ESTIMATING HOUSEHOLD WELFARE FROM DISAGGREGATE EXPENDITURES

ETHAN LIGON

Abstract. Existing models of life-cycle demand typically assume that expenditure shares for particular goods within a period are fixed, but this is at sharply at odds with strong empirical evidence, including Engel's Law. We show how one can exploit variation in the composition of expenditures to estimate demand systems that are flexible and may feature highly non-linear Engel curves; this same procedure yields an index of household welfare closely related to the marginal utility of expenditures within a period. We use these methods with repeated cross-sectional expenditure surveys from Uganda to estimate an incomplete demand system and household welfare in different periods.

ligon@berkeley.edu
University of California
Berkeley, CA 94720–3310

. *Date*: September 14, 2018.

. Much of this paper came together during a sabbatical stay at the UCL, and so I'd like to particularly thank Orazio Attanasio, Richard Blundell and their colleagues for their hospitality and encouragement. I must also thank Pascaline Dupas, Marcel Fafchamps, Laura Schechter, and Duncan Thomas for very helpful critical comments.

Measures of household or individual-level consumption expenditures are central to policy-relevant statistics such as poverty rates in most low income countries, and are also critical inputs to a wide variety of important research questions in many fields of economics, particularly in models involving risk, inequality, or life-cycle behavior (e.g., Angelucci and Giorgi 2009; Lise and Yamada 2017).[1] The household surveys used to collect these data almost invariably record disaggregate expenditures; that is, expenditures on many different kinds of goods or services. However, empirical work employing these data to measure changes in household welfare typically focuses on the sum of these disaggregate expenditures, divided by a price index, or *total* real household consumption expenditures.[2] Information on the *composition* of the household's consumption portfolio is usually discarded. The question of this paper: How can information on the composition of the consumption portfolio be exploited to measure households' material well-being?

To answer this question we introduce two main innovations. The first is to devise a new incomplete demand system which allows for variation in both the *scale* and *composition* of households' consumption portfolios to flexibly arise from changes in prices or households budgets. Second is to construct a simple estimator of this demand system which allows us to recover not only critical demand elasticities, but also household-specific latent variables which we show can be interpreted as the household-specific multipliers (λ) associated with households' within-period budget constraints.

We call our new demand system the Constant Frisch Elasticity (CFE) demand system. It is unique in that it is globally regular, has unrestricted rank (in the sense of Lewbel 1991), and can be shown to nest most of the theory-consistent demand systems in the empirical literature. We further devise a novel but simple method of estimation which exploits information in both the first and second moments of the distribution of log expenditures. These features allow us to estimate the first example of a globally-regular demand system with unrestricted rank which can be used to measure consumer welfare even if expenditures on only some goods and services are observed.

We next show how the demand system we derive can be estimated using one or more rounds of cross-sectional data on disaggregate household expenditures, in a specification involving logarithms of those disaggregate expenditures. This estimator delivers "Frisch elasticities" and estimates of the index λ of each household's marginal utility of expenditures (MUE), along with estimates of the effects of various observable household characteristics on demand. Importantly, unobserved household characteristics are naturally introduced in such a way that they do not bias estimates of the key demand elasticities.

1. See Attanasio and Pistaferri (2016) for a survey focused on the US case.

2. For example, while Browning, Crossley, and Winter (2014) provides an excellent recent survey of methods for measuring household consumption expenditures, the authors choose to focus exclusively on total household consumption expenditures.

Our estimator is implemented in two steps. The effects of prices and household characteristics on expenditures are obtained in a first seemingly-unrelated regression (SUR) step. A distinctive feature of our approach is that information regarding the composition of the consumption portfolio is then obtained using a singular value decomposition of the matrix of residuals, yielding estimates of Frisch elasticities and λs.

We observed above that measures of consumption expenditures are critical inputs to measuring poverty, inequality, and risk. We use our methods to explore these aspects of welfare using four rounds of data on household expenditures in Uganda, spanning the period of 2006–2012. These data are instances of the World Bank's *Living Standards Measurement Surveys* (Deaton 1997), which have now been conducted in many countries across many years, but the Ugandan data have particular interest because they span a period which includes the global "food price" crisis of 2008 during which the prices of important cereals more than doubled, as well as a "great recession" experienced by Uganda in 2010. With nothing more more than cross-sectional variation in expenditures in these data, we're able to obtain estimates of both the parameters of the demand system and the households' λs, and to measure the consequences of these shocks for rates of poverty, inequality, and risk aversion for differently situated households.

1. Related Literature

Our paper is related to several distinct threads of research, which we discuss in turn. The first is a body of research which attempts to leverage "Engel's Second Law" to obtain measures of welfare from disaggregate expenditures. As the name suggests, the idea of measuring welfare from disaggregate expenditures is as old as Engel (1857), but there's been recent interest in using the details of household expenditures or Engel curves for particular goods to construct welfare-related measures. For example, Almås (2012) estimates Engel curves for food from many different countries, and uses these results to correct bias in international PPP statistics and measures of cross-country inequality. In this paper we address a similar problem using micro-data to measure inequality and correct CPI statistics in Uganda, but without adopting Almås' assumption that there's a fixed relationship between food share and real income. Attanasio and Lechene (2014) estimate Engel curves for food using data from the Mexican Progresa experiment with the aim of testing the collective model of the household (Bourguignon, Browning, and Chiappori 2009). They assume that utility from food is additively separable from leisure and other consumption, a set of assumptions slightly weaker than those of Almås. Atkin et al. (2018) adopt a much weaker set of assumptions regarding separability and use estimates of what they call "relative Engel" curves to measure changes in welfare. But unlike our approach this requires the very strong assumption that there are no changes in relative prices which affect the slope of these curves; related, Almås, Beatty, and Crossley (2018) impose structure associated with the AIDS demand system and impose a sort of conditional separability. Finally, Young (2012) constructs a demand system meant to allow welfare comparisons across time for many different countries in Africa. But the actual goods used for estimation are mostly fairly fixed household assets and characteristics, and the demand system itself isn't theory consistent. The relation between what we're calling λ and consumer demand was first extensively considered by Ragnar Frisch (see esp. Frisch 1959, 1964, 1978), but wasn't developed empirically until demand systems which depend on prices and marginal utility were revived by Heckman and MaCurdy in the seventies. This empirical development allowed econometricians to deal with the difficulties of measuring permanent income in the context of life-cycle models (Heckman 1974, 1976; Heckman and MaCurdy 1980; MaCurdy 1981). This was pioneering, but central to identification in these early applications was the assumption that markets were complete so that risk-averse households could completely eliminate idiosyncratic risk. These papers also imposed much more structure on within-period demands than is necessary, requiring that all consumption goods have identical income and price elasticities, and featured demand systems which implicitly violated the symmetry characteristic of regular demands.[3] Nevertheless, the payoff from these assumptions was significant—by assuming a particular cardinalization of homothetic preference structures and constant MUEs it was possible to estimate these as fixed effects in a linear panel regression. A later generation of life-cycle models relaxed the assumption of full insurance, and made the idiosyncratic risk borne by households a focus of the analysis. One early example is Browning et al. (1994), which replaces the assumption of perfect insurance markets with an assumption of perfect credit markets and cleverly achieves identification of changes in households' $\log \lambda$ with a cardinalization and an assumption that these changes are normally distributed. A second important collection of papers replaces the goal of *estimating* models featuring homothetic utility and full insurance with the aim of *testing* for full insurance (Mace 1991; Deaton 1992; Cochrane 1991; Townsend 1994). More recent examples connect risk in labor earnings with persistence in consumption to the evolution of inequality (Low, Meghir, and Pistaferri 2010; Lise 2012; Blundell, Pistaferri, and Saporta-Eksten 2016; Arellano, Blundell, and Bonhomme 2017). The focus of all of these papers remains on highly aggregate forms of consumption and leisure; none of these exploit within-period consumer choices among disaggregate consumption goods, as in the present paper. The life-cycle models described so far use variation over both time and households to simultaneously estimate within-period demands and the the dynamic elements of model.

3. Somewhat later demands which depended on prices and λ were given the moniker "Frisch demands" by Martin Browning (Browning, Deaton, and Irish 1985), but estimation of these imposed not only the unnecessary restrictions of Heckman and MaCurdy, but also imposed linearity on Engel curves, and implausible restrictions on intertemporal marginal rates of substitution. Recognition of this latter problem (Browning 1986) seems to have led to a general abandonment of the approach for many years.

However, we are often much more confident that we understand the within-period allocation decisions made by households than we are in the exact nature of the frictions which shape households' intertemporal behavior. Blundell (1998) points out that if preferences are intertemporally separable then the problem can be tackled in two sequential steps, first using cross-sectional information to partially identify marginal utilities of expenditure (our λs), and then using variation in λ over time to estimate or interpret intertemporal behavior. Thus, our approach can be thought of as an implementation of the first step in this sequential approach, yielding estimates of the Lagrange multipliers associated with households' budget constraints using information about contemporary expenditures. This means that estimating an *index* of the 'true' marginal utility of expenditures exhausts the information available in contemporary expenditures; thus, though we estimate λs which are households' marginal utilities of expenditure for a particular cardinal utility function, one should think of these estimates as being an Index of the consumer's 'true' Marginal Utility of Expenditures, or what we'll call an "IMUE".

Flexibly estimating the IMUE sequentially has great value in part because the marginal utility of expenditures is a central object in models of risk and dynamics in both low- and high-income countries. One important connection is that when preferences are von Neumann-Morgenstern the elasticity of the MUE with respect to total expenditures can be interpreted as (minus) the household's relative risk aversion. A large number of recent papers featuring data from low-income countries assume that a household's marginal utility of expenditures can be modeled as the household's total real household consumption raised to a common negative exponent; examples include Kinnan (2017) and Karaivanov and Townsend (2014). Papers by Chiappori et al. (2014) and Laczó (2015) relax this by allowing different households to have different exponents. But this still involves assuming that utilities are homothetic, and requires the marginal utility of expenditures to depend on a single parameter which *also* governs the elasticity of intertemporal substitution. Non-parametric approaches such as that of Mazzocco and Saini (2011) are much less restrictive, but at the cost of not allowing for actual measurement of the IMUE.

2. Summary of empirical contribution

Our development of the CFE demand system along with methods for flexibly estimating both demand elasticities and IMUEs provides a useful new toolkit, and our use of this toolkit to understand Ugandan consumption expenditures over 2005–2012 yields four main novel empirical insights, with implications for policy. First, we contribute to literature on the estimation of globally regular demand systems (Cooper and McLaren 1996; Lewbel and Pendakur 2009) by demonstrating that the rank of consumer demand in Uganda is at least four. This is among the highest-rank globally-regular demand systems ever estimated

(Lewbel (2003) estimates a rank 4 system, and the system of Lewbel and Pendakur (2009) is "more than three"). A consequence is that one cannot use a single price index to adjust total consumption expenditures—in the Ugandan case, at least four such indices are necessary to measure changes in welfare due to changes in prices. Our measures of λ make these adjustments automatically, allowing us to cleanly relate our estimated IMUEs to traditional expenditure-based measures of headcount poverty. Second, since the rank of the demand system is greater than one, calculations of head-count poverty which rely on adjusting total consumption expenditures using a single price index must be incorrect. The World Bank and Uganda have produced estimates of poverty and inequality for Uganda that involve such incorrect calculations. We use our methods to correct these, and find that while the single-index approach yielded estimates of poverty rates that *fell* over the course of the 2008 food price crisis and subsequent 2010 recession, our estimates correctly capture the fact that quantities of most food consumed fell over this period for most households while nominal food expenditures rose, and thus provide a starkly different picture of changes in both the level and distribution of welfare in Uganda over this period. Third, our estimates of IMUE imply not only that many households in Uganda are poorer than previous estimates indicated, but also allow us to measure households' relative risk aversions. The key for identification of these up to unknown location and scale parameters is simply the requirement that preferences be von Neumann-Morgenstern (these two parameters aren't identified because we must allow for monotonic transformations of the utility function). We find strong evidence of heterogeneity in these relative risk aversions across households, and also find that relative risk aversion is decreasing in the expenditures we observe. This is a much stronger finding than the often asserted claim that *absolute* risk aversion is decreasing in expenditures, and implies that in fact absolute risk aversion is decreasing at a more than linear rate. Accordingly, the welfare consequences of risk for poorer households are greatly understated by the usual homothetic constant relative risk averse (CRRA) preferences typically assumed in the literature. Finally, one practical implication of our work is that it is possible to construct theory-consistent measures of household welfare much less expensively than is now the standard practice. That standard practice (Deaton 1997) is motivated by thinking about the arguments of the indirect utility function for a consumer with homothetic preferences. Since demands for such a consumer are rank one, it follows that welfare can be measured using data on total expenditures and a single price index. This is only two variables, but in practice measuring total consumption expenditures usually involves eliciting information on *all* expenditures—in the Ugandan case enumerators elicit information on 110 different items. In addition the data necessary to construct a price index (usually a Laspeyres index) is required. In our application we show that almost all the variation in observed IMUES can be captured by using data on a much smaller set of expenditures. In principle the number of distinct items need be no greater than the rank of the demand system (at least four); in practice a reduction to only twelve goods performs reasonably well.

The rest of this paper is organized as follows. We begin with a model of household behavior, and use this to constructively derive the globally-regular Constant Frisch Elasticity demand system. We then describe a method of estimating this demand system. We discuss the Ugandan data used for our application, and then present results on estimated demand elasticities and demand system rank; on distributions of $\log \lambda$, with implications for the measurement of poverty. We indulge in two distinct validation exercises, first, showing that our estimated demand system not only can predict estimated log expenditures well, but also both aggregate and individual expenditure shares. We next introduce a diagnostic measure to see whether non-randomly missing data might cause problems for our estimation, and conclude that it does not. We next present results related to aggregation across goods and the selection of a small number goods for measuring IMUEs, and then finally provide estimates of the distribution of relative risk aversions.

3. Model of Household Behavior

In this section we provide a simple model of household demand behavior, and use this model to derive a set of λ-constant (Heckman and MaCurdy 1980) or "Frischian" demands (Browning, Deaton, and Irish 1985). The class of demands we estimate are dual to Marshallian demands that do not generally have an explicit representation, but when separable can be regarded as an instance of the non-homothetic implicitly-additive Marshallian demands studied by Hanoch (1975) and recently exploited empirically by Comin, Lashkari, and Mestieri (2015). The key feature of our demand system is that it allows income elasticities to vary not only across goods, but also to vary with wealth and with prices in a manner which is both flexible and guaranteed to be theory consistent. The underlying preference structure is related to the "direct-addilog" utility first described by Houthakker (1960) or the "CRIE" preferences of Caron, Fally, and Markusen (2014), but allows for flexible substitution patterns across goods.

3.1. **The household's one-period consumer problem.** To fix concepts, suppose that in a particular period t a household with some vector of characteristics z_t faces a vector of prices p_t, and has budgeted a quantity of the numeraire good x_t to spend on contemporaneous consumption $c_t \in X \subseteq \mathbb{R}^n$.[4] The household's preferences over different consumption bundles are summarized by a utility function $U \in \mathcal{U}$, with $\mathcal{U}$ the set of utility functions mapping X into $\mathbb{R}$ which are increasing, concave, and continuously twice differentiable.

4. Note that if the vector z_t includes decisions the household has made about time use and the vector of prices p_t includes prevailing wage rates then this formulation of the problem can accommodate decision problems involving non-separable leisure.

Thus, the household solves the standard consumer's problem

$$V(x_t, p_t; z_t) = \max_{c \in X} U(c; z_t) \qquad \text{such that } \sum_{i=1}^{n} p_{it} c_i \leq x_t. \tag{1}$$

An interior solution to this problem is characterized by a set of n first order conditions for consumption goods which take the form

$$u_i(c; z_t) = \lambda_t^* p_{it}, \tag{2}$$

where u_i denotes the partial derivative of the momentary utility function U with respect to the ith good. In addition there's the budget constraint, with which the Lagrange multiplier λ_t^* is associated.

A solution to (1) takes the form of a set of demands $c(x_t, p_t; z_t)$. A classical identification result tells us that knowledge of these demands allows us to identify the utility function of a household with characteristics z_t up to a monotonic transformation $M : \mathbb{R} \to \mathbb{R}$, determining an equivalence class of utility functions $\mathcal{M}(U) = \{M(U) | c(x, p; z) = \arg\max_{c \in X} M(U(c; z)) \text{such that } p'c \leq x\}$. If $M(U) \in \mathcal{U}$ for all $U \in \mathcal{U}$ then we say that M is a *regular* transformation; in particular it preserves differentiability and concavity.

