Measuring Growth in Consumer Welfare with Income-Dependent Preferences

Nonparametric Methods and Estimates for the United States

Xavier Jaravel[†] Danial Lashkari[‡]

July 25, 2023

Abstract

How should we measure changes in consumer welfare given observed data on prices and expenditures? This paper proposes a nonparametric approach that holds under arbitrary preferences that may depend on observable consumer characteristics, e.g., when expenditure shares vary with income. Using total expenditures under a constant set of prices as our money metric for real consumption (welfare), we derive a principled measure of real consumption growth featuring a correction term relative to conventional measures. We show that the correction can be nonparametrically estimated with an algorithm leveraging the observed, cross-sectional relationship between household-level price indices and household characteristics such as income. We demonstrate the accuracy of our algorithm in simulations. Applying our approach to data from the United States, we find that the magnitude of the correction can be large due to the combination of fast growth and lower inflation for income-elastic products. Setting reference prices in 2019, we find that (i) the uncorrected measure underestimates average real consumption per household in 1955 by 11.5%, and (ii) the correction reduces the annual growth rate from 1955 to 2019 by 18 basis points, which is larger than the well-known "expenditure switching bias" over the same time horizon.

For stimulating conversations and thoughtful comments and feedback, we thank David Baqaee, Ariel Burstein, Ryan Chahrour, Ben Faber, Thibault Fally, Robert Feenstra, Shakeeb Khan, Arthur Lewbell, Marshall Reinsdorf, Robert Ulbricht, David Weinstein, and seminar participants at Berkeley, the Paris School of Economics, the NBER Summer Institute, and the Society for Economic Dynamics. We are grateful to Sylvia Tian for exceptional research assistance. This paper was first circulated in November 2021 under the title "Nonparametric Measurement of Long-Run Growth in Consumer Welfare." We are grateful to the Philip Leverhulme Trust for generous financial support. The empirical results of this paper rely in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at the University of Chicago Booth School of Business. The conclusions drawn from the NielsenIQ data are ours and do not reflect the views of NielsenIQ. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. Likewise, the views expressed here are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of New York or the Federal Reserve System.

[†]London School of Economics. Email: x.jaravel@lse.ac.uk . Website: https://www.xavierjaravel.com.

[‡]Federal Reserve Bank of New York. Email: danial.lashkari@ny.frb.org. Website: https://www.daniallashkari.com.

1 Introduction

How should we measure long-run changes in consumer welfare? Classical demand theory shows that intuitive index number formulas, which aggregate observed changes in consumed quantities and prices, may provide precise measures of the change in living standards. However, this powerful insight requires the crucial assumption that the composition of demand remains independent of consumer income (see, e.g., Diewert, 1993). This so-called homotheticity assumption runs counter to the empirical regularity that demand for many goods and services systematically depends on income, a fact known since at least Engel (1857). It also belies the growing empirical evidence on sizable differences in the rates of inflation in the cost-of-living experienced by different income groups in the United States, with lower inflation rates for higher-income groups.[1]

Despite this important and well-known theoretical limitation, classical price index formulas remain widely used in practice due to their simplicity, flexibility, and generality. Little is known about potential biases arising from the restrictive homotheticity assumption in the resulting measures of long-run growth in living standards. Current alternatives require us to specify and estimate the structure of the demand system, a task that leaves open many questions about the choices of functional forms and identification strategy. For instance, Baqaee and Burstein (2021) have recently offered an approach that relies on the knowledge of the elasticities of substitution across goods to construct measures of welfare growth (see also Samuelson and Swamy, 1974).

In this paper, we develop a novel approach for measuring welfare change that allows for flexible dependence of the patterns of demand on income and other sources of observed heterogeneity without the need for functional form assumptions. Compared to the standard setting, the only additional data requirement is access to a cross-section of product prices and quantities for consumers with heterogenous incomes. Such data is widely available through standard surveys of consumption expenditure. Our approach nonparametrically estimates the cross-sectional dependence of measured price index formulas on consumer income, which we show is sufficient to provide precise approximations for a theoretically consistent measure of real consumption. The approach remains valid for any continuously differentiable preferences under observable sources of heterogeneity.

We apply our method to account for the nonhomotheticity of demand in measuring growth in consumer welfare in the United States from 1955 to 2019. In addition to improving the measurement of long-run growth and inflation inequality, our new approach can have important policy implications, such as the indexation of the poverty line and a more efficient targeting of welfare benefits. This approach also provides a blueprint for distributional national accounts

[1] See, for example, Kaplan and Schulhofer-Wohl (2017), Jaravel (2019), Argente and Lee (2021), Klick and Stockburger (2021), and Jaravel (2021).

(Piketty et al., 2018) that allow for nonhomotheticity and inflation inequality.

We begin with the basic theory of the exact measurement of welfare change under stable preferences along a path of smoothly changing prices. We define real consumption as the expenditure required to achieve a certain level of welfare under constant prices fixed at a base period (money metric). Given this definition, there exists a mapping from real consumption to total consumer expenditure at any point in time. We show that we can recover this mapping as the solution to a differential equation defined in terms of the *Divisia function*. This function extends the notion of the Divisia index, which is a standard measure of the change in the cost-of-living. This measure is defined at any point in time for a given consumer as the mean of price growth across goods, weighted by the expenditure shares of the consumer. Since in our setting expenditure shares generically depend on income, it is natural to define the index as a function of total expenditure. Our results show that such a Divisia function summarizes all the information in the demand system that is relevant to recovering real consumption.

When preferences are homothetic, the Divisia function is constant in total expenditure and at any point in time equal to the Divisia index of any consumer. Our key differential equation in this case has a simple solution: the growth in real consumption is given by growth in total consumer expenditure, deflated by the value of the Divisia index. Since index formulas approximate the Divisa index for each consumer in the data, we can chain them over time to construct approximate measures of real consumption under homotheticity.

When preferences are nonhomothetic, the differential equation implies that we need to multiply the deflated total expenditure by a *nonhomotheticity correction* factor at any point in time. For each consumer, this correction is governed by the elasticity of the mapping between real consumption and total expenditure. Under homotheticity, since the mapping is always linear, the elasticity and the correction factor are both unity. With nonhomotheticity, however, the curvature of the mapping changes over time and the correction factor deviates from unity if price inflation varies as a function of income. Importantly, we show how this correction implies a systematic dependence of the measures of real consumption growth on the base vector of prices chosen to express them.

To see the intuition behind this correction, consider a setting where consumer welfare is rising over a time horizon during which inflation rates are lower for goods with higher income elasticity (luxuries). Fixing prices in the initial period as our base, real consumption is by definition linear in (and identical to) total expenditure in the initial period. As time passes, the *relative* cost of achieving higher levels of real consumption falls, since relative prices are falling for goods more heavily consumed by the rich. In other words, the mapping between real consumption and total expenditures becomes more concave over time.[2] Hence, a given rise in total expenditure translates

[2]One way to understand this change in concavity is that it accounts for the cumulative effect of the past inflation

into increasingly larger gains in real consumption as consumers become richer. The conventional approach assumes a linear mapping and thus ignores the gradual fall in its curvature, leading to an underestimation of the growth of real consumption under the initial base period in this case.[3] Our nonhomotheticity correction accounts for changes in the curvature of this mapping to accurately measure growth in terms of any base period.

Having characterized the properties of the exact mapping from real consumption to total expenditure, we next show how we can approximate it in settings where we only have discrete observations of consumer choices and where we do *not* know the underlying preferences. The key observation is that we can use the variations in the price index formulas across consumers/households with different levels of income to nonparametrically approximate the Divisia function. Using this insight, we can construct approximate solutions to our key differential equation to recover the values of real consumption. The main assumptions are that the preferences are smooth and identical across consumers.

Our baseline approximation algorithm is fairly simple and intuitive. In the base period, total expenditure by definition coincides with real consumption. This allows us to nonparametrically approximate the nonhomotheticity correction as the elasticity of the observed price index formulas of different consumers with respect to their total expenditure. Using this elasticity, we obtain approximations for the values of real consumption for each household in periods immediately before or after the base period. We can then recursively apply the same strategy in subsequent periods to approximate the values of nonhomotheticity correction and real consumption over the entire period of interest.

We provide several refinements and extensions for this baseline algorithm, depending on the alternative choices made about the nature of the approximations. For instance, we derive algorithms that integrate our key differential equation up to first or second orders of approximation in terms of the change in prices. We also consider alternative choices for the price index formula. Using geometric, Laspeyres, and Paasche indices, we can construct first-order approximations for the Divisia function, whereas by relying on Törnqvist, Fisher, or Sato-Vartia we can construct

inequality. Consumers who were previously poor may not have immediately benefitted from the fall in the price of income elastic products in the past, since they consumed very little of those luxuries. However, they benefit from those past price changes today if their nominal income rises and they begin to consume those luxuries. The rise in their nominal income now translates to higher real consumption growth because of the fall in luxury prices accumulated since the base period, which has led to a concave mapping between nominal and real consumption. See Oberfield (2022) for a manifestation of this idea in a model of growth featuring inflation inequality.

[3] If we instead express real consumption in terms of constant final period prices as our base, the same logic implies that conventional approach overestimates the growth in all preceding periods. In this case, since total consumer expenditure is identical to real consumption in the final period, it must be a convex function of real consumption in all prior periods. This leads to overestimating the growth of real consumption when using the final period as base. In Section 2.2, we show formally that the sign of the bias in growth measurement induced by the nonhomotheticity correction inherently depends on the choice of the base period.

second-order approximations.[4]

We demonstrate the accuracy of our different algorithms using a simulation with known preference parameters, relying on the nonhomothetic CES (nhCES) preferences of Comin et al. (2021). In an environment featuring growth in real consumption, we confirm that our procedure accurately recovers the evolution of the exact index using the observed cross-sectional data, without any knowledge of the underlying preference parameters.

In the empirical part of the paper, we apply our approach to data from the United States and quantify the magnitude of the bias in conventional measures of real consumption growth that ignore nonhomotheticity effects. We build a new linked dataset providing price changes and expenditure shares at a granular level from 1955 to 2019 across percentiles of the income distribution. This dataset combines several data sources, primarily drawing from disaggregated data series available from the Consumer Price Index (CPI) and the Consumer Expenditure Survey (CEX). This new linked dataset allows us to provide evidence on inflation inequality over a long time horizon, thus extending prior estimates that have focused on shorter time series. Computing inflation using group-specific price index formulas, we find that inflation inequality is a long-run phenomenon. Using a geometric index formula, we find that cumulative inflation from 1955 to 2019 varies from 700% at the top of the income distribution to 875% at the bottom.

Since richer households experience lower inflation rates in the data, our theory implies that, at any point other than the base period, consumers are actually *better off* than that suggested by conventional uncorrected measures. Intuitively, when we look into the past from the perspective of today's prices, we observe that (i) households were on average poorer 65 years ago, that is, they had stronger preferences for necessities; and (ii) necessities were cheaper. These empirical patterns imply that consumer welfare was higher 65 years ago when accounting for nonhomotheticity effects. Symmetrically, looking at today's economy from the perspective of prices in a distant period in the past, we observe that (i) households got on average richer and (ii) luxuries got cheaper, implying higher average welfare today if we account for nonhomotheticity effects.

Empirically, we find that the magnitude of the nonhomotheticity correction can be large. For example, taking base prices in 2019, we find that the uncorrected measure underestimates average real consumption (per household) in 1955 by about 11.5%.[5] The uncorrected measure

[4]Establishing the second-order equivalence of the Sato-Vartia index with superlative indices such as Fisher and Törnqvist constitutes another contribution of our paper. The order of approximation is given in terms of the annual growth in total expenditure and prices across goods, as discussed in Section 2.3.

[5]We find that the magnitude of the bias is similar across income percentiles. Note that our goal is to uncover the correct measures of real consumption at the income-percentile (synthetic household) level, without taking a stance on the aggregation of welfare. In other words, we report our measures of average real consumption only as a summary of the results across income percentiles, to make them comparable with the corresponding measures reported in the official statistics. Since our proxy for real consumption is a money metric utility, different approaches to aggregating these values across households yield different social welfare functions (see Blackorby and Donaldson, 1988; Slesnick, 1991; Bosmans et al., 2018, for the properties of money metric social welfare functions).

of cumulative real consumption growth is 270% over this period, or 2.07% growth annually. In contrast, with the nonhomotheticity correction and 2019 base prices, cumulative consumption growth falls to 232%, or an annualized growth rate of 1.89% per year.[6] Thus, in this case the nonhomotheticity correction reduces the annual growth rate from 1955 to 2019 by 18 basis points, which is larger than the difference of 11 basis points between the Laspeyres and Paasche indices over the same time horizon. These results show that the magnitude of the nonhomotheticity correction can be as large as the well-known "expenditure switching bias" (or "substitution bias") affecting the Laspeyres and Paasche indices, which demonstrates its quantitative relevance.

Finally, we show in an extension that our strategy generalizes to settings where preferences systematically vary with consumer characteristics, e.g., age, family size, education, etc. When these characteristics evolve over time, we need to adjust our measures using characteristic correction factors that capture the elasticity of the mapping from real consumption to total expenditure with respect to the changing characteristics. We characterize this mapping and provide algorithms to approximate the resulting corrections, using the cross-sectional variations in price index formulas and consumer characteristics. Empirically, we apply our algorithm to quantify the adjustment to average real consumption implied by consumer aging in the United States. We document a strong positive relationship between consumer age and inflation, which alters the measurement of real consumption because the average consumer age increases over time. We find that the implied adjustments to real consumption are economically meaningful but much smaller than the nonhomotheticity correction, which justifies our focus on the latter.

Prior Work Our paper builds on and contributes to three strands of the literature. First, we extend the literature on index number theory (e.g., Pollak, 1990; Diewert, 1993), which has enabled transparent and consistent comparisons of the aggregate measures of consumption and production over time and space relying only on observables. As emphasized by Samuelson and Swamy (1974), many classical results do not generalize beyond settings involving homotheticity in preferences. Under nonhomotheticity, Diewert (1976) has showed that one can still rely on the conventional price index formulas to measure changes in welfare locally. However, we show that these results do not generalize to welfare comparisons over long time horizons. We provide a detailed discussion of the contrast between our results and these classical results in Section 2.3.7.[7]

Second, we advance a growing literature raising the point that standard price index formulas suffer from a bias due to nonhomotheticities, whose magnitude is related to the covariance

[6] The sign and magnitude of the nonhomotheticity correction to the measurement of real consumption growth inherently depends on the choice of the base period, which we discuss further in Section 3.

[7] Our approach assumes utility maximization, and thus contrasts with the approach of Blundell et al. (2003), who rely on revealed preference inequalities to develop a test for the axioms of revealed preference, and propose lower and upper bounds on the true cost-of-living as a by-product of their strategy.

between income elasticities and price changes (e.g., Fajgelbaum and Khandelwal, 2016; Atkin et al., 2020; Baqaee and Burstein, 2021). In particular, Baqaee and Burstein (2021) have recently highlighted the failure of standard measures of real consumption to capture theoretically consistent welfare measures. They suggest relying on the estimates of the elasticities of substitution to account for the role of nonhomotheticity.[8] In contrast, we provide a nonparametric approach that does not require specifying the underlying demand functions. The importance of the covariance between income elasticities and inflation for measuring welfare change is also noted by Fajgelbaum and Khandelwal (2016) and Atkin et al. (2020). Fajgelbaum and Khandelwal (2016) measure changes in welfare gains from trade liberalization across different income groups in a parametric setting and under the assumption of an AIDS demand system (Deaton and Muellbauer, 1980).[9] Atkin et al. (2020) consider the problem of welfare measurement in the absence of reliable price data, and use separability assumptions on the structure of preferences to infer welfare from shifts in the Engel curves. For this procedure to hold without the need for estimation of structural elasticities of substitution, Atkin et al. (2020) rule out the types of covariance patterns that lead to large nonhomotheticity corrections in our framework. In summary, while this literature provides parametric corrections for the bias, our contribution is to provide a nonparametric correction that remains valid under arbitrary preferences where all consumer heterogeneity is in terms of observables.

Third, we contribute to the literature on the measurement of inflation inequality (e.g., Hobijn and Lagakos, 2005; McGranahan and Paulson, 2006; Kaplan and Schulhofer-Wohl, 2017; Jaravel, 2019; Argente and Lee, 2021). Prior work on inflation inequality has posited the existence of separate homothetic indices for different income groups. We apply our methodology to provide estimates of inflation inequality that are robust to potential biases arising from nonhomotheticities. Using our new linked dataset covering the period 1955-2019 in the United States, we apply our methodology to the measurement of short-, medium-, and long-run growth in real consumption, and we quantify the magnitude of the bias stemming from the nonhomotheticity correction.

The remainder of this paper is organized as follows: Section 2 presents our theory, approximation algorithms, and simulations. Section 3 reports the empirical analysis, and Section 4 generalizes our approach to settings where preferences vary with observable consumer characteristics. Several proofs and additional results are reported in the appendix.

[8] Baqaee and Burstein (2021) additionally study the consequences of the endogeneity of prices in general equilibrium, as well as unobserved heterogeneity, e.g. taste shocks. The latter effects have also recently been considered by Redding and Weinstein (2020). We note that, subsequent to our paper, Baqaee et al. (2022) proposed an alternative to our algorithms. We discuss the close connections between their approach and ours in Appendix B.3.

[9] An earlier literature showed how parametric AIDS specifications can be used to make welfare comparisons over time (Oulton, 2008) or across countries (Feenstra et al., 2009) in the presence of nonhomotheticities, estimating only income elasticities and without the need to estimate elasticities of substitution.

2 Measuring Welfare Changes under Nonhomotheticity

In this section, we present our theory for the exact measurement and empirical approximation of real consumption growth under preference nonhomotheticity. Section 2.1 introduces the notation and defines the main concepts used for the measurement of welfare, cost-of-living, and real consumption. Section 2.2 presents the theory for the exact measurement of welfare growth assuming the knowledge of a specific function that combines information on consumer demand with price changes. Section 2.3 derives our approximate results in terms of observable data. Finally, in Section 2.4 we perform a simulation to illustrate and validate the accuracy of our approach.

2.1 Definitions

2.1.1 Real Consumption and the True Price Index

Consider consumer preferences in the space of I products characterized by a utility function $U(\boldsymbol{q})$ where $\boldsymbol{q} \equiv (q_i)_{i=1}^{I}$ is the (nonnegative) vector of quantities consumed of each good. We assume that the corresponding expenditure function $E(u;\boldsymbol{p})$, characterizing expenditure required to achieve utility u under vector of prices $\boldsymbol{p} \equiv (p_i)_{i=1}^{I}$, is second-order continuously differentiable. Moreover, consider a path of prices $\boldsymbol{p}_t$ over the time interval $t \in [0, T]$, and let $\boldsymbol{s} = \boldsymbol{\omega}_t(y)$ denote the vector of expenditure shares across goods as a function of total expenditure y under these preferences at time t, with $y \equiv \sum_i p_i q_i$ and $s_i \equiv p_i q_i / y$. The function $\boldsymbol{\omega}_t(\cdot)$ characterizes the Marshallian demand for the vector of prices prevailing at time t.[10] Since we do not restrict the preferences to be homothetic, Marshallian demand depends on total spending y.

