

Nonparametric welfare and demand analysis with unobserved individual heterogeneity

Sam Cosaert^{*}

Thomas Demuynck[†]

Abstract

This paper combines elementary revealed preference principles and nonparametric estimation techniques to obtain nonparametric bounds on the distribution of the money metric utility over a population of heterogeneous households. Our approach is independent of any functional specification on the household utility functions, meaning that our results are robust against parametric specification errors. Our methodology can also be used to establish bounds on the distribution of the demand function in counterfactual price regimes. To demonstrate the relevance of our approach, we illustrate our findings using a repeated cross-sectional household consumption data set.

JEL-codes: D12, C14

Keywords: unobserved heterogeneity, stochastic revealed preference, household demand

1 Introduction

We present a framework to construct nonparametric bounds on the distribution of the money metric utility function while taking into account individual unobserved heterogeneity. Our ap-

^{*}K.U.Leuven, E.Sabbelaan 53, 8500 Kortrijk, Belgium, email: sam.cosaert@kuleuven.be

[†]Corresponding author, Maastricht University, Tongersestraat 53, 6711 LM Maastricht, Netherlands, email: t.demuynck@maastrichtuniversity.nl

proach combines elementary revealed preference concepts, in particular the Weak Axiom of Revealed Preference, with nonparametric estimation techniques. In this manner, our approach remains independent of any parametric specification on the underlying household utility functions or on the unobserved heterogeneity distribution. We further demonstrate how the framework can be used to establish bounds on the distribution of the demand functions in counterfactual price regimes. An illustration using the Consumer Expenditure Survey, a US cross-sectional household consumption data set, demonstrates the practical usefulness of our results.

Motivation Demand analysis provides a powerful tool to analyse behavioural responses and welfare effects due to price and income variations. In a typical demand study, the researcher first estimates the parameters of some parametric demand system,¹ and uses these estimates to calculate the associated indirect utilities. This ‘parametric’ approach has two major shortcomings. The first is that the outcome is sensitive to the specific functional structure chosen by the researcher. Imposing the wrong functional form can therefore severely bias the resulting analysis. A second shortcoming concerns the treatment of individual (unobserved) heterogeneity. In a typical consumer data set, we observe individuals or households only once. Given this data limitation, it is often assumed that similar looking individuals have similar preferences. Many demand studies therefore model a household’s demand to equal a rational systematic component, from a common utility function across all (similar looking) households, and a household specific additive error term capturing the unobserved heterogeneity or taste variation. By controlling for various observable characteristics (like household size), it is hoped that the issue of heterogeneity across the households is adequately addressed by including such additive error term. This assumption, however, disregards the finding that individuals who look very similar may actually differ dra-

¹Popular parametric demand systems are the Translog (Christensen, Jorgenson, and Lau, 1975), the Almost Ideal (Deaton and Muellbauer, 1980), or the Quadratic Almost Ideal (Banks, Blundell, and Lewbel, 1997) demand system.

matically in their actual choice behaviour.² As shown by Lewbel (2001), imposing additivity of the unobserved heterogeneity is a strong assumption. Its resulting implications come very close to enforcing a representative agent assumption.³ To summarize, we see that different people (although they may look the same) have different tastes and, consequentially, behave differently. In order to take this into account, it is crucial to allow for non-additive unobserved heterogeneity.

Literature overview In order to deal with aforementioned two problems, one can distinguish between two approaches. The first approach looks at the nonparametric differential ‘smooth’ restrictions that can still be established in a heterogeneous population. These usually take the form of population level generalizations of Slutsky symmetry, negativity and homogeneity. Recent examples that follow this approach are Hoderlein (2011), Blundell, Horowitz, and Parey (2013), Hausman and Newey (2013), Hoderlein and Vanhems (2013), and Dette, Hoderlein, and Neumeyer (2014). A second approach, followed in this paper, is to rely on revealed preference theory. Revealed preference theory was initiated by Samuelson (1938), Houthakker (1950) and further developed in several seminal contributions by Afriat (1967), Diewert (1973) and Varian (1982). The main aim of revealed preferences theory is to establish (combinatorial) restrictions on observed demand behaviour of a certain individual or household such that it is consistent with the classical model of utility maximization subject to a budget constraint. One of the main advantages of revealed preference theory is that it imposes no functional restrictions on the underlying utility function, except for some regularity conditions like local non-satiation.

Revealed preference theory, as it was initially developed, has two main problems. First, from an empirical point of view, the method does not really seem to provide very tight bounds. The main reason for this is that relative price variations usually tend to be quite small in comparison

²Unobserved heterogeneity is often seen as the main reason why demand estimations on cross sectional data typically have low r-squared values.

³See also Brown and Walker (1989) and McElroy (1987) for a discussion of other issues when taking into account unobserved heterogeneity.

to income variation. This implies that budget hyperplanes often do not cross. We refer to Bronars (1987) and Varian (1982) for a discussion of this problem. The second problem is that revealed preference theory is not well suited to deal with unobserved individual heterogeneity. As a result, most of its applications remain confined to a few panel consumption data sets, where the same household or individual is observed over multiple periods.

The first problem has been the subject of several recent studies that apply revealed preference theory to repeated cross sectional data by combining insights from revealed preference theory with nonparametric estimation techniques (see Blundell (2005); Blundell, Browning, and Crawford (2003, 2007, 2008) and Blundell, Browning, Cherchye, Crawford, De Rock, and Vermeulen (2015)). The main contribution from this literature is that it shows how to use nonparametric Engel curve demand estimates as an input for revealed preference analysis. If we assume that households in the same time period and location face the same relative prices, then the nonparametric Engel curves estimate the mean (or average) expansion paths for each price regime. The availability of these expansion paths greatly improves the nonparametric bounds on various welfare related concepts and on the counterfactual demand estimates that can be obtained using revealed preference techniques.

A remaining drawback of this approach is the way it deals with the issue of unobserved heterogeneity. Given that the Engel curve estimates are obtained from a mean regression, the methodology is subject to Lewbel (2001)'s critique: imposing revealed preference restrictions on the mean Engel curve estimates comes very close to imposing a representative consumer assumption. Given this, the approach does not fully address the individual heterogeneity problem. Moreover, despite the fact that the procedure has the potential to produce tight bounds on the 'representative' money metric utility and demand functions, it does not give us any information concerning the distribution of these functions across the heterogeneous population.

A useful extension of revealed preference theory that explicitly takes into account individual heterogeneity is Stochastic Revealed Preference Theory, initiated by McFadden and Richter (1971)

and Falmagne (1978).⁴ We refer to McFadden (2005) for an overview of the literature. Stochastic revealed preference takes as input the entire distribution of demand behaviour over a heterogeneous population of households for a finite number of budget sets.⁵ Therefore, it is well suited to deal with the issue of unobserved heterogeneity. The literature has put forward several rationality axioms (e.g. the Axiom of Revealed Stochastic Preference and the Weak Axiom of Stochastic Revealed Preference) that provide conditions on the distributions of choices such that a population of individuals is consistent with rational choice theory, which postulates that individuals are preference maximizers. Although the literature is mainly theoretical, several recent papers developed statistical tests to verify whether the stochastic revealed preference axioms are satisfied in reality. Hoderlein and Stoye (2014a) derive a statistical procedure to infer bounds on the fraction of the population that violates the Weak Axiom of Stochastic Revealed Preference. Kitamura and Stoye (2013) derive a statistical test to verify whether a population of heterogeneous households satisfies the Axiom of Stochastic Revealed Preference for a finite collection of budget sets, thereby explicitly taking into account transitivity of the preference relations. Finally, Kawaguchi (2012) derives several procedures to test the validity of various axioms of revealed stochastic preference. Interestingly, these studies find little evidence that the stochastic revealed preference restrictions are violated. The main difference between these papers and ours is the focus. While the existing contributions mainly deal with testing whether the axioms imposed by the stochastic revealed preference literature hold, we are more interested in the restrictions that the stochastic revealed preference axioms impose on the resulting distribution of the money metric utility and demand functions. In the terminology of Varian (1982): while above papers deal with *testing* the theory, we concentrate on the *recovery* of the underlying structure of the model.

⁴See also Block and Marschak (1959), McFadden (1975), Fishburn (1978), Cohen (1980), Barberá and Pattanaik (1986), Fishburn and Falmagne (1989), Cohen and Falmagne (1990), Fishburn (1992) and Bandyopadhyay, Dasgupta, and Pattanaik (1999) for other contributions.

⁵A second interpretation of stochastic revealed preference theory is that the demand behavior is generated by a single household with a random utility function.

Another closely related paper is Blundell, Kristensen, and Matzkin (2014). These authors focus on the issue of unobserved heterogeneity in a two goods setting. In particular, they tackle the problem of individual unobserved heterogeneity using nonparametric quantile demand estimates in combination with standard revealed preference tests (i.e. SARP). Hoderlein and Stoye (2014b) recently showed that in a two goods setting, imposing the usual revealed preference axioms on the quantile demands is equivalent to imposing the Axiom of Stochastic Revealed Preference on the entire data set.⁶ The analysis of Blundell, Kristensen, and Matzkin (2014) is based on an invertibility (monotonicity) condition on the unobserved heterogeneity term. In our framework, we abstain from imposing such condition.

Contribution The main contribution of this paper is to derive nonparametric bounds on the money metric utility functions and the demand functions without imposing any functional structure on the household utility functions and the unobserved heterogeneity structure. As such, we avoid the problem that our results might be biased because of a wrong functional specification or because the households do not satisfy the ‘representative agent’ condition. We establish our results by combining elementary stochastic revealed preference theory and nonparametric estimation techniques. Our framework not only allows us to derive bounds on the mean of the money metric utility and demand functions, but on the entire distribution of these functions over the heterogeneous population. This provides important additional information concerning the distribution of welfare and demand over the population.

In order to obtain our results, we exploit the Weak Axiom of Revealed Preferences (WARP) applied to a population of heterogeneous households. Although this axiom is weaker than the revealed preference axioms that exploit transitivity (e.g. the Strong Axiom of Revealed Preference), we nevertheless show that it is powerful enough to establish narrow bounds. We demonstrate the

⁶In a two-goods setting, the analysis is simplified by the fact that the Weak and Strong Axioms of (Stochastic) Revealed Preference coincide. In other words, imposing transitivity implies no additional testable implications, see Rose (1958).

usefulness of our results by applying it to the Consumer Expenditure survey, a US cross sectional consumption data set.

Outline In section 2, we set out our framework and we present the necessary notation, concepts and definitions for the remaining part of the paper. Section 3 establishes the theoretical results that provide the nonparametric bounds on the distribution of the money metric utility function and the demand functions. Section 4 contains our empirical application. We discuss estimation, statistical inference and we present several results. Section 5 concludes and points towards future research.

2 Notation and Definitions

We consider an economy with a large (infinite) number of households described by a probability space (J, Ω, P) . Household $h \in J$ is endowed with a utility function which we denote by $u_h(\mathbf{q}^h; \mathbf{a}^h)$. This function depends on a (column) vector of consumed goods $\mathbf{q}^h \in \mathbb{R}_+^n$, where n is the number of goods, and a vector of observable household specific attributes \mathbf{a}^h , e.g. household composition. Unobserved preference heterogeneity is captured by the fact that u_h depends on h . For a price vector $\mathbf{p} \in \mathbb{R}_{++}^n$ and an expenditure level $x \in \mathbb{R}_+$, we denote by (\mathbf{p}, x) , the budget set consisting of all bundles \mathbf{q} such that $\mathbf{p}\mathbf{q} \leq x$. In order to decide how much to consume, the household maximizes its utility function subject to a household budget constraint,⁷

$$\mathbf{q}_h(\mathbf{p}, x^h; \mathbf{a}^h) = \arg \max_{\mathbf{q}} u_h(\mathbf{q}; \mathbf{a}^h) \text{ s.t. } \mathbf{p}\mathbf{q} \leq x^h.$$

Utility functions are strictly quasi-concave and twice continuously differentiable in \mathbf{q} such that the demand functions are well defined and continuous in \mathbf{p} and x . For all income levels x and

⁷We abstract from the problem that households are typically composed of several individuals (Chiappori, 1988, 1992; Cherchye, De Rock, and Vermeulen, 2007).

prices \mathbf{p} and all measurable sets $A \in \Omega$ we assume that,

$$P(h \in A | \mathbf{p}, x, \mathbf{a}) = P(h \in A | \mathbf{a}).$$

This condition says that, conditional on all observable attributes, the unobserved heterogeneity is independent of prices and income. This ‘independence of budget sets’ condition is common in the literature.⁸ If we interpret unobserved heterogeneity as preference heterogeneity, it encompasses the idea, common in consumer demand, that preferences do not vary with prices and income. For notational convenience, we omit from now (until section 4.2) the dependence on the observable attributes \mathbf{a} , taking into account that every expression is valid conditional on a particular value of this vector.

