NBER WORKING PAPER SERIES

A NEW METHOD FOR MEASURING WELFARE WITH INCOME EFFECTS USING CROSS-SECTIONAL DATA

David Baqaee
Ariel Burstein
Yasutaka Koike-Mori

Working Paper 30549
http://www.nber.org/papers/w30549

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2022

We thank John Asker and Jon Vogel for comments. We acknowledge financial support from NSF grant No. 1947611 and the Alfred P. Sloan Foundation. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

A New Method for Measuring Welfare with Income Effects using Cross-Sectional Data
David Baqaee, Ariel Burstein, and Yasutaka Koike-Mori
NBER Working Paper No. 30549
October 2022
JEL No. E01,E31

ABSTRACT

We show how to recover the money-metric utility function, which converts income at one point in time into equivalent income at another point in time, using repeated cross-sectional household data. Our procedure allows unrestricted preferences, but requires that households' preferences be the same in both the cross-section and the time-series. In prior work, Jaravel and Lashkari (2022) provide a solution to this problem. We leverage a different theoretical insight to address this problem. Our idea is to trace out Hicksian (or compensated) demand curves through time by matching households on the same indifference curve at different points in time. Given Hicksian demand curves, we can construct cost-of-living indices and money-metric utility for every matched income level. We apply our method to household consumption survey data from the United Kingdom from 1974 to 2017. We find that the official annual inflation rate understates welfare-relevant inflation for the poorest households by around half a percentage point per year and overstates it for the richest households by around a quarter of a percentage point per year.

David Baqaee
Department of Economics
University of California at Los Angeles
Bunche Hall
Los Angeles, CA 90095
and CEPR
and also NBER
baqaee@econ.ucla.edu

Ariel Burstein
Department of Economics
Bunche Hall 8365
Box 951477
UCLA
Los Angeles, CA 90095-1477
and NBER
arielb@econ.ucla.edu

Yasutaka Koike-Mori
University of California, Los Angeles
yasutaka.koike.mori@gmail.com

1 Introduction

An ideal cost of living index is one that converts expenditures or income at one point in time into equivalent income at other points in time — that is, the income that makes households indifferent between their budget sets at both points in time. It is well known that standard price indices do not perform this task properly if preferences are non-homothetic. (See, e.g., Samuelson and Swamy, 1974). Jaravel and Lashkari (2022) develop a non-parametric approach to adjust standard price indices to correct for non-homotheticities using cross-sectional household-level data. In this paper, we develop a complementary solution to this problem that relies on a different theoretical insight.[1]

Our procedure, which builds on Baqaee and Burstein (2021), has an intuitive interpretation using Hicksian, or compensated, demand curves. To recover an ideal cost of living index or the money-metric utility function, one must integrate Hicksian demand curves with respect to changes in prices. More prosaically, this means that to construct a welfare-relevant inflation index, we must weigh changes in prices at each point in time using Hicksian demand for each good. In practice, one does not observe the Hicksian demand curves. In lieu of Hicksian demand, standard price indices (i.e. chain-weighted or Divisia) weigh changes in prices by observed — that is, Marshallian — demand curves instead. When preferences are homothetic, this practice is innocuous since for homothetic preferences income affects demand for all goods in the same way. However, when preferences are nonhomothetic, Marshallian and Hicksian budget shares are different and usual practice fails to recover a well-defined welfare measure.

To resolve this problem, we propose the following. Suppose we observe repeated cross-sections of households with identical preferences facing common prices. To construct the money-metric utility value in terms of some base time t_0 for a household with income I at time t, we must know the compensated demand of this household when prices were different. This can be done at each point in time $s \neq t$ by finding another household with a different income level $I' \neq I$ at time s who is on the same indifference curve as the household with income I at t. If we can find such a household at every point in time $t_0 \leq s < t$, then we can trace the compensated Hicksian demand curve through time, and by using these demand curves, rather than observed demand, we can calculate the money-metric utility function.[2] We show how to find these "compensated" matching

[1]For an alternative approach to measuring changes in welfare allowing for incomplete information on prices and imposing additional restrictions on preferences, see e.g. Hamilton (2001) and, more recently, Atkin et al. (2020).

[2]Hicksian demand can also be calculated given knowledge of elasticities of substitution. The procedure in this paper does not require nor allow one to estimate elasticities of substitution non-parametrically however. Intuitively, we only recover Hicksian demand curves evaluated at observed prices, whereas the

households, if they exist, by solving a recursive problem.

Like Jaravel and Lashkari (2022), our procedure is non-parametric in the sense that it works for any well-behaved preference relation. Similar to Jaravel and Lashkari (2022), the most important assumption we make is that all households in the sample have the same preference relation and that these preferences are not changing through time.[3] Although inspired by their paper, our approach is different to theirs. Rather than matching households on the same indifference curve through time, as we do, Jaravel and Lashkari (2022) use a correction term (which depends on the elasticity of the expenditure function with respect to utility) to correct the household-level chain-weighted index for non-homotheticity.[4]

One advantage of our approach is that we do not assume that the support of the cross-sectional distribution of (ex-ante unknown) utility is constant and unchanging through time. Instead our procedure endogenously determines the region of the income distribution where a money-metric utility function can be constructed given the available information. That is, our procedure endogenously determines the set of households for whom it is possible to find a matching/compensated household in the past. This is important in contexts economic growth shifts the support of utilities over time.

The outline of the paper is as follows. In Section 2, we define the cost-of-living index and money-metric utility and introduce a preliminary result taken from Baqaee and Burstein (2021). In Section 3, we show how this result can be used to recover cost-of-living and money-metric utility with the aid of cross-sectional data. We also discuss some extensions and limitations of our approach, including how our results can be used when some of our baseline assumptions are relaxed. For example, with enough data, we discuss how to handle idiosyncratic taste shocks that are uncorrelated with income. Similarly, we discuss how to apply our method if there is heterogeneity in preferences that is a function of observable characteristics. We also discuss how one can handle changes in quality that are not reflected in prices.

In Section 4, we apply our results to artificial data generated using popular functional forms for non-homothetic preferences. We illustrate that our procedure quickly converges to the truth as the number of households and frequency of observations grow. Our numerical examples are calibrated to match real-world data in terms of the frequency of

elasticities of substitution allow one to compute counterfactual Hicksian demand even for prices that are not observed.

[3]That is, in our baseline, we rule out demand/taste shocks in both the time-series and cross-section. We discuss below how this assumption may be relaxed.

[4]In Appendix C, we apply the Jaravel and Lashkari (2022) approach to our artificial data and the UK data.

observation and the rate at which prices and incomes are changing over time.

In Section 5, we apply our method to construct a money-metric utility function using household expenditure survey data from the United Kingdom from 1974 to 2017. We find that aggregate chain-weighted measures of inflation (following procedures of official statistics) understate the true inflation rate for all households below the 60th percentile of income in 2017 in our sample. In other words, for any income level in 2017 under the 60th percentile, the 1974 equivalent income is less than real income implied by an aggregate chain-weighted inflation index. The size of this gap is greatest for the poorest households, roughly 25 percentage points (0.5 percentage points per year on average), and declines to zero for households close to the 60th percentile. Conversely, aggregate chain-weighted measures of inflation overestimate the true inflation rate for households above the 60th level. For households in the 97th percentile of our sample, who spend around £82,000 per year, the inflation rate is overstated by 13 percentage points over the whole sample (0.25 percentage points per year on average). We are unable to compute the ideal inflation rate for the richest households in 2017 (97th percentile and above). The reason is that for the richest households in 2017, there did not exist equally well-off consumers in the past whose demand can be used in place of the compensated demand curves. We conclude in Section 6.

2 Money-Metrics and the Cost-of-Living

Consider a preference relation $\geq$ defined over consumption bundles c in $\mathbb{R}^N$. Suppose that we represent these preferences using a utility function $U(c)$ that maps consumption bundles to utility values. Given this utility function, we can define the *indirect utility function*

$$v(p, I) = \max_c \{U(c) : p \cdot c \leq I\},$$

mapping a vector of prices p and expenditures I to utility values. Define the *expenditure function* to be

$$e(p, u) = \min_c \{p \cdot c : U(c) \leq u\}.$$

The expenditure and indirect utility functions are useful because they can be used to construct money-metrics and cost-of-living indices.

Definition 1 (Money-Metric)**.** For some reference vector of prices $\bar{p}$, the *money-metric* function is

$$m(\bar{p}, p, I) = e(\bar{p}, v(p, I)).$$

The money-metric $m(\bar{p}, \cdot)$ is a specific cardinalization of the indirect utility function in the sense that a budget set $(p, I) \succeq (p', I')$ if, and only if, $m(\bar{p}, p, I) \geq m(\bar{p}, p', I')$. That $m(\bar{p}, p, I)$ expresses the value of the budget set (p, I) in terms of $\bar{p}$ prices. We use this cardinalization of utility throughout the rest of the paper.