The following gives an analogous identification result for the MUE for regular transformations:

Proposition 1. *Let $M : \mathbb{R} \to \mathbb{R}$ be a regular transformation. Then*

(1) *A household with utility $U \in \mathcal{U}$ which chooses c to solve (1) has a marginal utility of expenditures given by a function $\lambda^*(x, p; z) = \partial V / \partial x$. For any z this function is continuously differentiable in (x, p), strictly decreasing in x, and strictly increasing in p.*

(2) *A household which has a utility function $M(U)$ which chooses c to solve (1) has a marginal utility of expenditures given by a function $\lambda^M(x, p; z) = \lambda^*(x, p; z) M'(V(x, p; z))$; and*

(3) *The function $\lambda^M(x, p; z)$ varies one-to-one with $\lambda^*(x, p; z)$, and shares its properties of continuous differentiability, decreasing in x, and increasing in p.*

Proof. A solution to (1) is a set of demands $c(x, p; z)$. The assumed differentiability of U implies that $\lambda^* = \partial V / \partial x$ exists and is a continuously differentiable function of (x, p). The concavity of U along with standard comparative statics arguments imply that λ^* is decreasing in x and increasing in p, establishing (1). For (2), the household's indirect utility function is $M(V(x, p; z))$; since M is regular we can obtain the marginal utility of x by the chain rule. For (3), observe that for any z both λ^* and $M'(V(x, p; z))$ are positive but monotonically decreasing in x and montonically increasing in p; accordingly the product of the two is similarly monotone, establishing a one-to-one relationship between λ^* and λ^M. □

If we observe demands $c(x, p; z)$ we can identify the set of utility functions that would generate those demands. One upshot of Proposition 1 is that we can also identify the set of marginal utilities of expenditures $\mathcal{L}(V) = \{\lambda M'(V) | \lambda = \partial V/\partial x; M \text{ regular.}\}$ consistent with those demands.

If we can use data on demands c to obtain measures of some utility and MUE (U, λ) which are consistent with these demands, there's a strong sense in which it doesn't matter *which* (U, λ) we obtain, since with any such pair we can (i) characterize the entire set; and (ii) any U and corresponding λ consistent with observed demands has a one-to-one relationship with every other demand-consistent utility function and MUE.

3.2. **The household's intertemporal problem.** Since we are ultimately interested in the welfare of households in a stochastic, dynamic environment, we relate the solution of the static one-period consumer's problem above to a multi-period stochastic problem; at the same time we introduce a simple form of (linear) production (this could be easily generalized).

We assume the household has time-separable von Neumann-Morgenstern preferences, and that it weights future utility using a discount factor β_t (allowed to vary across periods). The resulting additive separability across dates and states means that can treat the household's global problem using a two-stage budgeting approach (W. M. Gorman 1959). As above, within a period t, a household is assumed to allocate funds for total expenditures in that period obtaining a total momentary utility described by the Marshallian indirect utility function $V(x_t, p_t; z_t)$, where x_t are time t expenditures, p_t are time t prices, and z_t are time t characteristics of the household. Note that the indirect utility function V inherits the cardinality of the utility function U; this is the household's "true" indirect utility function.

The household brings a portfolio of assets with total value $R_t b_t$ into the period, and realizes a stochastic income y_t. Given these, the household decides on investments b_{t+1} for the next period, leaving x_t for consumption expenditures during period t. More precisely, the household solves

$$\max_{\{b_{t+1+j}\}_{j=0}^{T-t}} \mathrm{E}_t \sum_{j=0}^{T-t} \beta_{t+j} V(x_{t+j}, p_{t+j}; z_{t+j})$$

subject to the intertemporal budget constraints

$$x_{t+j} = R_{t+j} b_{t+j} + y_{t+j} - b_{t+1+j}$$

for all $j = 0, \ldots, T - t - 1$, and taking the initial b_t as given.

The solution to the household's problem of allocating expenditures across time will satisfy the Euler equation

$$\frac{\partial V}{\partial x}(x_t, p_t; z_t) = \frac{\beta_{t+1}}{\beta_t} \mathrm{E}_t R_{t+1} \frac{\partial V}{\partial x}(x_{t+1}, p_{t+1}; z_{t+1}).$$

But by definition, these partial derivatives of the indirect utility function are equal to the functions λ^* evaluated at the appropriate prices and expenditures, so that we have

$$\lambda^*(p_t, x_t; z_t) = \frac{\beta_{t+1}}{\beta_t} \mathrm{E}_t R_{t+1} \lambda^*(p_{t+1}, x_{t+1}; z_{t+1}). \tag{3}$$

This expression tells us, in effect, that the household's marginal utility or marginal utility of expenditures λ_t^* satisfies a sort of martingale restriction, so that the current value of λ_t^* play a central role in predicting *future* values λ_{t+j}^*.

If we know the Frisch demand functions for a consumer with utility function U and observe prices and quantities demanded for some of these goods, then we can invert the demand relationship to obtain the consumer's λ_t^*.

3.3. **Differentiable Demand.** We now turn our attention to the practical problem of specifying a demand relation that can be estimated using the kinds of data we have available on disaggregated expenditures. Attfield and Browning (1985) take a so-called "differentiable demand" approach to a related problem; their method yields Frischian (aggregate) demands without requiring separability. These demands will, in general, depend on all prices, yet one need only estimate demand equations for a select set of goods.

Our analysis here resembles that of Attfield and Browning (1985) in outline, but where they arrive at a Rotterdam-like demand system in quantities, we obtain something importantly different in expenditures. This overcomes an important shortcoming of Attfield and Browning's demand system, which is that it is integrable only in the homothetic case.

It's easiest here to work with the consumer's profit function (W. Gorman 1976; Browning, Deaton, and Irish 1985). In this setting we can treat leisure as simply another consumption good and the wage as simply the price of this good, so that we have

$$\pi(p, r, z) = \max_{c \in X} rU(c; z) - p^\top c.$$

Here r has the interpretation of being the "price" of utility, while p the prices of consumption (and leisure). Let subscripts to the π function denote partial derivatives. Some immediate properties of importance: the price r is equal to the quantity $1/\lambda^*$ from our earlier analysis; the profit function is linearly homogeneous in p and r; by the envelope theorem $\pi_i(p, r, z) = -c_i$ for all $i = 1, \ldots, n$ and for any z; and (since we want to work with expenditures) $-p_i \pi_i = x_i$.

Using this last fact and taking the total derivative yields

$$dx_i = -\pi_i dp_i - p_i \sum_{j=1}^{n} \pi_{ij} dp_j - p_i \pi_{ir} dr - p_i \sum_{l=1}^{\ell} \pi_{iz_l} dz_l.$$

Since $d \log x = dx/x$ for $x > 0$, this can be written as

$$x_i d \log x_i = -\pi_i p_i d \log p_i - p_i \sum_{j=1}^{n} \pi_{ij} p_j d \log p_j - p_i \pi_{ir} r d \log -p_i \sum_{l=1}^{\ell} \pi_{iz_l} z_l.$$

Recalling that $-\pi_i p_i = x_i$

$$(4) \qquad d \log x_i = d \log p_i + \sum_{j=1}^{n} \frac{\pi_{ij}}{\pi_i} p_j d \log p_j + \frac{\pi_{ir}}{\pi_i} r d \log r + \sum_{l=1}^{\ell} \frac{\pi_{iz_l}}{\pi_i} z_l d \log z_l.$$

Now, let $\theta_{ij} = -\frac{\pi_{ij}}{\pi_i} p_j$ denote the (cross-) price elasticities of demand holding r constant (Frisch 1959, called these "want elasticities"); let δ_{il} denote the elasticity of demand for good i with respect to changes in the characteristic z_l; and let $\beta_i = \frac{\pi_{ir}}{\pi_i} r$ denote the elasticity of demand with respect to r. Using the fact that $1/r = \lambda^*$ we can rewrite this as

$$(5) \qquad d \log x_i = d \log p_i - \sum_{j=1}^{n} \theta_{ij} d \log p_j - \sum_{l=1}^{\ell} \delta_{il} d \log z_l - \beta_i d \log \lambda^*.$$

Using the linear homogeneity of the profit function, it follows that $\beta_i = \sum_{j=1}^{n} \theta_{ij}$.

Equation (5) gives us an exact description of how expenditures will change in response to infinitesimal changes in prices for a consumer with the utility function U and characteristics z.

Now we make an assumption which is important for reasons both involving principle and practice: that the elasticities $\Theta = [\theta_{ij}]$ (and so $\beta = [\beta_i]$) and $\delta = [\delta_{il}]$ are constant, and not functions of prices (p, r) or characteristics z. With this assumption, the matrix of parameters δ summarizes the effects of the consumer's characteristics z_t on demand; conditional on these characteristics, the term involving λ^* indicates the rate at which changes in welfare influence changes in expenditures on particular goods. Because the β_i are equal to the row sums of the matrix of elasticities $\Theta = (\theta_{ij})$, in this case the Θ matrix summarizes all the pertinent information for understanding changes in demand (conditional on changes in z); we call Θ the matrix of "Frisch elasticities," and refer to the result as the Constant Frisch Elasticity (CFE) demand system.

While our chosen parameterization is important, it is not arbitrary. With Θ and δ taken to be parameters (5) is separable in prices, characteristics, and $\log \lambda$, and some such separability is critical to identification. The alternative parameterization of Attfield and Browning (1985) is also separable, but is not integrable, and hence is unsuitable for measuring welfare. It is not clear whether there are any other separable parameterizations which are integrable.

Integration in our case is simple and direct; with Θ and δ constant we can integrate (5) to obtain an exact expression for the *level* of demand and expenditures. In particular, let α be an n-vector of constant parameters, which arise as constants of integration from (5). Then the demand for good i is given by

$$c_i = \alpha_i \exp(\delta_i^\top \log z) \left[\lambda^{\beta_i} \prod_{j=1}^{n} p_j^{\theta_{ij}} \right]^{-1}. \tag{6}$$

Note that when Θ is diagonal and all the β_i are equal then $\lambda = 1/x^{1/\beta}$, and we obtain constant elasticity of substitution (CES) demands as a special case. It's also possible to show that any PIGL demand system Muellbauer (1976), including the AIDS of Deaton and Muellbauer (1980), can be obtained when the different β_i take values of either one or β and observed quantities are linear combinations of the different c_i described by (6).

Thus it is that the CFE demand system nests most of the demand systems used in empirical work. But most of these demand systems are of low *rank* (Lewbel 1991), with CES or any other homothetic system having a rank of one, and AIDS and PIGL a rank of no more than two. For the CFE system the rank will be equal to the number of distinct values of β_i; thus the system may have a rank as large as n. To see this, note that the marginal utility of expenditures λ can be regarded as a function of total expenditures x and prices p. Then the budget constraint can be written in the form

$$\sum_{i=1}^{n} a_i(p)\lambda^{-\beta_i} = x,$$

with the function $\lambda(p, x)$ the solution to this equation. Using the same notation, expenditures for good i are $x_i(p, x) = a_i(p)\lambda(p, x)^{-\beta_i}$. Expressed in matrix form, the right hand side of this equation takes the form $\mathbf{a}(p)\mathbf{g}(p, x; z)$, with $\mathbf{g}(p, x; z)$ a diagonal matrix with rank equal to the number of distinct values of β_i. Using this result, it follows that the Engel curves of this demand system are a flexible functional form (Diewert 1971), with symmetry and homogeneity which can be tested or imposed by way of linear restrictions on the matrix Θ. Further, when these restrictions are satisfied the demand system is globally regular, and implies a simple parametric form for the direct utility function.

4. Estimation with (Possibly Repeated) Cross-Sectional Data

Suppose we have data on disaggregate expenditures for T cross-sections of households facing the same prices. We want to use these data to estimate the parameters of (6). However, those equations describe only the demand system for a single household. Adapting it, let $j = 1, \ldots, N$ index different households, and assume that household characteristics for the jth household at time t include both observable characteristics z_t^j and time-varying unobservable characteristics ϵ_{it}^j. Then we can write our structural estimating equation as

$$\log x_{it}^j = \log \alpha_i + \left(\log p_{it} - \sum_{k=1}^{n} \theta_{ik} \log p_{kt} \right) + \delta_i^\top \log z_t^j - \beta_i \log \lambda_t^j + \epsilon_{it}^j. \tag{7}$$

We assume that prices are unknown to the econometrician, but that all households face the same prices.[5] Expressed in a reduced form, we write

$$y_{it}^j = a_{it} + d_i^\top(\log z_t^j - \overline{\log z}_t) + b_i w_t^j + e_{it}^j, \tag{8}$$

where

$$\begin{aligned} y_{it}^j &= \log x_{it}^j \\ a_{it} &= \log \alpha_i + d_i^\top \overline{\log z}_t + \left[\log p_{it} - \sum_{k=1}^n \theta_{ij} \log p_{kt}\right] - \beta_i \overline{\log \lambda}_t + \beta_i \bar{\epsilon}_{it} \\ d_i &= \delta_i \\ e_{it}^j &= \epsilon_{it}^j - \bar{\epsilon}_{it} \\ b_i w_t^j &= -\beta_i(\log \lambda_t^j - \overline{\log \lambda}_t). \end{aligned}$$

Since we have one equation for each good that we observe expenditures for, (8) forms a seemingly-unrelated regression (SUR) system. In a first step we obtain the reduced form parameters (a_{it}, d_i) simply by using ordinary least squares to estimate (8), treating the a_{it} as a set of good-time effects, and then using the methods of Arellano (1987) to obtain robust estimates of the standard errors of the estimated parameters a_{it} and d_i.

4.1. **Identification of the Parameters of Interest.** The parameter estimates $d_i = \delta_i$ and the good-time effects have some independent interest, but the main objects we're interested in estimating are the demand elasticities β_i and the quantities $\log \lambda_t^j$, which are embedded in the residuals in the reduced form equations for log expenditures (8). Yet variation in residuals from these expenditure equations alone is enough to identify these up to single normalization.

4.1.1. *Frisch elasticities β_i and $\log \lambda$.* The residuals from (8) are equal to $b_i w_t^j + e_{it}^j$. The first term of this sum is what we're interested in. Arrange the estimated residuals as an $n \times NT$ matrix $\mathbf{Y}$. The first term in the equation captures the role that variation in marginal utility λ plays in explaining variation in expenditures. Because it's equal to the outer product of two vectors, this first term is at most of rank one. The second term captures the further role that other unobservables (e.g., unobservable household characteristics, measurement error) play in changes in demand; if there are m such unobservable factors, then this second term is of at most rank $\bar{m} = \min(m, n-1)$.

We proceed by considering the singular value decomposition (SVD) of $\mathbf{Y} = \mathbf{U}\Sigma\mathbf{V}^\top$, where $\mathbf{U}$ and $\mathbf{V}$ are unitary matrices, and where Σ is a diagonal matrix of the singular values of

5. This can be easily extended; for example, in our application below we can allow for different prices in different regions.

$\mathbf{Y}$, ordered from the largest to the smallest. Then the rank one matrix that depends on λ is $\mathbf{b}\mathbf{w}^\top = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^\top$, while the second matrix (of at most rank $\bar{m}$) is $\mathbf{d}\mathbf{Z}^\top = \sum_{k=2}^{\bar{m}} \sigma_k \mathbf{u}_k \mathbf{v}_k^\top$, where σ_k denotes the kth singular value of $\mathbf{Y}$, and where the subscripts on $\mathbf{u}$ and $\mathbf{v}$ indicate the column of the corresponding matrices $\mathbf{U}$ and $\mathbf{V}$. The sum of these matrices is equal to $\mathbf{Y}$, and the truncated sum of the first $k \leq \bar{m}$ matrices is the optimal k rank approximation to $\mathbf{Y}$, in the sense that by the Eckart-Young theorem this is the solution to the problem of minimizing the Frobenius distance between $\mathbf{Y}$ and the approximation. Accordingly, this is also the least-squares solution (Golub and Reinsch 1970).

The singular value decomposition thus identifies the structural parameters β_i and changes in log marginal utility up to an unknown scalar ϕ, so that we obtain estimates of $\phi\beta_i$ and of $\log \lambda_t^j/\phi$, with estimates of $\overline{\log \lambda_t})/\phi$ identified from changes in the reduced-form terms a_{it} over time.

4.1.2. *Missing Data.* Our demand system permits estimation of arbitrarily large and extremely disaggregate demand systems. In practice this raises the possibility that many different item expenditures may be "zeros" or missing for a given household. Thus, our SVD must somehow contend with missing data; the algorithm we've developed for doing this is described in Appendix A. But in addition to the issue of calculation, missing data raises the possibility of issues related to selection, biasing our estimates of the demand system (Vella 1998, gives a survey). To address this, we adapt methods developed in the psychometric literature to show that our estimator of Frisch elasticities and λs is unbiased under general conditions even with such missing data, and provide a simple test of these conditions.