We begin by defining our concept of real consumption as a money metric for consistent measurement of welfare over time.

Definition 1 (Real Consumption)**.** For a given vector of prices $\boldsymbol{p}_b$ (with $0 \leq b \leq T$), define *real consumption under constant time-b (base) prices* as a monotonic transformation $M_b(\cdot)$ of utility u given by

$$c^b = M_b(u) \equiv E(u;\boldsymbol{p}_b). \tag{1}$$

Equation (1) constitutes our money metric for welfare for a consumer with utility u, which gives the minimum expenditure needed to achieve that level of utility under the vector of prices prevailing at time b. Since real consumption is defined with reference to base time period b, we must include b in our notation for real consumption, c^b. For brevity, we will often drop the superscript to simplify the expressions whenever it is clear that the base b is fixed.

[10] From Shephard's lemma, we have $\omega_{i,t}(y) \equiv \partial \log E(u;\boldsymbol{p}_t) / \partial \log p_{i,t}$ subject to $y = E(u;\boldsymbol{p}_t)$.

Definition 1 constructs a fixed mapping from utility to real consumption that does not vary with time. We now define a time-dependent function $\chi_t^b(\cdot)$ that maps real consumption c under base period b to the value of the total expenditure required to achieve that level of real consumption under current prices $\boldsymbol{p}_t$. Formally, this function is given by

$$\chi_t^b(c) \equiv E\left(M_b^{-1}(c); \boldsymbol{p}_t\right), \tag{2}$$

where $M_b^{-1}(c)$ is the level of utility corresponding to real consumption c. Note that for a given consumer with real consumption c_t^b and total expenditure y_t at time t, we have $y_t = \chi_t^b\left(c_t^b\right)$. Moreover, by definition we have $c = \chi_b^b(c)$ for all c.

Corresponding to Definition 1, we define the growth in real consumption between periods t_0 and t under the base vector of prices at time b as the ratio $c_t^b/c_{t_0}^b$, which is also a (standard-of-living) quantity index. We also define an index for the inflation in the cost-of-living corresponding to the level of consumption c between periods t_0 and t.

Definition 2 (True Price Index)**.** Define the cost-of-living price index $\mathscr{P}_{t_0,t}^b(c)$ for a consumer with real consumption c (defined under base time period b) between periods t_0 and t ($0 \leq t_0, t \leq T$) as

$$\mathscr{P}_{t_0,t}^b(c) \equiv \frac{\chi_t^b(c)}{\chi_{t_0}^b(c)}. \tag{3}$$

Let us specifically consider the true price index defined between the base period b and the current period t, which satisfies $c \equiv \chi_t^b(c)/\mathscr{P}_{b,t}^b(c)$. Since $y = \chi_t^b(c)$, knowing this index allows us to find real consumption by deflating total expenditure. Using Definitions 1 and 2, we can write the following relationship between real consumption growth and the true price index between periods t_0 and t:

$$\frac{c_t^b}{c_{t_0}^b} = \frac{y_t/\mathscr{P}_{b,t}^b(c_t)}{y_{t_0}/\mathscr{P}_{b,t_0}^b\left(c_{t_0}\right)} = \frac{y_t/y_{t_0}}{\mathscr{P}_{t_0,b}^b\left(c_{t_0}\right) \times \mathscr{P}_{b,t}^b(c_t)}. \tag{4}$$

Equation (4) shows that the growth in real consumption for a consumer under any base period b is given by deflating the growth in the nominal consumer expenditure by a *composite* true price index. This composite price index is the product of the true price index between the initial period t_0 and the base period b, $\mathscr{P}_{t_0,b}^b\left(c_{t_0}\right)$, and the true price index between the base period b and the final period t, $\mathscr{P}_{b,t}^b(c_t)$. Crucially, the former index is evaluated at the initial level of real consumption c_{t_0} while the latter is evaluated at the final level of real consumption c_t.[11]

[11] In such a pairwise welfare comparison between periods t_0 and t, the specific choice of the initial year t_0 as base leads to the concept of equivalent variation (EV) as our measure of welfare growth, which we can write as

Homothetic Preferences Let us consider the restriction that the underlying preferences are homothetic; that is, the composition of demand does not depend on the level of utility. The utility function $U(\cdot)$ is homothetic if (and only if) we can write the expenditure function as $E(u;\boldsymbol{p}) = P(\boldsymbol{p}) \cdot F(u)$, for some unit expenditure function $P(\cdot)$ and some canonical homothetic cardinalization $F(\cdot)$ of utility (Diewert, 1993). Correspondingly, from Definition 2, the true price index $\mathscr{P}^{b}_{t_0,t}(c)$ between any two time periods t_0 and t takes the same value independent of the level of real consumption c and the choice of the base period b. Equation (4) then simplifies to[12]

$$\frac{c^b_t}{c^b_{t_0}} = \frac{y_t/y_{t_0}}{\mathscr{P}^{b}_{t_0,t}(c)}, \qquad \text{for any } c \text{ and for any } b, \tag{5}$$

implying that we can deflate nominal consumption growth by the true index between the initial and final periods for any level of real consumption.

2.1.2 Price Index Formulas

The indices defined in Section 2.1.1 are structural, in the sense that they require the knowledge of the underlying consumer preferences. In contrast, standard price index formulas can be computed only in terms of observed expenditures and prices. An index formula is a positive-valued function $\mathbb{P}\left(\boldsymbol{p}_{t_0}, \boldsymbol{s}_{t_0}; \boldsymbol{p}_t, \boldsymbol{s}_t\right)$ of a pair of initial and final vectors of prices and expenditure shares, which aggregates the changes into a single index. The most common examples include Laspeyres $\mathbb{P}_L$, Paasche $\mathbb{P}_P$, and geometric $\mathbb{P}_G$ indices, which only use one vector of expenditure shares in the initial or final periods:

$$\mathbb{P}_L \equiv \sum_i s_{i,t_0}\left(\frac{p_{i,t}}{p_{i,t_0}}\right), \qquad \mathbb{P}_P \equiv \left(\sum_i s_{i,t}\left(\frac{p_{i,t_0}}{p_{i,t}}\right)\right)^{-1}, \qquad \mathbb{P}_G \equiv \prod_i \left(\frac{p_{i,t}}{p_{i,t_0}}\right)^{s_{i,t_0}}, \tag{6}$$

where we have suppressed the arguments $\left(\boldsymbol{p}_{t_0}, \boldsymbol{s}_{t_0}; \boldsymbol{p}_t, \boldsymbol{s}_t\right)$ to avoid repetition. As is well-known, the above indices do not account for the substitution effects that change the composition of expenditure between the two periods. Important alternatives that use both initial and final expenditure shares and account for substitution effects include the Fisher $\mathbb{P}_F$, Törnqvist $\mathbb{P}_T$, and Sato-Vartia

$EV = c^{t_0}_t/c^{t_0}_{t_0} = \left(y_t/y_{t_0}\right)/\mathscr{P}^{t_0}_{t_0,t}\left(c^{t_0}_t\right)$. Alternatively, choosing the final period t as the base leads to the concept of compensating variation (CV), given as $CV = c^t_t/c^t_{t_0} = \left(y_t/y_{t_0}\right)/\mathscr{P}^{t}_{t_0,t}\left(c^t_{t_0}\right)$.

[12] Homotheticity is a necessary and sufficient condition for the true price index $\mathscr{P}^{b}_{t_0,t}(c)$ to be independent of c and for the growth in real consumption $c^b_t/c^b_{t_0}$ to be independent of the base b. Samuelson and Swamy (1974) refer to this result as the *homogeneity theorem.*

$\mathbb{P}_S$ index formulas defined as

$$\mathbb{P}_F \equiv (\mathbb{P}_P \cdot \mathbb{P}_L)^{\frac{1}{2}}, \qquad \mathbb{P}_T \equiv \prod_{i=1}^{I} \left(\frac{p_{i,t}}{p_{i,t_0}} \right)^{\bar{s}_{T,i}}, \qquad \mathbb{P}_S = \prod_i \left(\frac{p_{i,t}}{p_{i,t_0}} \right)^{\bar{s}_{S,i}}, \tag{7}$$

where the Fisher index is the geometric mean of the Laspeyres and Paasche, and where the Törnqvist weights are defined as $\bar{s}_{T,i} \equiv \frac{1}{2}\left(s_{i,t_0} + s_{i,t}\right)$ and the Sato-Vartia weights are proportional to $\bar{s}_{S,i} \propto \frac{s_{i,t}/s_{i,t_0}}{\log\left(s_{i,t}/s_{i,t_0}\right)}$ and sum to 1. As we will see in Section 2.3 below, we can rely on these index formulas to approximate the true price index and real consumption growth.

2.2 Exact Measurement of Welfare Change under Nonhomotheticity

In this section, we show how to construct the mapping $\chi_t^b(\cdot)$ from real consumption to total expenditure, given observable functions that characterize the evolution of expenditure shares $\boldsymbol{\omega}_t(\cdot)$ and prices $\boldsymbol{p}_t$. We first use the paths of prices and the expenditure share function to define a *Divisia function* $D_t(\cdot)$ of total expenditure at time t as

$$\log D_t(y) \equiv \sum_i \omega_{i,t}(y) \frac{d \log p_{it}}{dt}. \tag{8}$$

The following proposition shows that the knowledge of this function is sufficient to fully characterize the evolution of the mapping $\chi_t^b(c)$, and it thus summarizes all the information in the demand function that is relevant to recovering real consumption over time.[13]

Proposition 1. *Consider a path of prices $\boldsymbol{p}_t$ and preferences that lead to the Divisia function $D_t(\cdot)$ over the interval $[0, T]$. The mapping $\chi_t^b(\cdot)$ from real consumption to total expenditure is the solution to the following partial differential equation with the boundary condition $\chi_b^b(c) = c$:*

$$\frac{\partial \log \chi_t^b(c)}{\partial t} = \log D_t\left(\chi_t^b(c)\right). \tag{9}$$

In addition, for any path of total nominal expenditure y_t over the interval, the growth in real consumption, defined under period-b constant prices, at any point in time satisfies

$$\frac{d \log c_t^b}{dt} = \left(\frac{\partial \log \chi_t^b\left(c_t^b\right)}{\partial \log c_t^b} \right)^{-1} \times \left(\frac{d \log y_t}{dt} - \log D_t(y_t) \right). \tag{10}$$

[13] Appendix B.1 shows that Proposition 1 is a direct consequence of the integrability of the demand system characterized by the expenditure share function $\boldsymbol{\omega}_t(\cdot)$. For completeness, Appendix B.2 characterizes the inverse mapping from total expenditure to real consumption, which we call the real consumption function.

Proof. From Definition (2), we know that everywhere along the path, the total expenditure is equal to the mapping $\chi_t^b(\cdot)$ evaluated at the corresponding level of real consumption, i.e. $y_t = \chi_t^b\left(c_t^b\right) = E\left(M_b^{-1}\left(c_t^b\right); \boldsymbol{p}_t\right)$. Equation (9) follows from

$$\frac{\partial \log \chi_t^b(c)}{\partial t} = \sum_i \frac{\partial \log E\left(M_b^{-1}(c); \boldsymbol{p}_t\right)}{\partial \log p_{it}} \cdot \frac{d \log p_{it}}{dt} = \sum_i \omega_{i,t}\left(\chi_t^b(c)\right) \cdot \frac{d \log p_{it}}{dt},$$

where in the second equality we have used Shephard's lemma.

We can now write the full time derivative of the total expenditure as

$$\frac{d \log y_t}{dt} = \sum_i \frac{\partial \log E\left(M_b^{-1}\left(c_t^b\right); \boldsymbol{p}_t\right)}{\partial \log p_{it}} \cdot \frac{d \log p_{it}}{dt} + \frac{\partial \log E\left(M_b^{-1}\left(c_t^b\right); \boldsymbol{p}_t\right)}{\partial \log c_t^b} \cdot \frac{d \log c_t^b}{dt},$$

which leads to Equation (10) after rearranging terms, since the first term on the right-hand side equals $\log D_t(y_t)$. Intuitively, this equation shows that the change in nominal expenditure is the sum of two terms: (i) price changes holding real consumption constant; (ii) the change in real consumption interacted with the change in the curvature of the expenditure function as real consumption changes. □

To draw insights from Proposition 1, let us first consider the case of homothetic preferences. In this case, the composition of demand is independent of expenditure and we have $D_t(y) \equiv D_t$ for all y. Hence, Equation (9) implies that along the path we have

$$\log \mathscr{P}_{b,t}^b(c) = \log \chi_t^b(c) - \log c = \int_b^t \log D_\tau \, d\tau, \qquad \forall b, c.$$

The integral on the right-hand side defines the standard Divisia price index, which gives the true price index under the homotheticity assumption. Beyond the homothetic case, as is well-known, this integral does not necessarily offer a price index that is theoretically consistent (Hulten, 1973).[14] Proposition 1 shows that the theory-consistent way to recover the true price index under nonhomotheticity is to integrate the Divisia function using the differential equation (9):

$$\log \mathscr{P}_{t_0,t}^b(c) = \log \chi_t^b(c) - \log c = \int_{t_0}^t \log D_\tau\left(\chi_\tau^b(c)\right) d\tau, \qquad \forall b, c. \tag{11}$$

The second insight of Proposition 1 is to show that we can account for the contribution of nonhomotheticity using a simple multiplicative factor, which rescales the standard formula that

[14] For instance, the integral may take different values between the two initial and final periods depending on the path of expenditure shares considered between the two periods.

deflates nominal expenditure growth by the Divisia index, $\frac{d}{dt}\log y_t - \log D_t(y_t)$. Let us define the *nonhomotheticity correction function* $\Lambda_t^b(\cdot)$ as the elasticity of the true index to real consumption from the base period to the current period, that is,

$$\Lambda_t^b(c) \equiv \frac{\partial \log \mathscr{P}_{b,t}^b(c)}{\partial \log c} = \frac{\partial \log \chi_t^b(c)}{\partial \log c} - 1, \tag{12}$$

so that the multiplicative factor in Equation (10) is given by $\left(1+\Lambda_t^b(c_t)\right)^{-1}$. Under homothetic preferences, this nonhomotheticity correction is zero $\Lambda_t^b(c) \equiv 0$ and we recover the standard result. Otherwise, we have to account for the deviation of the nonhomotheticity correction function Λ_t from zero in Equation (10). Of course, even if prices do not change over time, we still find $\Lambda_t^b(c) = 0$.

As we move forward in time from the base period $t > b$, Equation (12) shows that the nonhomotheticity correction rises if the cost-of-living price index, from the base to the current period, is higher at higher levels of real consumption. In such cases, raising one's real consumption becomes more expensive over time, and thus the exact measure of real consumption growth is smaller than that with the uncorrected deflation of nominal consumption growth, $\frac{d}{dt}\log y_t - \log D_t(y_t)$. In contrast, if the true price index is higher at lower levels of real consumption, raising one's real consumption becomes less expensive over time, and thus the exact measure of real consumption growth exceeds what is suggested without correction.[15]

When does the nonhomotheticity correction require a sizable adjustment to the standard uncorrected approach? First, by definition the nonhomotheticity correction is small when the current period t is close to the base period b, so that the true index $\mathscr{P}_{b,t}^b(c)$ is small. Second, the dependence of the index on real consumption stems from systematic differences in price changes across goods as a function of their income elasticities. Indeed, we can rewrite the nonhomotheticity correction as[16]

$$\Lambda_t^b(c) = \int_b^t \sum_i \left(\omega_{i,\tau}\left(\chi_\tau^b(c)\right) \cdot \eta_{i,\tau}^b(c) \cdot \frac{d\log p_{i\tau}}{d\tau} \right) d\tau,$$

where $\eta_{i,t}^b(c) \equiv \frac{\partial \log \omega_{i,t}\left(\chi_\tau^b(c)\right)}{\partial \log c}$ denotes the elasticity of expenditure shares with respect to real consumption. Thus, the nonhomotheticity correction is zero if price inflation $\frac{d\log p_{i\tau}}{d\tau}$ is uncorrelated with income elasticities $\eta_{i,\tau}^b(c)$ across goods i, even if the average size of price inflation is large.

[15] We provide intuition for this result with examples at the end of this section.

[16] We note that the importance of the covariance between income elasticities and price changes for measuring welfare change in presence of nonhomotheticity has been highlighted in prior work (e.g., Fajgelbaum and Khandelwal, 2016; Atkin et al., 2020; Baqaee and Burstein, 2021). The main insight of our work is how to use this result to nonparametrically uncover the measures of welfare change based on cross-sectional data.

We conclude that the nonhomotheticity correction is likely to be sizable when preferences are nonhomothetic, price inflation is large and correlated with income elasticities across goods, and real consumption is expressed in terms of a base period that is distant from the current period.

Most importantly, Proposition 1 allows us to uncover real consumption over time by approximating the Divisia index function $\log D_t(y)$ using the cross-sectional variations in the price indices across households. Before presenting this result in Section 2.3, below we present a number of other theoretical implications of Proposition 1.

Real Consumption Growth and the Choice of Constant Prices How does the choice of the base period affect the measurement of growth in real consumption? The following lemma shows that there is a systematic relationship between the choice of the base period and the corresponding measure of real consumption.

Lemma 1. *Consider two base periods $b_1 < b_2$. At time t, the rate of growth in real consumption measured with constant prices in period b_2, relative to real consumption with constant prices in period b_1, satisfies*

$$\frac{d\log c_t^{b_2}}{d\log c_t^{b_1}} = 1 + \Lambda_{b_2}^{b_1}\left(c_t^{b_1}\right) = 1 + \frac{\partial \log \mathscr{P}_{b_1,b_2}^{b_1}\left(c_t^{b_1}\right)}{\partial \log c_t^{b_1}}. \tag{13}$$

Proof. See Appendix B.4. □

Lemma 1 shows that the sign of the bias induced by the nonhomotheticity correction inherently depends on the choice of the base period.[17] More specifically, it shows that the gap between measures of growth at time t using two different base periods, b_1 and b_2, depends on the nonhomotheticity correction between the two periods b_1 and b_2. For instance, assume $b_1 < b_2$, prices are on the rise, and price inflation negatively covaries with income elasticities across goods between periods b_1 and b_2. In this case $\Lambda_{b_2}^{b_1} < 0$, and by Equation (13) real consumption growth is lower when measured from the perspective of the later period b_2.

