

UC Berkeley

CEGA Working Papers

Title

Estimating Household Welfare from Disaggregate Expenditures

Permalink

<https://escholarship.org/uc/item/2cg4q0b9>

Author

Ligon, Ethan

Publication Date

2020-02-02

Series Name: WPS
Paper No.: 109
Issue Date: 4 February 2020

Estimating Household Welfare from Disaggregate Expenditures

Ethan Ligon



CEGA

Center for Effective Global Action

Working Paper Series

Center for Effective Global Action
University of California



This paper is posted at the eScholarship Repository, University of California. http://escholarship.org/uc/cega_wps Copyright © 2019 by the author(s).

The CEGA Working Paper Series showcases ongoing and completed research by faculty affiliates of the Center. CEGA Working Papers employ rigorous evaluation techniques to measure the impact of large-scale social and economic development programs, and are intended to encourage discussion and feedback from the global development community.

Recommended Citation:

Ligon, Ethan. (2020). Estimating Household Welfare from Disaggregate Expenditures. Working Paper Series No. WPS-109. Center for Effective Global Action. University of California, Berkeley.

ESTIMATING HOUSEHOLD WELFARE FROM DISAGGREGATE EXPENDITURES

ETHAN LIGON

ABSTRACT. Existing models of life-cycle demand typically relate household welfare to the scale of total consumption expenditures, deflated by a price index. This is only correct if preferences are homothetic and expenditure shares for particular goods within a period are fixed, and is sharply at odds with strong empirical evidence, including Engel's Law. Instead of using the scale of total expenditures we show how to exploit variation in the composition of a household's consumption portfolio 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, and analyze the effects of shocks such as the 2008 food price crisis on welfare in Uganda. Our results contrast sharply with those obtained using total 'real' expenditures, which cannot capture the differential impact of increased food prices across the expenditure distribution.

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

. *Date:* August 8, 2019.

. 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, Valerie Lechene, 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 the construction of 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 Total real consumption expenditures give us the *scale* of the household’s consumption, but tells us nothing about the *composition* of the household’s consumption portfolio. 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. Its form is uniquely determined by the requirement that it must be possible to infer a relative measure of well-being for one household by observing its consumptions relative to a second similar household, a property we call “Conditionally Independent Relative Consumptions for Similar Households” (CIRCSH). The second innovation is the construction of 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 show that a globally regular demand system will have (some) goods satisfying CIRCSH if and only if demands for those goods take a particular form we call Constant Frisch Elasticity (CFE). Aside from allowing us to draw inferences regarding welfare measures λ , this system has unrestricted rank (in the sense of Lewbel 1991), and can be shown to nest most of the globally regular demand systems commonly encountered in the empirical literature. We further devise a novel but simple method of estimation

1. Beyond these examples, 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 paper focuses exclusively on aggregate household consumption expenditures.

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 CFE demands can be estimated using one or more rounds of cross-sectional data on disaggregate household expenditures, in a specification using logarithms of those disaggregate expenditures. This estimator delivers “Frisch elasticities,” estimates of the index λ of each household’s marginal utility of expenditures (MUE), and 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.

Our estimator is implemented in two (or three) 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 in a second step using a singular value decomposition (SVD) of the matrix of residuals, yielding estimates of Frisch elasticities and λ s, up to a time-varying location parameter. When there are multiple rounds of data an optional third step exploits an auxiliary regression to estimate these location parameters.

In practice, disaggregate data on expenditures almost invariably contains many “zeros” or missing data – in our application no household reports positive consumption of all goods. This poses a practical challenge – how can one compute the SVD of a matrix with missing values? It also raises the familiar specter of selection bias. So after describing the steps of our estimation procedure, we first outline an efficient method for calculating the SVD of a matrix with missing values. We then extend (in Appendix A) a result from the psychometric literature, which establishes general conditions under which estimation will be unbiased. In our application these conditions will be satisfied if expenditures are “missing at random” *conditional* on (shadow) prices, observed household characteristics, and $\log \lambda$.³ Finally, the conditions for unbiasedness immediately suggest a set of exclusion restrictions which can be exploited

3. The usual explanations for observing zeros in such data include being at a corner (the consumer would like to consume a negative quantity at prevailing prices); stock-outs (the shadow price is higher than any observed price); or infrequency of purchases (the consumer either has an unobserved stock of the good, or the good is highly substitutable over short periods); and simple measurement error. Note that systematic variation in all of these is plausibly explained by prices, household characteristics and wealth.

to construct a simple test, which we describe informally in Section 4.2, in detail in Appendix A.3, and which we implement for our application in Section 6.3.

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 2005–2012. These data are instances of the World Bank’s *Living Standards Measurement Surveys* (LSMSs), which have now been conducted in many countries across many years (Deaton 1997). The Ugandan data are of 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 2010–11. With nothing 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’ λ s, and to measure the consequences of these shocks for rates of poverty, inequality, and risk aversion for differently situated households.

1. RELATED LITERATURE

The classical approach to welfare measurement associated with Hicks adopts some cardinal utility function (and the corresponding indirect utility function) as its lodestar. But it is simply not possible to use observed demand behavior across different sorts of households to infer the indirect utility function (Pollak and Wales 1979), and the present paper can be regarded as one of several recent approaches to devising measures of welfare which can be inferred from observed behavior. A number of papers attempt 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. Early applications include Costa (2001) and Hamilton (2001). More recently, 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 Progreso experiment with the aim of

testing a 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 along with 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.

Note that none of the approaches to welfare described in the previous paragraph is entirely consistent with the classical Hicksian approach to welfare – there is not, for example, a one-to-one mapping between the expenditure share of food and the indirect utility function. And although our measure does have a direct relationship to the indirect utility function (it can be interpreted as the partial derivative of indirect utility with respect to total expenditures), one cannot generally recover the indirect utility function from our measure λ . If one observed and estimated demands for *all* goods this system of demands could be ‘integrated’ to obtain a rationalizing indirect utility function, as in Lewbel and Pendakur (2009), but much of the appeal of our approach (as with the alternative approach of Browning (1999)) is that welfare inferences can be drawn using only an incomplete demand system.

The relation between 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 1970s. This empirical development allowed econometricians to deal with the difficulties of measuring unobserved permanent income in the context of life-cycle models (Heckman 1974, 1976; Heckman and MaCurdy 1980; MaCurdy 1981). This was pioneering, but an assumption of complete markets was central to identification in these early applications, so that risk-averse households would 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.⁴ Nevertheless, the payoff from these assumptions was significant by assuming a particular cardinalization (i.e., assuming a particular cardinal utility function within a period) of homothetic preference structures and constant MUEs, it was possible to estimate the MUEs 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 Blundell, Browning, and Meghir (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 thus far use variation over both time and households to simultaneously estimate within-period demands and the dynamic elements of the model. 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 λ s), 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

4. 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.

using information about contemporary expenditures. This means that estimating an *index* of the ‘true’ marginal utility of expenditures exhausts the information available in contemporaneous expenditures. Thus, though we estimate λ s 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 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. We use this toolkit to understand Ugandan consumption expenditures over 2005–2012; this 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 long-recognized consequence (Prais 1959) 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

relative prices. While others have tackled this problem (Muellbauer 1974; Pendakur 2002), our approach has much more modest data requirements. Further, 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 yields estimates of poverty rates that *fell* over the course of the 2008 food price crisis and subsequent 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 (up to unknown location and scale parameters). The key to identification of these is simply the requirement that preferences be von Neumann-Morgenstern. We find strong evidence of heterogeneity in relative risk aversion 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 homothetic constant relative risk averse (CRRA) preferences typically assumed in the literature.

The rest of this paper is organized as follows. We begin in Section 3 with a model of household behavior, and posit our CIRCSH condition, which attributes differences in consumption portfolios observed for otherwise similar households to differences in the households' budgets. We then show that this condition is satisfied if and only if demands are described by the globally regular Constant Frisch Elasticity system. In Section 4 we describe a method of estimating this demand system. Section 5 discusses the Ugandan data used for our application, and then in Section 6 we present results on estimated demand elasticities and demand system rank. We conduct two distinct validation exercises, first, showing that our estimated demand system not only can predict estimated log expenditures well, but also does well at predicting

both aggregate and individual expenditure shares. Second, we introduce a diagnostic measure to see whether non-randomly missing data might cause problems for our estimation, and conclude that it does not. This brings us to our main empirical finding, which in Section 7 exploits our estimates of IMUEs to understand the welfare consequences of the 2008 “food price” crisis for welfare in Uganda, with implications for the measurement of poverty. We contrast results from our approach with those obtained from the conventional approach, and finally extend our approach to estimate the relationship between welfare changes induced by the food price crisis and changes in distribution of risk aversion.

3. MODEL OF HOUSEHOLD BEHAVIOR

In this section we provide a simple model of household demand behavior,⁵ and use this model along with a requirement that welfare differences between similar households should be identifiable from relative consumptions to obtain a set of globally regular demands. These demands turn out to depend on an index of household welfare that can be interpreted as the logarithm of the households’ marginal utility of expenditures (λ). The index λ can be regarded as a function which will generally depend on total expenditures, prices, and household characteristics. By allowing total expenditures to vary we obtain a Marshallian demand system, but we are interested in applications in which we may not observe all expenditures. This suggests that we instead hold the function λ itself constant (Heckman and MaCurdy 1980), yielding a dual “Frischian” demand system (Browning, Deaton, and Irish 1985). Such demands are dual to Marshallian demands that do not generally have an explicit representation, and 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”

5. The discussion and application in this paper have focused on demand at the level of the *household*. It will be apparent that if we had data on expenditures at the *individual* level all the analysis of the paper would go through. Future research should consider what if we have household-level data, but inequality and differences in preferences across individuals within the household, as in recent papers such as Dunbar, Lewbel, and Pendakur (2013) or Calvi (Forthcoming).

preferences of Caron, Fally, and Markusen (2014), but allows for flexible substitution patterns across goods.

3.1. The household's one-period consumer problem. Suppose that in a particular period t a household with some vector of characteristics z_t faces a vector of prices p_t for \bar{n} goods, and has budgeted a quantity of the numeraire good x_t to spend on contemporaneous consumption $c_t \in X \subseteq \mathbb{R}_{++}^{\bar{n}}$.⁶ The household's preferences are assumed to be intertemporally separable, with orderings over consumption bundles within a period summarized by a utility function $U \in \mathcal{U}$, where \mathcal{U} is the set of utility functions mapping X into \mathbb{R} which are increasing, concave, and continuously twice differentiable. The set \mathcal{U} is said to be the set of *regular* utility functions.

At time t the household solves the standard consumer's problem of maximizing utility subject to a budget constraint, or

$$(1) \quad V(x_t, p_t; z_t) = \max_{c \in X} U(c; z_t) \quad \text{such that } p_t^\top c \leq x_t.$$

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

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

where u_i denotes the partial derivative of the momentary utility function U with respect to the i th good, and where p_{it} is the price of the i th good in period t . 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 vector of demands $c(x_t, p_t; z_t)$ and a function $\lambda_t^* = \lambda^*(x_t, p_t; z_t)$. If these demands satisfy (2) for some positive (x, p) we say that the demand system is *regular* at (x, p) ; if demands satisfy (2) for all positive (x, p) then these demands are said to be *globally regular*. A classical identification result tells us that knowledge of globally regular demands allows us to identify the utility function of a household with characteristics z_t up to a monotonic transformation $M : \mathbb{R} \rightarrow \mathbb{R}$, determining an equivalence class of utility functions $\mathcal{C}(U) = \{M(U) | c(x, p; z) \in \arg \max_{c \in X} M(U(c; z)) \text{ such that } p^\top c \leq x\}$. If $M(U) \in \mathcal{U}$ for all $U \in \mathcal{U}$ (i.e., it maps regular utility functions into regular utility functions) then we say that M is a *regular* transformation, and denote by \mathcal{M} the class of regular transformations.