For the remaining part of the paper, it will be more useful to work with the indirect utility function $v_h(\mathbf{p}, x^h)$ which gives the maximal utility that household h can obtain at prices \mathbf{p} and income x^h . The indirect utility function is defined from the direct utility function by,

$$v_h(\mathbf{p}, x^h) \equiv u_h(\mathbf{q}_h(\mathbf{p}, x^h)).$$

The indirect utility function is strictly increasing in the level of disposable income x^h . If we invert the indirect utility function $v_h(\mathbf{p}, x^h)$, with respect to x^h , we obtain the expenditure function $e_h(\mathbf{p}, u^h)$ which gives the minimal outlay for household h to reach utility level u^h at prices \mathbf{p} . Finally, using the expenditure function, we can define the money metric utility function,

$$\mu_h(\mathbf{p}_v, \mathbf{p}_t, x^h) \equiv e_h(\mathbf{p}_v, v_h(\mathbf{p}_t, x^h)).$$

The money metric utility $\mu_h(\mathbf{p}_v, \mathbf{p}_t, x^h)$ gives the minimal amount of expenditure that household

⁸See for example, Lewbel (2001, equation 4), Hausman and Newey (2013, Assumption1), Blundell, Kristensen, and Matzkin (2014, Assumption A.1) and Bhattacharya (2015).

h needs at prices \mathbf{p}_v to be equally well off as it would have been when facing prices \mathbf{p}_t and income x^h . The money metric utility lies at the basis of many cost of living indices. In particular, given two price vectors \mathbf{p}_t and \mathbf{p}_v and some reference budget (\mathbf{p}, x) , the Konüs cost of living index, describing the price increase from \mathbf{p}_t to \mathbf{p}_v , is defined as,

$$\frac{\mu_h(\mathbf{p}_v, \mathbf{p}, x)}{\mu_h(\mathbf{p}_t, \mathbf{p}, x)}$$

There are two natural choices for \mathbf{p} , namely \mathbf{p}_t and \mathbf{p}_v . Setting \mathbf{p} equal to the initial price \mathbf{p}_t gives the Laspeyres-Konüs cost of living index,

$$\frac{\mu_h(\mathbf{p}_v, \mathbf{p}_t, x)}{x}$$

If we set \mathbf{p} equal to the final price vector \mathbf{p}_v , we obtain the Paasche-Konüs cost of living index,

$$\frac{x}{\mu_h(\mathbf{p}_t, \mathbf{p}_v, x)}$$

Both indices are used to describe the increase in the cost necessary to maintain the same living standard over time. Their distributions can easily be constructed provided that we know the distribution of the money metric utility function. The money metric utility also provides a cardinalisation of the utility function in the sense that for any reference price vector \mathbf{p} and for any two budgets (\mathbf{p}_t, x) and (\mathbf{p}_v, y) :

$$\mu_h(\mathbf{p}, \mathbf{p}_t, x) \geq \mu_h(\mathbf{p}, \mathbf{p}_v, y) \iff v_h(\mathbf{p}_t, x) \geq v_h(\mathbf{p}_v, y).$$

As such, the difference in the money metric can be used as a measure for the welfare difference for two different budgets: if (\mathbf{p}_t, x) is the old budget and (\mathbf{p}_v, y) is the new one, then this welfare

change can be measured by,

$$\mu_h(\mathbf{p}, \mathbf{p}_v, y) - \mu_h(\mathbf{p}, \mathbf{p}_t, x)$$

Again, there are two obvious choices for the base price vector \mathbf{p} , namely \mathbf{p}_t or \mathbf{p}_v . The first leads to the equivalent variation,

$$EV = \mu_h(\mathbf{p}_t, \mathbf{p}_v, y) - x.$$

The second gives the compensating variation,

$$CV = y - \mu_h(\mathbf{p}_v, \mathbf{p}_t, x).$$

Revealed preferences The analysis in the following sections depends on a simple revealed preference idea. Fix a household h and consider two distinct budgets (\mathbf{p}_t, x) and (\mathbf{p}_v, y) . If the household is utility maximizing, then the following condition must hold,

$$\text{If } x \geq \mathbf{p}_t \mathbf{q}_h(\mathbf{p}_v, y) \text{ then } v_h(\mathbf{p}_t, x) > v_h(\mathbf{p}_v, y). \quad (1)$$

The reasoning behind the condition is simple, if $x \geq \mathbf{p}_t \mathbf{q}_h(\mathbf{p}_v, y)$, then the consumed bundle $\mathbf{q}_h(\mathbf{p}_v, y)$ at the budget (\mathbf{p}_v, y) was also feasible when $\mathbf{q}_h(\mathbf{p}_t, x)$ was chosen. Given that the household h is utility maximizing and that the budget sets are distinct, it follows that $u_h(\mathbf{q}_h(\mathbf{p}_t, x)) > u_h(\mathbf{q}_h(\mathbf{p}_v, y))$, or equivalently, $v_h(\mathbf{p}_t, x) > v_h(\mathbf{p}_v, y)$. It is easy to see that condition (1) implies the Weak Axiom of Revealed Preference (Samuelson, 1938), which states that for any two distinct budgets (\mathbf{p}_t, x) and (\mathbf{p}_v, y) ,

$$\text{If } x \geq \mathbf{p}_t \mathbf{q}_h(\mathbf{p}_v, y) \text{ then } y < \mathbf{p}_v \mathbf{q}_h(\mathbf{p}_t, x).$$

3 Nonparametric bounds

In this section we show how to use basic revealed preference restrictions, in particular condition 1, together with information on the distribution of $\mathbf{q}_h(\mathbf{p}_t, x)$ in order to establish bounds on the distribution of the money metric utility function and the mean demand functions. As a first partial result, we demonstrate the possibility to obtain bounds on the proportion of households in the economy that prefer a certain budget over another.

Observational assumptions We depart from the observational restrictions imposed by a repeated cross sectional household consumption dataset, where different households face the same prices in each cross section. This gives us a data structure with a limited set of price regimes, and for each price regime a large number of consumption bundles which are obtained from a random sample of households in the economy. We denote by $T = \{1, \dots, |T|\}$, the set of cross sections. The price vector corresponding to cross section $t \in T$ is denoted by \mathbf{p}_t .

Given that different households face distinct expenditure levels, it is possible to obtain the distribution of the random consumption bundles $\mathbf{q}_h(\mathbf{p}_t, x)$ for every cross sectional price vector \mathbf{p}_t , ($t \in T$) and for any level of expenditure x .⁹ We assume that $\mathbf{q}_h(\mathbf{p}_t, x)$ has a continuous density function which is strictly positive on its domain. We use the notation $\Pr[A]$ as a shorthand for the following probability,

$$\Pr[A] = \int \mathbb{1}[h \in A] dP(h),$$

where $\mathbb{1}[\cdot]$ is the binary indicator function which equals one if and only if the term between brackets is true. $\Pr[A]$ is the fraction of the households for which the statement A holds. Equivalently, it gives us the probability that A holds for a household h drawn at random from the population. We require sufficient variation of preferences and demand such that for any two distinct budgets

⁹Estimation will be discussed in section 3.

(\mathbf{p}_t, x) and (\mathbf{p}_v, y) ,

$$\Pr [x = \mathbf{p}_t \mathbf{q}_h(\mathbf{p}_v, y)] = 0, \text{ and}$$

$$\Pr [v_h(\mathbf{p}_t, x) = v_h(\mathbf{p}_v, y)] = 0$$

This will allow us to freely interchange strict and weak inequalities within the function $\Pr[\cdot]$.

3.1 Bounds on population preferences

Consider two budgets (\mathbf{p}_t, x) and (\mathbf{p}_v, y) . Given that all households are rational, we know from condition 1 that for all households h ,

$$\text{If } x \geq \mathbf{p}_t \mathbf{q}_h(\mathbf{p}_v, y) \text{ then, } v_h(\mathbf{p}_t, x) > v_h(\mathbf{p}_v, y).$$

This means that $\mathbb{1}[x > \mathbf{p}_t \mathbf{q}_h(\mathbf{p}_v, y)] \leq \mathbb{1}[v_h(\mathbf{p}_t, x) > v_h(\mathbf{p}_v, y)]$. Integrating both sides, and using the independence assumption, we obtain,

$$\begin{aligned} \Pr [x \geq \mathbf{p}_t \mathbf{q}_h(\mathbf{p}_v, y)] &\leq \Pr [v_h(\mathbf{p}_t, x) \geq v_h(\mathbf{p}_v, y)], \\ \iff r_{t,v}(x, y) &\leq \Pr [v_h(\mathbf{p}_t, x) \geq v_h(\mathbf{p}_v, y)]. \end{aligned} \quad (2)$$

This inequality provides a lower bound on the fraction of households that prefer the budget (\mathbf{p}_t, x) over the budget (\mathbf{p}_v, y) . Given inequality (2), and the fact that,

$$\Pr [v_h(\mathbf{p}_t, x) \geq v_h(\mathbf{p}_v, y)] + \Pr [v_h(\mathbf{p}_v, y) \geq v_h(\mathbf{p}_t, x)] = 1,$$

We immediately obtain the upper bound,

$$\Pr [v_h(\mathbf{p}_t, x) \geq v_h(\mathbf{p}_v, y)] \leq 1 - r_{v,t}(y, x).$$

For both lower and upper bounds to be valid, it should be the case that for all cross sections $t, v \in T$ and all expenditure levels x, y ,

$$r_{t,v}(x, y) + r_{v,t}(y, x) \leq 1.$$

This condition is equivalent to the Weak Axiom of Stochastic Revealed Preference applied to our setting (see Bandyopadhyay, Dasgupta, and Pattanaik (1999, 2002, 2004) and Matzkin (2007) for a similar inequality). Hoderlein and Stoye (2014a) and Kawaguchi (2012) recently developed (among other things) a statistical test that verifies whether this condition is satisfied. There are two potential issues that may arise. First of all, it may happen that $r_{t,v}(x, y) + r_{v,t}(y, x)$ is larger than one, in which case the bounds cannot be simultaneously satisfied. Alternatively, it may happen that $r_{t,v}(x, y) + r_{v,t}(y, x)$ is considerably smaller than one, in which case the range may be too large to contain much useful information. Whether one of those problems arises is obviously an empirical matter. Appendix A presents potential solutions to deal with both of these issues.

3.2 Bounds on money metric utility

Let us now show how to use inequality (2) to bound the distribution of the money metric utility function $\mu_h(\mathbf{p}_t, \mathbf{p}_0, x_0)$ for some price vectors \mathbf{p}_0 and \mathbf{p}_t corresponding to the prices of two cross sections in the data set and for a particular level of income x_0 . Let us first focus on the upper bounds.