Definition 2 (Cost-of-living)**.** For some reference budget constraint $(\bar{p}, \bar{I})$, the *cost-of-living* function is

$$r(p, \bar{p}, \bar{I}) = e(p, v(\bar{p}, \bar{I})).$$

Note that $r(p, \bar{p}, \bar{I})$ converts the value of budget constraint $(\bar{p}, \bar{I})$ into equivalent income under prices p.[5]

In other words, the key object is the function $e(p', v(p, I))$ which maps (p', p, I) into a scalar. The "money-metric" is the cross-section of this function that holds p' constant and the cost-of-living index is the cross-section that holds (p, I) constant. The money-metric is useful for converting different budget sets into a common price system for comparison. On the other hand, the cost-of-living index is useful for converting a common utility level, attained by $v(p, I)$, into equivalent income under different price systems. As its name suggests, the cost-of-living index is necessary for providing a consumer with a cost-of-living adjustment.

Denote the Hicksian budget share for good i to be $b_i(p, u)$ where p is a vector of prices and u is a utility level. The following proposition, which is a corollary of Lemma 1 from Baqaee and Burstein (2021), provides a characterization of both the cost-of-living index and the money-metric using Hicksian budget shares.

Proposition 1 (Money-Metric and Cost-of-Living)**.** *The money-metric of a budget set (p, I) in terms of $\bar{p}$ prices can be expressed as*

$$\log m(\bar{p}, p, I) = \log I - \int_{\bar{p}}^{p} \sum_{i \in N} b_i(\xi, v(p, I)) d \log \xi_i. \tag{1}$$

The cost-of-living for a budget set $(\bar{p}, \bar{I})$ in terms of p prices can be expressed as

$$\log r(p, \bar{p}, \bar{I}) = \log \bar{I} + \int_{\bar{p}}^{p} \sum_{i \in N} b_i(\xi, v(\bar{p}, \bar{I})) d \log \xi_i. \tag{2}$$

Intuitively, both the money-metric and the cost-of-living index can be expressed as integrals of Hicksian budget shares with respect to changes in prices. However, Hicksian demand curves are not directly observable, so operationalizing this result requires having a way to identify Hicksian budget shares. This is what we focus on in the next section.

[5] In index number theory, the cost-of-living index is also called the Konüs (1939) index.

3 Main Results

In this section, we discuss how Proposition 1 can be deployed to recover money-metric utility functions and cost-of-living indices if one has access to repeated cross-sectional data of consumers with common and stable preferences who all face common prices at each point in time but have different incomes. We start this section by introducing our main theoretical result. We then provide a simple numerical implementation. We end the section by discussing some extensions and limitations.

3.1 Theoretical Result

Suppose we observe a smooth path of prices p_t at each point in time $t \in [t_0, T]$ and, for consumers with income level $I \in [\underline{I}_t, \bar{I}_t]$ at time t we observe the vectors of expenditure shares $B(I, t)$ across all goods. The expenditure shares $B(I, t)$ can be thought of as *Marshallian* budget shares evaluated at income level I and prices p_t.

For any cardinalization of the indirect utility function and its associated Hicksian demand curves, the following identity holds

$$b_i(p_t, v(p_t, I)) = B_i(I, t).$$

Since $m(p_{t_0}, p_t, I)$ is the cardinalization of indirect utility that we are working with, we can write

$$b_i(p_t, m(p_{t_0}, p_t, I)) = B_i(I, t).$$

Using this identity, Proposition 1 can be rewritten as the following recursive integral equation.

Proposition 2 (Money-metric as Solution to Integral Equation)**.** *For $t \in [t_0, T]$, the money-metric is $m(p_{t_0}, p_t, I) = u(I, t)$, where $u(I, t)$ is a fixed point of the following integral equation*

$$\log u(I, t) = \log I - \int_{t_0}^{t} \sum_i B_i(u^{-1}(u(I, t), s), s) \frac{d \log p_{is}}{ds} ds \tag{3}$$

with boundary condition $u(I, t_0) = I$. Here, $u^{-1}(\cdot, s)$ is the inverse of u with respect to its first argument (income) given its second argument is equal to s.

In words, $u(I, t)$ converts the value of the budget constraint defined by prices p_t and income I into income under p_{t_0}. That is, $u(I, t) = e(p_{t_0}, v(p_t, I))$. This is the money-metric for (p_t, I) in terms of p_{t_0}. The solution to the integral equation above is $m(p_{t_0}, p_t, I)$ where p_{t_0} are prices at the initial condition. By varying the initial condition and evaluating $m(p_{t_0}, p_t, I)$ for different p_{t_0} for fixed (p_t, I) we recover the cost-of-living index $r(p_{t_0}, p_t, I)$.

Proposition 2 follows immediately from Proposition 1 once we recognize that in the integral equation above, $B_i(u^{-1}(\cdot, s), s) : \mathbb{R}_+ \to [0, 1]$ maps utility values to the budget share of good i at time s. That is, it is the Hicksian budget share of i.

If we can solve the integral equation in Proposition 2, then we can recover the money-metric and cost-of-living functions without direct knowledge of the elasticities of substitution or income elasticities. This is because we can compute the Hicksian budget shares $b(u_t, p_s)$ for utility level u_t at time t under prices p_s at time s by using the budget shares of a household on the same indifference curve at time s. That is, at time s, there is some household on the same indifference curve as u_t. The expenditures shares of this household, $B(u^{-1}(u_t, s), s)$, are equal to the Hicksian budget share $b(u_t, p_s)$. Hence, we can use the budget shares of this "matched" household to weigh prices changes at time s when constructing the money-metric value of the household with utility u_t.[6]

While the integral equation in Proposition 2 appears abstract, the intuition for it is quite simple and becomes more transparent once we go through a step-by-step procedure for solving it. We outline one such procedure below.

3.2 A Step-by-Step Procedure and More Intuition

The money metric is a fixed point of (3), which is a system of nonlinear equations, albeit an infinite-dimensional one. There are established numerical procedures for solving such equations. Here, we show a very simple iterative procedure that converges to the desired solution as we approach the continuous-time limit.[7]

For some interval of time $[t_0, T]$, suppose we have data on a grid of points $\{t_0, \ldots, t_M\}$ where $t_n < t_{n+1}$, with $t_M = T$. Then, use the following iterative procedure for each $n \in \{1, \ldots, M\}$ starting with $n = 1$:

$$\log u(I, t_n) \approx \log I - \sum_{s=0}^{n-1} b(u(I, t_{n-1}), t_s) \cdot \Delta \log p_{t_s}, \tag{4}$$

[6]Proposition 2 can be used to measure changes in welfare for observed changes in prices and income if we have repeated cross-sectional data and common stable preferences. However, we cannot use this procedure to answer counterfactual or macroeconomic welfare questions like those studied by Baqaee and Burstein (2021).

[7]The iterative procedure that we describe is useful for building intuition. However, one can also find a fixed point by solving directly the nonlinear system of equations. That is, we replace $b(u(I, t_{n-1}), t_s)$ in the right hand side of equation (4) with $b(u(I, t_n), t_s)$. When we use this refined algorithm in our artificial data in Section 4, we obtain even smaller errors (by three orders of magnitude) than the ones we report in the text. Moreover, the empirical results in Section 5 are roughly unchanged.

$$b(v, t_s) = B(u^{-1}(v, t_s), t_s). \tag{5}$$

with the boundary condition $u(I, t_0) = I$ and $b(u, t_0) = B(I, t_0)$. The summation in (4) above approximates the integral in (3) using a Riemnann sum and becomes exact in the continuous-time limit because the Riemann sum becomes an integral and $u(I, t_{n-1}) \rightarrow u(I, t_n)$.[8]

For those values of u that can be inverted, this procedure recovers welfare as the time interval shrinks to zero.[9] This procedure endogenously delineates those values of (I, t) for which $u(I, t)$ can be computed, and it does not require an assumption of full constant support over time on either the set of observed incomes or unobserved utilities.

To give more intuition, it helps to explicitly spell out the first few steps of this iterative procedure. We start with the boundary condition $u(I, t_0) = I$ since t_0-equivalent income in t_0 is just initial income. At time t_1, we use

$$\log u(I, t_1) \approx \log I - b(u(I, t_0), t_0) \cdot \Delta \log p_{t_1} = \log I - B(I, t_0) \cdot \Delta \log p_{t_1}$$

where the last equation uses the boundary condition, which implies $b(u(I, t_0), t_0) = B(I, t_0)$ and becomes exact in the continuous time limit as the gap between t_0 and t_1 shrinks to zero.

