

UC Berkeley

CUDARE Working Papers

Title

Risk sharing tests and covariate shocks

Permalink

<https://escholarship.org/uc/item/2zr503fq>

Author

Ligon, Ethan

Publication Date

2025-04-08

Data Availability

The data associated with this publication are available at:

<https://www.worldbank.org/en/programs/lrms/initiatives/lrms-ISA#8>

Peer reviewed

RISK SHARING TESTS AND COVARIATE SHOCKS

ETHAN LIGON

ABSTRACT. Full risk-sharing implies that marginal utilities of expenditure (MUEs) have a simple factor structure; Pareto weights divided by a common price. But MUEs must be inferred from expenditures. The standard approach to inferring MUEs and testing risk-sharing assumes homothetic preferences and unitary price elasticities. This works for idiosyncratic shocks, which affect budgets but not prices, but not "covariate" shocks which change relative prices. We obtain the complete class of preferences which allows one to infer MUEs from expenditures. We estimate these using Ugandan data. With these MUEs risk-sharing tests of covariate shocks deliver sensible results; the standard tests do not.

ligon@berkeley.edu

University of California

Berkeley, CA 94720–3310

Keywords. Risk sharing, Two-way fixed effects, Non-homothetic demands, Uganda, Marginal utility of expenditures, aggregate shocks

1. INTRODUCTION

It is well known that in a setting with risk-averse households and uncertainty, an efficient allocation will eliminate any “idiosyncratic” risk, in the sense that given this allocation the ratio of any two households’ marginal utilities will vary only in fixed proportion, regardless of the realization of the uncertain state. Suppose for example that Farmer A is unlucky and a hailstorm damages his rice crop. Aggregate resources are reduced by the size of the damage, even if all the damage accrued to A’s crops. So Farmer’s A’s marginal utility of rice will increase, but if efficiency prevails, then everyone else’s marginal utility of rice will increase proportionally, even if the shock was “idiosyncratic” in that it only affected the crops of one farmer.

The risk-sharing problem is often framed in terms of a planning problem, with different households assigned some Pareto weights *ex ante* and the planner choosing allocations decisions such that ratios of households’ marginal utilities of different goods (e.g., A’s marginal utility of rice to B’s marginal utility of rice) are equal to the ratio of their Pareto weights. But the problem immediately decentralizes in the sense that we can describe it in terms of prices and expenditures. In particular, we can say that an efficient allocation (a) keeps ratios of any two households’ marginal utilities of expenditures (MUEs) constant; and (b) that given these constant ratios, allocations (or household expenditures) depend *only* on common prices, faced by all the households efficiently sharing risk.¹

Townsend (1994) pioneered the idea of testing this “efficient risk-sharing” hypothesis in the development literature, and Deaton (1992) showed that one can implement the test as a two-way fixed effects regression.² The general approach involves regressing a measure of the (log) marginal utility of expenditures on the (log) average of households’ marginal utilities of expenditures, allowing for household-specific intercepts. The idea is to directly test the proportionality of marginal utilities of expenditures implied by efficiency. And in particular this proportionality implies a testable exclusion restriction: the event of the “idiosyncratic” shock should *not* affect individual

¹There are many different ways in which such efficient allocations could be implemented. Some of these would involve formal markets for contingent claims (Arrow & Debreu, 1954) or repeated exchange of some set of securities (Arrow, 1964; Arrow & Hahn, 1983) or forward markets (Townsend, 1978). These market may not even need to be complete (Levine & Zame, 2002). Neither must they be formal; various forms of informal exchange (Platteau & Abraham, 1987) or reciprocity (Cashdan, 1985) may well constitute a system of informal insurance approaches full efficiency, even in the complete absence of formal markets or legally enforceable contracts (Fafchamps, 1992; Ligon et al., 2002).

²Though there are important antecedents in the macro literature (Cochrane, 1991; Mace, 1991).

MUEs after controlling for its effect on the average MUE. Tests along these lines have become the stuff of textbooks (Bardhan & Udry, 1999; Ray, 1998) and have been conducted in settings all over the world, in hundreds of different studies.³ When the “idiosyncratic shock” is some measure of household income (or often deviations from mean income) such tests reliably reject the efficient risk-sharing hypothesis—household-level variation in income has a statistically significant effect on household consumption expenditures. Despite this rejection, many economists would say that the *magnitude* of this effect is not typically very large—empirically, the elasticity of consumption expenditures with respect to income is in the ballpark of 5-20%. So one might take the position that most idiosyncratic risk is shared.

But what about risk that is not idiosyncratic? After all, that hailstorm could damage the crops of not just one farmer in the village, but the crops of most farmers. Such a shock is sometimes said to be “covariate,” capturing the idea that if we condition on the shock affecting Farmer A then we’d expect some damage also to Farmer B. Notice that a covariate shock need not be universal—some households may be more affected than others.

If risk is shared efficiently, then how would we expect a covariate shock to affect outcomes? The central precept that marginal utilities of expenditure should vary in proportion will still hold, so if we were to regress log MUEs on the village average, allowing for household-specific intercepts (fixed effects), then efficient risk-sharing implies that the exclusion restriction should still hold. However, if the shock affects enough production in the village, then we might also expect the shock to have an effect on marginal utilities of expenditure because of its effect on local *prices*. Viewed from this perspective, a covariate shock is one that affects prices, while an idiosyncratic shock does not.

1.1. A puzzle. So, let us consider the effects of some covariate shocks. We take the Living Standards Measurement Survey (LSMS) panel data on household expenditures for Uganda, spanning 2005–2019, and construct a household “consumption aggregate” using the procedures of Deaton and Zaidi (2002). We are curious about the effects of “covariate” shocks on welfare, so we estimate the classic “consumption-smoothing” or risk-sharing regression (Deaton, 1992; Townsend, 1994) by regressing the log of the consumption aggregate on measures of droughts, floods, and pests (one at a

³Some recent prominent examples include Angelucci and De Giorgi (2009), Banerjee et al. (2015), Karlan et al. (2014), Kinnan (2022), Munshi and Rosenzweig (2016), and Santàeulalia-Llopis and Zheng (2018), while earlier studies are covered in surveys such as Alderman and Paxson (1994), Attanasio and Weber (2010), Dercon (2005), Morduch (1995), and Townsend (1995).

time) using two-way fixed effects and controlling for the demographic composition of the household. We add in (log) idiosyncratic household income as well, to serve as a benchmark. Here’s what we find,⁴ reporting the estimated coefficient on three different classes of shocks:

Drought	Floods	Pests	Income
0.043***	0.091***	0.094***	0.075***
(0.015)	(0.033)	(0.034)	(0.007)

So, in the conventional approach, we would reject full risk-sharing for each of these kinds of shocks. But those coefficients are positive! This would be expected for income, but does welfare really *increase* by roughly 4–8% when one of these (biblical!) covariate shocks is realized?

This sort of finding is not fragile, and survives various tweaks to the specification and different approaches to the calculation of standard errors. Neither is it confined to Uganda—“puzzling” coefficients are also found in every other LSMS panel with the requisite data that we’ve looked at (our search continues).⁵

And so, a stylized fact: These kinds of negative covariate shocks really can have a positive effect on measured consumption expenditures. That leads to the two questions this paper addresses.

- (1) How can we make sense of this? Our answer will be that using the consumption aggregate is a mistake. It’s used as a proxy for household’s marginal utility of expenditures (MUEs), but it’s only a *valid* proxy if utility functions are homothetic. Engel (1857) and a host of more recent evidence emphatically asserts (e.g., Jensen & Miller, 2008) that utility functions are *not* homothetic, and so the usual risk-sharing regression is mis-specified—even *with* full insurance we’d find these results. In particular, without homothetic preferences prices affect total expenditures in a non-separable way, so that the time effects of the TWFE regression cannot control for variation in relative prices.
- (2) How can we fix the specification, and obtain a valid test of risk-sharing against covariate shocks? To answer this, we go back to some basic consumer theory to find the broader class of preferences which permit us to infer MUEs using only data on expenditures (and household characteristics). We are able to

⁴There are lots of details, but these don’t belong in the introduction! See below for a complete discussion.

⁵Kurosaki (2014) also finds puzzling results along these lines when he examines the effects of droughts and floods in Pakistan.

completely describe this class by showing the expenditure system can be expressed in a form known as the generalized Pexider functional equation, and exploit results from the theory of functional equations to describe the complete set of solutions to these equations. The solutions take one of two forms: (i) a family of semiparametric demands which generalize demands associated with CRRA utility; and (ii) a second family of semiparametric demands which generalize Stone-Geary. Both families are non-homothetic, and together exhaust the class of demands from which MUEs can be constructed using expenditure data. The first family lends itself to linear tests of risk-sharing with two-way fixed effects; the second does not.

1.2. Organization. In this paper we first offer a diagnosis of the problems which give rise to the puzzling correlations described above. The diagnosis involves three distinct elements. First, covariate shocks cause variation in relative prices. Second, the actual structure of preferences isn't homothetic, so total expenditures depend on prices in a more complicated way than is assumed by the usual risk-sharing regression—in particular not all goods have unitary income and price elasticities. These first two elements are enough to generate correlations between covariate shocks and the consumption aggregate, even if risk is perfectly shared. Third, in practice the constructed sum of consumption expenditures is not complete, and excludes certain elastic goods or services, so that the constructed sum taken together has an income elasticity less than one. The consequence is that shocks that increase prices for observed inelastic goods will tend to increase observed expenditures, just as we observe in Uganda.

After our diagnosis, we offer a prescription. What does theory tell us about how to correctly test the hypothesis of full risk-sharing? Here the theory is remarkably clear on two basic points, both of which I believe to be novel. First, one cannot sum up expenditures on goods that have different income elasticities and use this sum to construct MUEs. Second, in order to use time effects to handle the effects of (possibly unobserved) prices on item-level expenditures, we must be able to express the system of expenditures as an additively separable function of the MUE and prices. The requirement that (some transformation of) item-level expenditures have this separability property (along with some standard regularity conditions) is equivalent to the utility function taking one of two particular semi-parametric forms, but only one of these forms is easily estimated using linear methods. Importantly, this class of utility functions nests the homothetic forms previously used in tests of risk sharing.

This knowledge of the utility function then dictates the form of demands and the risk-sharing test. Neither of these two points hinges on whether risk is actually efficiently shared.

