

NBER WORKING PAPER SERIES

A FIXED POINT APPROACH TO MEASURING WELFARE

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 December 2022

First circulated as “A New Method for Measuring Welfare with Income Effects using Cross-Sectional Data.” We thank Andy Atkeson, John Asker, Natalie Bau, Dave Donaldson, Pablo Fajgelbaum, 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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 Fixed Point Approach to Measuring Welfare
David Baqaee, Ariel Burstein, and Yasutaka Koike-Mori
NBER Working Paper No. 30549
October 2022, Revised December 2022
JEL No. E01,E31

ABSTRACT

We present a method for recovering welfare, measured via the money metric, using repeated cross-sectional household data for unrestricted preferences. Our approach is distinct from previous studies by Blundell et al. (2003) and Jaravel and Lashkari (2022) in that it converts the problem of measuring welfare into a fixed point problem. We show that, given Hicksian demand, the money-metric utility can be constructed by integration, and that if the money-metric is known, then Hicksian demand can be constructed by matching households in the cross-section who are on the same indifference curve at different points in time. With additional assumptions, our method also allows for missing or mismeasured prices, recasting the problem of inferring the missing prices as an additional fixed point problem. We apply our method to household consumption survey data from the United Kingdom from 1974 to 2017 and find that real consumption calculated using official aggregate inflation statistics overstates money-metric utility for the poorest households by around half a percent per year and understates it by around a quarter of a percentage point per year for the richest households. We discuss how our results change if the price of some service sectors is mismeasured.

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

Money-metric measures of welfare convert income under different sets of prices into equivalent income under a common base price vector. That is, a money-metric assigns cardinal values to budget constraints that are comparable to each other and have interpretable units. For this reason, money-metric utility is a backbone of welfare economics and a necessary ingredient for measuring welfare-relevant growth and inflation.

In general, money-metric utility can be calculated by integrating Hicksian or compensated demand with respect to changes in prices (see e.g. Hausman, 1981). In practice, compensated demand is not easily observed, so standard price deflators weigh changes in prices using uncompensated demand instead. This shortcut is justified if preferences are homothetic, as in this case, compensated and uncompensated expenditure shares coincide. However, in the more realistic case of non-homothetic preferences, money-metric utility is more difficult to calculate since Hicksian budget shares are not easily obtained (see Samuelson and Swamy, 1974).

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 function 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 revealed 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 If we can find such households, then we can calculate the money-metric utility function by integration. Conversely, to match households through time, we must know the money-metric utility function, since two households are on the same indifference curves if their money-metric utility values coincide. The insight is that this is a fixed point problem, and we propose a simple iterative algorithm that converges to the money-metric utility function in continuous time.

Our approach generalizes the standard practice of statistical agencies who weigh changes in prices over time using aggregate budget shares. When preferences are homothetic and stable in both the cross-section and the time-series, aggregate budget shares coincide with the compensated budget share of each individual in the sample. Hence, under these assumptions, conventional price deflators like the CPI or the PCE recover money-metric utility. However, when preferences are non-homothetic, the compensated

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.

budget shares of each household are distinct. Thus, instead of relying on aggregate budget shares to construct the price index, we need to use the budget shares of a unique corresponding income level in the past for each income today. Our paper provides a method for constructing this mapping.³

Our approach differs from an alternative that calculates Hicksian demand based on the estimated elasticities of substitution. Unlike the latter, our procedure does not require the estimation of non-parametric elasticities of substitution.⁴ Intuitively, our approach only recovers Hicksian demand evaluated at observed prices, whereas the elasticities of substitution determine how Hicksian demand will react to any price alteration, even those that have not been observed. For this reason, our procedure can measure changes in welfare for observed changes in prices and income, but it cannot address counterfactual welfare questions as those explored by Baqaee and Burstein (2021).

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

Blundell et al. (2003) bound the money-metric by using revealed choice arguments. For each income level at time t , Blundell et al. (2003) 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. Our approach has an advantage over Blundell et al. (2003) in that it provides a point estimate, rather than only bounds, for

³In general, chained-weighted indices as measured by statistical agencies are uninterpretable when preferences are non-homothetic. However, under additional assumptions chained indices do have meaningful interpretations. For example, Feenstra and Reinsdorf (2000) show that when the path of prices is linear in time then chained indices measure the cost-of-living price index for some intermediate utility level (between initial and final utility) under AIDS preferences. Caves et al. (1982) establish a similar result for Tornqvist price indices, up to a second order approximation. In this paper, we are interested in the money-metric utility function instead.

⁴In this sense, our approach is similar to Oulton (2012) who shows how to back out Hicksian budget shares by adjusting Marshallian budget shares using a Taylor series in income. He applies this methodology using the QAIDS demand system to estimate the cost-of-living index without needing to estimate price elasticities. Instead of relying on a Taylor series under a parametric functional form for demand, our approach purges income effects from substitution effects by matching households over time who are on the same indifference curve but face different prices.

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 to address non-homotheticity in household-level chain-weighted indices. Their approach requires that the support of the cross-sectional distribution of (ex-ante unknown) utilities remain constant and unchanging over time. In contrast, our approach does not make this assumption; instead, it endogenously determines the region of the income distribution for which a money-metric utility function can be constructed, given the available information. This is important when economic growth shifts the support of utilities over time. In Appendix D, we apply the Jaravel and Lashkari (2022) approach to our artificial and UK data and show that their algorithm can result in larger approximation errors when there is economic growth.

In contrast to these papers, we also extend our method to situations where some prices and expenditures are unobserved. This generalizes the influential Feenstra (1994) approach to imputing missing prices beyond the homothetic CES case. To do this, we require that preferences be indirectly separable between observed and unobserved goods. Under this additional assumption, we show that the money-metric utility can be recovered provided knowledge of the compensated elasticity of substitution between observed and unobserved goods.

This is possible because we can back out the change in the relative price of observed and unobserved goods using changes in the compensated expenditure share of the observed goods. For example, if the compensated budget share on observed goods is rising, and observed goods are net complements with unobserved goods, then this indicates that the relative price of unobserved goods is falling. This can then be used to calculate money-metric utility. However, knowing the change in the compensated budget share of observed goods requires knowing the money-metric utility function, making this another fixed-point problem. Importantly, we also show that the elasticity of substitution between observed and unobserved goods, which is required to infer missing prices, can be identified without knowledge of those missing prices.

Since our method can be extended to allow for unmeasured prices, our paper is also related to a strand of the literature that measures welfare allowing for incomplete information about prices. Notable examples include Costa (2001), Hamilton (2001) and, more recently, Atkin et al. (2020). These papers take advantage of horizontal shifts in Engel curves to identify money-metric utility changes. Our approach has distinct intuition, assumptions, and data requirements; we discuss these in more detail in Section 5 when we discuss the extension of our method to allow for missing prices.

Our approach can also be contrasted with more parametric approaches where welfare measures are computed using a fully-specified 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). In contrast, our approach provides a non-parametric way to compute welfare measures from the data without relying on low-dimensional functional forms.

The outline of the paper is as follows. In Section 2, we define the cost-of-living index and money-metric utility and show how they are related to Hicksian demand. Section 3 showcases how this can be applied to recover cost-of-living and money-metric utility with the help of cross-sectional data. We apply our method to artificial data generated using popular functional forms for non-homothetic preferences and show 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 this section, we also discuss several extensions and limitations of our approach, including how to use our results when some of our baseline assumptions are not met. For example, when sufficient data is available, we discuss how to deal with idiosyncratic taste shocks that are unrelated to income. Similarly, we explain how our method can be adapted to account for heterogeneity in preferences that are dependent on observable characteristics.

In Section 4, we apply our method to construct the money-metric utility function using household expenditure survey data from the United Kingdom from 1974 to 2017. We find that real consumption calculated by deflating income by aggregate chain-weighted inflation (as measured by official statistical agencies) overstates the money-metric utility for all households below the 60th percentile of income in 2017 in our sample. In other words, for any income level less than the 60th percentile, the 1974 equivalent income is less than real consumption. The size of this gap is greatest for the poorest households, roughly 20 percentage points (0.5 percentage points per year on average), and gradually diminishes until it reaches zero for households close to the 60th percentile. Conversely, real consumption calculated using aggregate inflation statistics understates the money-metric utility for households above the 60th level. For households in the 97th percentile of our sample, who spend around £81,000 per year, the size of this gap is 13 percentage points over the whole sample (0.25 percentage points per year on average).⁵ We are

⁵These results are consistent with Blundell et al. (2003), which report a relatively greater rise on the cost

unable to compute the money-metric for the richest households in 2017 (97th percentile and above). The reason is that for these households, there did not exist equally well-off consumers in the past whose demand can be used in place of the compensated demand curves.

In Section 5, we extend our methodology to allow for missing prices. To illustrate this extension, we assume that some service prices are mismeasured in the UK data, calibrate an elasticity, and apply our methodology. We find that the relative price of services has been rising faster than official data for rich but not poor households. We conclude in Section 6.

2 Money-Metrics and the Cost-of-Living

We start by defining the objects of interest: money-metric utility and the closely related cost-of-living function. Consider a rational preference relation \succeq defined over consumption bundles c in \mathbb{R}^N . Suppose that these preferences can be represented by a smooth utility function $U(c)$ that maps consumption bundles to utility values. Given this utility function, we can define the *indirect utility function*

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

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

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

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 $\bar{\mathbf{p}}$, the *money-metric* function maps (\mathbf{p}, I) to

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

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

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

The money-metric $e(\bar{\mathbf{p}}, v(\cdot))$ is itself an indirect utility function because a budget set of living for poorer households between 1975 and 1984 in the UK.