To gain intuition about the economics behind this result, let us consider a simple economy with two goods: burgers and mobile phones. Assume that mobile phones are more income elastic than burgers and that over a period of time, for example from 1970 to 2020, the relative price of mobile phones falls substantially relative to burgers. From the perspective of prices held constant at their 1970 level, real consumption growth over this 50-year period is larger when preference nonhomotheticity is taken into account. The reason is that consumers become richer over time, which leads to an increase in the propensity to spend on mobile phones, precisely when the relative price of mobile phones is falling. Thus, in this example conventional measures

[17] To the best of our knowledge, this point has not been made in prior work on measuring welfare change in the presence of preference nonhomotheticity.

of real consumption growth are biased downward because they do not account for the fact that the income-elastic goods become relatively cheaper at the same time that they become relatively more important from the point of view of consumer preferences.

In contrast, looking backward in time from the perspective of prices held fixed at a later period, for example 2020, real consumption growth during the period is smaller when accounting for the nonhomotheticity correction. Indeed, going backward in time, consumers become poorer and spend relatively more on the income-inelastic good, burgers, which become relatively cheaper. Thus, the fall in income is dampened by the fact that burgers are relatively cheaper while consumer demand for burgers has increased. Therefore, consumers in the past were richer than typically thought; that is, conventional measures of real consumption growth are biased upward.

These examples illustrate how the curvature of the mapping between welfare and our money metric depends inherently on the choice of the base period. Regardless of the choice of the base period, in the examples above the level of real consumption is always underestimated by the standard measures, all the more so as we move away from the base period.[18]

Characterization of the Real Consumption Function Proposition 1 characterizes the mapping from real consumption to total expenditure at any point in time. Since this mapping is monotonic, it also fully characterizes the inverse mapping $\widetilde{\chi}_t^b(y) \equiv \left(\chi_t^b\right)^{-1}(y)$, from total expenditure to real consumption, which we may refer to as the (indirect) real consumption function. The following lemma shows that the real consumption function provides a dual representation of the mapping from real consumption to total expenditure.

Lemma 2. *The real consumption function and the mapping from real consumption to expenditure satisfy the following relationship for all* t, $b \in [0, T]$ *and for all* $y > 0$*:*

$$\widetilde{\chi}_t^b(y) = \chi_b^t(y). \tag{14}$$

Proof. The real consumption function satisfies $\widetilde{\chi}_t^b(y) = E(v_t(y); \boldsymbol{p}_b)$, where we have defined the indirect utility $v_t(y)$ through $y = E(v_t(y); \boldsymbol{p}_t)$. Noting $v_t(y) = M_t^{-1}(y)$ for the money metric defined as in Equation (1) leads to the desired result. □

Appendix B.2 derives a direct characterization of the real consumption function as the solution to a first-order hyperbolic partial differential equation, and discusses its connection with the

[18] Another application of the insight that measured growth depends on the vector of fixed prices has recently been provided by Oberfield (2022). He constructs a general equilibrium growth model that features a U-shaped pattern of inflation inequality (as a function of household income) along the constant growth path. Along such paths, the rates of growth in real consumption, when measured in terms of a base period far in the past or one far in the future, are equal across households. In contrast, when measured in terms of the current base period, these rates take higher values and feature inequality across households.

differential equation (9). The appendix further discusses how we may use this representation of the differential equation (9) to construct other alternatives to our approach for approximating real consumption based on cross-sectional data.

2.3 Approximating Welfare Changes under Nonhomotheticity

Proposition 1 characterizes a theoretically consistent measure of real consumption as the solution to a differential equation expressed in terms of the Divisia function. This function in turn tells us how the true price index depends on total expenditure. In this section, we build a number of different approximate solutions to this differential equation using data on prices and repeated cross-sections of household expenditures. The key insight is that classical index number theory allows us to approximate the Divisia function for any underlying preferences, based on cross-sectional variations in price indices across households as a function of their total expenditure.

2.3.1 Setting for the Approximation

As in Section 2.1.1, we consider continuous paths for prices and total expenditure in some fixed time interval, but now additionally assume that the data provide us with only $T+1$ discrete observations along this path. Without loss of generality, we denote the end period by the integer T and let $t \in \{0, 1, \cdots, T\}$ denote the time index of each observation. Since the paths of prices and total expenditure are fixed, we assume that the following bounds on price inflation and nominal expenditure growth increasingly vanish as we increase the number of observations $T+1$:

$$\Delta_p \equiv \max_{i,t} \left\{ \left| \log\left(\frac{p_{i,t+1}}{p_{i,t}} \right) \right| \right\}, \qquad \Delta_y \equiv \max_{t} \left| \log\left(\frac{y_{i,t+1}}{y_{i,t}} \right) \right|. \tag{15}$$

We use the bounds above to introduce the concepts needed for constructing our approximation error bounds. Consider two sequences $\{f_t\}_{t=0}^{T}$ and $\{g_t\}_{t=0}^{T}$ defined as functions of the observed sequences of price and total expenditure along the path. As the number of observations $T+1$ and the bounds in Equation (15) change, the values of the two sequences also change.

Let us denote the corresponding mapping between the size of the bound Δ, where $\Delta = \max\{\Delta_p, \Delta_y\}$, and the values of the two sequences as $f_t \equiv f_t(\Delta)$ and $g_t \equiv g_t(\Delta)$.[19] Now, we define the sequence $\{f_t\}_{t=0}^{T}$ as an m-th order approximation of the sequence $\{g_t\}_{t=0}^{T}$, and denote this by $f_t - g_t = O\left(\Delta^{m+1}\right)$, if the differences between the values of the two sequences fall in magnitude with Δ^{m+1} as T grows. Formally, this relationship holds if $\lim_{\Delta \to 0} (f_t(\Delta) - g_t(\Delta)) \Delta^{-(m+1)} = a$ for some finite constant $a > 0$.

[19] Note that this definition involves a slight abuse of notation, since the sequence is a function of all observations of prices, total expenditures, and expenditure shares, not just of Δ.

For the key results presented in Section 2.3.4 below, we make the additional assumption that in each period we observe the composition of consumption expenditures for N consumers or households with identical preferences characterized by a continuously differentiable expenditure function, $E(u;\boldsymbol{p})$. They face the same sequence of prices and have heterogeneous levels of total expenditures, satisfying the bounds in Equation (15).

2.3.2 Index Formulas and Local Approximations of the True Index Function

We begin with a lemma showing that the sequences of geometric and Törnqvist price indices (between successive time points) provide approximations of the corresponding sequence of true price indices up to first and second orders, respectively.[20]

Lemma 3. *Assume that the underlying expenditure function $E(\cdot;\cdot)$ characterizing choices $(\boldsymbol{p}_t, \boldsymbol{s}_t, y_t)$ and $(\boldsymbol{p}_{t+1}, \boldsymbol{s}_{t+1}, y_{t+1})$ is third-order continuously differentiable in all its arguments. Then, if the corresponding changes in prices and total expenditures satisfy Equation (15), the geometric and Törnqvist price index formulas satisfy*

$$\log \mathscr{P}^{b}_{t,t+1}(c) = \log \mathbb{P}_G(\boldsymbol{p}_t, \boldsymbol{s}_t; \boldsymbol{p}_{t+1}, \boldsymbol{s}_{t+1}) + O(\Delta^2), \qquad \text{if } c \in \{c^b_t, c^b_{t+1}\}, \tag{16}$$

$$= \log \mathbb{P}_T(\boldsymbol{p}_t, \boldsymbol{s}_t; \boldsymbol{p}_{t+1}, \boldsymbol{s}_{t+1}) + O(\Delta^3), \qquad \text{if } c = \sqrt{c^b_t \cdot c^b_{t+1}}, \tag{17}$$

where $\Delta \equiv \max\{\Delta_p, \Delta_y\}$ and where $c^b_t = (\chi^b_t)^{-1}(y_t)$ denotes the level of real consumption corresponding to choice $(\boldsymbol{p}_t, \boldsymbol{s}_t, y_t)$.

Proof. See Appendix B.4. □

Recall that under homotheticity, the true price index does not depend on the level of real consumption c. As the proof of the lemma shows, under homotheticity the lemma holds for any level of real consumption c and with a tighter bound $\Delta \equiv \Delta_p$. In this case, the sequences of geometric and Törnqvist indices provide us with approximations of the Divisia index, which we can chain over time to integrate the Divisia index and approximate any true price index $\log \mathscr{P}^{b}_{t_0,t}(c)$. Thus, in the case of homothetic preferences, the error in the chained indices over the entire fixed interval, depending on whether the geometric or Törnqvist formula is used, is first or second order.[21]

In the presence of nonhomotheticity the lemma shows that approximations remain valid only for *local* levels of real consumption, in the sense that they are close to c^b_t and c^b_{t+1}. Thus, chaining

[20] As we discuss in Section 2.3.7, we can generalize this result for broader classes of index formulas defined in Section 2.1.2. Lemma 3 closely parallels the results of Diewert (1976), who shows that the Törnqvist price index is exact for the translog family of expenditure functions.

[21] The lemma implies the error bounds $O(T \cdot \Delta^2)$ and $O(T \cdot \Delta^3)$ for the chained geometric and Törnqvist formulas, respectively. Note that since we keep the interval and the overall true index fixed, we have $T^{-1} = O(\Delta)$.

geometric and Törnqvist indices does *not* lead to a theoretically consistent measure of the true price index over the entire interval. As we will see next, however, we can still rely on the insights of Proposition 1 to approximate the true price index.

2.3.3 Global Approximations for the True Index Function

Proposition 1 allows us to extend Lemma 3 to construct approximations for the true price index corresponding to arbitrary values of real consumption. This result is presented in the following lemma.

Lemma 4. *Assume that the conditions stated in Lemma 3 hold. Then, the true cost-of-living function* $\mathscr{P}^{b}_{t,t+1}(c) \equiv \chi^{b}_{t+1}(c)/\chi^{b}_{t}(c)$ *satisfies*

$$\log \mathscr{P}^{b}_{t,t+1}(c) = \pi^{+}_{t}\left(\chi^{b}_{t}(c)\right) + O\left(\Delta_p^2\right), \tag{18}$$

$$= \frac{1}{2}\left[\pi^{+}_{t}\left(\chi^{b}_{t}(c)\right) + \pi^{-}_{t+1}\left(\chi^{b}_{t+1}(c)\right)\right] + O\left(\Delta_p^3\right), \tag{19}$$

where $\Delta \equiv \max\left\{\Delta_p, \Delta_y\right\}$ *and where we have defined Laspeyres* $\pi^{+}_{t}(y)$ *and Paasche* $\pi^{-}_{t+1}(y)$ *geometric index functions as*

$$\pi^{+}_{t}(y) \equiv \sum_i \omega_{it}(y) \log\left(\frac{p_{it+1}}{p_{it}}\right), \tag{20}$$

$$\pi^{-}_{t+1}(y) \equiv \sum_i \omega_{it+1}(y) \log\left(\frac{p_{it+1}}{p_{it}}\right). \tag{21}$$

Proof. See Appendix B.4. □

Lemma 4 offers an approximate, discretized re-statement of Proposition 1. The two functions defined in Equations (20) and (21) allow us to approximate the Divisia index $D_t(y)$ as a function of total expenditure. As we will see below, we can nonparametrically estimate these functions using observed cross-sectional variations in index formulas across households.

We next use Lemma 4 to construct the central contributions of this paper, i.e., a number of algorithms that provide approximate real consumption over time using repeated cross-sectional data. These algorithms vary in the approaches we choose about how to use the lemma above to numerically integrate the differential equation (9) over time, starting from the base period b in which the mapping $\chi^{b}_{b}(c) = c$ is known.

2.3.4 First-Order Algorithms

We begin with our simplest algorithm that uses Equation (18) to approximate the true index function $\mathscr{P}^{b}_{t,t+1}(\cdot)$ and correspondingly the values of real consumption across households. First, we evaluate Equation (18) at the level of real consumption $c = c^{b,n}_{t}$ for each household n to find

$$\log \mathscr{P}^{b}_{t,t+1}\left(c^{b,n}_{t}\right) = \pi^{+}_{t}\left(y^{n}_{t}\right) + O\left(\Delta^{2}_{p}\right) = \pi^{+,n}_{t} + O\left(\Delta^{2}_{p}\right), \tag{22}$$

where we have used the fact that $\chi^{b}_{t}\left(c^{b,n}_{t}\right) = y^{n}_{t}$ and that $\pi^{+}_{t}\left(y^{n}_{t}\right) = \pi^{+,n}_{t} \equiv \log \mathbb{P}_{G}\left(\boldsymbol{p}_{t}, \boldsymbol{s}^{n}_{t}; \boldsymbol{p}_{t+1}, \boldsymbol{s}^{n}_{t+1}\right)$ coincides with the geometric index formula for this household. Next, we apply a Taylor series expansion of $\chi^{b}_{t+1}(\cdot)$ around c^{b}_{t} to write the left-hand side of Equation (22) as

$$\log \mathscr{P}^{b}_{t,t+1}\left(c^{b,n}_{t}\right) = \log \chi^{b}_{t+1}\left(c^{b,n}_{t}\right) - \log y^{n}_{t} = \log\left(\frac{y^{n}_{t+1}}{y^{n}_{t}}\right) - \frac{\partial \log \chi^{b}_{t+1}\left(c^{b,n}_{t}\right)}{\partial \log c^{b,n}_{t}} \log\left(\frac{c^{b,n}_{t+1}}{c^{b,n}_{t}}\right) + O\left(\Delta^{2}_{y}\right), \tag{23}$$

where we have used the definition $\mathscr{P}^{b}_{t,t+1}(c) \equiv \chi^{b}_{t+1}(c) / \chi^{b}_{t}(c)$ and the fact that $\chi^{b}_{t}\left(c^{b,n}_{t}\right) = y^{n}_{t}$.

Equations (22) and (23), along with the definition of the nonhomotheticity correction $\Lambda^{b}_{t}(\cdot)$ in Equation (12), allow us to derive an update rule for the values of real consumption across consumers from one period to the next:

$$\log \widehat{c}^{n}_{t+1} = \log \widehat{c}^{n}_{t} + \frac{1}{1 + \widehat{\Lambda}_{t+1}\left(\widehat{c}^{n}_{t}\right)}\left[\log\left(\frac{y^{n}_{t+1}}{y^{n}_{t}}\right) - \pi^{+,n}_{t}\right], \tag{24}$$

where we have omitted the superscript b indicating the base year to simplify notation, and where we have indicated our estimated value of real consumption at time t for household n by $\widehat{c}^{n}_{t}$.

The key remaining step is to estimate the value of the nonhomotheticity correction. The simplest approach is to once again rely on Equation (18), and the fact that $\log \mathscr{P}_{t,t+1}\left(c^{n}_{t}\right) \approx \pi^{+,n}_{t}$, to nonparametrically estimate $\log \mathscr{P}_{t,t+1}(\cdot)$ as a function of real consumption in each period. In particular, starting from the base period $t = b$, the real consumption for each consumer is equal to their observed total expenditure $\widehat{c}^{n}_{b} = c^{n}_{b} = y^{n}_{b}$. Thus, we can apply a nonparametric regression of the log geometric index formula $\pi^{+,n}_{b}$ on real consumption $\widehat{c}^{n}_{b}$ across households in this period; we thus recover an estimated true cost-of-living index $\log \widehat{\mathscr{P}}_{b,b+1}(\cdot)$ as a function of real consumption. We can now use the derivative of this function to evaluate the nonhomotheticity correction $\widehat{\Lambda}_{b+1}\left(\widehat{c}^{n}_{b}\right)$ for each household and apply the update rule in Equation (24) to find the value of real consumption for each household $\widehat{c}^{n}_{b+1}$ in the next period. This allows us to then nonparametrically estimate true cost-of-living index $\log \widehat{\mathscr{P}}_{b+1,b+2}(\cdot)$ in the next period as a function of real consumption, using a regression of the log geometric index formula $\pi^{+,n}_{b+1}$ on real consumption $\widehat{c}^{n}_{b+1}$ across households.

Applying the two steps above successively moving forward in time from the base period, we can recover the cumulative true cost-of-living $\log \widehat{\mathscr{P}}_{b,t+1}(c) \equiv \sum_{\tau=b}^{t} \log \widehat{\mathscr{P}}_{\tau,\tau+1}(c)$ as a function of real consumption for each period $t \geq b$. This function in turn allows us to estimate the nonhomotheticity correction and use Equation (24) to recover the values of real consumption in the next period.

Algorithm 1 formally states this procedure, using power series estimators,[22] moving either forward or backward in time from the base period:

Algorithm 1 (Baseline First-Order Algorithm)**.** *Consider a sequence of power functions* $\{f_k(z) \equiv z^k\}_{k=0}^{K_N}$ *for some* K_N*, where* N *is the number of consumers in the cross-section. Let* $\widehat{c}_b^n \equiv y_b^n$ *and for each* $t \geq b$*, successively apply the following two steps.*

1. *Nonparametrically fit the true price index between periods* t *and* $t+1$*:*

   Estimate the coefficients $\left(\widehat{\alpha}_{k,t}\right)_{k=0}^{K_N}$ *solving the following problem:*

$$\min_{\left(\alpha_{k,t}\right)_{k=0}^{K}} \sum_{n=1}^{N} \left(\pi_t^{+,n} - \sum_{k=0}^{K_N} \alpha_{k,t} f_k\left(\log \widehat{c}_t^n\right) \right)^2, \tag{25}$$

   where $\left\{\pi_t^{+,n}\right\}_n$ *are household-specific price index formulas at time* t *defined by*

$$\pi_t^{+,n} \equiv \log \mathbb{P}_G\left(\boldsymbol{p}_t, \boldsymbol{s}_t^n; \boldsymbol{p}_{t+1}, \boldsymbol{s}_{t+1}^n\right). \tag{26}$$

2. *Estimate the values of real consumption for consumers in period* $t+1$*:*

   Use Equation (24), where the approximate nonhomotheticity correction function is given by

$$\widehat{\Lambda}_{t+1}(c) \equiv \sum_{k=0}^{K_N} \left(\sum_{\tau=b}^{t} \widehat{\alpha}_{k,\tau} \right) f_k'(\log c). \tag{27}$$

To apply the algorithm backward in time for $t < b$*, simply re-label all the time indices in the data preceding the base period* b *such that* $t - \tau \to t + \tau$ *for all* $1 \leq \tau \leq b$ *and perform the same steps as above.*

In practice, the algorithm is easy to implement and consists of two steps: (i) running a sequence of period-by-period OLS regressions to recover the true cost-of-living index as a function

[22] One can apply alternative series-function approximations, using alternative basis functions such as Fourier, Spline, or Wavelets. The results here generalize to such alternative nonparametric methods subject to modified regularity assumptions on the expenditure function and the distribution of real consumption across consumers (Newey, 1997).

of real consumption; (ii) summing up period-specific OLS coefficients from the base to the current period, evaluating the nonhomotheticity correction, and applying the update in Equation (24).