In our application, no household in our sample reports positive expenditures for every good, and overall fewer than 40% of all possible item reported expenditures are positive. Where recorded values of consumption expenditure are equal to zero, we regard these as missing and drop them from the analysis. There are two reasons for this treatment of zeros: first, at an entirely practical level, our dependent variable is the logarithm of expenditures, which is undefined at zero. But second, if a household is at a corner when it chooses a particular consumption item, then the first order condition in (2) for that consumption good won't be correct (we'd be missing a multiplier related to non-negativity). By simply dropping observations for goods where consumption is zero we are effectively dropping observations where expenditures do not correctly reveal the index $\log \lambda$.

Our resolution is described in detail in the appendix, and is similar to the way that fixed effects estimation can address the problem of selection in unbalanced panels of households over time (e.g., Wooldridge 2002, Section 17.7.1). A simple practical test for selection bias can be constructed by using an approach related to one advocated by Wooldridge, in which we first regress an indicator for whether expenditures x_{it}^j are positive on the same right-hand side variables as appear in (8), obtaining residuals $\hat{r}^j_{it} = \mathbb{1}(x^j_{it} > 0) - \hat{a}^r_{it} - \hat{\delta}^{r\top}_i z^j_t$. We then use these residuals to augment the matrix of residuals from the first step estimation of (8), and use a singular value decomposition of the resulting $n \times 2NT$ matrix to obtain alternative estimates of elasticities $\hat{\beta}^r$ and the log λ^{jr}_t. If selection is an issue then these estimated values should differ from the values obtained from the decomposition of the unaugmented residual matrix.

5. Data

To illustrate some of the methods and issues discussed above, we use data from four rounds of surveys conducted in Uganda (in 2005–06, 2009–10, 2010–11, and 2011–12).[6] We first give a descriptive account of some of the data on household characteristics and expenditures from these surveys.

Table 1. Characteristics of households in Uganda. Figures in parentheses are standard deviations.

	N	Boys	Girls	Men	Women	Rural
2005	3115	1.48	1.48	1.12	1.24	0.72
		(1.45)	(1.44)	(0.89)	(0.86)	(0.45)
2009	2927	1.70	1.67	1.21	1.33	0.74
		(1.55)	(1.50)	(0.97)	(0.89)	(0.44)
2010	2639	1.77	1.78	1.26	1.40	0.78
		(1.57)	(1.56)	(1.01)	(0.95)	(0.41)
2011	2795	1.70	1.72	1.23	1.37	0.80
		(1.53)	(1.53)	(0.97)	(0.86)	(0.40)

5.1. **Summary Statistics.**

Table 1 gives some information on household characteristics. In each of four rounds, there are about 3000 households; of these, between 70–80% are rural. There is a panel aspect to these data. There are a total of 3727 distinct households observed across the four rounds; of these 2151 are observed in every round.

The average household size consists of 5.8 people; the average rural household is larger, at 5.9, while the average urban household consists of 5.5 people.[7]

6. These datasets are available at `http://go.worldbank.org/MO5MSKCQS0`, with documentation available at `http://go.worldbank.org/S233P3YC30`.

7. For our purposes a person is a household member if they've lived in the household for at least one month of the previous twelve. People identified as 'guests' who satisfy these criteria must also have spent the night prior to the interview.

5.2. **Expenditure Data.** Excluding durables, taxes, fees & transfers, there are 110 categories of expenditure in the data, of which 72 are different food items or categories, and 38 are other nondurables or services.

The Ugandan surveys we use elicit information on food consumption over the last seven days, with consumption quantities and values reported as being "out of purchases," "out of home produce," or "received in-kind/free;" consumption out of purchases "away from home" is also elicited for a selection of food items. Where consumption is "out of home produce" or "received in-kind/free" values as well as quantities are elicited. This recall period and approach to eliciting the source of consumption is typical of household consumption and expenditure surveys (Fiedler et al. 2012), and is designed to distinguish between the acquisition of stocks of food and consumption.

Appendix Table B.2 paints a picture of aggregate expenditure shares across these categories, listing mean and aggregate expenditure shares for all foods, ordered by the size of their aggregate expenditure share in 2005. Shares of aggregates are fairly stable across the period 2005–2011, with only a handful of goods changing their aggregate shares by as much as one percentage point (the only exceptions are cassava, sugar, and "other foods."). It should be noted, however, that stability of shares over time is not a prediction of theory, as it would be in a homothetic demand system—changes in incomes or relative prices can be expected to cause changes in shares.

However, while mean and aggregate shares are often stable over time, they differ dramatically for different goods. This does not seem consistent with a model in which consumers have homothetic utility. Such a model would predict equal aggregate and mean expenditure shares.

This general point is graphically borne out in Figure 1. For this figure we construct a statistic ρ_{it} which is the logarithm of aggregate shares minus the logarithm of mean shares, or, for good i at time t,

$$\rho_{it} = \log\left(\frac{\sum_{j=1}^N x_{it}^j}{\sum_{j=1}^N \sum_{k=1}^n x_{kt}^j}\right) - \log\left(\sum_{j=1}^N \frac{x_{it}^j}{\sum_{k=1}^n x_{kt}^j}\right).$$

We then produce a scatterplot of this statistic, ordered by the size of the statistic in 2005. Thus, each good (labelled on the left axis) has associated with it a statistic for each of four years, each with (overlapping) confidence intervals.

With homothetic preferences, this statistic must always be equal to zero, but we can reject this equality for most of the 41 goods in the figure. Instead, a positive value of the statistic identifies goods which play an outsized role in the consumption portfolios of wealthier (i.e., higher expenditure) households, and include passion fruit, bread, chicken, soda, and sweet bananas, among others. Conversely, when the statistic is negative we identify goods that are

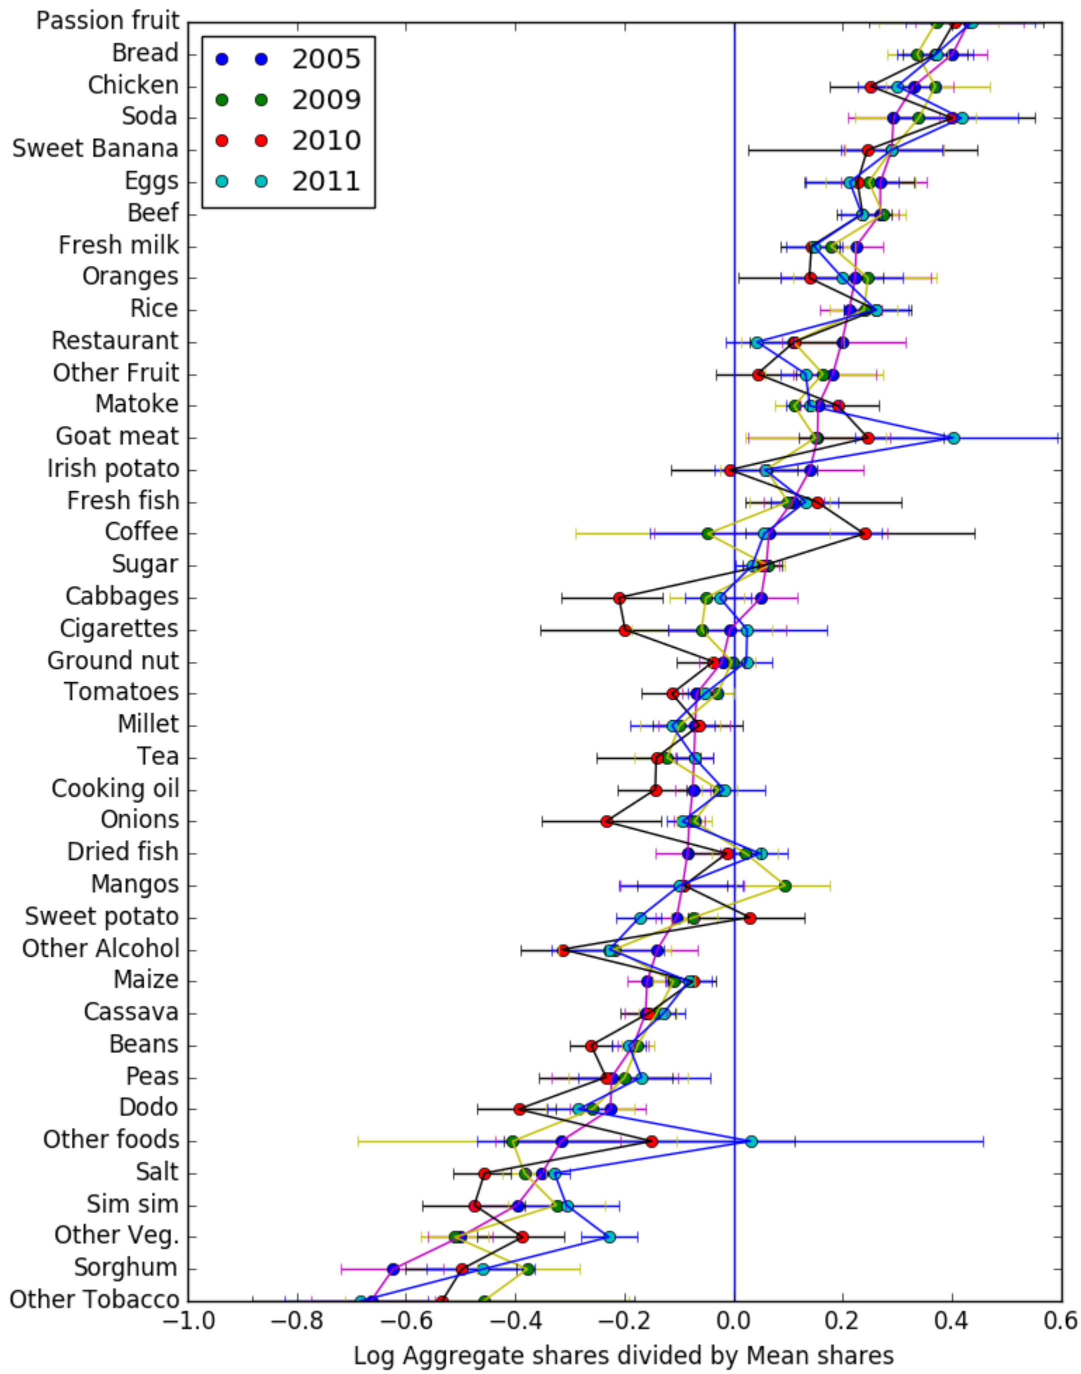

Figure 1. Log of mean shares minus log of aggregate shares for different years (ordered by ranking in 2005), with 95% confidence intervals.

particularly important in the portfolios of households with lower food expenditure. Here we see "other tobacco", sorghum, "other vegetables", sim sim (sesame) and salt.

The figure also argues against the usual *quasi*-homothetic specification of preferences, interpreted as though some positive "subsistence" level of the good is necessary for survival. Subsistence requirements of this case could account for the goods for which negative statistics are observed, such as salt, cassava, or maize, but as total expenditures increase, quasi-homothetic utility implies that budget shares should converge to a fixed constant. This implies that in Figure 1 the plotted statistics should converge to zero as one moves from bottom to top. This predicted pattern is not at all evident.

6. Results

6.1. **Estimates of Demand Elasticities.** We now turn to estimates of some of the parameters of the the demand system (7), estimated using the four rounds of data from Uganda discussed above. Table 2 presents results from our baseline specification. In this specification we obtain results for a system of 44 minimally aggregated consumption goods, assuming that all households face the same relative prices. We take as observable characteristics the number of men, women, boys and girls in each household, as well as the logarithm of total household size. We also include a dummy indicating a rural or urban location for the household. One interpretation of this dummy is that allows for differences in the general price level and average $\log \lambda$ between rural and urban areas, but it also can accommodate differences in the marginal utility functions across rural and urban areas (i.e., heterogeneity in the α_i parameters).

Table 2: Estimates of expenditure system assuming a single market. Controls include the numbers of boys, girls, men, and women in household, along with the log of household size. Figures in parentheses are estimated standard errors.

	β_i	$\log \alpha_i$	Rural	Boys	Girls	Men	Women	log Hsize
Passion Fruits	0.77^{***}	6.44^{***}	-0.37^{***}	-0.01	-0.00	0.03	0.09^{***}	0.22^{**}
	(0.06)	(0.19)	(0.06)	(0.03)	(0.03)	(0.03)	(0.03)	(0.10)
Oranges	0.73^{***}	5.78^{***}	-0.18^{***}	0.09^{***}	-0.00	0.11^{***}	0.16^{***}	0.04
	(0.07)	(0.25)	(0.06)	(0.03)	(0.03)	(0.04)	(0.04)	(0.13)
Coffee	0.64^{***}	4.67^{***}	-0.21^{***}	0.05	0.01	0.15^{***}	0.15^{***}	-0.06
	(0.09)	(0.24)	(0.08)	(0.03)	(0.03)	(0.04)	(0.04)	(0.13)
Other Fruits	0.63^{***}	5.89^{***}	0.04	0.07^{***}	0.03	0.10^{***}	0.04	0.30^{***}
	(0.08)	(0.19)	(0.07)	(0.03)	(0.03)	(0.03)	(0.03)	(0.11)

Continued on next page

Continued from previous page

	β_i	$\log \alpha_i$	Rural	Boys	Girls	Men	Women	log Hsize
Mangos	0.62***	5.89***	0.10*	0.00	−0.02	0.08**	−0.03	0.48***
	(0.06)	(0.23)	(0.06)	(0.03)	(0.03)	(0.03)	(0.03)	(0.13)
Sweet Bananas	0.61***	5.95***	−0.28***	−0.00	0.02	0.12***	0.06**	0.28***
	(0.05)	(0.17)	(0.05)	(0.02)	(0.03)	(0.03)	(0.03)	(0.09)
Ground nuts (shelled)	0.61***	6.02***	−0.01	0.08***	0.09***	0.12***	0.17***	−0.03
	(0.06)	(0.17)	(0.06)	(0.02)	(0.02)	(0.03)	(0.03)	(0.10)
Bread	0.60***	6.61***	−0.45***	0.01	0.03*	0.11***	0.09***	0.18**
	(0.04)	(0.17)	(0.04)	(0.02)	(0.02)	(0.02)	(0.02)	(0.07)
Soda	0.59***	7.13***	−0.29***	0.00	0.03	0.07**	0.08***	0.06
	(0.06)	(0.18)	(0.05)	(0.03)	(0.03)	(0.03)	(0.03)	(0.09)
Maize (cobs)	0.58***	5.83***	0.44***	0.03	0.07**	0.08**	−0.03	0.43***
	(0.07)	(0.22)	(0.07)	(0.03)	(0.03)	(0.03)	(0.04)	(0.11)
Fresh Milk	0.56***	7.00***	−0.30***	0.00	0.00	0.14***	0.05***	0.26***
	(0.04)	(0.16)	(0.03)	(0.01)	(0.01)	(0.02)	(0.02)	(0.06)
Eggs	0.50***	6.23***	−0.28***	−0.01	−0.02	0.09***	0.04*	0.22***
	(0.04)	(0.16)	(0.04)	(0.02)	(0.02)	(0.02)	(0.03)	(0.08)
Cooking oil	0.50***	6.18***	−0.44***	−0.01	0.01	0.09***	0.03**	0.21***
	(0.03)	(0.12)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.04)
Goat Meat	0.50***	7.30***	−0.18***	0.02	0.02	0.09**	0.07*	0.28**
	(0.07)	(0.18)	(0.06)	(0.03)	(0.03)	(0.03)	(0.04)	(0.12)
Tomatoes	0.49***	6.04***	−0.42***	−0.00	0.02**	0.08***	0.06***	0.13***
	(0.02)	(0.10)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.04)
Rice	0.48***	6.73***	−0.18***	0.04**	0.03**	0.08***	0.05***	0.34***
	(0.03)	(0.12)	(0.03)	(0.01)	(0.01)	(0.02)	(0.02)	(0.06)
Beans (fresh)	0.47***	6.29***	0.11**	0.03*	0.03	0.09***	0.04	0.32***
	(0.05)	(0.20)	(0.05)	(0.02)	(0.02)	(0.02)	(0.03)	(0.09)
Sugar	0.47***	6.80***	−0.38***	0.02*	0.03***	0.08***	0.08***	0.29***
	(0.02)	(0.10)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.04)
Irish Potatoes	0.47***	6.67***	0.21***	0.09***	0.05**	0.07***	0.05*	0.14
	(0.05)	(0.18)	(0.05)	(0.02)	(0.03)	(0.03)	(0.03)	(0.09)
Beer	0.46***	8.82***	−0.46***	0.03	0.09**	0.20***	0.08	−0.34**
	(0.09)	(0.26)	(0.09)	(0.04)	(0.04)	(0.05)	(0.05)	(0.16)
Dodo	0.46***	5.67***	−0.10***	0.06***	0.04***	0.05***	0.07***	0.08