6. 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.

By construction, regular transformations of the utility function don't affect demands, but of course they can change the marginal utility of expenditures. The following proposition enumerates some facts about the relationship between the MUE of a consumer with utility U and that of a consumer with utility $M(U)$ when M is regular.

Proposition 1. *Let $M : \mathbb{R} \rightarrow \mathbb{R}$ 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 $\lambda^*(x, p; z) = \partial V / \partial x$. 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 $M(U)$ and 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 being continuously differentiable, decreasing in x , and increasing in p .*

Proof. A solution to (1) gives a vector of demands $c(x, p; z)$. Since $U \in \mathcal{U}$ is differentiable, $\lambda^*(x, p; z) = \partial V / \partial x$ 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 both λ^* and $M'(V(x, p; z))$ are positive but monotonically decreasing in x and monotonically increasing in p ; accordingly the product of the two is similarly monotone, establishing a one-to-one relationship between λ^* and λ^M . \square

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

Since 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

will have a one-to-one relationship with every other demand-consistent utility function and MUE.

Because many different pairs of (U, λ) are consistent with any particular demand system we're free to make assumptions that pin down particular utility functions provided these don't restrict observed behavior. Such assumptions can be thought of as normalizations or perhaps 'cardinalizations,' as this amounts to choosing a particular cardinal (momentary) utility functions. Of course, for welfare comparisons across households it's important to adopt the *same* cardinalization.

3.2. Welfare and the demand system. Let $\aleph(x, p, z)$ be any household-level welfare measure which aggregates information on budgets, prices, and characteristics into a real-valued scalar. Interesting examples of such welfare measures include the indirect utility function (which plays a central role in Hicksian welfare analysis), total expenditures (perhaps deflated by a price index, which play a prominent role in measurements of poverty), the expenditure share of food (as suggested by Engel), or the MUE we've discussed above.

The idea of this paper is to explore what can be learned about such a welfare measure from the composition of the consumption portfolio, rather than its scale. Let us conduct a simple thought experiment, and suppose that there two different households with total budgets x and x' , respectively, which consume bundles c and c' . The two households are otherwise identical, and face the same prices; we call such households *similar*. We don't observe x and x' , but do observe quantities (or expenditures) for several different goods consumed by the two households (c_i, c'_i) , $i \in \mathcal{I} \subseteq \{1, 2, \dots, \bar{n}\}$. Then any scale-free measurement of the welfare of the two similar households should be based on relative demands c'_i/c_i for different goods i .

What restrictions on demands will allow us to draw inferences regarding welfare from a collection of relative consumptions for similar households? Since prices and characteristics are the same for both households, we want conditions which would allow us to attribute observed relative consumptions solely to differences between the two households' welfare. Further, since we may not observe all relevant prices or household characteristics, we want the mapping from relative consumptions for similar households into welfare to *not* depend on prices or characteristics. This need not be true of all goods, but for at least some different goods relative demands should satisfy a condition we call "Conditional Independence of Relative Consumptions for Similar Households" (CIRCSH):

Condition 1 (CIRCSH). *Good i exhibits Conditional Independence of Relative Consumptions for Similar Households if, for any positive x, x', p , and any z , demand c_i satisfies*

$$\frac{c_i(x', p; z)}{c_i(x, p; z)} = g_i(\aleph(x', p, z), \aleph(x, p, z))$$

for some non-constant scalar functions \aleph and g_i .

Let CIRCSH denote both the condition and also the set of goods having demands satisfying the condition. Thus, for any good $i \in \text{CIRCSH}$, the relative consumptions of similar households can depend on prices, characteristics, and budgets, but only via an aggregator \aleph which takes values which can vary across households, not across goods. The relationship between the aggregators for two similar households are related to relative consumption of good i by a good-specific function g_i . Knowledge of the function g_i along with observations on relative consumptions provides information about the relative welfare of the two households as measured by the welfare aggregators $\aleph(x', p, z)$ and $\aleph(x, p, z)$ for a given (p, z) .

How restrictive is the CIRCSH condition? Any globally regular demands that are “Generalized Linear” (Muellbauer 1975) will satisfy CIRCSH (and many demands are CIRCSH that are not generalized linear). Special cases of such demands described by Muellbauer can be obtained by placing additional restrictions on the aggregator \aleph . For example, globally regular “Price Independent Generalized Linear” (PIGL), “Price Independent Generalized Logarithmic” (PIGLog), or any homothetic demand system will satisfy CIRCSH with an aggregator $\aleph(x, p, z) = \aleph(x, z)$ (if one sets aside the requirement of global regularity then the claim goes through for some translation of demands, where the translation depends only on prices).

Aside from this, CIRCSH on its own is not enough to estimate demands or measure welfare; additional restrictions on the aggregator \aleph and good-specific functions g_i are required. However, the requirement of global regularity along with CIRCSH provides sufficient restrictions. The following proposition indicates the class of demands which are both CIRCSH and also globally regular. Further, it implies that there is no loss of generality in taking the welfare measure $\aleph(x, p, z)$ to be equal to the IMUE $\lambda(x, p, z)$.

Proposition 2. *Any globally regular demand system has demand for good i satisfying Condition CIRCSH if and only if demand for good i can be written in the form*

$$(3) \quad \log c_i(x, p, z) = \log \gamma_i(p) + \delta_i(z) - \beta_i \log \lambda(x, p, z),$$

for some function of prices $\gamma_i(p)$, some function of characteristics $\delta_i(z)$, and some positive scalar β_i , where the positive function λ is common across all goods.

Proof. It's immediate that any demand with form (3) satisfies Condition CIRCOSH, with $\aleph(x, p, z) = \lambda(x, p, z)$ and $g_i(\aleph(x', p, z), \aleph(x, p, z)) = (\aleph(x, p, z)/\aleph(x', p, z))^{\beta_i}$. To establish that these demands are part of a globally regular system, we have only to construct a utility function $U \in \mathcal{U}$ with the property that maximizing this utility function subject to a budget constraint yields demands of the form (3). Let $U(c; z) = \sum_i e^{\delta_i(z)} \frac{\beta_i}{\beta_i - 1} (c_i^{1-1/\beta_i} - 1)$. It's simple to verify that this function is in \mathcal{U} provided only that β_i is positive and not equal to one for all i . The corresponding marginal utilities are given by $u_i(c, z) = e^{\delta_i(z)} c_i^{-1/\beta_i}$. Substituting this into the first order condition for good i

$$u_i(c; z) = p_i \lambda(x, p, z)$$

and rearranging yields a demand function of the required form.

We establish the converse in four steps. First, note that since demand for good i is globally regular then demand for this good must satisfy the first order conditions $u_i(c; z) = p_i \lambda$, with λ the multiplier on the budget constraint in the consumer's problem, or in a vector notation $u(c; z) = p \lambda$, with c and p both \bar{n} -vectors. Since elements of \mathcal{U} are concave, it follows that $u(\cdot; z)$ is invertible, yielding $c = u^{-1}(p \lambda; z) \equiv f(p \lambda, z)$. For any good i , define

$$\Phi_i(x', x, p, z) = \left(\frac{f_i(p \lambda(x', p, z), z)}{f_i(p \lambda(x, p, z), z)} \right) / g_i(\aleph(x', p, z), \aleph(x, p, z)),$$

and observe that CIRCOSH holds if and only if $\Phi(x', x, p, z) = 1$ for all (x', x, p, z) ; this implies that $\Phi(x'', x, p, z) = \Phi(x', x, p, z)$. Then

$$\frac{f_i(p \lambda(x'', p, z), z)}{f_i(p \lambda(x', p, z), z)} = \frac{g_i(\aleph(x'', p, z), \aleph(x, p, z))}{g_i(\aleph(x', p, z), \aleph(x, p, z))}.$$

Since x doesn't appear on the left-hand side of this equation, it follows that $g_i(\aleph', \aleph) = h_i(\aleph')/h_i(\aleph)$ for some positive function h_i .

Second, rewrite the budget constraint in the consumer's problem (1) as $(p/x)^\top c \leq 1$; then the value of the multiplier associated with this transformed constraint is $x \lambda(1, p/x, z)$. Fixing z allows us to write this as $x \lambda(p/x)$ in the obvious abuse of notation. Substituting this for λ and p/x for p into the previous equation, then taking $x = 1$, we have

$$\frac{f_i(p \lambda(p/x'))}{f_i(p \lambda(p))} = \frac{h_i(\aleph(p/x'))}{h_i(\aleph(p))}.$$

Then with a further change of variable we can write

$$(4) \quad f_i(p\lambda(p/x)) = \gamma_i(p)h_i(\aleph(p/x))$$

for $\gamma_i(p) = f_i(p\lambda(p))/h_i(\aleph(p))$.

Because λ is strictly increasing in every price p_i , there exists a function $\theta(p)$ which maps a vector of prices to a positive scalar, so that $\bar{\lambda}(\theta(p/x)) \equiv \lambda(p/x)$ for all positive (p, x) . Letting all prices be equal to one in (4), let $\aleph(\theta)$ solve

$$f_i(\bar{\lambda}(\theta)) = \gamma_i(\mathbf{1})h_i(\bar{\aleph}(\theta)),$$

where $\bar{\aleph}(\theta)$ is defined analogously to $\bar{\lambda}(\theta)$, (i.e., with $x = 1$ and all prices equal to θ), and where $\mathbf{1}$ denotes an \bar{n} -vector of ones. Since $\bar{\lambda}(\theta)$ is strictly monotone in θ , a change of variables gives

$$f_i(\bar{\lambda}) = \gamma_i(\mathbf{1})h_i(\bar{\aleph}(\bar{\lambda})),$$

establishing a one-to-one relationship between $\bar{\lambda}$ and $\bar{\aleph}$. But since $\bar{\lambda}(\theta(p/x)) \equiv \lambda(p/x)$, substitution yields

$$f_i(\lambda(p/x)) = \gamma_i(\mathbf{1})h_i(\bar{\aleph}(\lambda(p/x)));$$

rearranging and inverting implies

$$\aleph(p/x) \equiv \bar{\aleph}(\lambda(p/x)) = h_i^{-1} \left[\frac{f_i(\lambda(p/x))}{g_i(\mathbf{1})} \right],$$

giving a one-to-one relationship between $\aleph(p/x)$ and $\lambda(p/x)$.

Exploiting this one-to-one relationship, let $b_i(\lambda) \equiv h_i(\aleph)$. Substituting into (4),

$$f_i(p\lambda) = \gamma_i(p)b_i(\lambda),$$

establishing the separability of f_i .

Third, use these results to re-write the original CIRCSH condition as

$$\frac{f_i(p\lambda(x', p, z), z)}{f_i(p\lambda(x, p, z), z)} = \frac{b_i(\lambda(x', p, z))}{b_i(\lambda(x, p, z))}.$$

Fixing (\bar{x}, \bar{p}) , define $d_i(z) = f_i(\bar{p}\lambda(\bar{x}, \bar{p}, z))$, so that we have $f_i(p\lambda, z) = \gamma_i(p)b_i(\lambda)d_i(z)$, with (γ_i, b_i, d_i) all positive functions.