The rest of our prescription is empirical: given the form demands must take if we're to control for prices using time effects, how can we construct estimates of MUEs? We use a simple estimator detailed in Ligon (2019).

Given our prescription, we next turn our attention to treatment: we use the aforementioned data from Uganda to construct estimates of MUEs, which are the natural objects to use in a test of risk-sharing. We conduct this test, and show that our use of these more general demand system in fact seems to resolve the puzzle we began with.

Finally, we conclude with a prognosis. Our construction of MUEs is independent of the risk-sharing hypothesis, and so these objects could be used in tests and estimation of the many dynamic models which put structure on the evolution of MUEs over time. We offer some thoughts and suggestions about ways in which these might proceed.

2. RISK-SHARING (THE GENERAL CASE)

Suppose that preferences are “regular” (i.e., can be represented by an increasing, concave, continuously differentiable utility function). More restrictively, assume that preferences are von Neumann-Morgenstern and intertemporally separable.

Following Townsend (1994), consider the problem facing a social planner in an environment with uncertainty; at date t state $s \in \{1, 2, \dots, S\}$ is realized with probability $\Pr_t(s)$. The planner maximizes a weighted sum of households' utilities, with the (Pareto) weight θ_i associated with household i 's utility. Utility at date t is discounted by some $\beta_t < 1$. At each date-state (t, s) the planner allocates a consumption *budget* to each household, but has to respect the aggregate resource constraint that the sum of all expenditures must be less than some given quantity $\bar{x}_t(s)$. For household i at date t in state s let the budget allocated be $x_{it}(s)$. Given this budget and taking prices as given the household then solves the usual consumer problem.

In general, indirect utility within the period will depend not only on the size of the budget, but also a complete vector of prices $p_t(s)$ and household characteristics z_{it} . Then the planner's intertemporal problem can be written

$$(1) \quad \max_{x_{it}(s)} \sum_i \theta_i \sum_t \beta_t \sum_s \Pr_t(s) V(x_{it}(s), p_t(s); z_{it}(s))$$

subject to the aggregate budget constraint $\sum_i x_{it}(s) = \bar{x}_t(s)$ for all t, s .

Let $\nu_t(s) = \mu_t(s)\beta_t\text{Pr}_t(s)$ be the multiplier associated with the aggregate resource constraint at time t in state s . Then the first order conditions associated with the assignment of $x_{it}(s)$ are

$$\theta_i \frac{\partial V}{\partial x}(x_{it}(s), p_t(s); z_{it}(s)) = \theta_i \lambda(x_{it}(s), p_t(s); z_{it}(s)) = \mu_t(s)$$

for all (i, t, s) , where $\lambda(x, p, z)$ is a function that can be interpreted as the household's marginal utility of expenditures (MUE), and where $\mu_t(s)$ is the shadow price associated with the aggregate budget within date-state (t, s) . Then taking logarithms and rearranging we have

$$(2) \quad \log \lambda_{it}(s) = \log \mu_t(s) - \log \theta_i.$$

Equation (2) is the hallmark of full risk-sharing, given time-separable von Neumann-Morgenstern preferences, expressing the simple factor structure of optimal allocations. It also immediately lends itself to testing: The right hand side can be estimated using panel data with two-way fixed effects (time and household), assuming only that prices are common.

Everything up to this point is standard and implied by Pareto optimality, provided only that agents are risk averse and have well-behaved preferences which are separable across dates and states.

There is, however, a key issue before taking the predictions of full risk sharing to data: one must take a stand on how to construct the MUE function $\lambda(x, p, z)$. Townsend adopted a representation of momentary household utility which depends only on the consumption aggregate, normalized by a price index and a scalar function of household characteristics; this is equivalent to assuming a homothetic utility function (Blackorby & Donaldson, 1988). The empirical literature has mostly followed his example. Typical risk-sharing tests following Townsend (or Deaton, 1992) assume homothetic or (less often) quasi-homothetic (e.g., Ogaki & Zhang, 2001; Zhang & Ogaki, 2004) preferences, for example assuming the household indirect utility function takes the ‘‘Constant Relative Risk Aversion’’ (CRRA) form

$$V(x, p, z) = \frac{(x/(\pi(p)g(z)))^{1-\gamma} - 1}{1 - \gamma},$$

where $\pi(p)$ is a scalar price index, and where $g(z)$ is a scalar function mapping household characteristics into ‘‘adult equivalents’’.⁶ Then the household's marginal

⁶Townsend (1994) actually works principally with exponential or CARA utility, with $V(x, p, z) = -\frac{1}{\sigma} \exp[-\sigma(x/\pi(p) - g(z))]$ which delivers a regression specified in levels rather than logs of total expenditures, but the subsequent literature has generally adopted the CRRA specification.

utility of expenditures is given by

$$\lambda(x, p, z) = \frac{x^{-\gamma}}{(\pi(p)g(z))^{1-\gamma}}.$$

Substituting this expression into (2) and re-arranging yields an estimating equation of the form exploited by Townsend and Deaton,

$$(3) \quad \log x_{it} = \frac{1}{\gamma} \log \theta_i - \left[\frac{1}{\gamma} \log \mu_t + \frac{1-\gamma}{\gamma} \log \pi(p_t) \right] - \frac{1-\gamma}{\gamma} \log g(z_{it}).$$

So we estimate this by regressing the log of the consumption aggregate on household fixed effects (which identify $\gamma^{-1} \log \theta_i$) and time (or perhaps village-time) effects (which identify the term in square brackets involving only prices), and some known function g of observed household characteristics.

From the comparison of the two equations (2) and (3) three important points emerge. First, that the Townsend/Deaton risk-sharing regression is a special case of the more general (2). Second, that the marginal utility of expenditures λ already automatically incorporates information on household characteristics that may affect demand, and does so much more flexibly than does the $g(z)$ function that appears in the Townsend approach (c.f., Lewbel, 2010). Third, in general λ also depends on the entire vector of prices p , while in the CRRA case prices affect total expenditures only via a single scalar price index $\pi(p)$.

Townsend-style risk-sharing tests are generally implemented by adding some measure of a “shock” to (3), and testing the exclusion restriction (idiosyncratic shocks shouldn’t affect MUEs). But even if perfect risk-sharing doesn’t hold (we can think of this as the Pareto weights θ_i varying with the state), total expenditures in this framework depend only on Pareto weights (reflecting households’ relative wealths), on prices p (capturing aggregate shocks to demand and supply), and on characteristics z (which may drive changes to the structure of household demands). So “shocks” can affect total expenditures only via one of these three channels.

Under a maintained hypothesis of full risk-sharing, idiosyncratic income variation will be insured, but variation in prices will still affect expenditures, implying a regression of the form

$$(4) \quad \log \lambda_{it} = \log \mu_t - \log \theta_i + \delta \text{Shock}_{it} + e_{it}.$$

When we observe λ_{it} , we can implement this regression simply by adding two-way fixed effects—time effects account for $\log \mu_t$, while household fixed effects account for $\log \theta_i$. This leaves a disturbance term e_{it} which can be interpreted as either (or

both of) measurement error in the dependent variable, or the effects of time-varying unobserved household characteristics on expenditures. In either case full insurance implies the exclusion restriction $\delta = 0$. (Here we also assume that the shock doesn't affect unobserved household characteristics.) Note that in this case full insurance implies $\delta = 0$ regardless of whether the shock is idiosyncratic (doesn't affect prices) or covariate (affects prices).

Now, consider the special case of CRRA preferences, for which MUEs can be expressed in the separable form $\log \lambda(x, p, z) = -\gamma \log \left(\frac{x}{\pi(p)g(z)} \right) - \log \pi(p) - \log g(z)$. Adding a "shock" to the right-hand side of (3) yields the CRRA risk-sharing regression

$$(5) \quad \log x_{it} = \frac{1}{\gamma} \log \theta_i - \frac{1}{\gamma} [\log \mu_t + (1 - \gamma) \log \pi(p_t)] + \delta \text{Shock}_{it} - \frac{1 - \gamma}{\gamma} g(z_{it}) + \frac{1}{\gamma} e_{it},$$

where now the *joint* hypothesis of full risk-sharing and CRRA preferences implies $\delta = 0$. But suppose that there's full risk-sharing but preferences are *not* CRRA. If the shock is idiosyncratic (and so doesn't affect prices, and also doesn't affect characteristics z) then we would still expect $\delta = 0$. But if the shock changes relative prices then the exclusion restriction will fail, because in this case the single index $\pi(p)$ can't account for the effects of changes in relative prices on the composition of expenditures. An immediate consequence is that the disturbance term e_{it} must be a function of those prices.

More particularly, in the face of increased prices expenditures on inelastic goods such as food will *increase*. And since most of the expenditure items we have household-level data on are different sorts of food, we might expect a measure of total food expenditures to be positively correlated with "covariate" negative shocks which increase local prices, such as drought, floods, pests, or changes in prices for agricultural inputs. Or putting a finer point on it, if the measure of x_{it} is not really total expenditures on all non-durable goods and services, but is expenditures on a subset of goods which have inelastic demands, then by definition expenditures on that subset will tend to increase with increases in prices of those goods. Per Engel (1857) the subset "food" would be a good example, with more recent evidence including McKenzie (2003) and Thomas and Frankenberg (2007). And in this case we would predict that any shock that causes increases in food prices will be associated with $\delta > 0$.

So, our diagnosis. Covariate shocks affect prices and hence expenditures in ways that aren't properly accounted for in a specification of the risk-sharing regression that assumes homothetic preferences (i.e., that prices have no effect on expenditures). Thus, even if there *is* full insurance covariate shocks may be correlated with total

expenditure. Further, if the constructed consumption aggregate excludes some elastic goods in particular, then we would expect the correlation between covariate shocks and the constructed consumption aggregate to be positive, just as observed in Uganda.

3. INFERRING MUEs FROM EXPENDITURE DATA

The problem we’ve identified is that the usual risk-sharing regression provides a joint test of full risk-sharing *and* CRRA utility. Assuming CRRA utility allows us to use nothing more than panel data on expenditures (and perhaps household demographics) to construct risk-sharing tests based on panel estimation of two-way fixed effects. In practice these are highly desirable properties. Are there non-homothetic utility structures which could more flexibly account for demand responses to changes in relative prices while at the same time preserving these desirable properties? In this section we establish that the answer is “yes”, and obtain the complete class of demand systems which (i) are implied by some regular utility function; and (ii) can be constructed using (just) data on consumption expenditures (and perhaps household demographics).

