

Measuring Welfare by Matching Households across Time

David R. Baqaee

UCLA

Ariel Burstein

UCLA

Yasutaka Koike-Mori

UCLA*

August 2023

Abstract

The money metric utility function is an essential tool for calculating welfare-relevant growth and inflation. We show how to recover it from repeated cross-sectional data without making parametric assumptions about preferences. We do this by solving the following recursive problem. Given compensated demand, we construct money metric utility by integration. Given money metric utility, we construct compensated demand by matching households over time whose money metric utility value is the same. We illustrate our method using 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 in 1974 pounds for the poorest households by around half a percent per year and understates it by around a third of a percentage point per year for the richest households. We extend our method to allow for missing or mismeasured prices, assuming preferences are separable between goods with well-measured prices and the rest. We discuss how our results change if the prices of some service sectors are mismeasured. *JEL Codes:* D11, E31.

*Previously circulated as “A New Method for Measuring Welfare with Income Effects using Cross-Sectional Data” and “A Fixed Point Approach to Measuring Welfare”. We thank Andy Atkeson, John Asker, Natalie Bau, Hector Chade, Dave Donaldson, Pablo Fajgelbaum, Andy Neumeyer, Jon Vogel, and Pierre-Olivier Weill. We are also grateful to the editor and six anonymous referees for their valuable suggestions. We acknowledge financial support from NSF grant No. 1947611 and the Alfred P. Sloan Foundation.

1 Introduction

Money metric utility functions are a backbone of welfare economics. They convert income under different prices into equivalent income under base prices. For example, the money metric for 1974 base prices is the hypothetical 1974 income a consumer needs to make her indifferent between her budget set in 1974 and another year. Money metrics are standard tools for quantifying economic growth and inflation because they cardinalize utility with interpretable units. They are also essential to answer a broad range of questions like how policies affect welfare or how social programs should be indexed.

One can calculate a money metric by deflating nominal income using a weighted average of changes in prices, where the weights are compensated (or Hicksian) budget shares (see e.g. Hausman, 1981). Since compensated budget shares are not directly observable, standard price deflators use uncompensated (or Marshallian) budget shares instead. This shortcut leads to the correct answer if preferences are homothetic, but fails when preferences are nonhomothetic. This is because when preferences are nonhomothetic, compensated and uncompensated budget shares are different, and using one in place of the other produces incorrect results.

In this paper, we show how to recover compensated budget shares and money metric utility without imposing homotheticity or parametric assumptions about preferences, and without estimating a demand system. To do this, consider repeated cross sections of households with identical preferences facing common prices. To construct the money metric utility function, in t_0 dollars, for a household with income I at time t , we must know the compensated demand of this household for every $s \in [t_0, t]$. This is revealed at each point in time s by the budget shares of another household with a different income level I' who is on the same indifference curve as the household with income I at t .¹

If we can find such households, then we can calculate the money metric utility function by integration. That is, if we know how to match households over time, we can recover money metric utility. Conversely, if we know money metric utility, then we can match households through time, since households are on the same indifference curves if, and only if, their money metric utility values coincide. The insight is that this is a fixed point problem in terms of observables that can be solved.

Our methodology endogenously identifies the set of households for which a money

¹Even though we assume that households have common preferences that are unchanged over time (we relax this assumption in Section 3.4), our matching approach is based on revealed-preference theory and is not based on interpersonal comparisons of “well-being”. That is, we match a household with income I under t prices with a household with income I' under s prices if the household at t is indifferent between these two budget sets. We do not need to postulate that two households are “equally well-off” if their utilities are the same.

metric value can be calculated reliably, without out-of-sample extrapolation. That is, our approach does not necessarily recover the money metric for all households in the sample because suitable matches may not exist. For example, if there is positive growth over time, then the richest household at any point in time is on an indifference curve that no other household was on in the previous periods. This means that for such a household, we cannot calculate compensated demand in the past and hence the money metric, unless we are prepared to extrapolate Engel curves out-of-sample.

Our method generalizes the standard practice of statistical agencies who weight changes in prices over time using aggregate budget shares. Conventional price deflators like the CPI or the PCE recover money metric utility under the assumptions of homothetic and stable preferences. However, when preferences are nonhomothetic, we show that one must use the budget shares of a unique corresponding income level in the past for each income today instead of aggregate budget shares.²

Our paper also provides a contrast to the popular but ad hoc approach of constructing price indices by household-income group using the budget shares of some fixed percentile of the income distribution in each period. This method lacks a theoretical foundation if percentiles of the income distribution do not remain on the same indifference curve over time, and the shape of the indifference curve varies as a function of income.³

Our approach differs from alternatives that calculate compensated demand based on estimated elasticities of substitution, as it does not require the estimation of non-parametric elasticities of substitution.⁴ Intuitively, our method only recovers compensated demand evaluated at observed prices, whereas the elasticities of substitution determine how compensated demand will react to any price change, even those that have not been observed. As a result, our procedure can measure changes in welfare for observed changes

²Chain-weighted indices measured by statistical agencies are generally uninterpretable when preferences are nonhomothetic. 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, chained indices measure the cost-of-living price index for some intermediate utility level under AIDS preferences. Caves et al. (1982) establish a similar result for Törnqvist price indices, up to a second-order approximation. But these are not money metrics. In this paper, we focus on the money metric utility function.

³National statistical agencies sometimes produce inflation statistics like this. For example, the UK's Office of National Statistics produces inflation indices by household expenditure groups (see <https://www.ons.gov.uk/economy/inflationandpriceindices/articles/inflationandthecostoflivingforhouseholdgroups/october2022>).

⁴In this respect, our approach resembles Oulton (2012), who demonstrates how to back out compensated budget shares by adjusting uncompensated budget shares using a Taylor series in income. He applies this methodology, using the Quadratic Almost Ideal Demand System of Banks et al. (1997), 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.

in prices and income but is not suited for addressing counterfactual welfare questions, such as those explored by Baqaee and Burstein (2023).

The paper is organized as follows. In Section 2, we define money metric utility and its dual, the cost-of-living index, and explain their relationship to compensated demand. In Section 3, we demonstrate how to recover the cost-of-living index and money metric utility given cross-sectional data when all prices are fully observed over time. We present two solution strategies, both of which exactly recover the money metric as long as the data are continuous in both the time series and the cross section. Using an artificial example, we show that both methods quickly converge to the truth as the number of households and the temporal frequency of observations increase.

We also discuss how our results change when there is preference heterogeneity in either the cross section or time series. To account for heterogeneity in preferences that depend on observable characteristics across households, we split the sample by observed characteristic and apply our method to each subsample separately. With unobservable taste shocks, there are certain cases in which our methodology produces reliable results. For example, our approach approximately recovers the true money metric as long as taste shocks are small and uncorrelated with price changes.

In Section 4, we illustrate our method by applying it to household expenditure survey data from the United Kingdom spanning from 1974 to 2017. We find that real consumption calculated by deflating income with aggregate chain-weighted inflation (as measured by official statistical agencies) overstates the money metric utility for all households below the 60th percentile of the spending distribution in 2017 in our sample. In other words, for expenditures below 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 percentile. For households in the 98th percentile of our sample, who spend around £94,000 per year, the size of this gap is 16 percentage points over the whole sample (0.36 percentage points per year on average).⁵ We are unable to compute the money metric for the richest households in 2017 (98th percentile and above). The reason is that for these households, there did not exist consumers in the past whose money metric utilities are high enough and whose observed demand can be used in place of the compensated budget shares.

⁵These results are consistent with Blundell et al. (2003), which report a relatively greater rise in the cost of living for poorer households between 1975 and 1984 in the UK.

Whereas real consumption calculated using the aggregate inflation rate has large errors relative to our true estimated money metric, a decile-specific chained deflator produces smaller errors in our UK dataset. Of course, one needs to compute the true money metric first, before knowing whether or not the ad hoc approach is a good approximation. Furthermore, computing quantile-specific chained deflators requires more information than our method.

In Section 5, we extend our methodology to allow for missing prices. To do this, we require the restriction that the expenditure function be separable between observed and unobserved prices. Under this additional assumption, we show that money metric utility can be recovered if we know the compensated elasticity of substitution between observed and unobserved goods. This generalizes the influential Feenstra (1994) approach to imputing missing prices beyond the homothetic CES case.

Specifically, we show how to back out the change in the relative price of observed and unobserved goods using changes in the compensated budget 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. 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.

We provide an empirical illustration of this extension in Section 6. We assume that some service prices are mismeasured, estimate elasticities of substitution between these services and other goods, and apply our methodology. We find that the price of the compensated bundle of services has been rising faster than official data for rich but not for poor households. This implies that the money metric is overstated for rich but not poor households.

We conclude in Section 7. Proofs and supplementary materials are in an Online Appendix.

Related Literature. Our paper is closely related to Blundell et al. (2003) and Jaravel and Lashkari (2024), both of which develop nonparametric approaches to measuring welfare for nonhomothetic preferences using cross-sectional household-level data. Although inspired by them, our approach is different and builds on Lemma 1 from Baqaee and Burstein (2023), which expresses the money metric as an integral of compensated demand curves. We discuss the alternative approaches of Blundell et al. (2003) and Jaravel and Lashkari (2024).

