

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, Revised November 2022

We thank John Asker, Natalie Bau, Dave Donaldson, Ezra Oberfield, 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.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by David Baqaee, Ariel Burstein, and Yasutaka Koike-Mori. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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, Revised November 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 rules out systematic taste shocks in both the cross-section and the time-series. We use a different theoretical insight to study this problem than previous studies by Blundell et al. (2003) and Jaravel and Lashkari (2022). Our idea is to trace Hicksian (or compensated) demand through time by matching households on the same indifference curve at different points in time. Given Hicksian demand, we can construct money-metric utility and cost-of-living indices 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

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

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

1 Introduction

Welfare is typically measured using money-metric utility, which converts income under one set of prices into equivalent income under a different set of prices. Money-metric utility can be constructed by integrating Hicksian (compensated) demand with respect to changes in prices (see e.g. Hausman, 1981). To do this, we need to know how households would spend their income under different prices if they were being compensated for those price changes. When preferences are homothetic, compensated expenditure shares coincide with uncompensated expenditure shares. This theoretically justifies standard chained price deflators, like the CPI or the PCE, which weigh price changes in each period using expenditure shares observed in that period. Outside of the homothetic special case, compensated and uncompensated budget shares do not coincide, and chaining is invalid (e.g., Samuelson and Swamy, 1974). Constructing money-metric utility for non-homothetic preferences is difficult because Hicksian budget shares are not easily obtained.

We resolve this problem as follows. Suppose we observe repeated cross-sections of households with identical preferences facing common prices.¹ To construct the money-metric utility in 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 .² Hence, to calculate the money-metric utility function, we need to match households and trace Hicksian demand through time.³ To match households through time, we need to know the money-metric utility function, since two households are on the same indifference curves if their money-metric utility values coincide. Our insight is that this is a fixed point problem, and we propose a simple iterative algorithm to find this fixed point.

Our approach generalizes standard practice by statistical agencies, who weigh changes in prices over time using aggregate budget shares. If preferences are homothetic and stable, these aggregate budget shares coincide with the compensated budget share of every individual in the sample. This is the assumption that justifies conventional price deflators. However, when preferences are non-homothetic, the compensated budget

¹In 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.

²Our approach is not based on interpersonal comparisons of utility. Instead, we match households that have the same Hicksian demand.

³Hicksian demand can also be calculated given knowledge of elasticities of substitution. The procedure in this paper does not require and is not equivalent to estimating elasticities of substitution non-parametrically. Intuitively, we only recover Hicksian demand evaluated at observed prices, whereas the elasticities of substitution allow one to compute counterfactual Hicksian demand even for prices that are not observed.

shares of each household are different. In this case, instead of using aggregate budget shares to construct the price index, we need to use for each income today the budget shares of some unique corresponding income level in the past. Our paper provides a way to construct this mapping.⁴

Our paper is closely related to Blundell et al. (2003) and Jaravel and Lashkari (2022) who also develop non-parametric approaches to measuring welfare for non-homothetic preferences using cross-sectional household-level data. Similar to both of these papers, and conventional practice in index number theory, 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.⁵ Although inspired by them, our approach, which integrates Hicksian demand curves, is different to theirs. We discuss these papers in turn.⁶

Blundell et al. (2003) bound the money-metric by using revealed choice arguments. For each income level at time t , Blundell et al. (2003) show how to construct a bundle that is strictly better and a bundle that is strictly worse in time $s \neq t$. The price of these two bundles then bound the true money-metric value. We exposit and implement an amended version of their methodology in Appendix C. We amend their algorithm for the lower-bound, since the version in their paper appears to have an error. Relative to Blundell et al. (2003), an advantage of our approach is that it provides a point estimate, rather than bounds, for the money-metric utility as long as the data is smooth and observed continuously. In the same appendix, we show that our estimates are within their bounds for both artificial and real-world data.

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. Relative to Jaravel and Lashkari (2022), our approach does 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

⁴Whereas our procedure can be used to measure changes in welfare for observed changes in prices and income, it cannot be used to answer counterfactual or macroeconomic welfare questions like those studied by Baqaee and Burstein (2021).

⁵These papers and ours also require that all prices are observed. For an alternative approach to measuring changes in welfare allowing for incomplete information on prices but imposing additional restrictions on preferences, see Hamilton (2001) and Atkin et al. (2020).

⁶In Appendices C and D we apply the Blundell et al. (2003) and Jaravel and Lashkari (2022) approaches to our artificial data and the UK data. We find that the point estimates of the money metric using our algorithm are always within the Blundell et al. (2003) bounds. Moreover, we show that the Jaravel and Lashkari (2022) algorithm has different small sample properties than our algorithm, giving larger approximation errors when applied to our artificial data examples and different point estimates when applied to our UK data.

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 if economic growth shifts the support of utilities over time.

Our approach can also be contrasted with more parametric approaches where welfare measures are computed using a fully-specified and estimated demand system (e.g. Deaton and Muellbauer 1980). Specific functional forms for non-homothetic preferences are used to understand phenomena as diverse as structural transformation (e.g. Boppart 2014, Comin et al. 2021, and Fan et al. 2022), international trade patterns (e.g. Matsuyama 2000, and Fajgelbaum et al. 2011), and savings behavior and inequality (e.g. Straub 2019). Our approach provides a non-parametric way to recover welfare measures from the data without imposing and estimating strong functional form assumptions.

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 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).⁷ 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 rational preference relation \succeq 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 and Cost-of-Living). Holding fixed some reference vector of prices \bar{p} , the *money-metric* function maps (p, I) to

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

Holding fixed some reference budget set (\bar{p}, \bar{I}) , the *cost-of-living* index maps p to

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

The money-metric $e(\bar{p}, v(\cdot))$ is itself an indirect utility function because a budget set (p, I) is preferred to another budget set (p', I') if, and only if, $e(\bar{p}, v(p, I)) > e(\bar{p}, v(p', I'))$. We use this money-metric cardinalization of utility throughout the rest of the paper. Whereas

⁷These results are consistent with Blundell et al. (2003), which report a relatively greater rise on the cost of living for poorer households between 1975 and 1984 in the UK.

the money-metric converts the value of different budget sets into equivalent value under some baseline prices, the cost-of-living function, $e(\cdot, v(\bar{p}, \bar{I}))$, converts the value of some baseline budget constraint (\bar{p}, \bar{I}) into equivalent income under different sets of prices.⁸

To summarize, the function $e(p', v(p, I))$, mapping (p', p, I) into a scalar, is an object of paramount interest. 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 ranking budget sets (i.e. measuring growth). 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 (i.e. measuring inflation).