(\mathbf{p}, I) is preferred to another budget set (\mathbf{p}', I') if, and only if, $e(\bar{\mathbf{p}}, v(\mathbf{p}, I)) > e(\bar{\mathbf{p}}, v(\mathbf{p}', I'))$. We use this money-metric cardinalization of utility throughout the rest of the paper. Whereas the money-metric converts the value of different budget sets into equivalent value under some baseline prices, the cost-of-living function, $e(\cdot, v(\bar{\mathbf{p}}, \bar{I}))$, converts the value of some baseline budget constraint $(\bar{\mathbf{p}}, \bar{I})$ into equivalent income under different sets of prices.⁶

To summarize, the function $e(\mathbf{p}', v(\mathbf{p}, I))$, mapping $(\mathbf{p}', \mathbf{p}, I)$ into a scalar, is an object of paramount interest. The “money-metric” is the cross-section of this function that holds \mathbf{p}' constant and the cost-of-living index is the cross-section that holds (\mathbf{p}, I) constant. The money-metric is useful for ranking budget sets (i.e. measuring growth).⁷ The cost-of-living index is useful for converting a common utility level, attained by $v(\mathbf{p}, I)$, into equivalent income under different price systems (i.e. measuring the cost of maintaining a fixed standard of living).

Denote the Hicksian budget share for good i to be $b_i(\mathbf{p}, u)$ where \mathbf{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 (\mathbf{p}, I) in terms of $\bar{\mathbf{p}}$ prices can be expressed as*

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

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

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

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.

⁶In index number theory, the cost-of-living index is also called the Konüs (1939) index.

⁷The equivalent and compensating variation are related to the money-metric and the cost-of-living index. Specifically, to measure the change in welfare from some initial budget set (\mathbf{p}, I) to some other budget set (\mathbf{p}', I') , the equivalent variation (in logs) is $\log e(\mathbf{p}, v(\mathbf{p}', I')) - \log I$ and the compensating variation is $\log I' - \log e(\mathbf{p}', v(\mathbf{p}, I))$.

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 $\mathbf{p}_t \in \mathbb{R}^N$ at each point in time $t \in [t_0, T]$.⁸ We also observe vectors of expenditure shares $\mathbf{B}(I, t) \in \mathbb{R}^N$ for consumers with preferences \succeq and income levels $I \in [\underline{I}, \bar{I}]$ at time t . The expenditure shares $\mathbf{B}(I, t)$ can be thought of as Marshallian budget shares evaluated at income level I and prices \mathbf{p}_t .

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

$$b_i(\mathbf{p}_t, v(\mathbf{p}_t, I)) = B_i(I, t).$$

Denote the money-metric utility function using reference prices \mathbf{p}_{t_0} and evaluated at budget set (\mathbf{p}_t, I_t) by $u(I, t)$. Using this cardinalization of indirect utility, we can write

$$b_i(\mathbf{p}_t, u(I, t)) = B_i(I, t).$$

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 $u(I, t) \equiv e(\mathbf{p}_{t_0}, v(\mathbf{p}_t, I))$ 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

⁸Prices 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 4). 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.

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

In words, $u(I, t)$ converts the value of the budget constraint defined by prices \mathbf{p}_t and income I into income under \mathbf{p}_{t_0} . By varying initial prices \mathbf{p}_{t_0} , for fixed (\mathbf{p}_t, I) , we can also recover the cost-of-living index.⁹

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 .

To better understand (3), observe the simplification that occurs when preferences are homothetic. In this case, budget shares do not depend on income levels, only on time. Therefore, when preferences are homothetic, (3) simplifies to

$$\log u(I, t) = \log I - \int_{t_0}^t \sum_i B_i(s) \frac{d \log p_{is}}{ds} ds, \quad (4)$$

which eliminates the need to find a fixed point. This equation, called a Divisia index, justifies the standard chain-weighting practices adopted in the national accounts for calculating price and quantity indices.

If we can solve (3), then we can compute the Hicksian budget shares $b(\mathbf{p}_s, u_t)$ for a utility level u_t at time t under prices \mathbf{p}_s at time s by using the budget shares of a different household on the same indifference curve at time s . The expenditures shares of this “matched” household, $B(u^{-1}(u_t, s), s)$, are equal to the Hicksian budget share $b(\mathbf{p}_s, u_t)$.

Proposition 2 provides a way to recover the money-metric and cost-of-living functions without needing direct knowledge of the potentially very high-dimensional demand system. Recall that the number of cross-price elasticities scales in the square of the number goods, and generically depends on both income and relative prices. Proposition 2 obviates the need to undertake this onerous estimation exercise by using the demand from other households and time periods in place of a counterfactual model of compensated demand.

The integral equation in Proposition 2 may initially appear abstract, but its underlying intuition is quite simple. We can better understand its meaning and how it can be used in practice using the following step-by-step solution procedure.

⁹ Alternatively, suppose that we have calculated $e(\mathbf{p}_{t_0}, v(\mathbf{p}_t, I))$ and we wish to obtain $e(\mathbf{p}_s, v(\mathbf{p}_t, I))$ for some $s \in [t_0, T]$. Find I' such that $e(\mathbf{p}_{t_0}, v(\mathbf{p}_t, I)) = e(\mathbf{p}_{t_0}, v(\mathbf{p}_s, I'))$. Then, $v(\mathbf{p}_t, I) = v(\mathbf{p}_s, I')$ and $e(\mathbf{p}_s, v(\mathbf{p}_t, I)) = e(\mathbf{p}_s, v(\mathbf{p}_s, I')) = I'$.

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.

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 $n = 1$:

$$\log u(I, t_n) \approx \log I - \sum_{s=0}^{n-1} \mathbf{B}(I_s^*, t_s) \cdot \Delta \log \mathbf{p}_{t_s}, \quad (5)$$

where I_s^* satisfies

$$u(I_s^*, t_s) = u(I, t_{n-1}), \quad (6)$$

with the boundary condition $u(I, t_0) = I$. The summation in (5) 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 $u(I, t_{n-1})$, in (6), converges to $u(I, t_n)$.¹⁰ 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. Abusing notation, let $b_i(u, t)$ be the Hicksian budget share of good i at prices \mathbf{p}_t for utility value u . For period t_1 , compute

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

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

¹⁰By uniqueness of limits, if (5) and (6) converge as the time interval shrinks to zero, this limit point is unique and must be the money-metric utility function. In our computations, we use the trapezoid rule rather than the left Riemann sum in equation (5) to approximate the integral in (3) since it is a better numerical approximation. Moreover, when computing equations (5) and (6) requires evaluating a function between two grid points, we use linear interpolation. We do not extrapolate.

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

$$\mathbf{b}(u(I, t_1), t_1) = \mathbf{B}(I, t_1).$$

That is, to each budget share $B_i(I, t_1)$, 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 $B_i(I, t_1)$ a utility value, which is precisely the Hicksian budget share. We now have Hicksian budget shares $\mathbf{b}(u, t_0)$ and $\mathbf{b}(u, t_1)$.

Next, calculate

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

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


$$\mathbf{b}(u(I, t_2), t_2) = \mathbf{B}(I, t_2).$$

That is, for each budget share $B_i(I, t_2)$ 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 $s \leq t$.

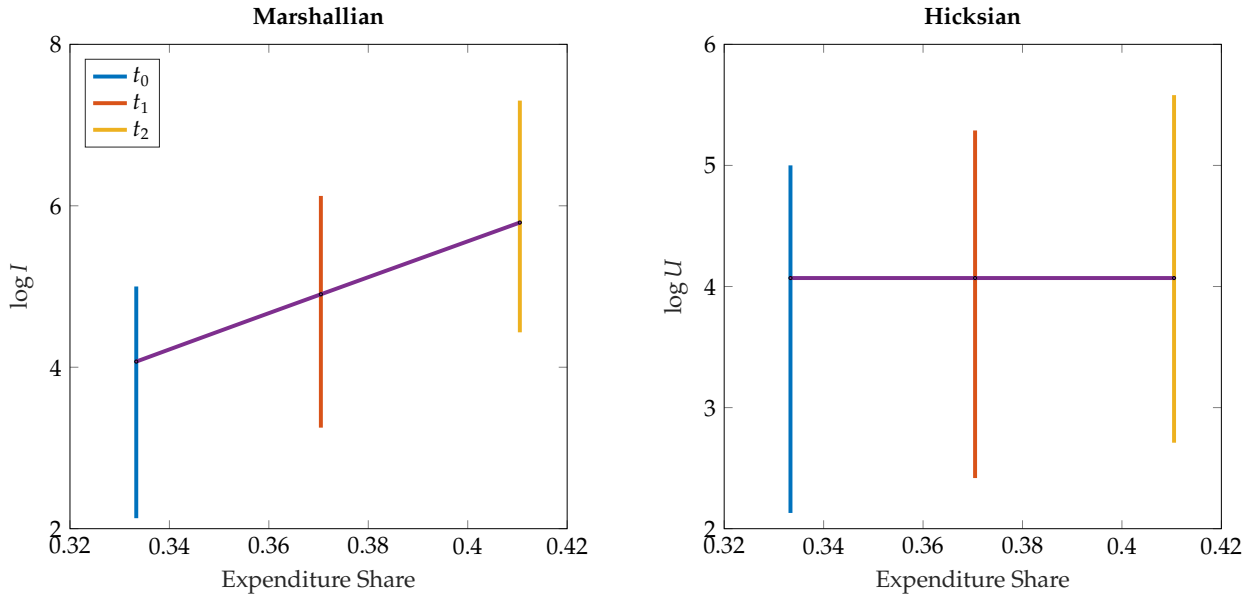
To see this procedure graphically, consider the left panel of Figure 1a showing the expenditure share on some good against nominal income for three different points in time. The fact that the lines are downward sloping means that higher incomes are associated with lower expenditure shares on the good. 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 for the



(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.

good, $B_i(u^{-1}(\bar{u}, s), s)$. 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 (5).