Upper bounds Fix a cross sectional price vector \mathbf{p}_0 and an income level x_0 . For any number $\pi \in (0, 1)$ and any cross section $t \in T$, let $h_t(\pi)$ be the π -th quantile of the variable $\mathbf{p}_t \mathbf{q}_h(\mathbf{p}_0, x_0)$,

$$\begin{aligned} \pi &= \Pr [h_t(\pi) \geq \mathbf{p}_t \mathbf{q}_h(\mathbf{p}_0, x_0)], \\ &= r_{t,0}(h_t(\pi), x_0) \end{aligned}$$

From inequality (2), we know that π is lower than the fraction of the households that prefer the budget $(\mathbf{p}_t, h_t(\pi))$ over the budget (\mathbf{p}_0, x_0) .

$$\begin{aligned}\pi &\leq \Pr [v_h(\mathbf{p}_t, h_t(\pi)) \geq v_h(\mathbf{p}_0, x_0)], \\ &= \Pr [h_t(\pi) \geq \mu_h(\mathbf{p}_t, \mathbf{p}_0, x_0)].\end{aligned}$$

The second line is obtained by inverting the indirect utility function $v_h(\mathbf{p}_t, h_t(\pi))$ with respect to its second argument. This can be done by the fact that the indirect utility function is strictly increasing in income.

Let us denote by $m_t(\pi)$ the quantile function of $\mu_h(\mathbf{p}_t, \mathbf{p}_0)$, i.e. for all $\pi \in (0, 1)$

$$\Pr [\mu_h(\mathbf{p}_t, \mathbf{p}_0, x_0) \leq m_t(\pi)] = \pi.$$

Then, using our previously established result, we have that,

$$\begin{aligned}\Pr [\mu_h(\mathbf{p}_t, \mathbf{p}_0, x_0) \leq m_t(\pi)] &= \pi \leq \Pr [\mu_h(\mathbf{p}_t, \mathbf{p}_0, x_0) \leq h_t(\pi)], \\ \iff m_t(\pi) &\leq h_t(\pi).\end{aligned}$$

The last line uses the assumption that the cumulative distribution function of $\mu_h(\mathbf{p}_t, \mathbf{p}_0, x)$ is strictly increasing on its support. This result shows that $h_t(\pi)$ is an upper bound on the π th quantile of the distribution of the money metric utility function. Using these upper bounds on the quantiles; we can also derive an upper bound on the mean value of the money metric utility. Let M be the mean of the function $\mu_h(\mathbf{p}_t, \mathbf{p}_0, x_0)$. We have that:

$$\begin{aligned}M &= \int_0^\infty \mu_h(\mathbf{p}_t, \mathbf{p}_0, x_0) dF(\mu_h(\mathbf{p}_t, \mathbf{p}_0, x_0)), \\ &= \int_0^1 m_t(\pi) d\pi \leq \int_0^1 h_t(\pi) d\pi.\end{aligned}$$

In practice, we compute the values of $h_t(\pi)$ for a finite grid of values $\pi_0, \pi_1, \dots, \pi_n$ with $\pi_0 = 0$ and $\pi_n = 1$.¹⁰ This allows us to approximate this upper bound by,

$$\int_0^1 h_t(\pi) d\pi \leq \sum_{n=1}^N (\pi_n - \pi_{n-1}) h_t(\pi_n).$$

The finer the grid, the better the approximation.

Lower bounds We use a similar procedure to compute lower bounds on the quantiles. For $\pi \in (0, 1)$, let $\ell_t(\pi)$ solve,

$$\begin{aligned} 1 - \pi &= \Pr [x_0 \geq \mathbf{p}_0 \mathbf{q}_h(\mathbf{p}_t, \ell_t(\pi))], \\ &= r_{0,t}(x_0, \ell_t(\pi)) \end{aligned}$$

Then,

$$\begin{aligned} 1 - \pi &\leq \Pr [v_h(\mathbf{p}_0, x_0) \geq v_h(\mathbf{p}_t, \ell_t(\pi))], \\ &= \Pr [\mu_h(\mathbf{p}_t, \mathbf{p}_0, x_0) \geq \ell_t(\pi)], \\ &= 1 - \Pr [\mu_h(\mathbf{p}_t, \mathbf{p}_0, x_0) \leq \ell_t(\pi)] \end{aligned}$$

As before, let $m_t(\pi)$ be the π th quantile of the distribution of the money metric utility $\mu_h(\mathbf{p}_t, \mathbf{p}_0, x_0)$.

We have that,

$$\begin{aligned} \Pr [\mu_h(\mathbf{p}_t, \mathbf{p}_0, x_0) \leq m_t(\pi)] &= \pi \geq \Pr [\mu_h(\mathbf{p}_t, \mathbf{p}_0, x_0) \leq \ell_t(\pi)], \\ \iff m_t(\pi) &\geq \ell_t(\pi) \end{aligned}$$

¹⁰The upper bound $h_t(1)$ can be set to the minimal income such that the budget hyperplane for $(\mathbf{p}_t, h_t(1))$ lies above the hyperplane for (\mathbf{p}_0, x_0) .

This shows that $\ell_t(\pi)$ is a lower bound for the quantile $m_t(\pi)$. The mean M is then bounded from below by the quantity $\int_0^1 \ell_t(\pi) d\pi$ which can be approximated by $\sum_{n=0}^{N-1} \ell_t(\pi_n)(\pi_{n+1} - \pi_n)$.¹¹

3.3 Bounds on demand functions

Now, we demonstrate how to adapt above framework in order to establish bounds on the quantiles of the demand functions for unobserved, counterfactual, price regimes \mathbf{p}_0 and expenditure levels x_0 , i.e. \mathbf{p}_0 does not necessarily correspond to a price vector of a certain cross section.

Consider a function $f : \mathbb{R}_+^n \rightarrow \mathbb{R} : \mathbf{q} \mapsto f(\mathbf{q})$. In this section, we will provide upper bounds on the quantiles of the distribution of the random variable $f(\mathbf{q}_h(\mathbf{p}_0, x_0))$. The function $f(\cdot)$ encompasses various interesting measures. For example, if we want to bound the expenditure share on one of the goods, we can use the function $f(\mathbf{q}) = \frac{1}{x_0} p_{0,j} q_j$, where $p_{0,j}$ is the price of good j in vector \mathbf{p}_0 , q_j is the quantity of good j in vector \mathbf{q} and x_0 is the expenditure level.

The focus on upper bounds is not restrictive given that we can always use information on upper bounds to construct lower bounds. In order to see this, let $-m(1 - \pi)$ be the $(1 - \pi)$ th quantile of the variable $-f(\mathbf{q}_h(\mathbf{p}_0, x_0))$ and let $-g(1 - \pi)$ be its upper bound. We then have that,

$$\begin{aligned}
1 - \pi &= \int \mathbb{1}[-f(\mathbf{q}_h(\mathbf{p}_0, x_0)) \leq -m(1 - \pi)] dP(h), \\
&\leq \int \mathbb{1}[-f(\mathbf{q}_h(\mathbf{p}_0, x_0)) \leq -g(1 - \pi)] dP(h), \\
\iff \pi &\geq 1 - \int \mathbb{1}[-f(\mathbf{q}_h(\mathbf{p}_0, x_0)) \leq -g(1 - \pi)] dP(h), \\
&= \int \mathbb{1}[-f(\mathbf{q}_h(\mathbf{p}_0, x_0)) > -g(1 - \pi)] dP(h), \\
&= \int \mathbb{1}[f(\mathbf{q}_h(\mathbf{p}_0, x_0)) \leq g(1 - \pi)] dP(h).
\end{aligned}$$

As such, we see that $g(1 - \pi)$ provides a lower bound on the π th quantile of $f(\mathbf{q}_h(\mathbf{p}_0, x_0))$. As an

¹¹The lower bound $\ell_t(0)$ can be set to the maximal income such that the hyperplane for the budget set $(\mathbf{p}_t, \ell_t(0))$ lies below the hyperplane for (\mathbf{p}_0, x_0) .

example, we can establish a lower bound on the π th quantile of $f(\mathbf{q}) = \frac{1}{x_0} p_{0,j} q_j$ by constructing an upper bound on the $(1 - \pi)$ th quantile of $-\frac{1}{x_0} p_{0,j} q_j (= \sum_{i \neq j} \frac{1}{x_0} p_{0,i} q_i - 1)$.

For every cross section t , we previously defined the value $\ell_t(1 - \pi)$ that satisfied the following condition,

$$\begin{aligned} \pi &= \Pr [x_0 \geq \mathbf{p}_0 \mathbf{q}_h(\mathbf{p}_t, \ell_t(1 - \pi))], \\ &= r_{0,t}(x_0, \ell_t(1 - \pi)). \end{aligned}$$

The value of $\ell_t(1 - \pi)$ can be obtained using information on x_0 , \mathbf{p}_0 and the distribution of $\mathbf{q}_h(\mathbf{p}_t, x)$ alone, which we assumed to be known. For the next step, we use the Weak Axiom of Stochastic Revealed Preference, which requires that,

$$\begin{aligned} r_{t,0}(\ell_t(1 - \pi), x_0) + r_{0,t}(x_0, \ell_t(1 - \pi)) &\leq 1, \\ \iff r_{0,t}(x_0, \ell_t(1 - \pi)) &\leq 1 - r_{t,0}(\ell_t(1 - \pi), x_0). \end{aligned}$$

Let $m(\pi)$ be the π th quantile of the distribution function of the random variable $f(\mathbf{q}_h(\mathbf{p}_0, x_0))$.

We have that,

$$\begin{aligned} \Pr [f(\mathbf{q}_h(\mathbf{p}_0, x_0)) \leq m(\pi)] &= \pi = r_{0,t}(x_0, \ell_t(1 - \pi)), \\ &\leq 1 - r_{t,0}(\ell_t(1 - \pi), x_0), \\ &= \Pr [\ell_t(1 - \pi) \leq \mathbf{p}_t \mathbf{q}_h(\mathbf{p}_0, x_0)] \\ &\leq \Pr \left[f(\mathbf{q}_h(\mathbf{p}_0, x_0)) \leq \max_{\mathbf{q}} f(\mathbf{q}) \text{ s.t } \ell_t(1 - \pi) \leq \mathbf{p}_t \mathbf{q} \text{ and } \mathbf{p}_0 \mathbf{q} = x_0 \right] \end{aligned}$$

The last inequality follows from the fact that whenever $\ell_t(1 - \pi) \leq \mathbf{p}_t \mathbf{q}_h(\mathbf{p}_0, x_0)$ holds, then $f(\mathbf{q}_h(\mathbf{p}_0, x_0)) \leq \max_{\mathbf{q}} f(\mathbf{q}) \text{ s.t } \ell_t(1 - \pi) \leq \mathbf{p}_t \mathbf{q} \text{ and } \mathbf{p}_0 \mathbf{q} = x_0$ must also hold. In order to see this, assume on the contrary that $f(\mathbf{q}_h(\mathbf{p}_0, x_0))$ is larger than $f(\mathbf{q})$ for all vectors \mathbf{q} where

$\mathbf{p}_0\mathbf{q} = x_0$ and $\ell_t(1 - \pi) \leq \mathbf{p}_t\mathbf{q}$. Then, given that $\mathbf{p}_0\mathbf{q}_h(\mathbf{p}_0, x_0) = x_0$, it must be that $\ell_t(1 - \pi) > \mathbf{p}_t\mathbf{q}_h(\mathbf{p}_0, x_0)$, a contradiction.

Above result shows that,


$$m(\pi) \leq \max_{\mathbf{q}} f(\mathbf{q}) \text{ s.t } \ell_t(1 - \pi) \leq \mathbf{p}_t\mathbf{q} \text{ and } \mathbf{p}_0\mathbf{q} = x_0,$$

for all cross sections t . In practice, we compute this right hand side for every cross section t and then take the lowest value across all cross sections as the upper bound. If f is a linear function, then the right hand side is a simple linear programming problem which can be solved efficiently.

The construction of the bounds in the simple two goods setting is illustrated in Figure 1. There are three budget lines corresponding to (\mathbf{p}_0, x_0) , $(\mathbf{p}_t, \ell_t(1 - \pi))$ and $(\mathbf{p}_v, \ell_v(\pi))$. The incomes $\ell_t(1 - \pi)$ and $\ell_v(\pi)$ are chosen such that the mass of households on the dashed line segment (where $x_0 > \mathbf{p}_0\mathbf{q}_h(\mathbf{p}_t, \ell_t(1 - \pi))$) is equal to π and the mass of households on the dotted line segment (where $x_0 > \mathbf{p}_0\mathbf{q}_h(\mathbf{p}_v, \ell_v(\pi))$) is equal to $(1 - \pi)$.