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*

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

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

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

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.

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

⁸In index number theory, the cost-of-living index is also called the Konüs (1939) index.

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 $p_t \in \mathbb{R}^N$ at each point in time $t \in [t_0, T]$.⁹ We also observe vectors of expenditure shares $B(I, t)$ across all goods demanded by consumers with preferences \succeq and income levels $I \in [I_t, \bar{I}_t]$ at time t . 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

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

Denote the money-metric utility function using base prices p_{t_0} and evaluated at budget set (p_t, I_t) by $u(I, t)$. Using this cardinalization of indirect utility, we can write

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

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 $u(I, t) \equiv e(p_{t_0}, v(p_t, I))$ is a fixed point of the following integral equation*

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

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 (time) is equal to s . That is $u^{-1}(\cdot, s)$ is a level of nominal income I^* in s such that $u(I^*, s) = u(I, t)$.

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

⁹Prices of sectoral aggregates tend to be smooth over time. This is the level of aggregation typically considered in work that documents non-homotheticities (including our application in Section 5). If the underlying prices within these aggregates contain jumps (due to, e.g. temporary sales or entry and exit of goods) then this needs to be taken into account when constructing these sectoral price aggregates. Doing so is beyond the scope of this paper.

solving it. We outline one such procedure below. But first, we discuss some implications of Proposition 2.

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, the solution to this integral equation is the money-metric, $u(I, t) = e(p_{t_0}, v(p_t, I))$, for (p_t, I) in terms of p_{t_0} prices. By varying initial prices p_{t_0} , for fixed (p_t, I) , we can also recover the cost-of-living index.

Proposition 2 follows immediately from Proposition 1 once we recognize that in the integral equation above, $B_i(u^{-1}(\cdot, s), s) : \mathbb{R}_+ \rightarrow [0, 1]$ 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 .

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.¹⁰

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

$$\log u(I, t_n) \approx \log I - \sum_{s=0}^{n-1} B(I_s^*, t_s) \cdot \Delta \log p_{t_s}, \quad (4)$$

¹⁰The 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 $u(I, t_{n-1})$ on the right hand side of equation (5) with $u(I, t_n)$. 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. Since the iterative procedure already has quite small errors, we do not show these alternative results.

where I_s^* satisfies

$$u(I_s^*, t_s) = u(I, t_{n-1}), \quad (5)$$

with the boundary condition $u(I, t_0) = I$. The summation in (4) above approximates the integral in (3) using a Riemann sum and becomes exact in the continuous-time limit because the Riemann sum becomes an integral and $u(I, t_{n-1}) \rightarrow u(I, t_n)$.¹¹ 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. Start with the boundary condition $u(I, t_0) = I$ since t_0 -equivalent income at t_0 is just initial income. For period t_1 , compute

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

where the last equation uses the boundary condition, which implies $b(u(I, t_0), t_0) = B(I, t_0)$. For values of I outside of $[L_0, \bar{I}_0]$, we cannot compute $u(I, t_1)$.

With $u(I, t_1)$ in hand, construct Hicksian budget shares for period t_1 :

$$b(u(I, t_1), t_1) = B(I, t_1).$$

That is, to each budget share $B(I, t_1)$, assign a utility value based on $u(I, t_1)$. Intuitively, we know budget shares as a function of income at t_1 , and we know utility as a function of income at t_1 . Since utility is monotone in income, this means that we can associate with each $B(I, t_1)$ a utility value, which is precisely the Hicksian budget share. We now have Hicksian budget shares $b(u, t_0)$ and $b(u, t_1)$.

Next, calculate

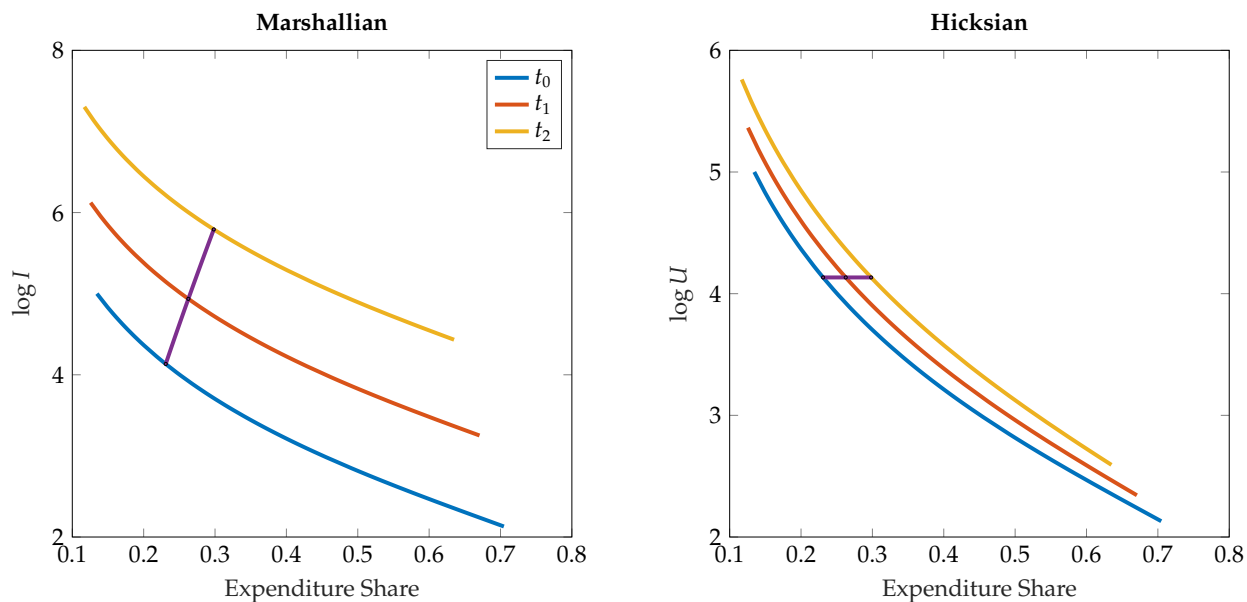
$$\log u(I, t_2) \approx \log I - b(u(I, t_1), t_1) \cdot \Delta \log p_{t_2} - b(u(I, t_1), t_0) \cdot \Delta \log p_{t_1},$$

and using $u(I, t_2)$, construct Hicksian budget shares for period t_2 :