Finally, to establish the form of the function b_i , take logs, obtaining

$$\log f_i(e^{\log \lambda + \log p}, z) = \log \gamma_i(p) + \log d_i(z) + \log b_i(\lambda).$$

Fixing $z = z_0$, this is an example of Pexider's functional equation $g(x + y) = h(x) + \ell(y)$, and so $b_i(\lambda) = \lambda^{-\beta_i}$ (Aczél and Dhombres 1989, Corollary 10, Chapter 4) for some constant $\beta_i > 0$, implying that the demand equation must take the form (3). \square

We call the form of demands in (3) Constant Frisch Elasticity (CFE) demands, as the coefficient β_i can be interpreted as the elasticity of demand with respect to λ , a way of thinking about demands proposed by Frisch (1959). The result establishes that globally regular demands satisfying the CIRCSH condition *must* be in the CFE class. Thus, having adopted the CIRCSH condition (which seems to be necessary for us to draw inferences regarding welfare from relative consumptions when our information on prices and characteristics is incomplete) and once we've required observed demands to be globally regular (which is important if we're to measure welfare in settings with risk) our hands are essentially tied: demands must take the CFE form (3).

4. ESTIMATION WITH INCOMPLETE DATA

This paper is motivated in part by an awareness of ignorance: If the econometrician does *not* observe all relevant prices, characteristics, and quantities, then what can be learned about household welfare? Proposition 2 provides necessary and sufficient conditions for developing a welfare measure that is conditionally independent of prices and characteristics, and further dictates that the welfare measure enter (log) demands in a (log) linear fashion, with the remaining influence of prices p and characteristics z on demand for good i summarized by functions $\gamma_i(p)$ and $\delta_i(z)$, respectively. In this section we provide methods for using one or more cross-sectional samples of data on households' expenditures and characteristics to estimate key parameters of the demand system, along with household-specific estimates of welfare. In addition to the requirements of Proposition 2, we adopt three sets of identifying assumptions. Some of these are harmless normalizations or 'accounting conventions', while others amount to the standard requirement that the disturbance term be mean independent of the explanatory variables in the demand equation. There's a strong similarity to the demand problem considered by data (2006); as in that case our estimating equation is "conditionally linear", and latent variables (price effects and $\log \lambda$ in our case, household fixed effects in theirs) play a key role in delivering the mean independence we require.

Suppose we have data from $T \geq 1$ rounds of cross-sectional observations, with J_t households randomly sampled in each round t , and a total number of observations $J = \sum_{t=1}^T J_t$. A given observation is then identified by a pair (j, t) naming the household and the time at which it was observed; the observation consists of a vector of expenditures and a set of household characteristics.

Let $\mathcal{I} \subseteq \text{CIRCSH}$ denote the index set of observed expenditures with demands satisfying Condition 1, and let the cardinality of \mathcal{I} be $n \leq \bar{n}$. In addition to not observing expenditures on all goods, we must allow for the fact that we don't observe all prices and characteristics. We begin by assuming that the econometrician observes a vector of characteristics \tilde{z} , a subset of all characteristics z . Further, we adopt the following assumptions regarding the relationship between z and \tilde{z} .

Assumption 1. *Let the effect of characteristics z on demand for good $i \in \mathcal{I}$ conditional on λ and prices p be governed by the function $\delta(z)$, which maps characteristics into an n -vector. This function is related to observable characteristics \tilde{z} by*

$$(5) \quad \delta_i(z) = \alpha_i + d_i^\top \tilde{z} + \zeta_i(z) \quad i \in \mathcal{I},$$

with d_i a vector of parameters. Regarding z as a vector of random variables, we assume

- (1) $\mathbb{E}(\zeta_i(z)|\lambda, p, \tilde{z}) = 0$ and $\mathbb{E}\delta_i(z) = \alpha_i$ for $i \in \mathcal{I}$;
- (2) $\text{Var}(\delta(z)|\lambda, p, \tilde{z}) = \text{Var}(\zeta(z)) = \Xi$, with Ξ finite and positive semi-definite.

The equation (5) on its own is without loss of generality, but Assumption 1.1 asserts that the unobserved residual $\zeta_i(z)$ is mean independent of (λ, p, \tilde{z}) , so that the conditional expectation is linear in \tilde{z} . This is less restrictive than it may appear, but some assumption along these lines is necessary for identification. An example of the issue is this: suppose that $z = \tilde{z}$ is simply household size, and that larger households tend to have larger expenditures on all goods. Assumption 1.1 accounts for this by attributing the higher levels of consumption to the observed household size, not to a possible correlation between household size and (unobserved) budget. More generally, if a particular set of characteristics z has the effect of changing demand for *all* goods in the same way that a change in total budget would, then this will be accounted for by the observable characteristics, by the latent index $\lambda(x, p, z)$, and by prices, not by the residual function $\zeta_i(z)$. Assumption 1.2 allows for arbitrary patterns of covariance in $\delta_i(z)$ across demand equations, but requires that the conditional variance not depend on any of (λ, p, \tilde{z}) . This is restrictive, but standard.

We next spell out some parallel assumptions on $\log \lambda(x, p, z)$.

Assumption 2. *For any positive p and \tilde{z} , conditional expectations and variances of $\log \lambda(x, p, z)$ exist and satisfy:*

- (1) $\mathbb{E}(\log \lambda(x, p, z)|p, \tilde{z}) = \mathbb{E}(\log \lambda(x, p, z)|p) = \mu(p)$;
- (2) $\text{Var}(\log \lambda(x, p, z)|p) = \sigma^2(p)$; and

(3) $\sigma^2(\mathbf{1}) = 1$, where $\mathbf{1}$ is a vector of prices all equal to one.

Assumption 2.1 defines a conditional expectation function $\mu(p)$ of $\log \lambda$, and asserts that this conditional expectation doesn't depend on the observed \tilde{z} . As with part (1) of the previous assumption this can be thought of requiring that variation in *budgets* be captured by λ , while variation in demand due to observable characteristics should be captured by δ . Assumption 2.2 simply defines the conditional variance function $\sigma^2(p)$, while Assumption 2.3 is a harmless normalization, which sets this variance to be unity if all prices are identically equal to one.

We next address the possibility that not only do we observe just a subset of expenditures, but that the expenditures we do observe may be measured with error. Let $x_i = p_i c_i$ denote 'true' expenditures on good i . Expenditures may be observed with random measurement error; we make the following assumption regarding these measurement errors.

Assumption 3. *Observed log expenditures on good $i \in \mathcal{I}$ are given by*

$$(6) \quad y_i = \log x_i + \epsilon_i,$$

where ϵ_i is the i th element of an n -vector of random measurement errors ϵ . We assume that $\mathbb{E}(\epsilon|p, \tilde{z}, \log \lambda) = 0$, and that $\mathbb{E}\epsilon\epsilon^\top = \Sigma$ is finite and positive semi-definite.

Allowing for the possibility of measurement error is not restrictive, of course, and we permit arbitrary patterns of correlation in measurement error across different goods. The assumption that errors are mean independent of prices, observable characteristics, and $\log \lambda$ is standard; note that we do not rule out the possibility of heteroskedasticity (i.e., the covariance of ϵ across goods may vary with prices, observed characteristics, or $\log \lambda$).

Now consider the 'reduced form' system of regression equations for observations of households at different periods

$$(7) \quad \mathbf{y}_i = \mathbf{S}a_i + \mathbf{Z}d_i + \mathbf{e}_i, \quad i = 1, \dots, n,$$

where bold-faced variables are vectors or matrices with J rows (i.e., variables that change their form with sample size). Thus, \mathbf{y}_i is a vector of J observed log expenditures on good i , for all observed households in all rounds, and $\mathbf{S} = (\mathbf{s}_1, \dots, \mathbf{s}_T)$ is a

$J \times T$ matrix

$$\mathbf{S} = \begin{bmatrix} \mathbf{1}_{J_1} & 0 & \cdots & 0 \\ 0 & \mathbf{1}_{J_2} & 0 & \vdots \\ \vdots & 0 & \ddots & 0 \\ 0 & \cdots & 0 & \mathbf{1}_{J_T} \end{bmatrix},$$

where $\mathbf{1}_k$ is a column vector of k ones, and the zeros in the matrix conform. The vector $a_i = (a_{i1}, a_{i2}, \dots, a_{iT})^\top$ is a set of good-specific time effects. Interpretation of estimated parameters is simpler if \mathbf{Z} and \mathbf{S} are orthogonal, so we define $\mathbf{Z} = (\mathbf{I} - \mathbf{S}(\mathbf{S}^\top \mathbf{S})^{-1} \mathbf{S}^\top) \tilde{\mathbf{z}}$ to be the $J \times \ell$ matrix of observed characteristics in deviations from their period-by-period means. Then d_i is the ℓ -vector of parameters which appears in (5); finally \mathbf{e}_i is a J -vector of residuals. More compactly, the equation can be written as

$$\mathbf{y}_i = \mathbf{X} b_i + \mathbf{e}_i,$$

with $\mathbf{X} = [\mathbf{S}, \mathbf{Z}]$ a $J \times (T + \ell)$ matrix and $b_i = [a_i; d_i]$ a $(T + \ell)$ vector of parameters.

This forms a system of n seemingly unrelated regressions (SUR). If we concatenate these equations horizontally and stack observations vertically we can write the system for our sample in the matrix form

$$\mathbf{Y} = \underset{J \times n}{\mathbf{S}} \underset{J \times TT \times n}{\mathbf{A}} + \underset{J \times \ell \ell \times n}{\mathbf{Z}} \underset{J \times \ell \ell \times n}{\mathbf{D}} + \underset{J \times n}{\boldsymbol{\mathcal{E}}},$$

or more compactly

$$(8) \quad \mathbf{Y} = \mathbf{X} \mathbf{B} + \boldsymbol{\mathcal{E}},$$

where $\mathbf{X} = [\mathbf{S}, \mathbf{Z}]$ and $\mathbf{B} = [\mathbf{A}^\top, \mathbf{D}^\top]^\top$. The usual rank conditions are presumed to be satisfied, with $\mathbf{X}^\top \mathbf{X}$ having full rank.

4.1. Estimation Steps. Putting together our description of the demand system with our assumptions regarding unobserved characteristics and measurement error, it follows that observed log expenditures on good $i \in \mathcal{I}$ at time t for a household j with observed characteristics z_t^j will take the form of the estimating equation:

$$(9) \quad y_{it}^j = a_{it} + d_i^\top \tilde{z}_t^j + \beta_i (\mu_t - \log \lambda_t^j) + \epsilon_{it}^j + \zeta_{it}^j.$$

Estimation of (9) proceeds in two (or three) steps. The first step estimates coefficients d_i and the good-time effects (related to prices) a_{it} . The second estimates elasticities β and $\mu_t - \log \lambda_t^j$. If we have more than a single cross-section a third step can exploit this, estimating the location parameters μ_t .

Step 1: OLS. The first step yields estimates of the effects of prices and observable characteristics on demand. We estimate the $n(T + \ell)$ parameters of (8) via ordinary least squares (OLS); then $\hat{B} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{Y}$ are the estimated coefficients, and $\hat{\boldsymbol{\mathcal{E}}} = (\mathbf{I} - \mathbf{X}(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top) \mathbf{Y}$ denotes the $J \times n$ matrix of OLS residuals. Note that though there is no restriction on the covariance matrix of the disturbances across equations, random sampling implies independence across observations. Combining these assumptions with the observation that each equation in the SUR system has the same right-hand side variables, it follows that \hat{B} is the best unbiased estimator of B (Zellner 1962).