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.⁶ Our approach has an advantage over Blundell et al. (2003) in that it provides a point estimate, rather than only bounds, for the money metric utility. On the other hand, in order for our methodology to recover point estimates for the money metric utility without approximation errors, the data must be observed continuously.⁷ We show in Online Appendix Section A.7 that our point estimates are always within their bounds in real-world data.

Jaravel and Lashkari (2024) use a correction term to address nonhomotheticity in household-level chain-weighted indices. Whereas our approach endogenously delineates a set of households for whom money metric utility can be calculated, without relying on out-of-sample extrapolation, the Jaravel and Lashkari (2024) method aims to uncover the money metric for all households observed at any point in time. That is, unlike our methodology, their approach does not provide a boundary on the set of households whose money metric values can be reliably computed. In Online Appendix Section A.8, we apply the Jaravel and Lashkari (2024) method to artificial examples. If the support of the cross-sectional distribution of utilities changes over time, then their algorithm can diverge or result in large errors (and these errors persist even as we increase the sample size and frequency of observation).

In contrast to both Jaravel and Lashkari (2024) and Blundell et al. (2003), we also extend our methodology to situations where some prices and expenditures are unobserved. Since our method can be extended to allow for unmeasured prices, our paper is also related to the literature that measures welfare allowing for incomplete information about prices. Most papers with nonhomothetic preferences follow the approach of Costa (2001) and Hamilton (2001). These papers take advantage of horizontal shifts in Engel curves to identify money metric utility changes. The frontier in this literature is Atkin et al. (2024), who show how to identify welfare changes assuming that preferences are quasi-separable between the measured and unmeasured goods.

Our paper, instead, generalizes the Feenstra (1994) method beyond the homothetic CES case. One advantage of our approach is that we do not need to make strong parametric

⁶We exposit and implement an amended version of their methodology in Online Appendix Section A.7, fixing a typographical error in their algorithm for the lower bound.

⁷We interpolate budget shares to turn discrete data continuous. If this interpolation is inaccurate, then this introduces approximation errors into our method. Such errors are not specific to our method and result from the fact that sums do not perfectly measure integrals. For example, even when preferences are homothetic, interpolation error affects the accuracy of standard chained-weighted price deflators.

assumptions within the set of observed prices. This is in contrast to Atkin et al. (2024) who need to fully model the demand system for the subset of goods with observed prices. This advantage of our approach comes at the cost that we require a stronger form of separability between the observed and unobserved prices than Atkin et al. (2024). We discuss these issues in more detail in Section 5.

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 nonhomothetic preferences are used to understand phenomena as diverse as structural transformation (e.g. Boppart 2014, Comin et al. 2021, and Fan et al. 2023), international trade patterns (e.g. Matsuyama 2000, and Fajgelbaum et al. 2011), and savings behavior and inequality (e.g. Straub 2019). Our approach provides a nonparametric way to compute welfare measures from the data without relying on low-dimensional functional forms.

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 utility function $\mathcal{U}(c)$ that maps consumption bundles to utility values. Given this utility function, we can define the *indirect utility function*

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

mapping a vector of prices p and expenditures I to utility values. We interchangeably refer to I as income, but in the data, we measure I using expenditures. Define the *expenditure function* to be

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

We assume that the expenditure function is continuously differentiable in all its arguments.

The expenditure and indirect utility functions are used to define money metrics and cost-of-living indices.

Definition 1 (Money Metric and Cost of Living). For a fixed reference vector of prices \bar{p} , the *money metric* function maps budget sets defined by (p, I) to

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

For a fixed reference budget set defined by (\bar{p}, \bar{I}) , the *cost-of-living index* maps prices, p , to

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

The money metric function, $e(\bar{p}, v(\cdot))$, converts the value of different budget sets (p, I) into equivalent dollars under some baseline prices \bar{p} . It is itself an indirect utility function because a budget set (p, I) is preferred to another budget set (p', I') if, and only if, $e(\bar{p}, v(p, I)) > e(\bar{p}, v(p', I'))$. Because the money metric ranks budget sets and assigns them an interpretable value, it is useful for measuring growth.⁸

The cost-of-living function, $e(\cdot, v(\bar{p}, \bar{I}))$, converts the value of some baseline budget constraint (\bar{p}, \bar{I}) into equivalent income under different sets of prices.⁹ Because the cost-of-living index converts a common utility level, $v(p, I)$, into equivalent income under different price systems, it is useful for measuring the cost-of-living adjustment to maintain a fixed standard of living.

In sum, the function $e(p', v(p, I))$, mapping (p', p, I) into a scalar, is an object of paramount interest. The money metric is the cross section of this function that holds p' constant and the cost-of-living index is the cross section that holds (p, I) constant. Our aim is to recover this object from the data.

Denote the compensated budget share for good i by $b_i(p, U)$ where p is a vector of prices and U is a utility level. The following lemma, which is a corollary of Lemma 1 from Baqaee and Burstein (2023), and follows from Shephard's lemma and the gradient theorem, provides a characterization of both the cost-of-living index and the money metric using compensated budget shares.

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

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

where C is any absolutely continuous path connecting \bar{p} to p .¹⁰ The cost of living for a budget set

⁸The equivalent and compensating variation are related to the money metric. Specifically, to measure the change in welfare from some initial budget set (p, I) to some other budget set (p', I') , the equivalent variation is $e(p, v(p', I')) - I$ and the compensating variation is $I' - e(p', v(p, I))$.

⁹In index number theory, the cost-of-living index is also called the Konüs (1939) index.

¹⁰Formally, the path integral in (1) is defined by $\int_{t_0}^{t_1} \sum_{i \in N} b_i(\xi_t, v(p, I)) \frac{d \log \xi_{it}}{dt} dt$ where $\{\xi_t : t \in [t_0, t_1]\}$ parameterizes the path C from \bar{p} and p as a function of a scalar t . The integrals in (1) and (2) are both path independent and only depend on the end points.

(\bar{p}, \bar{I}) in terms of p prices can be expressed as

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

According to Lemma 1, both the money metric and the cost-of-living index can be expressed as integrals of compensated budget shares with respect to changes in prices. However, compensated demand curves are not directly observable, so operationalizing this result requires having a way to identify compensated budget shares. This is what we focus on in the next section.

3 Recovering the Money Metric by Matching Households

In this section, we discuss how Lemma 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 two solution methods, and test them with artificial discrete data to assess their accuracy. We end the section by discussing how our results are affected by taste shocks and mismeasurement.

3.1 Theoretical Result

Consider an absolutely continuous path of prices $p_t \in \mathbb{R}^N$ as a function of time $t \in [t_0, T]$. Suppose we observe vectors of budget shares $B(I, t) \in \mathbb{R}^N$ for consumers with preferences \succeq and income levels $I \in [I_t, \bar{I}_t]$ for time $t \in [t_0, T]$.¹¹ Our aim is to recover the money metric utility function based on reference prices p_{t_0} evaluated at budget set (p_t, I) for $t \in [t_0, T]$ and $I \in [I_t, \bar{I}_t]$. We denote this function by $u(I, t) \equiv e(p_{t_0}, v(p_t, I))$.

The function $u(I, t)$ converts the value of the budget constraint defined by prices p_t and income I into income under base prices p_{t_0} . Once we are equipped with $u(I, t)$, it is also straightforward to compute the money metric for other base prices.¹² By varying base prices, for fixed (p_t, I) , we can also recover the cost-of-living index.

¹¹We can always produce an absolutely continuous path of prices by linearly interpolating between discrete-time observations. We can construct an associated budget share at each instant in time by linearly interpolating budget shares over time. See Footnote 15 for more details about interpolation.

¹²Suppose we wish to obtain the money metric for some other base prices: $\tilde{u}(I, t) = e(p_s, v(p_t, I))$ for some $s \in [t_0, T]$. The solution is $\tilde{u}(I, t) = I'$ where I' satisfies $u(I', s) = u(I, t)$. By construction, $v(p_t, I) = v(p_s, I')$, hence $\tilde{u}(I, t) = e(p_s, v(p_t, I)) = e(p_s, v(p_s, I')) = I'$.

Denote the uncompensated budget share of good i by B_i^M (the superscript M stands for Marshallian). For every good i ,

$$B_i^M(\mathbf{p}_t, I) = B_i(I, t)$$

whenever $t \in [t_0, T]$ and $I \in [I_t, \bar{I}_t]$. For any cardinalization of the indirect utility function and its associated compensated demand curves, the following identity between compensated and uncompensated budget shares also holds:

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

Using the money metric cardinalization of indirect utility, and slightly abusing notation, we can combine the previous two identities to obtain:¹³

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

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

Proposition 1 (Money metric as the Solution to the 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 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 $u(I^*, s) = u(I, t)$.