Continued on next page

Continued from previous page

	β_i	$\log \alpha_i$	Rural	Boys	Girls	Men	Women	log Hsize
	(0.04)	(0.14)	(0.03)	(0.01)	(0.01)	(0.02)	(0.02)	(0.06)
Beef	0.46***	7.54***	−0.18***	0.03***	0.02*	0.11***	0.07***	0.19***
	(0.02)	(0.10)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.05)
Onions	0.45***	5.21***	−0.40***	−0.01*	−0.00	0.08***	0.07***	0.14***
	(0.02)	(0.10)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.03)
Cassava (fresh)	0.45***	6.39***	0.22***	0.06***	0.04***	0.08***	0.01	0.28***
	(0.03)	(0.14)	(0.03)	(0.01)	(0.01)	(0.02)	(0.02)	(0.06)
Ground nuts (pounded)	0.45***	6.26***	−0.02	0.04***	0.02*	0.09***	0.09***	0.09*
	(0.04)	(0.14)	(0.03)	(0.01)	(0.01)	(0.02)	(0.02)	(0.05)
Fresh Fish	0.45***	7.08***	−0.11***	0.04**	0.02	0.12***	0.06**	0.12*
	(0.03)	(0.14)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.07)
Restaurant (food)	0.45***	8.80***	−0.58***	0.05**	0.03	0.19***	0.10***	−0.41***
	(0.05)	(0.21)	(0.04)	(0.02)	(0.02)	(0.03)	(0.03)	(0.07)
Other Alcoholic drinks	0.44***	7.31***	−0.37***	0.04*	0.01	0.17***	0.03	−0.06
	(0.07)	(0.21)	(0.07)	(0.02)	(0.02)	(0.03)	(0.03)	(0.09)
Pork	0.43***	7.25***	−0.31***	0.04*	0.02	0.09***	0.06**	0.16
	(0.05)	(0.18)	(0.06)	(0.03)	(0.02)	(0.03)	(0.03)	(0.10)
Dry/Smoked fish	0.41***	6.53***	−0.17***	0.03*	0.02	0.12***	0.05**	0.19**
	(0.04)	(0.17)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.07)
Cabbages	0.40***	5.92***	−0.14***	0.00	0.02	0.05***	0.03**	0.21***
	(0.03)	(0.12)	(0.03)	(0.01)	(0.01)	(0.02)	(0.02)	(0.06)
Other vegetables	0.40***	5.66***	−0.05	0.04**	0.06***	0.05**	0.07***	0.13*
	(0.04)	(0.14)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.07)
Matoke (heap)	0.39***	7.64***	−0.25***	0.00	−0.00	0.04	0.11***	0.33***
	(0.07)	(0.21)	(0.06)	(0.03)	(0.03)	(0.03)	(0.04)	(0.13)
Maize (flour)	0.39***	6.62***	0.07**	0.08***	0.08***	0.09***	0.05***	0.20***
	(0.03)	(0.13)	(0.03)	(0.01)	(0.01)	(0.02)	(0.02)	(0.05)
Restaurant (soda)	0.37***	7.53***	−0.29***	0.00	0.04	0.08**	0.03	−0.11
	(0.07)	(0.21)	(0.07)	(0.03)	(0.03)	(0.03)	(0.04)	(0.10)
Millet	0.36***	6.11***	0.15***	0.03	0.00	0.05*	0.05*	0.40***
	(0.05)	(0.18)	(0.05)	(0.02)	(0.02)	(0.02)	(0.03)	(0.09)
Chicken	0.34***	8.09***	−0.23***	−0.01	−0.02	0.07***	0.04	0.20**
	(0.04)	(0.13)	(0.04)	(0.02)	(0.02)	(0.02)	(0.02)	(0.09)

Continued on next page

Continued from previous page

	β_i	$\log \alpha_i$	Rural	Boys	Girls	Men	Women	log Hsize
Sweet Potatoes (fresh)	0.34***	6.49***	0.32***	0.07***	0.06***	0.07***	0.06***	0.32***
	(0.03)	(0.13)	(0.03)	(0.01)	(0.01)	(0.02)	(0.02)	(0.05)
Beans (dry)	0.33***	6.50***	−0.01	0.03***	0.04***	0.08***	0.04***	0.29***
	(0.02)	(0.11)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.04)
Sorghum	0.33***	5.99***	0.02	−0.03	−0.07**	−0.08**	−0.09**	0.75***
	(0.07)	(0.20)	(0.07)	(0.03)	(0.03)	(0.03)	(0.04)	(0.12)
Sim sim	0.33***	6.09***	0.02	0.01	−0.01	0.09***	0.06**	0.13
	(0.05)	(0.17)	(0.06)	(0.02)	(0.02)	(0.03)	(0.03)	(0.11)
Tea	0.31***	4.59***	−0.19***	0.02*	0.03***	0.12***	0.09***	0.10***
	(0.02)	(0.09)	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.04)
Salt	0.18***	4.18***	0.04***	0.03***	0.02***	0.05***	0.03***	0.24***
	(0.01)	(0.07)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
Cassava (dry/flour)	0.11**	6.66***	0.14***	0.03*	0.00	0.05**	−0.04*	0.53***
	(0.04)	(0.19)	(0.05)	(0.02)	(0.02)	(0.02)	(0.03)	(0.08)

In its first column Table 2 presents estimates of the Frisch elasticities in descending order, identified by the normalization that makes the standard deviation of the estimated $\log \lambda$ in the first round equal to one. The product of the Frisch elasticities with the households' relative risk aversions yields income elasticities, but these cross-sectional data cannot identify cardinal properties of the utility function such as risk attitudes. However, ratios of these estimated parameters can be interpreted as ratios of income elasticities. Thus, the most elastic goods are passion fruits followed closely by oranges; these have elasticities roughly twice that of millet, four times that of salt, and six times that of dry cassava. These orderings of elasticities seem to accord well with informal descriptions of what goods are more desirable.

All estimated elasticities (including those for unreported goods) are positive; thus, there is no evidence that any of these goods is inferior, with demand increasing as $\log \lambda$ decreases. Standard errors for these elasticities are obtained via a block bootstrap (Horowitz 2003).[8]

If all values of β_i were equal, preferences would be homothetic, and the rank of the demand system would be one. Further, as observed above, the rank r of the demand system for these goods is equal to the number of distinct values of β_i. To determine the rank, we adapt a machine-learning tool advocated by Pelleg, Moore, et al. (2000) involving the solution of a

8. We've also computed standard errors by calculating the inter-quartile range of the bootstrapped estimates, and scaling these up under the hypothesis of normality to provide an estimate of standard errors which is more robust to outliers; both estimators deliver very similar results.

set of k-means problems for k less than the number of goods, and the selection of r using the Bayesian Information Criterion; see Appendix C for details. This procedure leads to the inference that the rank of the demand system is at least four. One way of thinking about the demand system having rank four is that the data are telling us that we would need at least four different price indices to measure how changes in *total* expenditures effect household welfare (Lewbel 1991).

The second column of Table 2 gives estimates of $\log \alpha_i$, where α_i is a multiplicative preference parameter. With homothetic utility, i.e., $\beta_i = \beta$, α_i would be equal to n times the expenditure share of good i, and elasticities would be constant across all goods. In our non-homothetic case expenditure shares depend on the parameters α_i, elasticities β_i, and prices. The parameters $\log \alpha_i$ vary positively with expenditure share, and are set equal to mean log expenditures in our first round of data. Estimated standard errors for these parameters are simply equal to the standard deviation of residuals in 2005 divided by the square root of the number of observed positive expenditures in that year.

The third column of the table reports estimates of the effect of being a rural rather than an urban household. Associated standard errors are clustered by round, as are the standard errors associated with other household characteristics (Arellano 1987). The effect of being 'rural' is negative and significant for all but a few goods, consistent with the fact that total food expenditures are roughly 12% less than in urban areas. A handful of exceptions stand out: maize (cobs and flour), beans (fresh), Irish and sweet potatoes, cassava (fresh), millet, and salt expenditures are all significantly greater in rural areas, other things equal.

The next several columns report indicate how expenditures vary with household size and composition. Here we've included the log of household size, but also a count of the number of boys, girls (both under the age of 18), women, and men in the household. This allows for variation in expenditures to respond to household composition, but in a way which also allows for varying returns to scale. The reported coefficient on the logarithm of household size has the interpretation of an elasticity, while the coefficients on counts of boys, girls, men, and women are semi-elasticities. Adding a man to the household (holding total household size constant) has the largest effect on expenditures for beer, restaurant food, other alcoholic drinks, coffee, and fresh milk. Similarly, adding a women has the largest effects on shelled groundnuts, oranges, and coffee. For most goods the addition of an adult has a larger effect on household expenditures than does the addition of a child: if we take a simple average of semi-elasticities across goods we obtain 0.03 for boys, 0.02 for girls, 0.09 for men, and 0.06 for women. We can further identify particular "adult goods" where the difference in semi-elasticities between adults and children are greatest, such as coffee, soda, onions, eggs, bread, and food consumed in restaurants. But adult-child differences are smaller for staples such as millet, rice, and beans, and are even reversed for starchy staples such as maize, cassava, and sweet potatoes. There are also a handful of goods which seem to be differentially preferred by females: goods for which point estimates of elasticities are greater for women than for men, and for girls than for boys, are passion fruit, soda, "other vegetables", and shelled groundnuts.

6.2. **Estimates of** $\log \lambda$**.** The central aim of this paper is to extract measures of household-level welfare from data on consumption expenditures. Our approach uses information from expenditures themselves to separately identify changes in price levels from improvements in welfare. In contrast, the conventional way to do this is to construct the sum of expenditures on non-durable consumption and services, and then to make comparisons over time by deflating this total by a single price index obtained from some other source. It's well known that this is only justified if preferences are homothetic. Both our evidence (estimated values of β_i differ significantly across goods) and the stubborn empirical fact of Engel's Law rule out homotheticity, so the conventional approach must be invalid in principle. We use the episode of the 2008 food price crisis to show the severe problems the conventional approach can also have in practice.

Figure 2 presents histograms of the estimated $\log \lambda$ for each round of data. The mean and standard deviation of the distribution in the first year are normalized to zero and one, respectively; these normalizations suffice to identify not only the vector of elasticities β but also the distribution of $\log \lambda$ in subsequent years, which can otherwise vary in unrestricted ways, reflecting changes in the distribution of welfare relative to the base year.

So what can we say about changes in welfare in Uganda over this period? First, note that the average expenditure share for food in Uganda over this period exceeded 60%, so that most households were quite sensitive to changes in food prices. Second, note that the food price crisis led to large increases in prices that peaked shortly before the second survey we have, in 2009, but food prices in our 2009 data were still much higher than in 2005.[9] Of particular note is that the median prices of most staple starches (maize, millet, potatoes, sweet potatoes, cassava) had nominal prices twice as high in 2009 as in 2005. Such large price increases weren't confined to staple foods—the average increase (across different kinds of food) in median prices was 96%, and no foods *decreased* in price; broadly speaking food prices roughly doubled from 2005–2009.

that relative prices for most kinds of food increased during this period. Further, the large increase in relative food prices corresponds to sizeable decreases in the *quantities* of food consumed in our Ugandan data—on average across food items we observe a 3% decrease in quantities consumed between 2005 and 2009.

9. See also confirmation from Benson, Mugarura, and Wanda (2008) that food prices in Uganda increased dramatically during the food price crisis.

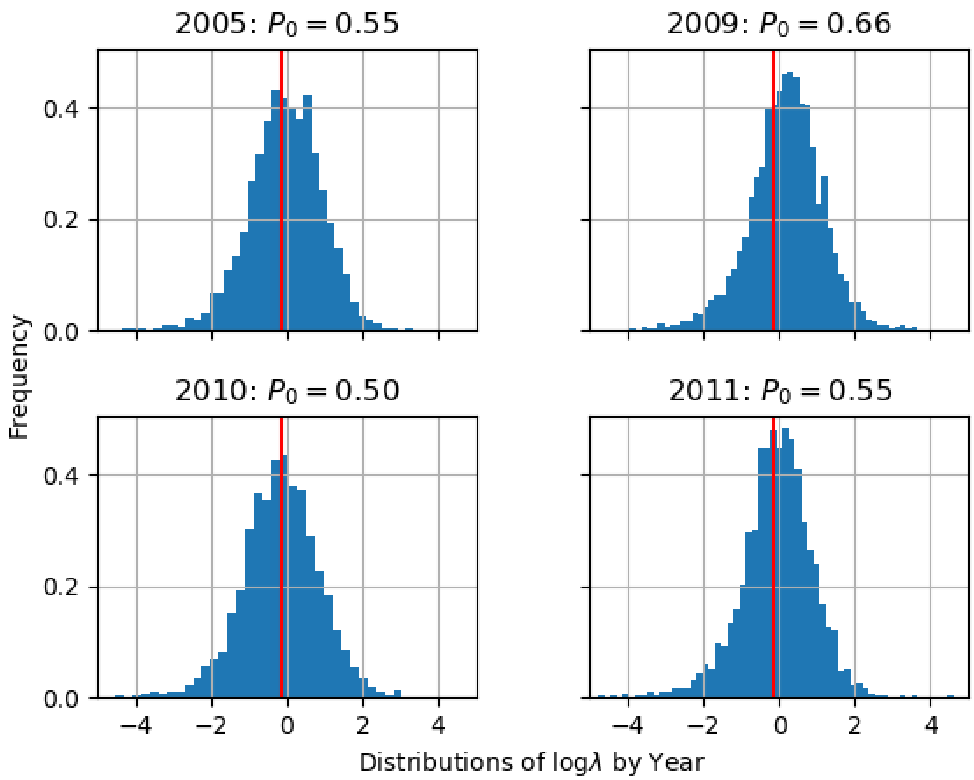

FIGURE 2. Distribution of $\log \lambda$ by Year.

If nominal food prices doubled, what happened to nominal expenditures on food? These increased, as one might expect for a neccessity, but by only 58% for the average household. Thus, either *quantities* of food consumed fell by roughly 40% or there were dietary shifts toward less expensive food. Both of these sorts of changes are evident in our data. Average quantities were lower in 2009 than in 2005 for about two-thirds of all foods. Bread consumption in 2009 was 25% of its level in 2005. Beer consumption fell by 23%, while consumption of "other alcoholic drinks" increased by 35%. Maize and cassava consumption fell by 10–25%, while consumption of matoke (a local starchy staple) increased by 38%. This matches contemporary news reports of near famine conditions in parts of the country, and is consistent with the rightward shift of 0.24 standard deviations in the distribution of $\log \lambda$ seen in Figure 2.[10] The mean value of $\log \lambda$ has some nice features as an aggregate welfare measure; adopting it we would say, roughly, that from 2005 to 2009 we observed a 24% reduction in social welfare for this population.

10. See Figure 1 below for evidence on the ability of the model to predict expenditures for different individual goods.

Prices go up by more than expenditures; there's a reduction in welfare captured by our measure. This all seems sensible. But how does it compare with conventional accounts of changes in poverty, deflating total expenditures by a single price index? Since the rank of the demand system is at least four division by any single index can't be correct, but different single indices can be wrong in different ways.

We follow what is perhaps the most obvious approach. The Ugandan Bureau of Statistics calculates a consumer price index (CPI) using methods which generally follow standard international procedures (Intersecretariat Working Group on Price Statistics 2004). In the Ugandan case, this involved using the 2005 expenditure data to construct weights similar in construction to the "aggregate shares" described in 5.2 (though for a coarser aggregation of goods). These data on shares are combined with data from monthly surveys of urban prices to construct a Laspeyres index.[11] For the critical period of 2005–2009 the CPI increased by 44%. Using this index leads to the inference that real per capita expenditures increased by about 24%, or about 0.25 standard deviations. This is the same magnitude as the change in $\log\lambda$, but in the wrong direction!

6.2.1. *Relation of* $\log\lambda$ *to poverty measures.* Consider the effects of this on measures of headcount poverty; these are summarized in Figure 3. In each year we plot a kernel density estimate of the distribution of log expenditures (deflated by the CPI) on the right, and the distribution of $-\log\lambda$ on the left (we've changed the sign so rightward shifts in both cases imply improvements in welfare). The two distributions shift in opposite directions not only in 2005–09, as noted above, but also in 2010–11.