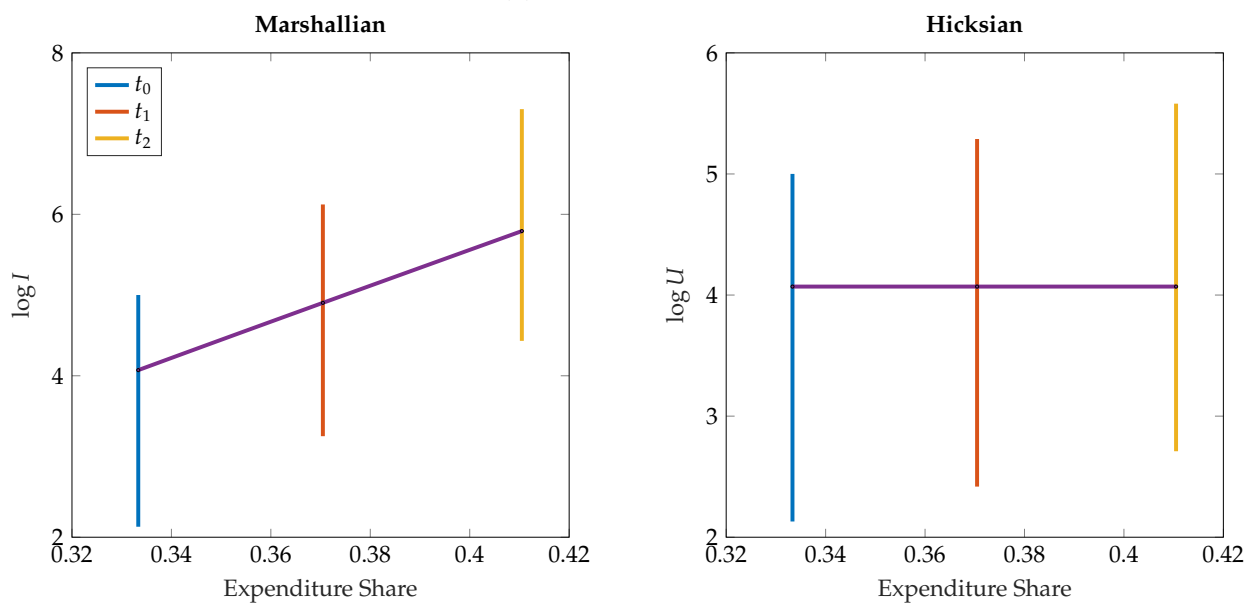
$$b(u(I, t_2), t_2) = B(I, t_2).$$

¹¹In 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. Moreover, when computing equations (4) and (5) requires evaluating a function between two grid points, we use linear interpolation.

That is, for each budget share $B(I, t_2)$ in t_2 , assign a utility value based on $u(I, t_2)$. 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 $s \leq t$.



(a) Non-homothetic



(b) Homothetic

Figure 1: 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 1a 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 1a displays the Hicksian budget shares for the same good. The purple line in the right panel of Figure 1a 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 1a 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).

To understand why Proposition 2 is unnecessary when preferences are homothetic, Figure 1b plots the same information as Figure 1a 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.¹² In this case, we can handle this by splitting the sample in two and applying our method to each sample separately.¹³

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).¹⁴ The expenditure function is

$$e(p, u) = \left(\sum_i \omega_i (u^{\varepsilon_i} p_i)^{1-\sigma(u)} \right)^{\frac{1}{1-\sigma(u)}}.$$

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

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

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

¹²This assumption is related to the assumption considered in Section 2.3 of Jaravel and Lashkari (2022).

¹³Similarly, 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.

¹⁴See also Hanoch (1975), Comin et al. (2021), and Matsuyama (2019) for more information on these preferences. In Appendix E we consider another example with an Almost Ideal Demand System (AIDS).

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

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

where v is the solution to $I = e(p_t, v)$. To evaluate the accuracy of our algorithm, we compare this expression for $u(I, T)$ 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 A.1 in Appendix A plots the paths of prices and incomes in our numerical example.

We consider two parameterizations. The first 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 second case assumes 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). Figure A.2 shows the difference between inflation rates calculated using the cost-of-living index for different income levels and the inflation rate implied by an aggregate chained consumption price deflator.

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.¹⁵ Figure 2 shows how this error varies as we vary the number of households

¹⁵As mentioned earlier, if we solve the fixed point problem described in Footnote 10, then the error is three orders of magnitude smaller. That is, 2×10^{-5} and 9×10^{-5} instead of 0.0078 and 0.0044. 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.

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.

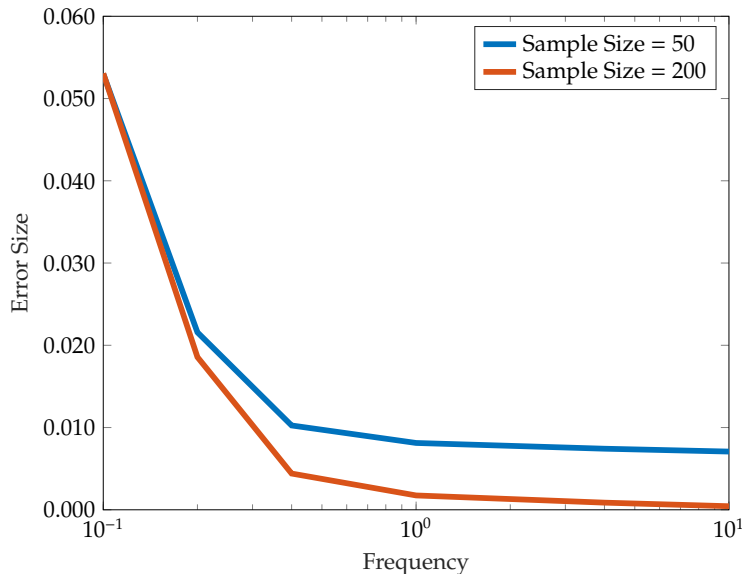


Figure 2: This figure displays the maximum error as a function of the frequency of observation, holding the path of price and income changes in Figure A.1 constant. This figure uses non-homothetic CES preferences with variable σ . Our baseline calibration is annual frequency corresponding to a value 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 cross-sectional household data. Our goal is to compare changes in welfare as measured by the money-metric with changes in real income. 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). 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 2017.¹⁶ The UK Family Expenditure Survey was used in Blundell

¹⁶Aggregate 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.

et al. (2003) and Blundell et al. (2008) to estimate Engel curves, test for deviations from revealed preference theory, and compute bounds for a true cost of living index.

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 B 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

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

where i is the good, h is the household, and t is the time period. The estimated regression line gives us $B(I, t)$.¹⁷ This regression is the only source of sampling uncertainty in our exercise. We calculate standard errors for our estimates of the money-metric by bootstrapping this regression. To do this, we redraw repeated samples with replacement. Although the Engel curves are estimated with considerable uncertainty, the standard errors for the money-metric are quite tight. This is due to the law of large numbers, since the money-metric combines many Engel curve estimates. For this reason, when we present our results, we do not report the bootstrapped standard errors.