Since the money metric exists, the integral equation (3) necessarily has a solution. Proposition A.1 in Online Appendix Section A.3 uses the contraction mapping theorem to show that the solution to this integral equation is also unique.

Proposition 1 follows immediately from Lemma 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 compensated 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.

¹³Our “abuse of notation” is that we do not index compensated budget shares by the utility cardinalization. This is to simplify notation, since we are interested in compensated budget shares only under the money metric cardinalization.

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 (1926) 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 compensated budget shares $b(\mathbf{p}_s, \bar{u})$ for a utility level \bar{u} 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 . That is, we “match” households with income I^* at time s to households with income I at time t if $u(I^*, s) = u(I, t)$. The budget shares of this “matched” household, $B(I^*, s)$, are equal to the compensated budget shares $b(\mathbf{p}_s, \bar{u})$.

Proposition 1 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 $B_i^M(\mathbf{p}_t, I)$. 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 1 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.

In the next section, we provide two solution methods for solving the integral equation in Proposition 1 with discrete data.

3.2 Two Solution Methods

The money metric is a fixed point of equation (3), which is a system of nonlinear equations, albeit an infinite-dimensional one. We provide two solution methods. The first is a simple iterative procedure that converges to the desired solution as we approach the continuous-time limit. The second is a recursive solution that is equivalent to the iterative one in the continuous-time limit but has better properties when the data are discrete.

For both methods, suppose that we have data on a grid of points $\{t_0, \dots, t_M\}$ where $t_n < t_{n+1}$, with $t_M = T$. For each t , we observe budget shares $B(I, t)$ for any income level $I \in [\underline{I}_t, \bar{I}_t]$.¹⁴

¹⁴In our empirical application, in Section 3.4, we fit a smooth curve through micro data to obtain $B(I, t)$ for $I \in [\underline{I}_t, \bar{I}_t]$ since cross-sectional household-level data on expenditures is noisy.

Iterative Solution. Use the following iterative procedure for each $n \in \{1, \dots, M\}$ starting with $n = 1$:

$$\log u(I, t_n) \approx \log I - \sum_{m=0}^{n-1} B(I_m^*, t_m) \cdot \Delta \log p_{t_m}, \quad (5)$$

where I_m^* satisfies

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

with the boundary condition $u(I, t_0) = I$. For any $m \leq n - 1$ for which we cannot find I_m^* satisfying (6), obtain a candidate I_m^* using a loglinear approximation in I . If the candidate I_m^* is not in the observed income distribution $[I_{t_m}, \bar{I}_{t_m}]$, then we cannot evaluate $B(I_m^*, t_m)$ without out-of-sample extrapolation, so we do not calculate $u(I, t_n)$. As a final check, we exclude $u(I, t_n)$ if there does not exist $I_m^* \in [I_{t_m}, \bar{I}_{t_m}]$ such that $u(I_m^*, t_m) = u(I, t_n)$ for all $m \leq n - 1$.

There are two approximation errors in the iterative solution. The first is that, in (5), we are approximating an integral using a discrete Riemann sum. The second is that, in (6), we are using $u(I, t_{n-1})$ rather than $u(I, t_n)$ on the right-hand side (since we do not know $u(I, t_n)$ in step n). However, as we approach the continuous-time limit, the estimates produced by (5) converges to the exact solution in (3). This is because the summation in (5) converges to an integral and $u(I, t_{n-1})$ in (6) converges to $u(I, t_n)$. Since (3) has a unique solution, the continuous-time limit of (5) converges to the money metric. To summarize, if data are continuous, then the result is an exact solution to the money metric that requires no estimation or interpolation.¹⁵

The iterative procedure that we describe is useful for building intuition. However, one can also find a fixed point by solving the system of equations directly. This gives a recursive variation of the iterative procedure described above. The two approaches are equivalent in the continuous-time limit.

¹⁵In practice, we use the trapezoid rule rather than the left-Riemann sum to approximate integrals. That is, we use $\frac{B(I_m^*, t_m) + B(I_{m+1}^*, t_{m+1})}{2}$ in place of $B(I_m^*, t_m)$ in (5). This numerical refinement is equivalent to linearly interpolating prices (in logs) and budget shares over time between discrete-time observations. If the true budget shares corresponding to the linearly interpolated path of prices are not themselves linear, then this will introduce an interpolation error into our results. This error disappears as the price shocks between any two consecutive periods become small.

Recursive Solution. Apply the iterative solution in (5) and (6) and call the resulting money metric $u_0(I, t)$. For each $i \geq 1$, and each $n \in \{1, \dots, M\}$, starting with $n = 1$, define

$$\log u_{i+1}(I, t_n) \approx \log I - \sum_{m=0}^{n-1} \mathbf{B}(I_m^*, t_m) \cdot \Delta \log \mathbf{p}_{t_m}, \quad (7)$$

where I_m^* satisfies

$$u_{i+1}(I_m^*, t_m) = u_i(I, t_n). \quad (8)$$

For any $m \leq n - 1$ for which we cannot find I_m^* satisfying (8), obtain a candidate I_m^* using loglinear approximations in I . If the candidate I_m^* is not in the observed income distribution $[L_{t_m}, \bar{L}_{t_m}]$, then we cannot evaluate $\mathbf{B}(I_m^*, t_m)$, so we do not calculate $u(I, t_n)$. Continue until $u_{i+1}(I, t) = u_i(I, t)$ for all feasible values of I and t . Then set $u(I, t) = u_i(I, t)$.

Once the recursive solution converges, it solves the fixed-point problem without any out-of-sample extrapolation of budget shares. The difference between the iterative and recursive solution is that we replace $u(I, t_{n-1})$ on the right hand side of equation (6) with $u(I, t_n)$ in (8). Proposition A.1 in Online Appendix Section A.3 shows that the continuous-time version of this recursive procedure is a contraction mapping and must necessarily converge to the unique solution (which is the money metric).

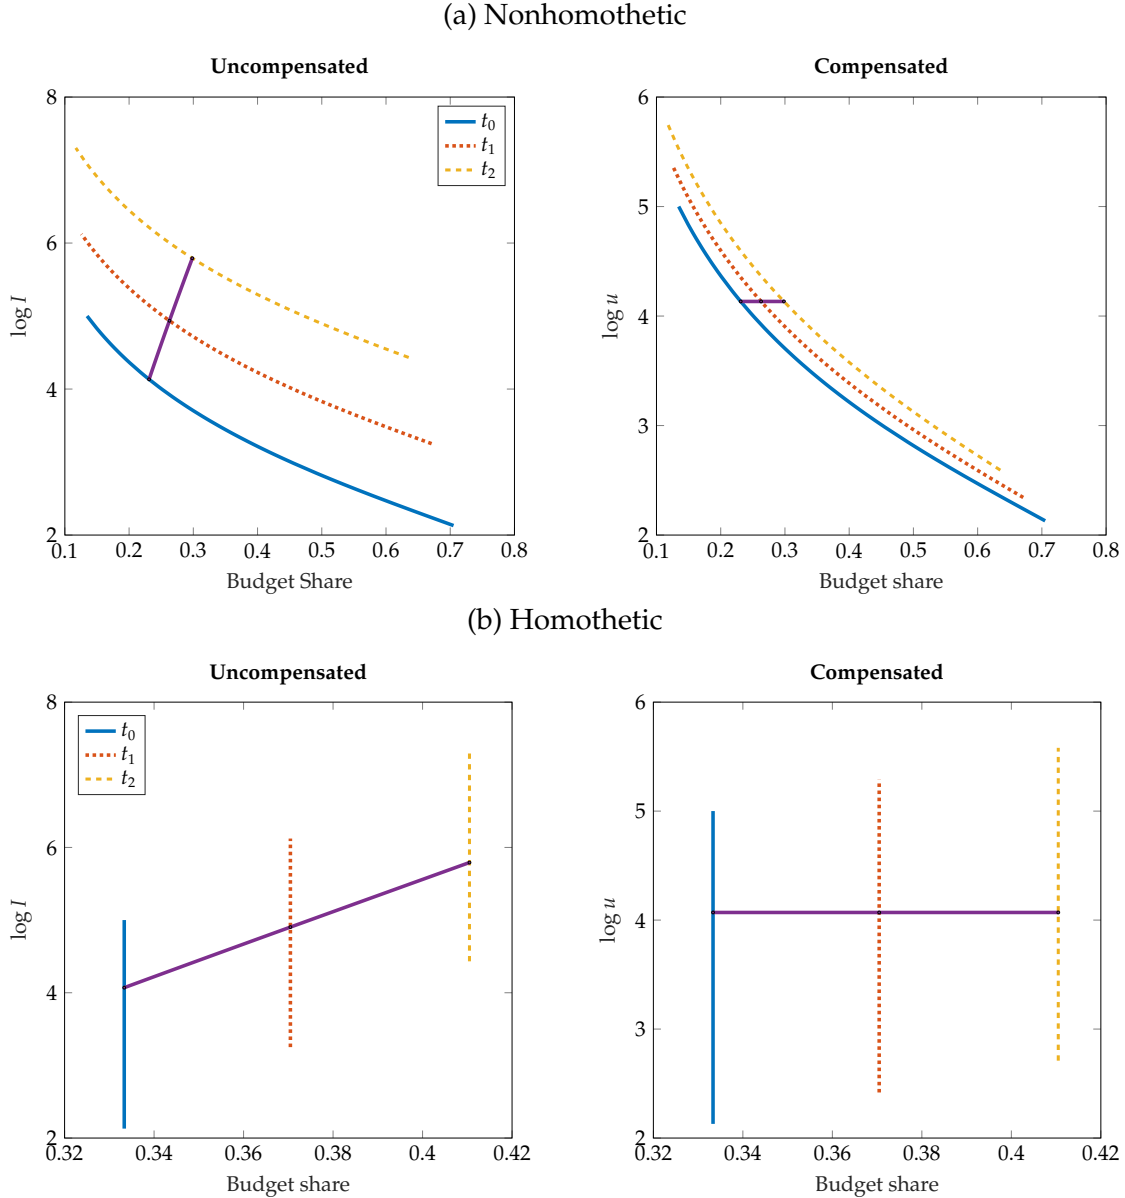
Both procedures endogenously delineate values of (I, t_n) for which $u(I, t_n)$ can be computed without extrapolating budget shares. We only compute money metric utility $u(I, t_n)$ if $u(I, t_n)$ is between the upper- and lower-bound of $u(\cdot, t_m)$ for every $m < n$. This ensures that the income level I_m^* that solves (6) or (8) is in the support of the income distribution where budget shares are observed.