Given $u(I, t_1)$, we construct Hicksian budget shares at t_1:

$$b(u, t_1) = B(u^{-1}(u, t_1), t_1)$$

for all u's for which $B(u^{-1}(\cdot, t_1), t_1)$ is observed. That is, to each budget share $B(I, t_1)$ in t_1, we assign a utility value based on $u(I, t_1)$. Hence, we now have Hicksian budget shares $b(u, t_0)$ and $b(u, t_1)$ for all values of $u \in [\underline{I}_0, \overline{I}_0]$. For $u(I, t_1)$ values outside of $[\underline{I}_0, \overline{I}_0]$, we cannot compute Hicksian budget shares in t_0 since there are no households in t_0 who are on the same indifference curve as $u(I, t_1)$.

Next, we construct

$$\log u(I, t_2) \approx \log I - b(u(I, t_1), t_1) \cdot \Delta \log p_{t_2} - b(u(I, t_1), t_0) \cdot \Delta \log p_{t_1},$$

[8]In our computations, we use the trapezoid rule rather than the left Riemann sum in equation (4) to approximate the integral in (3) since it is a better numerical approximation.

[9]Invertibility at (u, s) means that we observe an income level I such that $u(I, s) = u$. When applying the algorithm to the data in the next section, the value of the budget share in period s corresponding to $u\,(I, t_{n-1})$ is obtained by linear interpolation of the grid of u at the time s.

and given $u(I, t_2)$, we construct Hicksian budget shares in t_2:

$$b(u, t_2) = B(u^{-1}(u, t_2), t_2).$$

That is, for each budget share $B(I, t_2)$ in t_2, we assign a utility value based on $u(I, t_2)$. We continue this iterative process until t_M. Note that we can only calculate $u(I, t)$ for those I's for which $B(u^{-1}(u(I, t), s), s)$ is observed for all $s < t$.

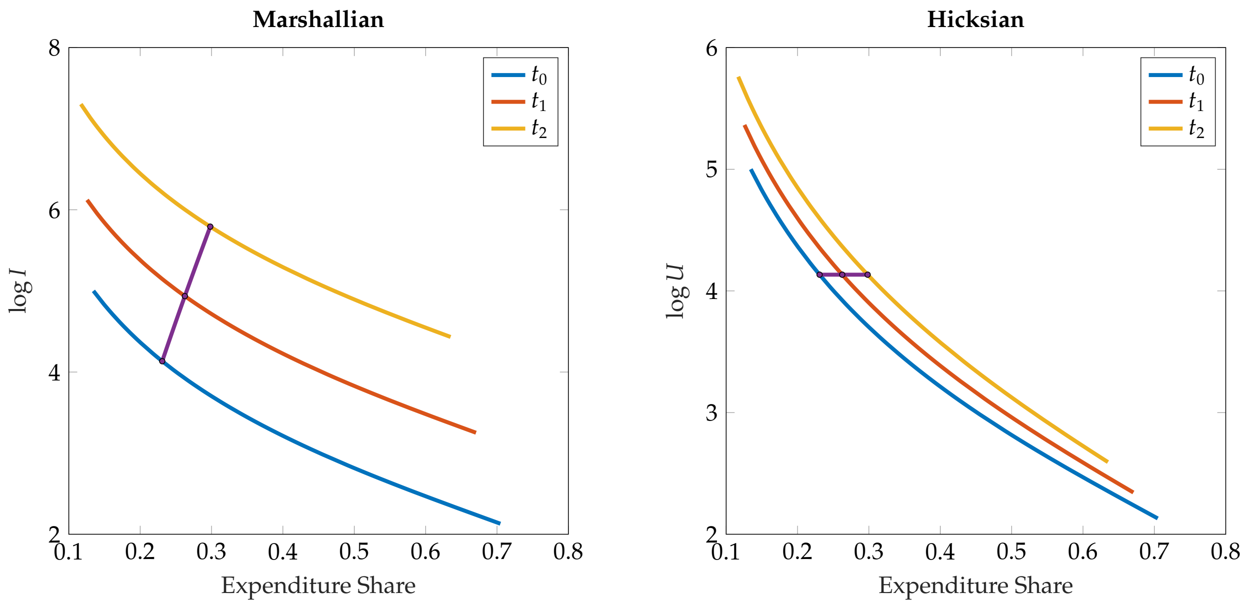

Figure 1: Non-homothetic preferences. Expenditure share for some good against log nominal income and log money-metric utility at three points in time.

To see this procedure graphically, consider the left panel of Figure 1 showing the expenditure share on some good against nominal income for three different points in time. The fact that the lines are downward sloping means that higher incomes are associated with lower expenditures on the good. In this example, incomes grow over time, so the range of nominal income levels shifts up over time.

In the data we observe budget shares as a function of income over time (Marshallian budget shares), but to construct the money-metric we require budget shares as a function of utility (Hicksian budget shares). The right panel of Figure 1 displays the Hicksian budget shares for the same good. The purple line in the right panel of Figure 1 shows for each period the Hicksian expenditure share for the good evaluated at some fixed utility level $\bar{u}$. The change in expenditures, holding utility constant, are pure substitution effects over time due to changes in relative prices. As implied by Proposition 1, multiplying the Hicksian budget shares by log price changes and summing over time gives the money-metric utility for the household with utility $\bar{u}$ at time t_2.

But, we cannot directly observe the figure on the right. How do we infer Hicksian budget shares? The purple line in the left panel of Figure 1 plots, for each period s, the income that gives the utility of $\bar{u}$, that is $u^{-1}(\bar{u}, s)$, and the associated budget share, $B_1(u^{-1}(\bar{u}, s), s)$. In other words, we can infer Hicksian budget shares for $\bar{u}$ by using the observed budget share along the purple line in the left panel. Then we can construct the mapping between income and utility at each point (the purple line) by iteratively applying the summation in (4).

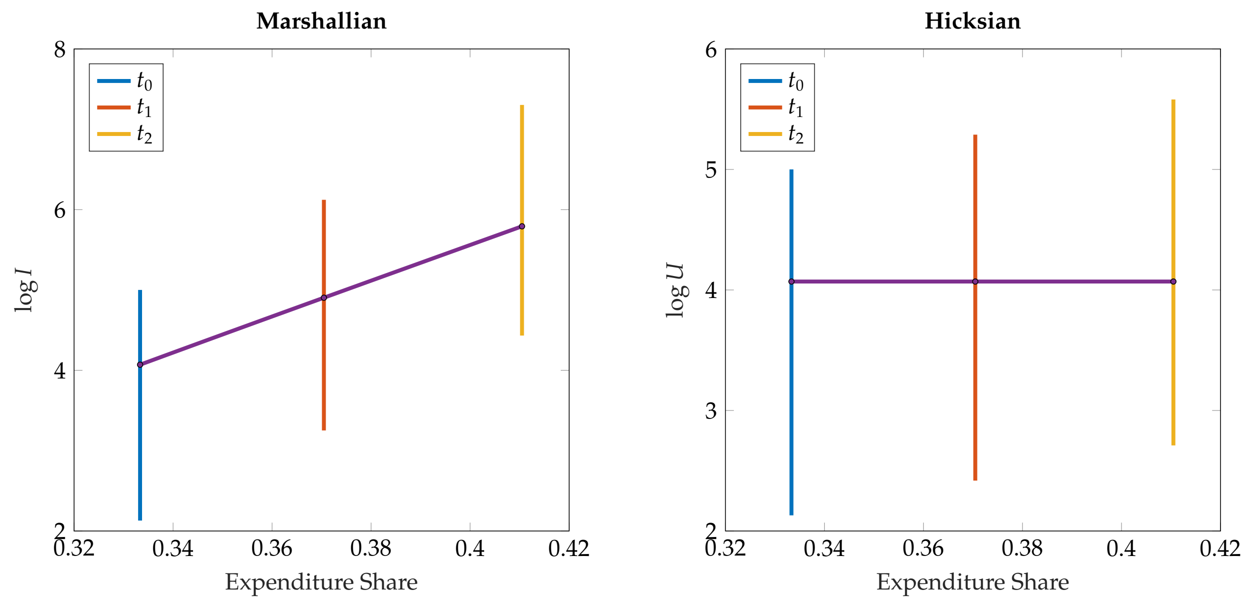

Figure 2: Homothetic preferences. Expenditure share for some good against log nominal income and log money-metric utility at three points in time.