The World Bank's online PovCalNet uses the same underlying datasets for calculating welfare statistics as we do, and recommends a PPP-adjustment of 946.89 (Ravallion, Chen, and Sangraula 2009). Using this adjustment and the recommended \$1.90 poverty line, the World Bank's figure for headcount poverty in 2005 is 55.4%. We use this figure to pin down a poverty line of 11.27 in (log) 2005 Ugandan shillings, and what we might call a corresponding $\log\lambda$-poverty line of -0.15.

What happens to poverty headcounts if we use our methods? Fixing the poverty line so that 55.4% of the population is below the poverty line in 2005, changes in the distribution of $\log\lambda$ across years imply an increase in the poverty rate to 66% in 2009. By 2010 the distribution of $\log\lambda$ improves by 10% relative to the base year, and the headcount poverty rate falls to 50%. The subsequent effects of a recession (Brunori, Palmisano, and Peragine 2015) then increases the mean $\log\lambda$ back to -0.03 and changed headcount poverty to 55%, both figures close to the original base-year values.

11. This broad description omits important details; see **oag14**.

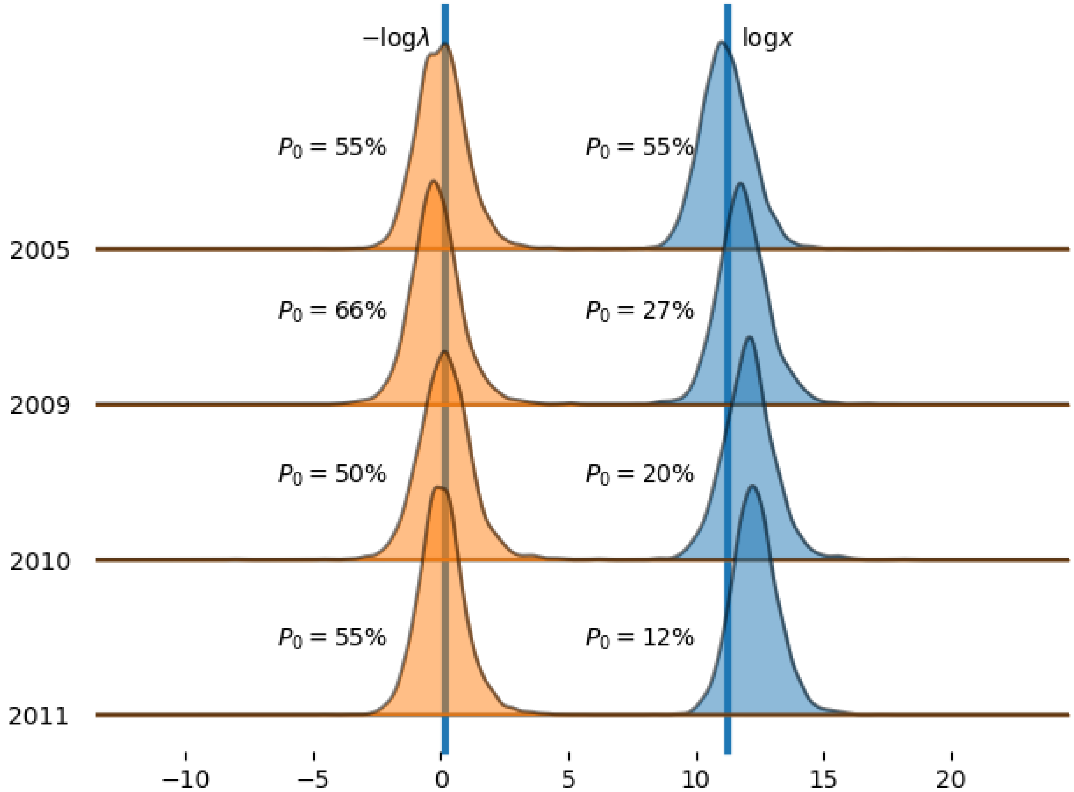

FIGURE 3. Changes in distributions of $-\log\lambda$ and deflated log expenditures over time, with implied headcount poverty statistics.

The conventional calculations tell a very different story. The same CPI adjustment that leads to an incorrect estimate of real expenditures is also incorrect for headcount poverty measures. Over the period in which food prices doubled the conventional approach suggests that poverty fell by half. The methods agree that poverty fell sharply from 2009 to 2010, but then during Uganda's "great recession" of 2010 the conventional calculation suggests a further dramatic fall in poverty, to 12%.

Using some index other than the CPI would yield different implications for headcount poverty. The World Bank seems to use an index which is CPI divided by aggregate per capita consumption expenditures, for example (though details of their calculations are not public). One sometimes encounters the recommendation that one should construct a Laspeyres index using the expenditure shares for households at the poverty line. The advice to use a single index misses the essential point that the demands and expenditure shares themselves vary in a way that can't be explained without using at least four different price indices. Though one could trivially devise a single price index that yielded the same poverty calculations as our approach, that price index would not work for any other realization of prices other than that observed *ex post*.

6.3. **Validation: Estimated Aggregate Shares versus Mean Shares.** In Figure 1 we used data on observed expeditures to produce a plot of a statistic equal to the logarithm of aggregate shares minus the logarithm of mean shares, ordered by the size of the statistic in 2005, and observed that the pattern of observed in that figure could not be generated by any demand system featuring homothetic preferences.

The question naturally arises: is the non-homothetic demand system we've estimated here capable of delivering the pattern of expenditure shares pictured in Figure 1?

This is a challenging test, because although we've used the observed data to estimate the demand system, our estimation procedure is designed to fit conditional expectations of log expenditures to the data, while the shares statistics we've constructed is built using logs of means and sums of expenditures. Jensen's inequality alone tells us that our ability to match the share statistics will depend not only on the estimated equation, but also on the distribution of residuals.

Let $h_i(p, \lambda, z) = \mathrm{E}(\log x_i | p, \lambda, z)$. Assume that the residuals e_{it}^j in the estimating equation (8) are independent and identically normally distributed for each good, with mean zero and variance σ_i^2. Then a simple estimator of $\mathrm{E}(x_{it}^j | p_t, \lambda_t^j, z_t^j)$ is $\exp(h_i(\hat{p}_t, \hat{\lambda}_t^j, z_t^j) + \hat{\sigma}_i^2/2)$, where $\hat{\sigma}_i^2$ is the maximum likelihood estimate of the variance of the residuals for good i, and where $\hat{p}_t$ and $\hat{\lambda}_t^j$ are estimates of price indices and $\log \lambda$ as described in Section 4.1.

We next simply substitute our estimates $\hat{x}_{it}^j$ into the expression defining the statistics ρ_i, and plot the values of these statistics predicted by our model of demand and estimates of prices and $\log \lambda$. The result is picture in the left hand panel of Figure 4.

The left panel of Figure 4 reproduces Figure 1, except with predicted rather than actual shares. The general pattern evidencing non-quasi-homotheticity is readily apparent. But beyond this, the ρ_i statistics calculated using our predicted expenditures have a Spearman correlation coefficient of 0.97 with statistics calculated using the observed data. The right-hand panel of Figure 4 provides a scatter plot of observed vs. predicted values of the statistic, along with a 45 degree line. The scatterplot confirms the success of our demand system at reproducing even patterns in the data that our estimator wasn't designed to fit.

6.4. **Validation: Testing for Selection Bias.** Earlier we described an approach to diagnosing selection problems that might be created by having missing or zero expenditures in our data; details are given in Appendix A. The idea is to estimate elasticities and $\log \lambda$s via a singular value decomposition of a residual matrix as described above, and then to re-estimate after augmenting this matrix with residuals that contain information about which observations are missing. If selection is important, then these estimates should be different.

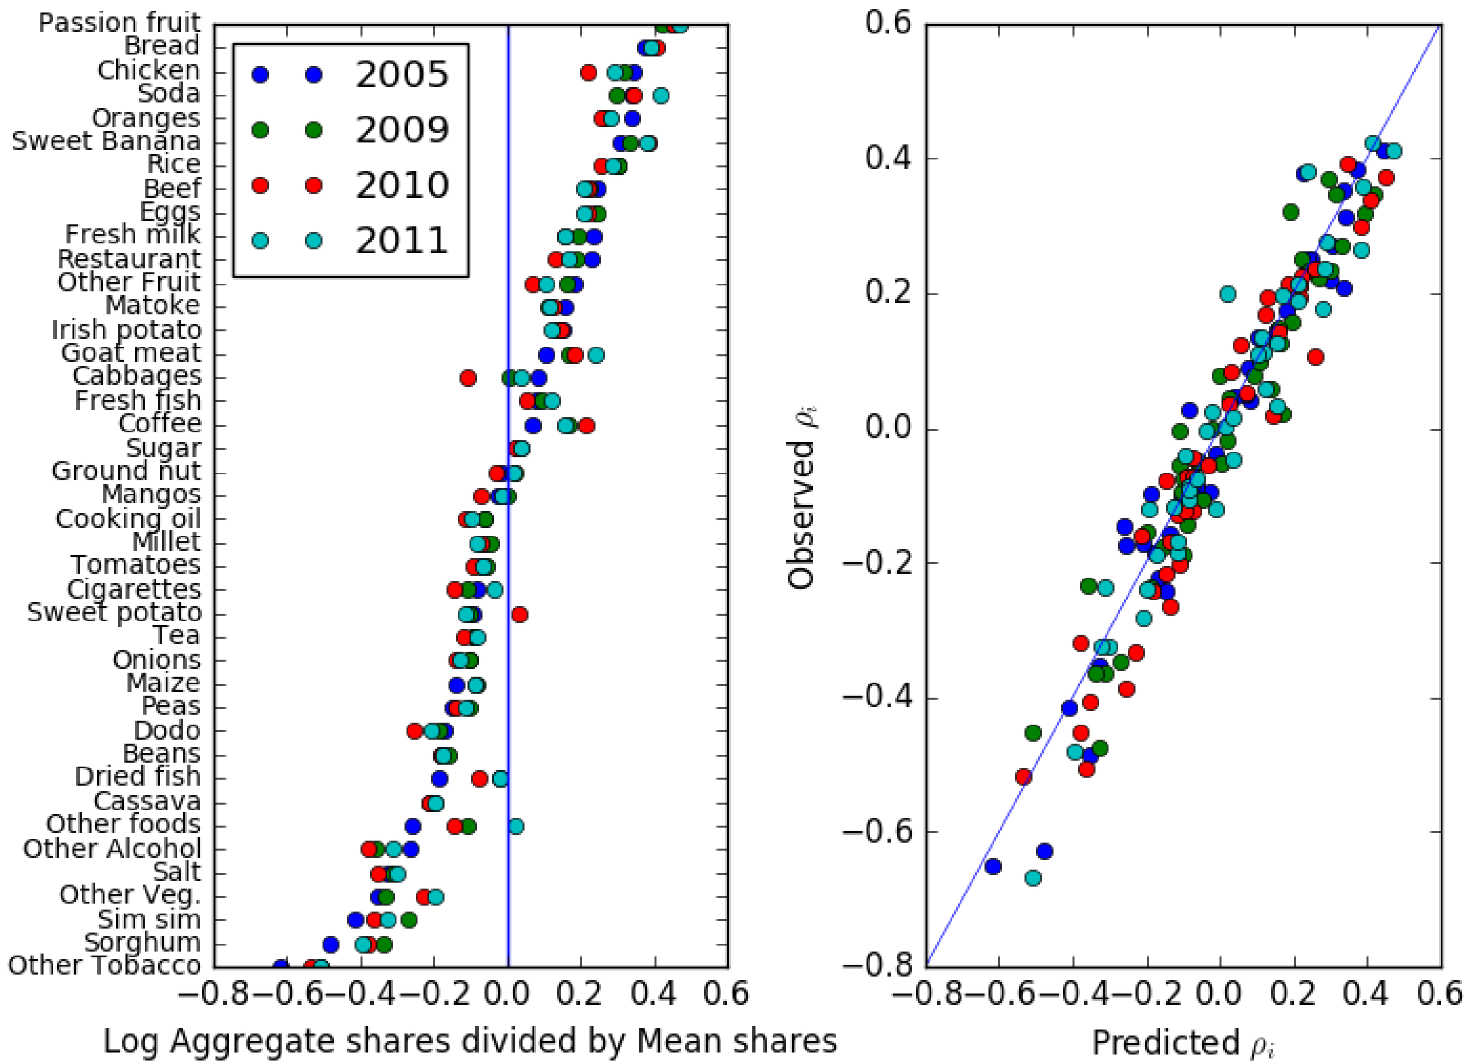

FIGURE 4. Left panel: Predicted log of mean shares minus log of aggregate shares for different years (ordered by ranking in 2005). Right panel: Predicted versus actual, with 45 degree line.

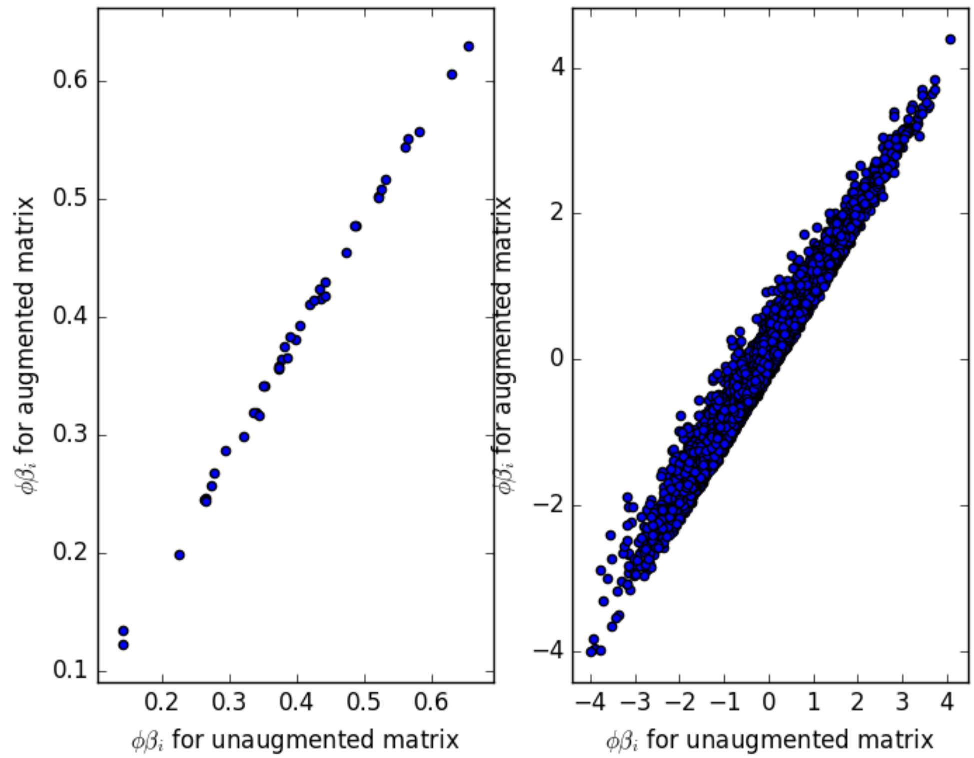

FIGURE 5. Test for selection bias. Scatterplots of β_i and $\log \lambda$ estimated with and without information on missing values.

Figure 5 gives an informal representation of the outcome of such a test, in the form of two scatter plots. The plot on the left shows the estimates of the $\phi\beta_i$ obtained in Table 2, versus the same vector obtained when we augment the residuals in (8) with the residuals from a regression of a missing indicator on the same right-hand side variables. The plot on the right does the same thing, but for estimates of the $\log \lambda_t^j$. As is evident from the figure, the relationships are quite tight, with no discernable evidence of selection bias, and almost no change in orderings (the Spearman correlation coefficients are respectively 1.00 and 0.99).

6.5. **Aggregation and Choice of Goods.** Table 3 provides an accounting of the variance of log expenditures for each of the 41 goods we consider (the final column), and reports the proportion of the total variance accounted for by log prices, household characteristics, and $\log \lambda$. The goods in this case are ordered so that those which have the highest proportion of variance accounted for by $\log \lambda$ come first.

Considering the different columns of Table 3, observe that by construction if a household has some observable characteristic (e.g., larger household size) which tends to led to higher expenditures on all goods then that will *ceteribus paribus* increase estimates of that household's $\log \lambda$. Also that by construction the $\log \lambda$ variables and the residuals will be both mutually orthogonal; however, other pairs of variables in the table need not be mutually orthogonal, and for this reason the sums of the proportions reported in the Table need not sum to one.