Sources of Error and Algorithm Refinement In addition to the first-order discretization error introduced in Equation (22), our baseline Algorithm 1 includes two additional sources of error. First, by performing a Taylor expansion of the true cost-of-living function $\mathscr{P}_{t,t+1}(\cdot)$ around c_t^n in Equation (23), we have introduced an additional error of the same order as the growth in real consumption (in turn of order Δ_y). While this choice simplifies the algorithm, we can refine the algorithm to remove this error from the analysis by, first, evaluating Equation (22) at points $c_{t+1}^{n,(\ell)}$ that become successively closer to c_{t+1}^n as ℓ increases and, second, applying the Taylor series expansion around the previous point $c_{t+1}^{n,(\ell-1)}$. Considering the limiting case $\chi_{t+1}\left(c_{t+1}^{n,(\ell)}\right) \to y_{t+1}^n$, we can write the latter's expansion as

$$\log \mathscr{P}_{t,t+1}\left(c_{t+1}^{n,(\ell)}\right) = \log y_{t+1}^n - \log \chi_t\left(c_{t+1}^{n,(\ell)}\right) = \log\left(\frac{y_{t+1}^n}{y_t^{n,(\ell-1)}}\right) - \frac{\partial \log \chi_t^b\left(c_{t+1}^{n,(\ell-1)}\right)}{\partial \log c_{t+1}^{n,(\ell-1)}} \log\left(\frac{c_{t+1}^{n,(\ell)}}{c_{t+1}^{n,(\ell-1)}}\right) + O\left(\epsilon^2\right), \tag{28}$$

where we have defined $y_t^{n,(\ell-1)} \equiv \chi_t^b\left(c_{t+1}^{n,(\ell-1)}\right)$ and the error $\epsilon \equiv \left|\log\left(c_{t+1}^{n,(\ell)}/c_{t+1}^{n,(\ell-1)}\right)\right|$ can now be made arbitrarily small as we iterate over ℓ.

The second source of error stems from the nonparametric estimation step. In particular, our baseline Algorithm 1 makes the simplifying choice to directly estimate the true cost-of-living index $\mathscr{P}_{t,t+1}(\cdot)$ as a function of real consumption in each period t. In applying this step, the algorithm combines two distinct sources of error: (i) the sampling error in estimating function $\pi_t^+(\cdot)$, due to the finite-sample cross-sectional data; and (ii) the error in the estimates of real consumption in periods away from the base, which leads to a measurement error problem (error-in-variables) in estimation. We can refine the algorithm by separating the two stages: first, nonparametrically estimate the geometric index function $\pi_t^+(\cdot)$ in each period using the cross-sectional data on geometric indices; second, nonparametrically fit the mapping $\chi_t(\cdot)$ for each successive period as a function of real consumption estimates recovered in that period. Algorithm A.1 in Appendix A.1.2 combines this strategy with the Taylor series approximation in Equation (28) to provide a refinement of our first-order algorithm. Appendix A.1.2 further discusses the sources of approximation error in this algorithm.

2.3.5 Second-Order Algorithms

In parallel to the approach laid out in Section 2.3.4, relying on the first-order approximation of our key differential equation in Equation (18), we can similarly construct algorithms that instead rely on the second-order approximation in Equation (19). Algorithm A.3 in Appendix A.2 uses an iterative structure to achieve this second-order approximation. Recall that Algorithm 1 evaluates the nonhomotheticity correction only at the current period's level of real consumption, $\widehat{\Lambda}_{t+1}(\widehat{c}_t)$, to approximate the real consumption growth c_{t+1}/c_t. In contrast, our second-order algorithm additionally evaluates the nonhomotheticity correction function at the next period's level of real consumption, $\widehat{\Lambda}_{t+1}(\widehat{c}_{t+1})$. As a result, the algorithm further involves solving for a fixed-point problem in each period to update the value of real consumption in successive periods. Algorithm A.4 in Appendix A.2.2 provides a refinement of the second-order algorithm, along the same lines as we discussed above.

2.3.6 Other Extensions

In this section, we discuss three additional extensions of our baseline and refined algorithms.

Alternative Algorithms Both Algorithms 1 and A.3 approximate the nonhomotheticity correction by nonparametrically estimating the elasticity of the mapping $\chi_t^b(c)$ from expenditure to real consumption. We favor this approach, since it intuitively and transparently establishes the link between cross-household inequality in cumulative inflation and the nonhomotheticity correction. However, an alternative approach is to approximate the nonhomotheticity correction by nonparametrically estimating the elasticity of the inverse mapping, $\widetilde{\chi}_t^b(y)$. Algorithm A.2 (in Appendix A.1.3) and Algorithm A.5 (in Appendix A.2.3) provide first- and second-order schemes based on this alternative approach. In Section 2.4 below, we provide a comparison of the approximation errors among all of our alternative algorithms.

Finally, as another example, subsequent to our work Baqaee et al. (2022) have presented a different alternative to our benchmark first-order algorithm. In Appendix B.3, we establish the tight theoretical connection between their approach and ours. We provide evidence using both synthetic and real-world data that in practice their approach leads to results that are similar to those produced by our baseline Algorithm 1.[23]

Alternative Price Index Formulas We can generalize the results of Lemma 3, and thus the first- and second-order Algorithms 1 and A.3, to index formulas beyond geometric and Törnqvist. The following proposition states this result formally.

[23] For the results in the case of synthetic data, see Appendix C.2, and for those in the case of real U.S. data, see Section 3 and Appendix Figure E.9.

Proposition 2. *If the expenditure function $E(\cdot;\cdot)$ is second-order continuously differentiable in all its arguments, then the price index formulas defined in Section 2.1.2 satisfy*

$$\log \mathbb{P}_G\left(\boldsymbol{p}_t, \boldsymbol{s}_t; \boldsymbol{p}_{t+1}, \boldsymbol{s}_{t+1}\right) = \log \mathbb{P}_I\left(\boldsymbol{p}_t, \boldsymbol{s}_t; \boldsymbol{p}_{t+1}, \boldsymbol{s}_{t+1}\right) + O\left(\Delta^2\right), \qquad I \in \{P, L, T, F, S\},$$
$$\log \mathbb{P}_T\left(\boldsymbol{p}_t, \boldsymbol{s}_t; \boldsymbol{p}_{t+1}, \boldsymbol{s}_{t+1}\right) = \log \mathbb{P}_I\left(\boldsymbol{p}_t, \boldsymbol{s}_t; \boldsymbol{p}_{t+1}, \boldsymbol{s}_{t+1}\right) + O\left(\Delta^3\right), \qquad I \in \{F, S\},$$

where $\Delta \equiv \max\left\{\Delta_y, \Delta_p\right\}$ *with* Δ_y *and* Δ_p *defined as in Equation (15).*

Proof. See Appendix B.4. □

One implication of Proposition 2 is the classification of price index formulas into two groups: the first group (composed of geometric, Laspeyres, and Paasche index formulas) provides a first-order approximation to the true price index, while the second group (composed of Törnqvist, Fisher, and Sato-Vartia) provides a second-order approximation. To reflect the accuracy of the approximations for each group, we refer to the first group of index formulas as *first-order index formulas* and to the second group as *second-order index formulas*.

It follows that the results of Lemma 3 for first- and second-order approximations extend to any formulas in the first- and second-order family of indices, respectively. For instance, the Sato-Vartia or the Fisher index between periods t and $t+1$ approximates the true price index between these two points for the corresponding level of real consumption specified in Lemma 3. Moreover, we can replace the Törnqvist index with the Sato-Vartia or Fisher index in our second-order algorithm.We rely on these extended results in our empirical exercise in Section 3 where, due to data limitations, the most natural choice for a second-order index is the Fisher index.

Observable Heterogeneity in Consumer Characteristics Our method requires us to infer the relationship between the true price index and total expenditure from the cross-household relationship between price index formulas and total expenditures (e.g., Step 1 of Algorithm 1). However, the observed relationship between household-level price indices and household expenditures may in principle be confounded by other factors, for example household age or education. To alleviate this potential concern, we can (nonparametrically) control for observable covariates in this step of the algorithm. However, to build a theoretically consistent account of the potential dependence of consumer preferences on characteristics beyond income, we need to generalize our concept of real consumption. As we will discuss in Section 4 below, such a generalization leads to further corrections in our standard measures of real consumption, beyond the nonhomotheticity correction, in order to account for the impact of potential changes in consumer characteristics on consumer welfare over time.[24]

[24] Empirically, we find that the results from our baseline algorithm are robust to this extension.

2.3.7 Discussion

As discussed above, Lemma 3 and Proposition 2 together classify common price index formulas into two first- and second-order groups, based on the accuracy of the approximations they provide for true price indices under arbitrary underlying preferences. Our approach thus differs from the standard treatment of index formulas, which classifies index formulas based on the underlying family of preferences for which they provide exact measures of true price indices (Diewert, 1993). For instance, the Törnqvist price index is exact for the family of preferences that lead to a translog unit cost function.[25] Unlike our approach, the concept of exact price indices requires specifying the underlying form of the preference functions.

One crucial step is to define, as in Diewert (1976), the Fisher and Törnqvist price indices as "superlative" price indices, on the grounds that they are exact for families of preferences that can provide a second-order approximation to other homothetic preferences, namely the quadratic and the translog family, respectively. In line with this insight, Diewert (1978) has shown that alternative choices of superlative indices, when chained, lead to very similar estimates for the changes in cost-of-living and real consumption in practice. Lemma 3 and Proposition 2 formalize these classical insights and generalize them to include the Sato-Vartia index. Instead of establishing the exactness of different index formulas for distinct families of preferences that may approximate general preferences, the lemma provides bounds on the approximation error of the reduced-form indices for arbitrary preferences.[26]

As mentioned, these classical results do not allow us to provide precise approximations of real consumption growth over long time horizons beyond the case of homothetic preferences.[27] By solving this problem, our approach offers a substantial generalization of index number theory to nonhomothetic preferences.

2.4 Simulation

In this section, we perform a simple simulation to illustrate and validate the accuracy of our algorithms in accounting for the effect of nonhomotheticity when measuring real consumption growth.

Comin et al. (2021) have shown that the nonhomothetic CES (nhCES) preferences lead to a demand system compatible with the cross-sectional relationship between household income

[25] As for other examples, the Laspeyres and Paasche indices are exact for Leontief utility functions, and the geometric and Sato-Vartia index formulas are exact for Cobb-Douglas and CES utility functions. The Fisher price index is exact for the family of preferences that lead to quadratic unit cost functions.

[26] In line with Equation (17), Diewert (1976) shows that the Törnqvist index is exact for the family of nonhomothetic preferences characterized by a translog expenditure function, for the true index under the level of real consumption specified in Lemma 3.

[27] Samuelson and Swamy (1974) provide examples showing how classical price indices fail under nonhomotheticity.

and the composition of expenditure among three main sectors of the economy: *a*griculture, *m*anufacturing, and *s*ervices. Following their specification, we assume that the expenditure function satisfies:

$$E(u;\boldsymbol{p}_t) \equiv \left(\sum_{i \in \{a,m,s\}} \psi_i \left(u^{\varepsilon_i} p_{i,t} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}. \tag{29}$$

We use the same parameters as in Comin et al. (2021): $(\sigma, \varepsilon_a, \varepsilon_m, \varepsilon_s) = (0.26, 0.2, 1, 1.65)$, implying that services are luxuries (income elasticities exceeding unity) and agricultural goods are necessities (income elasticities lower than unity). We consider a population of a thousand households with an initial distribution of expenditure that is log-normal, with a mean corresponding to the average U.S. per-capita nominal consumption expenditure of $3,138 in 1953 and a standard deviation of log expenditure of 0.5 (Battistin et al., 2009). We consider a horizon of 70 years and assume that over this horizon nominal expenditure grows at the constant rate of 4.48% per year, in line with the U.S. data for the period 1953-2019. In each of the cases discussed below, we choose the fixed sectoral demand shifters ψ_i in Equation (29) in such a way that in the first period the composition of aggregate expenditure fits the US average shares of sectoral consumption in the three sectors in 1953.[28]

To examine the role of the covariance between price inflation and income elasticities, we consider a simple, purely illustrative simulation. We set the inflation rate in the manufacturing sector to be 3.19%, to match the average inflation rate in the US over the period 1953-2019. We then consider two illustrative cases featuring either positive or negative covariances between inflation and income elasticities. To study the case with a positive covariance, the inflation rate is set to be 1pp higher in service and 1pp lower in agriculture compared to manufacturing, leading to the inflation rates of 4.19% in services and of 2.19% in agriculture. To illustrate the case of a negative covariance, we reverse these parameters, setting inflation rates to 2.19% in services and 4.19% in agriculture. The resulting rates of growth in average real consumption in the simulated data in the positive, zero, and negative covariance cases are 0.7%, 1.3%, and 1.9% per year, respectively.

Given the known structure of the underlying preferences, this example allows us to compute the true values of real consumption for each household and assess the accuracy of our algorithms. Relying only on the simulated data, we also apply the standard uncorrected deflation of nominal consumption expenditure for each household to assess the magnitude of the bias in the uncorrected measures.

Figures 1a-1d report the results. We compare the evolution of the average measures of real consumption across the simulated population over time with the two different approximations. First, we see that the conventional approach based on chaining uncorrected measures of nominal

[28] The corresponding shares in the US based on the BLS data are 0.14, 0.27, and 0.59 for agriculture, manufacturing, and services, respectively.

Figure 1: Illustrative Simulation of the Evolution of Average Real Consumption

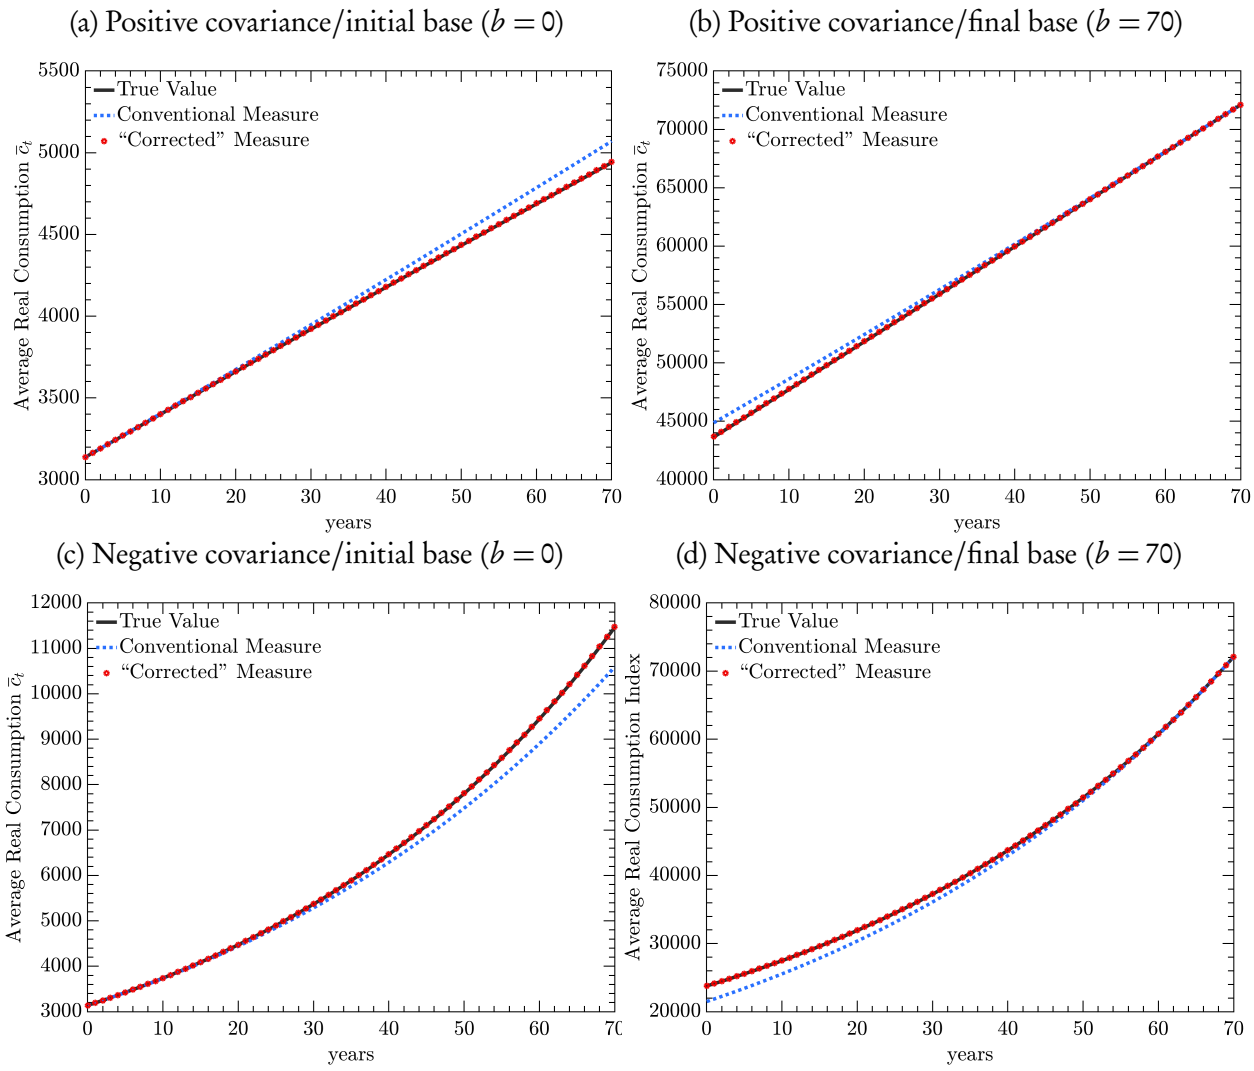

Note: The figures compare the evolution of the true value of average real consumption with two different approaches to approximating this value: 1) the average of the uncorrected nominal real consumption growth deflated by household-specific geometric price indices, and 2) applying the nonhomotheticity correction using the first-order algorithm. The panels show the resulting series for the choices of base period (a) $b = 0$ and (b) $b = 70$ with a positive income elasticity-inflation covariance and (c) $b = 0$ and (d) $b = 70$ with a negative covariance.

expenditure growth deflated by the Törnqvist index leads to sizable bias depending on the choice of the base period and/or the covariance between price inflation and income elasticities. While errors accumulate in the uncorrected chained values, applying our first-order nonhomotheticity correction yields results that are virtually indistinguishable from the true evolution of real consumption based on the underlying preferences. Thus, our approach accurately recovers the evolution of the true index without the knowledge of the parameters of the demand system.

In Appendix C, we provide an illustration of the evolution of the expenditure function in our simulation over time and compare it against a homothetic benchmark. This analysis demonstrates how changes in the curvature of the expenditure function translate into biases in the uncor-

rected measures of real consumption growth. The appendix further provides a detailed analysis of the size of the approximation error under our alternative algorithms, and extends the simulation to a wider range of values for the covariance between price inflation and income elasticities.