We apply our procedure sequentially from 1974 to 2017 using the UK cross-sectional data constructed in the manner described above. Computing $u(I, t)$ 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 3 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

While these sample coverage issues affect aggregate nominal growth rates, they do not affect our results, which are at the household-level.

¹⁷We also estimated $B_i(I, t)$ using locally weighted scatterplot smoothing (LOWESS) and obtained very similar results.

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 3 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.

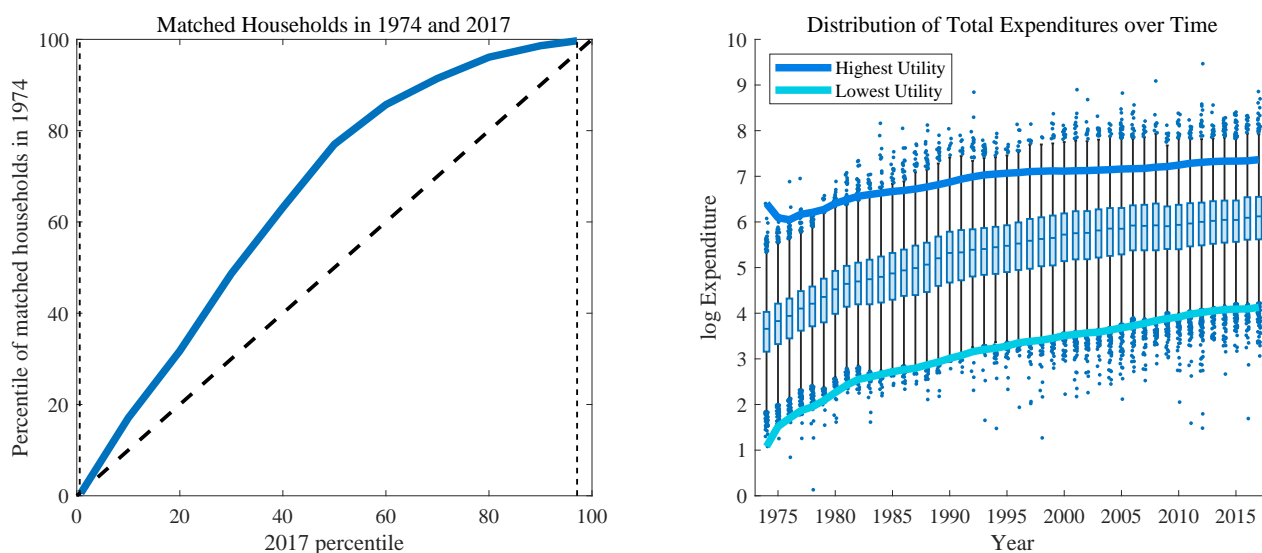


Figure 3: The figure on the left shows, for each income percentile in 2017, the income percentile in 1974 of the matched household that is on the same indifference curve as the 2017 household. The figure on the right shows the sample distribution of (weekly) log expenditures from 1974 to 2017. The upper 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 4 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. Hence, the money-metric is higher than real income for richer households and lower for poorer households, and the size of the gap 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 5 displays the log difference between the red and blue lines in Figure 4. 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, as revealed by the histograms in Figure 5, that inequality across households is larger based on the money-metric than based on real income.

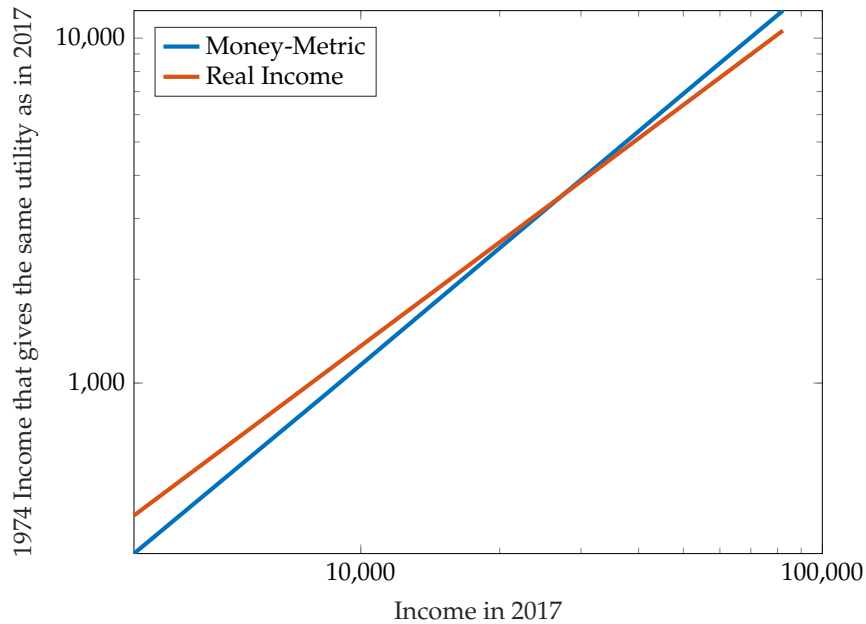


Figure 4: Money-metric $e(p_{1974}, v(p_{2017}, I_{2017}))$ and real income using aggregate chain-weighted inflation between 1974 to 2017 (annualized pounds, log scale).