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, Marshallian and Hicksian budget shares coincide. Therefore, we can construct the money-metric using a price index based on Marshallian budget shares by good.

The iterative procedure that we describe is useful for building intuition. However, one can also find a fixed point by solving directly the system of equations. That is, we replace $u(I, t_{n-1})$ on the right hand side of equation (6) with $u(I, t_n)$. A simple way to find such a fixed point is to start with (5) and (6) and call the resulting money-metric utility $u^0(I, t)$. Then apply (5) and (6) again but this time use $u^0(I, t_n)$ in place of $u(I, t_{n-1})$ in (6) and call the resulting money-metric utility $u^1(I, t)$. Iterate on this until convergence. When we use this refined algorithm in our artificial data in Section 3.3, we obtain even smaller errors (by three orders of magnitude). The empirical results in Section 4 using real-world are roughly unchanged. Since the iterative procedure already has quite small errors, we do not always show these alternative results.

3.3 Quantitative Illustration with 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).¹¹ The expenditure function is

$$e(\mathbf{p}, u) = \left(\sum_i \omega_i (u^{\varepsilon_i} p_i)^{1-\gamma(u)} \right)^{\frac{1}{1-\gamma(u)}}. \quad (7)$$

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

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

and Marshallian budget shares are $B(I, t) = \mathbf{b}(\mathbf{p}_t, u)$ where u solves $I = e(\mathbf{p}_t, u)$. Income elasticities can vary across goods and the elasticity of substitution can vary across indifference curves (but is constant along any indifference curve, as under standard CES). As

¹¹See also Hanoch (1975), Comin et al. (2021), and Matsuyama (2019) for more information on these preferences. In Appendix E we consider another example with the Almost Ideal Demand System (AIDS) of Deaton and Muellbauer (1980).

shown by Baqaee and Burstein (2021), the money-metric function for t_0 reference prices is

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

where v is the solution to $I = e(\mathbf{p}_t, v)$. To evaluate the accuracy of our algorithm, we compare this exact 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 households facing a common price vector over forty years. 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 the highest inflation rate. Figure 2 plots the paths of prices and incomes in our numerical example.




Figure 2: Exogenous price and income paths for the artificial examples.

We consider two parameterizations. The first case allows for income effects but the elasticity of substitution does not depend on u . We follow Comin et al. (2021), and set $\varepsilon_1 = 0.2$, $\varepsilon_2 = 1$, $\varepsilon_3 = 1.65$, and $\gamma = 0.25$. 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 $\gamma(u) = 10 - 2 \log u$, with the intercept value ensuring that elasticities of substitution remain higher than unity. The share parameter ω is calibrated separately

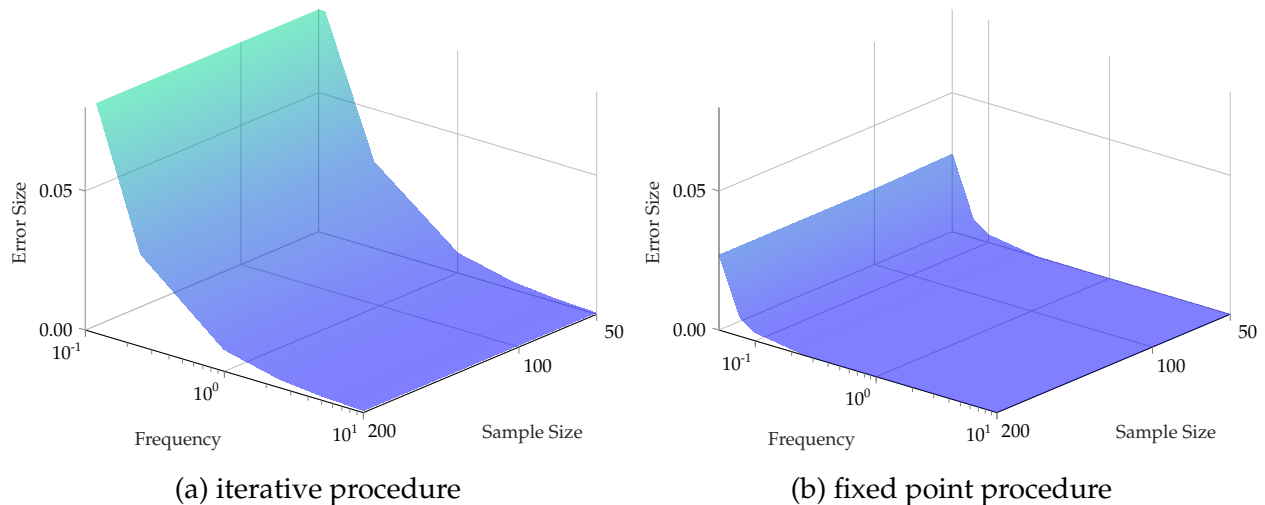


Figure 3: Maximum error as a function of the frequency of observation and number of households in the sample, holding the path of price and income changes constant. This figure uses non-homothetic CES preferences with variable γ . Our baseline calibration is annual frequency corresponding to a value 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 left panel uses the iterative procedure outlined in Section 3.2 and the right panel iterates on the iterative procedure until convergence to a fixed point.

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).

Since this data is artificial, we are only interested in the errors in our procedure compared to the exact solution. However, for completeness, Figure A.1 in Appendix A shows the difference between the money-metric and real consumption calculated using aggregate inflation (as defined in Section 4) for different income levels in this artificial economy.

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 in the final period. Under both parameterizations, constant and variable γ , the error is very small. For example, with 100 households and annual data, the maximum error in the final period is 0.0078 and 0.0044 for variable and constant γ preferences. This is equivalent to roughly two thirds of 1% of income. If instead of using the iterative procedure in Section 3.2, we solve the fixed point problem, 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. Figure 3 shows how this error varies as we vary the number of households and the frequency of observations using the variable γ non-homothetic specification. 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.

3.4 Discussion

In practice, data is imperfect and noisy. For example, 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. In Section 5, we extend our method to allow for unobserved changes in quality of some goods under stronger assumptions. However, if there are unobservable shocks to preferences themselves that are not idiosyncratic (i.e. taste shocks that 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 UK Application

In this section, we apply our algorithm to long-run cross-sectional household data. Our goal is to compare welfare as measured by the money-metric with real consumption. We define *real consumption* consistently with how it is constructed by statistical agencies in the

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.

national accounts: nominal expenditures deflated by a chain-weighted price index that reflects observed aggregate budget shares.¹⁴

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 2017.¹⁵ The UK Family Expenditure Survey was also used in Blundell 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 measure prices using 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 17 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.

We pool all households in our sample and assume that they have the same stable preference relation over the 17 categories of goods and services for which we have price data. To investigate the validity of this assumption, we can split the sample by observable characteristics as discussed in Section 3.4. We provide examples using marital status and age in Appendix A. This added flexibility comes at the expense of shrinking the boundaries over which the money-metric can be computed, since households with different characteristics (e.g. married and unmarried households) cannot be matched to one another through time. We do not find marked differences in the money-metric function by age or marital status, so these results are relegated to the appendix.

4.1 Mapping Data to the Model

Our procedure requires the expenditures I and the budget shares $\mathbf{B}(I, t)$ at time t across all goods. To deal with measurement error and idiosyncratic noise, we fit a smooth curve

¹⁴The analog to real consumption in our theoretical model is $\log RC(I, t) = \log I - \int_{t_0}^t \sum_{i=1}^N \bar{B}_i(t) \frac{d \log p_{is}}{ds} ds$, where $\bar{B}_i(t)$ is the income-weighted average budget share of good i in period t . Note that the price deflator is common for all households. If preferences are homothetic, then real consumption for every household coincides with money-metric utility since budget shares are the same for all income levels at each point in time.

¹⁵Aggregate 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.

to the budget share of each good i at time t as a function of income. We use these curves as $B(I, t)$. More precisely, 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

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

where i is the good, h is the household, and t is the time period. The estimated regression line gives us $B(I, t)$.¹⁶ This regression is the only source of sampling uncertainty in our exercise. We can 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 fairly tight. This is due to the law of large numbers, since the money-metric combines many Engel curve estimates. For this reason, and to make the figures less cluttered, when we present our results, we do not report the bootstrapped standard errors.

We calculate the money-metric utility for 1974 base prices by applying 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 $b(p_s, u(I, t))$. That is, for each income level I at time t , we must be able to find consumers at $s < t$ who were on the same indifference curve as the one delivered by I at time t . The left panel of Figure 4 illustrates how households in 2017 are matched with households in 1974 in order to estimate $b(p_{1974}, u(I, 2017))$. For example, households in the 50th percentile of income in 2017 are matched with households in the 78th 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 (6). The right panel of Figure 4 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

¹⁶We also estimated $B_i(I, t)$ using locally weighted scatterplot smoothing (LOWESS) and obtained very similar results.

indifference curve as the richest 97th percentile of households in 2017.