Step 2: Singular Value Decomposition. In the second step, we estimate the n -vector of Frisch elasticities β and the household IMUEs up to a location parameter. This is accomplished by computing the singular value decomposition (SVD) of the residual matrix $\hat{\boldsymbol{\mathcal{E}}}$; this provides what we need to estimate elasticities β and $w_t^j = \mu(p_t) - \log \lambda_t^j$. The SVD of $\hat{\boldsymbol{\mathcal{E}}}$ can be written

$$\hat{\boldsymbol{\mathcal{E}}} = \sum_{k=1}^n \sigma_k \mathbf{u}_k \mathbf{v}_k^\top,$$

where the σ_k are the singular values of $\hat{\boldsymbol{\mathcal{E}}}$ (ordered so that $\sigma_k \geq \sigma_{k+1}$); the \mathbf{u}_k are the corresponding left singular vectors, having length n ; and the \mathbf{v}_k are the corresponding right singular vectors, with length J . Some important and well-known properties of the SVD are that $\mathbf{u}_k^\top \mathbf{u}_k = \mathbf{v}_k^\top \mathbf{v}_k = 1$ for $k = 1, \dots, n$, and that $\mathbf{u}_k^\top \mathbf{u}_{k'} = \mathbf{v}_k^\top \mathbf{v}_{k'} = 0$ for all $k \neq k'$; that is, the \mathbf{u}_k and \mathbf{v}_k are orthonormal vectors.

Calculation of the SVD is a standard basic matrix operation. With the SVD in hand, we simply let $\hat{\beta} = \sigma_1 \mathbf{v}_1 \phi$, the largest singular value of $\hat{\boldsymbol{\mathcal{E}}}$ multiplied by its corresponding right singular vector and a non-zero constant ϕ ; and let $\hat{\mathbf{w}} = \mathbf{u}_1 / \phi$ denote the corresponding left singular vector divided by the same constant. Then we have

$$\hat{\boldsymbol{\mathcal{E}}} = \hat{\mathbf{w}} \hat{\beta}^\top + \sum_{k=2}^n \sigma_k \mathbf{u}_k \mathbf{v}_k^\top = \hat{\mathbf{w}} \hat{\beta}^\top + \boldsymbol{\varepsilon}.$$

Note that the property of orthonormality then implies that $\mathbf{w}^\top \boldsymbol{\varepsilon} = 0$. By the Eckart-Young theorem the outer product $\hat{\mathbf{w}} \hat{\beta}^\top$ is the rank-one matrix closest to $\hat{\boldsymbol{\mathcal{E}}}$ in a least-squares sense (i.e., under the Frobenius norm). We exploit both this and the orthonormality property of the left and right singular vectors in the following proposition to develop a useful property of the estimator $(\hat{\beta}, \hat{\mathbf{w}})$.

Proposition 3. *If demands take the CFE form (3), at every time $t = 1, \dots, T$ all households face the same prices p_t , and Assumptions 1–3 are satisfied, then $\mathbb{E}\hat{\beta}_i \hat{w}_t^j = \beta_i w_t^j$ for all $i \in \mathcal{I}$, and all households $j = 1, \dots, J_t$.*

Proof. Consider the conditional expectation of y_i given (\tilde{z}, p, λ) ; this is

$$\begin{aligned} \mathbb{E}(y_i|\tilde{z}, p, \lambda) &= \mathbb{E}(\log p_i + \log c_i + \epsilon_i|\tilde{z}, p, \lambda) \\ &= \mathbb{E}(\log p_i + \log c_i|\tilde{z}, p, \lambda) \quad \text{by Assumption 3;} \\ &= \mathbb{E}(\log p_i + \gamma_i(p) + \delta_i(z) - \beta_i \log \lambda|\tilde{z}, p, \lambda) \quad \text{from (3)} \\ &= \log p_i + \gamma_i(p) + \alpha_i + d_i^\top \tilde{z} - \beta_i \log \lambda \quad \text{by Assumptions 1 \& 2.} \end{aligned}$$

Define $a_i(p) = \alpha_i + \log p_i + \gamma_i(p) - \beta_i \mu(p)$. Then we have

$$\mathbb{E}(y_i|\tilde{z}, p, \lambda) = a_i(p) + d_i^\top \tilde{z} + \beta_i(\mu(p) - \log \lambda).$$

Now, consider the expected value of y_i conditional on the information set (\tilde{z}, p) ; by the law of iterated expectations and Assumption 2 we obtain

$$\mathbb{E}(y_i|\tilde{z}, p) = a_i(p) + d_i^\top \tilde{z}.$$

By assumption at time t everyone faces the same prices p_t , so we can write $a_{it} = a_i(p_t)$, with $a_i = (a_{i1}, \dots, a_{iT})$ corresponding to the same vector as in (7), and

$$\mathbb{E}(y_i|\tilde{z}, a_i) = a_i + d_i^\top \tilde{z}.$$

The vector of errors in (7) is

$$\mathbf{e}_i = \mathbf{y}_i - \mathbb{E}(\mathbf{y}_i|\mathbf{Z}, \mathbf{a}) = \mathbf{y}_i - \mathbf{a}_i - d_i^\top \mathbf{Z} = \mathbf{w}\beta_i + \boldsymbol{\epsilon}_i + \boldsymbol{\zeta}_i,$$

where \mathbf{w} is a J -vector consisting of elements $w_t^j = \mu(p_t) - \log \lambda_t^j$, and where $\boldsymbol{\epsilon}_i$ and $\boldsymbol{\zeta}_i$ are J -vectors of measurement errors ϵ_{it}^j and $\zeta_i(z_t^j)$.

OLS provides an unbiased estimate of the matrix of coefficients B , so that $\mathbb{E}(\hat{\boldsymbol{\mathcal{E}}|\boldsymbol{\mathcal{E}}) = \boldsymbol{\mathcal{E}}$. Conditioning on \mathbf{w} gives

$$\mathbb{E}(\hat{\boldsymbol{\mathcal{E}}|\mathbf{w}) = \mathbb{E}(\boldsymbol{\mathcal{E}}|\mathbf{w}) = \mathbf{w}\beta^\top.$$

The SVD of the matrix of OLS residuals yields

$$\hat{\boldsymbol{\mathcal{E}}} = \hat{\mathbf{w}}\hat{\beta}^\top + \hat{\boldsymbol{\epsilon}},$$

with $\hat{\mathbf{w}}^\top \hat{\boldsymbol{\varepsilon}} = 0$ by construction, so

$$(10) \quad \mathbb{E}[\hat{\mathbf{w}}\hat{\boldsymbol{\beta}}^\top + \hat{\boldsymbol{\varepsilon}}|\mathbf{w}] = \mathbb{E}[\hat{\mathbf{w}}\hat{\boldsymbol{\beta}}^\top|\mathbf{w}] = \mathbf{w}\boldsymbol{\beta}^\top.$$

□

Thus, using OLS to estimate (7) and then calculating the SVD of the OLS residuals gives an unbiased estimate of the outer product of elasticities $\boldsymbol{\beta}$ and welfare measures \mathbf{w} up to the scalar constant ϕ . Our choice of ϕ amounts to a choice of units, and is consequently somewhat arbitrary. One natural choice would be to take $\phi = 1$; this would then imply the estimated variance of w (pooled across all observations) $\hat{\sigma}^2 = J^{-1}\hat{\mathbf{w}}^\top\hat{\mathbf{w}} = 1$. However, our Assumption 2.2 implies a different normalization, which sets the estimated variance of w in period one to unity. This normalization implies that $\phi = \sqrt{J_1^{-1} \sum_{j=1}^{J_1} u_{11}^j{}^2}$ is the maximum likelihood estimate of the standard deviation of the elements of the left-singular vector \mathbf{u}_1 observed in the first period. We adopt this normalization henceforth.

Some remarks are in order. First, while the outer product $\hat{\mathbf{w}}\hat{\boldsymbol{\beta}}^\top$ is an unbiased estimator of $\mathbf{w}\boldsymbol{\beta}^\top$, $\hat{\boldsymbol{\beta}}$ and $\hat{\mathbf{w}}$ are not unbiased estimators of $\boldsymbol{\beta}$ and \mathbf{w} . However, $\hat{\boldsymbol{\beta}}$ is \sqrt{J} -consistent, with the immediate consequence that $\hat{\mathbf{w}}$ is asymptotically unbiased. Second, in the homoskedastic Gaussian case in which $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2 I)$ and $\zeta(z) \sim \mathcal{N}(0, \sigma_\zeta^2 I)$ these are the maximum likelihood estimators (MLE), and so are efficient (Anderson 1963; Tipping and Bishop 1999). We choose to not adopt these distributional assumptions, but the estimators $\hat{\boldsymbol{\beta}}$ and $\hat{\mathbf{w}}$ can nevertheless be regarded as quasi-MLE estimators.

Step 3: Auxiliary Regression. In step one of the estimation procedures we obtained estimates of a_{it} and d_i ; in step two we obtained estimates of elasticities $\boldsymbol{\beta}$ and $w_t^j = \mu(p_t) - \log \lambda_t^j$. But to identify $\log \lambda_t^j$ we need an estimate of the location parameter $\mu(p_t)$. If $T = 1$ this is no problem, because $\mu(p_t)$ is just a constant we can choose via some normalization (e.g., set it to zero).

However, if we have $T > 1$ and we wish to be able to make welfare comparisons across periods then we must estimate the different $\mu(p_t)$ for $t = 1, \dots, T$. Our approach to this begins by recalling from the proof of Proposition 3 that $\mu(p_t)$ appears not only in w , but also in each of the terms a_{it} for $i = 1, \dots, n$, with

$$(11) \quad a_{it} = \alpha_i - \beta_i \mu(p_t) + \log p_{it} + \gamma_i(p_t).$$

We re-write the terms involving prices as $-\beta_i \mu(p_t) + \log p_{it} + \gamma_i(p_t) = -\beta_i \mu(p_t) + \log \bar{p}_t + r_i(p_t)$, where \bar{p}_t serves as a measure of the price level, and where the $r_i(p_t)$ are

then a measure of the log relative price of good i within period t . This then gives us

$$(12) \quad a_{it} = \alpha_i - \beta_i \mu_t + \log \bar{p}_t + r_{it},$$

where we exploit the assumption of common prices to define $\mu_t = \mu(p_t)$ and $r_{it} = r_i(p_t)$. Since we can choose the units in which different prices are measured we're free to adopt the following normalizations: (i) $\bar{p}_1 = p_{i1} = 1$ for all i ; (ii) $r_i(\mathbf{1}_n) = 0$ for all i ; (iii) $\mu_1 = 0$; and finally we require (iv) that $r_i(p_t)$ be uncorrelated with α_i , β_i , and $\log \bar{p}_t$. Beyond these defining restrictions we assume that prices may be unobserved by the econometrician. But even without prices these normalizations and restrictions identify or define all the unknown quantities on the right-hand side of (12). The term α_i is a good-specific preference parameter; μ_t is the mean value of $\log \lambda_t^j$ at time t , and can be thought of as a measure of aggregate welfare at time t ; \bar{p}_t is the general price level at time t ; and r_{it} is the log of the relative price of good i at time t .

We then construct an auxiliary regression. From step one we have nT estimates \hat{a}_{it} of the left-hand-side variable a_{it} as well as estimates \hat{V} of the $nT \times nT$ covariance matrix of the \hat{a}_{it} . From step two we have estimates $\hat{\beta}_i$ of the elasticities that appear on the right-hand side of (12). Then regarding this as a two-way error component balanced panel model, we treat (12) as a regression of our previously estimated good-time effects on n good fixed effects, T time effects, and on the generated regressors $\hat{\beta}_i$ interacted with a collection of time effects, or

$$(13) \quad \hat{a}_{it} = \alpha_i + \eta_t + \mu_t \hat{\beta}_i + r_{it}, \quad \text{with } \sum_t \eta_t = 0.$$

This regression could be estimated via OLS, which would be unbiased. But if $\text{Var}(\hat{a}) \neq \sigma_a^2 I$ then the residuals r_{it} will be heteroskedastic, and OLS inefficient. A simple feasible GLS estimator is available, since we already have estimates of the covariance matrix of the \hat{a} .