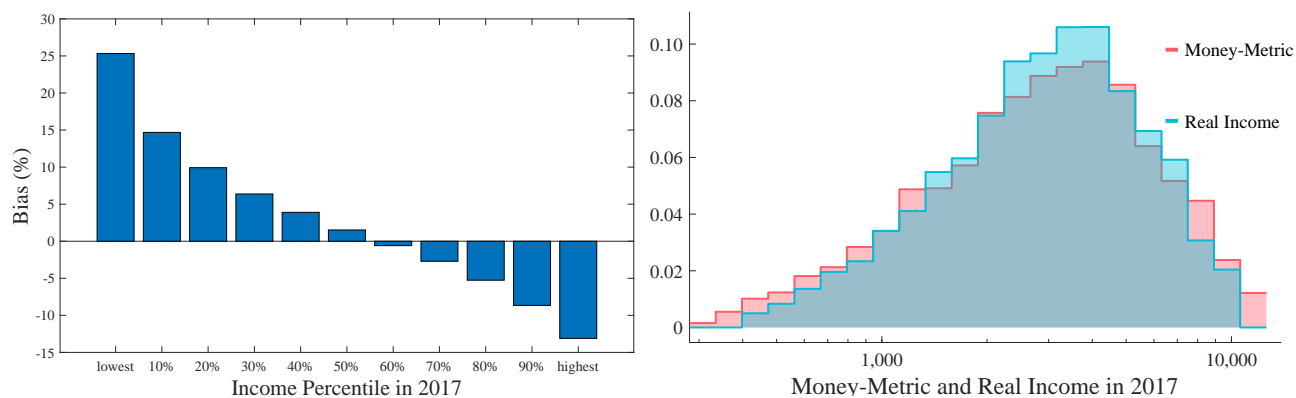


Figure 5: The left panel shows the log difference between real income and money-metric in 1974 for different percentiles of the income distribution in 2017. Highest and lowest correspond to the utilities and their percentiles in Figure 3. The right panel is a histogram (using household weights) of money-metric $e(p_{1974}, v(p_{2017}, I_{2017}))$ and real income using aggregate chain-weighted inflation (annualized pounds, log scale). The distributions are truncated at the upper and lower bounds of Figure 3.

The reason for the patterns we document is the following. 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. Since the aggregate inflation rate is expenditure- rather than population-weighted, the average inflation rate tends to put more weight on the expenditures of relatively rich households and hence performs better for them than relatively poor households.

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.
- Baqee, D. and A. Burstein (2021). Welfare and output with income effects and taste shocks. Technical report, National Bureau of Economic Research.
- Blundell, R., M. Browning, and I. Crawford (2008). Best nonparametric bounds on demand responses. *Econometrica* 76(6), 1227–1262.
- Blundell, R. W., M. Browning, and I. A. Crawford (2003). Nonparametric Engel curves and revealed preference. *Econometrica* 71(1), 205–240.
- Boppart, T. (2014). Structural change and the Kaldor facts in a growth model with relative price effects and non-Gorman preferences. *Econometrica* 82(6), 2167–2196.
- Comin, D., D. Lashkari, and M. Mestieri (2021). Structural change with long-run income and price effects. *Econometrica* 89(1), 311–374.
- Deaton, A. and J. Muellbauer (1980). An almost ideal demand system. *The American Economic Review* 70(3), 312–326.
- Fajgelbaum, P., G. M. Grossman, and E. Helpman (2011). Income distribution, product quality, and international trade. *Journal of Political Economy* 119(4), 721–765.
- Fally, T. (2022). Generalized separability and integrability: Consumer demand with a price aggregator. *Journal of Economic Theory* 203, 105471.
- Fan, T., M. Peters, and F. Zilibotti (2022). Growing like India: The unequal effects of service-led growth. Technical report, Technical report, National Bureau of Economic Research.
- 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.
- Hausman, J. (1981). Exact consumer's surplus and deadweight loss. *American Economic Review* 71(4), 662–76.
- 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. (2000). A ricardian model with a continuum of goods under nonhomothetic preferences: Demand complementarities, income distribution, and north-south trade. *Journal of political Economy* 108(6), 1093–1120.
- 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.
- Straub, L. (2019). Consumption, savings, and the distribution of permanent income. *Unpublished manuscript, Harvard University*.

Online Appendix

A Additional Figures for Section 4

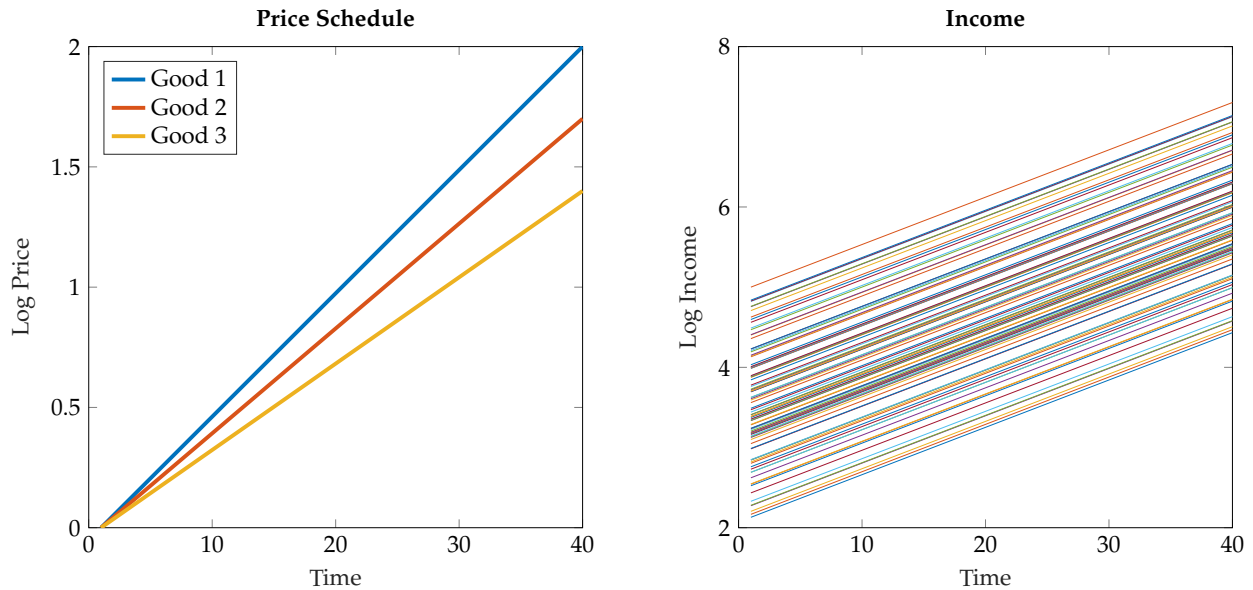


Figure A.1: Exogenous price and income paths for the artificial examples in Section 4

B Additional details of the UK data used in Section 5

We use two different datasets. One is a household-level expenditure survey and the other is data on prices of different categories of goods. The first data set is *Family Expenditure Survey and Living Costs and Food Survey Derived Variables*, which is a dataset of annual household expenditures with demographic information compiled from various household surveys conducted in the UK. Each sample includes about 5,000-7,000 households. The spending categories in the survey correspond to RPI (Retail Price Index) categories. We have continuous data from 1974 to 2017. Starting in 1995, the data are split into separate files for adults and children, so we merge them into households by adding up their expenditures.