3.1. Preliminaries. In this section we set aside any explicit consideration of household characteristics. There’s no loss of generality in this. With time-separable von-Neumann-Morgenstern preferences we can simply describe utility functions and demand implicitly conditioning on those characteristics.

Thus, let $\mathcal{U}_n$ be the set of strictly increasing, strictly concave, twice-continuously differentiable functions mapping $\mathbb{R}_+^n$ into $\mathbb{R}$, and call $\mathcal{U}_n$ the set of *regular* utility functions over $\mathbb{R}_+^n$.

For a household with a utility function $U \in \mathcal{U}_n$ with a total budget $\bar{x} > 0$ facing prices $p \in \mathbb{R}^n$, a Lagrangian formulation of the *consumer’s problem* is to solve $\max_{c \in \mathbb{R}_+^n} U(c) + \lambda(\bar{x} - p^\top c)$, with λ the Lagrange multiplier or MUE that we would like to obtain for our risk-sharing test.

We can express a solution to the consumer’s problem in terms of demand functions which depend on prices and λ . This form of demands was advocated by Ragnar Frisch, so we might say that Frischian demands map the product of positive quantity λ and n prices into n quantities demanded. We say that

Condition 1. *An n -vector of Frisch demands $f(p, \lambda)$ is rationalized by U if there exists a $U \in \mathcal{U}_n$ such that*

$$(6) \quad u_j(f(p, \lambda)) := \frac{\partial U}{\partial c_j}(f(p, \lambda)) = p_j \lambda$$

for all p in any open subset of $\mathbb{R}_+^n$ and $\lambda > 0$ for $j = 1, \dots, n$.

Similarly, we say a given f is *rationalizable* if there exists a $U \in \mathcal{U}_n$ which rationalizes f .

Condition 1 basically requires that demands be interior solutions to the problem of maximizing some regular utility function subject to a budget constraint. If a consumer has a utility function U , and solutions to that consumer's problem are characterized by the first order conditions (6), then these demands will also be solutions to this consumer's problem.

From the form of the right-hand side of (6) it's apparent that Frischian demands must be homogeneous of degree zero in $(p, 1/\lambda)$. It follows that we can simplify the expression of Frischian demands, as we have $f(p, \lambda) = f(p\lambda, 1)$, which can then simply be written as $f(p\lambda)$. This same point follows from the observation that at the optimum the vector of marginal utilities u can be inverted to yield $f(p, \lambda) = u^{-1}(p\lambda)$, so that we know we can write Frischian demands as simply $f(p\lambda)$.

3.2. MUEs for risk-sharing tests. The problem: we observe consumption expenditures $\{x_j\}$ for some (but perhaps not all) goods j . From these data we wish to infer values for MUEs; further, we want to be able to use these MUEs in a risk-sharing test that can be implemented using two-way fixed effects.

The fundamental risk-sharing equation is

$$\log \lambda = \log \mu(p) - \log \theta,$$

but we don't directly observe λ , only expenditures. So if we're to preserve the risk-sharing test we need to be able to write some transformation of expenditures as an additively separable function of prices and $\log \lambda$. Note that CRRA utility does exactly this as expenditures for good j in the CRRA case satisfy $-\gamma \log x_j = \log \lambda + (1 - \gamma) \log p_j$, or (summing over goods) $\log x = \frac{-1}{\gamma} \log \lambda + \log \pi(p)$ for some linearly homogeneous price index $\pi(p)$. However, other more general preferences will also work. What are these preferences?

In general, for any good j we need functions (ϕ_j, a_j) such that

$$(7) \quad \phi_j(x_j) = \log \lambda + a_j(p).$$

If there exist functions (ϕ_j, a_j) satisfying (7), then we can substitute into the fundamental risk-sharing equation, obtaining

$$\phi_j(x_j) = [\log \mu(p) - a_j(p)] - \log \theta.$$

Since the term in brackets varies only with prices this can serve as the basis of the kind of risk-sharing test that we're after, with time-effects identifying $\log \mu(p) - a_j(p)$ and household fixed effects identifying $\log \theta$. The key property is (7); this is a special case of a more general property we'll call λ -separability.⁷

Condition 2. *The Frischian expenditures on good j , $x_j(p, \lambda) \equiv p_j f_j(p, \lambda)$ are λ -separable if there exist functions (ϕ_j, a_j, b_j) such that*

$$(8) \quad \phi_j(x_j(p, \lambda)) = a_j(p) + b_j(\lambda),$$

with ϕ_j continuously differentiable and a_j either non-constant or zero.

Note that while rationalizability is a property of the entire system of demands and expenditures, (λ) -separability is a property of a particular good. In particular it's possible that some but not all demands or expenditures are (λ) -separable.

3.3. Demands and utilities when expenditures are λ -separable. Exploiting the fact that expenditures must be linearly homogeneous, it turns out that one can write any rationalizable λ -separable expenditures in the form

$$k(p + \lambda) = g(\lambda)\ell(p) + h(p),$$

which is called the *generalized Pexider* equation. This gives us a single functional equation in two variables, which can be solved for the four functions g , h , k , and ℓ . Exploiting this allows us to describe all rationalizable demands and utilities when expenditures are λ -separable:

Theorem 1. *If expenditures for some good i satisfy Condition 1 and Condition 2 with ϕ_j increasing; $a_j(p)$ either non-constant or zero, and continuous at a point; and with b_j continuous at a point, then transformation functions ϕ_j , Frischian demands f_j and rationalizing marginal utility u_j must satisfy one of the following two cases for positive constants α_j , β_j , and σ_j :*

- (1) *(Constant Frisch Elasticity): $\phi_j(x_j) = \log(x_j)$; $f_j(p, \lambda) = (\alpha_j/(\lambda p_j))^{\beta_j}$; and $u_j(c) = \alpha_j c_j^{-1/\beta_j}$.*
- (2) *(Generalized Stone-Geary): $\phi_j(x_j) = x_j^{\sigma_j}$; $f_j(p, \lambda) = [(\beta_j/(\lambda p_j))^{\sigma_j} + \alpha_j]^{1/\sigma_j}$; and $u_j(c) = \beta_j (c_j^{\sigma_j} - \alpha_j)^{-1/\sigma_j}$.*

Proof. See Appendix A.2. □

⁷This property generalizes what Browning et al. (1985) calls “Case 2” demands, discussed in detail in Ligon (2016b).

Rationalizing Utility Functions. The labels of the different cases in Theorem 1 indicate names for the rationalizing utility function U having marginal utilities $u_j(c)$; an example of “Constant Frisch Elasticity” (CFE) utility can be written as $U(c) = \sum_{j=1}^n \alpha_j \beta_j \frac{c_j^{1-1/\beta_j} - 1}{\beta_j - 1}$.⁸ The CFE system generalizes the Constant Elasticity of Substitution (CES) system (take $\beta_j = \beta$), of which the Cobb-Douglas system is a limiting case (take $\beta \rightarrow 1$, applying L’Hôpital’s rule). Both the CES and Cobb-Douglas cases are homothetic, and consistent with the CRRA indirect utility function. Finally, the “Generalized Stone-Geary” case gives what is, to the best of my knowledge, a marginal utility function which has not previously appeared in the literature. This case gives demands which are not linear in parameters, which may limit its usefulness in applied empirical work. However, when $\sigma_j = 1$ one obtains the quasi-homothetic Stone-Geary utility function, which suggests that it could be used to explore the behavior of Engel curves, perhaps exploiting a Box-Cox approach to estimation.

3.4. Estimating $\log \lambda$. Using Theorem 1, the condition that expenditures be λ -separable implies that expenditures must take one of two forms:

- (1) $\log x_j = a_j(p) - \beta_j \log \lambda$; or
- (2) $\log x_j = \frac{1}{\sigma_j} \log \left[\left(\frac{\beta_j}{\lambda} \right)^{\sigma_j} + a_j(p) \right]$.

The second form (generalized Stone-Geary) does not easily allow us to estimate $\log \lambda$ using data on expenditures. But the first does, using standard “interactive fixed effects” panel methods (Bai, 2009). But here, instead of the panel dimensions varying over households and time periods, they vary over households and items of consumption expenditure. In particular, indexing goods by j and households by i we have for *every* period t

$$(9) \quad \log x_{it}^j = g_j(z_{it}) + a_j(p_t) - \beta_j \log \lambda_{it} + \epsilon_{it}^j,$$

where z_{it} are observable household characteristics, such as household size and composition, $a_j(p_t)$ measures the effect of common time t prices p_t on expenditures, where $\beta_j \log \lambda_{it}$ gives us the effect of λ on expenditures for good j , and where ϵ_{it}^j can be regarded as measurement error in expenditures, or perhaps the effects of unobserved household characteristics on expenditures. The term $a_j(p_t)$ we can account for using good-time effects. Different strategies may be employed to estimate the functions g_j ; our preference here is perhaps the simplest, which is to simply assume that $g_j(z)$ is

⁸Ligon (2019) gave this name to a general form of this utility function, but special cases include the “direct addilog” of Houthakker (1960) or the “constant relative income elasticity” form of Caron et al. (2014).

linear in the vector of characteristics z . And then the tools of factor analysis can be used to simultaneously estimate the parameters β_j and $\log \lambda_{it}$, up to an unknown (and unimportant) factor of proportionality. (Details are given in Ligon, 2019, while code to compute estimates is provided by Ligon (2023).) We call the system (9) the *Constant Frisch Elasticity* (CFE) expenditure system, as it is an example of the systems considered by Frisch (1959) with its chief distinguishing characteristic the fact that the elasticities (the β_j) of expenditures with respect to MUE are constant.

4. DATA

We use eight rounds of LSMS data from Uganda, with expenditures on 41 different goods (mostly foods). We also use some demographic data: counts of men, women, boys, and girls; the log of total household size (to allow for effects of scale); a “Rural” dummy variable. Our test of risk-sharing requires panel data on households over time. The Ugandan LSMS covers eight waves of data collection, spanning the years 2005–2020. Waves were not always collected at regular intervals, however, and only a subset of households appear multiple times. In particular, while we observe usable data from 5797 distinct households over this period, 980 of these households appear only once, and so contribute nothing of value to the risk sharing test (though data on these households’ consumption expenditures is still useful for estimating the CFE demand system). Table 1 describes the number of households observed in each round, and how many households observed in a particular round also turn up in another. For example, looking at the first row of the table, we have data for 3122 households in 2005, and of these 1288 are also observed in 2019-20. Overall, we have an eight year unbalanced panel with a total of 22,500 usable household-year observations.