So, our step three involves two within transformations W_i and W_t which respectively eliminate the good and time effects in (13) (incidentally yielding estimates of the preference parameters α and the log of the price level), and estimates

$$\hat{\mu} = (\tilde{X}^\top \hat{V}^{-1} \tilde{X})^{-1} \tilde{X}^\top \hat{V}^{-1} \tilde{a},$$

where $\tilde{X} = W_i W_t I_T \otimes \hat{\beta}$ and $\tilde{a} = W_i W_t a$. This feasible GLS estimator has the usual properties of being consistent and asymptotically efficient.

4.2. 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 elasticities and $\log \lambda$ s 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 expenditures are reported to be 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 hold (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 Appendix A, 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). Roughly, just as fixed effects addresses selection if selection conditional on the fixed effects is random, in our case if data is missing at random conditional on $(a, \tilde{z}, \log \lambda)$ then selection won’t be an issue.

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 (7), and compute residuals from this regression. We then use these residuals to augment the matrix of residuals from the first step estimation of (7), and use a singular value decomposition of the resulting $n \times 2J$ matrix to obtain alternative estimates $\hat{\beta}^r$ and $\hat{\mathbf{w}}^r$ of elasticities and the household specific welfare measures \mathbf{w} . If selection is an issue then these estimated values will differ from the values obtained from the decomposition of the unaugmented residual matrix.

5. DATA

We use the methods described above to measure households’ welfare in Uganda over an interval of time which encompasses the 2008 “food price crisis” and a recession in 2010–11. We have data from four rounds of the World Bank’s *Living Standards Measurement Surveys* (LSMSs) conducted in Uganda (in 2005–06, 2009–10, 2010–11, and 2011–12).⁷ For the sake of concision we will sometimes refer to the 2005–06 round as “2005”, and so on. We first give a descriptive account of some of the data on household characteristics and expenditures from these surveys.

TABLE 1. Mean characteristics of households in Uganda by year. Figures in parentheses are standard deviations.

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

Table 1 provides information on the household characteristics we employ in our application. In each of four rounds, there are about 3000 households; of these, between 70–80% are rural. 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.⁸ Household members eighteen years or younger are taken to be boys or girls. Excluding durables, taxes, and “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.⁹

7. These datasets are public use datasets available from the World Bank.

8. 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.

9. There are some minor discrepancies in food items across periods; see Appendix B for details and the way in which we’ve addressed these discrepancies.

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 selected 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. To avoid issues with different expenditures being elicited with different recall periods (Tarozzi 2007) in our estimation we restrict attention to goods with seven day recall period, which effectively restricts attention to different kinds of food.

Appendix B.1 reports 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. Aggregate shares 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, save 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 fairly stable over time, for some goods mean and aggregate shares differ dramatically from each other. On its face this is not 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 ,

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

We then produce a scatterplot of this statistic, ordered by the size of the statistic in 2005. Thus, each good i (labelled on the left axis) has an associated statistic ρ_{it} for each period t , 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 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. Such goods include beer, soda,

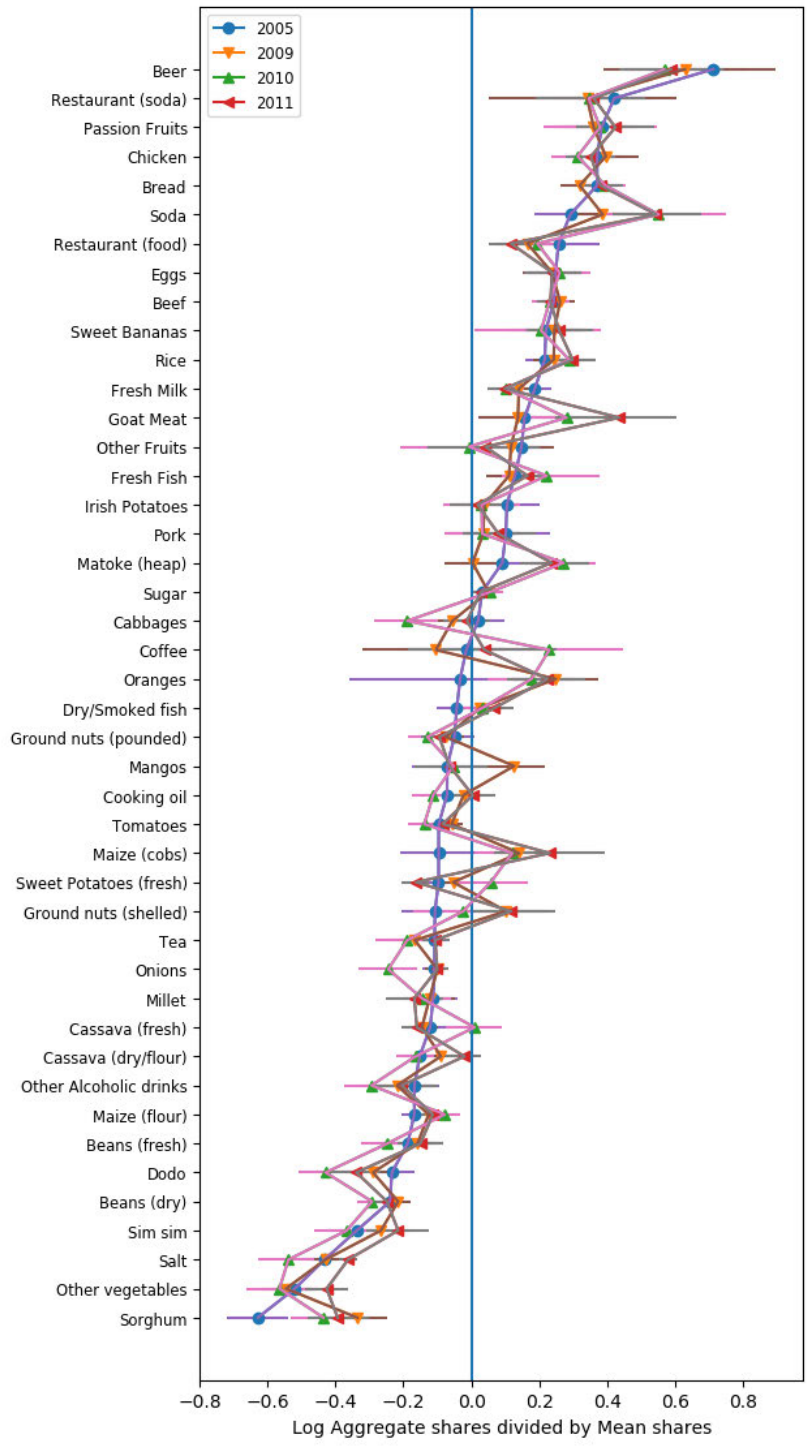


FIGURE 1. Log of aggregate shares minus log of mean shares for different years (ordered by ranking in 2005), with 95% confidence intervals.

passion fruits, chicken, bread, spending at restaurants, eggs, and beef, among others. Conversely, when the statistic is negative we identify goods that are particularly important in the portfolios of households with lower food expenditure. Here we see sorghum, “other vegetables”, salt, sim sim (sesame) and dry beans.

Figure 1 and this *prima facie* case against homotheticity implies that the composition of households’ consumption portfolios will vary systematically with total budget. This kind of variation is exactly what we wish to use to draw inferences about welfare even absent information on all expenditures.

6. RESULTS

6.1. Estimates of Demand Elasticities. We now turn to estimates of some of the parameters of the the demand system (9), 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 indicator which takes the value one for rural households. This allows for possible differences in the the α_i taste parameters in equations (5) across rural and urban households (i.e., α_i for urban households; $\alpha_i + \delta_{i,rural}$ for rural households).

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

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

Continued on next page

Continued from previous page

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

Continued on next page

Continued from previous page

	β_i	α_i	Rural	Boys	Girls	Men	Women	log Hsize
Dodo	(0.09) 0.46***	(0.26) 5.67***	(0.09) -0.10***	(0.04) 0.06***	(0.04) 0.04***	(0.05) 0.05***	(0.05) 0.07***	(0.16) 0.08
Beef	(0.04) 0.46***	(0.14) 7.54***	(0.03) -0.18***	(0.01) 0.03***	(0.01) 0.02*	(0.02) 0.11***	(0.02) 0.07***	(0.06) 0.19***
Onions	(0.02) 0.45***	(0.10) 5.21***	(0.02) -0.40***	(0.01) -0.01*	(0.01) -0.00	(0.01) 0.08***	(0.01) 0.07***	(0.05) 0.14***
Cassava (fresh)	(0.02) 0.45***	(0.10) 6.39***	(0.02) 0.22***	(0.01) 0.06***	(0.01) 0.04***	(0.01) 0.08***	(0.01) 0.01	(0.03) 0.28***
Ground nuts (pounded)	(0.03) 0.45***	(0.14) 6.26***	(0.03) -0.02	(0.01) 0.04***	(0.01) 0.02*	(0.02) 0.09***	(0.02) 0.09***	(0.06) 0.09*
Fresh Fish	(0.04) 0.45***	(0.14) 7.08***	(0.03) -0.11***	(0.01) 0.04**	(0.01) 0.02	(0.02) 0.12***	(0.02) 0.06**	(0.05) 0.12*
Restaurant (food)	(0.03) 0.45***	(0.14) 8.80***	(0.03) -0.58***	(0.02) 0.05**	(0.02) 0.03	(0.02) 0.19***	(0.02) 0.10***	(0.07) -0.41***
Other Alcoholic drinks	(0.05) 0.44***	(0.21) 7.31***	(0.04) -0.37***	(0.02) 0.04*	(0.02) 0.01	(0.03) 0.17***	(0.03) 0.03	(0.07) -0.06
Pork	(0.07) 0.43***	(0.21) 7.25***	(0.07) -0.31***	(0.02) 0.04*	(0.02) 0.02	(0.03) 0.09***	(0.03) 0.06**	(0.09) 0.16
Dry/Smoked fish	(0.05) 0.41***	(0.18) 6.53***	(0.06) -0.17***	(0.03) 0.03*	(0.02) 0.02	(0.03) 0.12***	(0.03) 0.05**	(0.10) 0.19**
Cabbages	(0.04) 0.40***	(0.17) 5.92***	(0.03) -0.14***	(0.02) 0.00	(0.02) 0.02	(0.02) 0.05***	(0.02) 0.03**	(0.07) 0.21***
Other vegetables	(0.03) 0.40***	(0.12) 5.66***	(0.03) -0.05	(0.01) 0.04**	(0.01) 0.06***	(0.02) 0.05**	(0.02) 0.07***	(0.06) 0.13*
Matoke (heap)	(0.04) 0.39***	(0.14) 7.64***	(0.03) -0.25***	(0.02) 0.00	(0.02) -0.00	(0.02) 0.04	(0.02) 0.11***	(0.07) 0.33***
Maize (flour)	(0.07) 0.39***	(0.21) 6.62***	(0.06) 0.07**	(0.03) 0.08***	(0.03) 0.08***	(0.03) 0.09***	(0.04) 0.05***	(0.13) 0.20***
Restaurant (soda)	(0.03) 0.37***	(0.13) 7.53***	(0.03) -0.29***	(0.01) 0.00	(0.01) 0.04	(0.02) 0.08**	(0.02) 0.03	(0.05) -0.11
	(0.07) 0.37***	(0.21) 7.53***	(0.07) -0.29***	(0.03) 0.00	(0.03) 0.04	(0.03) 0.08**	(0.04) 0.03	(0.10) -0.11