In artificial examples with discrete data, the recursive solution has smaller errors than the iterative solution, although, both methods work well. When we use real data from the UK, in Section 4, the results are almost unchanged between the iterative and recursive methods. Since the iterative procedure is simpler and faster to compute, we only show results for the iterative method for our empirical results.

Figure 1 illustrates the outcome of our procedure. The left panel of Figure 1a shows the budget share on some good against nominal income for three different points in time. The fact that the lines are downward sloping means that this good is a necessity. 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 (uncompensated

Figure 1: Budget share for some good against nominal income and money metric utility in different periods.



budget shares), but to construct the money metric we require budget shares as a function of utility (compensated budget shares). The right panel of Figure 1a (color version available online) displays the compensated budget shares for the same good. The horizontal solid purple line in the right panel of Figure 1a shows for each period the compensated budget share for the good evaluated at some fixed utility level \bar{u} . The change in budget shares, holding utility constant, are pure substitution effects over time due to changes in relative prices. As implied by Lemma 1, multiplying the compensated 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 compensated budget shares? The upward-sloping solid 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 good, $B_i(u^{-1}(\bar{u}, s), s)$. In other words, we can infer compensated 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 1 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, uncompensated and compensated budget shares coincide. Therefore, we can construct the money metric using a price index based on uncompensated budget shares by good.

3.3 Example with Artificial Discrete Data

To illustrate how our method fares when faced with discrete data, rather than continuous data, we consider a simple artificial example. Suppose the expenditure function is nonhomothetic CES

$$e(\mathbf{p}, U) = \left(\sum_i \omega_i (U^{\varepsilon_i} p_i)^{1-\gamma} \right)^{\frac{1}{1-\gamma}}. \quad (9)$$

The money metric function for t_0 reference prices is

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

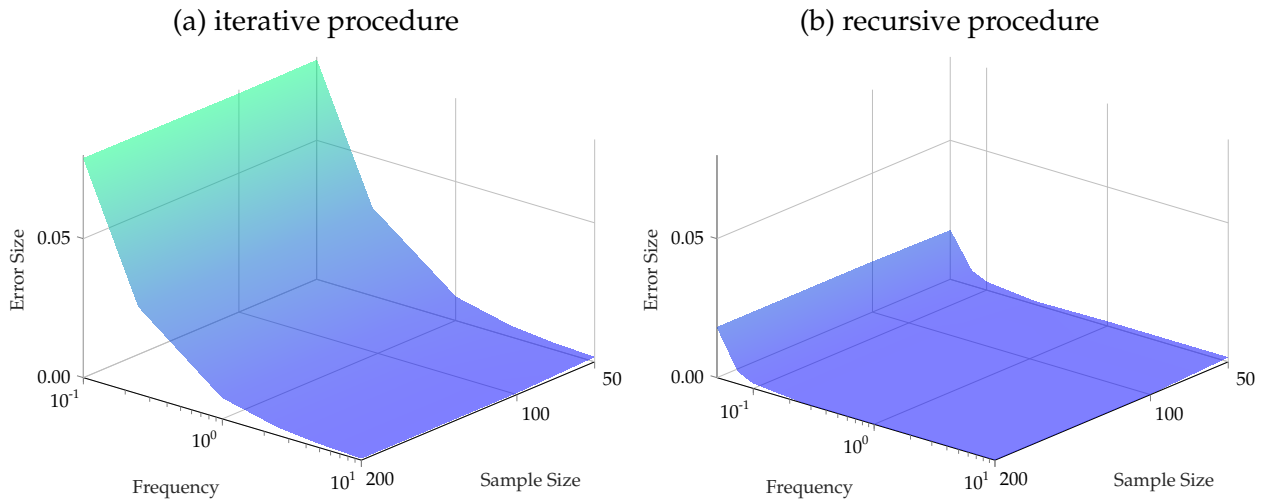
where V is the indirect utility function and solves $I = e(\mathbf{p}_t, V)$.¹⁶ We evaluate the accuracy of our algorithm by comparing this exact expression for $u(I, T)$ with the results of our numerical procedure applied to artificial data generated using these preferences.

For illustration, we set $\gamma = 0.25$, $\varepsilon_1 = 0.2$, $\varepsilon_2 = 1$, $\varepsilon_3 = 1.65$, which are values taken

¹⁶See Hanoch (1975), Matsuyama (2019), and Comin et al. (2021) for more information on these preferences. Baqaee and Burstein (2023) show that the money metric for a nonhomothetic CES has a closed-form expression in terms of observable budget shares and the elasticity of substitution: $u(I, t) = I \times \left(\sum_i B_i(I, t) \left(\frac{p_{i,t_0}}{p_{i,t}} \right)^{1-\gamma} \right)^{\frac{1}{1-\gamma}}$. Note that the budget shares depend on the difference in ε_i and not their overall level. This implies that ε_i are identified, and matter, up to an additive constant.

from Comin et al. (2021). We generate repeated cross-sectional data on income and budget shares over three goods for a finite number of 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). The share parameters are calibrated so that the budget share of each good for the median household in the first period are uniform. All incomes and prices grow exponentially, at different rates, over the sample period.¹⁷

Figure 2: Maximum error as a function of the frequency of observation and sample size



Notes: Throughout, we hold the path of prices and incomes constant. Our baseline calibration is annual frequency corresponding to a value of $10^0 = 1$ observation 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 and the right panel the recursive solution method in Section 3.2.

We apply both the iterative and recursive solution methods. We use linear interpolation to evaluate budget shares for I between two observed income levels. To assess the accuracy of our procedure, we use the infinity norm—the maximum absolute value of the log difference between the true money metric function and our estimate in the final period. The error is very small. For example, with 100 households and annual data, the maximum error in the final period is 0.0077 for the iterative procedure and 4×10^{-5} for the recursive procedure. So, the error is less than 1% of income for the iterative procedure and around 1/250th of 1% for the recursive procedure. Figure 2 shows how this error varies as we vary the number of households and the frequency of observations. As expected, the error

¹⁷The ratio of top to bottom nominal expenditures every period is around 18. The annual (log) growth rate of nominal expenditures is 5.7%, and the annual growth rates of prices of each good are 5%, 4.25%, and 3.5%, respectively.

converges to zero as we approach the continuous-time limit. The error also falls as the number of households in the sample increases.¹⁸

3.4 Tastes Shocks and Mismeasured Expenditures

In practice, data are imperfect and noisy. There are two potential sources of error in the data: (1) expenditures by good may be mismeasured or affected by taste shocks; (2) prices may be missing or mismeasured. In this section, we focus on mismeasured expenditures due to measurement error and or taste shocks. We address missing or mismeasured prices in Section 5, where we impose stronger assumptions on preferences.

If there are arbitrary unobservable shocks to preferences or measurement error, then our methodology cannot be used reliably. However, there are certain tractable cases with shocks where we can still apply our methodology. In this section, we discuss these cases. We begin by considering the straightforward scenario where preferences vary as a function of observable characteristics—for instance, households with children have distinct tastes compared to those without.¹⁹

Proposition 2 (Tastes Vary by Observed Characteristics). *If there are differences in preferences that are functions of observable characteristics, then split the sample by characteristic and apply Proposition 1 to each subsample separately.*²⁰

Next, we consider the more difficult case where observed expenditures depend on unobservable taste shocks or measurement error. Suppose that observed budget shares are

$$\tilde{B}(I, t|\kappa) = B(I, t) + \kappa\epsilon(I, t),$$

where $B(I, t)$ are the *true* expenditure shares. That is, $B(I, t)$ are expenditure shares generated by the preferences that we wish to construct the money metric for. However, we cannot observe $B(I, t)$ because the data feature either taste shocks or mismeasurement. The term $\kappa\epsilon(I, t)$ is defined to be the difference between observed budget shares and budget shares generated by the preferences that we are interested in.²¹ The scalar $\kappa \geq 0$ controls the magnitude of these errors.