The quantity \bar{q}_2 is the maximum value of good 2 that corresponds to a bundle on the budget (\mathbf{p}_0, x_0) (where $\mathbf{p}_0\mathbf{q} = x_0$) and $\ell_t(1 - \pi) \leq \mathbf{p}_t\mathbf{q}$. From the result above, we know that this value gives an upper bound on the π th quantile of the distribution of $q_{2,h}(\mathbf{p}_0, x_0)$. Given that there are only two goods, this upper bound immediately gives a lower bound on the $(1 - \pi)$ th quantile of $q_{1,h}(\mathbf{p}_0, x_0)$, given by \underline{q}_1 . Similarly, \bar{q}_1 gives an upper bound on the $(1 - \pi)$ th quantile of $q_{1,h}(\mathbf{p}_0, x_0)$, while \underline{q}_2 gives a lower bound on the π th quantile of $q_{2,h}(\mathbf{p}_0, x_0)$. As such, the π th quantile of $q_{2,h}(\mathbf{p}_0, x_0)$ is bounded by the quantities \bar{q}_2 and \underline{q}_2 . Given these bounds on the quantiles of the demand functions, we can compute bounds on the mean of the demand function by using a similar procedure as for the money metric utility function.

Figure 1: Illustration of the construction of the bounds



4 Application

In this section, we discuss the empirical implementation of the theoretical bounds that were established in the previous section. We first present our estimation procedure for the measures $r_{t,v}(x, y)$, $l_t(\pi)$ and $h_t(\pi)$. Next, we discuss how we control for observed heterogeneity and endogeneity of the total expenditures. We also very briefly discuss the issue of statistical inference on bounds. Finally, we present some empirical results.

4.1 Estimation procedure

The construction of the bounds in the previous section assumed that we know the distribution of the variables $q_h(\mathbf{p}_t, x)$ for every cross sectional price \mathbf{p}_t and every income level x . Given these distributions it is fairly easy to obtain the quantities $r_{t,v}(x, y) = \Pr[x \geq \mathbf{p}_t \mathbf{q}_h(\mathbf{p}_v, y)]$, which form the main building blocks for our bounds. In practice, however, these probabilities need to be estimated. We propose a kernel estimator.

Consider the v th cross section, $v \in T$. Assume that this cross section contains a sample of n observed household demand bundles $\{\mathbf{q}_{v,i}\}_{i \leq n}$ where i corresponds to a particular observation. We denote by $\{z_{v,i}\}_{i \leq n}$ the log of the expenditure levels ($z_{v,i} = \ln(\mathbf{p}_v \mathbf{q}_{v,i})$). We assume that the sample $\{\mathbf{q}_{v,i}\}_{i \leq n}$ is i.i.d drawn from the random vector \mathbf{q}_v with distribution $F(\cdot)$. We denote by z_v the random variable $\ln(\mathbf{p}_v \mathbf{q}_v)$. Finally, let \mathbf{s}_v be the random vector of normalized consumption, $\mathbf{s}_v = \mathbf{q}_v/x_v$ and denote the realizations of \mathbf{s}_v by $\mathbf{s}_{v,i} = \mathbf{q}_{v,i}/x_{v,i}$

We have that,

$$\begin{aligned} r_{t,v}(x, y) &= \int \mathbb{1}[x \geq \mathbf{p}_t \mathbf{q}_h(\mathbf{p}_v, y)] dP(h), \\ &= \int \mathbb{1}[x \geq \mathbf{p}_t \mathbf{q}_v] dF(\mathbf{q}_v | z_v = \ln(y)), \\ &= \int \mathbb{1}[x \mathbf{p}_v \mathbf{s}_v \geq \mathbf{p}_t y \mathbf{s}_v] dF(\mathbf{s}_v | z_v = \ln(y)), \\ &= \int \mathbb{1}[(x \mathbf{p}_v - y \mathbf{p}_t) \mathbf{s}_v \geq 0] dF(\mathbf{s}_v | z_v = \ln(y)). \end{aligned}$$

Here we used the identity $\mathbf{p}_v \mathbf{s}_v = 1$ and the fact that, conditional on $z_v = \ln(y)$, $\mathbf{q}_v = y \mathbf{s}_v$. We can estimate this value using the Nadaraya-Watson estimator,

$$\frac{\frac{1}{nh} \sum_{i=1}^n \mathbb{1}[(x \mathbf{p}_v - y \mathbf{p}_t) \mathbf{s}_{v,i} \geq 0] k\left(\frac{z_{v,i} - \ln(y)}{h}\right)}{\frac{1}{nh} \sum_{i=1}^n k\left(\frac{z_{v,i} - \ln(y)}{h}\right)}.$$

where h is the bandwidth and $k(\cdot)$ is a symmetric kernel function that satisfies $\int k(v) dv = 1$ and $\int v k(v) dv = 0$.¹²

If for $n \rightarrow \infty$, $h \rightarrow 0$ and $nh \rightarrow \infty$, then the estimators $\hat{r}_{t,v}(x, y)$ consistently estimate $r_{t,v}(x, y)$. Under suitable conditions¹³ the estimator $\sqrt{nh} [\hat{r}_{t,v}(x, y) - r_{t,v}(x, y)]$ is asymptotically normally distributed (see Li and Racine (2007)).

¹²In our application, we use the Gaussian kernel.

¹³In particular, (i) $\mu(y) > 0$, (ii) $r_{t,v}(x, y) \in (0, 1)$, (iii) $\mu(y)$ and $r_{t,v}(x, y)$ have continuous second order derivatives with respect to y . and (iv) $nh^7 \rightarrow 0$

The estimators for $h_t(\pi)$ and $\ell_t(\pi)$ are computed as the solution to the following equations,

$$\begin{aligned}\pi &= \hat{r}_{t,0}(\hat{h}_t(\pi), x_0), \\ 1 - \pi &= \hat{r}_{0,t}(x_0, \hat{\ell}_t(\pi)),\end{aligned}$$

using standard binary search algorithms. In order for this algorithm to work, we assume that $\hat{r}_{0,t}(x_0, \hat{\ell}_t(\pi))$ is decreasing in $\hat{\ell}_t(\pi)$. This assumption is (asymptotically) valid if all goods are normal (i.e. all demand functions are increasing in income).¹⁴

The estimators $\hat{h}_t(\pi)$ are equivalent to a conditional quantile kernel estimator. These estimators are consistent and asymptotically normal for $h_t(\pi)$ as long as for $n \rightarrow \infty$, $h \rightarrow 0$ and $nh \rightarrow \infty$ (Li and Racine, 2007, section 6.3). Using a proof similar to the one of Cai (2002, Theorem 2), we show in the appendix that for $n \rightarrow \infty$, $h \rightarrow 0$ and $nh \rightarrow \infty$, the estimators $\hat{\ell}_t(\pi)$ are consistent for $\ell_t(\pi)$ and that, given some additional conditions,¹⁵ they are asymptotically normally distributed. As usual with kernel estimators, each of these estimators will have an asymptotic bias which does not disappear asymptotically when using the optimal bandwidth. One possible solution is to undersmooth.

The estimators for the bounds on the demand functions are computed by substituting the estimated values $\hat{\ell}_t(\pi)$ for the values of $\ell_t(\pi)$ in the linear programming problems. The resulting estimators are determined as the minimum over a finite set of values which are themselves the solution of a linear maximization problem that contains the estimates $\hat{\ell}_t(\pi)$ as a parameter. From the continuous mapping theorem, it follows that these are also consistent.

Observable heterogeneity and endogeneity We adjust the kernel estimators $\hat{r}_{t,v}(x, y)$ by including a semi-parametric specification. We have two reasons to do this. First of all, given the data

¹⁴See also Blundell, Browning, and Crawford (2003) for a similar assumption.

¹⁵In particular, we require that $r_{t,v}(x_0, \ell_t(\pi))$ has strict negative partial derivative with respect to $\ell_t(\pi)$ and that $\sqrt{nh^3} \rightarrow \infty$.

limitations, we would like to allow our estimators to depend on the vector of observed covariates, \mathbf{a}^h , without fully conditioning on each of its values. Next, we need to take into account the fact that total expenditures are probably endogenous. We follow Blundell, Browning, and Crawford (2008), and consider the following semiparametric modification,

$$\mathbb{E} [\mathbb{1}[(x\mathbf{p}_t - y\mathbf{p}_v) \mathbf{s}_{v,i} \geq 0]] = g(z_{v,i} - \phi(\mathbf{a}'_{v,i}\theta)) + \mathbf{a}'_{v,i}\gamma + \varepsilon_{v,i},$$

where $\mathbf{a}_{v,i}$ be the observed household composition in cross section v for household i . The function $\phi(\mathbf{a}'_{v,i}\theta)$ can be interpreted as the log of a general equivalence scale for the household, and $\mathbf{a}'_{v,i}\gamma$ documents the way in which observable demographic differences across households impact on the left hand side. Similar to Blundell et al. (2008) we use an estimate of the general equivalence scale $\phi(\mathbf{a}'_{v,i}\theta)$ taken from the Organisation for Economic Co-operation and Development (OECD) scales.

In order to control for the endogeneity of z_v , Blundell, Browning, and Crawford (2008) suggest to use a control function approach based on the two step semiparametric estimator (this estimator is based on the procedure set out by Newey, Powell, and Vella (1999)). In a first step, we obtain the residuals from a regression of the log of total expenditure on all exogenous variables in the model and on an excluded instrument. We take the log of (equivalent) labor income as an instrument. In the second step, we conduct a semiparametric regression of $\mathbb{1}[(x\mathbf{p}_t - y\mathbf{p}_v) \mathbf{s}_{v,i} \geq 0]$ on $g(z_{v,i} - \phi(\mathbf{a}'_{v,i}\theta))$, $\mathbf{a}'_{v,i}\gamma$ and $\hat{\delta}_{v,i}$, where $\hat{\delta}_{v,i}$ are the residuals from the first stage regression.

Inference on bounds The methodology outlined in section 3 provides nonparametric bounds on various parameters of interest (e.g. the quantiles of the money metric utility). Given that the bounds are based on finite sample estimates, we are confronted with the issue of statistical inference, in particular, the construction of confidence intervals. Given that our estimates only provide bounds, this problem fits in the literature that deals with the construction of confidence intervals for partially identified estimators. We refer to the several recent papers by Imbens and Manski

(2004); Chernozhukov, Hong, and Tamer (2007); Stoye (2009); Bugni (2010); Chernozhukov, Lee, and Rosen (2013) and in particular to the recent paper of Hoderlein and Stoye (2014a) who consider the problem of constructing confidence intervals in a setting which is similar to ours.

In principal, we need to deal with two issues. First, the construction of CI for the estimates of the bounds and next, the construction of the CI for the identified set itself. For the latter we choose Bonferroni type intervals which provide conservative inference. For the former, we notice that the estimates of our bounds are obtained as the maximum or minimum of a number of estimators that are computed using samples from different cross sections. It can be shown that in such cases, the usual bootstrap procedure is not valid (Andrews, 2000). In order to obtain asymptotic valid inference we use the subsampling procedure which is discussed in detail by Politis, Romano, and Wolf (1999). Subsampling is similar to bootstrap but the samples taken are smaller and draws are obtained without replacement. The subsampling procedure is valid under very weak assumptions, in particular for all of our estimators (see also Appendix D for more information on the implementation of our subsampling).