Figure 4: 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 dashed diagonal line is the 45 degree line. The vertical dotted lines are the boundaries for households that can be matched. 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.

4.2 Empirical Results

The blue line in Figure 5 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 income in 1974 if all households faced the same effective inflation rate, as given by the chain-weighted aggregate inflation rate. When the red line is above the blue line, this means that real consumption 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 consumption 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 implied by using an aggregate price deflator calculated as in 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: This figure plots $e(p_{1974}, v(p_{2017}, I_{2017}))$ and real consumption using aggregate chain-weighted inflation between 1974 to 2017 (annualized pounds, log scale). This figure converts income in 1974 into equivalent income in 2017 and vice versa.




Figure 6: The left panel shows the log difference between real consumption 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 4. The right panel is a histogram (using household weights) of money-metric $e(p_{1974}, v(p_{2017}, I_{2017}))$ and real consumption using aggregate chain-weighted inflation (annualized pounds, log scale). The distributions are truncated at the upper and lower bounds of Figure 4.

Figure 6 displays the log difference between the red and blue lines in Figure 5. As expected, the difference is positive for poor households, meaning that real consumption calculated using aggregate inflation is upward biased, and negative for rich households, meaning that real consumption is downward biased. The size of the bias is around 20% for

the poorest households. This means that over the 43 year sample, annual inflation rates calculated as in the official statistics understate the welfare-relevant inflation implied by the money-metric 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 6, that inequality across households is larger based on the money-metric than based on real consumption.

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.

5 Incorporating Unobserved Price Changes

In this section, we extend our methodology to allow for the possibility of missing or unreliably measured price changes. This may occur because the infrastructure for collecting comprehensive price data is absent, as in developing country contexts, or because changes in some prices are inherently difficult to measure, for example those of services and new goods. We present our theoretical results in Section 5.1 and provide an empirical illustration in Section 5.2.

5.1 Theoretical Results

To compute welfare without data on some prices, we must make stronger assumptions about preferences.

Assumption 1 (Indirect Separability). Partition the set of goods into X and Y . Suppose that preferences are *indirectly separable* in the sense that the expenditure function can be written as

$$e(\mathbf{p}, u) = e(e^X(\mathbf{p}^X, u), e^Y(\mathbf{p}^Y, u), u), \quad (8)$$

where \mathbf{p}^X and \mathbf{p}^Y are vectors of prices in X and Y , and e^X and e^Y are non-decreasing in and homogeneous of degree one in prices.

We assume that prices and expenditure shares of goods in X are observed, but prices and expenditure shares in Y are unobserved. Assumption 1 does not restrict cross-price elasticities for goods in X or Y but does restrict cross-price effects between X and Y . Let σ_{ij} be the Hicksian cross-price elasticity of demand between $i \neq j$. When Assumption 1 holds, then for any $i \in X$ and $j, k \in Y$, we have

$$\frac{\sigma_{ij}}{\sigma_{ik}} = \frac{b_j}{b_k}.$$

That is, the ratio of the cross-price elasticity of i with respect to j and k depends only the budget share of j and k . By symmetry, the same holds if we swap the role of X and Y . As an example, a CES aggregator is (additively) separable in every partition of its arguments.

Denote the Hicksian budget share of X goods by

$$b_X = \sum_{i \in X} b_i(\mathbf{p}, u) = b_X(e^X(\mathbf{p}^X, u)/e^Y(\mathbf{p}^Y, u), u),$$

where the second equality uses Assumption 1 and the fact that e is homogenous of degree one in prices. Hence, the budget share on X goods is pinned down, for a fixed u , by a single relative price of the X and Y bundles $e^X(\mathbf{p}^X, u)/e^Y(\mathbf{p}^Y, u)$. This implies that raising all prices in X by the same amount, holding utility constant, changes the share of spending on X by

$$\sum_{i \in X} \frac{\partial \log b_X}{\partial \log p_i} = (1 - b_X)(1 - \sigma(\mathbf{p}, u)),$$

for some scalar-valued function $\sigma(\mathbf{p}, u)$. We can think of $\sigma(\mathbf{p}, u)$ as the (compensated) elasticity of substitution between X and Y goods. By Shephard's lemma, we can also consider how the compensated budget share on X changes if a single price $i \in X$ changes:

$$\frac{\partial b_X}{\partial \log p_i} = (1 - b_X)b_i(1 - \sigma(\mathbf{p}, u)), \quad (i \in X).$$

This elasticity of substitution is disciplined by the curvature of the upper-nest of the expenditure function

$$\sigma(\mathbf{p}, u) = 1 - \frac{1}{(1 - b_X)b_X} \frac{\partial^2 \log e}{(\partial \log e^X)^2}.$$

In general, $\sigma(\mathbf{p}, u)$ depends on the entire vector of prices, some of which are unobserved. The following assumption ensures that σ can be expressed as a function of b_X and u , rather than as a function of all prices.

Assumption 2 (Monotone Budget Share). Suppose that $b_X(\mathbf{p}, u)$ is strictly monotone in the price of some $i \in X$.

This assumption is different to the monotonicity assumption of budget shares in u required in Atkin et al. (2020). They require that some budget share be strictly monotone in income (since they infer money-metric utility by inverting budget shares). Assumption 2 allows every budget share to be non-monotone (or even invariant) in income, but it requires that the compensated budget share of X be monotone in at least one price (unitary price elasticities for all goods is not allowed).

Lemma 1. *Assumption 2 implies that we can write $\sigma(\mathbf{p}, u)$ as $\sigma(b_X(\mathbf{p}, u), u)$. By abusing notation, we denote this function by $\sigma(b_X, u)$.*

Lemma 1 allows us to express the elasticity of substitution σ as a function of two scalars: utility and the compensated budget share of X goods.

Denote the *relative* Marshallian and Hicksian expenditure share on $i \in X$ by

$$B_{Xi}(I, t) = \frac{B_i(I, t)}{B_X(I, t)}, \quad \text{and} \quad b_{Xi}(\mathbf{p}, u) = \frac{b_i(\mathbf{p}, u)}{b_X(\mathbf{p}, u)}.$$

The following proposition extends Proposition 2 to account for unmeasured prices.

Proposition 3 (Money-Metric with Missing Prices). *Under Assumptions 1 and 2, the money-metric $u(I, t)$ solves the following integral equation*

$$\log u(I, t) = \log I - \int_{t_0}^t \sum_{i \in X} b_{Xi}(\mathbf{p}_s, u(I, t)) \frac{d \log p_{is}}{ds} ds - \int_{t_0}^t \frac{d \log b_X(\mathbf{p}_s, u(I, t)) / ds}{\sigma(b_X(\mathbf{p}_s, u(I, t)), u(I, t)) - 1} ds, \quad (9)$$

where

$$b_{Xi}(\mathbf{p}_s, u(I, t)) = B_{Xi}(u^{-1}(u(I, t), s), s), \quad b_X(\mathbf{p}_s, u(I, t)) = B_X(u^{-1}(u(I, t), s), s).$$

If we know the shape of the function $\sigma(b_X, u)$, Proposition 3 can be used to obtain the money-metric utility function using a similar algorithm to the one in Section 3.2. Proposition 3 is a consequence of Proposition 2. To derive it, we use changes in the compensated budget share of X goods, $d \log b_X(\mathbf{p}_s, u(I, t)) / ds$, to infer the Hicksian-budget-share-weighted changes in prices for the unobserved goods $\sum_{i \in Y} b_i(\mathbf{p}_s, u(I, t)) d \log p_{is} / ds$ given the elasticity of substitution $\sigma(b_X, u)$. Plugging this into (3) yields Proposition 3.

Compared to Proposition 2, the fixed point in Proposition 3 has some additional terms. First, the compensated elasticity of substitution $\sigma(b_X, u(I, t))$ on the right-hand side

depends on $u(I, t)$, and since $u(I, t)$ depends on the compensated elasticity of substitution, there is a fixed-point in this term. Second, the changes in the expenditure share on X goods, $d \log b_X(\mathbf{p}_s, u(I, t))/ds$, are compensated. To compute these changes, we must use the money-metric utility function, $u(I, t)$, to match households on the same indifference curve through time and use changes in the expenditure shares of matched households over time. Hence, there is also fixed-point in this term.

To better understand Proposition 3, it helps to consider the homothetic special case.

Example 1 (Homothetic preferences). Suppose that preferences are homothetic. In this case, Proposition 3 simplifies to

$$\log u(I, t) = \log I - \int_{t_0}^t \sum_{i \in X} B_{Xi}(\mathbf{p}_s) \frac{d \log p_{is}}{ds} ds - \int_{t_0}^t \frac{d \log B_X/ds}{\sigma(B_X(s)) - 1}. \quad (10)$$

When preferences are homothetic, there is no longer a fixed-point problem since budget shares and elasticities of substitution do not depend on utility. If we also assume that the upper-nest $e(x, y, u)$ is CES, then $\sigma(b_X(\mathbf{p}_s))$ is a constant and we get

$$\log u(I, t) = \log I - \int_{t_0}^t \sum_{i \in X} b_{Xi}(\mathbf{p}_s) \frac{d \log p_{is}}{ds} ds - \frac{\log b_X(\mathbf{p}_t) - \log b_X(\mathbf{p}_{t_0})}{\sigma - 1}. \quad (11)$$

Equation (11) is a version of the popular Feenstra (1994) formula.¹⁷ This formula is near-ubiquitous in the macroeconomics and trade literatures for adjusting price indices to account for missing price changes (typically those of new goods). Relative to Feenstra (1994), Proposition 3 allows the elasticity of substitution to vary as a function of prices, allows for non-homotheticities, and allows unrestricted non-parametric preferences among the X and Y goods.