Our algorithm does not require a representative sampling of 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 samples from the entire income distribution except for top earners

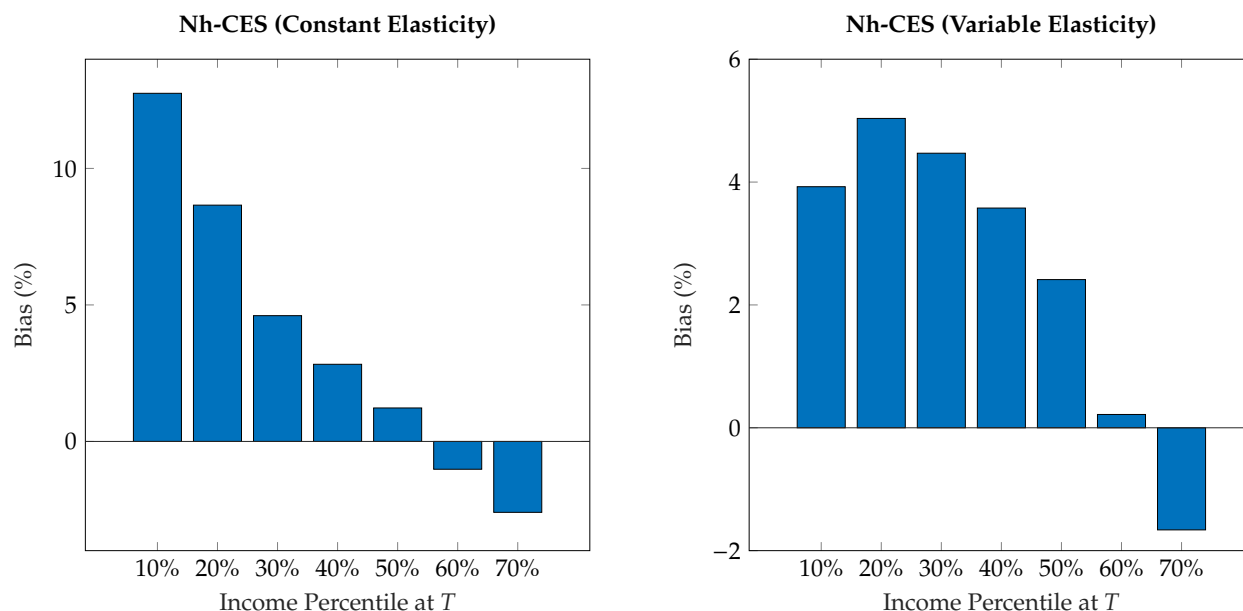


Figure A.2: Non-homotheticity bias: log difference between real income and the money-metric for artificial examples in Section 4. For poor households, the money metric is lower than real income calculated using aggregate inflation because the inflation rate is lower for income elastic goods.

and some pensioners. In order to correct for possible nonresponse bias, household weights are provided since 1997.¹⁸ 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 for the money-metric does not use household weights.¹⁹

For the prices, we use the underlying data for the consumer price index (CPI) and the retail price index (RPI). 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).²⁰ By comparing the RPI and CPI with the consumption deflator provided by the Office of National Statistics, 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, “al-

¹⁸Prior 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.

¹⁹We also use weights to calculate the percentiles in the left panel of Figure 3, the histograms in Figure 5, and the quintiles for Figure A.5 in the Appendix.

²⁰https://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.

cohol” 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 A.1 summarizes how we integrated the CPI 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 A.1: RPI and CPI Correspondence Table

C Comparison with Blundell et al. (2003)

In this appendix, we exposit and apply the welfare bounds in Blundell et al. (2003) to artificial and real data. We start by discussing how we implement their methodology, since, due to an inconsistency in their equations, we do not exactly implement their procedure.

C.1 Description of Bounding Algorithm

To bound the cost-of-living, Blundell et al. (2003) provide an algorithm for an upper-bound and a lower-bound. Following the notation in their paper, let $q_t(I)$ be bundle of goods

consumed by a household with income I in period t . Blundell et al. (2003) assume that $q_t(I)$ is an injective function (each I maps to a unique bundle of quantities in each period).

Algorithm A (Upper-bound). To recover an upper-bound for $e(p_s, v(p_t, I_t))$, start by defining $q^* = q_t(I_t)$ and let T be the set of periods for which we have data.

- (1) Set $i = 0$ and $F^{(i)} = \{q_s^i = q_s(p_s \cdot q^*)\}_{s \in T}$.
- (2) Set $F^{(i+1)} = \{q_s^{i+1} = q_s(\min_{q \in F^{(i)}} p_s \cdot q)\}_{s \in T}$.
- (3) If $F^{(i+1)} = F^{(i)}$, then set $Q_B(q^*) = F^{(i)}$ and stop. Else set $i = i + 1$ and go to step (2).

We have that $e(p_s, v(p_t, I_t)) \leq \min_q \{p_s \cdot q : q \in Q_B(q^*)\}$.

Intuitively, the cost of living in period s associated with q^* , $e(p_s, v(p_t, I_t))$, is weakly less than $p_s \cdot q^*$. Hence, for every s , we must have that $q_s^0 = q_s(p_s \cdot q^*)$ is weakly preferred to q^* . This collection of bundles, $\{q_s^0\}_{s \in T}$, all of which are preferred to q^* , is $F^{(0)}$ defined in step (1). In step (2), we search across all of these bundles to find the cheapest one in each period s . We update each q_s^i to be the bundle that households with that level of income actually picked in each period (which is still better than q^*). We continue this indefinitely until this procedure converges, at which point we have our upper-bound.

As mentioned in the text, the lower-bound algorithm provided by Blundell et al. (2003) is not correct. We provide an amended version below.

Amended Algorithm B (Lower-bound). To recover a lower-bound for $e(p_s, v(p_t, I_t))$, start by defining $q^* = q_t(I_t)$ and let T be the set of periods for which we have data.