4.2 Data Description

We illustrate our approach by using a data sample from the Consumer Expenditure Survey (CEX), a repeated cross section. We use data on consumption decisions by US households from 1994 to 2007 (14 years). It is important to note that the consumer expenditures are derived from the diary survey (and not from the interview data). The diary data seem well-suited for (static) demand analysis. First of all, given that we focus on non-durable consumption, which is customary in static demand analysis, information on the purchase of big, durable items is unnecessary. Second, for non-durable commodities, the diary survey invites respondents to indicate their consumption in a two-week period. Because this period is relatively short, respondents should be able to recall their expenditures. We follow Blundell et al. (2008) by focusing our attention to three broad ex-


penditure categories, namely, food, other non-durables and services.¹⁶ As the diary survey reports expenditures on a two-week basis, we convert these to yearly equivalents. Converting two-week expenditures to yearly data poses an important problem of seasonality. Therefore, we deseasonalize using a dummy regression approach. Specifically, the expenditures on each category (reported for two weeks) are regressed on month dummies. Residuals from this regression (which can be interpreted as the variation in expenditures which can not be explained by seasonality or by months) are added to the mean expenditures for each category in order to construct deseasonalized expenditures. Observations with negative total expenditures are dropped. As mentioned above, we also take into account that variation in expenditures can be driven by the household composition, e.g. the number of adults or the number of kids living in the family. Therefore, we deflate total expenditures as well as total income by an OECD equivalence scale.

For the empirical analysis, we restrict attention to (i) households who have completed the two-week diary, (ii) households who are not living in student housing, (iii) households who are vehicle owners (to include fuel expenses), (iv) households where both members work at least 17 hours, (v) households in which both members are not self-employed, (vi) households in which the age of the reference person is at least 21 and finally we restrict attention to (vii) households that consist of a husband, a wife and possibly children. As a final step we also remove some outlier observations.¹⁷ On average, we are left with 2163 observations per cross-section with a minimum of 1775 observations in 1994 and a maximum of 2379 observations in 2007. The left pane of Figure 2 plots the evolution of the mean consumption shares of the three goods over the considered periods. Price information is obtained from the Bureau of Labor Statistics. The right pane of Figure 2 gives the evolution of these prices, normalized at the 1994 level.

¹⁶See Appendix C for a list of the different goods used for the construction of the aggregates.

¹⁷In particular, we removed observations for which rescaled total expenditures or expenditure shares are not within 3 standard deviations from the mean and observations for which rescaled total expenditures are among the 5 per cent lowest or 5 per cent highest expenditures or for which the expenditure shares are close to 0.

Figure 2: Evolution of average consumption shares and prices



4.3 Empirical results

In this section, we provide the results of several exercises. Due to limited space, we need to restrict our analysis to some particular base years and some reference income levels. Additional results are available from the authors upon request.

Bounds on the mean cost of living Let us first show how our bounds perform with respect to the computation of the mean of the Laspeyres-Konüs cost of living index,

$$\int \frac{\mu_h(\mathbf{p}_t, \mathbf{p}_0, x_0)}{x_0} dP(h) = \frac{1}{x_0} \int \mu_h(\mathbf{p}_t, \mathbf{p}_0, x_0) dP(h).$$

The Laspeyres-Konüs price index measures the income that one would need, relative to the income in period 0, in order to be equally well off as in the initial period. We take 1994 as the reference year which means that \mathbf{p}_0 corresponds to the price vector in the year 1994. We choose x_0 as the (OECD equivalence scale deflated) median expenditure level in 1994. The bounds on the cost of living that we obtain using our procedure are given in the last column of Table 1. The table

also reports values for various other prices indices like the Laspeyres (L), the Paasche (P) and the Tornqvist price index (T).¹⁸ We also provide information on three other nonparametric bounds. The first are the Lerner bounds which are obtained from the fact that:

$$\min_j \left\{ \frac{p_{t,j}}{p_{0,j}} \right\} \leq \frac{\mu(\mathbf{P}_t, \mathbf{P}_0, x_0)}{x_0} \leq \max_j \left\{ \frac{p_{t,j}}{p_{0,j}} \right\}.$$

The bounds by (Pollak, 1971) improve upon this by replacing the upper bound by the Laspeyres price index.

$$\min_j \left\{ \frac{p_{t,j}}{p_{0,j}} \right\} \leq \frac{\mu(\mathbf{P}_t, \mathbf{P}_0, x_0)}{x_0} \leq \frac{\mathbf{P}_t \mathbf{Q}_0}{x_0}.$$

The second to last column gives the bounds that are obtained by using the procedure set out by Blundell, Browning, and Crawford (2003). This method first estimates nonparametric Engel curves and subsequently uses these estimates in combination with revealed preference restrictions to establish nonparametric bounds. We would like to emphasize that there is a clear conceptual difference between the bounds of Blundell, Browning, and Crawford (2003) (and Pollak), and ours. Their procedure provides bounds on the cost of living that correspond to some kind of ‘representative individual’ whose demand functions equal the mean demand functions over the population. Our bounds, on the other hand, correspond to bounds on the mean cost of living over all households within the population. Although in our case, both procedures give similar results, this does not have to be the case in general. One partial explanation for this fact might be that the distribution of the cost of living over the households is quite symmetric (see below). Moreover, we show in our next exercise that the cost of living for particular households (e.g. households at the 10th or 90th percentile) may considerably diverge from the mean cost of living.

Also notice that the bounds for both the BBC approach and our approach are quite narrow. This shows that, although we relax the representative agent assumption, we are still able to get fairly reliable results.


¹⁸These are computed on the basis of nonparametric Engel curve estimates.

Table 1: Bounds on the mean Laspeyres Konüs cost of living index

year	Price indices			Nonparametric Bounds			bounds
	L	P	T	Lerner	Pollak	BBC	
1994	1.0000	1.0000	1.0000	[1.0000, 1.0000]	[1.0000, 1.0000]	[1.0000, 1.0000]	[1.0000 1.0000]
1995	1.0275	1.0271	1.0273	[1.0086, 1.0357]	[1.0086, 1.0275]	[1.0250, 1.0275]	[1.0249 1.0292]
1996	1.0604	1.0596	1.0600	[1.0358, 1.0708]	[1.0358, 1.0604]	[1.0591, 1.0604]	[1.0574 1.0621]
1997	1.0860	1.0844	1.0852	[1.0483, 1.1019]	[1.0483, 1.0860]	[1.0830, 1.0860]	[1.0819 1.0875]
1998	1.0972	1.0929	1.0951	[1.0327, 1.1236]	[1.0327, 1.0972]	[1.0932, 1.0972]	[1.0900 1.0983]
1999	1.1242	1.1212	1.1227	[1.0709, 1.1470]	[1.0709, 1.1242]	[1.1205, 1.1242]	[1.1180 1.1256]
2000	1.1716	1.1712	1.1714	[1.1480, 1.1886]	[1.1480, 1.1716]	[1.1689, 1.1716]	[1.1692 1.1739]
2001	1.2066	1.2048	1.2057	[1.1456, 1.2437]	[1.1456, 1.2066]	[1.2025, 1.2066]	[1.2025 1.2086]
2002	1.2206	1.2154	1.2181	[1.1301, 1.2742]	[1.1301, 1.2206]	[1.2143, 1.2201]	[1.2122 1.2222]
2003	1.2618	1.2562	1.2591	[1.1659, 1.3263]	[1.1659, 1.2618]	[1.2556, 1.2607]	[1.2534 1.2636]
2004	1.3094	1.3066	1.3080	[1.2243, 1.3679]	[1.2243, 1.3094]	[1.3048, 1.3089]	[1.3042 1.3115]
2005	1.3666	1.3698	1.3682	[1.3115, 1.4247]	[1.3115, 1.3666]	[1.3648, 1.3665]	[1.3658 1.3706]
2006	1.4181	1.4206	1.4194	[1.3400, 1.4839]	[1.3400, 1.4184]	[1.4157, 1.4178]	[1.4164 1.4202]
2007	1.4655	1.4679	1.4667	[1.3966, 1.5276]	[1.3966, 1.4655]	[1.4633, 1.4655]	[1.4644 1.4691]

Distribution of the cost-of-living Let us now have a look at the bounds on the quantiles of this cost of living index over the population, a feature which is only identifiable using our results. Figure 3 provides upper and lower bounds on the quantiles of the Laspeyres-Konüs cost of living index, for the 10th (dashed), 50th (solid) and 90th (dotted) percentile. Again the base year is 1994 and the reference income is given by the median expenditure level in this year. From Table 1 we already saw that the bounds on the mean price index were quite narrow. This narrowness is also found for the bounds on the quantiles as can be seen from Figure 3. The width of the distribution for a particular year depends to a large extent on the difference in relative slopes between the base year (\mathbf{p}_0/x_0) and the evaluation year (\mathbf{p}_t/x_t). The closer the relative prices, the narrower the difference between the largest and smallest cost of living for the particular year. The reason is that the distribution is naturally bounded between the minimum and maximum values of y/x_0 for which the budget hyperplanes corresponding to (\mathbf{p}_t, y) and (\mathbf{p}_0, x_0) do not intersect. We see that the distribution is narrow in the year 2000 and the widest in the year 2002 giving differences in cost of living up to more than 5 percentage points between the 10th and 90th percentile. One

Figure 3: Distribution of the cost of living




noticeable feature about the figure is that there seems to be a considerable amount of heterogeneity in the population although the width of the distribution remains more or less constant for the latter 5 years. Confidence intervals (using the subsampling procedure) and estimates for other quantiles are given in Appendix D.

Figure 4 gives another illustration of the kind of questions that can be answered given the framework in this paper. The figure gives bounds on the average of the Laspeyres-Konüs cost of living for different starting quantiles,

$$\int \frac{\mu_h(\mathbf{p}_t, \mathbf{p}_0, x_{0,i})}{x_{0,i}} dP(h).$$

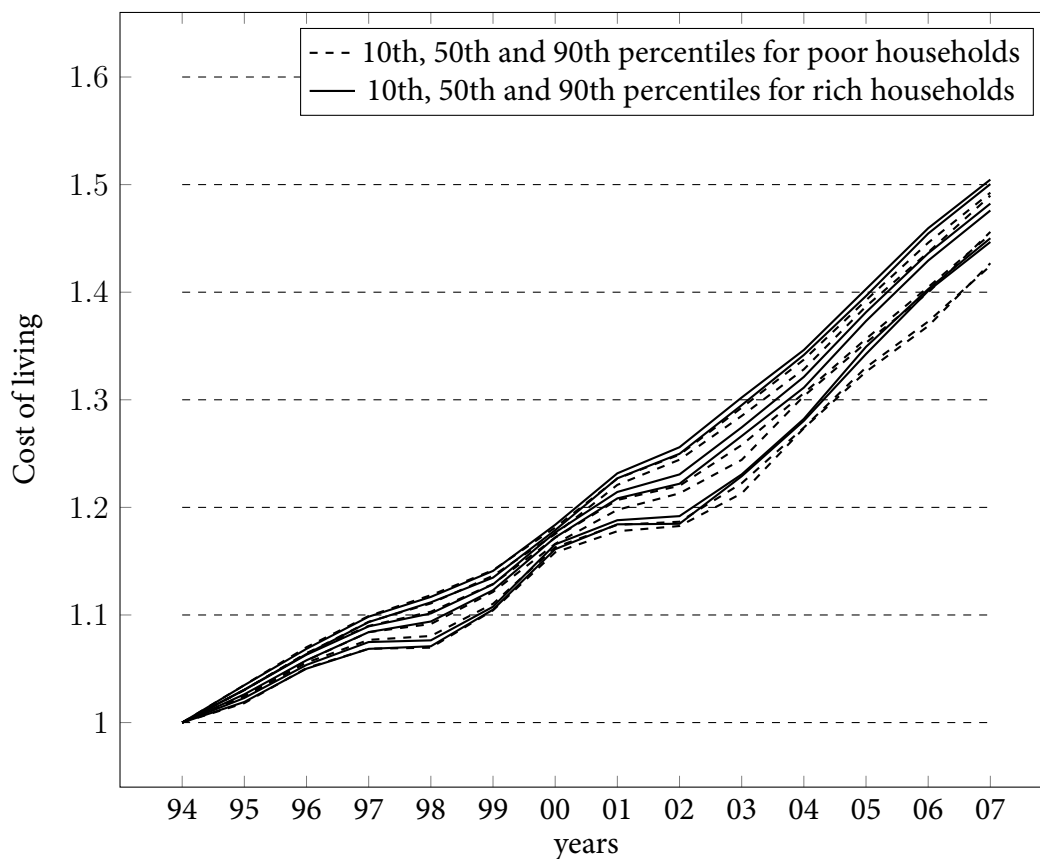
Figure 4: Change in mean cost of living 1994-2007 for different starting quantiles of income



Here, $x_{0,i}$ represents the income at the i th quantile of the income distribution in 1994, \mathbf{p}_0 are the prices in 1994 and \mathbf{p}_t is the price vector for 2007. The figure gives an idea of the average price increase (over the heterogeneous population) for households starting at different quantiles of the income distribution in 1994. On average one sees an increase over the quantiles, which means that (on average) the cost of living for households starting at the lower end of the income distribution in 1994 was lower than for household starting at the higher end of the income distribution. In other words, the households that started at the lower end of the income distribution had (on average) a lower increase in the cost of living. Also, notice that the upper bound for the lowest quantile is below the lower bound for the upper quantile. This shows that the average cost of living values are significantly different (although the numbers are very close to each other in absolute terms).