Continued on next page

Continued from previous page

	β_i	α_i	Rural	Boys	Girls	Men	Women	log Hsize
Millet	0.36*** (0.05)	6.11*** (0.18)	0.15*** (0.05)	0.03 (0.02)	0.00 (0.02)	0.05* (0.02)	0.05* (0.03)	0.40*** (0.09)
Chicken	0.34*** (0.04)	8.09*** (0.13)	-0.23*** (0.04)	-0.01 (0.02)	-0.02 (0.02)	0.07*** (0.02)	0.04 (0.02)	0.20** (0.09)
Sweet Potatoes (fresh)	0.34*** (0.03)	6.49*** (0.13)	0.32*** (0.03)	0.07*** (0.01)	0.06*** (0.01)	0.07*** (0.02)	0.06*** (0.02)	0.32*** (0.05)
Beans (dry)	0.33*** (0.02)	6.50*** (0.11)	-0.01 (0.02)	0.03*** (0.01)	0.04*** (0.01)	0.08*** (0.01)	0.04*** (0.01)	0.29*** (0.04)
Sorghum	0.33*** (0.07)	5.99*** (0.20)	0.02 (0.07)	-0.03 (0.03)	-0.07** (0.03)	-0.08** (0.03)	-0.09** (0.04)	0.75*** (0.12)
Sim sim	0.33*** (0.05)	6.09*** (0.17)	0.02 (0.06)	0.01 (0.02)	-0.01 (0.02)	0.09*** (0.03)	0.06** (0.03)	0.13 (0.11)
Tea	0.31*** (0.02)	4.59*** (0.09)	-0.19*** (0.02)	0.02* (0.01)	0.03*** (0.01)	0.12*** (0.01)	0.09*** (0.01)	0.10*** (0.04)
Salt	0.18*** (0.01)	4.18*** (0.07)	0.04*** (0.01)	0.03*** (0.01)	0.02*** (0.01)	0.05*** (0.01)	0.03*** (0.01)	0.24*** (0.02)
Cassava (dry/flour)	0.11** (0.04)	6.66*** (0.19)	0.14*** (0.05)	0.03* (0.02)	0.00 (0.02)	0.05** (0.02)	-0.04* (0.03)	0.53*** (0.08)

In its first column Table 2 presents estimates of the Frisch elasticities β_i in descending order. It's possible to show that the product of the 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 consistent with informal descriptions of what goods are more desired by Ugandan consumers.

All estimated elasticities 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).¹⁰ The overall fit of the demand system measured via the R^2 statistic is 0.65, which compares favorably to the fit achieved by other full-rank demand systems such as the EASI demand system of Lewbel and Pendakur (2009), and is considerably better than the fits of the AIDS, QUAIDS, or Rotterdam demand systems estimated by Decoster, Vermeulen, et al. (1998).

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 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). A corollary is that any welfare measure that uses only one price index (such as ‘real’ total expenditures) cannot account properly for the welfare effects of changes in relative prices.

The second column of Table 2 gives estimates of α_i , where α_i is the preference parameter introduced in Assumption 1. With homothetic utility (i.e., $\beta_i = \beta$ for all i) the α_i parameters would be proportional to the expenditure share of good i , and Frisch, income, and price 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 α_i vary positively with expenditure share, and are normalized so as to be 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

10. 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.

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 (both fresh and dried), millet, and salt expenditures are all significantly greater for rural households, other things equal.

The next several columns 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, 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, and coffee. Similarly, adding a woman 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. Validation: Estimated Aggregate Shares versus Mean Shares. In Figure 1 we used data on observed expenditures 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 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

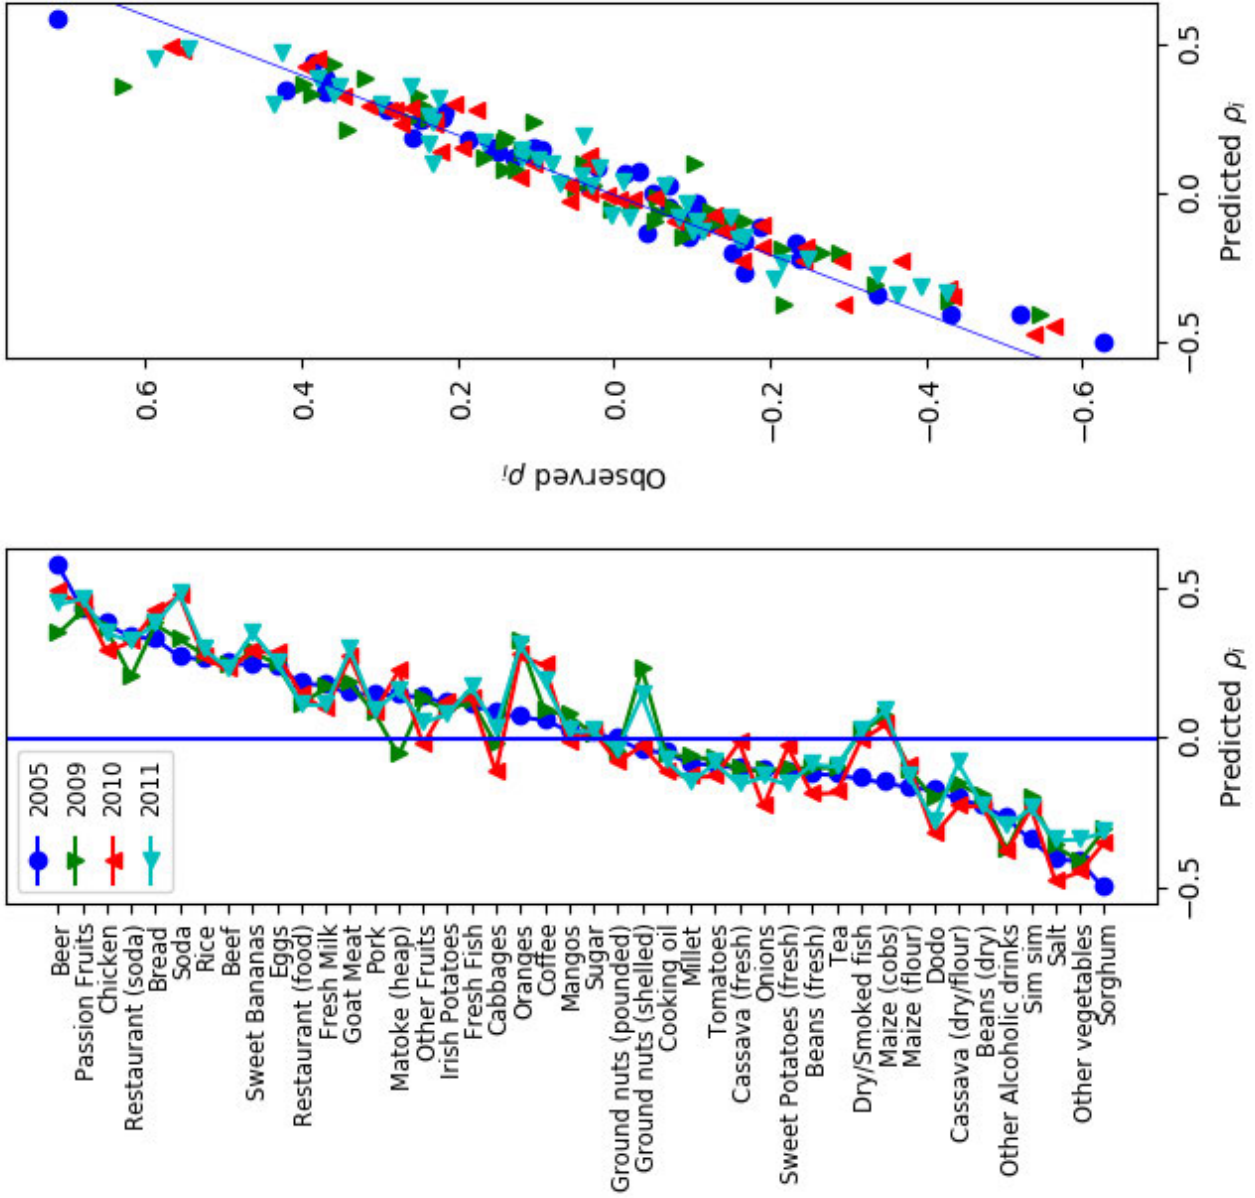


FIGURE 2. Left panel: Predicted log of mean shares minus log of aggregate shares for different years (ordered by ranking in 2005). Right panel: Predicted versus observed, with 45 degree line.

expectations of log expenditures to the data, while the shares statistic we’ve constructed is built using logs of means and sums of expenditures. Jensen’s inequality alone implies that our ability to match the share statistics will depend not only on the estimated equation, but also on the distribution of residuals.

Let $h_i(p, \lambda, z) = \mathbb{E}(\log x_i | p, \lambda, z)$. Assume that the residuals e_{it}^j in the estimating equation (7) are independent and identically normally distributed for each good, with mean zero and variance σ_i^2 . Then a simple estimator of $\mathbb{E}(x_{it}^j | p_t, \lambda_t^j, z_t^j)$ 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 λ as described in Section 4.

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 2.

The left panel of Figure 2 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 2 provides a scatter plot of observed vs. predicted values of the statistic, along with a 45 degree line. The scatterplot confirms the remarkable success of our demand system at reproducing even patterns in the data that our estimator wasn’t designed to fit.

6.3. Validation: Testing for Selection Bias from Missing Data. 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 3 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 β_i obtained in Table 2, versus the same vector obtained when we augment the residuals in (7) with the residuals from a regression of a missing indicator on the same right-hand side variables. Our estimates of elasticities are basically unchanged, save for small difference in scale. 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

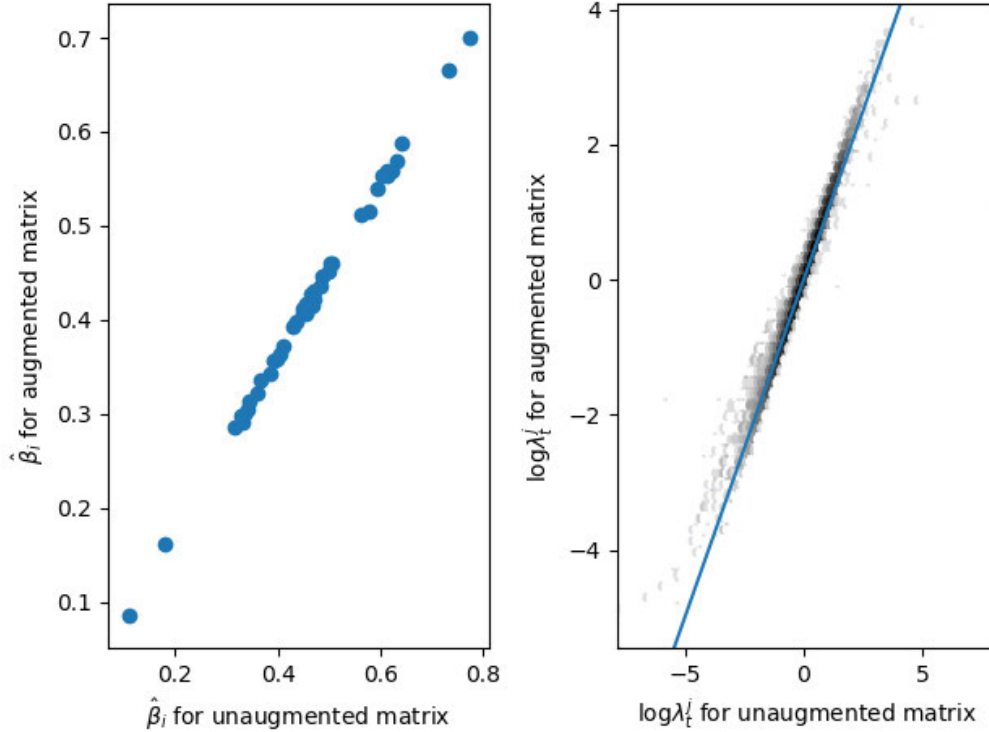


FIGURE 3. Test for selection bias. Scatterplots of β_i and $\log \lambda$ estimated with and without information on missing values.

no discernable evidence of selection bias, and almost no change in orderings (the Spearman correlation coefficients are respectively 1.00 and 0.99).