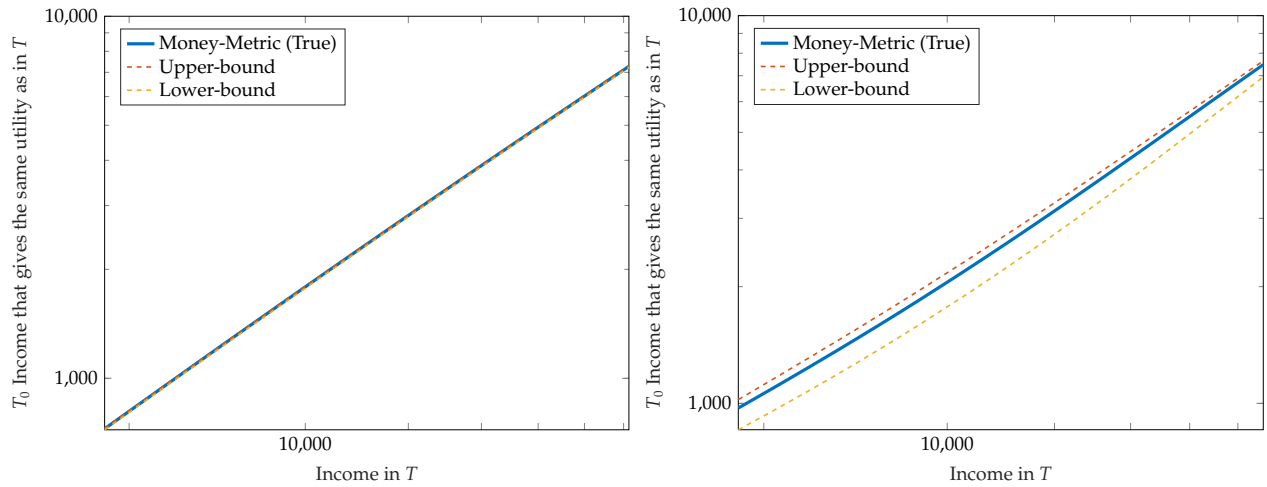
- (1) Set $i = 0$, and let $F^{(i)} = \{I_s^i : p_t \cdot q_s(I_s^i) = I_t\}_{s \in T}$.
- (2) Set $F^{(i+1)} = \{\max_{I_k \in F^{(i)}} \{I_s^{i+1} : I_k = p_k \cdot q_s(I_s^{i+1})\}\}_{s \in T}$.
- (3) If $F^{(i+1)} = F^{(i)}$, then set $Q_W(q^*) = \{q_s(I_s^i)\}_{s \in T}$ and stop. Else set $i = i + 1$ and go to step (2).

We have that $\max_{q_s \in Q_W(q^*)} p_s \cdot q_s \leq e(p_s, v(p_t, I_t))$.

Intuitively, in step (1), for each period s , we find the income level I_s^0 such that $p_t \cdot q_s(I_s^0) = I_t$. The bundle $q_s(I_s^0)$ was affordable at t but was not purchased. Hence, the true cost-of-living in period s must be greater than I_s^0 . The collection of income levels constructed in this step is $F^{(0)}$ and all are less than the true cost-of-living. In step (2), for each period s , we search over I_k^i and find the maximum level of income I_s^{i+1} such that $I_k^i = p_k \cdot q_s(I_s^{i+1})$ is satisfied. The new I_s^{i+1} is weakly greater than I_s^i but we still know that I_s^{i+1} is less than the

true cost-of-living. We continue this indefinitely until this procedure converges, at which point we have our lower-bound.

C.2 Results with Artificial & UK Data



(a) Non-homothetic CES with constant elasticity (b) Non-homothetic CES with variable elasticity

Figure A.3: Upper- and lower-bound using the amended Blundell et al. (2003) algorithm for the artificial economies in Section 4. Our algorithm results are indistinguishable from the blue line.

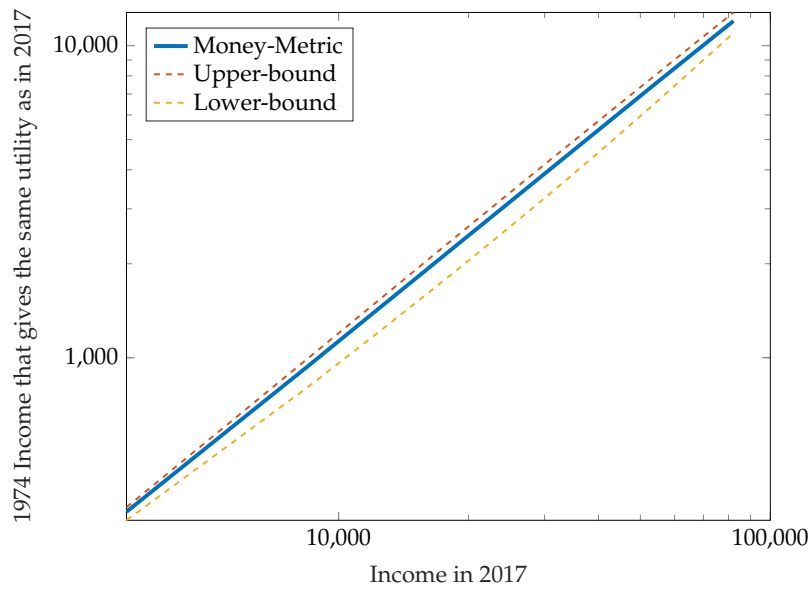


Figure A.4: Upper- and lower-bound using the amended Blundell et al. (2003) algorithm for the UK data in Section 5. Our algorithm produced the blue line.

D 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. Note that the definition of real consumption in our paper is not the same as theirs. For brevity, we do not include in this appendix a description of these algorithms.

D.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 A.2 shows that the error is very low for non-homothetic CES (with constant elasticity of substitution). On the other hand, the approximation error is larger for non-homothetic CES with variable elasticity of substitution.

	Nh-CES(Constant)	Nh-CES(Variable)
First Order	0.0491	0.1422
Second Order	0.0028	0.1130

Table A.2: $\max_h |\log u(I, T) - \log u(I, T)^{TRUE}|$: results of the first/second order algorithm of Jaravel and Lashkari (2022) with their K parameter set to two applied to the artificial data in Section 4.

D.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 A.5. We next apply their algorithms to the underlying disaggregated data that we use in our empirical results; results are displayed in Figure A.6.