Relative to the homothetic special case in (10), the additional complication in (9) is that changes in the budget share of X and the elasticity of substitution must both be compensated. To see the issue, restate (9) using Marshallian budget shares as

$$\log u(I, t) = \log I - \int_{t_0}^t \sum_{i \in X} B_{Xi}(I_s^*, s) \frac{d \log p_{is}}{ds} ds - \int_{t_0}^t \frac{d \log B_X(I_s^*, s)/ds}{\sigma(B_X(I_s^*, s), u(I, t)) - 1} ds,$$

¹⁷The only (relatively inconsequential) difference between (11) and Feenstra (1994) is the assumption that e^X and e^Y also be homothetic CES aggregators.

where I_s^* is implicitly defined by $u(I_s^*, s) = u(I, t)$. The necessary compensation satisfies

$$\underbrace{\frac{d \log B_X(I_s^*, s)}{ds}}_{\Delta \text{ budget share due to prices holding fixed utility}} = \underbrace{\frac{\partial \log B_X(I_s^*, s)}{\partial s}}_{\Delta \text{ budget share due to prices holding fixed income}} + \underbrace{\frac{\partial \log B_X(I_s^*, s)}{\partial I} \frac{dI_s^*}{ds}}_{\Delta \text{ budget share due to compensating income}}.$$

When preferences are homothetic, the second term disappears and there is no fixed point problem.

We provide an example of non-homothetic separable preferences that satisfy Assumptions 1 and 2 below.

Example 2 (Indirect Addilog). Suppose that the expenditure function $e(\mathbf{p}, u)$ is indirectly separable and is implicitly defined as

$$f(u) = \frac{\omega_X}{\sigma_X - 1} \left(\frac{e(\mathbf{p}, u)}{e^X(\mathbf{p}^X, u)} \right)^{\sigma_X - 1} + \frac{\omega_Y}{\sigma_Y - 1} \left(\frac{e(\mathbf{p}, u)}{e^Y(\mathbf{p}^Y, u)} \right)^{\sigma_Y - 1}, \quad (12)$$

where $f(\cdot)$ is an increasing function. CES is the special case where $\sigma_X = \sigma_Y$. For this demand system, the compensated elasticity of substitution between X and Y is

$$\sigma(b_X(\mathbf{p}, u), u) = b_X(\mathbf{p}, u)\sigma_Y + (1 - b_X(\mathbf{p}, u))\sigma_X,$$

which varies both as a function of utility and as a function of prices.

We now show that $\sigma(b_X(\mathbf{p}_s, u), u)$, the unknown term required to apply Proposition 3, can be expressed non-parametrically in terms of elasticities that are estimable using only data on prices in X . This is an important result as it demonstrates that, in general, recovering $\sigma(b_X(\mathbf{p}_s, u), u)$ does not require data on unobserved prices.

Proposition 4 (Identifying Substitution Elasticity of X and Y). *Let $\eta_i(I, t) - 1 = \partial \log B_i(I, t) / \partial \log I$ be the income elasticity of demand for each $i \in X$ at time t . Let $1 - \epsilon_{Xk}(I, t) = \partial \log B_X(I, t) / \partial \log p_k$ be the uncompensated elasticity of the budget share of X with respect to the price of any good $k \in X$ at time t . Then, we have*

$$\sigma(b_X(\mathbf{p}_s, u(I, t)), u(I, t)) = 1 - \frac{1 - \epsilon_{Xk}(I_s^*, s) + B_k(I_s^*, s) \sum_{i \in X} (\eta_i(I_s^*, s) - 1) B_{Xi}(I_s^*, s)}{(1 - B_X(I_s^*, s)) B_{Xk}(I_s^*, s)},$$

where I_s^* is implicitly defined by $u(I_s^*, s) = u(I, t)$.

Proposition 4 shows that if we know income elasticities for all goods in X and the uncompensated price elasticity ϵ_{Xk} for one good $k \in X$ in each period, then we can recover

the relevant elasticity of substitution and apply Proposition 3. Estimating the income elasticities, η_i for $i \in X$, is relatively straightforward since we simply need to fit a curve that relates the budget share of i to income in each period. Estimating the price elasticity ϵ_{Xk} is more challenging, but we only require a single elasticity per income group and period. This is reasonably straightforward to estimate if one of the observed prices is exogenous. Of course, if one puts more structure on the demand system, then even less information is required. We provide one example below.

Example 3 (Non-homothetic CES). Suppose that $e(e^X(\mathbf{p}^X, u), e^Y(\mathbf{p}^Y, u), u)$ in (8) takes the form

$$e(x, y, u) = \left(\omega_X u^{\xi_X} x^{1-\gamma(u)} + \omega_Y u^{\xi_Y} y^{1-\gamma(u)} \right)^{\frac{1}{1-\gamma(u)}}.$$

In this case, $\sigma(b_X, u) = \gamma(u)$. That is, $\sigma(b_X, u)$ varies as a function of utility but not as a function of relative prices. According to Proposition 4, the function $\sigma(\cdot)$ is determined by the following expression:

$$\sigma(I) = 1 - \frac{1 - \epsilon_k(I, t_0) + B_k(I, t_0) \sum_{i \in X} (\eta_i(I, t_0) - 1) B_{Xi}(I, t_0)}{(1 - B_X(I, t_0)) B_{Xk}(I, t_0)}, \quad (13)$$

where $1 - \epsilon_k(I, t_0)$ is the uncompensated elasticity of B_X with respect to the price of k and $\eta_k(I, t_0)$ is the income elasticity for $k \in X$ at time t_0 for households with income I . Since σ is not a function of relative prices, Proposition 4 needs to be applied only in the initial period, t_0 , to recover the shape of the σ function.

In writing (13), we assume that $\epsilon_k(I, t_0)$ and $\eta_i(I, t_0)$ are known at t_0 . This is without loss of generality since Proposition 3 can be applied with time running forward $t > t_0$ and backward $t < t_0$. Furthermore, once we apply Proposition 3 to obtain the money metric with t_0 reference prices, we can easily obtain the money-metric at $t_s \in [t_0, T]$ base prices, as described in footnote 9.

Relation to previous literature. When price data is unavailable or unreliable, a strand of the literature, building on Hamilton (2001) and Costa (2001), estimates changes in welfare by inverting Engel curves. The frontier in this literature is Atkin et al. (2020), who show how to identify welfare changes assuming that preferences are quasi-separable between the measured and unmeasured goods. Their procedure requires that some expenditure share be strictly monotone in income (i.e. homothetic preferences are ruled out). To apply their method, one also needs to estimate a Hicksian demand sub-system for the set of goods where prices are measured. In their applications, they either rely on first-order approximations or use a CES sub-system.

In contrast, we make stronger assumptions about preferences (separability rather than quasi-separability). In exchange, we do not require that expenditure shares be strictly monotone in income (i.e. homothetic preferences are allowed). More importantly, our approach only requires a single Marshallian price elasticity as a function of income in each period (rather than a Hicksian system). Given estimates of this elasticity, we can non-parametrically and non-linearly back out the elasticity of substitution between the measured and unmeasured goods and use this to non-linearly solve for welfare changes.

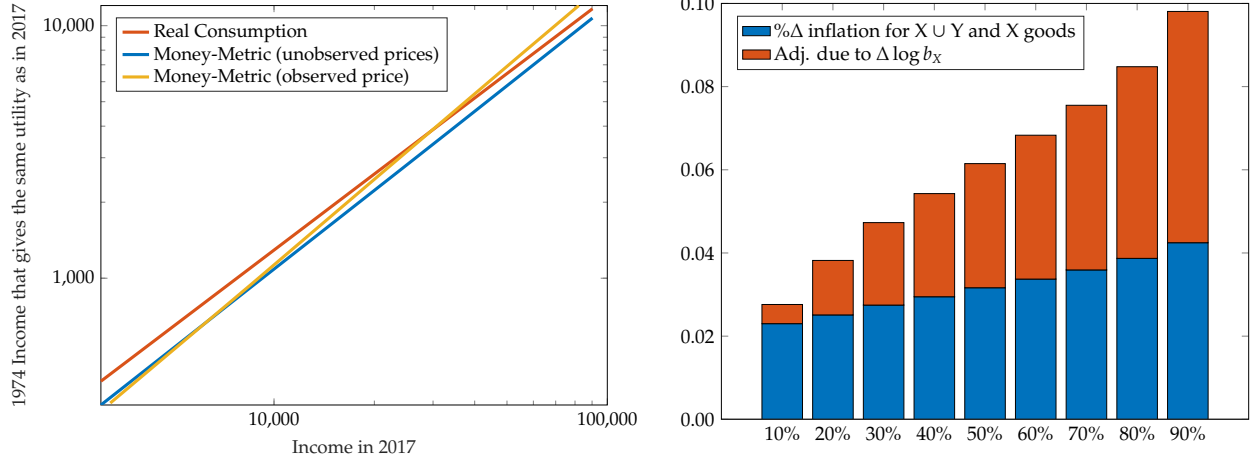
5.2 Empirical Illustration

As an illustration, we apply Proposition 3 to the UK data that we used in Section 4. Since service prices are difficult to measure, as a test case, we partition the consumption bundle and assume that prices for “recreation & culture”, “education”, “accommodation”, and “leisure goods & services” are not reliably observed. These are the Y goods, which in our data account for roughly 10–30% percent of spending. We assume that prices for all other categories of spending are measured accurately. These other categories are the X goods. We impose Assumptions 1 and 2 and apply Proposition 3.