The bounds in Figure 4 restrict the average of the Laspeyres-Konüs cost of living. We can also

Figure 5: Distribution of the cost of living (poor households: 10th percentile of the income distribution in 1994; rich households: 90th percentile of the income distribution in 1994)



bound the distribution of this cost of living index conditional on the income of households. Figure 5 gives bounds on the distribution of the Laspeyres-Konüs cost of living for poor households, corresponding to the 10th percentile of the income distribution in 1994, and rich households, corresponding to the 90th percentile of the income distribution in 1994.

From 1994 to 2002, we find that the distribution of the cost of living is similar for households starting at the lower and upper ends of the income distribution. However, the results in Figure 5 also indicate that the cost of living increased more sharply for richer households after 2002. This confirms our results in Figure 4. Interestingly, the difference in (increase in) cost of living between poor and rich households is somewhat less outspoken at the upper end (i.e. the 90th percentile)

than at the lower end of the cost of living distribution (especially by 2007).


Distribution of the compensating variation Figure 6 shows the distribution of the compensating variation,

$$x_t - \mu_h(\mathbf{p}_t, \mathbf{p}_0, x_0)$$

Here, x_0 is taken to be the median income in 2000 and x_t is the median income in cross section t . This compensating variation gives the difference between the median income in year t and the minimum income that would be necessary in order to obtain the welfare level at budget (\mathbf{p}_0, x_0) . Values above zero indicate a welfare gain for a household at the median income in year t compared to a household at the median income in year 2000. We see that all quantiles are below zero for the years 1994-1999 and 2001 and quantiles are above zero for the years 2005-2007. Once again, there seems to be quite a lot of heterogeneity present in the population. For many years, the range between the 10th and 90th percentile is around \$400 per year which is substantial. The large increase in 2005 is mainly due to a sharp increase in the median expenditure level in that year.

Distribution of demand As a last exercise, let us have a look at the bounds on the demanded consumption shares for counterfactual price regimes. To keep focus, we restrict ourselves to the computation of bounds for the own price effect for the food aggregate. We construct normalized prices by dividing all cross sectional prices by the median income in the corresponding year, and we take the mean of these normalized prices as a reference point. The reference income level x_0 is set at 1. We let the price for food range from 0.9 times its reference level to 1.05 times its reference level. The prices of all others goods are held constant. Figure 7 presents the results for three quantiles, the 10th (dashed), the median (solid) and the 90th (dotted). Again we see a lot of heterogeneity in the demand curves over the population although the price responses look very similar across the three quantiles. As is customary in revealed preference analysis, it is only

Figure 6: Distribution of compensating variation, baseyear 2000

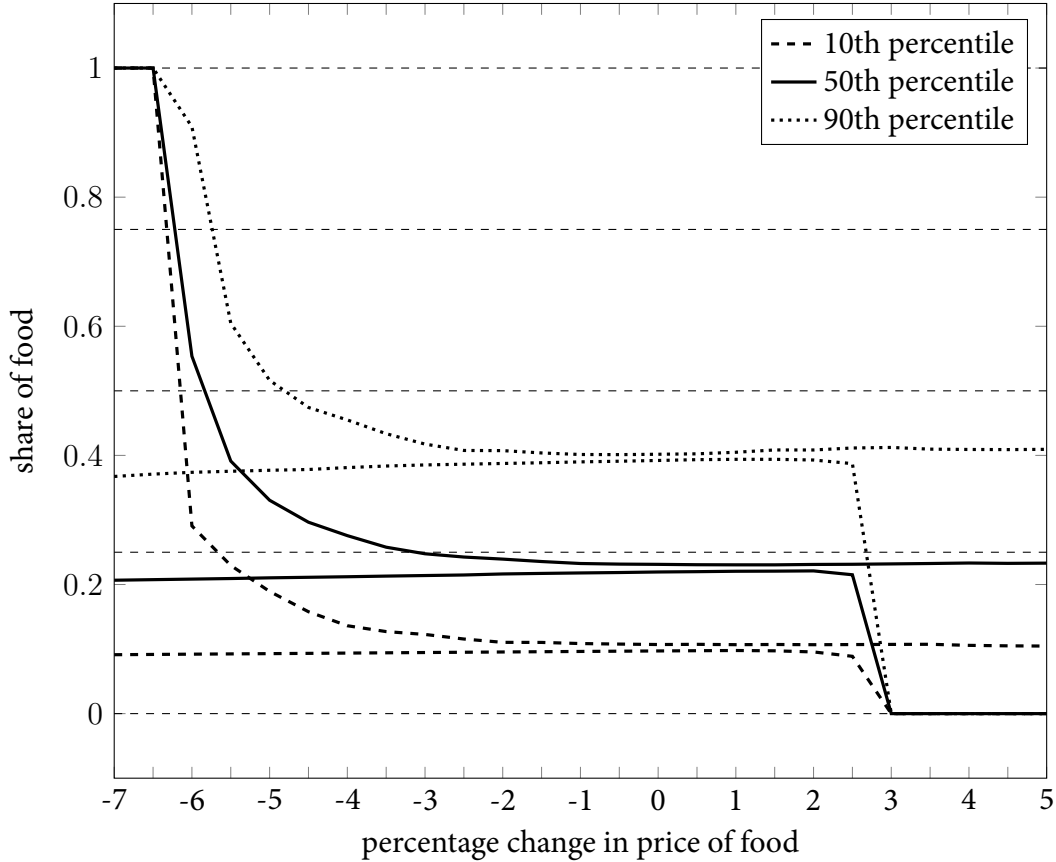


possible to construct bounds on the counterfactual demands for prices in the convex hull of the observed prices. This explains the large and simultaneous drop of all lowerbounds at the 3% price increase, and the large and simultaneous jump of all upperbounds when prices are 6% below the reference level.

5 Conclusions

In this paper, we used elementary revealed preference techniques together with nonparametric estimation techniques in order to bound the distribution of the money metric utility and the demand functions over a population of heterogeneous households. Our methodology has two attractive

Figure 7: Bounds on the demand shares



features. First of all, the results are entirely nonparametric which means that they are independent on any functional form imposed on the underlying utility functions. Second, we impose minimal conditions on the structure of the individual, unobserved heterogeneity. When we apply our techniques to data from the US consumer expenditure survey, we find that our method generates narrow and informative bounds on the quantiles of the money metric utility function. Our results also demonstrate that individual heterogeneity creates considerable variation in welfare between households in the population (conditional on the same level of expenditure). We also demonstrate how our results can be used to obtain informative bounds on the distribution of the demand functions in counterfactual price-income situations.

There are several avenues for follow up research. First of all, we only briefly touched upon the highly relevant topic of statistical inference. However, given that our data is obtained from a random sample, measurement error and small sample biases influence our bounds, and statistical inference becomes relevant. Next, it would be interesting to see how our methodology extends to discrete choice settings. One way to incorporate discrete choices would be to consider a setting where individuals make discrete choices in addition to continuous choices. Many of the results from this paper readily extend to such setting. Alternatively one could imagine a setting where all choices are discrete (see Manski (2007) and Sher, il Kim, Fox, and Bajari (2011) for a theoretical account of stochastic revealed preferences recovery in such setting). It would be interesting to look how the methodology developed in this paper transfers to such discrete choice setting. Finally, it would be interesting to see how other (more strict) stochastic revealed preference axioms that explicitly take into account transitivity may even further improve our bounds.

References

- Afriat, S. N., 1967. The construction of utility functions from expenditure data. *International Economic Review* 8, 67–77.
- Afriat, S. N., 1973. On a system of inequalities in demand analysis: An extension of the classical method. *International Economic Review* 14, 460–472.
- Andrews, D. W. K., 2000. Inconsistency of the bootstrap when a parameter is on the bound of the parameter space. *Econometrica* 68, 399–405.
- Bandyopadhyay, T., Dasgupta, I., Pattanaik, P. K., 1999. Stochastic revealed preference and the theory of demand. *Journal of Economic Theory* 84, 95–110.
- Bandyopadhyay, T., Dasgupta, I., Pattanaik, P. K., 2002. Demand aggregation and the weak axiom of stochastic revealed preference. *Journal of Economic Theory* 107, 483–489.

- Bandyopadhyay, T., Dasgupta, I., Pattanaik, P. K., 2004. A general revealed preference theorem for stochastic demand behavior. *Economic Theory* 23, 589–599.
- Banks, J., Blundell, R., Lewbel, A., 1997. Quadratic engel curves and consumer demand. *The Review of Economics and Statistics* 79, 527–539.
- Barberá, S., Pattanaik, P. K., 1986. Falmagne and the rationalizability of stochastic choices in terms of random orderings. *Econometrica* 54, 707–715.
- Bhattacharya, D., 2015. Nonparametric welfare analysis for discrete choice. *Econometrica* 83, 617–649.
- Block, H. D., Marschak, J., 1959. Random orderings and stochastic theories of response. Discussion Paper 66, Cowles Foundation.
- Blundell, R., 2005. Presidential address: How revealing is revealed preference? *Journal of the European Economic Association* 3, 211–235.
- Blundell, R., Browning, M., Cherchye, L., Crawford, I., De Rock, B., Vermeulen, F., 2015. Sharp for SARP: Nonparametric bounds on the behavioral and welfare effects of price changes. *American Economic Journal: Microeconomics* 7, 43–60.
- Blundell, R., Browning, M., Crawford, I., 2003. Nonparametric engel curves and revealed preference. *Econometrica* 71, 205–240.
- Blundell, R., Browning, M., Crawford, I., 2007. Improving revealed preference bounds on demand responses. *International Economic Review* 48, 1227–1244.
- Blundell, R., Browning, M., Crawford, I., 2008. Best nonparametric bounds on demand responses. *Econometrica* 76, 1227–1262.

- Blundell, R., Horowitz, J., Parey, M., 2013. Nonparametric estimation of a heterogeneous demand function under the Slutsky inequality restriction. Tech. rep., cemmap.
- Blundell, R., Kristensen, D., Matzkin, R., 2014. Bounding quantile demand functions using revealed preference inequalities. *Journal of Econometrics* 179, 112–127.
- Bronars, S. G., 1987. The power of nonparametric tests of preference maximization. *Econometrica* 55, 693–698.
- Brown, B. W., Walker, M. B., 1989. The random utility hypothesis and inference in demand systems. *Econometrica* 57, 815–829.
- Bugni, F. A., 2010. Bootstrap inference in partially identified models defined by moment inequalities: Coverage of the identified set. *Econometrica* 78, 735–753.
- Cai, Z., 2002. Regression quantiles for time series. *Econometric Theory* 18, 169–192.
- Cherchye, L., De Rock, B., Vermeulen, F., 2007. The collective model of household consumption: a nonparametric characterization. *Econometrica* 75, 553–574.
- Chernozhukov, V., Hong, H., Tamer, E., 2007. Estimation and confidence regions for parameter sets in econometric models. *Econometrica* 75, 1243–1284.
- Chernozhukov, V., Lee, S., Rosen, A. M., 2013. Intersection bounds: Estimation and inference. *Econometrica* 81, 667–737.
- Chiappori, P.-A., 1988. Rational household labor supply. *Econometrica* 56, 63–89.
- Chiappori, P.-A., 1992. Collective labor supply and welfare. *Journal of Political Economy* 100, 437–467.
- Christensen, L. R., Jorgenson, D. W., Lau, L. J., 1975. Transcendental logarithmic utility functions. *American Economic Review* 65, 367–383.