¹⁸In Online Appendix Section A.8, we apply the Jaravel and Lashkari (2024) algorithm to this example and find similarly small errors. However, in the same appendix, we provide other examples where the Jaravel and Lashkari (2024) approach yields large errors or diverges.

¹⁹This assumption is similar to that considered in Section 2.3 of Jaravel and Lashkari (2024).

²⁰Similarly, 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.

²¹See Baqaee and Burstein (2023) for a detailed analysis of how welfare should be defined when preferences are subject to taste shocks. In general, in the presence of taste shocks, $B(I, t)$ need not correspond to the preferences of any individual in the cross section.

Define $\tilde{u}(I, t|\kappa)$ to be the solution to the integral equation

$$\log \tilde{u}(I, t|\kappa) = \log I - \int_{t_0}^t \sum_i \tilde{B}_i(\tilde{u}^{-1}(\tilde{u}(I, t|\kappa), s|\kappa), s|\kappa) \frac{d \log p_{is}}{ds} ds. \quad (10)$$

Proposition 1 assumes that $\kappa = 0$. That is, $\tilde{u}(I, t|0) = u(I, t)$.

When there is idiosyncratic (mean-zero) noise at the level of individual households, averaging over households ensures that $\kappa = 0$ as long as the law of large numbers holds. In such situations, we can apply Proposition 1 without concerns about taste shocks and recover the money metric for preferences in the absence of the idiosyncratic noise. However, if the errors do not average out, they could potentially impact the results. To analyze the extent of this influence, we derive a first-order approximation of $\tilde{u}(I, t|\kappa)$ with respect to the error term κ . The general form of this first-order approximation can be found in Lemma A.1 in Online Appendix Section A.2. Within the main text, we highlight two tractable and salient special cases.

Proposition 3 (Taste Shocks Uncorrelated with Price Shocks). *Suppose that for all I and $s \leq t$, we have $\text{Cov}\left(\epsilon(I, s), \frac{d \log p}{ds}\right) = 0$. Then, to a first-order approximation around $\kappa \approx 0$,*

$$\tilde{u}(I, t|\kappa) \approx u(I, t),$$

where the remainder term is order κ^2 .

In words, if the shocks are uncorrelated with price changes, then the money metric we construct by solving the wrong integral equation is, to a first-order approximation, correct. This approximation assumes that κ is small, but does not require that t be close to t_0 .²²

Next, we consider how taste shocks that are correlated with price changes affect our results.

Proposition 4 (Engel Curve Slopes Uncorrelated with Price Shocks). *Suppose that for all I and $s \leq t$, the slope of Engel curves is uncorrelated with price changes $\text{Cov}\left(\frac{\partial B(I, s)}{\partial I}, \frac{d \log p}{ds}\right) = 0$.*

²²This result bears a superficial resemblance to previous results, for example by Baqaee and Burstein (2023), that Divisia indices approximately measure welfare correctly when taste shocks are uncorrelated with price changes. However, Proposition 3 is different since it characterizes the solution to an integral equation and not the Divisia index. The results about the Divisia index are based on a second-order approximation that requires that t be close to the base year t_0 . On the other hand, Proposition 3 is based on a first-order approximation in κ , and t can be far from t_0 .

Then, to a first-order approximation around $\kappa \approx 0$,

$$\tilde{u}(I, t|\kappa) - u(I, t) \approx -\kappa \int_{t_0}^t \text{Cov}\left(\epsilon(u(I, t), s), \frac{d \log p}{ds}\right) ds,$$

where the remainder term is order κ^2 .

If the slope of Engel curves is uncorrelated with price shocks (necessarily the case when preferences are homothetic), then the money metric we construct is biased according to how the taste/measurement shocks $\epsilon(I, t)$ covary with price shocks. That is, although our methodology will have errors, the sign and magnitude of these errors can be linked to the underlying shocks in a straightforward way. In particular, if the mismeasured expenditures are biased upwards for goods whose relative price rose, then the constructed money metric will be biased downwards.

We need the requirement that the slope of Engel curves be uncorrelated with price shocks because otherwise, as we solve the integral equation forward, errors in prior values of $\tilde{u}(I, s|\kappa)$, for $s \leq t$, contaminate the matching process in a systematic way and induce additional biases in $\tilde{u}(I, t|\kappa)$.

4 Empirical Illustration

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 national accounts: nominal expenditures deflated by a chain-weighted price index that reflects observed (either aggregate or decile-specific) budget shares.²³ When preferences are homothetic, then real consumption for every household coincides with money metric utility.

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 subcategories of goods and services from 1974 to 2017.²⁴ The

²³The 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 some average budget share of good i in period t . If we use the aggregate budget shares, then the price deflator is common for all households. Alternatively, we can group households by quantiles of the spending distribution and use average budget shares by quantile. We compare our results with both aggregate and decile-specific price deflators.

²⁴Aggregate 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,

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 (ONS), 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.²⁵ Between 1974 and 2017 prices rose relatively less for product categories that are disproportionately consumed by richer households, such as leisure goods and services. Even though we consider product categories that are more aggregated than the official data, our data tracks the official inflation figures from the ONS fairly well.²⁶

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 (following Proposition 2). We provide examples using marital status and age in Online Appendix Section A.4. 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.

4.1 Mapping Data to the Model

Our procedure requires expenditures I and budget shares $B(I, t)$ at time t across all goods. To deal with idiosyncratic noise, we fit a smooth curve to the budget share of each good i at time t as a function of I . We use these curves as $B(I, t)$. More precisely, we estimate the true $B_i(I, t)$ function for each good i by fitting the following curve for each t using ordinary

which are at the household level.

²⁵See Online Appendix Section A.5 for details about our concordance table. We also calculate our results using more disaggregated spending categories, using only CPI data, from 2001 to 2017. Online Appendix Figure A.3 compares these results to what we get if we instead use the more aggregated 17 spending categories instead for the same time period. The gaps relative to the chain-weighted inflation index are qualitatively similar but moderately larger when we use more disaggregated spending categories. Unfortunately, the more disaggregated data are not available for the full sample, so we use the more aggregated data for our benchmark. In principle, one should apply our methodology to the most disaggregated spending categories possible in order to minimize aggregation bias.

²⁶See Online Appendix Figure A.5 and Online Appendix Table A.1 for comparisons of our data with aggregate inflation and inflation by decile of expenditures as reported by the ONS.

least squares

$$B_{iht} = \alpha_{0it} + \alpha_{1it} \log I_{ht} + \alpha_{2it} (\log I_{ht})^2 + \varepsilon_{iht},$$

where h is the household and t is the time period. The estimated regression line gives us $B(I, t)$, which we then normalize to ensure that budget shares add up to one across goods for every income level and time period. Importantly, we only evaluate the estimated $B(I, t)$ in-sample to avoid out-of-sample extrapolation errors. As mentioned before this potentially limits the set of values for which we can construct the money metric, but ensures that our estimates are more reliable. Our results are virtually unchanged if we estimate the Engel curves nonparametrically (i.e. using locally weighted scatterplot smoothing, LOWESS) instead of quadratic functions (see Figure A.2 in Online Appendix Section A.4).

Since this regression is the only source of sampling uncertainty in our exercise, we calculate standard errors for our estimates of the money metric by bootstrapping this regression. To do this, we redraw repeated samples with replacement. Although the Engel curves are estimated with considerable uncertainty, the standard errors for the money metric are 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 money metric utility using 1974 base prices by applying our procedure sequentially from 1974 to 2017 to the UK data.²⁷ Computing $u(I, t)$ requires that for each time $s < t$, we can estimate the compensated budget share $b(p_s, u(I, t))$. That is, for each expenditure 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 3 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 expenditures in 2017 are matched with households in the 78th percentile of expenditures in 1974. The dashed diagonal line is the 45-degree line and is what we would get if we matched households by percentile of the distribution. This is how price deflators by spending group are typically calculated by statistical agencies (we compare our results with such a measure below).

The right panel of Figure 3 plots the distribution of log expenditures in our data and the solid lines show the sample of households for which we can calculate $u(I, t)$. Our algorithm can recover the money metric up to about the 98th percentile of households in 2017. For

²⁷Given the money metric at some base prices, we can easily obtain the money metric at any other base prices in $t_m \in [t_0, T]$, as explained in Footnote 12.

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 expenditures is highly fat-tailed, which means that in 1974, there are households who are on the same indifference curve as the richest 98th percentile of households in 2017.

Figure 3: Results of the matching process



Notes: The figure on the left shows, for each expenditure percentile in 2017, the expenditure 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 compensated budget share can be computed as a function of time. The lower and upper bounds in 2017 represent the 0.2th and 98th percentile, respectively, of the spending distribution.