To understand why Proposition 2 is unnecessary when preferences are homothetic, Figure 2 plots the same information as Figure 1 but for homothetic preferences. Since there are no income effects, budget shares at a point in time do not vary with household income or utility. That is, observed and Hicksian budget shares coincide. Therefore, we can construct the money-metric using a price index based on observed budget shares by good.

3.3 Extensions and Limitations

In practice, data is imperfect and noisy. Specifically, recorded expenditure shares can change through time for reasons other than changes in observed prices and income. Under some additional assumptions, our procedure can be modified to account for some of these issues.

For example, if there is classical measurement error or idiosyncratic taste shocks at the individual consumer level, uncorrelated with any observable, then we can eliminate this noise by averaging over multiple households with the same (or similar) income level. If the noise is caused by idiosyncratic taste shocks, then our money-metric utility function will apply to preferences in the absence of the taste shocks.

At the opposite extreme, suppose that there are persistent differences in preferences that are functions of observable characteristics, for example households with children have different preferences to those without.[10] In this case, we can handle this by splitting the sample in two and applying our method to each sample separately.[11]

If there are unobservable demand shifters that affect the entire distribution of households, then we cannot deal with that by averaging or conditioning on observable characteristics. This happens if there are aggregate taste shocks that affect the entire distribution of households, or if there are changes in quality over time that are not reflected in prices. If there are changes in quality, then our method can be applied to the quality-adjusted version of prices (following standard quality-adjustment practice) without issue. However, if there are unobservable shocks to preferences that are not idiosyncratic, and cannot be eliminated by averaging, then our methodology cannot be used. An example is if household preferences over time are systematically different to preferences in the past in ways we cannot model.

4 Illustrative Example Using Artificial Data

In this section, we illustrate and evaluate our algorithm using artificial data from fully parameterized preferences. We consider generalized non-homothetic CES preferences from Fally (2022).[12] The expenditure function is

$$e\left(p,u\right)=\left(\sum_{i}\omega_{i}\left(u^{\varepsilon_{i}}p_{i}\right)^{1-\sigma(u)}\right)^{\frac{1}{1-\sigma(u)}}.$$

[10] This assumption is related to the assumption considered in Section 2.3 of Jaravel and Lashkari (2022).

[11] Similarly, if we observe two groups of households that face different prices at a point in time (e.g. households living in different locations), then we can apply our method to each sample separately.

[12] See also Hanoch (1975), Comin et al. (2021), and Matsuyama (2019) for more information on these preferences. In Appendix B we consider another example with an Almost Ideal Demand System (AIDS).

By Shephard's lemma, Hicksian budget shares $b(p,u)$ are

$$b_i(p,u) = \frac{\omega_i \left(u^{\varepsilon_i} p_i\right)^{1-\sigma(u)}}{\sum_j \omega_j \left(u^{\varepsilon_j} p_j\right)^{1-\sigma(u)}},$$

and Marshallian budget shares are $B(I,t) = b(p_t,u)$ where u solves $I = e(p_t,u)$. Income elasticities can vary across goods and the elasticity of substitution σ can vary across indifference curves (but is constant along any indifference curve, as under standard CES). As shown by Baqaee and Burstein (2021), the money-metric function for t_0 reference prices is

$$m(p_{t_0}, p, I) = \left(\sum_i \omega_i \left(u^{\varepsilon_i} p_{i,t_0}\right)^{1-\sigma(u)}\right)^{\frac{1}{1-\sigma(u)}}$$

where u is the solution to $I = e(p,u)$. To evaluate the accuracy of our algorithm, we compare this closed-form expression for $m(p_{t_0}, p_T, I)$ with the results of our numerical procedure applied to artificial data generated using these preferences.

We generate repeated cross-sectional data on income and expenditure shares over 3 goods for 100 households that face a common price vector for $T = 40$ periods. The distribution of income in the first period is lognormal (parameterized to match the distribution of household expenditures in the 1974 UK household survey, described in the next section). All incomes grow by a factor of ten over the sample period at a constant annual growth rate. Good 1 has the lowest income elasticity (and hence we refer to it as a necessity) and the highest inflation rate. Figure 3 displays the paths of household income and price data in our illustrative example.

We consider three parameterizations. The first one is the homothetic case with $\varepsilon_i = \varepsilon$ and $\sigma = 0.25$. The second case allows for income effects but the elasticity of substitution is independent of u. We follow Comin et al. (2021), and set $\varepsilon_1 = 0.2$, $\varepsilon_2 = 1$, $\varepsilon_3 = 1.65$, and $\sigma = 0.25$. The third case further assumes that the elasticity of substitution is a log-linear decreasing function of u, consistent with estimates in Auer et al. (2021). We set $\sigma(u) = 10 - 2\log u$, with the intercept value ensuring that elasticities of substitution remain higher than unity. The share parameter ω is calibrated separately in each case so that the budget shares of each good for the median household in the first period are all the same (equal to one third for each good).

We define *real income* to be income deflated by a chain-weighted price index based on observed aggregate budget shares (as is standard practice in national income accounting). Define the non-homotheticity bias to be the log difference between real income and the

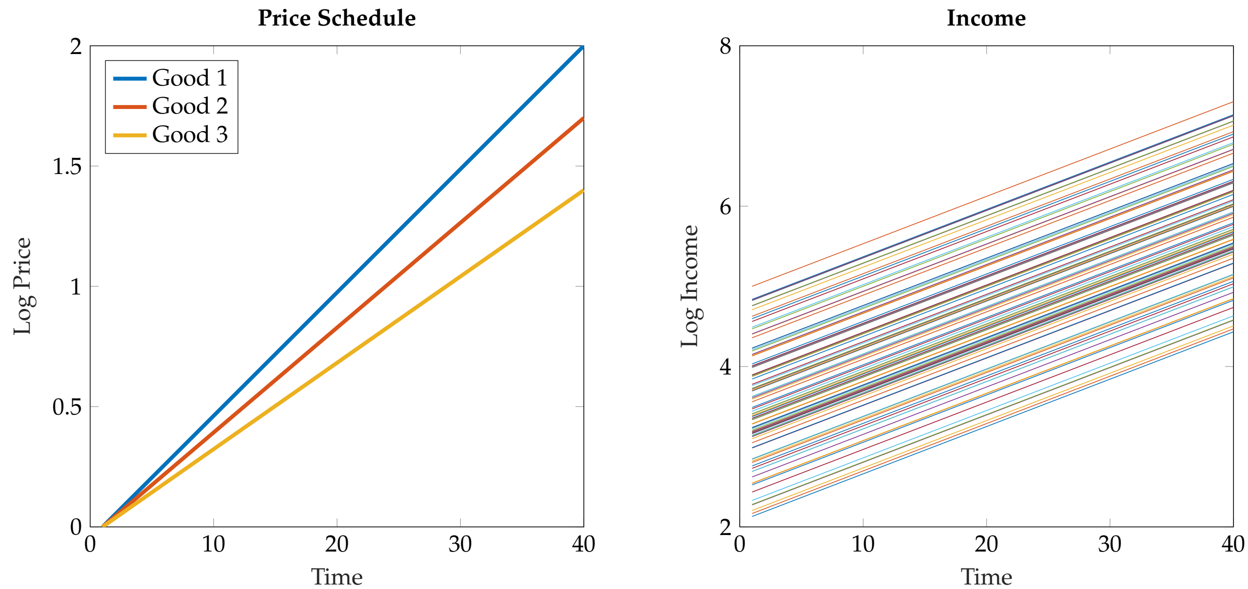

Figure 3: Exogenous Price and Income Path.

money-metric. Figure 4 plots the gap between real income and the money-metric in our examples under fixed and variable elasticities of substitution. We do not plot the homothetic case since the bias is zero. In both cases, since inflation is higher for necessities, and poor households are more reliant on necessities, the bias is larger for poorer households. We do not report the bias for households above the 70th percentile because for those values of income, there do not exist similarly well-off households in the past whose demand can be used (i.e. the invertibility condition fails for those households).

To assess the accuracy of our procedure, we use the infinity norm — that is, the maximum absolute value of the log difference between the true money-metric function and our estimate at time T. Under both parameterizations, constant and variable σ, the error is very small: 0.0078 and 0.0044. This is equivalent to roughly two thirds of 1% of income.[13] Figure 5 shows how this error varies as we vary the number of households and the frequency of observations using the variable σ non-homothetic specification as an example. As expected, the error converges to zero as we approach the continuous-time limit. The error also falls as the number of households in the sample increases.

[13] As mentioned earlier, if we solve the fixed point problem described in Footnote 7, then the error is three orders of magnitude smaller. We prefer to use the iterative procedure because it is more intuitive to describe and makes little difference when we apply it to real data.