TABLE 1. Households’ Attrition in the Ugandan LSMS

Round	2005	2009	2010	2011	2013	2015	2018	2019
2005-06	3122	2606	2386	2363	1566	1470	1356	1288
2009-10	—	2974	2617	2581	1685	1578	1451	1376
2010-11	—	—	2685	2487	1603	1506	1389	1322
2011-12	—	—	—	2843	1712	1610	1481	1410
2013-14	—	—	—	—	3117	2870	2613	2468
2015-16	—	—	—	—	—	3305	2968	2796
2018-19	—	—	—	—	—	—	3241	3007
2019-20	—	—	—	—	—	—	—	3076

Table 2 reports some characteristics of the sample households, by year. These characteristics are important because the household-level marginal utility of expenditures

seems likely to depend not only on household size, but also on its composition. Our

TABLE 2. Mean characteristics of households by year.

(t/N_t)	Girls	Boys	Women	Men	Rural	log HSize
2005-06	1.55	1.56	1.17	1.05	0.72	1.48
3122	(1.48)	(1.49)	(0.81)	(0.82)	(0.45)	(0.68)
2009-10	1.73	1.77	1.26	1.13	0.74	1.60
2974	(1.53)	(1.59)	(0.82)	(0.90)	(0.44)	(0.64)
2010-11	1.85	1.86	1.31	1.16	0.78	1.66
2685	(1.60)	(1.62)	(0.88)	(0.93)	(0.41)	(0.63)
2011-12	1.79	1.77	1.29	1.15	0.79	1.63
2843	(1.57)	(1.58)	(0.80)	(0.90)	(0.40)	(0.62)
2013-14	1.58	1.59	1.25	1.12	0.74	1.54
3117	(1.50)	(1.50)	(0.77)	(0.86)	(0.44)	(0.64)
2015-16	1.28	1.29	1.17	1.05	0.75	1.34
3304	(1.39)	(1.40)	(0.77)	(0.88)	(0.43)	(0.74)
2018-19	1.46	1.42	1.23	1.11	0.75	1.48
3176	(1.40)	(1.36)	(0.72)	(0.86)	(0.44)	(0.63)
2019-20	1.42	1.37	1.23	1.09	0.76	1.46
3076	(1.40)	(1.32)	(0.71)	(0.85)	(0.43)	(0.64)
Pooled	1.57	1.57	1.23	1.11	0.75	1.52
24297	(1.49)	(1.49)	(0.79)	(0.87)	(0.43)	(0.66)

estimates of MUEs relies on preferences *not* being homothetic—we are using data on the composition of households’ consumption portfolios, and since we observe only a subset of all consumption goods changes in relative prices make efforts to use consumption aggregates ill-advised (Christiaensen et al., 2022). Figure 1 uses the fact that homothetic preference structures imply that *aggregate* expenditure shares (of the sort that appear in national income and product accounts) will be the same as *average* expenditure shares. The figure takes logs of each of these, and then differences. It is apparent that these shares are not the same at all, with goods to the right (e.g., beer) playing a larger role in the consumption portfolios of households with larger budgets, and goods to the left (e.g., sorghum) playing a more important role for poorer households.

The data we have from the LSMS surveys includes self-reported data on a variety of “shocks” the household may have experienced (Heltberg et al., 2015, discuss the collection of this sort of data in a variety of different household surveys), generally elicited using the prompt “Did you experience [SHOCK] in the last twelve months?”⁹

⁹There is some variation in the elicitation in different rounds, and the first 2005-06 round in particular uses a longer reporting period.




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

Where the answer is “Yes”, the respondent is asked about the timing of the shock (in what month did the shock occur, and for how long did it last). There is some modest variation across rounds in the language used to describe different sorts of shocks (see Appendix A.1), but one can distinguish two classes. The first is idiosyncratic, shocks which directly involve the household. Frequently reported idiosyncratic shocks include health issues (serious illness or accident), thefts of property, and death (death of “income earners” is reported separately from the death of other household members). The second is covariate, shocks which seem likely to affect many households within a local area, though not necessarily equally. Frequently reported covariate shocks are drought (or “irregular rains”), agricultural pests, floods, and adverse agricultural prices (unusually expensive inputs or unusually low prices for output). Table 3 reports the incidence of these shocks across different rounds.

To get a notion of frequency from Table 3, there are about 2800 households observed per year. Drought is by far the most frequently reported shock. Meteorologists regard 2005–08 as a period of major drought for Uganda, with 2010–11 and 2014–15 periods

TABLE 3. Reported incidence of different kinds of shocks by year.

Shock	2005	2009	2010	2011	2013	2015	2018	2019	Total
Health	82	377	301	156	133	88	190	197	1524
Theft	349	233	96	55	76	62	75	83	1029
Death	423	74	58	35	66	35	44	49	784
Death of earner	99	27	17	19	30	19	20	19	250
Drought	1234	1344	710	560	914	598	736	529	6625
Floods	426	61	102	148	98	62	74	117	1088
Pests	475	219	77	92	71	53	130	84	1201
Prices	71	113	54	65	67	12	78	29	489

of minor drought (Byakatonda et al., 2021), consistent with the household reports in Table 3. Health shocks (both illness and accidents) are the second most frequently reported shocks, followed by pests (which includes both crop pests and livestock disease), floods, death, theft, and adverse changes in agricultural prices.

While there’s important variation across years in the proportion of households that report drought, we would like to know that the data are consistent with those shocks being covariate, in the sense that if a person at one location reports drought, then another nearby person is likely to report the same. Figure 2 documents all the drought reports across the eight different LSMS-ISA countries over all the years (2005–20) for which we have data. Dots correspond to sample clusters, and the “redder” the dot, the higher the proportion of households which reported drought.

Casual observation of the figure suggests not only that there are many dark dots (high proportions of households within a cluster agreeing that there’s drought), but also that darker dots tend to be close to each other, consistent with the notion that these reports really are reflecting covariate shocks.

5. ESTIMATION OF THE CFE EXPENDITURE SYSTEM

With these data we estimate the β_j parameters and the values of $\log \lambda_{it}$ for every household-year. Figure 3 shows our estimates of the β_j elasticities; note that we can easily reject the hypothesis that these are all equal, as they would be in the nested CRRA case—there’s a wide range of income elasticities across different goods, even for the same household, and the least elastic (those with low values of β_j) are what we might expect (starchy staples, salt). Similarly, goods such as fresh milk, sweet bananas, coffee, oranges, and passion fruit all exhibit high elasticities. The ratios of the β coefficients are equal to ratios of income elasticities, implying that passion fruit is roughly three times as income elastic as cassava.




FIGURE 3. Estimates of elasticities β_j . These can be interpreted as price elasticities, and (by Pigou's Law) are proportional to income elasticities. Horizontal bars are 95% confidence intervals.

6. TESTS OF RISK SHARING IN THE FACE OF COVARIATE SHOCKS

In this section we want to actually conduct the risk-sharing tests described above. We'll use two different approaches. The first is the classic test assuming CRRA utility: for these the dependent variable is the log of total expenditures. The second is the test assuming our more general CFE preferences; for these the dependent variable is $w = -\log \lambda$.

6.1. Effects of shocks on welfare. The risk-sharing test we're concerned with is meant to provide a test of the null hypothesis that $E(\log \lambda | p, d) = E(\log \lambda | p, d, \text{Shocks})$. If demands are consistent with the CFE specification, then the welfare measures w we've constructed are estimates of $-\log \lambda$, and we are justified in using w as the dependent variable in the two-way fixed effects regressions we've described. When




FIGURE 4. Distribution of w by year.

$\beta_j = \bar{\beta}$ for all goods j , then CFE demands will coincide with the special case of CRRA, and we will have $-\log \lambda$ proportional to $\log x$.

Table 4 describes the effects of different shocks on both w (equation 4, the CFE risk-sharing regression) and on $\log x$ (equation 5, the CRRA risk-sharing regression). Reported standard errors here and elsewhere are robust (Arellano, 1987). These estimates provide strong additional evidence favoring the general CFE over the CRRA specifications of the risk-sharing regression. Because the elasticities in the CFE demand system take different values for different goods, the disturbance term in the standard CRRA risk-sharing regression must be a function of relative prices, so that a regression of $\log$ total food expenditures on shocks which increased food expenditures would yield positive estimated values of δ . The same would not be true for a regression with w as the dependent variable.

The results of Table 4 confirm this reasoning rather dramatically. In fact, *every* covariate shock has a significant positive coefficient in the CRRA regressions, while *no* covariate shock has a significant coefficient in the CFE regressions. This is strong




FIGURE 5. Engel Pie: Estimated expenditure shares as a function of $\log x$. Inset circles indicate the 0.01, 0.5, and 0.99 quantiles of the distribution of $\log x$. Food labels are ordered alphabetically clockwise (labels which won't easily fit omitted).

evidence in favor of the view that these negative covariate shocks increase local consumption expenditures by causing an increase in local prices, and that the CRRA specification does not control for these price effects correctly.

Among the idiosyncratic shocks, income is significant, has the expected sign, and has a magnitude similar to that found in other Townsend-style tests of full insurance. No other coefficients in either specification are significant, with the sole exception of Health for CFE. This last should not be a surprise, and neither should it be interpreted as a rejection of the risk-sharing model, as our estimation of the CFE demands did not include any information on health as a household characteristic. If health affects demands, in our specification it can do so either via a shock to the budget and thus w (which would be at odds with full insurance, but not the CFE demand specification) or via the disturbance term in the demand equations, which explicitly depends on unobserved household characteristics such as health.

TABLE 4. Effects of different shocks on welfare measures w and $\log x$. Shocks in the top panel are idiosyncratic, and in the lower panel are covariate. All regressions control for year-market effects, household fixed effects, and a vector of household demographics.