To apply Proposition 3, we calibrate the elasticity of substitution $\sigma(b_X, u)$ between X and Y . For simplicity, as in Comin et al. (2021), we assume that σ is constant in both the time-series and the cross-section. Following their estimates, we set $\sigma = 0.25$. As explained above, Proposition 3 can also be applied when σ varies in both the time series (i.e. prices) and the cross-section (i.e. income). However, estimating the $\sigma(b_X, u)$ function is beyond the scope of this paper. Our more limited goal is to show how knowledge of this elasticity can be used to back-out the money-metric utility function in the absence of comprehensive price information.

Figure 7 compares the results from using Propositions 2 and 3. The left panel displays the resulting money-metric functions as well as real consumption. For low-income households, the two money-metrics are quite similar and both are below real consumption. However, for households with high incomes, the money-metric calculated using Proposition 3 is lower than the one calculated using Proposition 2. For high-income households, the money-metric implied by Proposition 3 is closer to real consumption instead. The fact that the blue line is lower than the yellow line for rich households suggests that, for these households, prices in Y have risen more than the official price data suggest

The right panel of Figure 7 shows the percent difference between the two money-metrics for different income deciles as well as the breakdown of the difference into two terms. The first is the difference between overall inflation and inflation for X goods



(a) Money-metric $e(\mathbf{p}_{1974}, v(\mathbf{p}_{2017}, I))$ and real consumption as a function of income in 2017. (b) Percent difference in money-metric values for different percentiles of the income distribution.

Figure 7: These figures compare money-metric values when using Propositions 2 and 3. The left panel plots the money-metric and real consumption values in annualized pounds on a log scale. The right panel shows the percent difference in money-metric values using Propositions 2 and 3. This difference is decomposed into (i) the difference between overall inflation and inflation for X goods and (ii) the adjustment due to changes in the compensated expenditures on X goods.

implied by the two methods:

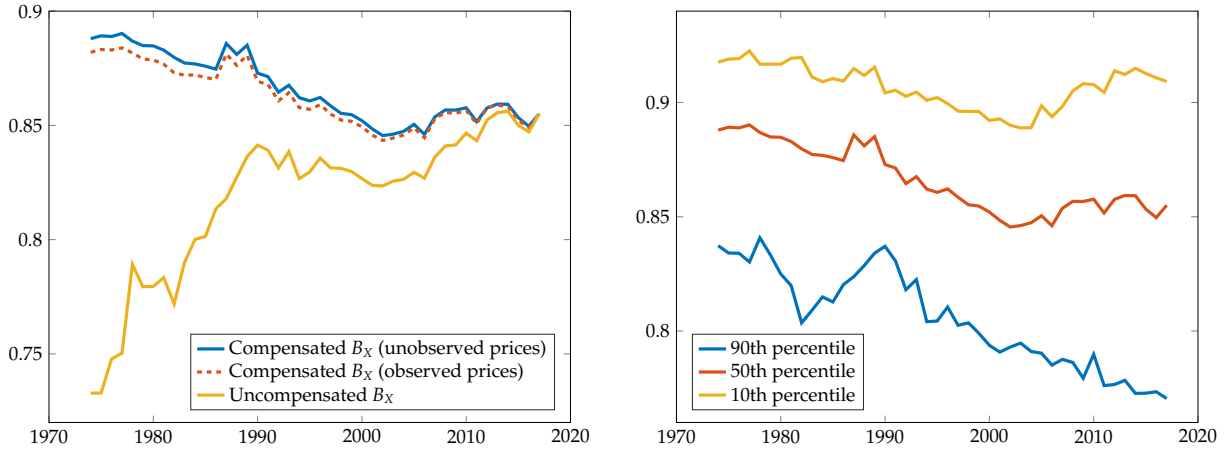
$$\frac{\int_{1974}^{2017} \sum_{i=1}^N b_i(\mathbf{p}_s, u) (d \log p_{is} / ds) ds - \int_{1974}^{2017} \sum_{i \in X} b_{Xi}(\mathbf{p}_s, u) (d \log p_{is} / ds) ds}{\int_{1974}^{2017} \sum_{i=1}^N b_i(\mathbf{p}_s, u) (d \log p_{is} / ds) ds}.$$

These are the blue bar graphs in the right panel of Figure 7. The remainder is the adjustment due to changes in the expenditure share of X goods, similar to the Feenstra (1994) adjustment:

$$\frac{\frac{1}{\sigma-1} \int_{1974}^{2017} (d \log b_X(\mathbf{p}_s, u) / ds) ds}{\int_{1974}^{2017} \sum_{i=1}^N b_i(\mathbf{p}_s, u) (d \log p_{is} / ds) ds}. \quad (14)$$

These are the orange bar graphs in the right panel of Figure 7. This decomposition shows that inflation among X goods has tended to be higher than among all goods by roughly the same amount (around 2% points) for all income deciles. However, the change in compensated expenditures on X goods has been very different. Compensated expenditures on X goods have been falling much more quickly for rich households than poor.

To better understand why the money-metric utility is different for rich households,



(a) Budget shares of X for the household with the median level of income in 2017. (b) Compensated budget share of X according to Proposition 3 for households with different 2017 income levels.

Figure 8: Compensated and uncompensated budget share of X for households of differing income levels plotted against time.

we investigate how compensated expenditures on X goods have been changing. Figure 8 plots changes in expenditures on X over time for different households. The left panel displays how compensated and uncompensated expenditures on X goods have changed as a function of time for a household with the median level of 2017 income. The blue line shows how compensated expenditures have changed according to Proposition 3 and the dotted red line shows how they have changed according to Proposition 2. Compensated expenditure have changed by roughly the same amount in both cases, falling slightly from around 88% to around 85%. Compensated expenditures capture pure substitution effects over time — changes in expenditures due to changes in relative prices. The yellow line in the left panel shows how uncompensated, or Marshallian, expenditures on X goods have changed for a household with the median level of income in 2017. These expenditures have been rising rapidly from around 70% to around 85%. That is, over time, the household whose nominal income is equal to the median in 2017 has spent significantly more on X goods over time. This is because inflation makes this household poorer over time and poor households tend to spend more of their income on X goods (which exclude services like recreation and education).

The right panel of Figure 8 shows the compensated budget share on X goods for households at three different points in the income distribution: the 10th, 50th and 90th percentiles in 2017. For poor households, there has been almost no change on expenditures on X goods. This explains why the adjustment term in (14) is small for these households.

For the median household, there has been a modest decrease in the share of spending on X goods. Since X and Y are complements ($\sigma < 1$), this indicates that the relative price of Y goods has been rising relative to X goods for these households. Finally, for the richest households, there has been a fairly dramatic reduction in their spending on X goods from around 83% to around 76%. This suggests that for these households, the relative price of Y goods has been rising fairly rapidly compared to X goods. This explains why the adjustment term, (14), for these households is large and negative. It also explains why the money-metric according to Proposition 3 (the blue line in Figure 7) has a flatter slope than the money-metric calculated according to Proposition 2 (the yellow line in Figure 7).

6 Conclusion

In this paper, we propose a straightforward and intuitive approach to construct money-metric representations of utility using repeated cross-sectional data. Our method does not require any estimation when the data on prices is comprehensive, aside from cross-sectional interpolation of how budget shares vary with income.

If the data on prices is incomplete, the method can still be used, but stronger assumptions on preferences and knowledge of one Marshallian elasticity for each income level in each period is required. In both cases, the unifying idea is that money-metric utility must satisfy a fixed point equation in terms of observable variables.

Despite its advantages, our approach relies on the assumption that preferences are not systematically and unobservably heterogeneous in either the cross-section or the time-series 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.
- Caves, D. W., L. R. Christensen, and W. E. Diewert (1982). The economic theory of index numbers and the measurement of input, output, and productivity. *Econometrica: Journal of the Econometric Society*, 1393–1414.
- Comin, D., D. Lashkari, and M. Mestieri (2021). Structural change with long-run income and price effects. *Econometrica* 89(1), 311–374.
- Costa, D. L. (2001). Estimating real income in the united states from 1888 to 1994: Correcting cpi bias using engel curves. *Journal of political economy* 109(6), 1288–1310.
- 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.
- Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *The American Economic Review*, 157–177.
- Feenstra, R. C. and M. B. Reinsdorf (2000). An exact price index for the almost ideal demand system. *Economics Letters* 66(2), 159–162.
- 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*.
- Oulton, N. (2012). How to measure living standards and productivity. *Review of Income and Wealth* 58(3), 424–456.
- 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




Figure A.1: Non-homotheticity bias: log difference between real consumption and the money-metric for artificial examples in Section 3.3. For poor households, the money metric is lower than real consumption calculated using aggregate inflation because the inflation rate is lower for income elastic goods.