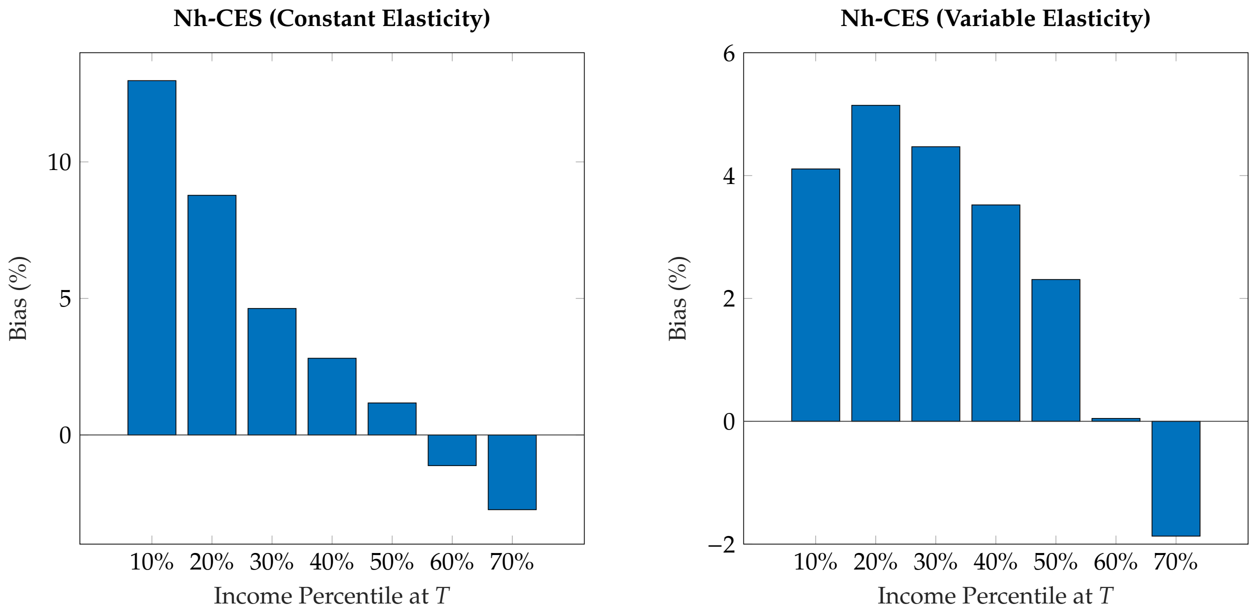

Figure 4: Non-homotheticity bias: Log difference between real income and the money-metric

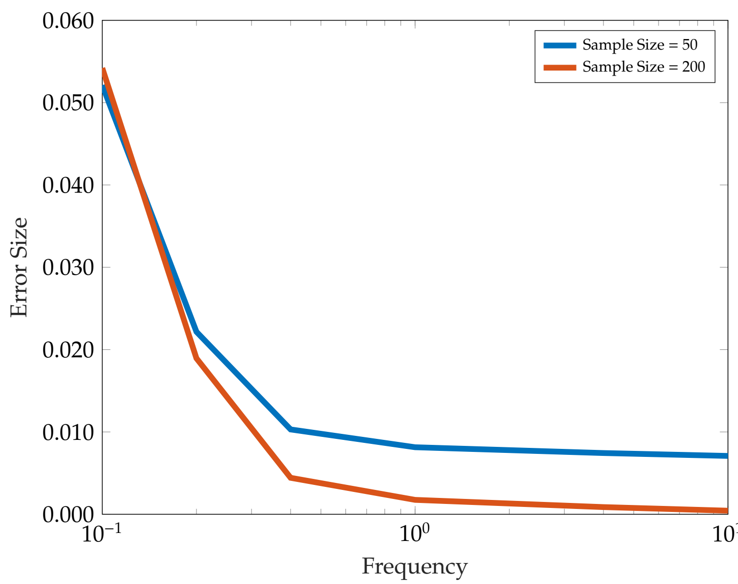

Figure 5: This figure displays the maximum error as a function of the frequency of observation, holding the path of price and income changes in Figure 3 constant. Our baseline calibration is annual frequency corresponding to a value of $10^0 = 1$ observations per year on the x-axis. If we observe the data once every decade, then the frequency is 1/10, and if we observe the data every month, then the frequency is 12. The blue line is the case where only 50 households can be observed, while the red line corresponds to a case where 200 households can be observed.

5 Empirical Results

In this section, we apply our algorithm to long-run household cross-section data. Our goal is to compare changes in welfare as measured by the money-metric, with changes in real income as measured by income deflated by a chain-weighted inflation index. For this purpose, we use the *Family Expenditure Survey and Living Costs and Food Survey Derived Variables* for the UK (see Oldfield et al., 2020), which is a repeated cross-section of UK household expenditures over different sub-categories of goods and services from 1974 to 2017.[14]

Following the practice of the Office of National Statistics, we use the retail price index (RPI) in the period 1974-1998 and the consumer price index (CPI) in the period 1998-2017. To concord the RPI, CPI, and household expenditure data, we assemble nine aggregate product categories that can be used consistently over the entire period of analysis. See Appendix A for additional details. Between 1974 and 2017 prices rose relatively less for product categories such as leisure goods and services, that are disproportionately consumed by richer households and experienced a rise in expenditure shares over time.

5.1 Mapping Data to the Model

Our procedure requires the income I and the budget shares $B(I,t)$ at time t across all goods. To deal with the measurement error and idiosyncratic noise, we fit a smooth curve for each good at each time point t and use these curves as $B(I,t)$. That is, we estimate the true $B_i(I,t)$ function for some good i by fitting the following curve for each t using ordinary least squares

$$B_{iht} = \alpha_{it} + \beta_{it} \log I_{ht} + \gamma_{it} \left(\log I_{ht}\right)^2 + \varepsilon_{iht},$$

where i is the good, h is the household, and t is the time period. The estimated regression line gives us $B(I,t)$.[15]

We apply our procedure sequentially from 1974 to 2017 using the UK cross-sectional data constructed in the manner described above. Computing $u(I,t) \equiv m(p_{t_0}, p_t, I)$ requires that for each time $s < t$, we can estimate the Hicksian budget share $b(u(I,t),s)$. That is, for each income level I at time t, we must be able to find consumers at $s < t$ whose utility values were the same as that delivered by I at time t. The left panel of Figure

[14] Aggregate nominal consumption growth in our sample is lower than that in the UK national accounts. According to the UK Office for National Statistics, this difference is due to differences in sample coverage. While these sample coverage issues affect aggregate nominal growth rates, they do not affect our results, which are at the household-level.

[15] We also estimated $B_i(I,t)$ using locally weighted scatterplot smoothing (LOWESS) and obtained very similar results.

6 illustrates how households in 2017 are matched with households in 1974 in order to estimate $b(u(I, 2017), 1974)$. For example, households in the 50th percentile of income in 2017 are matched with households in the 77th percentile of income in 1974.

Our algorithm naturally implies that we can only compute $u(I, t)$ if $u(I, t)$ is less than the upper bound and more than the lower bound of utility levels at all past times $s < t$. Otherwise, we cannot carry out the inversion in (5). The right panel of Figure 6 plots the distribution of log expenditures in our data and the solid lines show the sample of households for which we can calculate $u(I, t)$. Our algorithm can recover the money-metric up to about the 97th percentile of households in 2017. For the richest households, we are unable to compute $u(I, t)$ because there are no households in our sample that were on the same indifference curve in the past. Nevertheless, our algorithm covers a significant range of households. Our sample coverage is high because the distribution of spending is highly fat-tailed, which means that in 1974, there are households who are on the same indifference curve as the richest 97th percentile of households in 2017.