Shock	w (CFE)	$\log x$ (CRRA)
Death	0.023 (0.050)	0.020 (0.050)
Death of earner	0.033 (0.089)	0.007 (0.089)
Health	-0.057** (0.024)	-0.028 (0.024)
Theft	0.023 (0.033)	0.022 (0.033)
Income	0.056*** (0.007)	0.075*** (0.007)
Drought	0.007 (0.015)	0.043*** (0.015)
Floods	0.030 (0.033)	0.091*** (0.033)
Pests	0.042 (0.034)	0.094*** (0.034)
Prices	-0.044 (0.038)	0.099*** (0.038)

The data we have on expenditures is for the past week; the data we have on shocks is for the past year. Could this somehow cause the apparently aberrant effects of positive shocks on expenditures, rather than something to do with prices? What if we considered only shocks that were in close temporal proximity to the expenditures? We construct a dummy variable which takes the value one if there's a reported covariate shock within the m months prior to the interview, and then re-estimate our risk-sharing regressions allowing m to vary from zero months up to twelve months.

Results are reported in Figure 6.

Figure 6 provides even stronger evidence favoring the CFE over the CRRA risk-sharing specifications. For the CRRA ($\log x$) specification, having had any covariate shock within the last year has a significant positive effect on log consumption, in every month. In contrast, for the CFE (w) specification, having any covariate shock within the last month has a rather large and significant negative effect on w . In subsequent months the magnitudes of the effects of shocks on w are uniformly smaller than




FIGURE 6. Effects of any covariate shock within the last m months on welfare measures, from a two-way panel regression. Scale is in standard deviations of the dependent variable (w or $\log x$). Error bars cover a span of two standard errors about the point estimate.

the CRRA case, though some are positive and statistically significant. The full risk-sharing hypothesis holds that covariate shocks should influence w only through prices, so we could interpret the significant negative effect in the first month as evidence against the full risk-sharing hypothesis.

To try to get a better understanding of these patterns, Table 5 reports the same kind of regressions described by Figure 6 which allow for shock windows of zero through twelve months, but keeps the different kinds of shocks separate, as in Table 4. As in that table, most of the covariate shock coefficients (the last four columns) are significant and positive in the CRRA specification, while only two are significant for the CFE specification reported in the subsequent Table 6.

Using a two-way fixed effects regression we've seen that covariate shocks have at most a small effect on w , save for the first month. We've argued that this is because the main mechanism by which covariate shocks affect MUEs is via prices, and the two-way fixed effects regression controls for these. What if we *don't* control for prices? A panel

TABLE 5. Effects of different shocks within the last m months on log consumption expenditures.

Months	Health	Theft	Death	Death of earner	Drought	Floods	Pests	Prices
0	0.013 (0.092)	0.058 (0.116)	0.088 (0.309)	0.779* (0.454)	-0.069 (0.152)	0.283 (0.210)	0.291* (0.152)	-0.220 (0.172)
1	-0.008 (0.050)	0.007 (0.065)	-0.024 (0.147)	0.216 (0.204)	-0.016 (0.063)	0.104 (0.104)	0.203** (0.098)	0.068 (0.153)
2	-0.030 (0.042)	-0.010 (0.053)	0.015 (0.107)	0.345 (0.214)	-0.022 (0.042)	0.119 (0.074)	0.176** (0.078)	0.090 (0.115)
3	-0.019 (0.038)	-0.011 (0.048)	0.042 (0.090)	0.294* (0.176)	-0.001 (0.031)	0.172*** (0.059)	0.155** (0.062)	0.104 (0.087)
4	-0.032 (0.035)	-0.021 (0.045)	0.042 (0.075)	0.104 (0.159)	0.016 (0.025)	0.137*** (0.049)	0.134** (0.056)	0.122* (0.068)
5	-0.034 (0.032)	-0.022 (0.041)	0.048 (0.069)	0.046 (0.138)	0.050** (0.022)	0.116*** (0.044)	0.158*** (0.051)	0.097* (0.056)
6	-0.037 (0.030)	-0.007 (0.037)	0.026 (0.064)	0.037 (0.123)	0.059*** (0.019)	0.080* (0.041)	0.142*** (0.047)	0.120** (0.051)
7	-0.024 (0.029)	0.012 (0.037)	0.005 (0.062)	0.021 (0.116)	0.046*** (0.018)	0.061 (0.038)	0.130*** (0.043)	0.109** (0.048)
8	-0.022 (0.028)	-0.001 (0.035)	0.018 (0.058)	0.048 (0.108)	0.051*** (0.016)	0.068* (0.037)	0.129*** (0.042)	0.102** (0.044)
9	-0.019 (0.027)	0.004 (0.034)	0.026 (0.055)	0.030 (0.101)	0.056*** (0.016)	0.065* (0.036)	0.121*** (0.039)	0.097** (0.042)
10	-0.018 (0.026)	0.009 (0.034)	0.028 (0.053)	0.020 (0.096)	0.049*** (0.015)	0.074** (0.035)	0.118*** (0.037)	0.107*** (0.040)
11	-0.025 (0.025)	0.014 (0.033)	0.032 (0.051)	0.026 (0.093)	0.046*** (0.015)	0.072** (0.034)	0.104*** (0.036)	0.098** (0.038)
12	-0.028 (0.024)	0.022 (0.033)	0.020 (0.050)	0.007 (0.089)	0.043*** (0.015)	0.091*** (0.033)	0.094*** (0.034)	0.099*** (0.038)

regression incorporating only household fixed effects would then reveal the effects that shocks have on MUEs via prices. Of course, attribution becomes problematical, since the incidence of covariate shocks across years and markets will be correlated not only with prices, but possibly other shocks. Accordingly, Figure 7 shows the results from a series of regressions of w and $\log x$ on the number of reported covariate shocks within the last m months, just as in Figure 6, but in this case only as a “one-way” panel estimator, controlling for household fixed effects but *not* time-market effects.

Here w falls with reported “negative” covariate shocks (drought, floods, pests, prices) in the way we would expect, with estimated coefficients significantly negative for both the month of the shock and at more distant intervals (anything from

TABLE 6. Effects of different shocks within the last m months on current $w = -\log \lambda$

Months	Health	Theft	Death	Death of earner	Drought	Floods	Pests	Prices
0	−0.057 (0.092)	0.239** (0.116)	−0.072 (0.309)	1.013** (0.454)	−0.367** (0.152)	0.357* (0.210)	−0.087 (0.152)	−0.528*** (0.172)
1	0.014 (0.050)	0.045 (0.065)	−0.079 (0.147)	0.406** (0.204)	−0.061 (0.063)	0.095 (0.104)	0.125 (0.098)	0.059 (0.153)
2	0.018 (0.042)	0.060 (0.053)	−0.087 (0.107)	0.485** (0.214)	0.005 (0.042)	0.062 (0.074)	0.114 (0.078)	0.057 (0.115)
3	0.015 (0.038)	0.005 (0.048)	0.037 (0.090)	0.367** (0.176)	−0.002 (0.031)	0.102* (0.059)	0.111* (0.062)	−0.086 (0.087)
4	−0.027 (0.035)	−0.008 (0.045)	0.020 (0.075)	0.107 (0.159)	0.008 (0.025)	0.115** (0.049)	0.076 (0.056)	−0.064 (0.068)
5	−0.033 (0.032)	−0.014 (0.041)	0.027 (0.069)	0.082 (0.138)	0.018 (0.022)	0.077* (0.044)	0.104** (0.051)	−0.096* (0.056)
6	−0.042 (0.030)	−0.007 (0.037)	0.036 (0.064)	0.066 (0.123)	0.015 (0.019)	0.038 (0.041)	0.089* (0.047)	−0.077 (0.051)
7	−0.036 (0.029)	0.006 (0.037)	0.015 (0.062)	0.060 (0.116)	−0.000 (0.018)	0.014 (0.038)	0.053 (0.043)	−0.069 (0.048)
8	−0.036 (0.028)	0.002 (0.035)	0.033 (0.058)	0.102 (0.108)	0.009 (0.016)	0.007 (0.037)	0.065 (0.042)	−0.071 (0.044)
9	−0.036 (0.027)	0.003 (0.034)	0.073 (0.055)	0.084 (0.101)	0.006 (0.016)	0.020 (0.036)	0.065* (0.039)	−0.072* (0.042)
10	−0.034 (0.026)	0.021 (0.034)	0.057 (0.053)	0.053 (0.096)	0.006 (0.015)	0.021 (0.035)	0.068* (0.037)	−0.052 (0.040)
11	−0.049* (0.025)	0.019 (0.033)	0.049 (0.051)	0.078 (0.093)	0.008 (0.015)	0.021 (0.034)	0.054 (0.036)	−0.043 (0.038)
12	−0.057** (0.024)	0.023 (0.033)	0.023 (0.050)	0.033 (0.089)	0.007 (0.015)	0.030 (0.033)	0.042 (0.034)	−0.044 (0.038)

seven to twelve months). In contrast, “negative” covariate shocks have a uniformly *positive* effect on expenditures, significant at all but the shortest intervals.

7. CONCLUSION

Using standard risk-sharing regressions to test for insurance against covariate shocks (shocks that affect relative prices) can yield very surprising results—in the example of Uganda, droughts, floods, pests, and adverse changes in prices appear to be at least partially uninsured, but to *improve* welfare as measured by “real” consumption expenditures.

We argue that these surprising results are a consequence of the standard risk-sharing regressions actually being a *joint* test of full insurance along with preferences




FIGURE 7. Effects of any covariate shock within the last m months on welfare measures with household fixed effects. Scale is in standard deviations of the dependent variable (w or $\log x$). Error bars cover a span of two standard errors about the point estimate.

being homothetic—in this case changes in relative prices affect welfare only via a single scalar price index. There is extremely strong evidence against utility being homothetic, starting with Engel (1857), since homothetic utility implies that all demands must have an income elasticity of one. This doesn’t matter much when one tests risk-sharing with respect to idiosyncratic shocks, since almost by definition these won’t affect prices. But it can matter very much when prices change, since increases in prices can very easily increase total “real” expenditures while utility actually falls.

There are two great virtues of the risk-sharing regression framework. The first is its theoretical simplicity—it’s really a model of households’ marginal utilities of expenditure (MUEs), which in a world with full insurance have a simple factor structure which is easily tested using panel methods with two-way fixed effects. The second is that assuming homothetic utility allows us to write the MUE as a simple function of nothing but total expenditures (typically $x^{-\gamma}$). And there is plenty of carefully

collected data on household expenditures which one can use to construct these MUEs for dozens of countries over many years.