- Cohen, M., Falmagne, J.-C., 1990. Random utility representation of binary choice probabilities: A new class of necessary conditions. *Journal of Mathematical Psychology* 34, 88–94.
- Cohen, M. A., 1980. Random utility systems – the infinite case. *Journal of Mathematical Psychology* 22, 1–23.
- Deaton, A., Muellbauer, J., 1980. *Economics and Consumer Behavior*. Cambridge University Press.
- Detle, H., Hoderlein, S., Neumeyer, N., 2014. Testing multivariate economic restrictions using quantiles: The example of Slutsky negative semidefiniteness. Tech. rep.
- Diewert, W. E., 1973. Afriat and revealed preference theory. *The Review of Economic Studies* 40, 419–425.
- Falmagne, J.-C., 1978. A representation theorem for finite random scale systems. *Journal of Mathematical Psychology* 18, 52–72.
- Fishburn, P. C., 1978. Choice probabilities and choice functions. *Journal of Mathematical Psychology* 18, 205–219.
- Fishburn, P. C., 1992. Induced binary probabilities and the linear ordering polytope: A status report. *Mathematical Social Sciences* 23, 67–80.
- Fishburn, P. C., Falmagne, J.-C., 1989. Binary choice probabilities and rankings. *Economics Letters* 31, 113–117.
- Guilbaud, G. T., 1953. Sur une difficulté de la théorie du risque. *Colloques Internationaux du Centre National de Recherche Scientifique, Econometrie* 40, 19–25.
- Hausman, J. A., Newey, W. K., 2013. Individual heterogeneity and average welfare. Tech. rep., cemmap.

- Hoderlein, S., 2011. How many consumers are rational. *Journal of Econometrics* 164, 294–309.
- Hoderlein, S., Stoye, J., 2014a. Revealed preference in a heterogeneous population. *Review of Economics and Statistics* 96, 197–213.
- Hoderlein, S., Stoye, J., 2014b. Testing stochastic rationality and predicting stochastic demand: The case of two goods. *Economic Theory Bulletin*.
- Hoderlein, S., Vanhems, A., 2013. Estimating the distribution of welfare effects using quantiles. Tech. rep.
- Houthakker, H. S., 1950. Revealed preference and the utility function. *Economica* 17, 159–174.
- Imbens, G. W., Manski, C. F., 2004. Confidence interval for partially identified parameters. *Econometrica* 72, 1845–1857.
- Kawaguchi, K., 2012. Nonparametric test of axioms of revealed stochastic preference. Tech. rep., LSE.
- Kitamura, Y., Stoye, J., 2013. Nonparametric analysis of random utility models. CeMMMap Working Paper CWP36/.13.
- Lewbel, A., 2001. Demand systems with and without errors. *American Economic Review* 91, 611–618.
- Li, Q., Racine, S., 2007. *Nonparametric Econometrics*. Princeton University Press, Princeton.
- Manski, C. F., 2007. Partial identification of counterfactual choice probabilities. *International Economic Review* 48, 1393–1410.
- Marschak, J., 1960. *Stanford Symposium on Mathematical Methods in the Social Sciences*. Stanford University Press, Ch. Binary Choice Constraints on Random Utility Indicators.

- Matzkin, R. L., 2007. Handbook of Econometrics. Elsevier, Ch. Nonparametric Identification.
- McElroy, M. B., 1987. Additive general error models for production, cost, and derived demand or share systems. *Journal of Political Economy* 95, 737–757.
- McFadden, D., 1975. Tchebycheff bounds for the space of agent characteristics. *Journal of Mathematical Economics* 2, 225–242.
- McFadden, D., Richter, M. K., 1971. On the extension of a set function on a set of events to a probability on the generated boolean σ -algebra. Tech. rep., University of California, Berkeley.
- McFadden, D. L., 2005. Revealed stochastic preference: a synthesis. *Economic Theory* 26, 245–264.
- Newey, W. K., Powell, J. L., Vella, F., 1999. Nonparametric estimation of triangular simultaneous equations models. *Econometrica* 67, 565–603.
- Politis, D. N., Romano, J. P., Wolf, M., 1999. *Subsampling*. Springer.
- Pollak, R. A., 1971. The theory of the cost of living index. Research Discussion Paper 11, Bureau of Labor Statistics, Washington DC.
- Rose, H., 1958. Consistency of preference: The two-commodity case. *Review of Economic Studies* 25, 124–125.
- Samuelson, P. A., 1938. A note on the pure theory of consumer's behavior. *Economica* 5, 61–71.
- Sher, I., il Kim, K., Fox, J. T., Bajari, P., 2011. Partial identification of heterogeneity in preference orderings over discrete choices. Tech. rep., NBER.
- Stoye, J., 2009. More on confidence intervals for partially identified parameters. *Econometrica* 77, 1299–1315.

Tucker, H., 1967. A Graduate Course in Probability. Academic Press, New York.

Varian, H. R., 1982. The nonparametric approach to demand analysis. *Econometrica* 50, 945–974.

A The Weak Axiom of Stochastic Revealed Preference

If $r_{t,v}(x, y) + r_{v,t}(y, x)$ is larger than one, we see two possible solutions. A first solution is to allow a certain fraction of the population to violate the Weak Axiom of Revealed Preference, i.e. a certain subset of the population is considered to be irrational. Applying this solution would amount to subtracting a certain percentage, that equals the fraction of irrational households, from $r_{t,v}(x, y)$ and $r_{v,t}(y, x)$, thereby widening the range of possible values for $\Pr[v_h(\mathbf{p}_t, x) \geq v_h(\mathbf{p}_v, y)]$. A second solution is to relax the rationality constraints for all households simultaneously. In order to do this, we can consider an adaptation of an early proposal of Afriat (1973) for revealed preference tests in a non-stochastic setting to our specific setting. In particular, we capture optimization error by a so-called Afriat index $e \in [0, 1]$. For a given value of e , the new rationality criterion adjusts condition (1) in the following way,

$$\text{If } e \cdot x \geq \mathbf{p}_t \mathbf{q}_h(\mathbf{p}_v, y), \text{ then } v_h(\mathbf{p}_t, x) > v_h(\mathbf{p}_v, y).$$

Intuitively, we only check whether behaviour is rational while allowing the household to waste as much as $(1 - e)$ of the income x by making irrational choices. Using this Afriat index, we can construct the following probabilities,

$$r_{t,v}^e(x, y) = \Pr[e \cdot x \geq \mathbf{p}_t \mathbf{q}_h(\mathbf{p}_v, y)].$$

The number $r_{t,v}^e(x, y)$ is increasing in e and $r_{t,v}^0(x, y) = 0$. Given this, there will always be a value of $e \in [0, 1]$ such that,

$$r_{t,v}^e(x, y) + r_{v,t}^e(y, x) \leq 1.$$

The analysis could then proceed by replacing $r_t(x, y)$ by the numbers $r_t^{e^*}(x, y)$ where e^* is either fixed a priori or coincides with the largest number for which this inequality holds.

A second problem arises if $r_{t,v}(x, y) + r_{v,t}(y, x)$ is considerably below 1. In such cases, the bounds are not very informative. A solution is to impose a stronger stochastic revealed preference condition. In the construction of $r_{t,v}(x, y)$ above, we only used information concerning the two budget sets (\mathbf{p}_t, x) and (\mathbf{p}_v, y) . In some cases, however, it is possible to include information on additional budget sets and transitivity of the preference relation to obtain tighter bounds. One such tightening relies on the fact that for any three distinct numbers a, b and c it is always the case that,

$$\Pr(a > c) \geq \Pr(a > b) + \Pr(b > c) - 1.$$

Indeed, the probability that c is larger than b is given by $1 - \Pr(b > c)$. As such, $\Pr(a > b > c)$ is bounded from below by $\Pr(a > b) - (1 - \Pr(b > c))$. The conclusion then follows from the fact that $\Pr(a > c) \geq \Pr(a > b > c)$. Rewriting above condition shows that it is equivalent to the famous triangle inequality.

$$\Pr(b > c) \leq \Pr(b > a) + \Pr(a > c).$$

The triangle inequality has first been noted by Guilbaud (1953) and has been popularized by Marschak (1960). The inequality is one of the key conditions in the literature on binary probability systems. This literature, which is closely related to the literature on stochastic revealed

preference theory, tries to characterize all collections of binary probabilities over a finite set of alternatives that are induced by probability distributions over the family of linear orders (preference relations) on this set.

If we apply above condition to our setting and use the previously established bounds, we obtain that for all $t, v, w \in T$ and all incomes x, y, z ,

$$\begin{aligned} \Pr [v_h(\mathbf{p}_t, x) \geq v_h(\mathbf{p}_v, y)] &\geq \Pr [v_h(\mathbf{p}_t, x) \geq v_h(\mathbf{p}_w, z)] + \Pr [v_h(\mathbf{p}_w, z) \geq v_h(\mathbf{p}_v, y)] - 1, \\ &\geq r_{t,w}(x, z) + r_{w,v}(z, y) - 1. \end{aligned}$$

In cases where $r_{t,v}(x, y)$ is lower than

$$\max_{w,z} \{r_{t,w}(x, z) + r_{w,v}(z, y) - 1\},$$

this improves the lower bound on $\Pr[v_h(\mathbf{p}_t, x) \geq v_h(\mathbf{p}_v, y)]$. Of course this tightening of the bounds can be iterated until no further improvements are possible. If the range is still too wide, further tightening could still be obtained by using other, though more elaborate ‘binary probability system’ conditions (see, for example, Fishburn (1992) for an overview of the various kinds of conditions that could be imposed).

B Proof of consistency and asymptotic normality of $\hat{\ell}_t(\pi)$

In this appendix we demonstrate the consistency and asymptotic normality of $\hat{\ell}_t(\pi)$. The proofs are similar to Cai (2002). For ease of notation, we write $\hat{\ell}$ for $\hat{\ell}_t(\pi)$ and ℓ for $\ell_t(\pi)$. We first show that $\hat{\ell} \xrightarrow{P} \ell$.

Consistency Given that $\hat{r}_{t,v}(x, y) \xrightarrow{P} r_{t,v}(x, y)$ and both functions are monotone in y and bounded, it follows from Tucker (1967, Theorem 1) that,

$$\sup_y |\hat{r}_{t,v}(x, y) - r_{t,v}(x, y)| \xrightarrow{P} 0.$$

For any $\varepsilon > 0$, set,

$$\delta(\varepsilon) = \min\{\pi - r_{t,v}(x, \ell + \varepsilon), r_{t,v}(x, \ell - \varepsilon) - \pi\} > 0.$$

This uses the fact that $r_{t,v}(x, y)$ is strictly decreasing in y at (x, y) .

Lemma 1. *If $|\ell - \hat{\ell}| \geq \varepsilon$ then $|r_{t,v}(x, \hat{\ell}) - \hat{r}_{t,v}(x, \hat{\ell})| \geq \delta(\varepsilon)$.*

Proof. Assume first that $\ell \geq \hat{\ell}$. Then we have that $\hat{\ell} \leq \ell - \varepsilon$. This implies $r_{t,v}(x, \hat{\ell}) \geq r_{t,v}(x, \ell - \varepsilon)$.