3 Empirics

In this section, we apply our approach to data from the US and quantify the magnitude of the bias in conventional measures of real consumption growth.

3.1 Data and Descriptive Statistics

Data To assess the empirical importance of the nonhomotheticity correction, we build a dataset providing total expenditures and expenditure shares at a granular level, across 598 items from the consumer expenditure survey (CEX). These items, called universal classification codes (UCC), are defined by the BLS and cover the entire consumption basket of households in the U.S. We obtain price changes for each item using CPI price series combined with the official concordance provided by the BLS for active UCCs, which we extend manually in prior years for UCCs that were discontinued. Appendix D provides a complete description of the steps we take in the construction of the data.

Using the CEX micro-data, we obtain expenditure patterns and socio-demographic characteristics at the household level. We then aggregate the household-level data to the level of pre-tax income percentiles. We thus obtain expenditure patterns that vary across income percentiles, which we will use to compute the income elasticity of inflation. We also use this dataset to measure consumption growth rates across income percentiles. To ensure that the patterns of consumption are consistent with national accounts at the aggregate level, we reweigh the data series so that aggregates match the official aggregate personal consumption expenditures provided by the Bureau of Economic Analysis (BEA).[29] Our analysis is thus fully consistent with macroeconomic aggregates and extends the logic of the distributional national accounts (Piketty et al. (2018)) to a setting allowing for the computation of inflation inequality.

Prior to 1984, the data require special treatment since CEX household-level data and CEX expenditure summary tables by product category and socio-demographic groups are no longer available, except in two years, 1972 and 1960. We use these two data points to interpolate the data for missing years. Prior to 1960, we use our first-order approximation to the correction for nonhomotheticities to extrapolate expenditure shares back to 1955, and we obtain the growth

[29] See Appendix D for a detailed description of this step. As described in Appendix D, we also ensure that our dataset perfectly matches the official CEX summary tables published by the BLS by product categories and income quintiles.

rate of aggregate consumption expenditures from the BEA.[30] Given the data limitations prior to 1984, we present two sets of results, first focusing on the period from 1984 to 2019 for which high-quality CEX data is available annually, and then a longer historical analysis going back to 1955.

Descriptive Inflation Statistics This new linked dataset allows us to provide evidence on inflation inequality over a long time horizon, thus extending prior estimates that have focused on much shorter time series. Computing inflation using group-specific price indices, we find that inflation inequality is a long-run phenomenon. Panels (a) and (b) of Figure 2 report aggregate and heterogeneous inflation patterns between 1984 and 2019, using chained geometric price indices. While Panel (a) shows that the cumulative inflation rate with aggregate expenditure shares is about 120%, panel (b) reports that inflation was higher for lower-income groups, ranging from 140% at the bottom to 110% at the top. Thus, over the course of these 35 years, a gap of around 30 percentage points has opened up in the chained geometric indices between the lowest and highest income groups. This finding is consistent with the growing literature on "inflation inequality," the fact that inflation rates are higher for lower-income households (e.g., Kaplan and Schulhofer-Wohl, 2017; Jaravel, 2019; Argente and Lee, 2021). While this literature focuses on post-2000 patterns, our data shows that this trend persists over several decades.

Furthermore, panels (c) and (d) extend the analysis back to 1955, showing that inflation inequality also existed over this longer time horizon. We find that on average over the 1955-2019 period, the annual inflation rate was about 35 basis points lower for the top relative to the bottom of the income distribution. This sustained difference in inflation leads to a gap of about 175 percentage points in cumulative inflation over the period, which varies from 700% at the top to 875% at the bottom of the income distribution. To the best of our knowledge, this paper is the first to build a dataset with disaggregated consumption patterns providing evidence on inflation inequality for a period of nearly 65 years.

Online Appendix Figure E.1 reports additional descriptive patterns on the dynamics and magnitude of inflation inequality over time.[31] Inflation inequality was strongest after 1995, weak between 1984 and 1995, and significant between 1955 and 1984.[32]

[30] See Appendix D for a detailed description of this step.

[31] Note that, although the cumulative level of inflation inequality shown in Figure 2 is economically meaningful, it is smaller than the deviations we considered in the illustrative example in Section 2.4.

[32] Explaining these patterns of inflation inequality falls beyond the scope of this paper, but we note that they are consistent with several mechanisms that were proposed in recent work. For example, demand-driven theories of directed innovation can lead to inflation inequality in periods of sustained economic growth, such as the postwar period, with a stronger effect when inequality is rising, as in the 1990s and 2000s (see Jaravel (2019)).

Figure 2: Descriptive Inflation Statistics

(a) Inflation with aggregate expenditures, 1984-2019

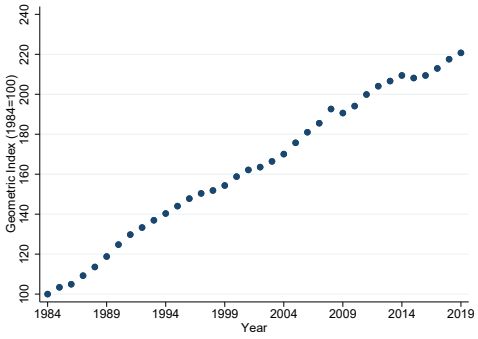

(b) Inflation by income percentiles, 1984-2019

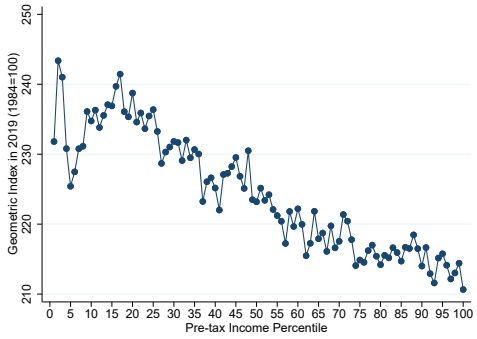

(c) Inflation with aggregate expenditures, 1955-2019

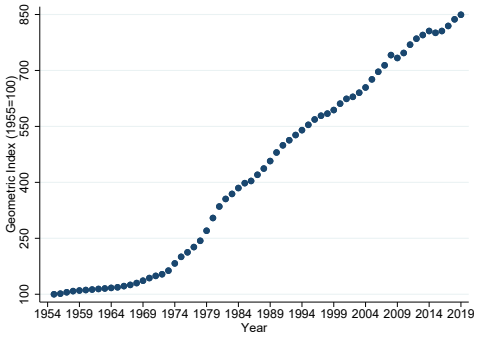

(d) Inflation by income percentiles, 1955-2019

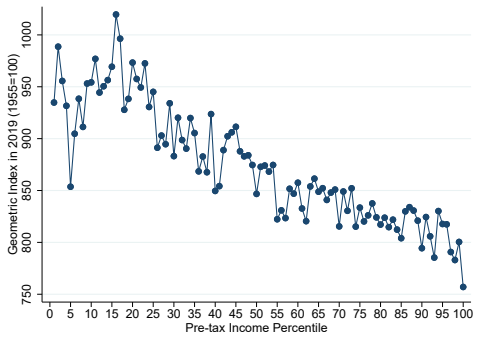

Note: This figure describes inflation patterns in our data. Panel (a) reports inflation from 1984 to 2019 using aggregate expenditure shares. Panel (b) shows heterogeneity in cumulative inflation rates between 1984 and 2019 by pre-tax income percentiles. In this panel, price indices are built using expenditure shares that are specific to each pre-tax income percentile. Panels (c) and (d) repeat the analysis for a longer period, from 1955 to 2019. All panels use chained geometric price indices.

3.2 Main Estimates

Analysis from 1984 to 2019 We first implement Algorithm 1 using our main dataset and the geometric price index formulas, leveraging the observed expenditure patterns and prices for each income percentile from 1984 to 2019. As we saw in Section 2, the negative covariance between household income and price indices shown in Figure 2 implies that the uncorrected measures of real consumption should underestimate the values of real consumption under any fixed base period. Indeed, this is what we find in panel (a) of Figure 3, which reports the bias in the average level of average real consumption, absent the nonhomotheticity correction, both under the initial and the final periods as the base.[33,34]

Using 1984 prices as the base, we find that the level of average real consumption (per house-

[33] Algorithm 1 is implemented using each pre-tax income percentile cell as one observation in the cross-section, and we then average the results. We use a second-order polynomial ($K = 2$) and show in sensitivity analyses below that the results remain similar for any $K \geq 1$.

[34] As already mentioned, we report the measures of average real consumption across households as a way of summarizing the results, without taking a stance on a welfare function (see footnote 5).

Figure 3: Nonhomotheticity Correction and Bias in Average Real Consumption, 1984-2019

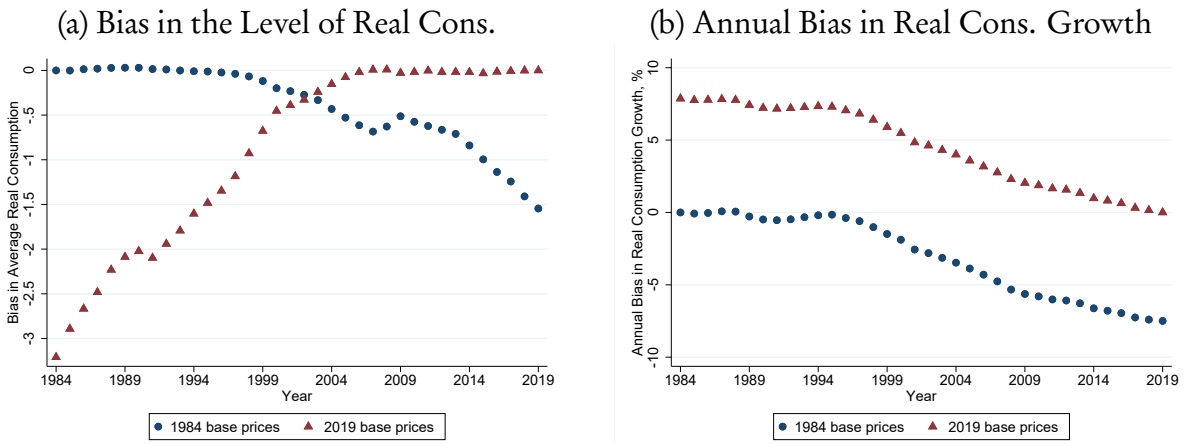

Note: This figure reports the biases in the level of average real consumption per household, in panel (a), and in annual growth in real consumption per household, in panel (b). The bias is computed by applying Algorithm 1 to obtain the nonhomotheticity correction. We then compare conventional measures of real consumption to corrected measures. In panel (b), the bias is expressed as a percentage of the standard homothetic measure of current-period growth. Algorithm 1 is applied to our main dataset at the level of pre-tax income percentiles, using geometric price indices. We then average percentile-level results to obtain average real consumption per household.

hold) is underestimated by about 1.5% in 2019. Mechanically, the bias in the level of real consumption is very small in the first few years after 1984. It grows gradually as the negative covariance between inflation and household income leads to a gradual change in the curvature of the expenditure function relative to the base year. Likewise, the panel shows that, using 2019 prices as the base, the level of real consumption in 1984 is underestimated by about 3.2%. Thus, due to the nonhomotheticity correction, at any point other than the base period we find that consumers are actually *better off* than what is implied by standard uncorrected measures. Intuitively, when we look into the past from the perspective of today's prices, we observe that (i) households were poorer 30 years ago and (ii) necessities were cheaper, which implies that consumer welfare 30 years ago was higher than that reported in conventional measures that ignore changes in the relative price of necessities and luxuries. Symmetrically, looking at today's economy from the perspective of prices in a distant period in the past, we observe that (i) households got richer and (ii) luxuries got cheaper; therefore welfare is higher than that reported in conventional measures that do not account for nonhomotheticity.

As shown in Panel (a) of Figure 3, the nonhomotheticity bias affecting the *level* of real consumption has the same sign regardless of the base year for prices. In contrast, the nonhomotheticity bias in the *growth* of real consumption does depend on the choice of base year. To see why, note that with 1984 prices as the base, real consumption growth is underestimated, since real consumption in the future is underestimated by the conventional measure without the nonhomotheticity correction. Symmetrically, with 2019 prices as the base, growth is overestimated since the level of real consumption is underestimated in all past periods. Panel (b) of Figure 3

reports these results, expressing the size of the bias as a share of measured growth.[35] With 1984 prices as the base, the conventional measure *under*estimates real consumption growth by about 7.5% in 2019. Taking 2019 prices as the base, the conventional measure *over*estimates real consumption growth by approximately 7.5% in 1984.

It is also instructive to examine the disaggregated patterns for the nonhomotheticity correction across pre-tax income percentiles. Figure 4 plots these results. Panels A report the bias in annual growth in real consumption for each income percentile. Panel A(i) focuses on growth in 2019, with 1984 prices as the base.[36] We find that the correction is larger for low-income groups: the annual growth in real consumption in 2019 is underestimated by 10% at the bottom of the income distribution, and only by 4% at the top. Symmetrically, Panel A(ii) shows that, with 2019 prices as the base, annual growth in 1984 is overestimated by about 9% at the bottom of the income distribution compared with 6% at the top.

Panels B of Figure 4 consider the biases for the levels of real consumption. The two panels show that the nonhomotheticity correction in levels is very similar across all income percentiles, with some noise inherent in survey data on expenditures. The effects in levels take into account the combination of annual corrections and percentile-specific growth rates, as accumulated over the full period.

Thus, the first key takeaway from our analysis is that the nonhomotheticity correction can be sizable and, given the observed patterns of inflation inequality, it generally implies that welfare over time is higher than commonly thought. The extent of the resulting bias in the level of real consumption is similar across income percentiles. Online Appendix Figure E.2 confirms this finding by reporting the chained index formula, $\Pi_t \pi_t^n$, compared with the corrected nonhomothetic deflator, y_t^n/c_t^n: the correction is similar in magnitude for all pre-tax income percentiles. To assess the quantitative relevance of the nonhomotheticity correction, it is instructive to compare its size to other sources of bias. In Online Appendix Figure E.3, we find that the size of the nonhomotheticity correction is of the same order of magnitude as the divergence between percentile-specific homothetic indices and the average homothetic index, which highlights the quantitative relevance of the nonhomotheticity correction.

[35] For each income percentile n, the annual bias in real consumption is defined as the difference between the uncorrected measure, $\Delta \log y_t^n - \pi_t^n$, and the corrected measure, $\Delta \log c_t^{b,n}$. Using the approximation in Equation (24), we thus define the bias as $\lambda_t^n \equiv \frac{\Delta \log y_t^n - \pi_t^n - \Delta \log c_t^{b,n}}{\Delta \log y_t^n - \pi_t^n} = \frac{\Lambda_t^b(c)}{\Lambda_t^b(c)+1}$. We compute the bias for each percentile and then average over all income percentiles.

[36] The biases are expressed as a share of measured growth, as given by λ_t^n defined in footnote 35 above for each percentile n in 2019.

Figure 4: Nonhomotheticity Correction and Biases in Real Consumption by Income Percentiles

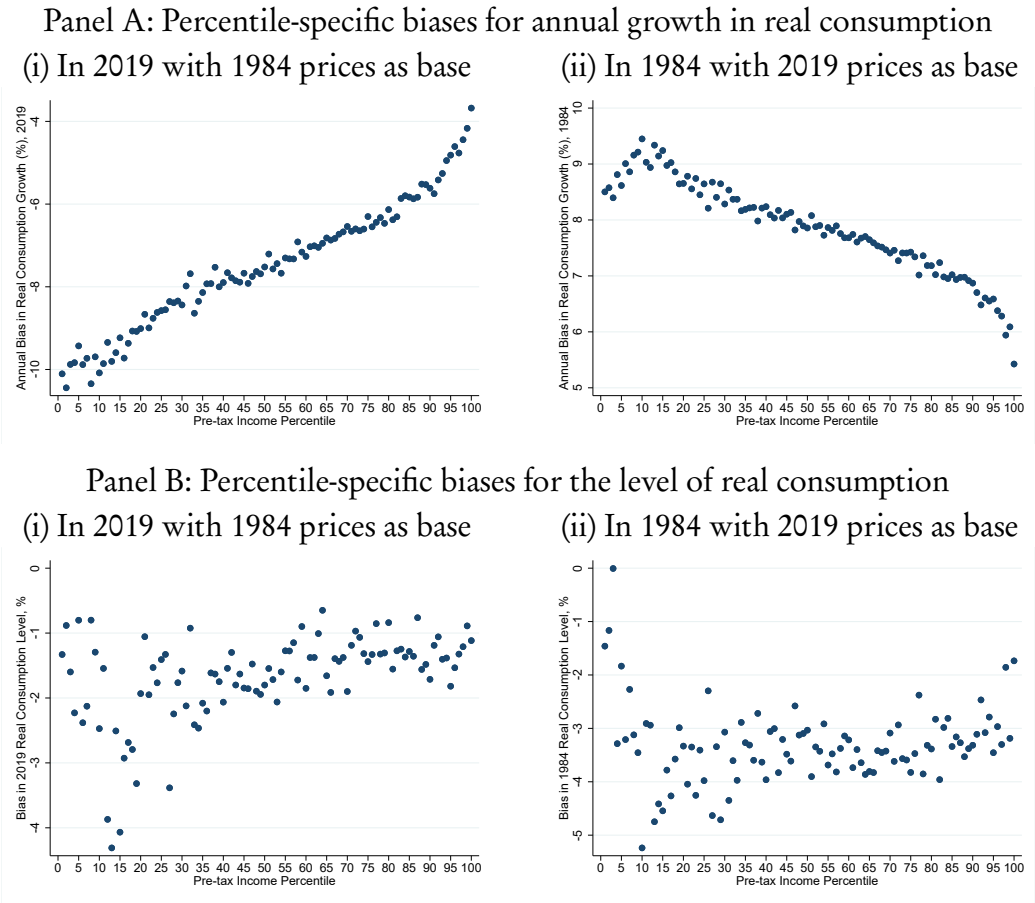

Note: This figure reports the biases in measures of real consumption due to the nonhomotheticity correction. The results for the annual growth in real consumption are depicted using 1984 prices as the base in panel A(i) and 2019 prices as the base in panel A(ii). Panel B reports the result for the bias in the level of real consumption. All panels use geometric price index formulas.

Analysis from 1955 to 2019 Next, we extend the analysis back to 1955, reporting the results in Figure 5.[37] Panel (a) reports the bias in levels; the patterns are identical to Figure 3 after 1984. With 1984 prices as the base, we find that the level of real consumption is underestimated by about 2% in both 1955 and 2019 due to the nonhomotheticity correction. As a result, the conventional measure of cumulative real consumption growth between 1955 and 2019 is not meaningfully affected by the nonhomotheticity correction, simply because the two biases in levels in 2019 and 1955 turn out to be of the same magnitude.[38]

[37] As explained in Appendix D, due to data limitations (i) we assume the expenditure shares observed in 1960 remain constant for the period 1955-1960, and (ii) we interpolate expenditure shares between years 1960 and 1972, and between 1972 and 1984.