However, if we want to understand insurance against covariate shocks or price changes, the evidence that we need to abandon homothetic utility is overwhelming. So are there other ways to construct estimates of MUE that depend only on expenditures? We show that there are. If we use data on *item*-level expenditures, instead of adding these up to obtain a total, there’s information based on the composition of these expenditures which can be used to estimate the MUE λ . A condition for being able to use expenditure data to estimate λ is that the expenditure system must be separable in λ and prices. When we imposing this separability and exploit the homogeneity of expenditures in prices, we show that the expenditure system can be written in a form called a generalized Pexider functional equation.

We exploit results from the theory of functional equations to obtain the entire class of possible demand systems consistent with inferring the MUE from nothing more than expenditures. These solutions fall into two families of semiparametric demands—one corresponds to a generalization of CRRA utility we call “Constant Frisch Elasticity” (CFE), while the other corresponds to a generalization of Stone-Geary utility. Both admit non-homothetic preferences and very flexible responses to changes in relative prices, but only the CFE system is easily estimated.

We use an eight round panel dataset from Uganda to estimate MUEs from expenditure data. We also obtain estimates of elasticities which emphatically reject the hypothesis of unitary income elasticities (a feature of CRRA demands). Using self-reported data on both covariate and idiosyncratic shocks, we estimate the risk-sharing regressions, using as dependent variables (a) the logarithm of total expenditures, as in the usual CRRA case; and (b) the values of $w = -\log \lambda$ estimated from the CFE expenditure system in a two-way panel regression (which also includes household demographics). As theory predicts, in this specification covariate shocks and adverse price changes have significant effects on log total expenditures (even if perfectly observed)¹⁰, and *no* significant effect on the CFE estimates of MUE, as these are constructed to account for any changes in relative prices.

This paper has focused on the case of full risk-sharing. However, the construction of estimates of MUE is independent of the efficient risk-sharing hypothesis, so estimated

¹⁰If in addition the expenditures we observe disproportionately omit certain goods or services with high income elasticities, then covariate shocks will tend to have a positive effect on total expenditures, as observed here.

MUEs could be used to estimate and test any of a large variety of dynamic life-cycle models, along lines suggested by Blundell (1998). The usual consumption Euler equation is, after all, a statement about MUEs across time, and ratios of MUEs across periods give us a way to calculate intertemporal marginal rates of substitution, free of the usual homotheticity assumptions.

8. REFERENCES

- Aczél, J., & Dhombres, J. (1989). *Functional equations in several variables* (Vol. 31). Cambridge University Press.
- Alderman, H., & Paxson, C. H. (1994). Do the poor insure? A synthesis of the literature on risk and consumption in developing countries. In *Economics in a changing world: Volume 4: Development, trade and the environment* (pp. 48–78). Springer.
- Anderson, T. W., & Rubin, H. (1956). Statistical inference in factor analysis. *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability: Held at the Statistical Laboratory, University of California, December, 1954, July and August, 1955, 1*, 111.
- Angelucci, M., & De Giorgi, G. (2009). Indirect effects of an aid program: How do cash transfers affect non-eligibles’ consumption. *The American Economic Review*, 99(1), 486–508.
- Arellano, M. (1987). Computing robust standard errors for within-groups estimators. *Oxford Bulletin of Economics and Statistics*, 49(4), 431–34.
- Arrow, K. J. (1964). The role of securities in the optimal allocation of risk bearing. *Review of Economic Studies*, 31, 91–96.
- Arrow, K. J., & Debreu, G. (1954). Existence of an equilibrium for a competitive economy. *Econometrica*, 265–290.
- Arrow, K. J., & Hahn, F. (1983). *General competitive analysis*. Elsevier.
- Attanasio, O., & Weber, G. (2010). Consumption and saving: Models of intertemporal allocation and their implications for public policy. *Journal of Economic Literature*, 48, 693–751.
- Bai, J. (2003). Inferential theory for factor models of large dimensions. *Econometrica*, 71(1), 135–171.
- Bai, J. (2009). Panel data models with interactive fixed effects. *Econometrica*, 77(4), 1229–1279.
- Banerjee, A., Duflo, E., Glennerster, R., & Kinnan, C. (2015). The miracle of micro-finance? Evidence from a randomized evaluation. *American economic journal: Applied economics*, 7(1), 22–53.
- Bardhan, P., & Udry, C. (1999). *Development microeconomics*. Oxford University Press.

- Bartholomew, D. J., Knott, M., & Moustaki, I. (2011). *Latent variable models and factor analysis: A unified approach*. John Wiley & Sons.
- Blackorby, C., & Donaldson, D. (1988). Money metric utility: A harmless normalization? *Journal of Economic Theory*, 46(1), 120–129.
- Blundell, R. (1998). Consumer demand and intertemporal allocations: Engel, Slutsky, and Frisch. *Econometric Society Monographs*, 31, 147–168.
- Browning, M., Deaton, A., & Irish, M. (1985). A profitable approach to labor supply and commodity demands over the life cycle. *Econometrica*, 53, 503–543.
- Byakatonda, J., Openy, G., Sempewo, J. I., & Mucunguzi, D. B. (2021). Over a century evidence of historical and recent dryness/wetness in sub-humid areas: A Uganda, East African case. *Meteorological Applications*, 28(5), e2028.
- Caron, J., Fally, T., & Markusen, J. R. (2014). International trade puzzles: A solution linking production and preferences. *The Quarterly Journal of Economics*, 129(3).
- Cashdan, E. (1985). Coping with risk: Reciprocity among the Basarwa of northern Botswana. *Man*, 20, 454–474.
- Christiaensen, L., Ligon, E., & Sohnesen, T. P. (2022). Consumption subaggregates should not be used to measure poverty. *The World Bank Economic Review*, 36(2), 413–432. <https://escholarship.org/uc/item/9b9929jh>
- Cochrane, J. H. (1991). A simple test of consumption insurance. *Journal of Political Economy*, 99, 957–976.
- Connor, G., & Korajczyk, R. A. (1986). Performance measurement with the arbitrage pricing theory: A new framework for analysis. *Journal of financial economics*, 15(3), 373–394.
- Deaton, A., & Zaidi, S. (2002). *Guidelines for Constructing Consumption Aggregates for Welfare Analysis*. World Bank Publications. https://www.google.com/books/edition/Guidelines_for_Constructing_Consumption/wpZgrh650wgC
- Deaton, A. (1992). Household saving in LDCs: Credit markets, insurance, and welfare. *Scandinavian Journal of Economics*, 94, 253–273.
- Dercon, S. (2005). Risk, insurance, and poverty: A review. *Insurance against poverty*, 9–37.
- Engel, E. (1857). Die productions-und consumtionsverhältnisse des königreichs sachsen. *Zeitschrift des Statistischen Biureaus des Königlich Sachsichen Ministeriums des Innern*, (8 & 9), 1–54.

- Fafchamps, M. (1992). Solidarity networks in preindustrial societies: Rational peasants with a moral economy. *Economic Development and Cultural Change*, 41(1), 147–174.
- Frisch, R. (1959). A complete scheme for computing all direct and cross demand elasticities in a model with many sectors. *Econometrica*, 27(2), 177–196.
- Hansen, B. (2022). *Econometrics*. Princeton University Press.
- Heltberg, R., Oviedo, A. M., & Talukdar, F. (2015). What do household surveys really tell us about risk, shocks, and risk management in the developing world? *The Journal of Development Studies*, 51(3), 209–225.
- Houthakker, H. S. (1960). Additive preferences. *Econometrica*, 28, 244–257. <https://www.jstor.org/stable/pdf/1907719.pdf>
- Jensen, R. T., & Miller, N. H. (2008). Giffen behavior and subsistence consumption. *American Economic Review*, 98(4), 1553–77.
- Karlan, D., Osei, R., Osei-Akoto, I., & Udry, C. (2014). Agricultural decisions after relaxing credit and risk constraints. *Quarterly Journal of Economics*, 129(2).
- Kinnan, C. (2022). Distinguishing barriers to insurance in Thai villages. *Journal of Human Resources*, 57(1), 44–78.
- Kurosaki, T. (2014). Vulnerability of household consumption to floods and droughts in developing countries: Evidence from Pakistan. *Environment and Development Economics*, 20(2), 209–235.
- Levine, D. K., & Zame, W. R. (2002). Does market incompleteness matter? *Econometrica*, 70(5), 1805–1839.
- Lewbel, A. (2010). Shape-invariant demand functions. *The Review of Economics and Statistics*, 92(3), 549–556.
- Ligon, E. (2016a). All λ -separable demands and rationalizing utility functions. *Economics Letters*, 147, 16–18. <http://dx.doi.org/10.1016/j.econlet.2016.08.010>
- Ligon, E. (2016b). Some λ -separable frisch demands with utility functions. *Economics Bulletin*, 36(1), A8. <http://escholarship.org/uc/item/1s06c2zp>
- Ligon, E. (2019). *Estimating household welfare from disaggregate expenditures* [Unpublished manuscript]. <https://escholarship.org/uc/item/3ts0g5tn>
- Ligon, E. (2023). **Cfedemands**: Tools for estimating and computing Constant Frisch Elasticity (CFE) demands. <https://pypi.python.org/pypi/CFEDemands/>
- Ligon, E., Thomas, J. P., & Worrall, T. (2002). Informal insurance arrangements with limited commitment: Theory and evidence from village economies. *Review of*

- Economic Studies*, 69(1), 209–244. <http://restud.oxfordjournals.org/content/69/1/209.full.pdf+html>
- Mace, B. J. (1991). Full insurance in the presence of aggregate uncertainty. *Journal of Political Economy*, 99, 928–956. <http://www.jstor.org/stable/2937653>
- McKenzie, D. J. (2003). How do households cope with aggregate shocks? Evidence from the Mexican peso crisis. *World Development*, 31(7), 1179–1199.
- Morduch, J. (1995). Income smoothing and consumption smoothing. *Journal of Economic Perspectives*, 9(3), 103–114.
- Munshi, K., & Rosenzweig, M. (2016). Networks and misallocation: Insurance, migration, and the rural-urban wage gap. *American Economic Review*, 106(01), 46–98.
- Ogaki, M., & Zhang, Q. (2001). Decreasing relative risk aversion and tests of risk sharing. *Econometrica*, 69(2), 515–526.
- Platteau, J.-P., & Abraham, A. (1987). An inquiry into quasi-credit contracts: The role of reciprocal credit and interlinked deals in small-scale fishing communities. *Journal of Development Studies*, 23(4), 461–490.
- Ray, D. (1998). *Development economics*. Princeton University Press.
- Santàeulalia-Llopis, R., & Zheng, Y. (2018). The price of growth: Consumption insurance in China 1989–2009. *American Economic Journal: Macroeconomics*, 10(4), 1–35.
- Thomas, D., & Frankenberg, E. (2007). Household responses to the financial crisis in Indonesia: Longitudinal evidence on poverty, resources, and well-being. In *Globalization and poverty* (pp. 517–560). University of Chicago Press.
- Townsend, R. M. (1978). On the optimality of forward markets. *The American Economic Review*, 68(1), 54–66. <https://www.jstor.org/stable/pdf/1809687.pdf>
- Townsend, R. M. (1994). Risk and insurance in village India. *Econometrica*, 62(3), 539–591.
- Townsend, R. M. (1995). Consumption insurance: An evaluation of risk-bearing systems in low-income economies. *The Journal of Economic Perspectives*, 9(3), 83–102.
- Williams, B. (2020). Identification of the linear factor model. *Econometric Reviews*, 39(1), 92–109.
- Zhang, Q., & Ogaki, M. (2004). Decreasing relative risk aversion, risk sharing, and the permanent income hypothesis. *Journal of Business and Economic Statistics*, 22(4), 421–430.