4.2 Results

The solid blue line in the left panel of Figure 4 plots the expenditure function $e(p_{1974}, v(p_{2017}, I))$ for different values of I . This expresses different income levels in 2017 (x -axis) in terms of 1974 pounds (y -axis) — the money metric utility function with 1974 base prices. We can also use this figure to convert different income levels in 1974 (y -axis) in terms of 2017 pounds (x -axis) — the cost-of-living function.²⁸

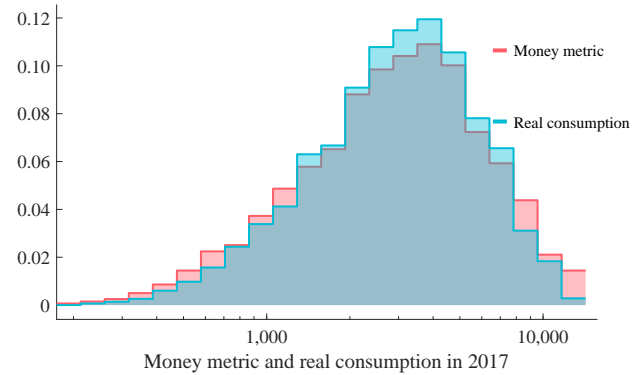
²⁸That is, pick an I' on the y -axis, and find the associated I on the x -axis. Then, since $v(p_{1974}, I') = v(p_{2017}, I)$, it must be that $e(p_{2017}, v(p_{1974}, I')) = e(p_{2017}, v(p_{2017}, I)) = I$. In words, I is the cost-of-living adjustment needed to keep a household with budget set (p_{1974}, I') on the same indifference curve in 2017 as in 1974.

For comparison, the dotted red line shows the equivalent incomes in 1974 if all households faced the same effective inflation rate, as given by the chain-weighted aggregate inflation rate. When the dotted red line is above the solid 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 the distribution and then widens suggesting that the richest households are better off in 1974 pounds than what is implied by official statistics. Accordingly, the histograms in the right panel of Figure 4 show that inequality across households is larger based on money metric values than based on real consumption.

Figure 4: Comparison of money metric with chain-weighted real consumption.



(a) Real consumption using aggregate chain-weighted inflation between 1974 to 2017 (annualized pounds, log scale) and the money metric $e(p_{1974}, v(p_{2017}, I_{2017}))$. This figure converts income in 1974 into equivalent income in 2017 and vice versa.



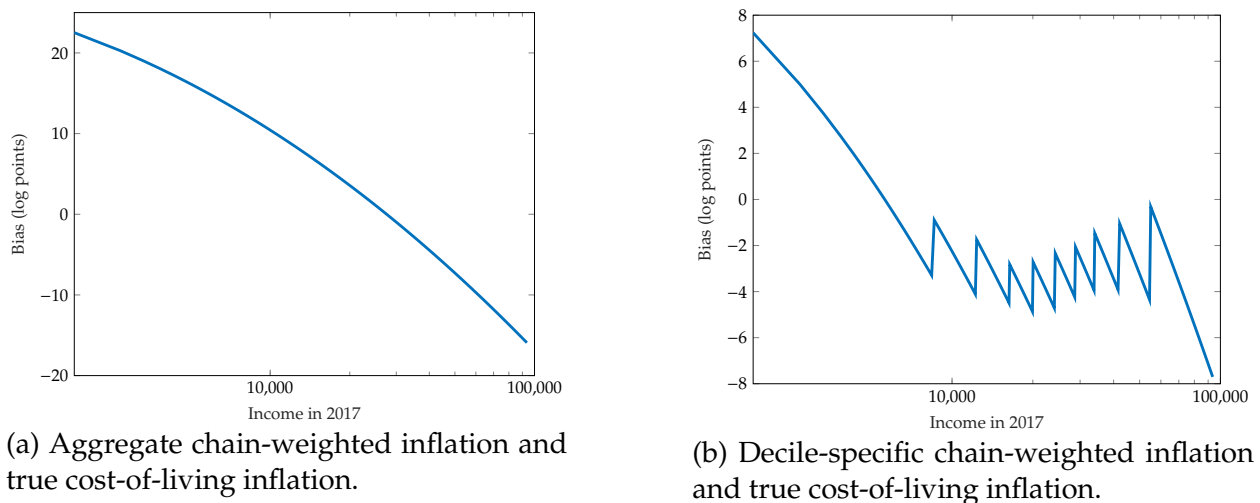
(b) Histogram (using household weights) of money metric $e(p_{1974}, v(p_{2017}, I_{2017}))$ and real consumption using aggregate chain-weighted inflation (annualized pounds, log scale). The distributions are truncated at the upper and lower bounds of Figure 3.

The left panel of Figure 5 displays the log difference between the dotted red and solid blue lines in Figure 4. 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 20 log points for the poorest households. This means that over the 44-year sample, annual inflation rates calculated as in the official statistics understate the true welfare-relevant inflation (i.e. the deflator implied by the money metric and cost-of-

living functions in Lemma 1) for these households by around 0.5 percentage points per year. On the other hand, for the richest households, the official inflation rate overstates the true inflation by 0.36 percentage points per year on average.

The right panel of Figure 5 shows the errors between the true inflation rate and chain-weighted decile-specific inflation. The errors are much smaller, but not zero. We stress that this does not guarantee that quantile-specific chained deflators always approximate the true money metric well. We expect that in contexts where growth is more rapid, the differences can be larger. The data requirements for constructing the money metric, following our method, are slightly less demanding than the ones required for constructing quantile-specific chained deflators.²⁹

Figure 5: Log difference between chain-weighted inflation and true cost-of-living inflation



Notes: Results are reported in log points (i.e. 100 times the log difference). The sample is from 1974 to 2017.

5 Extension with Partially Observed Prices

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

²⁹Whereas quantile-specific chained deflators require a representative sampling of the entire distribution of households, our methodology 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, as explained in Online Appendix Section A.5. Otherwise, the data requirements of the two methodologies are the same.

goods. The results in this section generalize Feenstra (1994) beyond the homothetic CES case.

To compute welfare without data on some prices, we impose the following assumption about preferences throughout this section.

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

$$e(\mathbf{p}, U) = e(e^X(\mathbf{p}^X, U), e^Y(\mathbf{p}^Y, U), U), \quad (11)$$

where U is utils, \mathbf{p}^X and \mathbf{p}^Y are vectors of prices in X and Y , and e^X and e^Y are nondecreasing in and homogeneous of degree one in prices.

We assume that prices and budget shares of goods in X are observed, but prices and budget shares in Y are unobserved. Assumption 1 does not restrict cross-price elasticities for goods within X or Y but does restrict cross-price effects between X and Y . CES aggregators, used by Feenstra (1994), are separable in every partition of their arguments, so our separability assumption is much weaker. Separability can be tested using the Leontief-Sono conditions, see Blackorby et al. (1998).³⁰ We provide a proof-of-concept illustration for our empirical application below.

Denote the compensated budget share of X goods by

$$b_X = \sum_{i \in X} b_i(\mathbf{p}, U) = b_X \left(\frac{e^X(\mathbf{p}^X, U)}{e^Y(\mathbf{p}^Y, U)}, U \right),$$

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 scalar, $\frac{e^X(\mathbf{p}^X, U)}{e^Y(\mathbf{p}^Y, U)}$, which we can interpret as the relative price of the X and Y bundles.

Define the compensated elasticity of substitution between X and Y goods to be

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

That is, $\sigma(\mathbf{p}, U)$ captures how spending on X goods changes relative to Y goods if the price of all X goods rises by the same amount, holding U constant.³¹

³⁰The Leontief-Sono conditions, which are necessary and sufficient for separability, imply that, for each $i, j \in X$ and $k \in Y$, we must have $\frac{\partial \log \left(\frac{b_i}{b_j} \right)}{\partial \log p_k} = 0$, where b_i and b_j are both compensated budget shares. The same must hold if we swap X and Y .

³¹This elasticity of substitution is disciplined by the curvature of the upper nest of the expenditure

Denote the *relative* uncompensated and compensated budget 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 1 to account for unmeasured prices.

Proposition 5 (Money metric with Missing Prices). *Suppose Assumption 1 holds. For $t \in [t_0, T]$, the money metric $u(I, t)$ is a fixed point of the following integral equation as long as $\sigma(\mathbf{p}_s, u(I, t)) \neq 1$ almost everywhere for $s \in [t_0, t]$:*

$$\log u(I, t) = \log I - \int_{t_0}^t \sum_{i \in X} b_{Xi}(\mathbf{p}_s, u(I, t)) \frac{d \log p_{is}}{ds} ds - \int_{t_0}^t \frac{d \log b_X(\mathbf{p}_s, u(I, t))}{ds} \frac{1}{\sigma(\mathbf{p}_s, u(I, t)) - 1} ds, \quad (12)$$

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 $\sigma(\mathbf{p}_s, u)$ for $s \in [t_0, T]$, Proposition 5 can be used to obtain the money metric utility function using similar procedures to the ones in Section 3.2. Proposition 5 is a consequence of Proposition 1. To derive it, we use the rate of change in the compensated budget share of X goods, $\frac{d \log b_X(\mathbf{p}_s, u(I, t))}{ds}$, to infer the compensated-budget-share-weighted rate of change in prices for the unobserved goods $\sum_{i \in Y} b_i(\mathbf{p}_s, u(I, t)) \frac{d \log p_{is}}{ds}$ given the elasticity of substitution $\sigma(\mathbf{p}_s, u(I, t))$. We require that $\sigma(\mathbf{p}_s, u(I, t)) \neq 1$ almost everywhere in order to do this, since when $\sigma(\mathbf{p}_s, u(I, t)) = 1$, the compensated share of b_X does not respond to the relative price of X and Y goods.