We see considerable variation across goods in terms of the proportion of variance accounted for by variation in $\log \lambda$. The goods for which $\log \lambda$ is particularly important involve proportions of variance explained of over 25%, and include "other tobacco", passion fruit, bread, oranges, coffee, "other fruit", soda, rice, sweet bananas, tomatoes, beef, chicken, and eggs. Note that while a considerable proportion of the variation in these goods is accounted for by variation in $\log \lambda$, it doesn't follow that these are all goods principally consumed by wealthier households—in Figure 1 we saw that in fact "other tobacco" was the good most disproportionately represented in the consumption portfolio of poorer households. Many goods have variation in log expenditures which isn't isn't closely related to variation in $\log \lambda$, including salt, sorghum, cassava, sweet potatoes, and millet—for all of these the proportion of total variance accounted for by $\log \lambda$ is less than 10%. And again a low proportion of variance accounted for by $\log \lambda$ doesn't mean that these are goods for poor people—meals in restaurants shows up in Figure 1 as a good that is disproportionately represented in the consumption portfolios of wealthier households, despite the fact that the proportion of variance accounted for by $\log \lambda$ is only 0.13 in this case.

Table 3: Analysis of Variance. The final column gives the total variance of expenditure for the good. The remaining columns give the proportion of variance attributable to each factor. Sorted by the proportion of variance explained by $\log \lambda$.

	$\log \lambda$	Characteristics	Prices	Residual	Total var
Other Tobacco	0.40	0.02	0.01	0.66	0.91
Passion fruit	0.38	0.07	0.02	0.54	0.88
Bread	0.34	0.13	0.01	0.51	0.86
Oranges	0.30	0.08	0.01	0.57	1.01
Coffee	0.30	0.06	0.03	0.67	0.96
Other Fruit	0.30	0.08	0.02	0.62	1.06
Soda	0.29	0.05	0.01	0.64	0.75
Rice	0.28	0.21	0.00	0.58	0.58

Continued on next page

Continued from previous page

	$\log \lambda$	Characteristics	Prices	Residual	Total var
Sweet Banana	0.28	0.12	0.00	0.62	0.82
Tomatoes	0.28	0.11	0.01	0.59	0.59
Beef	0.27	0.17	0.00	0.59	0.45
Chicken	0.26	0.11	0.00	0.73	0.32
Eggs	0.26	0.09	0.00	0.66	0.59
Onions	0.25	0.10	0.00	0.62	0.56
Goat meat	0.25	0.16	0.01	0.64	0.56
Cooking oil	0.25	0.11	0.00	0.61	0.69
Sugar	0.25	0.18	0.01	0.56	0.65
Fresh milk	0.23	0.10	0.01	0.62	0.96
Fresh fish	0.23	0.11	0.00	0.69	0.57
Matoke	0.23	0.19	0.01	0.58	0.86
Peas	0.23	0.09	0.01	0.73	0.63
Other foods	0.22	0.03	0.03	0.70	1.38
Cigarettes	0.22	0.07	0.02	0.69	1.01
Cabbages	0.21	0.09	0.02	0.73	0.46
Ground nut	0.19	0.09	0.00	0.71	0.74
Other Veg.	0.19	0.07	0.01	0.77	0.82
Irish potato	0.17	0.14	0.01	0.68	0.74
Dodo	0.17	0.07	0.02	0.77	0.70
Other Alcohol	0.16	0.05	0.01	0.77	1.03
Mangos	0.15	0.10	0.03	0.72	1.07
Dried fish	0.15	0.10	0.01	0.70	0.71
Restaurant	0.13	0.09	0.01	0.75	1.03
Tea	0.13	0.11	0.09	0.68	0.58
Sim sim	0.13	0.06	0.01	0.83	0.58
Beans	0.11	0.16	0.00	0.73	0.68
Maize	0.11	0.16	0.00	0.71	0.99
Millet	0.10	0.14	0.00	0.74	0.89
Sweet potato	0.08	0.23	0.01	0.70	0.88
Salt	0.06	0.19	0.04	0.74	0.34
Sorghum	0.06	0.11	0.01	0.83	1.05

Continued on next page

Continued from previous page

	$\log\lambda$	Characteristics	Prices	Residual	Total var
Cassava	0.02	0.17	0.00	0.83	0.95

We discussed above the aggregation we've performed to reduce 72 different food items or categories to 49; in our results above we dropped a further eight of these items on the grounds that the number of available observations was quite small. We think our decisions about aggregation and dropping are sensible, but there's no denying that they are also somewhat *ad hoc*.

A natural question to ask is just how much these *ad hoc* decisions matter? We address this by re-estimating the entire system of 72 goods without any aggregation, and without dropping any goods at all. This yields a new set of estimated $\log\lambda$s and elasticities. We compute a correlation of 0.94 between the $\log\lambda$s estimated using the 41-good system and the 72-good system, and a correlation of 0.95 between the $\phi\beta_i$ estimated for the goods in common between the two systems. We conclude that our aggregation is not important to the analysis.

This raises another question: since we can go from 72 goods to 41 and still obtain good estimates of households' IMUE, can we reduce the number of goods further? The proportion of variance in expenditures toward the bottom of Table 3 explained by variation in $\log\lambda$ is very small, suggesting that these goods have little value in estimating $\log\lambda$. Dropping such goods has no particular value in our present analysis, using already collected data, but if it is possible to perform our analysis with fewer goods that could reduce the costs of future data collection in Uganda.

Suppose we left out goods for which the proportion of variance accounted for by $\log\lambda$ was small? While identifying some 'optimal' set of goods is a topic for another paper, here we report the results of an experiment which drops goods sequentially from our 72-good system, ordered according to the proportion of variance accounted for by $\log\lambda$ (but always including rice, our numéraire). For each good dropped we calculate the Spearman correlation coefficient between the values of $\log\lambda$ estimated using the entire 72 good system with the smaller system; we then reorder this list of goods according the successive changes in the correlation coefficient.

Results from this exercise are reported in Figure 6. We find that 29 of 72 goods seem to have almost no value at all in helping to calculate the $\log\lambda$, in the sense that the ordering of the IMUEs is completely unchanged if we drop these goods. We can drop down to 47 goods before the Spearman correlation coefficient falls below 0.99, and to 33 before it drops below 0.95. If we restrict ourselves to just a dozen goods of the original 72 goods

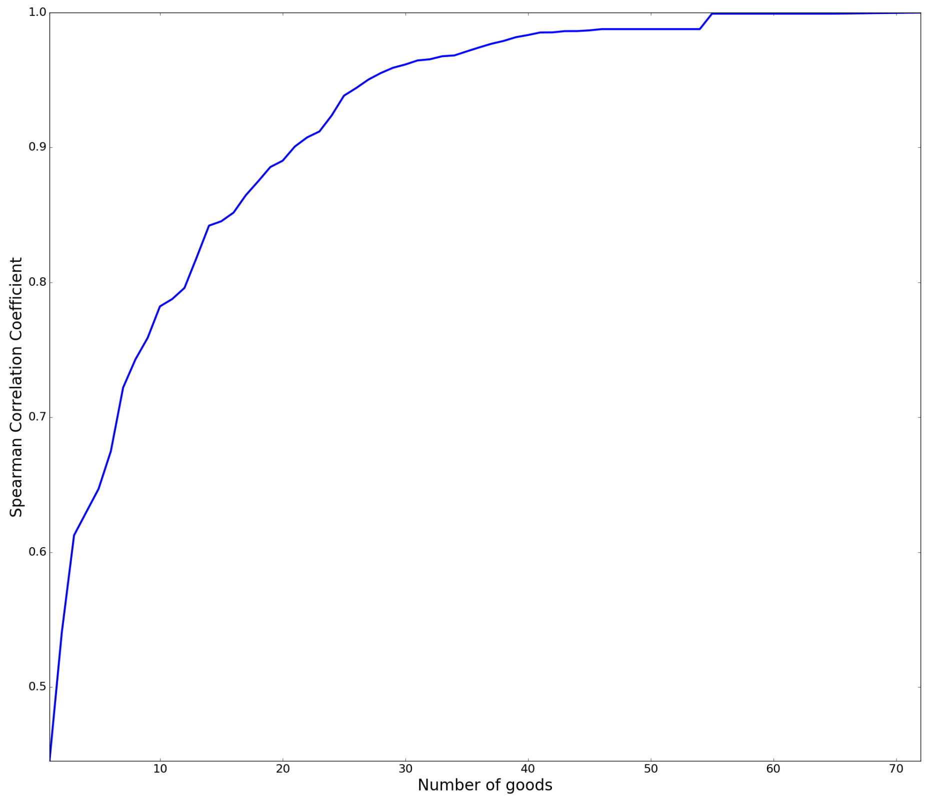

Figure 6. Spearman correlation between IMUE estimated using full 72-good system with reduced systems

(sugar, maize (flour), restaurant (food), cassava (fresh), beans (dry), fresh milk, soda, beef, onions, tomatoes, cooking oil, and passion fruits) the Spearman correlation between the $\log \lambda$ estimated in the 72-good system and the 12-good system falls to 0.80.

If our aim is to estimate IMUEs with a minimum of data this may seem encouraging, but what is the cause of the observed reduction in correlation? There turn out to be two reasons for the change. The first is the obvious one: when we drop goods we're throwing away at least *some* information about the IMUEs. Our rule of thumb for discarding goods means we tend to throw out the goods with less information first, which is why there's almost no reduction in our ability to estimate the $\log \lambda$ when we discard less elastic goods such as cassava, salt, and millet. The second reason is related to missing data: when we restrict ourselves to a small number of commodities, there are fewer households with enough non-missing data on expenditures, and some households drop out of the estimation entirely, affecting our estimates of both elasticities and of $\log \lambda$.

There are ways to address this second problem were we to collect new data on a smaller set of goods. Using a longer recall period ought to result in observing more positive expenditures, for example, or simply providing the training and incentives for enumerators to minimize item non-response in the expenditure module are both measures which we might expect to increase the proportion of positive item expenditures we observe (indeed, the reduction of both enumerator and respondent burden associated with having a smaller expenditure module seems likely to improve item response rates).

Finally, our rule for choosing goods is also sensible but *ad hoc*. Both the criterion of Spearman's correlation and the sequence in which we've dropped goods could either be justified or improved upon, but this is a subject for another paper.

6.6. **Measuring Heterogeneity in Risk Attitudes.** When the household-specific λs we've calculated are equal to the marginal utility of expenditures, then $\omega = \partial \log \lambda / \partial \log x$ is what Frisch called the household's "money flexibility", or what we might think of as the expenditure elasticity of marginal utility. If preferences are also von Neumann-Morgenstern then $-\omega$ can also be interpreted as the household's relative risk aversion, and its negative reciprocal as the elasticity of intertemporal substitution. Further, with data on x we can calculate the quantity ω for each household.

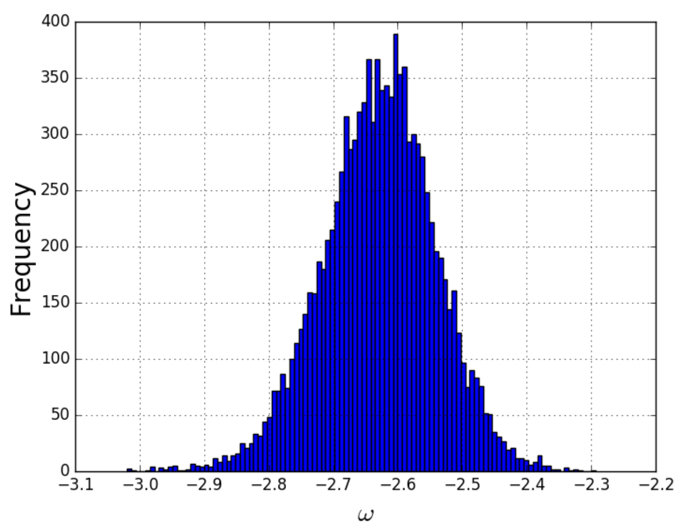

FIGURE 7. Distribution of estimated ω in 2005. Standard deviation of pooled estimates is 0.09; kurtosis is 0.43.

Results of calculating households' values of ω in 2005 are shown in the histogram of Figure 7 (calculating these values in later years would require us to measure how relative prices had changed). The mean of this distribution is -2.63, with a standard deviation of 0.09. Households are assumed to have identical β_i parameters and to face identical prices, so

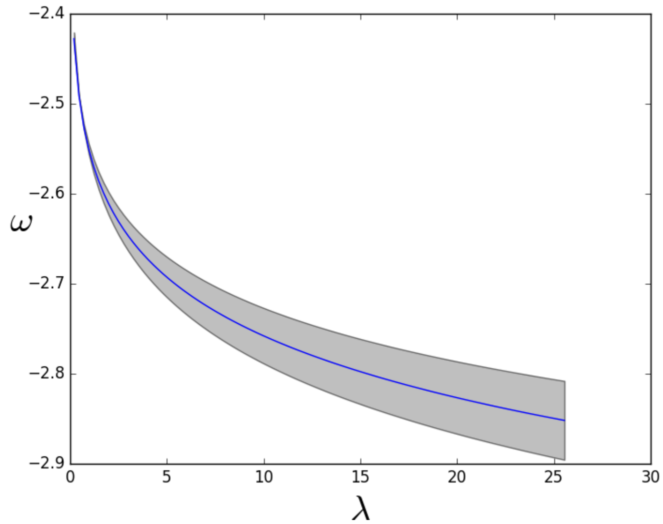

FIGURE 8. Calculated ω versus λ. Varying all estimates of β_i by plus or minus two standard errors yields the shaded area.

differences in ω (and hence in risk aversion, under the separability assumptions given above) are driven entirely by differences in $\log \lambda$. The estimated values of ω in Figure 7 are well within the range of plausible relative risk aversions, though the range is smaller than one might suppose based on the evidence of Chiappori et al. (2014).

The relationship between ω and $\log \lambda$ is illustrated in Figure 8, where λ is on the horizontal axis. The central plotted line offers calculations of ω given the point estimates of β_i elasticities given in Table 2 (taking $\phi = 1$); the shaded region allows each point estimate to vary by plus or minus one standard error, giving some sense of how sensitive our estimates are to imprecisely estimated Frisch elasticities. From the figure, one can see that ω is decreasing in λ, and so increasing in total expenditures. Since ω is negative, and has the interpretation of the elasticity of λ with respect to total expenditures, the figure illustrates the point that the utility of wealthier households is less sensitive to variation in total expenditures than is the utility of poorer households. Translated into statements about risk aversion, Figure 8 indicates that households have decreasing relative risk aversion with respect to expenditures on food.

Knowing the quantity ω is what we need for estimating within-period demands or indifference curves. However, if we're interested in measuring risk attitudes or more generally the curvature of the momentary utility function this isn't enough—the purely cross-sectional demand behavior we observe simply can't non-parametrically identify the momentary utility function, because any monotonic transformation of utility (say $M(U)$) would generate exactly the same intra-temporal demands.

To see this, let us suppose that the "true" (momentary) indirect utility function is not $V(p,x)$, but a monotonic transformation $V^*(p,x) = M(V(p,x))$. We've estimated $\lambda = \partial V/\partial x$, but if utility is $M(U)$ then the marginal utility of expenditures isn't λ, but rather $\lambda M'(U)$, where M' is the derivative of the monotone transformation. Without knowledge of the transformation M we're limited in what we can say about ω, risk attitudes, or intertemporal substitution.

However, with modest additional assumptions it's possible to estimate the empirical *distribution* of households' ω, up to two unknown parameters.

Recall that the Arrow-Pratt measure of relative risk aversion for a household with indirect utility $V^*(p,x)$ is given by the negative of

$$\omega^*(p,x) = x\frac{\partial^2 V^*/\partial x^2}{\partial V^*/\partial x}.$$

With $V^*(p,x) = M(V(p,x))$, we have $\partial V^*/\partial x = M'(V(p,x))\lambda(p,x)$ (recalling that $\lambda = \partial V/\partial x$). Differentiating again and applying the chain rule allows us to write

$$\omega^*(p,x) = x\frac{M''}{M'}\lambda(p,x) + \frac{\partial \log \lambda(p,x)}{\partial \log x}.$$

The quantities in the second term are the ω elasticities shown in Figure 7, though the value will depend on the unknown factor of proportionality ϕ. The first term involves the first and second derivatives of the unknown function M. A judicious parameterization is

$$M(U) = \frac{U^{1-\sigma} - 1}{1-\sigma};$$

this matches related assumptions used by MaCurdy (1983) or Browning, Deaton, and Irish (1985). With this parameterization of M we have the first term $-x\lambda\frac{M''}{M'} \approx \sigma$ to a first order approximation, so that we have

$$\omega^*(p,x) \approx -\sigma + \omega(p,x)/\phi.$$

This then allows us to identify households' ω up to the unknown constants σ and ϕ. If households all have a common transformation M, then the distribution of ω in Figure 7 will be approximately the same as the distribution of (minus) households' relative risk aversions.