[38] More generally, the biases in uncorrected measures are likely to vanish for *some* base period between any given initial and final periods in environments in which inflation always varies monotonically in income in the cross-section and nominal expenditure growth and inflation rates are stable over time. In such case, just like the case in Figure 5(a), the nonhomotheticity correction changes sign before and after the base period (see Figure 5(b)), and thus cancels out when the base period is somewhere in the middle of two periods under consideration. However,

With 2019 prices as the base, the nonhomotheticity correction becomes particularly large as we go back in time, because inflation inequality exists throughout the entire period and the nonhomotheticity correction accumulates over time. In 1955, average real consumption (per household) is underestimated by about 11.4% by the uncorrected measure. This finding shows that the nonhomotheticity correction can become large over long time horizons, depending on the choice of base prices.

Furthermore, Panel (b) of Figure 5 documents the bias in annual growth due to the nonhomotheticity correction. With 1984 prices as the base, the bias before and after 1984 changes sign. Specifically, it ranges from a positive bias of 5% in 1955 to a negative bias of -7% in 2019. In contrast, with 2019 prices as the base, the bias in annual consumption growth is always positive and becomes large as we go back in time, approaching 15% in 1955.

To better appreciate the magnitude of the nonhomotheticity correction, Panel (c) of Figure 5 reports cumulative consumption growth per household between 1955 and 2019; Panel (d) reports the same patterns by annualizing consumption growth. The standard uncorrected measure of cumulative consumption growth is 270% over this period, or 2.07% growth annually. With 1984 prices as the base, the nonhomotheticity correction leaves these patterns almost unchanged, implying a cumulative consumption growth of 267%. However, with 2019 prices as the base, the difference becomes large: cumulative consumption growth falls to 232%, or an annualized growth rate of 1.89% per year. Intuitively, from today's perspective, consumer welfare in the past was higher than conventionally thought, because income was lower in the past and necessities were relatively cheaper. Hence, real consumption growth was smaller than conventionally thought.

With 2019 prices as the base, the nonhomotheticity correction reduces the annual growth rate by 18 basis points, which is larger than the observed difference of 11 basis between Laspeyres and Paasche indices over the same time horizon. Online Appendix Figure E.4 reports the patterns for the Laspeyres and Paasche indices. Cumulative real consumption growth was 277% with the Paasche index, compared with 254% with Laspeyres, or a gap of 23 percentage points. By comparison, the nonhomotheticity correction induces a gap of 38 percentage points relative to the conventional measure. These results show that the magnitude of the nonhomotheticity correction can be as large as the well-known "expenditure switching bias" (or "substitution bias") affecting the Laspeyres and Paasche indices, which demonstrates its quantitative relevance.

note that this bias-free base period varies depending on the specific choices of these initial and final periods. In this example, while there is little bias for comparing average real consumption between 1955 and 2019, the comparison between 1955 and 1984 leads to an overestimation of the growth in real consumption.

Figure 5: Nonhomotheticity Correction and Bias in Average Real Consumption, 1984-2019

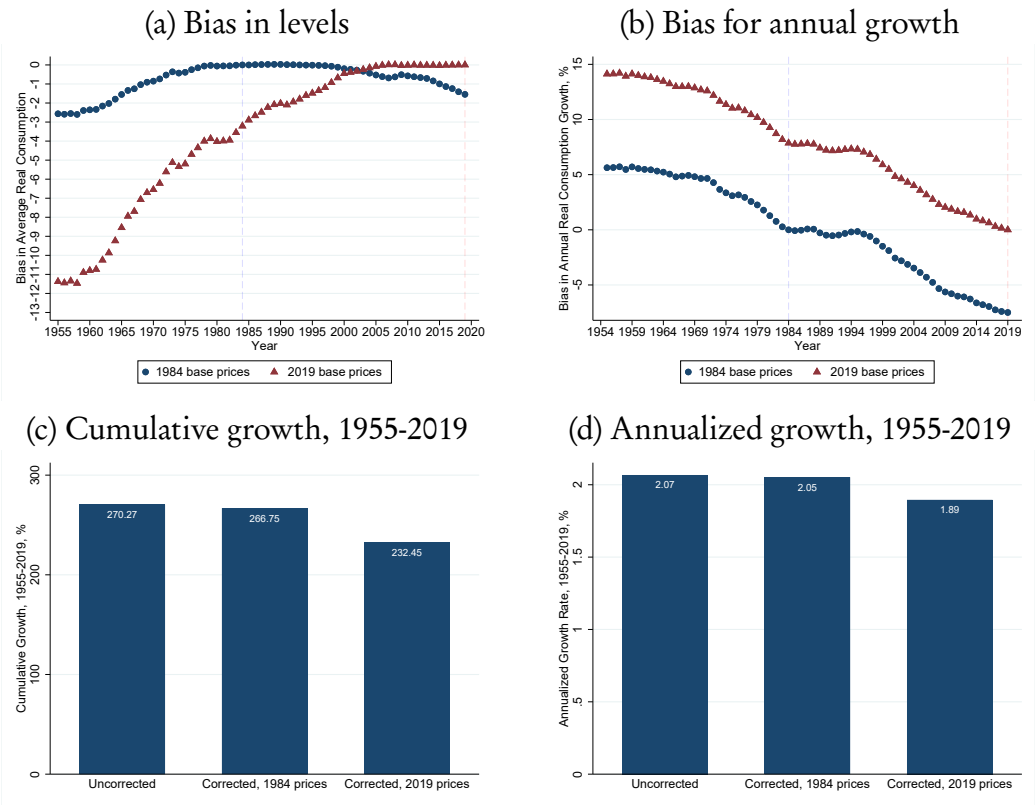

Note: This figure reports the biases in the level of average real consumption per household (Panel (a)) and in annual growth in real consumption per household (Panel (b)). The bias is computed by applying Algorithm 1 to obtain the nonhomotheticity correction at the level of pre-tax income percentiles; we then average percentile-level results to obtain average real consumption per household. Panels (c) and (d) report patterns of cumulative real consumption growth depending on the price index. All panels use geometric price indices.

3.3 Sensitivity Analysis

We now conduct several tests to assess the robustness of our findings. We first examine the sensitivity of our results to alternative price indices, the second-order algorithm, and the inclusion of controls, using the same dataset as in our baseline specifications. We then build alternative datasets to assess the stability of the results depending on data construction choices and the level of aggregation of expenditure data.[39]

Alternative Algorithms, Indices, and Controls We implement several sensitivity tests using the same datasets as in our baseline specifications. First, we assess the stability of the results when using a Fisher price index formula along with our first-order Algorithm 1, instead of using the geometric index formula. We also examine whether the results change when we use Algorithm A.3,

[39]In additional robustness checks, we find that the results remain similar when using higher-order polynomials to estimate the income elasticity of inflation, when keeping expenditure shares fixed at the 1984 or 2019 levels, and with quarterly instead of annual data (not reported).

which implements a second-order approximation. The results are shown in Panels A(i) and A(ii) of Figure 6: the patterns remain unchanged with the Fisher index as well as with the algorithm providing a second-order approximation.

Next, we assess whether the patterns remain similar when including controls. We implement Algorithm 1 as in Section 3.2, but we now add controls in the estimation of the income elasticity of inflation in constructing the nonhomotheticity correction. We first control for education, age, and race, reporting the results in Panel B(i) of Figure 6. We then introduce additional controls for region (Midwest, Northeast, West, South), rural vs. urban area, gender, and city population size, reporting the patterns in Panel B(ii). The patterns remain similar to those with our baseline specification without controls. Likewise, Online Appendix Figure E.5 shows that the annual bias in growth measurement remains almost unchanged when controls are included.[40]

Sensitivity Analysis with Alternative Datasets To assess the sensitivity of our findings to data construction choices, we build and study four alternative datasets.[41]

To document whether our results are sensitive to aggregation choices, we build two alternative datasets that closely follow our main dataset but use different levels of aggregation, grouping UCCs into broader categories. First, we create a version of the dataset at the level of the 32 product categories from CE summary tables, which are available from 1984 to 2019. Online Appendix Figure E.10 reports the results, applying Algorithm 1 to this dataset. The results are very similar to those obtained with our main dataset, with slightly smaller magnitudes due to the higher level of aggregation.[42]

Second, we manually group the 598 UCCs into 114 mutually exclusive product categories that are continuously available from 1984 to 2019. The results are reported in Online Appendix Figure E.11, showing that at this level of aggregation the results are almost indistinguishable from those obtained with the dataset in our main analysis.

[40]The Online Appendix reports additional sensitivity analyses. First, we assess the sensitivity of our results to the choice of the degree of the polynomial, K, when implementing Algorithm 1. Because the empirical relationship between the household-level inflation rate, π_t^n, and log real consumption, $\log \hat{c}_t^n$, is approximately log-linear during the period we study, we obtain very similar results for any $K \geq 1$. Appendix Figure E.6 reports the results for $K = 1$ and $K = 3$. Second, we analyze the data using the alternative algorithms described in the appendix, specifically the first-order refined algorithm (Algorithm A.1), the second-order refined algorithm (Algorithm A.4), and the algorithms based on estimation of the real consumption function to the first order (Algorithm A.2) and second order (Algorithm A.5). The results are very similar, as reported in Appendix Figure E.7. Third, we consider a specification controlling for state fixed effects and obtain similar results (Appendix Figure E.8). Fourth, we apply the algorithm of Baqaee et al. (2022) to our data (Appendix Figure E.9): we find that our baseline algorithm and their algorithm deliver very similar results, in line with our theoretical results showing that the two algorithms are equivalent to the first order.

[41]Online Appendix D provides a complete description of the data construction steps.

[42]The fact that the results are slightly weakened with more aggregated data was expected since inflation inequality is weaker when working with more aggregated product categories (Jaravel, 2019).

Figure 6: Sensitivity Analysis

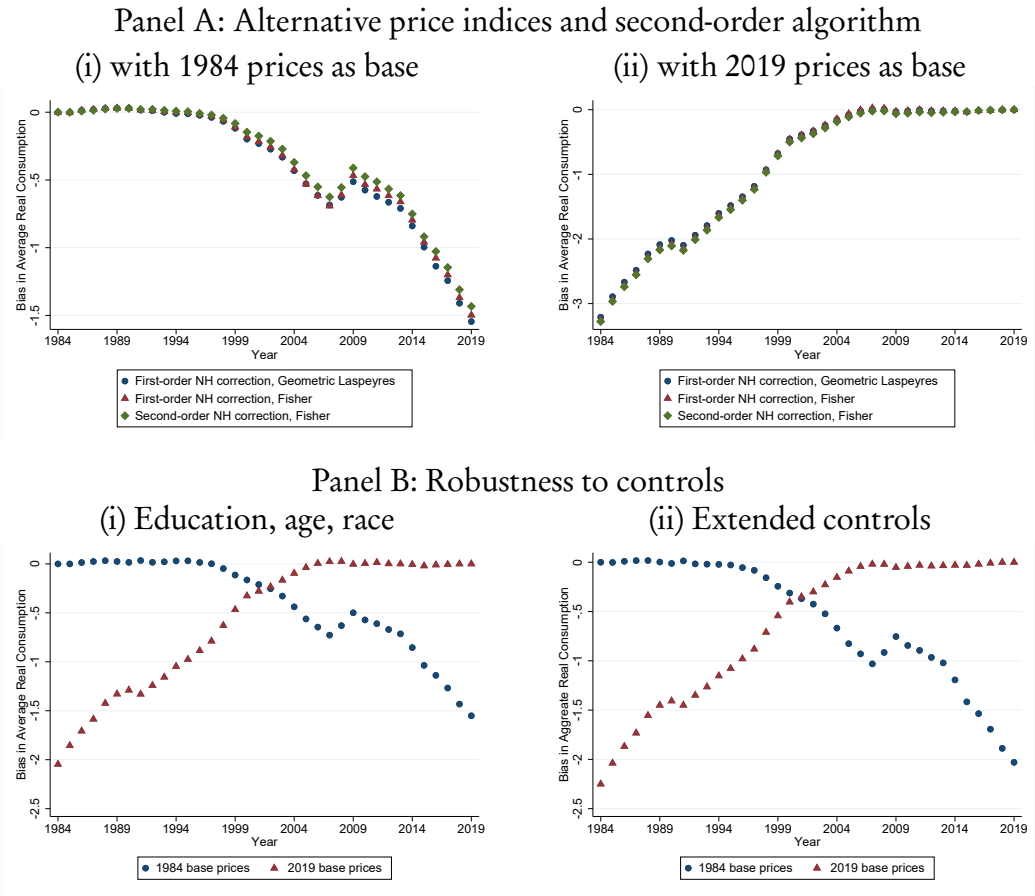

Note: This figure reports the biases in the level of average real consumption per household due to the nonhomotheticity correction under different specifications. Panel A reports the results under alternative price indices, geometric or Fisher, with the first-order algorithm, as well as with the second-order algorithm. Panel A(i) uses 1984 prices as the base, while Panel A(ii) uses 2019 prices. Panel B reports the results with controls, using the geometric index and the first order algorithm. Panel B(i) controls for education, age, and race in the estimation of the income elasticity of inflation. Panel B(ii) controls for region (midwest, northeast, west, south), urban vs. rural area, gender, and city population size, in addition to education, age, and race.

Moreover, to document the magnitude of the nonhomotheticity correction with highly disaggregated data, we implement our algorithm for a subset of expenditures for which product-level data is available, using Nielsen data covering consumer packaged goods, or about 15% of aggregate expenditure. This robustness check is motivated by prior work showing that most of the heterogeneity in inflation rates arises at the product level, within detailed product categories (Jaravel, 2019). We assess whether using product-level data meaningfully affects the size of the bias we estimate, at the cost of restricting attention to a subset of total expenditure. To implement this robustness check, we work with the Nielsen data from 2004 to 2014. Although the data cover a shorter time horizon, the annual level of inflation inequality is larger and the impact of the nonhomotheticity correction is stronger, as shown in Online Appendix Figure E.12 . The magnitude of the annual bias in real consumption growth increases faster than in our alternative

datasets, reaching 3% of the uncorrected measure after only a decade.[43]

Finally, we implement a robustness test inspired by the distributional national accounts of Piketty et al. (2018): we discipline our household-level data such that aggregate expenditure shares match exactly the official CPI consumption weights used by the Bureau of Labor Statistics (BLS) for eight product categories. Indeed, the BLS makes available the aggregate consumption weights used when calculating the CPI, which may differ from the expenditure shares in the CEX micro-data.[44] These weights are available at the level of eight consistent product categories from 1955 to 2019. We discipline our household-level CEX micro-data by introducing scaling factors, which are uniform across households but are allowed to vary across the eight categories, such that aggregate expenditure shares from our micro-data match exactly the aggregate consumption weights used by the BLS for the eight product categories.[45] This robustness check thus allows us to infer whether our results are sensitive to data construction choices about expenditure patterns. We obtain results very similar to those using our baseline dataset, as shown in Online Appendix Figure E.15. For example, using 2019 prices as the base, the average level of real consumption per household is underestimated by 11.7% in this robustness check, compared to 11.4% in the baseline specification.

Overall, these robustness checks show that the findings obtained with our baseline dataset are not sensitive to data construction choices. Moreover, the finding that the correction is stronger with more disaggregated data highlights the importance of using micro-data to accurately measure growth in consumer welfare with income-dependent preferences.

[43]To provide a precise comparison of the magnitude of the biases obtained with the Nielsen data, we repeat the analysis with our main CEX-CPI dataset restricted to the product categories covered in the Nielsen data between 2004 and 2014. The restricted CPI-CEX sample covers 44 UCC items belonging to the following categories: alcoholic beverage; food at home; personal care products; pets, toys, hobbies, and playground equipment; sewing machines, fabric and supplies; tools, hardware, outdoor equipment and supplies. The results are reported in Appendix Figure E.13: we find that the patterns remain qualitatively similar but are attenuated when we use the more aggregate CEX-CPI data. Taking 2004 prices as the base, the bias in the level of real consumption in 2014 is -0.056% with the Nielsen data, and -0.016% with the CEX-CPI sample; the bias in annual real consumption growth in 2014 is -2.83% with the Nielsen data, and -0.78% with the CEX-CPI sample. Thus, the biases are about 3.5 times larger with the detailed Nielsen data.The divergence between estimates is similar when we take 2014 prices as the base. Finally, we run an additional specification accounting for the welfare effect of new products in the Nielsen data. We account for the welfare effects of changes in product variety using a CES price index, which we compute for each of the 9131 Nielsen product categories using the methodology of Feenstra (1994), which was applied to scanner data in Broda and Weinstein (2010) and Jaravel (2019). The biases become larger because new goods create larger benefits for higher-income households, lowering their price indices and making the income elasticity of inflation more negative. Taking 2004 prices as the base, the bias in the level of real consumption in 2014 is -0.17% when accounting for changes in product variety; the bias in annual real consumption growth in 2014 is -7.99% (Appendix Figure E.14). Thus, compared with our baseline Nielsen estimates, the biases are about three times larger when accounting for new goods.

[44]The official CPI consumption weights are available at https://www.bls.gov/cpi/tables/relative-importance/home.htm.

[45]See Appendix D for a detailed description of this step.

4 Measuring Welfare Changes with Observed Heterogeneity

In this section, we extend the results of Section 2.2 to a setting that includes additional sources of observed consumer characteristics that change over time, beyond income. Examples of such characteristics include the age and education of consumers, or the number of household members. Focusing in particular on the case of age, we use our theory to quantify the correction to aggregate real consumption implied by consumer aging in the United States.

4.1 Correction for Change in Consumer Characteristics

Assume that we observe a vector of consumer characteristics (covariates) $\boldsymbol{x}_t \equiv (x_{dt})_{d=1}^{D} \in \mathbb{R}_+^D$ at time t.[46] We assume that consumer preferences are characterized by a well-behaved utility function $u = U(\boldsymbol{q}; \boldsymbol{x})$ that depends on the consumer characteristics. We let $y = E(u; \boldsymbol{p}, \boldsymbol{x})$ denote the corresponding expenditure function. As before, we assume a path of prices $\boldsymbol{p}_t$ and let $\omega_{i,t}(y; \boldsymbol{x})$ denote the expenditure share on good i for a consumer facing prices $\boldsymbol{p}_t$, with total expenditure y and characteristics $\boldsymbol{x}$. We first define our generalized concept of real consumption in this environment.