7. ESTIMATES OF $\log \lambda$ AND WELFARE EFFECTS OF THE FOOD PRICE CRISIS

In this section we give a characterization of the values of the IMUEs ($\log \lambda$) we’ve obtained for every observed household. One convenient way to look for variation in these measures is to see how the distribution of these varies across years. As it happens there were some large aggregate shocks within the timeframe covered by our sample, with an approximate doubling of food prices in 2008 (the “food price crisis”, Abbott and Borot de Battisti (2011)), and a subsequent recession. We measure changes in the distribution of $\log \lambda$ in the face of these shocks, and discuss the consequences for rates of poverty. We find that both shocks are associated with large increases in poverty rates. While perhaps this does not seem surprising, it is sharply at odds with conventional approaches to the measurement of poverty, which interpret sharply increased food expenditures as evidence of *improved* welfare. Going beyond

measurement of poverty rates, we find evidence that these aggregate shocks are also associated with increases in households' relative risk aversion.

7.1. 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 (aggregate and mean shares are different for the same good; 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. An analysis of the episode of the 2008 food price crisis shows the severe problems the conventional approach can also have in practice.

Figure 4 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 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 2005.¹¹ Of particular note is that of most staple starches (maize, millet, potatoes, sweet potatoes, cassava) had nominal median 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 nominal food prices doubled from 2005–2009.

If nominal food prices doubled, what happened to nominal expenditures on food? These increased, as one would expect for a necessity, but by only 58% for the average household. Thus, either *quantities* of food consumed fell by roughly 40% or there

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




FIGURE 4. Distribution of $\log \lambda$ by Year.

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 which experienced only modestly higher prices) 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 0.29 standard deviations in the distribution of $\log \lambda$ seen in Figure 4. 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 29% reduction in welfare for this population, followed by a two percent improvement in 2010, with no subsequent change in the mean of the distribution in 2011.

Prices go up by more than expenditures from 2005 9; 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, which simply deflate 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.

Relation of $\log \lambda$ to poverty measures. 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 (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.¹² 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 about the same magnitude as the change in $\log \lambda$, but in the wrong direction!

Consider the effects of this on measures of headcount poverty; these are summarized in Figure 5. 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 do we, 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 (determined by the fact that 55.4% of the population has values of $\log \lambda$ which exceed -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 69% in 2009. By 2010 the mean of the distribution of $\log \lambda$ shifts back by 13%, almost back to the level of the base year, and the headcount poverty rate falls to 56%. The subsequent effects of a recession (Brunori, Palmisano, and Peragine 2015) slightly

12. This broad description omits important details; see Uganda Office of the Auditor General (2014).




FIGURE 5. Changes in distributions of $-\log \lambda$ and deflated log expenditures over time, with implied headcount poverty statistics.

shifts and changes the shape of the distribution, resulting in an increase in headcount poverty to 59%.

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 recession of 2010-11 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*.

7.2. Measuring Heterogeneity in Risk Aversion. In this section we explore the connection between λ and total expenditures. The IMUE λ is decreasing in total expenditures x by construction, but it’s also of interest to know the rate at which it’s decreasing. This leads to consideration of the expenditure elasticity of λ , or $\omega = \partial \log \lambda / \partial \log x$. This is what Frisch called the household’s “money flexibility,” and provides information on the responsiveness of the welfare measure λ to changes in expenditures. Under a cardinalization of utility which makes λ the MUE then $-\omega$ can also be interpreted as relative risk aversion.

In a model which is nearly a special case of ours, Chiappori et al. (2014) adopt this cardinalization and estimate heterogeneous relative risk aversion using data from Thailand. In addition to the cardinalization they adopt two key identifying assumptions. First, they assume that for a given household j all good elasticities $\beta_i = \beta^j$ are common, so that preferences are homothetic (but can vary across households). Second, they assume that there’s full insurance, so that a given household’s λ is fixed over time.

In this section we show that it’s possible to estimate risk aversion under generally weaker assumptions than those adopted by Chiappori et al. (2014). That paper assumes that risk attitudes differ across households because homothetic preferences differ across households. In contrast, we assume that households have common preferences, but risk aversion varies with resources because preferences are non-homothetic. We entirely drop the very strong assumption of full insurance. If we maintain the Chiappori et al. (2014) cardinalization then we can estimate the distribution of relative aversions; if we weaken the cardinalization assumption to allow a flexible parametric class of regular transformations of momentary utility we can estimate the distribution of relative risk aversions up to a location parameter.

We begin by exploring the relationship between ω and λ , where no cardinalization is required. The relationship is illustrated in Figure 6, where λ is on the horizontal axis. The central plotted line offers calculations of ω given the point estimates of α_i and the β_i elasticities given in Table 2; 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




FIGURE 6. Calculated relative risk aversions versus λ . Varying all estimates of β_i by plus or minus two standard errors yields the shaded area.

see that this elasticity is decreasing in λ , and so increasing in total expenditures. The figure illustrates the point that the utility of wealthier households is less sensitive to variation in total food expenditures than is the utility of poorer households.

Knowing ω is what we need for estimating within-period Engel curves. If we adopt the same cardinalization as Chiappori et al. (2014) then λ is the MUE. If preferences are also von Neumann-Morgenstern, then $-\omega$ can also be interpreted as the household's relative risk aversion. If in addition utility is time separable with exponential discounting, then its negative reciprocal is the elasticity of intertemporal substitution. With the cardinalization, the elasticity ω gives us a link between λ and risk aversion (or more generally the curvature of the momentary utility function). However, without the cardinalization this isn't enough — the purely cross-sectional demand behavior we observe which identifies λ and ω 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.

Let us replace the cardinalization assumption with the weaker assumption that all households have momentary utility identified up to a common regular twice

continuously-differentiable transformation M . Then the “true” (momentary) indirect utility function is not $V(p, x)$, but $V^*(p, x) = M(V(p, x))$. 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 risk aversion, or intertemporal substitution.

However, with modest additional assumptions it’s possible to estimate the *distribution* of households’ relative risk aversions, up to an unknown location parameter. Recall that the Arrow-Pratt measure of relative risk aversion for a household with indirect utility $V^*(x, p)$ is given by

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

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

$$\text{RRA}(p, x) = -\frac{\partial \log \lambda(p, x)}{\partial \log x} - x \frac{M''}{M'} \lambda(x, p)$$

The quantities in the first term on the right-hand side are the relative risk aversions shown in Figure 7. The second term involves the first and second derivatives of the unknown transformation M . A judicious parameterization of M 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

$$\text{RRA}(p, x) \approx \sigma - \omega(p, x).$$

This then allows us to identify households’ relative risk aversions up to the unknown constant σ .

Results of calculating households’ values of relative risk aversion in different years are shown in the densities estimated in Figure 7. The global mean is 1.83, with a standard deviation of 0.05. However, there are significant shifts in the distribution over time, corresponding to shifts in the distribution of the $\log \lambda$. In 2005, the mean is 1.82. The increase in IMUEs in 2009 is associated with a statistically significant increase in mean relative risk aversion to 1.84; a (significant) fall in 2010 to 1.832; and finally rather large shift during the recession year of 2011 to 1.85. Households

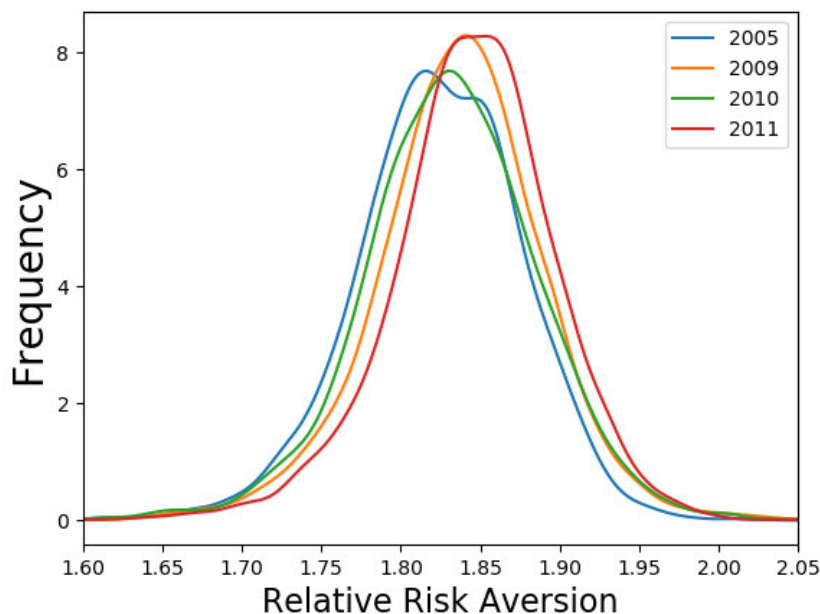


FIGURE 7. Distribution of estimated relative risk aversion in different years. Standard deviation of pooled estimates is 0.09; kurtosis is 0.22.

are assumed to have identical β_i parameters and to face identical prices, so differences in risk aversion within a period are driven entirely by differences in $\log \lambda$; differences in risk aversion across periods are due to a combination of differences in $\log \lambda$ and changes in relative prices. The estimated values of relative risk aversion in Figure 7 are well within the range of plausible relative risk aversions.

The finding that households in Uganda have heterogeneous relative risk aversions echoes the findings of Chiappori et al. (2014) for households in Thailand. However, their identification strategy assumes homothetic utility and relies 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 relative risk aversions (up to an unknown location) are just the much weaker requirements that elasticities are constant (but may vary across goods) and that the household maximizes utility within the period subject to a budget constraint.

8. 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. Rather than the usual

approach of trying to measure total expenditures, we instead focus on the information available from the *composition* of a household’s consumption portfolio.

We establish that if differences in household welfare can be inferred from differences in relative consumption of goods across otherwise similar households, then any globally regular demand system must take a particular semi-parametric form we call a Constant Frisch Elasticity (CFE) demand system.

The CFE demand system features highly flexible Engel curves, and a measure of household welfare which is closely related to the household’s marginal utility of expenditures. We devise a simple estimator which allows us to estimate both critical demand parameters as well as household-specific measures of welfare we call $\log \lambda$.

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. In particular, our methods are designed for use by an ignorant econometrician: they will work even if we observe only expenditures on some goods, and some household characteristics. Observing total expenditures and prices is entirely unnecessary.

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 demands for 44 different kinds of food. We estimate both household $\log \lambda$ and Frischian elasticity parameters from this expenditure system, in addition to other demand parameters. We use the estimated distributions of $\log \lambda$ over time to characterize changes in the poverty rate related to the 2008 food price crisis, and to a 2010–11 “great recession.” Our results contrast sharply with other accounts, which, given evidence of increased food expenditures, concluded that poverty had fallen sharply during both episodes. A separate analysis allows us to characterize the distribution of households’ relative risk aversions. This distribution is identified only up to a location parameter, but we find convincing evidence of heterogeneity, and evidence that both the food price crisis and the 2010–11 recession increased not only rates of poverty, but also increased the average households’ sensitivity to risk variation in food expenditures, as measured by relative risk aversion.