The finding that households have heterogeneous relative risk aversions echoes recent findings for households in Thailand; as here, Chiappori et al. (2014) use observed data on expenditures to estimate the distribution of relative risk aversion up to an unknown parameter. However, their estimates assume homothetic utility and rely on a maintained hypothesis that households are fully insured. We are able to avoid these strong assumptions entirely; the analogous assumptions which allow us to identify the distribution of risk attitudes (up to unknown location and scale parameters) are just the much weaker requirements that elasticities are constant (instead of all equal to one, and may vary across goods) and that the household maximizes utility within the period subject to a budget constraint.

7. Conclusion

In this paper we've outlined some of the key methodological ingredients needed in a recipe to estimate a simple measure of household welfare. This measure is closely related to the household's marginal utility of expenditures; it differs only in that it controls for household characteristics and adopts a particular cardinalization of utility.

The methods described are theoretically coherent, in the sense that they're consistent with a particular utility-derived demand system. Further, our approach lends itself to straightforward statistical inference and hypothesis testing, and is very parsimonious in its data requirements.

Our approach involves estimating an incomplete demand system of a new sort which features a highly flexible relationship between total expenditures and demand. The methods described here involve using one or more cross-sections of data on household expenditures on different nondurable goods and/or services. The limited data requirements suggest that these methods may be useful in constructing programs to inexpensively *measure* and *monitor* households' welfare over both different environments and across time.

In an application of these methods we use four rounds of data from Uganda. We focus on food expenditures in this dataset, estimating a system of 41 demands. We estimate both household $\log \lambda$ and Frischian elasticity parameters from this expenditure system, in addition to other demand parameters. A separate analysis allows us to characterize the distribution of households' relative risk aversions; we find convincing evidence of heterogeneity, though the distribution is not fully identified.

8. References

Ahn, Seung Chan, Young Hoon Lee, and Peter Schmidt. 2001. GMM estimation of linear panel data models with time-varying individual effects. *Journal of Econometrics* 101:219–255.

Almås, Ingvild. 2012. International income inequality: measuring PPP bias by estimating engel curves for food. *American Economic Review* 102 (2): 1093–1117.

Almås, Ingvild, Timothy K.M. Beatty, and Thomas Crossley. 2018. *Lost in translation: what do engel curves tell us about the cost of living?* CESifo Working Paper Series 6886. CESifo Group Munich. `https://EconPapers.repec.org/RePEc:ces:ceswps:_6886`.

Angelucci, Manuela, and Giacomo De Giorgi. 2009. Indirect effects of an aid program: how do cash transfers affect non-eligibles' consumption. *The American Economic Review* 99, no. 1 (March): 486–508.

Arellano, M. 1987. Computing robust standard errors for within-groups estimators. *Oxford Bulletin of Economics and Statistics* 49 (4): 431–34.

Arellano, Manuel, Richard Blundell, and Stéphane Bonhomme. 2017. Earnings and consumption dynamics: a nonlinear panel data framework. *Econometrica* 85 (3): 693–734.

Atkin, David, Ben Faber, Thibault Fally, and Marco Gonzalez-Navarro. 2018. A new engel on the gains from trade: theory and evidence within and across countries. Manuscript.

Attanasio, Orazio P, and Valerie Lechene. 2014. Efficient responses to targeted cash transfers. *Journal of Political Economy* 122 (1): 178–222.

Attanasio, Orazio P, and Luigi Pistaferri. 2016. Consumption inequality. *Journal of Economic Perspectives* 30 (2): 3–28.

Attfield, Clifford LF, and Martin J Browning. 1985. A differential demand system, rational expectations and the life cycle hypothesis. *Econometrica* 53 (1): 31–48.

Benson, Todd, Samuel Mugarura, and Kelly Wanda. 2008. Impacts in uganda of rising global food prices: the role of diversified staples and limited price transmission. *Agricultural Economics* 39 (s1): 513–524.

Blundell, Richard. 1998. Consumer demand and intertemporal allocations: engel, slutsky, and frisch. *Econometric Society Monographs* 31:147–168.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. 2016. Consumption inequality and family labor supply. *American Economic Review* 106 (2).

Bourguignon, F., M. Browning, and P. A. Chiappori. 2009. Efficient intra-household allocations and distribution factors: implications and identification. *Review of Economic Studies* 76 (2): 503–528.

Browning, Martin. 1986. The costs of using frisch demand functions that are additive in the marginal utility of expenditure. *Economics Letters* 21 (3): 205–207.

Browning, Martin, Francois Bourguignon, Pierre-Andre Chiappori, and Valérie Lechene. 1994. Income and outcomes: a structural model of intrahousehold allocation. December, *Journal of Political Economy* 102 (6): 1067–1096.

Browning, Martin, Thomas F. Crossley, and Joachim Winter. 2014. The measurement of household consumption expenditures. *Annual Review of Economics* 6:475–501.

Browning, Martin, Angus Deaton, and Margaret Irish. 1985. A profitable approach to labor supply and commodity demands over the life cycle. *Econometrica* 53:503–543.

Brunori, Paolo, Flaviana Palmisano, and Vito Peragine. 2015. *Inequality of opportunity during the great recession in uganda.* Working Paper 2015/039. WIDER.

Caron, Justin, Thibault Fally, and James R. Markusen. 2014. International trade puzzles: a solution linking production and preferences. *The Quarterly Journal of Economics* 129 (3). doi:10.1093/qje/qju010.

Chiappori, Pierre-André, Krislert Samphantharak, Sam Schulhofer-Wohl, and Robert M. Townsend. 2014. Heterogeneity and risk sharing in village economies. *Quantitative Economies* 5, no. 1 (March): 1–27. http://onlinelibrary.wiley.com/doi/10.3982/QE131/epdf.

Cochrane, John H. 1991. A simple test of consumption insurance. *Journal of Political Economy* 99:957–976.

Comin, Diego A, Danial Lashkari, and Martí Mestieri. 2015. *Structural change with long-run income and price effects.* Technical report. National Bureau of Economic Research. http://www.nber.org/papers/w21595.pdf.

Cooper, Russel J, and Keith R McLaren. 1996. A system of demand equations satisfying effectively global regularity conditions. *The Review of Economics and Statistics:* 359–364.

Deaton, Angus. 1992. Household saving in LDCs: credit markets, insurance, and welfare. *Scandinavian Journal of Economics* 94:253–273.

———. 1997. *The analysis of household surveys.* Baltimore: The John Hopkins University Press.

Deaton, Angus, and Margaret Irish. 1984. Statistical models for zero expenditures in household budgets. *Journal of Public Economics* 23 (1-2): 59–80.

Deaton, Angus, and John Muellbauer. 1980. An almost ideal demand system. *The American Economic Review* 70, no. 3 (June): 312–326.

Development Research Group of the World Bank. Downloaded 2017. *PovcalNet: the on-line tool for poverty measurement.* http://iresearch.worldbank.org/PovcalNet/.

Diewert, W. E. 1971. An application of the shephard duality theorem: a generalized leontief production function. *Journal of Political Economy* 79:461–507.

Engel, Ernst. 1857. Die productions-und consumtionsverhaltnisse des königreichs sachsen. *Zeitschrift des Statistischen Biireaus des Koniglich Sachsischen Ministeriums des Innern,* nos. 8 & 9: 1–54.

Fiedler, John L, Keith Lividini, Odilia I Bermudez, and Marc-Francois Smitz. 2012. Household consumption and expenditures surveys (hces): a primer for food and nutrition analysts in low-and middle-income countries. *Food and Nutrition Bulletin* 33 (3_suppl2): S170–S184.

Frisch, Ragnar. 1959. A complete scheme for computing all direct and cross demand elasticities in a model with many sectors. *Econometrica* 27, no. 2 (April): 177–196.

———. 1964. Dynamic utility. *Econometrica:* 418–424.

———. 1978. *New methods of measuring marginal utility.* Porcupine Press.

Golub, G. H., and C. Reinsch. 1970. Singular value decomposition and least squares solutions. *Numerische Mathematik* 14 (5): 403–420.

Gorman, W. M. 1959. Separable utility and aggregation. *Econometrica* 27:469–81.

Gorman, W.M. 1976. Essays in economic analysis. Chap. Tricks with utility functions, 211–243. Cambridge University Press. `http://users.ox.ac.uk/~econ0211/writings/gorman.pdf`.

Hanoch, Giora. 1975. Production and demand models with direct or indirect implicit additivity. *Econometrica* 43 (3): 395–419.

Heckman, James J. 1974. Shadow prices, market wages, and labor supply. *Econometrica* 42 (4): 679–694. `https://www.jstor.org/stable/pdf/1913937.pdf`.

———. 1976. A life-cycle model of earnings, learning, and consumption. *Journal of political economy:* S9–S44. `https://www.jstor.org/stable/pdf/1831101.pdf`.

Heckman, James J., and Thomas E. MaCurdy. 1980. A life cycle model of female labour supply. *The Review of Economic Studies:* 47–74.

Horowitz, Joel L. 2003. Bootstrap methods for markov processes. *Econometrica* 71 (4): 1049–1082.

Houthakker, Hendrik S. 1960. Additive preferences. *Econometrica:* 244–257. `https://www.jstor.org/stable/pdf/1907719.pdf`.

Intersecretariat Working Group on Price Statistics. 2004. *Consumer price index manual: theory and practice.* Technical report. https://www.ilo.org/public/english/bureau/stat/download/cpi/cpi_manual_en.pdf.

Karaivanov, Alexander, and Robert M. Townsend. 2014. Dynamic financial constraints: distinguishing mechanism design from exogenously incomplete regimes. *Econometrica* 82 (3): 887–959.

Kinnan, Cynthia. 2017. Distinguishing barriers to insurance in thai villages. Unpublished Ms.

Laczó, Sarolta. 2015. Risk sharing with limited commitment and preference heterogeneity: structural estimation and testing. *Journal of the European Economic Association* 13 (2): 265–292.

Lewbel, Arthur. 1991. The rank of demand systems: theory and nonparametric estimation. *Econometrica* 59 (3): 711–730.

———. 2003. A rational rank four demand system. *Journal of Applied Econometrics* 18 (2): 127–135. http://www.jstor.org/stable/30035200.

Lewbel, Arthur, and Krishna Pendakur. 2009. Tricks with hicks: the easi demand system. *The American Economic Review* 99, no. 3 (June): 827–863.

Lise, Jeremy. 2012. On-the-job search and precautionary savings. *Review of Economic Studies* 80 (3): 1086–1113.

Lise, Jeremy, and Ken Yamada. 2017. Household sharing and commitment: evidence from panel data on individual expenditures and time use. Manuscript.

Little, Roderick, and Donald Rubin. 2002. *Statistical analysis with missing data.* 2nd. Wiley Series in Probability and Statistics. John Wiley & Sons.

Low, Hamish, Costas Meghir, and Luigi Pistaferri. 2010. Wage risk and employment risk over the life cycle. *American Economic Review* 100 (4): 1432–67.

Mace, Barbara J. 1991. Full insurance in the presence of aggregate uncertainty. *Journal of Political Economy* 99:928–956. http://www.jstor.org/stable/2937653.

MaCurdy, Thomas E. 1981. An empirical model of labor supply in a a life-cycle setting. *Journal of Political Economy* 89 (61): 1059–1085.

———. 1983. A simple scheme for estimating an intertemporal model of labor supply and consumption in the presence of taxes and uncertainty. *International Economic Review* 24 (2): 265–289.

Mazzocco, Maurizio, and Shiv Saini. 2011. Testing efficient risk sharing with heterogeneous risk preferences. *American Economic Review* 102 (1): 428–468.

Meredith, William. 1993. Measurement invariance, factor analysis and factorial invariance. *Psychometrika* 58 (4): 525–543.

Muellbauer, John. 1976. Community preferences and the representative consmer. *Econometrica* 44 (5): 979–999.

Pelleg, Dan, Andrew W Moore, et al. 2000. x-means: extending k-means with efficient estimation of the number of clusters. In *Proceedings of the 17th international conference on machine learning,* 1:727–734.

Ravallion, Martin, Shaohua Chen, and Prem Sangraula. 2009. Dollar a day revisited. *The World Bank Economic Review* 23 (2): 163–184.

Townsend, Robert M. 1994. Risk and insurance in village India. *Econometrica* 62 (3): 539–591.

Vella, Francis. 1998. Estimating models with sample selection bias: a survey. *The Journal of Human Resources* 33:127–169.

Wedel, Michel, and Wagner A Kamakura. 2001. Factor analysis with (mixed) observed and latent variables in the exponential family. *Psychometrika* 66 (4): 515–530.

Wooldridge, Jeffrey M. 2002. *Econometric analysis of cross section and panel data.* Cambridge, MA: MIT Press.

Young, Alwyn. 2012. The african growth miracle. *Journal of Political Economy* 120 (4): 696–739.

Appendix A. For Online Publication: Issues with Missing Data

We consider two issues related to the possibility that certain item expenditures may be zero or missing for some households in our dataset. The first is a practical issue, having to do with the calculation of the singular value decomposition (SVD) of a matrix when some of the elements of that matrix are missing. The resolution of this problem involves relying just on observed data. This is clearly okay if the data is "missing at random" (Little and Rubin 2002), but if there's selection on unobservables then the matter is less clear. Thus, our second issue involves obtaining conditions under which selection doesn't compromise our SVD calculation.

A.1. **A method for computing the singular value decomposition of random matrices with missing elements.** Consider a random matrix $\mathbf{X}$ having a continuous distribution with support over some subset of $\mathbb{R}^{n\times m}$ and a second random matrix $\mathbf{M}$ of the same dimension with elements either 1 or 0. Assume without loss of generality that $n \leq m$, and that the matrix $\mathbf{X}$ is "low rank" in the sense that its rank is strictly less than n. We observe only a matrix $\mathbf{A} = \mathbf{X} \otimes \mathbf{M}$, where $\otimes$ is the Hadamard product. Any zero element of $\mathbf{A}$ is said to be "missing;" we assume that both the row and column sums of $\mathbf{M}$ are greater than zero.

We wish to construct a matrix $\hat{\mathbf{X}}$ close to $\mathbf{X}$ in the Frobenius norm. If we assume that the rank of $\mathbf{X}$ is some known number r then we have the compact singular value decomposition of the matrix $\mathbf{X} = \mathbf{U}^*\mathbf{\Sigma}^*\mathbf{V}^{*\top}$, with $\mathbf{U}^*$ $n \times r$, $\mathbf{\Sigma}^*$ $r \times r$ and diagonal, and $\mathbf{V}^*$ $m \times r$.

Our strategy involves first using the m columns of $\mathbf{A}$ to estimate the $n \times r$ matrix $\mathbf{U}^*\mathbf{\Sigma}^*$. We construct a matrix

$$\mathbf{P} = m\mathbf{A}\mathbf{A}^\top \oslash (\mathbf{M}\mathbf{M}^\top),$$

where $\oslash$ is the Hadamard (element-by-element) division operator. $\mathbf{P}$ can be interpreted as the matrix product $\mathbf{A}\mathbf{A}^\top$ scaled to ignore zero elements; note that it's positive semi-definite by construction. Then the square root of the eigenvalues of $\mathbf{P}$ is an estimator for the diagonal matrix $\mathbf{\Sigma}^*$, while the corresponding eigenvectors $\mathbf{U}$ estimate $\mathbf{U}^*$.

With $\mathbf{U\Sigma}$ in hand we proceed row by row to construct an estimate of $\mathbf{V}^*$: suppose a is a column vector from the matrix $\mathbf{A}$, and that $\mathbf{A}$ has the compact singular value decomposition $\mathbf{U\Sigma V}^\top$. The vector a can be partitioned into two parts y and x, while a matrix $\mathbf{U\Sigma}_x$ can be constructed by selecting just the rows of $\mathbf{U\Sigma}$ corresponding to the x elements of the vector a. Then the row of $\mathbf{V}^*$ corresponding to a can be estimated by

$$v = (\mathbf{U\Sigma}_x)^+ x,$$

where the + operator here indicates the Penrose-Moore pseudo-inverse. Iterating over all m columns of $\mathbf{A}$ then yields the desired matrix $\hat{\mathbf{X}}$.

A.2. **Effects of non-random selection.** Let (X, M) be corresponding columns of the random matrices $\mathbf{X}$ and $\mathbf{M}$, each an n-dimensional random variable, where X has realizations x, and M observable realizations m. If X and M are independent we say that the data is "missing at random" and the methods describe above deliver what we want, but otherwise there may be a selection problem. We want to establish conditions on the joint distribution $F(X, M)$ which ensure that $\mathrm{E}(\mathbf{U}\mathbf{\Sigma}) = \mathbf{U}^*\mathbf{\Sigma}^*$, provided the columns of $\mathbf{X}$ and $\mathbf{M}$ are uncorrelated across households.