Definition 3 (Generalized Real Consumption)**.** For reference prices $\boldsymbol{p}_b$ (with $0 \leq b \leq T$), define *real consumption under period-b constant prices* for a consumer with utility u and characteristics $\boldsymbol{x}$ as a monotonic transformation $M_b(u, \boldsymbol{x})$ of utility given by

$$c^b = M_b(u; \boldsymbol{x}) \equiv E(u; \boldsymbol{p}_b; \boldsymbol{x}). \tag{30}$$

Definition 3 generalizes Definition 1 to a setting in which preferences potentially depend on consumer characteristics. We cannot compare welfare across consumers with different characteristics since they have distinct preferences. However, we can still compare the expenditure required by consumers with such distinct preferences for any level of welfare when they face identical prices. Therefore, we can state that the real consumption of a consumer with preferences $\boldsymbol{x}_t$ with utility u_t is higher than that of a consumer with preferences $\boldsymbol{x}_{t_0}$ and utility u_{t_0} by the amount $c_t^b - c_{t_0}^b \equiv M_b(u_t; \boldsymbol{x}_t) - M_b\left(u_{t_0}; \boldsymbol{x}_{t_0}\right)$, using reference prices $\boldsymbol{p}_b$.

Let us investigate the definitions above under two special cases. First, if consumer preferences do not change, i.e., $\boldsymbol{x}_t \equiv \boldsymbol{x}_{t_0}$, then the definition above reduces to our Definition 1, under homogeneous preferences. Second, if prices do not change, i.e., $\boldsymbol{p}_t \equiv \boldsymbol{p}_{t_0}$, the growth in real consumption simply accounts for the growth in nominal expenditure even if consumer characteristics change, $c_t^b / c_{t_0}^b \equiv y_t / y_{t_0}$.

[46] The assumption that the elements of the vector are positive valued is without loss of generality, as we can always transform the characteristic space in such a way that this condition holds.

In parallel to the definitions introduced in Section 2.1.1, we denote by $\chi_t^b(c;\boldsymbol{x}) \equiv E\left(M_b^{-1}(c;\boldsymbol{x});\boldsymbol{x}\right)$ the mapping from real consumption to expenditure at time t for a consumer with characteristic vector $\boldsymbol{x}$. The following proposition generalizes Proposition 1 to account for potential changes in consumer characteristics.

Proposition 3. *Consider a path of prices $\boldsymbol{p}_t$ and preferences that lead to the generalized Divisia index function $D_t(y;\boldsymbol{x}) \equiv \sum_i \omega_{i,t}(y;\boldsymbol{x}) \frac{d\log p_{it}}{dt}$ over the interval $[0,T]$. The mapping from real consumption to total expenditure $\chi_t^b(\cdot;\cdot)$ at time t is the solution to the following differential equation with initial condition $\chi_b^b(c;\boldsymbol{x}) = c$ for all $\boldsymbol{x}$:*

$$\frac{\partial \log \chi_t^b(c;\boldsymbol{x})}{\partial t} = \log D_t\left(\chi_t^b(c;\boldsymbol{x});\boldsymbol{x}\right). \tag{31}$$

In addition, for any path of total nominal expenditure y_t and vector of characteristic $\boldsymbol{x}_t$ over the interval, the growth in real consumption, defined under period-b constant prices, at any point in time satisfies

$$\frac{d\log c_t^b}{dt} = \frac{1}{1+\Lambda_t^b(c_t;\boldsymbol{x}_t)}\left[\frac{d\log y_t}{dt} - \log D_t(y_t;\boldsymbol{x}_t) - \sum_d \Gamma_{d,t}^b(c_t;\boldsymbol{x}_t)\frac{d\log x_{dt}}{dt}\right], \tag{32}$$

where the nonhomotheticity correction function $\Lambda_t(c;\boldsymbol{x})$ and the characteristic-d correction function $\Gamma_{dt}(c;\boldsymbol{x})$ are given by

$$\Lambda_t^b(c;\boldsymbol{x}) \equiv \frac{\partial \log \chi_t^b(c;\boldsymbol{x})}{\partial \log c} - 1, \qquad \Gamma_{d,t}^b(c;\boldsymbol{x}) \equiv \frac{\partial \log \chi_t^b(c;\boldsymbol{x})}{\partial \log x^d}. \tag{33}$$

Proof. See Appendix B.4. □

Proposition 3 extends the same insight behind Proposition 1 to the case with preferences that depend on consumer characteristics. It shows that the knowledge of the Divisia function is sufficient to uncover the mapping between real consumption and total consumption expenditure. The main difference is that we now need to know how the Divisia function depends both on total consumer expenditure and on consumer characteristics.

Let us now define the true price index $\mathscr{P}_{t_0,t}^b(c;\boldsymbol{x})$ under characteristic-dependent preferences:

$$\mathscr{P}_{t_0,t}^b(c;\boldsymbol{x}) \equiv \frac{\chi_t^b(c;\boldsymbol{x})}{\chi_{t_0}^b(c;\boldsymbol{x})}, \tag{34}$$

which is a generalization of the definition in Equation (3). This index measures the growth from period t_0 to t in the cost-of-living corresponding to a constant level of real consumption c for a

consumer with a constant vector of characteristics $\boldsymbol{x}$. As before, we can express the true price index as $\log \mathscr{P}^{b}_{t_0,t}(c;\boldsymbol{x}) = \int_{t_0}^{t} \log D_\tau\left(\chi^{b}_{\tau}(c;\boldsymbol{x});\boldsymbol{x}\right) d\tau$. By characterizing the mapping $\chi^{b}_{t}(c;\boldsymbol{x})$, Proposition 3 also fully characterizes the true price index in terms of the generalized Divisia function.

Proposition 3 further characterizes the instantaneous growth in real consumption. In addition to the nonhomotheticity correction, defined just like before, we also need the characteristic correction function index $\Gamma^{b}_{d,t} \equiv \frac{\partial \log \chi^{b}_{t}}{\partial \log x} \equiv \frac{\partial \log \mathscr{P}^{b}_{b,t}}{\partial \log x}$, which captures the elasticity of the true price index with respect to consumer characteristics. This index allows us to account for the effect of changing consumer preferences (through changes in observable characteristics) on real consumption. Similar to the nonhomotheticity correction function, these characteristic correction functions account for the cumulative cross-product covariance between price inflations and the elasticities of demand with respect to each characteristic:

$$\Gamma^{b}_{d,t}(c;\boldsymbol{x}) = \int_{b}^{t} \left[\sum_{i=1}^{I} \omega_{i,\tau}\left(\chi^{b}_{\tau}(c);\boldsymbol{x}\right) \zeta^{b}_{i,d,\tau}(c;\boldsymbol{x}) \frac{d \log p_{i,\tau}}{d\tau} \right] d\tau,$$

where $\zeta_{i,d,t}(c;\boldsymbol{x}) \equiv \frac{\partial \log \omega_{i,\tau}\left(\chi^{b}_{\tau}(c);\boldsymbol{x}\right)}{\partial \log x_d}$ accounts for the elasticity of the expenditure share of good-i with respect to characteristic d.

To see the intuition behind these results, consider an aging consumer and assume that inflation is on average higher for goods that are elastic with respect to age. In this case, over time there is an increase in the level of expenditure required to maintain the same level of real consumption for this consumer, due to the aging-induced reallocation of expenditure toward goods with prices that are rising faster. Holding prices fixed as in the initial period, Equation (32) shows that we need to deflate the growth in nominal expenditure by an additional term, $\frac{\partial \log \mathscr{P}^{b}_{b,t}(c_t;\boldsymbol{x}_t)}{\partial age_t} \frac{dage_t}{dt}$, to account for the effect of aging on real consumption growth. Thus, when reference prices are set as the initial base period, conventional measures of real consumption growth are biased upward because they do not account for the fact that, as people age, the relative prices of the products they favor increase. As in the case of nonhomotheticity, the sign of the bias inherently depends on the choice of the base period for prices. Holding prices fixed in the final period to express real consumption, conventional measures of real consumption growth are now biased downward since, going backward in time, consumers are getting younger and the relative prices of the products they favor is falling.

4.2 Approximating the Characteristic Correction Function

We generalize Algorithm 1 to account for variations in observable consumer characteristics and to approximate the characteristic correction function introduced in Section 4.1. Algorithms A.6 and A.7 in Appendix A.3 achieve these generalizations based on first-order and second-order price index formulas, respectively.

The idea underlying our approach is similar to that of Algorithm 1: starting in the base period, we nonparametrically estimate the relationship between the measured price index formulas across consumers and their total expenditures and other characteristics. We then use the estimated relationship with total expenditure and with other characteristics to approximate the corresponding correction functions.

4.3 Application to the Measurement of Real Consumption in the U.S. with Consumer Aging

In this section, we apply our approach to data from the US on aging and quantify the magnitude of the bias in conventional measures of real consumption growth.

Data and Summary Statistics To study the impact of consumer aging on real consumption growth, we build another version of the dataset in our main analysis where cells now correspond to age and income deciles, rather than income percentiles. Specifically, using the CEX data, in each year we define ten deciles of the (pre-tax) income distribution and, within each income decile, we compute ten age deciles. We then compute average age within each of these cells.[47]

Using this dataset, we compute inflation rates across age groups and find higher inflation rates for older households, as shown in Panel (i) of Figure 7. This panel reports the cumulative inflation rate by age deciles, using the geometric index between 1955 and 2019. The age elasticity of inflation is positive, especially for older ages. Between 1955 and 2019, cumulative inflation rates diverge by about 200 percentage points between the first and tenth age deciles. Thus, the relative prices of products purchased by younger households have been falling over time. To the best of our knowledge, this paper is the first to provide evidence on inflation inequality across age groups over a long time horizon. Online Appendix Figure E.16 reports additional patterns of inflation across groups, showing that the age elasticity of inflation is higher at older ages in all periods.

As reported in Online Appendix Figure E.17, average household age has been on the rise in the U.S., especially from 1970 onward. Therefore, by the logic of Section 4.1, conventional measures

[47] As in our main dataset, we use the data for 1960 and 1972 to interpolate expenditure shares in other years. Online Appendix D provides a complete description of the data construction steps.

Figure 7: Consumer Aging and Real Consumption

(a) Inflation by Age Decile, 1955-2019

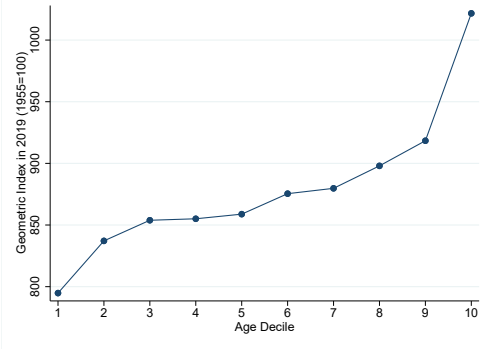

(b) Bias in the Level of Real Cons.

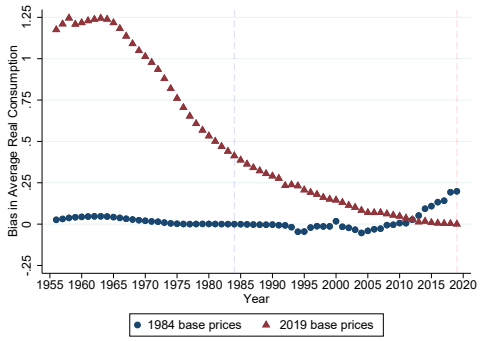

Note: Panel (a) of this figure reports the cumulative geometric laspeyres index, from 1955 to 2019, for each age decile. Panel (b) reports the bias in the level of real consumption per household due to the aging correction, relative to the nonhomothetic specification without the aging correction. Algorithm A.6 is applied to our dataset at the level of "age decile by income decile" units, using geometric Laspeyres price indices. We then average the results to obtain average real consumption per household with the aging correction.

of real consumption must be biased upward. We now proceed to quantify the magnitude of this bias.

Aging Correction for Average Real Consumption We apply Algorithm A.6 to quantify the adjustment to average real consumption implied by consumer aging. Panel (b) of Figure 7 reports the results. Specifically, we report the deviation in the level of average real consumption when accounting for both aging and nonhomotheticities, relative to the benchmark measure with only the nonhomotheticity correction.[48]

Using 2019 prices as the base, we find a meaningful aging correction: in 1955, the benchmark measure overestimates real consumption by about 1.2%. Intuitively, households in 1955 were on average younger than in 2019, and the prices of product categories purchased predominantly by younger households were higher. Therefore, society as a whole had lower real consumption in 1955 than commonly thought; that is, the conventional measure that does not account for consumer aging is biased upward.

Using 1984 as the base, the correction becomes much smaller, although it has the same sign. The benchmark measure overestimates real consumption by about 30 basis points in 2019. Intuitively, households are on average older in 2019 than in 1984 and the relative prices of goods purchased by older households have increased over time; that is, society is worse off in 2019 relative to conventional measures without the aging correction.[49]

[48] In the dataset with age-by-income cells used for our analysis in this section, the effect of the nonhomotheticity correction (relative to the standard homothetic real consumption measure) is close in magnitude to the bias shown in Section 3 with our baseline dataset using income percentiles.

[49] To understand the difference in the magnitude of the aging correction depending on the choice of base years,

In sum, these patterns illustrate that changes in consumer characteristics such as age can have a meaningful effect on the measurement of average real consumption, depending on the choice of base prices. In the case of aging, the adjustments are economically meaningful but much smaller than the nonhomotheticity correction, which justifies our focus on the latter. While there is a strong relationship between age and inflation, the correction to average real consumption implied by aging is smaller than the nonhomotheticity correction primarily because the change in average household age over time is relatively slow.

5 Conclusion

In this paper, we extended the results of the classical index number theory to settings in which the composition of demand depends on income (nonhomotheticity) and other consumer characteristics. We developed a procedure for nonparametric measurement of consumer welfare based on price index formulas, imposing minimal restrictions on the underlying preferences. This approach remains valid under any observable household heterogeneity in preferences, and requires only data on spending patterns in a cross-section of households.

We showed the practical relevance of the correction for nonhomotheticities when computing long-run growth in consumer welfare. With our correction taking 2019 prices as base, growth in consumer welfare is significantly attenuated in the United States in the post-war era, due to the combination of fast growth and lower inflation for income-elastic products. The correction reduces the annual growth rate from 1955 to 2019 by 18 basis points, which is larger than the "expenditure switching bias" affecting Laspeyres and Paasche indices over the same time horizon. Extending this analysis to other countries and time periods, as well as to the measurement of purchasing power parity (PPP) indices across countries with preference heterogeneity, is a promising direction for future research.

Our results may have important implications for the way in which national statistical agencies around the world construct measures of real economic value. Our approach has the potential to be widely adopted for at least three reasons. First, our approach has a light data requirement. It combines standard price data with information from surveys of consumer expenditures, which are typically available to statistical agencies as these surveys are already used in the construction of homothetic price indices. In Section 3, we have offered a blueprint for how the Bureau of Labor Statistics (BLS) can use data that is already available to construct improved measures of real consumption growth and inequality in the US. Second, our approach has a light computational burden. In its first-order renditions, our algorithms simply require one cross-sectional regression

note that the speed of consumer aging is slower before the 1980s, and that the covariance between inflation and household age is also weaker before the 1980s, as shown in Online Appendix Figures E.16 and E.17.

per period to construct the required corrections, irrespective of the number of products considered. Our second-order algorithms also converge in a few steps per period when applied to U.S. data. Third, our approach closely follows the standard practice for constructing real economic values by deflating year-on-year growth in nominal values by price index formulas. Our algorithms construct first- and second-order corrections to these standard formulas to account for the role of income-dependence in preferences. Our approach thus allows statistical agencies to transparently examine the contribution of the nonhomotheticity adjustments to their measures. The tight connection between the nonhomotheticity corrections and observed inflation inequality in the cross-section of households further strengthens the transparency of our procedure.

Due to its light computational and data requirements, our approach can readily be used by statistical agencies to generate distinct series of real consumption (or panels across different income quantiles) for all base years for which cross-sectional data is available. Depending on their goals, different data users may opt to rely on data expressed in terms of a different base periods. For instance, if a government program has determined in a certain year that households should be eligible for some benefits if they are below a given consumption threshold (e.g., the poverty line), then this year constitutes a suitable base year for tracking household consumption and potential changes in this threshold in all future years. In contrast, if the goal is to evaluate consumption growth over long time horizons in a way that could be best understood by households today, using today's prices offers a suitable base to express measures of real consumption, insofar as households are likely to better understand money metrics based on the prices they currently face. We believe these and other applications of our framework are fruitful directions for statistical agencies going forward.