APPENDIX A. APPENDICES

A.1. Data on Shocks. Data on self-reported shocks comes from a “shocks” module in different rounds of the Ugandan LSMS surveys. There are minor differences in elicitation and in the list of shocks reported in certain years. Table A.1 reports the slight aggregation of shocks we use.

TABLE A.1. Different shock labels across eight rounds of Ugandan data, with harmonized labels.

Existing Label	Label
Conflict/Violence	Conflict
Death of Income Earner(s)	Death of earner
Death of Other Household Member(s)	Death
Drought	Drought
Drought/Irregular Rains	Drought
Erosion	Erosion
Fire	Fire
Floods	Floods
Irregular Rains	Drought
Landslides	Erosion
Landslides/Erosion	Erosion
Loss of Employment of Previously Employed Household Member(s) (Not Due to Illness or Accident)	Lost Earnings
Other (Specify)	Other
Reduction in the Earnings of Currently (Off-Farm) Employed Household Member(s)	Lost Earnings
Serious Illness or Accident of Income Earner(s)	Health
Serious Illness or Accident of Other Household Member(s)	Health
Theft	Theft
Theft of Agricultural Assets/Output (Crop or Livestock)	Theft
Theft of Money/Valuables/Non-Agricultural Assets	Theft
Unusually High Costs of Agricultural Inputs	Prices
Unusually High Level of Crop Pests & Disease	Pests
Unusually High Level of Crop Pests & Disease	Pests
Unusually High Level of Livestock Disease	Pests
Unusually Low Prices for Agricultural Output	Prices

TABLE A.2. Incidence of shocks by round

Shock	2005	2009	2010	2011	2013	2015	2018	2019
Conflict/Violence	268	34	27	46	11	12	22	25
Death of Income Earner(s)	99	27	17	19	30	19	20	19
Death of Other Household Member(s)	423	74	58	35	66	35	44	49
Drought	1234	—	—	—	735	526	556	360
Drought/Irregular Rains	—	1344	710	560	—	—	—	—
Erosion	—	—	—	—	15	3	19	13
Fire	105	26	21	20	17	18	12	9
Floods	426	61	102	148	98	62	74	117
Irregular Rains	—	—	—	—	179	72	180	169
Landslides	—	—	—	—	1	—	2	2
Landslides/Erosion	—	21	5	17	—	—	—	—
Loss of Employment of Previously Employed Household Member(s) (Not Due to Illness or Accident)	—	9	10	8	5	10	5	6
Other (Specify)	111	101	60	57	62	45	50	75
Reduction in the Earnings of Currently (Off-Farm) Employed Household Member(s)	—	28	3	10	6	4	15	15
Serious Illness or Accident of Income Earner(s)	82	189	152	91	86	56	107	119
Serious Illness or Accident of Other Household Member(s)	—	188	149	65	47	32	83	78
Theft	349	—	—	—	—	—	—	—
Theft of Agricultural Assets/Output (Crop or Livestock)	—	127	48	21	35	27	41	40
Theft of Money/Valuables/Non-Agricultural Assets	—	106	48	34	41	35	34	43
Unusually High Costs of Agricultural Inputs	71	60	19	27	49	7	10	12
Unusually High Level of Crop Pests & Disease	—	137	40	61	54	35	—	—
Unusually High Level of Crop Pests & Disease	292	—	—	—	—	—	120	79
Unusually High Level of Livestock Disease	183	82	37	31	17	18	10	5
Unusually Low Prices for Agricultural Output	—	53	35	38	18	5	68	17

A.2. Proof of Theorem 1. In this appendix we provide a proof of Theorem 1. We first supply a lemma pertaining to the homogeneity of λ -separable expenditure systems, and then provide solutions to the generalized Pexider equation. With these preliminary results in hand we establish that any rationalizable system of expenditures which is λ -separable takes the form of the generalized Pexider equation, and map the general solutions of the equation into corresponding demands and utilities.

A Lemma Pertaining to Homogeneity.

Lemma A.1. *If demand for good i satisfies Condition 1 and Condition 2, then the functions ϕ_i , a_i and b_i are either all logarithmic or ϕ_i and a_i are both positive homogeneous of some degree σ_i , while b_i is positive homogeneous of degree $-\sigma_i$.*

Proof. Any rationalizable expenditures x_i must be homogeneous of degree one in $(p, 1/\lambda)$. Exploiting Condition 2 then implies that

$$x_i = \phi_i^{-1}(a_i(p) + b_i(\lambda))$$

is similarly homogeneous of degree one. The function ϕ_i must then either be homogeneous of degree σ_i , with $\phi_i(x_i) = x_i^{\sigma_i}$, or else $\phi_i(x_i) = \log(x_i)$. In either case Frisch quantities can be written as

$$c_i = f_i(p\lambda) = \frac{1}{p_i} \phi_i^{-1}(a_i(p) + b_i(\lambda)) - d_i(p)$$

for some function d_i homogeneous of degree zero.

We consider the power and logarithmic cases in turn.

First suppose that $\phi_i(x) = x^{\sigma_i}$. Then the sum $a_i + b_i$ must also be homogeneous of degree σ_i in $(p, 1/\lambda)$, and the individual functions a_i and b_i respectively either homogeneous of degree σ_i and $-\sigma_i$ or else the zero function. It follows that $f_i(p, r) = \phi_i^{-1}(a_i(p)/p_i^{\sigma_i} + b_i(\lambda)/p_i^{\sigma_i}) - d_i(p)$, and that $a_i(p)/p_i^{\sigma_i}$ and $b_i(\lambda)/p_i^{\sigma_i}$ are either zero or positive homogeneous of degree zero, so that

$$(a_i(p\theta) + b_i(\lambda/\theta)) = \theta^{\sigma_i}(a_i(p) + b_i(\lambda)) = \theta^{\sigma_i}a_i(p) + \theta^{\sigma_i}b_i(\lambda)$$

for any positive scalar θ . Differentiating this with respect to $1/\lambda$ establishes that b'_i is homogeneous of degree $\sigma_i - 1$, so that b_i is homogeneous of degree σ_i (by Euler's theorem of positive homogeneous functions). A similar argument involving the gradient with respect to p establishes the same for a_i .

For the logarithmic case, $\phi(x_i) = \log(x_i) = \log(p_i) + \log(c_i)$ implies that

$$f_i(p, \lambda) + d_i(p) = \exp(a_i(p) + b_i(\lambda) - \log(p_i))$$

which must be positive homogeneous of degree zero in $(p, 1/\lambda)$. This implies that for any $\theta > 0$

$$a_i(\theta p) + b_i(\lambda/\theta) - \log(\theta p_i) = a_i(p) + b_i(\lambda) - \log(p_i),$$

which in turn implies that

$$a_i(\theta p) + b_i(\lambda/\theta) = a_i(p) + b_i(\lambda) + \log(\theta),$$

implying that both a_i and b_i are linear in logs of (p, λ) . □

Generalized Pexider Equation Applied to Vector Spaces. We now introduce our main tool for solving the functional equations implied by separability and rationalizability; this tool is an application of what is called the generalized Pexider equation, when the domain of application is limited to real vector spaces.

Consider the generalized Pexider equation

$$(10) \quad k(x + y) = g(x)l(y) + h(y)$$

where

$$(11) \quad g(x) = \frac{k(x) - h(0)}{l(0)}$$

$$(12) \quad \varphi(y) = \frac{l(y)}{l(0)}$$

$$(13) \quad \psi(y) = h(y) - h(0) \frac{l(y)}{l(0)}$$

$$(14) \quad k(x + y) = k(x)\varphi(y) + \psi(y)$$

$$(15) \quad \kappa(x) = k(x) - k(0)$$

$$(16) \quad \kappa(x + y) = \kappa(x)\varphi(y) + \kappa(y).$$

Next we give statements of two related lemmata. The first is just a statement of the solution of the well-known functional equation of Cauchy applied to real vector spaces; the second is a statement of the solution to what is sometimes called Cauchy's exponential equation, again for real vector spaces.

Lemma A.2. *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, with f continuous at a point. Then if*

$$(17) \quad f(x + y) = f(x) + f(y)$$

then $f(x) = Cx$ for some constant $m \times n$ matrix C .

Also

Lemma A.3. *Let $h : \mathbb{R}^n \rightarrow \mathbb{R}^m$. If*

$$h(x + y) = h(x)h(y)$$

then either $h(x) = 0$ or $h(x) = e^{f(x)}$, where f is an arbitrary solution to Cauchy's equation (17).

Corollary 1. *Any solution to the functional equation of Lemma A.3 which is continuous and non-constant is of the form*

$$h(x) = \exp(Cx),$$

where C is a constant matrix and the exp operator is element by element.

The following is just a restatement of Theorem 15.1 of Aczél and Dhombres, 1989, and describes all solutions to the generalized Pexider equation (10) over the general domain of Abelian groupoids.