9. REFERENCES

- Abbott, Philip, and Adeline Borot de Battisti. 2011. Recent Global Food Price Shocks: Causes, Consequences and Lessons for African Governments and Donors. *Journal of African Economies* 20 (Supplement): i12 i62.
- Aczél, J., and J. Dhombres. 1989. *Functional equations in several variables*. Vol. 31. Encyclopedia of Mathematics and its Applications. Cambridge University Press.
- 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.
- Anderson, Theodore Wilbur. 1963. Asymptotic theory for principal component analysis. *Annals of Mathematical Statistics* 34 (1): 122 148.
- 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 (): 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.
- 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, R., M. Browning, and C. Meghir. 1994. Consumer demand and the life-cycle allocation of household expenditures. *Review of Economic Studies* 61:57–80.
- 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.
- . 1999. *Modelling commodity demands and labour supply with m-demands*. Discussion paper 99-08. University of Copenhagen.
- 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.
- Calvi, Rossella. Forthcoming. Why are older women missing in India? The age profile of bargaining power and poverty. *Journal of Political Economy*.

- 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).
- Chiappori, Pierre-André, Krislert Samphantharak, Sam Schulhofer-Wohl, and Robert M. Townsend. 2014. Heterogeneity and risk sharing in village economies. *Quantitative Economics* 5, no. 1 (): 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.
- Costa, Dora L. 2001. Estimating real income in the United States from 1888 to 1994: correcting CPI bias using Engel curves. *Journal of political economy* 109 (6): 1288–1310.
- data, Estimating demand response with panel. 2006. Sebastien lecocq and jean-marc robin. *Empirical Economics* (): 1043–1060.
- 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.
- Decoster, André, Frederic Vermeulen, et al. 1998. *Evaluation of the empirical performance of two-stage budgeting aids, quads and rotterdam models based on weak separability*. Kath. Univ., Department Economie, Center for Economic Studies.

- Development Research Group of the World Bank. Downloaded 2017. *PovcalNet: the on-line tool for poverty measurement*. <http://iresearch.worldbank.org/PovcalNet/>.
- Dunbar, Geoffrey R, Arthur Lewbel, and Krishna Pendakur. 2013. Children's resources in collective households: identification, estimation, and an application to child poverty in Malawi. *The American Economic Review* 103 (1): 438–471.
- 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 (): 177–196.
- . 1964. Dynamic utility. *Econometrica*: 418–424.
- . 1978. *New methods of measuring marginal utility*. Porcupine Press.
- Hamilton, Bruce W. 2001. Using Engel's law to estimate CPI bias. *The American Economic Review* 91, no. 3 (): 619–630.
- 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.

- Heckman, James J, and Richard Robb. 1986. Alternative methods for solving the problem of selection bias in evaluating the impact of treatments on outcomes. In *Drawing inferences from self-selected samples*, edited by Howard Wainer, 63–107. Springer.
- 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 (): 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. 1974. Prices and inequality: the united kingdom experience. *The Economic Journal* 84 (333): 32–55.
- . 1975. Aggregation, income distribution and consumer demand. *Review of Economic Studies* 42:525–543.
- 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.
- Pendakur, Krishna. 2002. Taking prices seriously in the measurement of inequality. *Journal of Public Economics* 86 (1): 47–69.
- Pollak, Robert A., and Terence J. Wales. 1979. Welfare comparisons and equivalence scales. *The American Economic Review* 69 (2): 216–221. <http://www.jstor.org/stable/1801646>.
- Prais, Sigbert J. 1959. Whose cost of living? *The Review of Economic Studies* 26 (2): 126–134.

- Ravallion, Martin, Shaohua Chen, and Prem Sangraula. 2009. Dollar a day revisited. *The World Bank Economic Review* 23 (2): 163–184.
- Tarozzi, Alessandro. 2007. Calculating comparable statistics from incomparable surveys, with an application to poverty in India. *Journal of Business and Economic Statistics* 25 (3): 314–336.
- Tipping, Michael E, and Christopher M Bishop. 1999. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 61 (3): 611–622.
- Townsend, Robert M. 1994. Risk and insurance in village India. *Econometrica* 62 (3): 539–591.
- Uganda Office of the Auditor General. 2014. *A value for money audit report on the production of price statistics by the Uganda Bureau of Statistics*. Technical report. <http://www.oag.go.ug/wp-content/uploads/2016/07/PRICE-INDICES.pdf>.
- 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.
- Zellner, A. 1962. An efficient method of estimating seemingly unrelated regression equations and tests for aggregation bias. *Journal of the American Statistical Association* 57:348–368.

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 $\mathbf{U} \mathbf{\Sigma}$ 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 $\mathbf{U} \mathbf{\Sigma} \mathbf{V}^\top$. The vector a can be partitioned into two parts y and x , while a matrix $\mathbf{U} \mathbf{\Sigma}_x$ can be constructed by selecting just the rows of $\mathbf{U} \mathbf{\Sigma}$ 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 $\mathbb{E}(\mathbf{U}\Sigma) = \mathbf{U}^*\Sigma^*$, provided the columns of \mathbf{X} and \mathbf{M} are uncorrelated across households.

To this end, we modestly extend an idea from the psychometric literature 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 need not compromise consistent estimation (Heckman and Robb 1986). 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 4.** (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 tell 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 $\mathbb{E}_s(g(X)|w, l)$ denote the conditional expectation and $\text{Var}_s(g(X)|w, l)$ 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 $\mathbb{E}(g(X)|w, l)$ and $\text{Var}(g(X)|w, l)$ correspond to the population distribution $F(x, w, l)$ without any selection.

Proposition 4. *Given Assumption 4, if $\mathbb{E}(X|w, l) = \mathbb{E}(X|w)$ and $\text{Var}(X|w, l) = \text{Var}(X|w)$ for all (w, l) then $\mathbb{E}_s(X|w, l) = \mathbb{E}(X|w)$ and $\text{Var}_s(X|w, l) = \text{Var}(X|w)$ 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

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

for all (x, 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 $\mathbb{E}_s(g(X)|w, l) = \mathbb{E}(g(X)|w, l) = \mathbb{E}(g(X)|w)$. The first part of the proposition is then established by taking $g(X) = X$, while the second is established taking $g(X) = g(X|w) = \mathbb{E}(XX^\top|w) - \mathbb{E}(X|w)\mathbb{E}(X|w)^\top$ for any w . \square

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 4 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 J 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 4. Suppose that we have $X = WA + U_X$, $L = WB + U_L$, and $M = LC + U_M$. Then we have

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

Under the null hypothesis that $\mathbb{E}(X|W, L) = \mathbb{E}(X|W)$, 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 (7), 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 $\mathbf{U}\mathbf{\Sigma}$ 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

Continued on next page

Continued from previous page

Code	Preferred Label	Aggregate Label
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
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

Continued on next page

Continued from previous page

Code	Preferred Label	Aggregate Label
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
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 alcoholic 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 44 categories.¹³

Table B.2: Aggregate and mean expenditures shares for 44 minimally aggregated goods in selected years.

Goods	Aggregate 2005	Aggregate 2011	Mean 2005	Mean 2011
Restaurant (food)	0.084	0.084	0.065	0.075
Sweet Potatoes (fresh)	0.076	0.068	0.083	0.080
Maize (flour)	0.072	0.081	0.086	0.090
Beef	0.065	0.067	0.051	0.053
Beans (dry)	0.062	0.068	0.079	0.087
Sugar	0.062	0.050	0.060	0.049
Fresh Milk	0.047	0.046	0.039	0.042
Cassava (dry/flour)	0.044	0.052	0.052	0.053
Cassava (fresh)	0.042	0.054	0.048	0.064
Rice	0.031	0.031	0.025	0.023
Fresh Fish	0.027	0.026	0.024	0.022
Beer	0.024	0.012	0.012	0.007
Cooking oil	0.024	0.025	0.025	0.025
Dry/Smoked fish	0.022	0.025	0.023	0.023
Chicken	0.021	0.030	0.015	0.021
Tomatoes	0.021	0.019	0.023	0.021
Other Alcoholic drinks	0.020	0.019	0.024	0.023

Continued on next page

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

Continued from previous page

Goods	2005	2011	2005	2011
Bread	0.017	0.017	0.012	0.012
Millet	0.017	0.016	0.019	0.018
Ground nuts (pounded)	0.015	0.017	0.016	0.019
Other Fruits	0.014	0.002	0.012	0.002
Goat Meat	0.013	0.017	0.012	0.011
Maize (cobs)	0.013	0.014	0.015	0.011
Irish Potatoes	0.013	0.013	0.012	0.013
Beans (fresh)	0.013	0.017	0.015	0.020
Other vegetables	0.012	0.008	0.020	0.013
Soda	0.012	0.009	0.009	0.005
Matoke (heap)	0.011	0.011	0.010	0.008
Sorghum	0.010	0.013	0.019	0.020
Dodo	0.009	0.009	0.012	0.012
Onions	0.009	0.011	0.010	0.013
Passion Fruits	0.008	0.004	0.006	0.002
Ground nuts (shelled)	0.008	0.007	0.009	0.007
Mangos	0.007	0.005	0.007	0.005
Sweet Bananas	0.007	0.005	0.005	0.004
Pork	0.007	0.011	0.006	0.010
Salt	0.006	0.006	0.010	0.009
Eggs	0.006	0.006	0.005	0.004
Cabbages	0.006	0.007	0.006	0.007
Restaurant (soda)	0.006	0.005	0.004	0.004
Tea	0.006	0.004	0.006	0.004
Sim sim	0.005	0.005	0.007	0.006
Oranges	0.003	0.003	0.003	0.002
Coffee	0.001	0.001	0.001	0.001

B.1. Aggregate and Mean Shares.

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 is k then the statistic C_k is distributed χ_{n-k}^2 . We compute the corresponding likelihood L_k . Using the Bayesian Information Criterion (BIC)

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

gives in our case an estimate of rank four for the demand system. Figure 8 presents the result of these calculations; one can see that the minimum is achieved at $k = 4$, with $b = (0.18, 0.36, 0.47, 0.62)$.

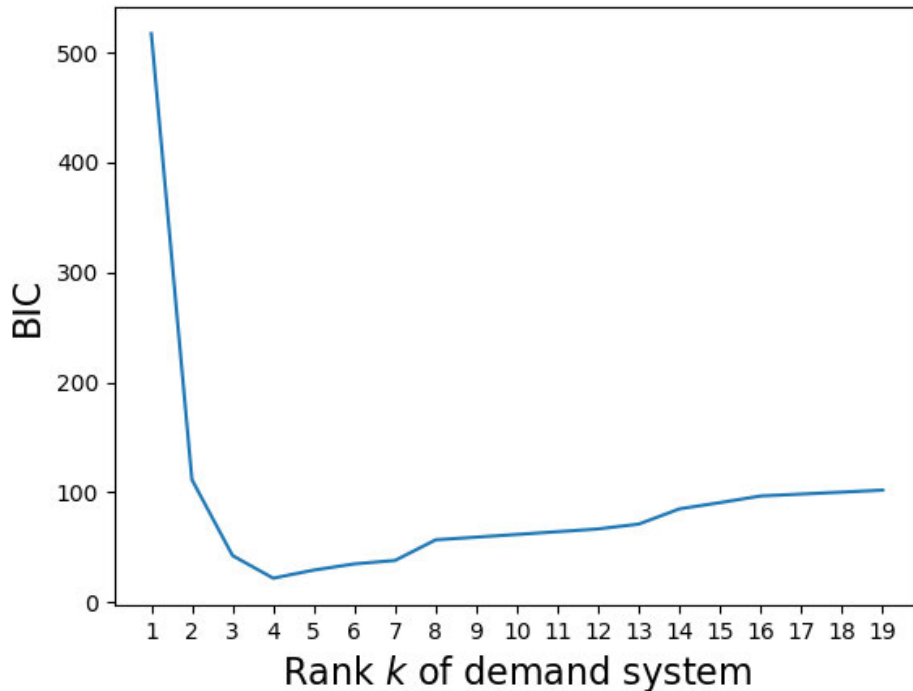


FIGURE 8. Using the BIC criterion to estimate the rank of the demand system.