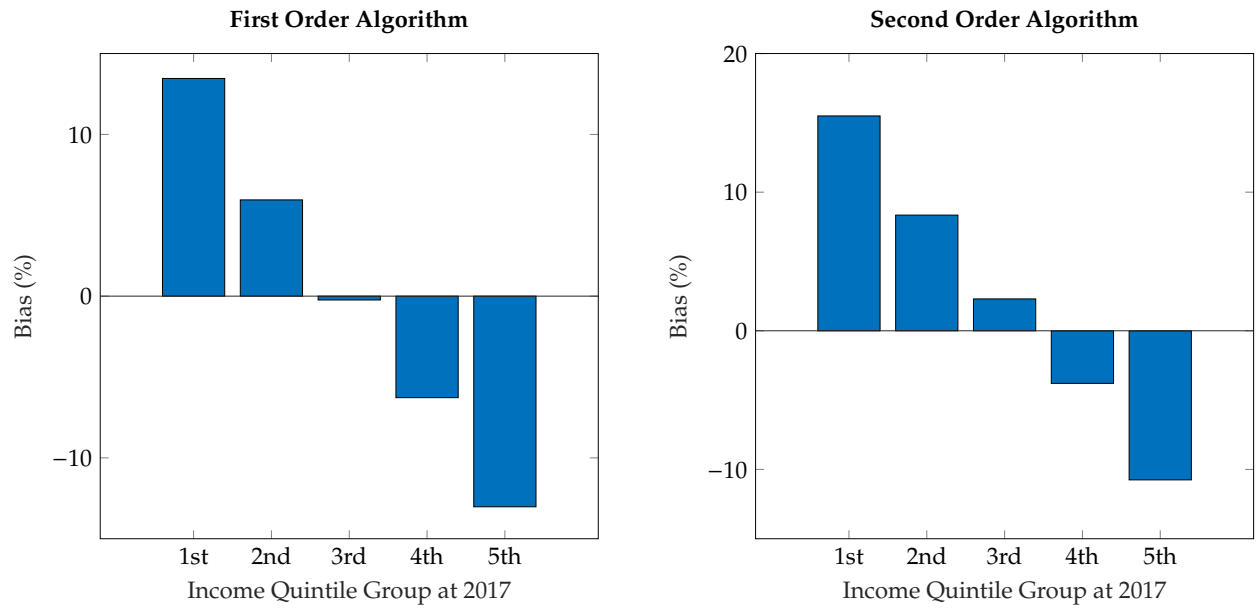


Figure A.5: Results of the first/second order algorithm of Jaravel and Lashkari (2022) with $K = 1$ to the aggregated (by quintile) UK household data: Log difference between real income and the money-metric

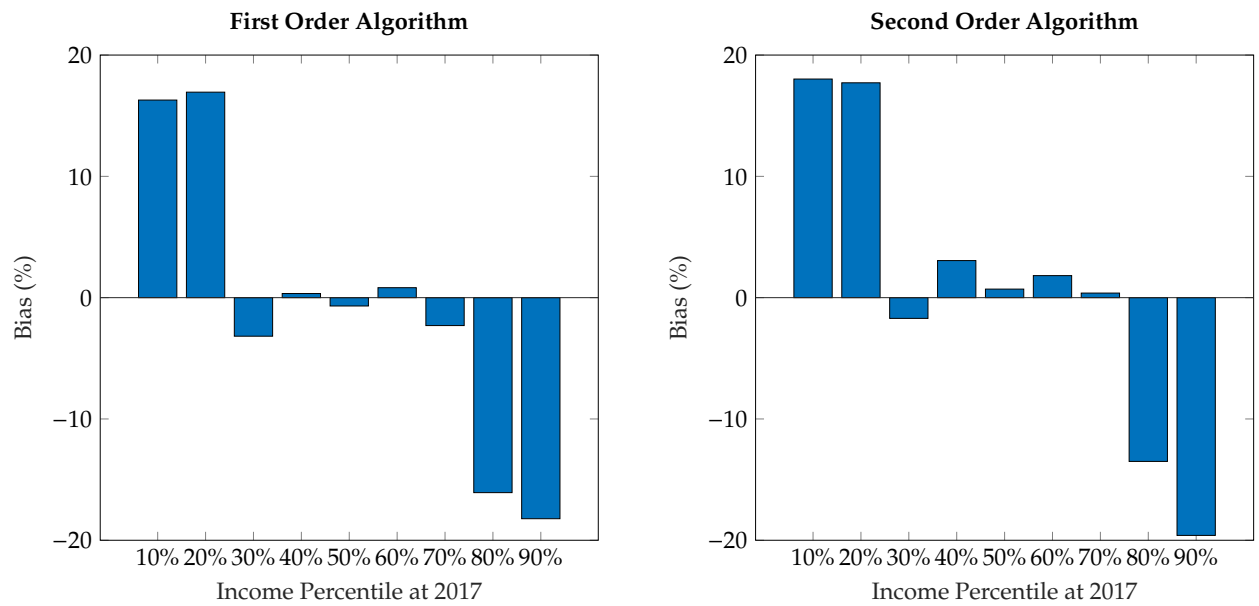


Figure A.6: Results of the first/second order algorithm of Jaravel and Lashkari (2022) with $K = 1$ to the disaggregated UK household data: Log difference between real income and the money-metric

E 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: the Almost Ideal Demand System (AIDS) due to Deaton and Muellbauer (1980). 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 is²¹

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

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.

$$\begin{bmatrix} \alpha_0 & \alpha_1 & \alpha_2 & \alpha_3 & \beta_1 & \beta_2 & \beta_3 & \gamma_{11} & \gamma_{22} & \gamma_{33} & \gamma_{12} & \gamma_{13} & \gamma_{23} \\ 2 & 1/3 & 1/3 & 1/3 & -0.15 & -0.05 & 0.2 & -1/4 & -1/4 & -1/4 & 1/8 & 1/8 & 1/8 \end{bmatrix}$$

The approximation error, $\max_i |\log u(I, T) - \log u(I, T)^{TRUE}|$, is 0.0011 when we use the iterative procedure and 7.0×10^{-7} when we use the fixed point procedure (see Footnote

²¹To obtain the expression for the money-metric, we use $e(p_0, v(p, I)) = e(p_0, u)$, where $I = c(p) u^{d(p)}$.

10). For completeness, Figure A.7 presents the non-homotheticity bias in our artificial AIDS example.

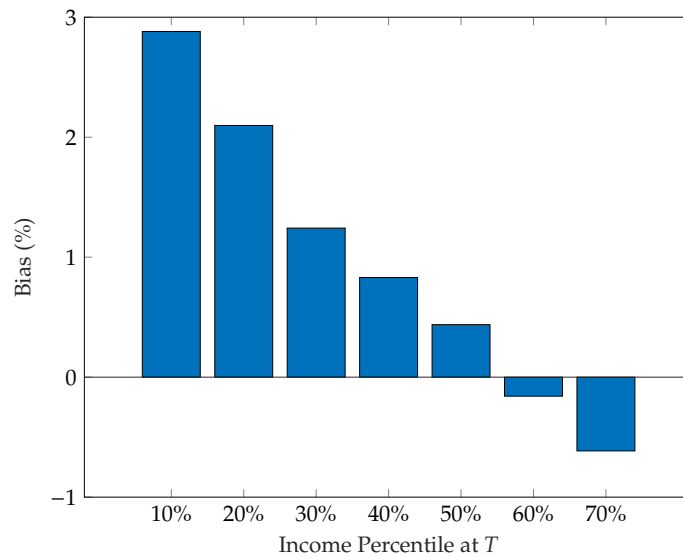


Figure A.7: Log difference between real income and the money-metric using our algorithm and the Almost Ideal Demand System.