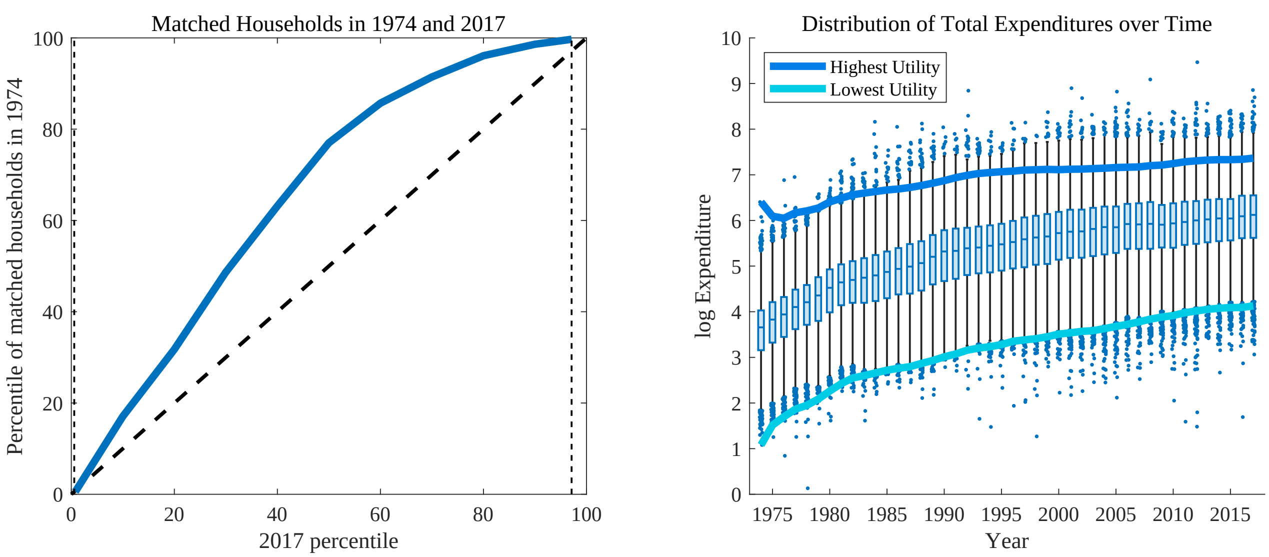

Figure 6: The figure on the left shows, for each income percentile in 2017, the income percentile in 1974 of the matched household that is on the same indifference curve as the 2017 household. The figure on the right shows the sample distribution of (weekly) log expenditures from 1974 to 2017. The upper and lower blue boxes represent the 75th and 25th percentiles, respectively. The solid lines indicate the upper and lower bounds of the sample for whom the Hicksian budget share can be computed as a function of time. The lower and upper bounds in 2017 represent the 0.8th and 97th percentile, respectively, of the income distribution (vertical lines in left panel).

5.2 Empirical Results

The blue line in Figure 7 plots the expenditure function $e(p_{1974}, v(p_{2017}, I))$ for different values of income. This expresses different incomes in 2017 in terms of 1974 pounds. For comparison, the red line shows the equivalent if all households faced the same effective inflation rate, as given by the Tornqvist chain-weighted aggregate inflation rate. When the red line is above the blue line, this means that real income based on chain-weighted aggregate inflation is higher than equivalent income using the money-metric for households in the sample, and the size of the bias is largest for the poorest households. That is, the poorest households are not as well off as what is implied by relying on the official statistics. Conversely, the gap reverses around the 60th percentile of income and then widens suggesting that the richest households are better off in 1974 pounds than what is implied by official statistics.

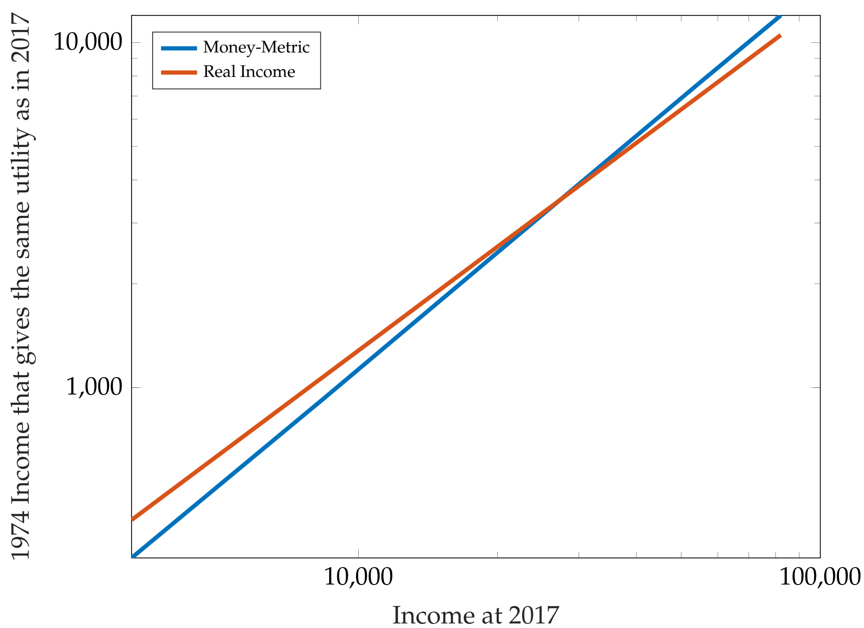

Figure 7: Money-metric $m(p_{1974}, p_{2017}, I_{2017})$ and real income using aggregate chain-weighted inflation between 1974 to 2017 (annualized pounds, log scale).

Figure 8 displays the log difference between the red and blue lines in Figure 7. As expected, the difference is positive, meaning that real income as calculated in the official statistics is upward biased for poor households and downward biased for rich households. The size of the bias is around 25% for the poorest households. This means that over the 43 year sample, official annual inflation rates understate the welfare-relevant inflation for these households by around 0.5 percentage points per year. On the other hand, for the richest households, the official inflation rate is overstated by around 0.25 percentage points per year on average. This implies that inequality across households is larger based on the money metric than based on real income, as revealed by the histograms in Figure 9.

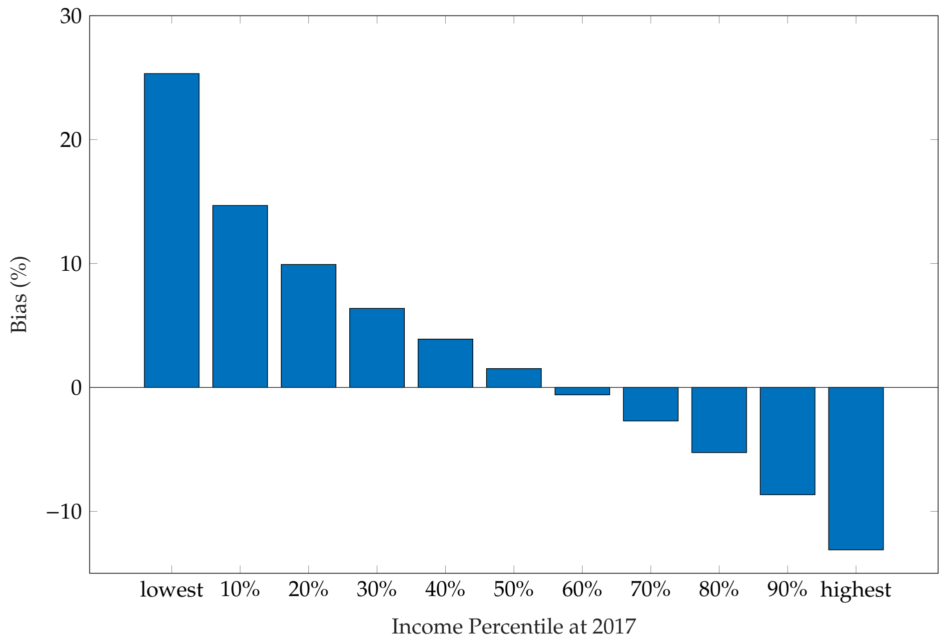

Figure 8: Log difference between real income and money-metric in 1974 for different percentiles of the income distribution in 2017. Highest and lowest correspond to the utilities and their percentiles in Figure 6.

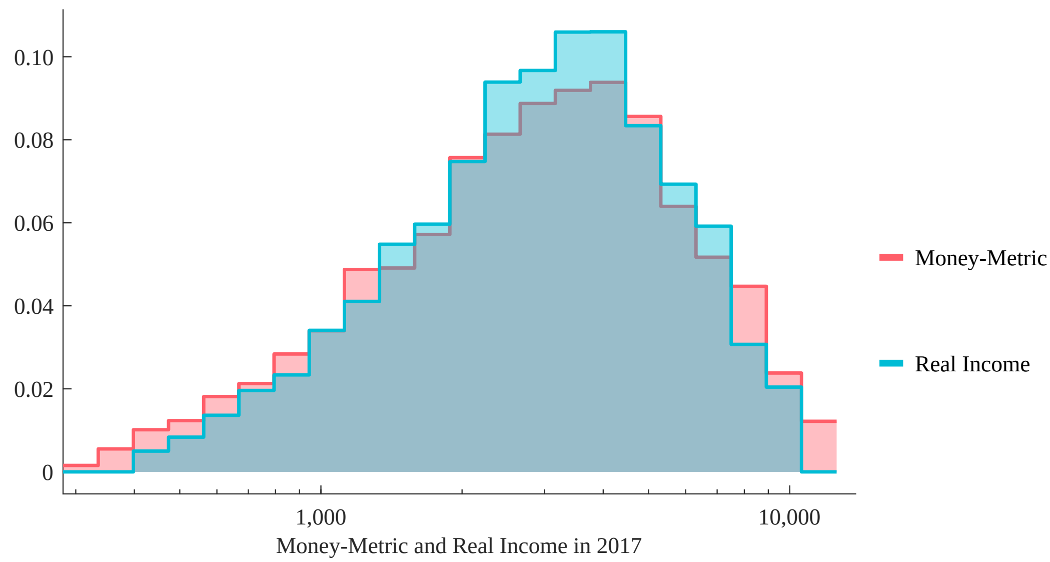

Figure 9: Histogram (using household weights) of money-metric $m(p_{1974}, p_{2017}, I_{2017})$ and real income using aggregate chain-weighted inflation (annualized pounds, log scale). The distributions are truncated at the upper and lower bounds of Figure 6.