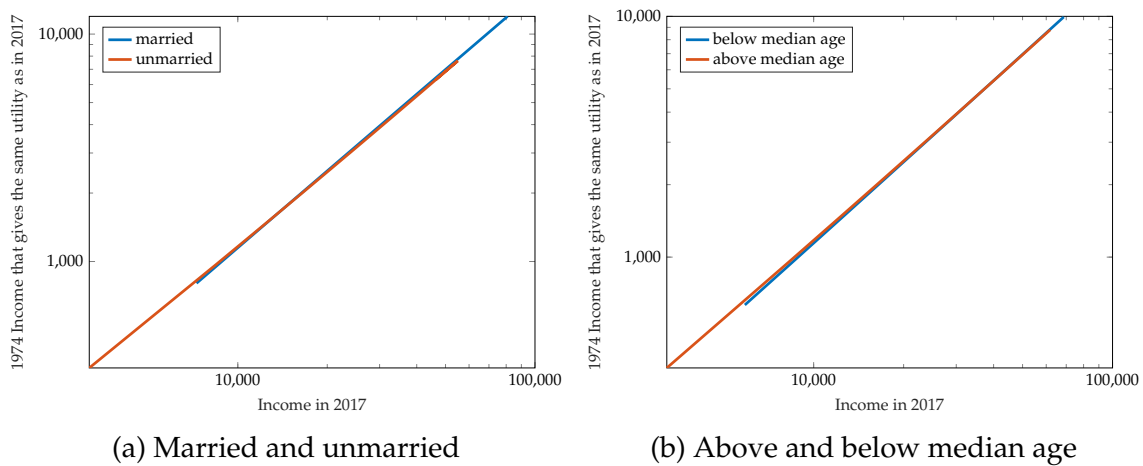


Figure A.2: Money-metric $e(\mathbf{p}_{1974}, v(\mathbf{p}_{2017}, I_{2017}))$ by household characteristic (annualized pounds, log scale) for the UK data in Section 4.

B Additional details of the UK data used in Section 4

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 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 consumption 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, “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 17 categories that are available for

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 4, the histograms in Figure 6, 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.

the entire period for both RPI and CPI. Table A.1 summarizes how we integrated the CPI and RPI baskets.

Integrated Categories	RPI	CPI
Bread & Cereals	Bread, Cereals and Biscits	Bread & cereals
Meat & Fish	Meat, Fish, Beef, Lamb and Pork	Meat & fish
	Poultry and Other meat	-
Milk & Eggs	Butter, Cheese and Eggs	Milk, cheese & eggs
	Fresh milk and Milk products	-
Oils & fats	Oils & fats	Oils & fats
Fruit	Fruit	Fruit
Vegetable	Potatoes and Other vegeables	Vegetables including potatoes & other tubers
Other food	Sweets & Chocolates	Food Products
	Other Foods	Suger, jam, honey, syrups, chocolate & confectonery
Non-Alcoholic Beverages	Tea and Soft drinks	Non-Alcoholic Beverages
	Coffee & other hot drinks	
Tobacco	Cigarettes & tobacco	Tobacco
Catering	Catering	Catering services
	Alcoholic drink	Alcoholic beverage
Household & Fuel	Housing except morgage interest	Housing, water and fuels
	Fuel & light	
	(-)Dwelling insurance & ground rent	
Clothing	Clothing & footwear	Clothing & footwear
Household Goods & Service	Household goods	Furniture and household equipment & routine repair of house
	domestic services	
Postage & Telecom	Postage	Communicaion
	Telephones & Telemessages	
Personal Goods & Service	Personal goods & services	Health
	Fees & subscriptions	Miscellaneous goods and service
	Dwelling insurance & ground rent	-
Transport	Motoring expenditure	Transport
	Fares & other travel costs	-
Leisure Goods & Service	Leisure goods	Recreation & culture
	Leisure services	Education
	-	Accomodation service

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 $F^{(i)} = \{I_s^i : p_t \cdot q_s(I_s^i) = I_t\}_{s \in T}$.

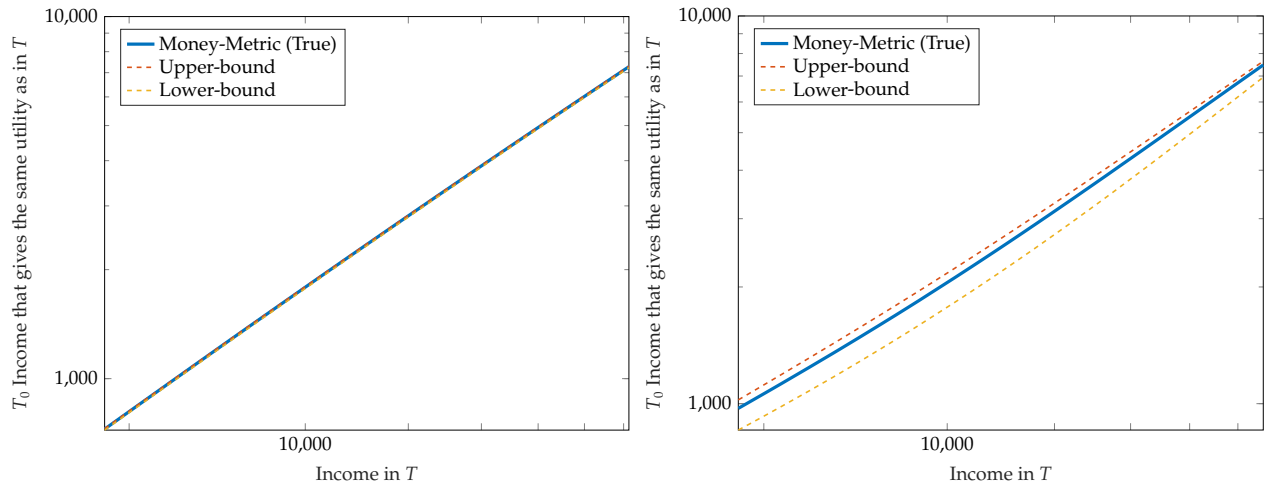
(2) Set $F^{(i+1)} = \{\max_{I_k \in F^{(i)}} \{I_s^{i+1} : I_k = p_k \cdot q_s(I_s^{i+1})\}\}_{s \in T}$.

(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 3.3. 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 4. 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 3.3. 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.

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

	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 3.3.

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.

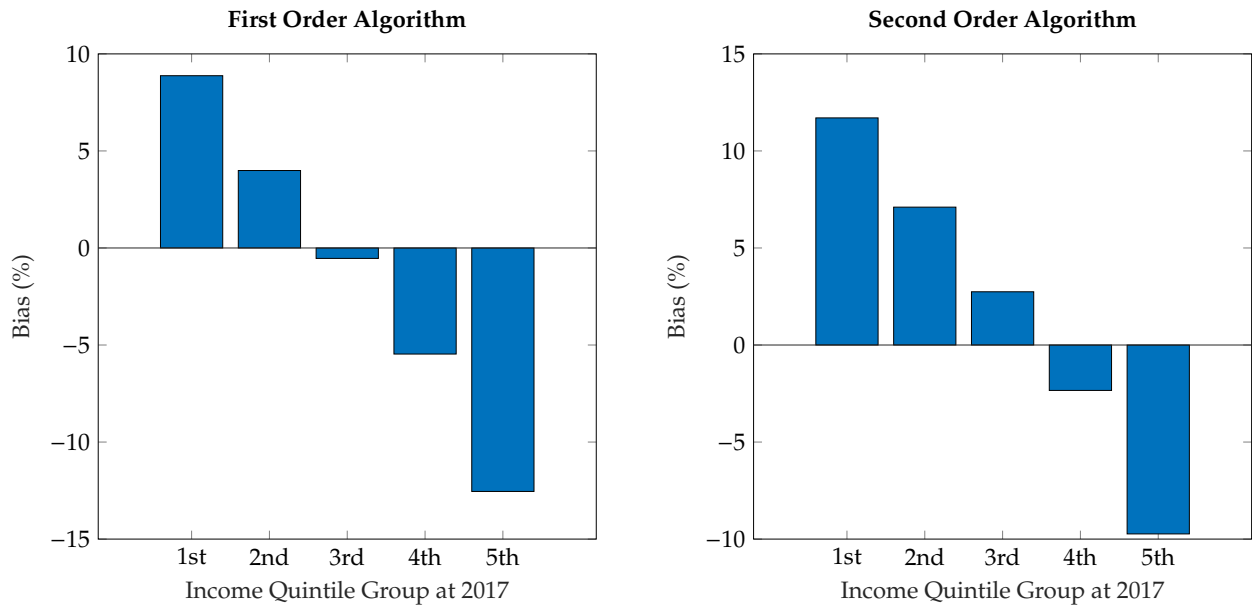


Figure A.5: Results of the first/second order algorithm of Jaravel and Lashkari (2022) with $K = 1$ to the aggregated (by quintile) UK household data: Log difference between real consumption and the money-metric




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 consumption and the money-metric

E Using artificial data from Almost Ideal Demand System

In this appendix, we redo our analysis of Section 3.3 using another popular form of non-homothetic preferences: the Almost Ideal Demand System (AIDS) due to Deaton and Muellbauer (1980). The expenditure function is

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

where $c(\mathbf{p})$ and $d(\mathbf{p})$ are given by:

$$c(\mathbf{p}) = \exp \left(\alpha_0 + \sum_{i=1}^I \alpha_i \log p_i + \frac{1}{2} \sum_{i=1}^I \sum_{j=1}^I \gamma_{ij} \log p_i \log p_j \right)$$

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

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 $b(\mathbf{p}, u)$ are

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

The money-metric function for t_0 reference prices is²¹

$$e(\mathbf{p}_0, v(\mathbf{p}, I)) = c(\mathbf{p}_0) \left(\frac{I}{c(\mathbf{p})} \right)^{\frac{d(\mathbf{p}_0)}{d(\mathbf{p})}}.$$

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 3.3. 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, $\max_h |\log u(I, T) - \log u(I, T)^{TRUE}|$, is 0.0011 when we use the iterative procedure and 7.0×10^{-7} when we use the fixed point procedure. For completeness, Figure A.7 presents the non-homotheticity bias in our artificial AIDS example.