To this end, we modestly extend an idea due to Meredith (1993) (Theorem 5). This idea in turn can be thought of as extending the usual result from the program evaluation literature that selection on observables doesn't compromise consistent estimation in the program evaluation literature. Instead of assuming that selection depends only on a set of observable variables, we allow selection to depend on a set of latent variables. (This also generalizes the idea from fixed effects estimation in unbalanced panels that selection isn't a problem if the probability of a household being missing in a given period can be "explained" by a fixed effect.)

In particular, let W be a p-dimensional latent variable with realizations w that underlies X; and let L be an q-dimensional random variable with realizations l that underlies M.

A collection of functions $s_i(l)$ comprise a *selection rule* that gives the marginal probability that $m_i = 1$ for an individual with attributes l, $i = 1, \dots, n$. We adopt the following assumptions:

Assumption 1. (1) *The conditional distribution of X given w is non-degenerate, and first and second moments exist;*

(2) *W and L are not independent; and*

(3) *The joint probability distribution of M is assumed to be conditionally independent across i, so we have*

$$s(m|L) = \prod_{i=1}^{n} s_i(L)^{m_i}(1 - s_i(L))^{1-m_i}.$$

The first assumption is standard. The second is necessary for knowledge of w to tells us something about the probability of selection. The third could perhaps be relaxed, but allowing conditional dependence in selection would considerably complicate estimation.

Now, let $\mathrm{E}_s(g(X)|w, l)$ denote the conditional expectation and $\mathrm{Var}_s(g(X)|w, l)$ the conditional covariance matrix of $g(X)$ conditional on w given the selection rule $s(m|l)$ for any Lebesgue-measurable function g. Otherwise expectations and covariances $\mathrm{E}(g(X)|w, l)$ and $\mathrm{Var}(g(X)|w, l)$ correspond to the population distribution $F(x, w, l)$ without any selection.

Proposition 2. *Given Assumption 1, if* $E(X|w,l) = E(X|w)$ *and* $Var(X|w,l) = Var(X|w)$ *for all* (w,l) *then* $E_s(X|w,l) = E(X|w)$ *and* $Var_s(X|w,l) = Var(X|w)$ *for all* w*, all selection functions* s*, and all* l *such that* $s_i(l) > 0$ *for some* $i = 1, \dots, n$*.*

Proof. Let $F(x,w,l,m)$ be the joint distribution of (x,w,l,m) in the population; then the selected joint distribution is given by the Lebesgue integral

$$F_s(x,w,l) = \int s(m|l)dF(x,w,l) / \int s(m|l)dF(l)$$

for all $(x \otimes m, w, l)$ such that $s_i(l) > 0$ for some $i = 1, \dots, n$, and then $F_s(x|w,l) = F(x|w,l)$ over the same set; thus $\mathrm{E}_s(g(X)|w,l) = \mathrm{E}(g(X)|w,l) = \mathrm{E}(g(X)|w)$. The first part of the proposition is then established by taking $g(X) = X$, while the second is established taking $g(X) = g(X|w) = \mathrm{E}(XX^\top|w) - \mathrm{E}(X|w)\mathrm{E}(X|w)^\top$ for any w. ☐

The key to the result is that the factors which determine selection directly affect X only via their interaction with the latent variables W; this is related to but weaker than a condition sometimes called "missing at random conditional on X" (Wedel and Kamakura 2001).

A special case which satisfies the conditions of Proposition 2 occurs when the selection rule doesn't depend on L at all, so that we have $s(m|L) = s(m|L')$ for all (L, L'). In this case data are "missing at random", in the lexicon of Little and Rubin (2002), and we are able to obtain unbiased estimates of $(\mathbf{U}^*, \mathbf{\Sigma}^*, \mathbf{V}^*)$. In the more general case in which the selection rule varies with L but the conditions of the proposition are satisfied we can expect to obtain unbiased estimates only of $(\mathbf{U}^*, \mathbf{\Sigma}^*)$, as the sample of households for which we observe expenditures may be systematically different from the population.

A.3. **A simple test for selection bias.** Our case is related to the case in which there's attrition in a panel of households, but in which the factors that cause attrition can be captured by fixed effects. In our case suppose that we're using only a single cross-section, with a large number of households N and a smaller fixed number of goods n; now the fact that expenditures on some goods are zero or missing is just like observations for some households being missing in the unbalanced panel case. The idea for fixed effects panel estimation is that the selection process depends on fixed household characteristics. Our latent variables aren't household fixed effects, but instead latent interactions which capture variation in prices (a_{it}) and household resources ($\log \lambda_t^j$), but these are also exactly the features we might be concerned would lead to selection problems were we estimating demands for a single good, as in Deaton and Irish (1984). If, after conditioning on these, expenditures are zero or missing simply because of the timing of stocking decisions or some other random censoring process then we can expect our linear model with interactive latent variables a_{it} and $\beta_i \log \lambda_t^j$ to deliver consistent results, much as in the case of Ahn, Lee, and Schmidt (2001).

Here we describe a simple test of the central conditions required in Proposition 2. Suppose that we have $X = WA + U_X$, $L = WB + U_L$, and $M = LC + U_M$. Then we have

$$\begin{aligned}[\mathbf{X}, \mathbf{M}] &= [\mathbf{WA}, \mathbf{LC}] + [\mathbf{U}_X, \mathbf{U}_L] \\ &= \mathbf{W}[\mathbf{A}, \mathbf{BC}] + [\mathbf{U}_X, \mathbf{U}_L\mathbf{C} + \mathbf{U}_M].\end{aligned}$$

Under the null hypothesis that $\mathrm{E}(X|W, L) = \mathrm{E}(X|W)$, a singular value decomposition of $[\mathbf{X}, \mathbf{M}]$ allows us to estimate both $\mathbf{W}$ and $\mathbf{A}$. On the other hand, if the hypothesis doesn't hold and $\mathbf{W}$ is not orthogonal to $\mathbf{U}_L\mathbf{C} + \mathbf{U}_M$ then the SVD of this augmented matrix will estimate quantities which depend on $\mathbf{U}_L\mathbf{C} + \mathbf{U}_M$.

In the application of this paper the matrix $\mathbf{X}$ is obtained from the residuals of (8), delivering a matrix which is orthogonal to household characteristics and good-time effects. Thus, we similarly filter the dummy variables $\mathbf{M}$ by regressing these on the same right-hand side, and using estimated residuals from this regression. A maintained assumption of our application is that the rank of the matrix $\mathbf{W}$ is one; thus the test proposed here becomes a question of whether the decomposition of the augmented matrix yields significantly different estimates of the rank one $\mathbf{U\Sigma}$ and rank one $\mathbf{V}$ what we obtain when we simply decompose $\mathbf{X}$.

Appendix B. For Online Publication: Food items across rounds and aggregation

Food codes and items are fairly consistently recorded across rounds, but not perfectly so; further, some are clearly sensibly treated as substitutes (e.g., different size bunches of matoke). Other food items are treated separately in some rounds (e.g., "Watermelon" in 2010 and 2011) but assigned to an aggregate (e.g., "Other Fruits") in other rounds, necessitating the use of the coarser aggregate to achieve consistency across rounds. Table B.1 gives a precise accounting of all codes and aggregation. We supply a "Preferred Label" column and an "Aggregate Label." The "preferred" label eliminates minor differences in spelling or word usage across rounds (e.g., "Matoke" versus "matooke"). The "Aggregate Label" need not be unique; expenditures for all items with the same "Aggregate Label" will be summed together, yielding what we call a "minimally-aggregated" set of data of food expenditures. This minimal aggregation confines itself to combining expenditures on different food items which seem to obviously be very close substitutes. Sometimes these differences are just in units: we aggregate "clusters" and "heaps" of Matoke, for example. Othertimes the form of the good is somewhat different: fresh and dried cassava are aggregated, for example.

Table B.1: Labels for various food items in different rounds, with “Preferred” and “Aggregate” labels.

Code	Preferred Label	Aggregate Label
100	Matoke (??)	Matoke
101	Matoke (bunch)	Matoke
102	Matoke (cluster)	Matoke
103	Matoke (heap)	Matoke
104	Matoke (other)	Matoke
105	Sweet Potatoes (fresh)	Sweet Potatoes
106	Sweet Potatoes (dry)	Sweet Potatoes
107	Cassava (fresh)	Cassava
108	Cassava (dry/flour)	Cassava
109	Irish Potatoes	Irish Potatoes
110	Rice	Rice
111	Maize (grains)	Maize
112	Maize (cobs)	Maize
113	Maize (flour)	Maize
114	Bread	Bread
115	Millet	Millet
116	Sorghum	Sorghum
117	Beef	Beef
118	Pork	Pork
119	Goat Meat	Goat Meat
120	Other Meat	Other Meat
121	Chicken	Chicken
122	Fresh Fish	Fresh Fish
123	Dry/Smoked fish	Dry/Smoked fish
124	Eggs	Eggs
125	Fresh Milk	Fresh Milk
126	Infant Formula	Infant Formula
127	Cooking oil	Cooking oil
128	Ghee	Ghee
129	Margarine, Butter, etc	Margarine, Butter, etc
130	Passion Fruits	Passion Fruits

Continued on next page

Continued from previous page

Code	Preferred Label	Aggregate Label
131	Sweet Bananas	Sweet Bananas
132	Mangoes	Mangoes
133	Oranges	Oranges
134	Other Fruits	Other Fruits
135	Onions	Onions
136	Tomatoes	Tomatoes
137	Cabbages	Cabbages
138	Dodo	Dodo
139	Other vegetables	Other Vegetables
140	Beans (fresh)	Beans
141	Beans (dry)	Beans
142	Ground nuts (in shell)	Ground nuts
143	Ground nuts (shelled)	Ground nuts
144	Ground nuts (pounded)	Ground nuts
145	Peas	Peas
146	Sim sim	Sim sim
147	Sugar	Sugar
148	Coffee	Coffee
149	Tea	Tea
150	Salt	Salt
151	Soda	Soda
152	Beer	Beer
153	Other Alcoholic drinks	Other Alcoholic drinks
154	Other drinks	Other drinks
155	Cigarettes	Cigarettes
156	Other Tobacco	Other Tobacco
157	Restaurant (food)	Restaurant (food)
158	Restaurant (soda)	Soda
159	Restaurant (beer)	Beer
160	Other juice	Other juice
161	Other foods	Other foods
162	Peas (dry)	Peas
163	Ground nut (paste)	Ground nuts

Continued on next page

Continued from previous page

Code	Preferred Label	Aggregate Label
164	Green pepper	Other Vegetables
165	Pumpkins	Other Vegetables
166	Avocado	Other Fruits
167	Carrots	Other Vegetables
168	Eggplant	Other Vegetables
169	Watermelon	Other Fruits
170	Pineapple	Other Fruits
171	Pawpaw	Other Fruits

The aggregation in Table B.1 results in total of 49 different items. Most of these are straightforward types of food, such as peas, mangoes, ground nuts, maize, or sugar. Food consumed in restaurants is a category of its own, however, and may be thought of an aggregate bundle of food and services. Alcoholic beverages account for two additional categories, "beer" and "other alcoholic drinks." And then finally there are two non-food categories included, "cigarettes" and "other tobacco." Altogether there are five categories which are explicitly undifferentiated aggregates: "other fruits", "other vegetables", "other alchoholic drinks", "other drinks", and "other tobacco." Other categories may be implicitly aggregated: for example, "ground nut" includes nuts shelled, unshelled, and made into paste. Finally, even after aggregation some of these categories contain very few positive observations in at least some years; dropping these yields a total of 41 categories.[12]

B.1. **Aggregate and Mean Shares.** Aggregate and mean shares for 41 "aggregated" goods.

Goods	Aggregate 2005	Aggregate 2011	Mean 2005	Mean 2011
Matoke	0.113	0.114	0.097	0.099
Maize	0.079	0.084	0.092	0.091
Cassava	0.077	0.092	0.090	0.105
Restaurant	0.074	0.073	0.061	0.070
Sweet potato	0.070	0.063	0.078	0.075

Continued on next page

12. Excluded goods include beer, infant formula, butter & margarine, ghee, ground nuts, pork, other juice, other drinks, and other meat.

Continued from previous page

Goods	2005	2011	2005	2011
Beans	0.066	0.074	0.080	0.090
Beef	0.058	0.058	0.044	0.046
Sugar	0.055	0.044	0.051	0.042
Fresh milk	0.041	0.040	0.033	0.034
Rice	0.027	0.027	0.022	0.021
Fresh fish	0.024	0.022	0.022	0.020
Ground nut	0.022	0.024	0.023	0.024
Cooking oil	0.021	0.022	0.022	0.022
Dried fish	0.020	0.022	0.022	0.021
Chicken	0.019	0.026	0.014	0.019
Tomatoes	0.019	0.017	0.020	0.017
Other Alcohol	0.018	0.016	0.021	0.020
Soda	0.016	0.012	0.012	0.008
Bread	0.015	0.015	0.010	0.010
Other foods	0.015	0.003	0.021	0.003
Millet	0.015	0.014	0.016	0.015
Other Fruit	0.012	0.019	0.010	0.017
Goat meat	0.012	0.015	0.010	0.010
Irish potato	0.012	0.012	0.010	0.011
Other Veg.	0.011	0.016	0.018	0.020
Sorghum	0.009	0.012	0.017	0.018
Cigarettes	0.009	0.004	0.009	0.004
Dodo	0.008	0.007	0.010	0.010
Onions	0.008	0.010	0.009	0.011
Passion fruit	0.007	0.003	0.005	0.002
Mangos	0.006	0.004	0.007	0.004
Sweet Banana	0.006	0.005	0.004	0.003
Salt	0.006	0.005	0.008	0.007
Eggs	0.006	0.005	0.004	0.004
Cabbages	0.005	0.006	0.005	0.007
Tea	0.005	0.003	0.005	0.004
Peas	0.005	0.005	0.006	0.006
Sim sim	0.004	0.004	0.007	0.006

Continued on next page

Continued from previous page

Goods	2005	2011	2005	2011
Oranges	0.002	0.002	0.002	0.002
Other Tobacco	0.002	0.001	0.004	0.002
Coffee	0.001	0.001	0.001	0.000

Appendix C. For Online Publication: Estimating the rank of the demand system

Using the fact that our estimates are normally distributed, we can evaluate the likelihood that the rank of the system is $k \leq n$ by solving the k-means problem

$$C_k = \min_{b \in \mathbb{R}^k} [min_k |\beta_i - b_k|]_i^\top \mathbf{V}^{-1} [min_k |\beta_i - b_k|]_i,$$

where $\mathbf{V}$ is the estimated covariance matrix of our estimates of β. Note that under the null hypothesis that the rank of the matrix k then the statistic C_k is distributed χ^2_{n-k}. We compute the corresponding likelihood L_k. Using the Bayesian Information Criterion (BIC)

$$\mathrm{BIC}_k = k \log n - 2 \log L_k$$

gives in our case an estimate of rank four for the demand system. Figure 9 presents the result of these calculations; one can see that the minimum is achieved at $k = 4$, with $b = (0.18, 0.36, 0.47, 0.62)$.

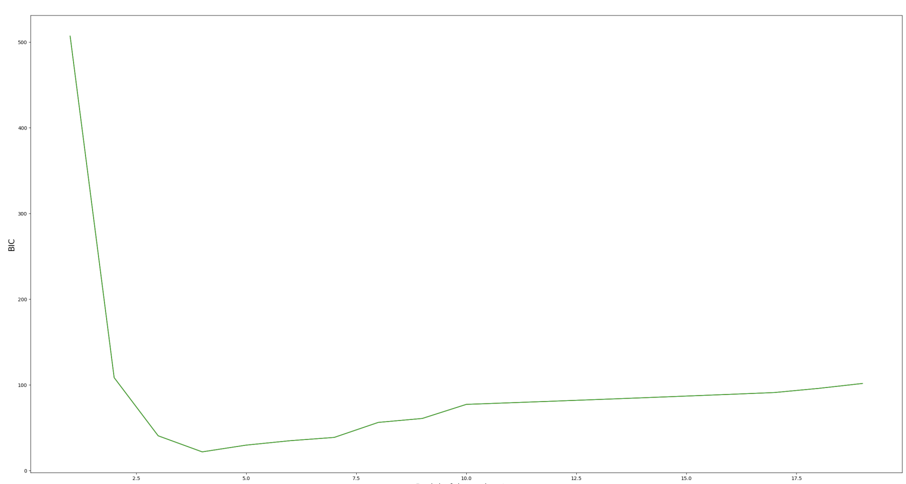

Figure 9. Using the BIC criterion to estimate the rank of the demand system.