Intuitively, for a given relatively poor household in 2017, consumers with the same utility level who lived in earlier years spent on average relatively more on sectors with higher inflation rates than consumers as a whole. Therefore, the inflation rate for these consumers is higher than the aggregate inflation rate.

6 Conclusion

In this paper, we provide a simple and intuitive procedure for constructing money-metric representations of utility using repeated cross-sectional data. Our insight is that one can trace out Hicksian, or compensated, demand curves through time by matching households whose incomes are equivalent in utility terms. Although our approach is non-parametric, it relies on the assumption that preferences are the same in both the cross-section and the time-series dimensions and that all consumers face common prices. Relaxing these assumptions is an interesting avenue for future work.

References

Atkin, D., B. Faber, T. Fally, and M. Gonzalez-Navarro (2020, March). Measuring Welfare and Inequality with Incomplete Price Information. NBER Working Papers 26890, National Bureau of Economic Research, Inc.

Auer, R., A. Burstein, S. Lein, and J. Vogel (2021). Unequal expenditure switching: Evidence from switzerland.

Baqaee, D. and A. Burstein (2021). Welfare and output with income effects and taste shocks. Technical report, National Bureau of Economic Research.

Comin, D., D. Lashkari, and M. Mestieri (2021). Structural change with long-run income and price effects. *Econometrica 89*(1), 311–374.

Fally, T. (2022). Generalized separability and integrability: Consumer demand with a price aggregator. *Journal of Economic Theory 203*, 105471.

Hamilton, B. W. (2001, June). Using engel's law to estimate cpi bias. *American Economic Review 91*(3), 619–630.

Hanoch, G. (1975). Production and demand models with direct or indirect implicit additivity. *Econometrica: Journal of the Econometric Society*, 395–419.

Jaravel, X. and D. Lashkari (2022). Nonparametric measurement of long-run growth in consumer welfare.

Konüs, A. A. (1939). The problem of the true index of the cost of living. *Econometrica: Journal of the Econometric Society*, 10–29.

Matsuyama, K. (2019). Engel's law in the global economy: Demand-induced patterns of structural change, innovation, and trade. *Econometrica 87*(2), 497–528.

Oldfield, Z., J. Banks, P. Levell, and A. Leicester (2020). Family expenditure survey and living costs and food survey derived variables, 1968-2017.[data collection]. *UK Data Service*.

Samuelson, P. A. and S. Swamy (1974). Invariant economic index numbers and canonical duality: survey and synthesis. *The American Economic Review 64*(4), 566–593.

Online Appendix

A Additional details of the UK data used in Section 5

The *Family Expenditure Survey and Living Costs and Food Survey Derived Variables.* is an annual survey that records information on about 5,000-7,000 households. It is a continuous cross-sectional survey on expenditures and demographics, combining the annual *Living Costs and Food Survey* (and its predecessors) in a way that allows for time-series comparisons. The main dataset we use is the RPI category expenditure basket, which provides the most continuous data from 1974 to 2017. Starting in 1995, the data are split into separate files for adults and children, so we merge them.

Our algorithm does not require information on the entire distribution of households, and can recover the money-metric for a subsample of observed households, even if that subsample does not sample incomes at the same frequency as the population. The expenditure survey is sampled on the entire UK, except for top earners and some pensioners. In order to correct for possible nonresponse bias, household weights are provided since 1997.[16] We use these weights to calculate the chained aggregate price index, which we use to calculate real income as in the official statistics. However, our approach does not require household weights.[17]

To construct the consumption deflator in the national accounts, the Office of National Statistics switched from the Retail Price Index (RPI) to the Consumer Price Index (CPI).[18] By comparing the RPI and CPI with the consumption deflator provided by the ONS, we identify the switching point as 1998 and do the same for our price data.

Because the CPI and RPI consider different baskets of goods and services, we merged various sub-categories to obtain a consistent set of categories over time. For example, "alcohol" in the RPI includes some items served outdoors, which is included in "restaurants" in the CPI. In this case, we merged "Catering and Alcohol" in the RPI and matched it with "Restaurant and Alcohol" in the CPI. We end up with nine categories that are available for the entire period for both RPI and CPI. Table 1 summarizes how we integrated the CPI

[16]Prior to 1997, benefit unit weights are provided instead of household weights. Since a benefit unit is a single person or a couple with any dependent children, there can be more than one benefit unit weight in a household. For example, if a couple with their children and the father's parents live together, then two benefit unit weights are recorded. In this case, we use the simple average as the household weight.

[17]We also use weights to calculate the percentiles in the left panel of Figure 6, the histograms in Figure 9, and the quintiles for Figure 11 in the Appendix.

[18]`https://webarchive.nationalarchives.gov.uk/ukgwa/20151014001957mp_/http://www.ons.gov.uk/ons/guide-method/user-guidance/prices/cpi-and-rpi/mini-triennial-review-of-the-consumer-prices-index-and-retail-prices-index.pdf`.

and RPI baskets.

Integrated Categories	RPI	CPI
Food	Total food	Food and non-alcoholic beverages
Tobacco	Cigarettes & tobacco	Tobacco
Clothing	Clothing & footwear	Clothing & footwear
Household&Fuel	Housing except morgage interest	Housing, water and fuels
	Fuel & light	
	(-)dwelling insurance & ground rent	
Household Goods	Household goods	Furniture and household equipment & routine repair of house
Personal Goods & Service	Personal goods & services	Health
	Household services	Miscellaneous goods and service
	dwelling insurance & ground rent	-
Transport	Motoring expenditure	Transport
	Fares & other travel costs	
Leisure Goods & Service	Leisure goods	Communication
	Leisure services	Recreation & culture
	-	Education
	-	Accomodation service
Catering	Catering	Catering services
	Alcoholic drink	Alcoholic beverage

Table 1: RPI and CPI Correspondence Table

B Using artificial data from Almost Ideal Demand System

In this appendix, we redo our analysis of Section 4 using another popular form of non-homothetic preferences: Almost Ideal Demand System (AIDS). The expenditure function is

$$e(p,u) = c(p)\, u^{d(p)}$$

where $c(p)$ and $d(p)$ are given by:

$$c(p) = \exp\left(a_0 + \sum_{i=1}^{I} a_i \log p_i + \frac{1}{2}\sum_{i=1}^{I}\sum_{j=1}^{I} \gamma_{ij} \log p_i \log p_j\right)$$

$$d(p) = \exp\left(\sum_{i=1}^{I} \beta_i \log p_i\right)$$

where $\sum \alpha_i = 1$, $\sum \beta_i = \sum \gamma_{ij} = 0$ and $\gamma_{ij} = \gamma_{ji}$ for all i and j.

By Shephard's lemma, Hicksian budget shares $b(p,u)$ are

$$b_i(p,u) = \alpha_i + \sum_{j=1}^{I} \gamma_{ij} \log p_j + \beta_i d(p) \log u.$$

The money-metric function for t_0 reference prices is[19]

$$m(p_0, p, I) = c(p_0) \left(\frac{I}{c(p)} \right)^{\frac{d(p_0)}{d(p)}}.$$

In assigning parameter values, we assume that the expenditure share is decreasing in utility for good 1 and increasing for good 3, as in the non-homothetic CES example in Section 4. Specifically, we consider the following parameter values, that also ensure that the expenditure share on all goods is positive in all periods in the artificial dataset.

α_0	α_1	α_2	α_3	β_1	β_2	β_3	γ_{11}	γ_{22}	γ_{33}	γ_{12}	γ_{13}	γ_{23}
2	1/3	1/3	1/3	−0.15	−0.05	0.2	−1/4	−1/4	−1/4	1/8	1/8	1/8

Table 2: Parameters for AIDS

Figure 10 presents the non-homotheticity bias. The approximation error of the money-metric in period T, $\max_h \left|\log u(I,T) - \log u(I,T)^{TRUE}\right|$, is 0.0011.