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

²¹To obtain the expression for the money-metric, we use $e(\mathbf{p}_0, v(\mathbf{p}, I)) = e(\mathbf{p}_0, u)$, where $I = c(\mathbf{p}) u^{d(\mathbf{p})}$.

F Proofs

Proof of Proposition 1. By definition,

$$\begin{aligned}\log e(\mathbf{p}, v(\bar{\mathbf{p}}, \bar{I})) &= \log e(\bar{\mathbf{p}}, v(\bar{\mathbf{p}}, \bar{I})) + \log e(\mathbf{p}, v(\bar{\mathbf{p}}, \bar{I})) - \log e(\bar{\mathbf{p}}, v(\bar{\mathbf{p}}, \bar{I})) \\ &= \log \bar{I} + \log e(\mathbf{p}, v(\bar{\mathbf{p}}, \bar{I})) - \log e(\bar{\mathbf{p}}, v(\bar{\mathbf{p}}, \bar{I})).\end{aligned}$$

To finish, rewrite

$$\log e(\mathbf{p}, v(\bar{\mathbf{p}}, \bar{I})) - \log e(\bar{\mathbf{p}}, v(\bar{\mathbf{p}}, \bar{I})) = \int_{\bar{\mathbf{p}}}^{\mathbf{p}} \sum_{i \in N} \frac{\partial \log e(\xi, v(\bar{\mathbf{p}}, \bar{I}))}{\partial \log \xi_i} d \log \xi_i,$$

and use Shephard's lemma to express the price elasticity of the expenditure function in terms of budget shares, and obtain (1). To obtain (2), switch \mathbf{p} and $\bar{\mathbf{p}}$ as well as I and \bar{I} . ■

Proof of Proposition 2. This follows immediately from the definition of $u^{-1}(\cdot, s)$ which maps incomes at t_0 to equivalent income at time s . Hence, for some amount of t_0 income, say $u(I, t)$, the equivalent income at time s is $u^{-1}(u(I, t), s)$. The Marshallian budget share $B(u^{-1}(u(I, t), s), s)$ is just $b(u(I, t), s)$. ■

Proof of Lemma 1. Start by assuming that b_X is increasing in p_k for some $k \in X$. Then, we have that $\frac{\partial \log b_X}{\partial \log p_k} = (1 - b_X)(1 - \sigma)b_{Xk} > 0$. That is, $\sigma(\mathbf{p}, u) < 1$. This implies that $b_X(e^X(\mathbf{p}^X, u)/e^Y(\mathbf{p}^Y, u), u)$ is increasing in its first argument. In other words, we can write $e^X(\mathbf{p}^X, u)/e^Y(\mathbf{p}^Y, u) = f(b_X, u)$. Hence, we can write $\sigma(\mathbf{p}, u) = \sigma(e^X(\mathbf{p}^X, u)/e^Y(\mathbf{p}^Y, u), u) = \sigma(f(b_X, u), u)$ as needed. A symmetric argument applies when b_X is decreasing in p_k for some $k \in X$. ■

Proof of Proposition 3. By Euler's theorem of homogeneous functions, we know that

$$\frac{\partial \log e}{\partial \log e^X} + \frac{\partial \log e}{\partial \log e^Y} = 1.$$

Differentiating this identity with respect to e^X and e^Y yields the following equations

$$\frac{\partial^2 \log e}{(\partial \log e^X)^2} = -\frac{\partial^2 \log e}{\partial \log e^X \partial \log e^Y} = \frac{\partial^2 \log e}{(\partial \log e^Y)^2}.$$

Next, we know that

$$b_X = \sum_{i \in X} b_i = \sum_{i \in X} \frac{\partial \log e}{\partial \log e^X} \frac{\partial \log e^X}{\partial \log p_i} = \frac{\partial \log e}{\partial \log e^X} \sum_{i \in X} \frac{\partial \log e^X}{\partial \log p_i} = \frac{\partial \log e}{\partial \log e^X}$$

Hence, fixing utility, the total derivative of b_X with respect to prices is

$$\begin{aligned}
b_X d \log b_X &= \frac{\partial^2 \log e}{(\partial \log e^X)^2} \sum_{i \in X} \frac{\partial \log e^X}{\partial \log p_i} d \log p_i + \frac{\partial^2 \log e}{\partial \log e^Y \partial \log e^X} \sum_{i \in Y} \frac{\partial \log e^Y}{\partial \log p_i} d \log p_i \\
&= \frac{\partial^2 \log e}{(\partial \log e^X)^2} \left[\sum_{i \in X} \frac{\partial \log e^X}{\partial \log p_i} d \log p_i - \sum_{i \in Y} \frac{\partial \log e^Y}{\partial \log p_i} d \log p_i \right] \\
&= \frac{\partial^2 \log e}{(\partial \log e^X)^2} \left[\sum_{i \in X} b_{X_i} d \log p_i - \sum_{i \in Y} b_{Y_i} d \log p_i \right]
\end{aligned}$$

Using the fact that

$$\sigma(p, u) = 1 - \frac{1}{(1 - b_X) b_X} \frac{\partial^2 \log e}{(\partial \log e^X)^2},$$

we can rewrite this as

$$d \log b_X = (1 - b_X)(1 - \sigma) \left[\sum_{i \in X} b_{X_i} d \log p_i - \sum_{i \in Y} b_{Y_i} d \log p_i \right],$$

where we suppress the fact that σ is a function of prices and utility. Rearranging this gives

$$-\frac{d \log b_X}{1 - \sigma} + (1 - b_X) \sum_{i \in X} b_{X_i} d \log p_i + b_X \sum_{i \in X} b_{X_i} d \log p_i = \sum_{i \in X} b_i d \log p_i + \sum_{i \in Y} b_i d \log p_i,$$

or

$$-\frac{d \log b_X}{1 - \sigma} + \sum_{i \in X} b_{X_i} d \log p_i = \sum_{i \in X} b_i d \log p_i + \sum_{i \in Y} b_i d \log p_i.$$

Plug this back into Proposition 2 to get the desired result. It is important to note however that $d \log b_X$ in the expression above is the compensated change in the budget share of X. ■

Proof of Proposition 4. Consider a perturbation to p_k for $k \in X$ holding fixed utils:

$$\begin{aligned}
\frac{\partial \log b_X}{\partial \log p_k} &= \frac{1}{b_X} \frac{\partial}{\partial \log p_k} \left[\sum_{i \in X} \frac{\partial \log e}{\partial \log e^X} \frac{\partial \log e^X}{\partial \log p_i} \right] \\
&= \frac{1}{b_X} \frac{\partial}{\partial \log p_k} \left[\sum_{i \in X} \frac{\partial \log e}{\partial \log e^X} b_{X_i} \right] \\
&= \frac{1}{b_X} \left[\sum_{i \in X} \frac{\partial}{\partial \log p_k} \frac{\partial \log e}{\partial \log e^X} b_{X_i} + \sum_{i \in X} \frac{\partial \log e}{\partial \log e^X} \frac{\partial b_{X_i}}{\partial \log p_k} \right]
\end{aligned}$$

$$\begin{aligned}
&= \frac{1}{b_X} \left[\sum_{i \in X} \frac{\partial^2 \log e}{(\partial \log e^X)^2} b_{Xk} b_{Xi} + \sum_{i \in X} \frac{\partial \log e}{\partial \log e^X} \frac{\partial b_{Xi}}{\partial \log p_k} \right] \\
&= \frac{1}{b_X} \left[\sum_{i \in X} \frac{\partial^2 \log e}{(\partial \log e^X)^2} b_{Xk} b_{Xi} + \frac{\partial \log e}{\partial \log e^X} \frac{\partial \sum_{i \in X} b_{Xi}}{\partial \log p_k} \right] \\
&= \frac{1}{b_X} \frac{\partial^2 \log e}{(\partial \log e^X)^2} b_{Xk},
\end{aligned}$$

where the last line uses the fact that $\frac{\partial \sum_{i \in X} b_{Xi}}{\partial \log p_k} = 0$. Using the following relationship

$$\frac{\partial^2 \log e}{(\partial \log e^X)^2} = b_X \frac{\partial \log b_X}{\partial \log e^X} = b_X(1 - b_X)(1 - \sigma(\mathbf{p}, u)),$$

the compensated change in expenditures on X in response to a change in the price of $k \in X$ is given by

$$\frac{\partial \log b_X}{\partial \log p_k} = (1 - b_X)(1 - \sigma(\mathbf{p}, u))b_{Xk}.$$

On the other hand, the uncompensated (Marshallian) response is

$$\frac{\partial \log B_X}{\partial \log p_k} = 1 - \epsilon_{Xk} = (1 - b_X)(1 - \sigma(\mathbf{p}, u))b_k^X - \sum_{i \in X} (\eta_i - 1)b_k b_{Xi},$$

where the second equation is analogous to the Slutsky equation. Rearranging this for $\sigma(\mathbf{p}, u)$ yields the desired result

$$1 - \frac{(1 - \epsilon_{Xk}) + \sum_{i \in X} (\eta_i - 1)b_k b_i / b_X}{(1 - b_X)b_k / b_X} = \sigma(\mathbf{p}, u).$$

■