Compared to Proposition 1, the fixed point in Proposition 5 has some additional terms. First, the compensated elasticity of substitution $\sigma(\mathbf{p}_s, u(I, t))$ 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 rate of change in the budget share of X goods, $\frac{d \log b_X(\mathbf{p}_s, u(I, t))}{ds}$, 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 budget shares of matched households over time. Hence, there is also a fixed point in this term.

To better understand Proposition 5, it helps to consider the homothetic special case.

Example 1 (Homothetic preferences). Suppose that preferences are homothetic. In this

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

case, Proposition 5 simplifies to

$$\log u(I, t) = \log I - \int_{t_0}^t \sum_{i \in X} B_{Xi}(\mathbf{p}_s) \frac{d \log p_{is}}{ds} ds - \int_{t_0}^t \frac{d \log B_X}{ds} \frac{1}{\sigma(\mathbf{p}_s) - 1} ds. \quad (13)$$

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 expenditure function is CES, then $\sigma(\mathbf{p}_s)$ is a constant and we get

$$\log u(I, t) = \log I - \int_{t_0}^t \sum_{i \in X} B_{Xi}(\mathbf{p}_s) \frac{d \log p_{is}}{ds} ds - \frac{\log B_X(t) - \log B_X(t_0)}{\sigma - 1}. \quad (14)$$

Equation (14) is a version of the popular Feenstra (1994) formula.³² This formula is commonly used in the macroeconomics and trade literatures for adjusting price indices to account for missing price changes (typically those of new goods). Relative to this CES case, Proposition 5 allows the elasticity of substitution to vary as a function of prices, allows for nonhomotheticities, and does not impose parametric assumptions on preferences among the X goods and among the Y goods.

Relative to the homothetic special case in (13), the additional complication in (12) is that changes in the budget share of X and the elasticity of substitution must both be compensated. To see the issue, restate (12) using uncompensated budget shares as

$$\log u(I, t) = \log I - \int_{t_0}^t \sum_{i \in X} B_{Xi}(I_s^*, s) \frac{d \log p_{is}}{ds} ds - \int_{t_0}^t \frac{d \log B_X(I_s^*, s)}{ds} \frac{1}{\sigma(\mathbf{p}_s, u(I, t)) - 1} ds, \quad (15)$$

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

With more structure on the demand system, this expression can be further simplified. For example, suppose that the expenditure function in (11) can be written as

$$e(\mathbf{p}, U) = \left(\omega_X U^{\xi_X} e^X(\mathbf{p}^X, U)^{1-\gamma(U)} + \omega_Y U^{\xi_Y} e^Y(\mathbf{p}^Y, U)^{1-\gamma(U)} \right)^{\frac{1}{1-\gamma(U)}} \quad (16)$$

for any level of utility U . Let $V(\mathbf{p}, I)$ be the indirect utility function associated with (16). In this example, the elasticity of substitution varies as a function of utility but not as a function of relative prices, as in Fally (2022). With this restriction, the second integral in

³²The only (relatively inconsequential) difference between (14) and Feenstra (1994) is the assumption that e^X and e^Y also be homothetic CES aggregators.

equation (15) can be evaluated explicitly, and the expression simplifies to

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

where $\sigma(u(I, t)) = \gamma(V(\mathbf{p}_t, I))$. Of course, if the elasticity of substitution σ is also constant as a function of utility, then the denominator becomes just σ .

We now show that the compensated elasticity of substitution σ , which is the unknown term required to apply Proposition 5, can be expressed nonparametrically 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 $\sigma(\mathbf{p}_s, u)$ does not require data on unobserved prices. Denote by $\epsilon_X(I, s)$ the uncompensated elasticity of the budget share of X with respect to the price of the X bundle. That is, let $\epsilon_X(I, s)$ be the scalar that satisfies the following equation for each level of income I at each time s :

$$\sum_{k \in X} \frac{\partial \log B_X(I, s)}{\partial \log p_k} d \log p_k = \epsilon_X(I, s) \sum_{k \in X} B_{Xk}(I, s) d \log p_k.$$

Proposition 6 shows that the compensated elasticity of substitution between X and Y can be deduced given knowledge of $\epsilon_X(I, s)$ and income elasticities.

Proposition 6 (Identifying Substitution Elasticity of X and Y). *Suppose Assumption 1 holds. Let $\eta_i(I, t) - 1 = \frac{\partial \log B_i(I, t)}{\partial \log I}$ be the income elasticity of demand for each $i \in X$ at time t . Then, we have*

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

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

Proposition 6 shows that if we know income elasticities for all goods in X and can estimate the uncompensated elasticity of B_X with respect to prices in X , ϵ_X , then we can recover the relevant elasticity of substitution and apply Proposition 5. 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_X is more challenging, but we only require a single elasticity per income group and period. That is, the number of elasticities that needs to be estimated does not depend on the number of goods.

With more structure on the demand system, then even less information is required. We provide one example below.

Example 2 (Generalized nonhomothetic CES). Consider the case where the expenditure

function takes the form in (16). According to Proposition 6, the function $\sigma(\cdot)$ is determined by the following expression

$$\sigma(I) = 1 - \frac{\epsilon_X(I, t_0) + B_X(I, t_0) \sum_{i \in X} (\eta_i(I, t_0) - 1) B_{Xi}(I, t_0)}{(1 - B_X(I, t_0))}, \quad (17)$$

Since σ is not a function of relative prices, Proposition 6 needs to be applied only in the initial period, t_0 , to recover the shape of the σ function.³³ If σ also does not vary with utility, as in the example in Section 3.3, then equation (17) can still be used but only needs to be applied for one income group.

Relation to previous literature. When price data are unavailable or unreliable, a large strand of the literature relies on Feenstra (1994), which our method generalizes. A different strand, building on Hamilton (2001) and Costa (2001), estimates changes in welfare by inverting Engel curves. This procedure requires that relative budget shares be strictly monotone in income (i.e. homothetic preferences are ruled out). Atkin et al. (2024) provide a recent micro-founded treatment of this idea. To apply their method, one needs to estimate a compensated demand subsystem for the set of goods where prices are measured, a task that can suffer from a curse of dimensionality if the number of goods with observed prices is large. In their applications, they either rely on first-order approximations or use a CES subsystem to keep the estimation challenges manageable.

In contrast, we make stronger assumptions about preferences (separability rather than quasi-separability). In exchange, we do not require that budget shares be strictly monotone in income. More importantly, without making further assumptions, our approach only requires a single uncompensated price elasticity as a function of income in each period (rather than a compensated system). Given estimates of this elasticity, we can nonparametrically and nonlinearly back out the elasticity of substitution between the measured and unmeasured goods and use this to nonlinearly solve for welfare changes.

6 Empirical Illustration with Partially Observed Prices

As an illustration, we apply Proposition 5 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

³³In writing (17), we assume that $\epsilon_X(I, t_0)$ and $\eta_i(I, t_0)$ are known at t_0 . This is without loss of generality since Proposition 5 can be applied with time running forward $t > t_0$ and backward $t < t_0$. Furthermore, once we apply Proposition 5 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 Section 3.1.

into a subset of luxury services and the rest. That is, we assume that prices for leisure goods & services and catering are not reliably observed (see Table A.2 in Online Appendix Section A.5 for a description of these categories). These are the Y goods, which in our data account for roughly 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 Assumption 1 throughout this section.³⁴

Table 1: Elasticity of budget share of X with respect to price index of X

	OLS (1)	IV (2)	OLS (3)	IV (4)
$\sum_{i \in X} B_{Xi}(h, t) \Delta \log p_{it}$	0.144** (0.069)	0.073*** (0.019)	0.146** (0.069)	0.061*** (0.021)
$\sum_{i \in X} B_{Xi}(h, t) \Delta \log p_{it} \times \mathbf{1}(h \geq \text{median})$			0.005 (0.007)	0.025 (0.039)
F-stat		403,945		177,760
Quantile FE	Y	Y	Y	Y
Year FE	Y	N	Y	N
Obs	41,000	41,000	41,000	41,000

Notes: Columns (2) and (4) use the log difference in world oil prices as an instrument. All lags are two-year differences (results are similar for annual and triennial differences). The sample years are 1974–2017. Standard errors are clustered at the household quantile level (we have 1,000 quantiles). Two and three stars indicate statistical significance at the 5% and 1% level respectively.

