0207)	(0.0273)	(0.0286)	(0.0284)	(0.0239)
Height		0.00413				0.00444***		
		(0.00355)				(0.00154)		
Malaria last 4 wks			-0.00471				-0.00469	
			(0.00348)				(0.00307)	
Diarrhea last 4 wks			0.00368				-0.0338**	
			(0.00974)				(0.0137)	
Sample restriction	None	None	None	Excl. outliers	None	None	None	Excl. outliers
Observations	7980	7971	7980	7422	8651	8639	8651	8086
R ²	0.484	0.492	0.485	0.569	0.273	0.263	0.273	0.217
Craig-Donald F	8.909	9.076	8.926	7.909	9.060	9.073	8.768	8.530

All models include year-district interaction terms, seasonal (quarterly) dummies, age-trend interactions and fixed effects at the individual level. Cluster-adjusted standard errors in parenthesis. Dependent variable: log of body weight in kilos. Excluded instrument is mean monthly rainfall the year preceding the interview (log), interacted with the relevant age-categories

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$




Figure 1: Weight for age z-score. KHDS, Wave 1. Children (Age<18.)




Figure 2: Body mass index (BMI). KHDS Wave 1. Adults (Age>17).




Figure 3: Deviation in lagged rainfall (log; x-axis) versus deviation in household per capita expenditure (log; y-axis). Cluster (village) means.




Figure 4: Deviation in lagged rainfall (log; x-axis) versus deviation in household per capita food expenditure (log; y-axis). Cluster (village) means.

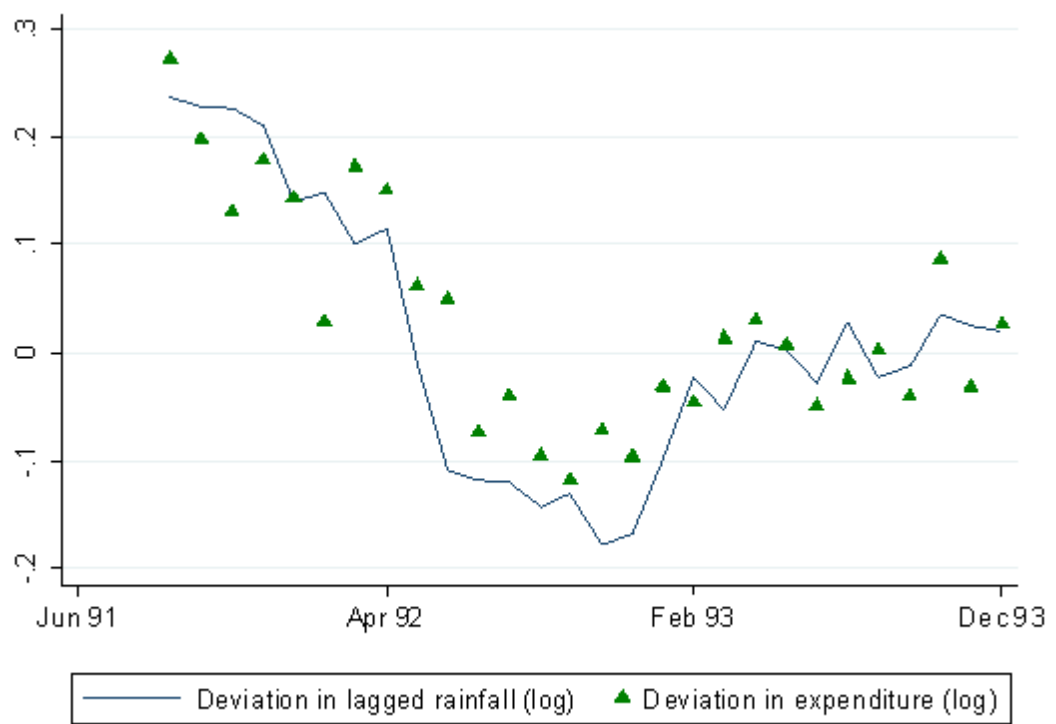


Figure 5: Rainfall and per capita expenditure in Kagera, 1991-1994.




Figure 6: Rainfall and per capita food expenditure in Kagera, 1991-1994.




Figure 7: Rainfall and per capita income in Kagera, 1991-1994.