Theorem 2. *For any x, y in an Abelian groupoid, solutions to (10) will satisfy one of:*

- (1) *If $\varphi(x) = 1$ for all x , then $\kappa(x)$ is an arbitrary function; $\psi(x) = \kappa(x)$; and $k(x) = \kappa(x) + B$. Or;*
- (2) *if $\varphi(x_0) \neq 1$ for some x_0 , then we have $C = \frac{\kappa(x_0)}{\varphi(x_0)-1}$; and $\kappa(x) = C[\varphi(x) - 1]$; and two sub-cases:*
 - (a) *$C = 0$; $\kappa(x) = 0$; $\varphi(x)$ arbitrary; $k(x) = B$; $\psi(y) = B(1 - \varphi(y))$; or*
 - (b) *$C \neq 0$; $k(x) = C\varphi(x) + B$; $\psi(x) = B(1 - \varphi(x))$; where $\varphi(x)$ satisfies $\varphi(x + y) = \varphi(x)\varphi(y)$ (Cauchy's exponential equation); and where $\kappa(x)$ satisfies $\kappa(x + y) = \kappa(x) + \kappa(y)$ (Cauchy's equation).*

If we restrict the domain under consideration to a real vector space, then we can give explicit solutions to (10), as follows:

Proposition 1. *For any $x, y \in \mathbb{R}^n$, solutions to (10) will satisfy one of:*

- (1) *If $\varphi(x) = 1$ for all x , then $\kappa(x) = \psi(x) = Cx$ and $k(x) = Cx + B$, where $B \in \mathbb{R}^m$. Or;*
- (2) *if $\varphi(x_0) \neq 1$ for some x_0 , then we have $C = \frac{\kappa(x_0)}{\varphi(x_0)-1}$; and $\kappa(x) = C[\varphi(x) - 1]$; and two sub-cases:*
 - (a) *$C = 0$; $\kappa(x) = 0$; $\varphi(x)$ arbitrary; $k(x) = B$; $\psi(y) = B(1 - \varphi(y))$; or*
 - (b) *$C \neq 0$; $k(x) = C\varphi(x) + B$; $\psi(x) = B(1 - \varphi(x))$; $\kappa(x) = Cx$; and one of:*
 - (i) $\varphi(x) = 0$;

- (ii) $\varphi(x) = \exp(Ax)$; or
- (iii) $\varphi(x) = \exp(f(x))$, f nowhere continuous.

Proof. Just a specialization of Theorem 2 to the case in which domain is a real vector space, which then allows subsequent application of Lemma A.3 and Lemma A.2. $\square$

Proof of Theorem.

Proof. First, Lemma A.1 establishes that (ϕ_j, a_j, b_j) in (8) are all either logarithmic or positive homogeneous of some degree $(-)\sigma_j$.

In the logarithmic case the logarithm of demand for good j can be written $\log f_j(\lambda p) = [-\log p_j + a_j(p)] + b_j(\lambda)$, which not only has expenditures λ -separable, but also quantities λ -separable. Then the main result of Ligon (2016a) applies, with $\phi_j = \log$, yielding the result that $\log(f_j(\lambda p)) = \tilde{\alpha}_j - \beta_j \log(p_j \lambda)$. Let $c_j = f_j(\lambda p)$ and solve for $p_j \lambda$, obtaining $p_j \lambda = \alpha_j c_j^{-1/\beta_j}$, where $\alpha_j = e^{\tilde{\alpha}_j}$ must be positive.

In the homogeneous case, we have

$$(p_j f_j(\lambda p))^{\sigma_j} = a_j(p) + b_j(\lambda),$$

or

$$f_j(\lambda p)^{\sigma_j} = p_j^{-\sigma_j} a_j(p) + p_j^{-\sigma_j} b_j(\lambda).$$

This takes the form of the generalized Pexider equation (10), with $x = \log \lambda$ and y the vector $\log p$, when the vector-valued function $k(x + y) = [f_j(\exp(x + y))^{\sigma_j}]$, $h(y) = [a_j(\exp(y))e^{-\sigma_j y_j}]$, $g(x) = [b_j(\exp(x))]$, and $\ell(y) = [e^{-\sigma_j y_j}]$. Now, we seek to apply Proposition 1, which gives solutions to the system of functional equations 10–16. Part of this system is the function $\varphi(y)$. Using our knowledge that $\ell_j(y) = e^{-\sigma_j y_j}$ and (13), it follows that in this equation the function $\varphi(y) = \ell(y) = [e^{-\sigma_j y_j}]$. Now, consulting the different possible cases of Proposition 1 we see that with this solution of φ the only cases that can apply are the cases indicated by 2a and 2bii. The former implies that $f_j(\lambda p)^{\sigma_j}$ is a constant, so that (to be consistent with the properties of Frisch demands) $\sigma_j = 0$. But then the function ϕ_j isn't increasing, and the only solutions that are relevant to our problem are the solutions 2bii. These imply that $k(z) = Cz + B$, with C and B constant matrices, and $\varphi(y) = \exp(Cy)$. Thus C is a diagonal matrix, with diagonal elements $-\sigma_j$. Using equations (14) and (15) we obtain $k(x + y) = (Cx + B)e^{Cy} + [h(y) - h(0)\varphi(y)]$; then using our definition of $h(y)$ in terms of p gives us $f_j(\lambda p)^{\sigma_j} = \alpha_j / (p_j \lambda)^{\sigma_j} + \beta_j$. Noting that $u_j(c) = p_j \lambda$ and solving for this gives us the solution for marginal utilities. $\square$

A.3. Details of Estimation. Here we describe the steps involved in our estimation of the demands given by (9). We assume that observed household characteristics z_{it} affect log expenditures linearly, and that all households within a given region-year in Uganda face the same relative prices.¹¹ The latter assumption allows us to estimate the influence of prices by using region-year dummy variables, while the former allows us to express the influence of household characteristics on demand for good j as $\gamma_j z_{it}$, where γ_j is a vector of parameters to be estimated, and where t is understood to index region-years. Substitute $w_{it} = -\log \lambda_{it}$, and let $y_{it}^j = \log x_{it}^j$. Then (9) becomes

$$(18) \quad y_{it}^j = \gamma_j' z_{it} + a_{jt} + \beta_j w_{it} + \epsilon_{it}^j.$$

If we observed either β_j or w_{it} this would be an entirely standard linear regression problem. We do not, making this what Hansen (2022) calls a factor model with additional regressors.¹²

We make four standard assumptions. First, that (z_{it}, w_{it}) are orthogonal to the error ϵ_{it}^j , the usual identifying assumption in a linear regression context. Second, that w_{it} is orthogonal to z_{it} . While not necessary for identification (see Williams, 2020), this assumption simplifies both the interpretation and estimation of w_{it} . Third, that errors ϵ_{it}^j are homoskedastic across goods, though we allow clustering at the region-year level, so that $\text{Var}(\epsilon_{it}^j) = \text{Var}(\epsilon_{it}^{j'})$. Both the second and third assumptions are testable, and can be relaxed. Finally, as with any factor analysis some normalization is required for identification; here it's convenient to take $\text{Var}(w_{it}) = 1$.

Estimation is straight-forward, but requires three steps:

- (1) Regress log expenditures on observed characteristics and time-region dummies

$$y_{it}^j = a_{jt} + \gamma_j' z_{it} + e_{it}^j,$$

yielding estimates of the vector of the parameter vector γ and residuals $\hat{e}_{it}^j$.

- (2) The residuals $\hat{e}_{it}^j$ provide an unbiased estimate of $\beta_j w_{it} + \epsilon_{it}^j$. Let $\hat{e}_{it}$ denote the J -vector of residuals for household i at t , and let $\Sigma = \text{E } e_{it} e_{it}^\top$ be its covariance matrix. From (18) and our assumptions above this implies

$$\Sigma = \beta \beta^\top + \sigma \mathbf{I}.$$

¹¹Our data includes four regions over eight years, resulting in a total of 32 different relative price vectors. If *levels* of prices differ within a region this will be reflected in marginal utilities of expenditure.

¹²Our estimation approach follows Hansen's procedure, though by assuming that w_{it} is orthogonal to z_{it} we avoid the need for iteration.

This is a classical factor analysis problem, and the (quasi-) maximum likelihood estimate of β is proportional to the first principal component of $\hat{\Sigma}$ (Anderson & Rubin, 1956). Our normalization $\text{Var}(w_{it}) = 1$ supplies the missing factor of proportionality.

- (3) Finally, we use a standard “regression” method (Bartholomew et al., 2011) to estimate the w_{it} . We return to (18), but now treat the w_{it} as unknown parameters to be estimated in the regression

$$y_{it}^j = a_{jt} + \gamma_j' z_{it} + \hat{\beta}_j w_{it} + \tilde{\epsilon}_{it}^j,$$

where the difference from the original (18) is that now we can include the estimated $\hat{\beta}_j$ as (generated) regressors.

Observations.

- Of the three steps outlined above, two are least squares regressions, the properties of which will be familiar to economists. The second step employs factor analysis—a widely used procedure in other fields but less common in economics. A key consideration in factor analysis concerns the assumed structure of the disturbance covariance matrix, say $\Psi = \frac{1}{NT} \text{E} \sum_{i,t} \epsilon_{it} \epsilon_{it}^\top$. Here we assume a form of homoskedasticity, with $\Psi = \sigma \mathbf{I}$, in which case the first principal component of $\hat{\Sigma}$ provides an unbiased and efficient estimate of β . Some alternative assumptions:
 - If Ψ is diagonal, the standard approach (Anderson & Rubin, 1956) uses maximum likelihood estimation assuming normally distributed disturbances. Even when this distributional assumption isn’t satisfied, the (quasi-) maximum likelihood estimator retains good properties and can be implemented through a simple iterative procedure.
 - If Ψ is unrestricted, Bai (2003) provides conditions under which our principal components estimator still possesses desirable properties. The caveat is that the argument is asymptotic, requiring both J and TN to approach infinity. In our application, while TN is reasonably large, J is measured only in dozens.
 - Connor and Korajczyk (1986) establish consistency of our principal components estimator under a weaker sufficient and (Bai, 2003) necessary condition than our assumption that $\Psi = \sigma \mathbf{I}$. Instead, they require only

an asymptotic version of this, requiring that errors across goods be uncorrelated and homoskedastic as NT grows large, but holding the number of goods J fixed.

- Williams (2020) establishes identification results for linear factor models for fixed J and very modest restrictions on Ψ . However, to our knowledge, no practical estimators have yet been developed for this case.
- With a consistent estimator of Σ , we can estimate β . Notably, it is possible to estimate Σ consistently even when some food expenditure data is missing, as is the case in our application.

UNIVERSITY OF CALIFORNIA, BERKELEY