To apply Proposition 5, we must estimate the compensated elasticity of substitution between X and Y . To do this, we group households into a thousand groups by quantiles of the spending distribution. We run the following regression

$$\Delta \log B_X(h, t) = \epsilon_X \sum_{i \in X} B_{Xi}(h, t) \Delta \log p_{it} + \text{controls} + \text{error},$$

where t is time, h is the quantile of the spending distribution, and ϵ_X measures the uncompensated elasticity of the budget share of X goods with respect to the price of X goods. To check for heterogeneity, we allow this elasticity to depend on whether quantile h is above or below the median.

³⁴As a proof-of-concept, we provide a test for separability between X and Y goods in Online Appendix Section A.6. To do this, we estimate whether the relative compensated budget shares of $i, j \in Y$ respond to changes in the price of $k \in X$.

We estimate this regression by OLS. Given that we include year fixed effects, identification comes from variation across households in the price change of the X bundle. We also instrument the right-hand side variable using world oil prices (in which case we cannot include year fixed effects). The identification strategy requires that oil price shocks exogenously move the price of goods versus services. We view our exercises as a proof of concept rather than a full-fledged elasticity estimation.

The results of this regression are reported in Table 1. The first two columns assume that ϵ_X does not vary as a function of expenditures, and the last two columns allow for the possibility that ϵ_X varies as a function of expenditures. Since the second row is insignificant with small coefficients, we assume ϵ_X does not vary by quantile. We also assume that ϵ_X does not vary as a function of time (we check for subsample stability by rerunning the regression on the first and second half of the time period).

The OLS and IV point estimates are $\epsilon_X = 0.14$ and $\epsilon_X = 0.07$, though with overlapping confidence intervals. For concreteness, we take $\epsilon_X = 0.14$ and apply Proposition 6 to recover an estimate of the compensated elasticity of substitution σ for each value of I and in each time period.³⁵ The results, for the 25th, 50th, and 75th percentile of the expenditures distribution are plotted in Figure 6. The estimated elasticity is below one, so X and Y are complements, and increasing in income. Richer households are more willing to substitute between X and Y goods than poorer households.

Figure 7 uses these estimates of the compensated elasticity and computes the money metric. The resulting money metric is plotted against the money metric from Section 4 when we assumed that all prices are perfectly observed. For low-income households, the two money metrics are quite similar and both are below real consumption (computed using an aggregate chain-weighted price deflator assuming that all prices are observed). However, for households with high incomes, the money metric calculated using Proposition 5 is lower than the one calculated using Proposition 1. The fact that the dotted blue line is lower than the dashed yellow line for rich households suggests that, for these households, prices in Y have risen more than the official price data suggest.

Figure 8 shows the percent difference between the money metric with observed prices and the money metric with unobserved prices for different deciles of expenditures as well as the breakdown of the difference into two terms. The first is the difference between

³⁵Figure A.4 in the Online Appendix displays results using the IV point estimates instead. Results are qualitatively similar, but the adjustment to the money metric values for rich households is larger than in Figure 7 because the implied elasticity of substitution σ is closer to one for richer households. Figure A.4 also shows that the results are very similar to our baseline if we calibrate the compensated elasticity of substitution between X and Y goods to be constant in both the time series and the cross-section and equal to 0.5.

Figure 6: Compensated elasticity of substitution between X and Y .




Figure 7: Money metric $e(p_{1974}, v(p_{2017}, I))$ and real consumption.

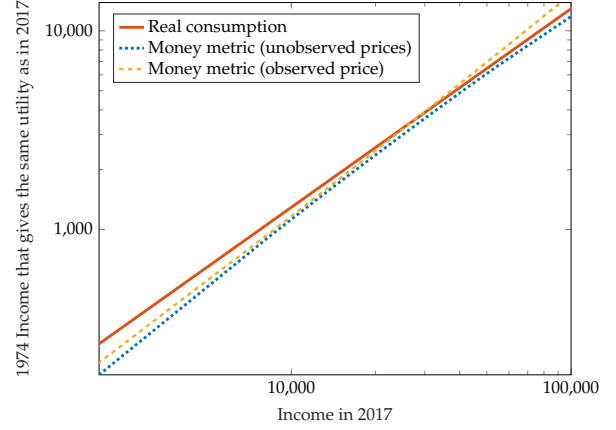



Figure 8: Log difference in estimated money metrics under observed and unobserved prices by decile of the expenditures distribution.



overall inflation and inflation for X goods implied by the two methods:

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

These are the dark blue bar graphs in Figure 8. The remainder is the adjustment due to changes in the budget share of X goods, similar to the Feenstra (1994) adjustment:

$$\frac{\int_{1974}^{2017} \frac{1}{\sigma(p_s, u) - 1} \frac{d \log b_X(p_s, u)}{ds} ds}{\int_{1974}^{2017} \sum_{i=1}^N b_i(p_s, u) \frac{d \log p_{is}}{ds} ds}. \quad (18)$$

These are the light yellow bar graphs in Figure 8. This decomposition shows that inflation

among X goods has tended to be higher than among all goods by roughly the same amount (around 1 percentage point) for all 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.

Figure 9: Compensated changes in the budget share B_X for different I percentiles in 2017.





Figure 10: Compensated and uncompensated changes in the budget share B_X for the median household in 2017.



To better understand this, we investigate how compensated expenditures on X goods have changed over time. Figure 9 shows the compensated budget share on X goods for households at three different points in the distribution: the 10th, 50th, and 90th percentiles in 2017. For poor households, there was almost no change on expenditures on X goods. This explains why the adjustment term in (18) is small for these households. For the median household, there was a modest decrease in the share of spending on X goods. Since X and Y are complements, this indicates that the relative price of Y goods rose relative to X goods for these households. Finally, for the richest households, there was a fairly dramatic reduction in their spending on X goods from around 74% to around 63%. This suggests that for these households, the relative price of Y goods rose fairly rapidly compared to X goods. This explains why the adjustment term, (18), for these households is large and negative. Furthermore, since the elasticity of substitution for rich households is closer to one, the implied difference in the relative price of X and Y goods is larger. This explains why the money metric according to Proposition 5 (the dotted blue line in Figure 7) has a flatter slope than the money metric calculated according to Proposition 1 (the dashed yellow line in Figure 7).

These difference in compensated expenditures are not mirrored in uncompensated expenditures. Figure 10 compares the compensated and uncompensated changes in expen-

ditures for the median household. Whereas, for the median household, the compensated expenditures on X goods declined somewhat over time, the uncompensated expenditures on X goods increased very strongly. Intuitively, a household in 1974 with nominal expenditures equal to the median of the expenditure distribution in 2017 is actually fairly rich. Such a household spends relatively less on goods (X) and relatively more on services (Y). As we roll time forward, such a household is effectively becoming poorer, due to inflation, and this causes the expenditures on the X goods to rise due to income effects. That is, the income effect overwhelms the substitution effect.

7 Conclusion

In this paper, we propose a method to construct money metric representations of utility—an essential input to measuring welfare-relevant growth—using repeated cross-sectional data. Our method does not require any estimation when the data on prices is comprehensive, aside from interpolation of how budget shares vary with income and time. If the data on prices is incomplete, the method can still be used under a separability assumption on preferences and knowledge of one uncompensated elasticity of substitution.

Whether prices are fully or partially observed, the unifying idea in both cases is that money metric utility can be calculated using observed demand of matched households in the cross-sectional distribution over time. Doing so involves solving a simple fixed point equation in terms of observable variables.

Despite its advantages, our approach does not allow for preferences to vary in arbitrary and unobserved ways in the cross section or the time series, and requires that all consumers face common prices. Furthermore, we have abstracted from intertemporal choice. Relaxing these assumptions is an interesting avenue for future work.

References

- Atkin, D., B. Faber, T. Fally, and M. Gonzalez-Navarro (2024, February). Measuring welfare and inequality with incomplete price information. *Quarterly Journal of Economics* 139(1).
- Banks, J., R. Blundell, and A. Lewbel (1997). Quadratic engel curves and consumer demand. *Review of Economics and statistics* 79(4), 527–539.
- Baqaei, D. R. and A. Burstein (2023). Welfare and output with income effects and taste shocks. *The Quarterly Journal of Economics* 138(2), 769–834.

- Blackorby, C., D. Primont, and R. R. Russell (1998). Separability: A survey. *Handbook of utility theory* 1, 51–92.
- 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.
- Divisia, F. (1926). L’indice monétaire et la théorie de la monnaie. *Revue d’économie politique* 40(1), 49–81.
- 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 (2023). Growing like india—the unequal effects of service-led growth. *Econometrica* 91(4), 1457–1494.
- 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 (2024). Measuring growth in consumer welfare with income-

- dependent preferences: Nonparametric methods and estimates for the united states. *Quarterly Journal of Economics* 139(1).
- 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.
- Straub, L. (2019). Consumption, savings, and the distribution of permanent income. *Unpublished manuscript, Harvard University*.