C Comparison with Jaravel & Lashkari (2022)

In this appendix, we apply the first-order and second-order algorithms described in Jaravel and Lashkari (2022) on our artificial and real data. By setting the base year in the Jaravel and Lashkari (2022) algorithm to t_0, their definition of real consumption matches our money-metric. For brevity, we do not include in this appendix a description of these algorithms.

C.1 Results with Artificial Data

We first calculate the approximation error when applying the Jaravel and Lashkari (2022) algorithms to the artificial data that we use in Section 4. Table 3 shows that the error is very low for non-homothetic CES (with constant elasticity of substitution) and AIDS. On

[19] To obtain the expression for the money-metric, we use $m(p_0, p, I) = e(p_0, u)$, where $I = c(p)\, u^{d(p)}$.

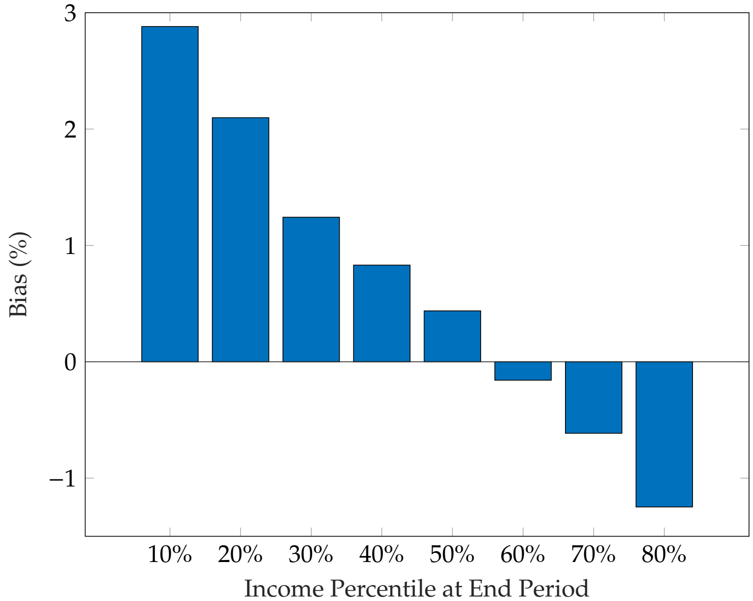

Figure 10: Log difference between real income and the money-metric

the other hand, the approximation error is larger for non-homothetic CES with variable elasticity of substitution.

	Nh-CES(Constant)	Nh-CES(Variable)	AIDS
First Order	0.0491	0.1422	0.0401
Second Order	0.0028	0.1130	6.60×10^{-6}

Table 3: $\max_h \left|\log u(I,T) - \log u(I,T)^{TRUE}\right|$: Results of the first/second order algorithm of Jaravel and Lashkari (2022) with $K = 2$ to the artificial data in Section 4.

C.2 Results with UK Household Data

We next apply the Jaravel and Lashkari (2022) algorithms to the UK household data. In the main application in Jaravel and Lashkari (2022), the algorithms are applied to households in the US CEX by quintile group. Following this, we first apply their algorithms to the quintile data; results are displayed in Figure 11. We next apply their algorithms to the underlying disaggregated data that we use in our empirical results; results are displayed in Figure 12.

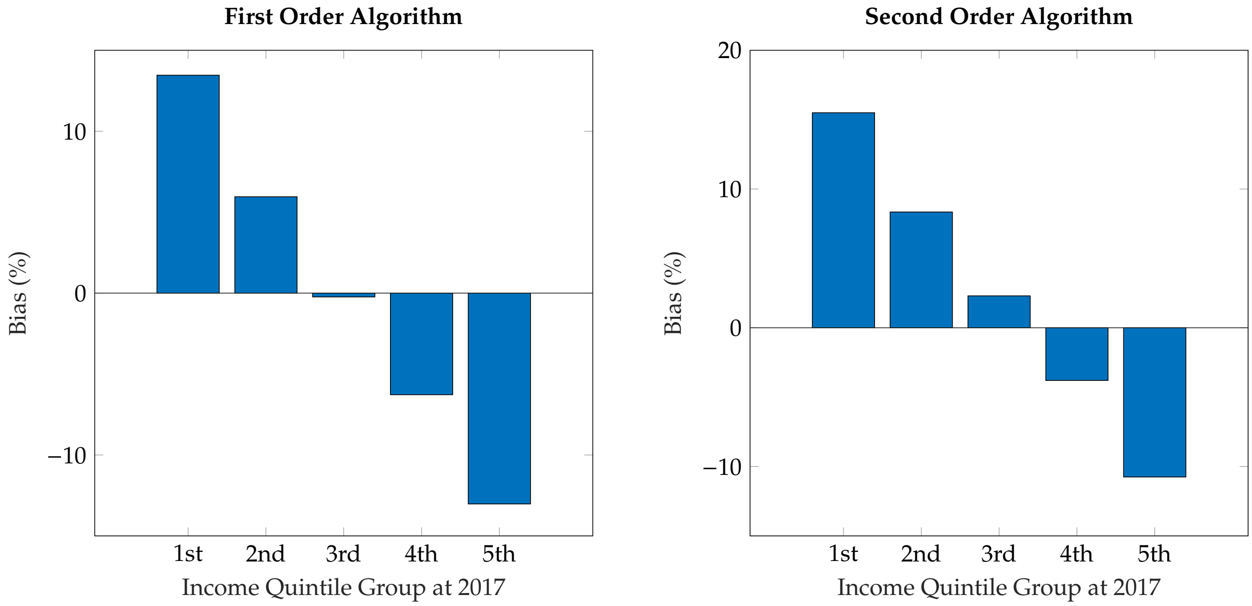

Figure 11: Results of the first/second order algorithm of Jaravel and Lashkari (2022) with $K = 1$ to the aggregated (by quintile) UK household data: Log difference between real income and the money-metric

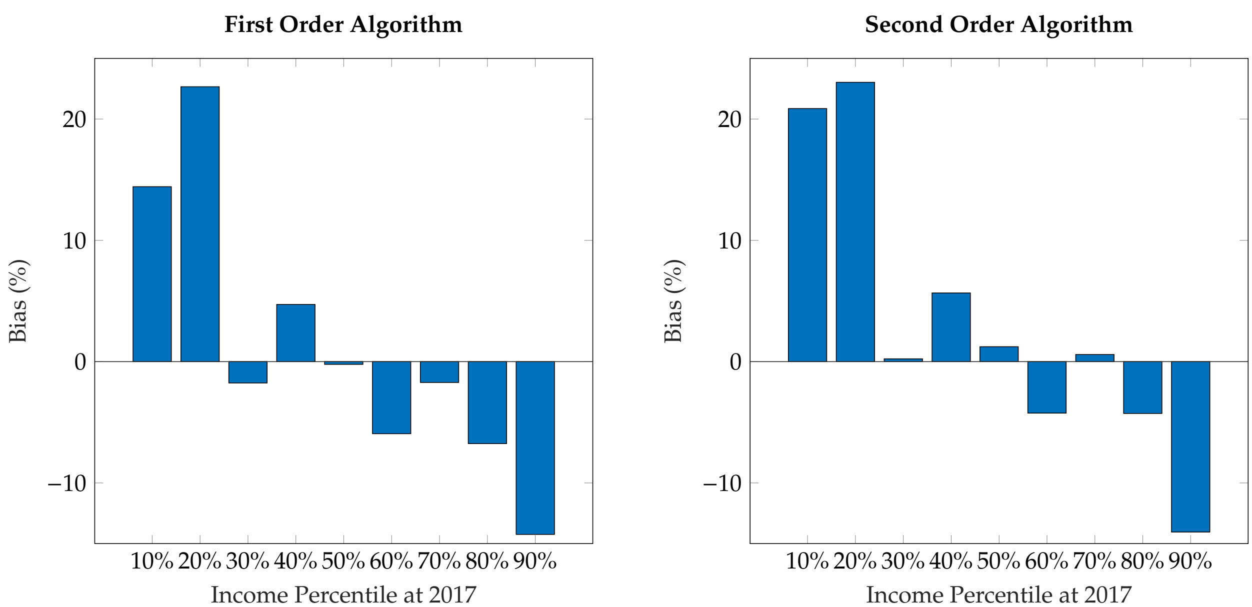

Figure 12: Results of the first/second order algorithm of Jaravel and Lashkari (2022) with $K = 1$ to the disaggregated UK household data: Log difference between real income and the money-metric