As such,

$$r_{t,v}(x, \hat{\ell}) - \hat{r}_{t,v}(x, \hat{\ell}) \geq r_{t,v}(x, \ell - \varepsilon) - \pi \geq \delta(\varepsilon).$$

This shows that $|r_{t,v}(x, \hat{\ell}) - \hat{r}_{t,v}(x, \hat{\ell})| \geq \delta(\varepsilon)$.

If $\ell < \hat{\ell}$ we have that $\hat{\ell} > \ell + \varepsilon$ and consequentially, $r_{t,v}(x, \hat{\ell}) \leq r_{t,v}(x, \ell + \varepsilon)$. As such,

$$\hat{r}_{t,v}(x, \hat{\ell}) - r_{t,v}(x, \hat{\ell}) \geq \pi - r_{t,v}(x, \ell + \varepsilon) \geq \delta(\varepsilon).$$

Again, we see that $|r_{t,v}(x, \hat{\ell}) - \hat{r}_{t,v}(x, \hat{\ell})| \geq \delta(\varepsilon)$. As such, we see that for every $\varepsilon > 0$,

$$\Pr\left(|\ell - \hat{\ell}| \geq \varepsilon\right) \leq \Pr\left(|\hat{r}_{t,v}(x, \hat{\ell}) - r_{t,v}(x, \hat{\ell})| \leq \delta(\varepsilon)\right).$$

As the right hand side goes to zero by consistency of $\hat{r}_{t,v}(x, y)$ for $n \rightarrow \infty$, $h \rightarrow 0$ and $nh \rightarrow \infty$, the left hand side also goes to zero. Given this, we see that $\hat{\ell} \xrightarrow{P} \ell$. \square

Asymptotic normality Let us now look at the asymptotic distribution of $\hat{\ell}_n$. We know that $\hat{r}_{t,v}(x, y)$ has the following limiting distribution,

$$\frac{\sqrt{nh}}{V(x, y)^{1/2}}(\hat{r}_{t,v}(x, y) - r_{t,v}(x, y) - B(x, y)) \rightarrow \mathcal{N}(0, 1),$$

Where $B(x, y) = O(h^2)$ is the asymptotic bias and $V(x, y)$ is the asymptotic variance. We first consider two lemma's.

Lemma 2. *If $\varepsilon_n/h = o(1)$, then,*

$$\hat{r}_{t,v}(x, y + \varepsilon_n) - \hat{r}_{t,v}(x, y) = \frac{\partial r_{t,v}(x, y)}{\partial y} \varepsilon_n + o_P(\varepsilon_n) + o_P(h^2) + o_P((nh)^{-1/2}).$$

Proof. The proof involves straightforward but long and tedious computation. As such, it is available upon request. □

Lemma 3. *Fix a number $v < \infty$ and define $\varepsilon_n = \frac{vV(x, y)^{1/2} - B(x, y)}{\left| \frac{\partial r_{t,v}(x, y)}{\partial y} \right|}$. Then if $h \rightarrow 0$, $nh^3 \rightarrow \infty$ and $nh^5 = O(1)$, then,*

1. $\varepsilon_n/h = o(1)$.
2. $\sqrt{nh} o(\varepsilon_n) = o(1)$.

Proof. Easy. □

Now we are ready to demonstrate the following,

$$\sqrt{nh} \frac{\left| \frac{r_{t,v}(x, \ell)}{\partial \ell} \right|}{V(x, \ell)^{1/2}} \left(\hat{\ell}_n - \ell + \frac{B(x, \ell)}{\left| \frac{\partial r_{t,v}(x, \ell)}{\partial \ell} \right|} \right) \rightarrow^d N(0, 1).$$

Proof. Denote by $\Phi(v)$ the normal cumulative distribution function, then

$$\begin{aligned}
& \Pr \left((nh)^{1/2} \frac{\left| \frac{\partial r_{t,v}(x,\ell)}{\partial \ell} \right|}{V(x,\ell)^{1/2}} \left(\hat{\ell}_n - \ell + \frac{B(x,\ell)}{\left| \frac{\partial r_{t,v}(x,\ell)}{\partial \ell} \right|} \right) \leq v \right), \\
&= \Pr \left(\hat{\ell}_n - \ell \leq \frac{\frac{vV(x,\ell)^{1/2}}{(nh)^{1/2}} - B(x,\ell)}{\left| \frac{\partial r_{t,v}(x,\ell)}{\partial \ell} \right|} \right), \\
&= \Pr \left(\hat{\ell}_n - \ell \leq \varepsilon_n \right), \\
&= \Pr \left(\hat{\ell}_n \leq \varepsilon_n + \ell \right), \\
&= \Pr \left(\hat{r}_{t,v}(x, \hat{\ell}_n) \geq \hat{r}_{t,v}(x, \ell + \varepsilon_n) \right), \\
&= \Pr \left(\hat{r}_{t,v}(x, \hat{\ell}_n) \geq \hat{r}_{t,v}(x, \ell) + \frac{\partial r_{t,v}(x, \ell)}{\partial \ell} \varepsilon_n + o_P(\varepsilon_n) + o_P(h^2) + o_P((nh)^{-1/2}) \right), \\
&= \Pr \left(\hat{r}_{t,v}(x, \hat{\ell}_n) - \hat{r}_{t,v}(x, \ell) \geq - \left| \frac{\partial r_{t,v}(x, \ell)}{\partial \ell} \right| \varepsilon_n + o_P(\varepsilon_n) + o_p(h^2) + o_p((nh)^{-1/2}) \right), \\
&\sim \Pr \left((nh)^{1/2} \frac{1}{V(x, y)^{1/2}} \left(\hat{r}_{t,v}(x, \ell) - \hat{r}_{t,v}(x, \hat{\ell}_n) - B(x, \ell) \right) \geq -v \right) \\
&= \Pr \left((nh)^{-1/2} \frac{1}{V(x, y)^{1/2}} \left(\hat{r}_{t,v}(x, \ell) - r_{t,v}(x, \ell) - B(x, \ell) \right) \geq -v \right) \rightarrow^d \Phi(v).
\end{aligned}$$

□

C Construction of aggregates

Food is an aggregate of cereals, bakery products, beef, pork, poultry, seafood, other meat, eggs, milk products, other dairy products, fresh fruit, fresh vegetables, processed fruit, processed vegetables, sweets, fat and oils, non-alcoholic beverages, prepared food, snacks and condiments.

Other non-durables contain expenditures on alcohol consumption, tobacco, clothes (for all household members), footwear, reading material, stationery, school supplies, cleaning products, garden supplies, household textile, non-durable housewares, medical products, personal care products, audio-visual equipment, recreational goods, pet goods and vehicle expenses.

Services include utilities, media bills, repair services, insurance, postal services, gasoline, vehicle expenses (services), public transportation, medical care services, personal care services, recreational services, home services, rental services, membership fees, school fees, other fees, pet services and care services.

D Confidence intervals

To compute the Bonferroni intervals we followed a subsampling procedure (in line with Politis, Romano, and Wolf (1999)). Subsampling is similar to the bootstrap procedure but the samples are smaller and the draws are obtained without replacement. Consider a dataset of size n and an estimator \hat{g}_n which converges at a rate such that $\tau_n(\hat{g} - g)$ converges to a non-degenerate asymptotic distribution for $n \rightarrow \infty$. In our case, $\tau_n = \sqrt{nh_n}$ where h_n is the bandwidth. The subsampling procedure proceeds by taking (without replacement) subsamples of size m and compute the associated estimator g_m^* . Then, under very weak conditions, it can be shown that for $m \rightarrow \infty$, $m/n \rightarrow 0$ and $\tau_m/\tau_n \rightarrow 0$ as $n \rightarrow \infty$, the statistic $\tau_m(g^* - \hat{g})$ converges to the same asymptotic distribution as $\tau_n(\hat{g} - g)$.

We apply the subsampling approach to the distribution of the (Laspeyres-Konüs) cost of living index using 999 subsamples of size $m \approx \sqrt{n}$.¹⁹ Table 2 presents our original estimates of upper and lower bounds on the quantiles of the Laspeyres-Konüs cost of living index and the corresponding asymptotic 95% confidence intervals for the setting in Figure 3.

¹⁹Other values of m give similar results.

Table 2: 95% confidence bounds on the distribution of the Laspeyres-Konüs cost of living index - bounds: sample estimates of lower and upper bounds on cost of living.

year	10th percentile		30th percentile		50th percentile		70th percentile		90th percentile	
bounds 1994	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
CI 1994	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
bounds 1995	1.0186	1.0235	1.0228	1.0279	1.0255	1.0303	1.0277	1.0321	1.0303	1.0342
CI 1995	1.0176	1.0246	1.0219	1.0286	1.0247	1.0308	1.0270	1.0326	1.0295	1.0347
bounds 1996	1.0506	1.0549	1.0554	1.0607	1.0584	1.0634	1.0606	1.0659	1.0638	1.0685
CI 1996	1.0491	1.0561	1.0543	1.0615	1.0575	1.0640	1.0599	1.0665	1.0630	1.0691
bounds 1997	1.0697	1.0761	1.0784	1.0850	1.0832	1.0896	1.0881	1.0933	1.0934	1.0974
CI 1997	1.0678	1.0781	1.0767	1.0862	1.0819	1.0905	1.0867	1.0942	1.0922	1.0983
bounds 1998	1.0697	1.0787	1.0845	1.0940	1.0933	1.1018	1.0998	1.1082	1.1090	1.1155
CI 1998	1.0661	1.0823	1.0815	1.0961	1.0909	1.1035	1.0976	1.1096	1.1067	1.1171
bounds 1999	1.1028	1.1094	1.1128	1.1218	1.1198	1.1285	1.1271	1.1339	1.1343	1.1394
CI 1999	1.1004	1.1120	1.1109	1.1235	1.1180	1.1297	1.1255	1.1350	1.1327	1.1407
bounds 2000	1.1590	1.1646	1.1654	1.1704	1.1694	1.1749	1.1730	1.1779	1.1782	1.1819
CI 2000	1.1579	1.1660	1.1643	1.1714	1.1684	1.1756	1.1720	1.1785	1.1769	1.1828
bounds 2001	1.1799	1.1873	1.1934	1.2021	1.2036	1.2111	1.2126	1.2185	1.2230	1.2274
CI 2001	1.1774	1.1904	1.1910	1.2041	1.2008	1.2127	1.2099	1.2199	1.2206	1.2291
bounds 2002	1.1798	1.1902	1.1987	1.2131	1.2135	1.2262	1.2262	1.2367	1.2449	1.2497
CI 2002	1.1766	1.1949	1.1955	1.2161	1.2107	1.2286	1.2230	1.2390	1.2407	1.2523
bounds 2003	1.2157	1.2276	1.2405	1.2519	1.2545	1.2676	1.2684	1.2798	1.2901	1.2948
CI 2003	1.2117	1.2324	1.2366	1.2553	1.2516	1.2705	1.2651	1.2823	1.2855	1.2975
bounds 2004	1.2735	1.2792	1.2918	1.3004	1.3044	1.3147	1.3174	1.3259	1.3350	1.3397
CI 2004	1.2697	1.2832	1.2887	1.3035	1.3015	1.3172	1.3147	1.3281	1.3317	1.3425
bounds 2005	1.3360	1.3402	1.3537	1.3581	1.3659	1.3710	1.3779	1.3819	1.3931	1.3946
CI 2005	1.3325	1.3434	1.3506	1.3607	1.3631	1.3732	1.3749	1.3839	1.3901	1.3973
bounds 2006	1.3849	1.3890	1.4033	1.4083	1.4171	1.4223	1.4308	1.4350	1.4430	1.4486
CI 2006	1.3809	1.3926	1.4000	1.4115	1.4139	1.4250	1.4268	1.4371	1.4404	1.4520
bounds 2007	1.4328	1.4385	1.4523	1.4571	1.4640	1.4696	1.4754	1.4813	1.4933	1.4952
CI 2007	1.4298	1.4421	1.4490	1.4600	1.4612	1.4722	1.4728	1.4835	1.4899	1.4979