References

Antonelli, Giovanni B, "Sulla teoria matematica della economia politica (Pisa, Folchetto; republished with an introduction by G. Demaria and Comments by G. Ricci, Milan: Falfari, 1952)," 1886. A12

Argente, David and Munseob Lee, "Cost of Living Inequality during the Great Recession," *Journal of the European Economic Association*, 2021, *19* (2), 913–952. 2, 7, 28

Atkin, David, Benjamin Faber, Thibault Fally, and Marco Gonzalez-Navarro, "Measuring Welfare and Inequality with Incomplete Price Information," 2020. 7, 13

Baqaee, David and Ariel Burstein, "Welfare and Output with Income Effects and Demand Instability," 2021. 2, 7, 13

Baqaee, David R, Ariel T Burstein, and Yasutaka Koike-Mori, "Measuring Welfare with Income Effects using Cross-Sectional Data," 2022. 7, 22, 35, A0, A13, A14, A29, A30, A42

Battistin, Erich, Richard Blundell, and Arthur Lewbel, "Why Is Consumption More Log Normal than Income? Gibrat's Law Revisited," *Journal of Political Economy*, 2009, *117* (6), 1140–1154. 25

Blackorby, Charles and David Donaldson, "Money metric utility: A harmless normalization?," *Journal of Economic Theory*, 1988, *46* (1), 120–129. 5

Blundell, Richard W, Martin Browning, and Ian A Crawford, "Nonparametric Engel curves and revealed preference," *Econometrica*, 2003, *71* (1), 205–240. 6

Bosmans, Kristof, Koen Decancq, and Erwin Ooghe, "Who's afraid of aggregating money metrics?," *Theoretical Economics*, 2018, *13* (2), 467–484. 5

Broda, Christian and David E Weinstein, "Product Creation and Destruction: Evidence and Price Implications," *American Economic Review*, 2010, *100* (3), 691–723. 37, A45

Comin, Diego, Danial Lashkari, and Marti Mestieri, "Structural change with long-run income and price effects," *Econometrica*, 2021, *89* (1), 311–374. 5, 24, 25

Deaton, Angus and John Muellbauer, "An almost ideal demand system," *American Economic Review*, 1980, *70* (3), 312–326. 7

Diewert, W Erwin, "Exact and superlative index numbers," *Journal of Econometrics*, 1976, *4* (2), 115–145. 6, 17, 24

_ , "Superlative Index Numbers and Consistency in Aggregation," *Econometrica*, 1978, *46* (4), 883–900. 24

Diewert, Walter E, "The economic theory of index numbers: a survey," in Walter E Diewert and Alice O Nakamura, eds., *Essays in Index Number Theory*, Emerald Group Publishing Ltd, 1993, chapter 7, pp. 163–208. 2, 6, 10, 24

Engel, Ernst, "Die Produktions- und Consumtionsverhältnisse des Königreichs Sachsen," *Zeitschrift des Statistischen Büreaus des Königlich Sächischen Ministeriums des Innern*, 1857, *8-9*, 1–54. 2

Fajgelbaum, Pablo D and Amit K Khandelwal, "Measuring the Unequal Gains from Trade," *Quarterly Journal of Economics*, 2016, *131* (3), 1113–1180. 7, 13

Feenstra, Robert C, "New Product Varieties and the Measurement of International Prices," *American Economic Review*, 1994, *84* (1), 157–77. 37, A45

_ , **Hong Ma, and DS Prasada Rao**, "Consistent comparisons of real incomes across time and space," *Macroeconomic Dynamics*, 2009, *13* (S2), 169–193. 7

Hobijn, Bart and David Lagakos, "Inflation inequality in the United States," *Review of Income and Wealth*, 2005, *51* (4), 581–606. 7

Hulten, Charles R, "Divisia Index Numbers," *Econometria*, 1973, *41* (6), 1017–1025. 12

Jaravel, Xavier, "The unequal gains from product innovations: evidence from the US retail sector," *Quarterly Journal of Economics*, 2019, *134* (2), 715–783. 2, 7, 28, 35, 36, 37, A37, A45

_ , "Inflation Inequality: Measurement, Causes, and Policy Implications," *Annual Review of Economics*, 2021, *13*, 599–629. 2

Kaplan, Greg and Sam Schulhofer-Wohl, "Inflation at the household level," *Journal of Monetary Economics*, 2017, *91*, 19–38. 2, 7, 28

Klick, Josh and Anya Stockburger, "Experimental CPI for lower and higher income households," 2021. 2

McGranahan, Leslie and Anna Paulson, "Constructing the Chicago Fed Income Based Economic Index–Consumer Price Index: Inflation Experiences by Demographic Group: 1983-2005," 2006. 7

Newey, Whitney K., "Convergence rates and asymptotic normality for series estimators," *Journal of Econometrics*, 1997, *79* (1), 147–168. 20, A4

Oberfield, Ezra, "Inequality and Measured Growth," 2022. 4, 15

Oulton, Nicholas, "Chain indices of the cost-of-living and the path-dependence problem: An empirical solution," *Journal of Econometrics*, 2008, *144* (1), 306–324. 7

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman, "Distributional national accounts: methods and estimates for the United States," *Quarterly Journal of Economics*, 2018, *133* (2), 553–609. 3, 27, 37, A36

Pollak, Robert A, "The theory of the cost-of-living index," in W Erwin Diewert, ed., *Price Level Measurement*, Elsevier North Holland Inc., 1990, pp. 5–77. 6

Redding, Stephen J. and David E. Weinstein, "Measuring Aggregate Price Indices with Taste Shocks: Theory and Evidence for CES Preferences," *Quarterly Journal of Economics*, 2020, *135* (1), 503–560. 7

Samuelson, Paul A, "The problem of integrability in utility theory," *Economica*, 1950, *17* (68), 355–385. A12

_ **and S Swamy**, "Invariant Economic Index Numbers and Canonical Duality: Survey and Synthesis," *American Economic Review*, 1974, *64* (4), 566–593. 2, 6, 10, 24

Slesnick, Daniel T, "Aggregate deadweight loss and money metric social welfare," *International Economic Review*, 1991, pp. 123–146. 5

Appendix to "Measuring Growth in Consumer Welfare with Income-Dependent Preferences"

Xavier Jaravel, London School of Economics
Danial Lashkari, Federal Reserve Bank of New York

June 2023

Contents

A Alternative Algorithms

A.1 First-Order Algorithms without Observable Heterogeneity

A.1.1 Baseline First-Order Algorithm

See Section 2.3.4 in the main text for the formal statement of our baseline first-order algorithm. We can motivate the above algorithm based on the results of Lemma B.1 in Appendix B.4.2 below, and in particular the combination of Equations (B.17) and (B.19). See further discussions in Section 2.3.4.

A.1.2 Refined First-Order Algorithm

To use the Taylor series expansion in Equation (28), we need to also apply a similar expansion to the right hand side of Equation (18):

$$\pi_t^+\left(\chi_t\left(c_{t+1}^{n,(\ell)}\right)\right)=\pi_t^+\left(\chi_t\left(c_{t+1}^{n,(\ell-1)}\right)\right)+\frac{\partial\,\pi_t^+\left(y_t^{n,(\ell-1)}\right)}{\partial\log y_t^{n,(\ell-1)}}\cdot\frac{\partial\log\chi_t\left(c_{t+1}^{n,(\ell-1)}\right)}{\partial\log c_{t+1}^{n,(\ell-1)}}\log\left(\frac{c_{t+1}^{n,(\ell)}}{c_{t+1}^{n,(\ell-1)}}\right)+O\left(\epsilon^2\right),$$

where $y_t^{n,(\ell-1)}$ and ϵ are defined as in Equation (28). Combining the expansions on the two sides of Equation (18) leads to the following update rule for the estimate of real consumption $\widehat{c}_{t+1}^{n,(\ell)}$:

$$\log\widehat{c}_{t+1}^{n,(\ell)}=\log\widehat{c}_{t+1}^{n,(\ell-1)}+\frac{1}{1+\widehat{\Lambda}_{t+1}^{n,(\ell-1)}}\left[\log\left(\frac{y_{t+1}^n}{y_t^{n,(\ell-1)}}\right)-\widehat{\pi}_t^+\left(y_t^{n,(\ell-1)}\right)\right],\tag{A.1}$$

where the adjusted nonhomotheticity correction is given by

$$\widehat{\Lambda}_{t+1}^{n,(\ell-1)}\equiv\frac{\partial\log\widehat{\chi}_t\left(c_{t+1}^{n,(\ell-1)}\right)}{\partial\log c_{t+1}^{n,(\ell-1)}}\left(1+\frac{\partial\,\pi_t^+\left(y_t^{n,(\ell-1)}\right)}{\partial\log y_t^{n,(\ell-1)}}\right)-1,$$

where $\widehat{\chi}_t(\cdot)$ and $\widehat{\pi}_t^+(\cdot)$ denote the estimates of the mapping from real consumption to expenditure and the geometric index function at time t.

Using the above result, we now state the refined Algorithm A.1.

Algorithm A.1 (Refined First-Order Algorithm)**.** *Consider a sequence of power functions* $\{f_k(z)\equiv z^k\}_{k=0}^{K_N}$ *for some* K_N*, where* N *is the number of consumers in the cross-section. Let* $\widehat{c}_b^n\equiv y_b^n$ *and* $\widehat{\chi}_b(c)=\chi_b(c)=c$*. For each* $t\geq b$*, successively apply the following two steps.*

1. *Nonparametrically fit the true price index between periods t and $t+1$:*

   Let $\widehat{\pi}_t^+(y) \equiv \sum_{k=0}^{K_N} \alpha_{k,t}^+ f_k(\log y)$, based on the estimates of the coefficients $\left(\widehat{\alpha}_{k,t}\right)_{k=0}^{K_N}$ that solve the following problem:

$$\min_{\left(\alpha_{k,t}^+\right)_{k=0}^{K}} \sum_{n=1}^{N} \left(\pi_t^{+,n} - \sum_{k=0}^{K_N} \alpha_{k,t}^+ f_k\left(\log y_t^n\right) \right)^2, \tag{A.2}$$

   where $\left\{\pi_t^{+,n}\right\}_n$ are household-specific price index formulas at time t defined by Equation (26).

2. *Initialize the values of real consumption in period $t+1$ using the baseline first-order algorithm:*

   Evaluate $\widehat{c}_{t+1}^{n,(0)}$ using Algorithm 1.

3. *Iteratively find the real consumption in period $t+1$:*

   Iterate over the following steps over $\ell \in \{0, 1, \cdots\}$ until convergence for some tolerance $\epsilon \ll 1$.

   (a) *Update the values of real consumption for consumers in period $t+1$ using Equation (A.1).*

   (b) *Stopping criterion: if $\max_n \left| \widehat{c}_{t+1}^{n,(\ell)} - \widehat{c}_{t+1}^{n,(\ell-1)} \right| < \epsilon$ and set $\widehat{c}_{t+1}^{n} \equiv \widehat{c}_{t+1}^{n,(\ell)}$.*

4. *Estimate the values of real consumption for consumers in period $t+1$:*

   Let $\widehat{\chi}_{t+1}(c) \equiv \sum_{k=0}^{K_N} \beta_{k,t} f_k(\log c)$, based on the estimates of the coefficients $\left(\widehat{\beta}_{k,t+1}\right)_{k=0}^{K_N}$ that solve the following problem:

$$\min_{\left(\beta_{k,t}\right)_{k=0}^{K}} \sum_{n=1}^{N} \left(y_{t+1}^n - \sum_{k=0}^{K_N} \beta_{k,t+1} f_k\left(\log c_{t+1}^n\right) \right)^2, \tag{A.3}$$

To apply the algorithm backward in time for $t < b$, simply re-label all the time indices in the data preceding the base period b such that $t - \tau \to t + \tau$ for all $1 \le \tau \le b$ and perform the same steps as above.

Analysis of Error To analyze the resulting errors in the estimates, first we present the following lemma that states the uniform convergence rates for the approximation of the geometric index function in Equation (A.2).

Lemma A.1. *Assume that the expenditure function $E(u; \boldsymbol{p})$ is continuously differentiable of order $m \ge 3$. Moreover, assume that household specific price indices are observed with zero mean noise, and that, at time t, the distribution of household expenditure y^n is i.i.d, and with a probability density function that is bounded away from zero over a compact interval $\left[\underline{y}_t, \overline{y}_t\right]$. Then, the estimated*

Laspeyres and Paasche geometric index functions converge uniformly to the corresponding true function as $N \to \infty$ *and* $K_N \to \infty$ *such that* $K_N^3/N \to 0$*, at the rate given by*

$$\widehat{\pi}_t^{+,-}(y) = \pi_t^{+,-}(y) + O_p\left(K_N\left(\sqrt{\frac{K_N}{N}} + K_N^{-(m-1)}\right)\right).$$

Proof. The result follows from Theorem 1 of Newey (1997) for the case of power series. The fact that the expenditure function is continuously differentiable of order m implies that the expenditure shares $\omega_{it}(y)$ are differentiable of order $m-1$. Thus, the assumptions of the lemma correspond to Assumptions 8 and 9 of Newey (1997), ensuring the desired result.[A1] □

As the sample size in the cross-section in each period grows toward infinity, Lemma A.1 shows that the nonparametric approximations of the Laspeyres and Paasche geometric index functions converge toward the true underlying functions. The remaining error in the estimates of $\widehat{c}_{t+1}^n$ generated by the algorithm in each period stems from (i) the first-order approximation in Equation (18), and (ii) the error in the fitted function $\widehat{\chi}_t(\cdot)$ in the previous period.

A.1.3 Algorithm Based on the Indirect Real Consumption Function

Our baseline first-order and second-order Algorithms 1 and A.3 approximate the nonhomotheticity correction as the inverse of the elasticity of the mapping from expenditure to real consumption $\chi_t^b(c)$. In this section, we provide parallels to these algorithm that rely on the elasticity of the inverse of that mapping, which is the indirect real consumption function $\widetilde{\chi}_t^b(y)$. Unlike our baseline algorithms, this approach does not allow us to draw an intuitive and transparent connection between the nonhomotheticity correction and the observed heterogeneity in inflation across households. Nevertheless, as we see in the simulation exercise in Appendix C, the evolution of approximation error in the approach based on the real consumption function could be distinct from, and at times improve over, that in our baseline algorithms.

Algorithm A.2 (First-Order Algorithm based on the Real Consumption Function (RCF))**.** *Consider a sequence of power functions* $\{f_k(z) \equiv z^k\}_{k=0}^{K_N}$ *for some* K_N*, where* N *is the number of consumers in the cross-section. Let* $\widehat{c}_b^n \equiv y_b^n$ *and for each* $t \geq b$*, successively apply the following two steps.*

1. *Nonparametrically fit the real consumption function* $\widetilde{\chi}_t(\cdot)$ *at time* t*:*

   Estimate the coefficients $\left(\widehat{\alpha}_{k,t}\right)_{k=0}^{K_N}$ *solving the following problem:*

$$\min_{(\alpha_{k,t})_{k=0}^{K}} \sum_{n=1}^{N}\left(\log \widehat{c}_t^n - \sum_{k=0}^{K_N} \alpha_{k,t} f_k(\log y_t^n)\right)^2, \tag{A.4}$$

[A1] In the notation of Newey (1997), this case corresponds to $r \equiv 1$, $d = 0$, $s = m-1$.

where $\{\widehat{c}_t^n\}_n$ *are the current period's household-specific values of real consumption.*

2. *Estimate the nonhomotheticity correction and the values of real consumption for consumers in period* $t+1$*:*

   Use Equation (24) to compute the values of each household's real consumption, where the estimate of the nonhomotheticity correction for household n *at time* t *is given by*

$$\widehat{\Lambda}_{t+1}^n \equiv \left[\sum_{k=0}^{K_N} \widehat{\alpha}_{k,t} f_k'(\log y_t^n)\right]^{-1} - 1. \tag{A.5}$$

We can also motivate the above algorithm based on the results of Lemma B.1 in Appendix B.4.2 below, and in particular the combination of Equations (B.17) and (B.19).

A.2 Second-Order Algorithms without Observable Heterogeneity

A.2.1 Baseline Second-Order Algorithm

In this Section, we provide an extension of our baseline algorithm that allow us to approximate changes in real consumption growth to the second order of approximation.

Algorithm A.3 (Baseline Second-Order Algorithm)**.** *Let* $\widehat{c}_b^n \equiv y_b^n$ *, and consider the sequence of power functions* $\{f_k(z) \equiv z^k\}_{k=0}^{K_N}$ *where* K_N *grows with* N*, the number of consumers in the cross-section. For each* $t \geq b$*, apply the following steps:*

1. *Initialize the values of the real consumption in period* $t+1$ *using the first-order algorithm:*

   Evaluate $\widehat{c}_{t+1}^{n,(0)}$ *using Algorithm 1.*

2. *Iteratively find the real consumption in period* $t+1$*:*

   Iterate over the following steps over $\ell \in \{0, 1, \cdots\}$ *until convergence for some tolerance* $\epsilon \ll 1$*.*

   (a) *Nonparametrically fit a first-order term needed for finding the true price index between periods* t *and* $t+1$*:*
   Solve for the coefficients $\left(\widehat{\alpha}_{k,t}^{\dagger,(\ell)}\right)_{k=0}^{K}$ *in the following problem:*

$$\min_{\left(\alpha_{k,t}^{\dagger}\right)_{k=0}^{K}} \sum_{n=1}^{N} \left(\pi_t^{+,n} - \sum_{k=0}^{K} \alpha_{k,t}^{\dagger} f_k\left(\log \widehat{c}_{t+1}^{n,(\ell)}\right)\right)^2, \tag{A.6}$$

   where $\pi_t^{+,n}$ *are the geometric indices defined by Equation (26). .*

(b) *Nonparametrically fit the true price index between periods t and $t+1$:*
Solve for the coefficients $\left(\widehat{\beta}_{k,t}^{(\ell)}\right)_{k=0}^{K}$ *in the following problem:*

$$\min_{(\beta_{k,t})_{k=0}^{K}} \sum_{n=1}^{N}\left(\pi_t^{*,n}+\rho_t^{n,(\ell)}-\sum_{k=0}^{K}\beta_{k,t}f_k\left(\log \widehat{c}_t^{\,n}\right)\right)^2, \tag{A.7}$$

with $\rho_t^{n,(\ell)}$ *is defined as:*

$$\rho_t^{n,(\ell)} \equiv \frac{1}{4}\sum_{k=0}^{K}\widehat{\alpha}_{k,t}^{\dagger,(\ell)}\left[f_k'\left(\log \widehat{c}_t^{\,n}\right)+f_k'\left(\log \widehat{c}_{t+1}^{\,n,(\ell)}\right)\right]\log\left(\frac{\widehat{c}_{t+1}^{\,n,(\ell)}}{\widehat{c}_t^{\,n}}\right). \tag{A.8}$$

and where $\pi_t^{*,n}$ *is the value of a second-order Törnqvist price index for consumer n, that is,* $\pi_t^{*,n} \equiv \log \mathbb{P}_T\left(\boldsymbol{p}_t, \boldsymbol{s}_t^n; \boldsymbol{p}_{t+1}, \boldsymbol{s}_{t+1}^n\right)$.

(c) *Estimate the values of real consumption for consumers in period $t+1$:*
Update the real consumption in the next period for each consumer

$$\log \widehat{c}_{t+1}^{\,n,(\ell+1)} = \log \widehat{c}_t^{\,n,(\ell)} + \frac{1}{1+\frac{1}{2}\left[\widehat{\Lambda}_t\left(\widehat{c}_t^{\,n}\right)+\widehat{\Lambda}_{t+1}^{(\ell)}\left(\widehat{c}_{t+1}^{\,n,(\ell)}\right)\right]}\left(\log\left(\frac{y_{t+1}^n}{y_t^n}\right)-\pi_t^{*,n}\right), \tag{A.9}$$

where we have defined the approximate nonhomothetic correction function as:

$$\widehat{\Lambda}_{t+1}^{(\ell)}(c) \equiv \widehat{\Lambda}_t(c)+\sum_{k=0}^{K}\widehat{\beta}_{k,t}^{(\ell)}f_k'(\log c). \tag{A.10}$$

(d) *Stopping criterion: if* $\max_n\left|\widehat{c}_{t+1}^{\,n,(\ell+1)}-\widehat{c}_{t+1}^{\,n,(\ell)}\right| < \epsilon$ *and set* $\widehat{c}_{t+1}^{\,n} \equiv \widehat{c}_{t+1}^{\,n,(\ell+1)}$ *and* $\widehat{\Lambda}_{t+1}(c) \equiv \widehat{\Lambda}_{t+1}^{(\ell)}(c)$.

We can motivate the above algorithm based on the results of Lemma B.1 in Appendix B.4.2 below, and in particular the combination of Equations (B.18) and (B.20). The key complication stems from the results presented in Equation (19), which show that for a second-order approximation of the true price index function $\mathscr{P}_{b,t+1}^{\,b}(c)$ defined in Equation (12), we need to keep the real consumption constant between the two periods. In step (b) of the algorithm above, we use an adjustment of the Törnqvist index to address this point. As a result, function $\widehat{\mathscr{P}}_{b,t+1}(c) \equiv \sum_{k=0}^{K}\left(\sum_{\tau=b}^{t}\widehat{\beta}_{k,\tau}\right)f_k(\log c)$ provides a second-order approximation for the true price index function $\mathscr{P}_{b,t+1}^{\,b}(c)$. Equation (A.9) then updates our current guess $\widehat{c}_{t+1}^{\,n,(\ell)}$ about next-period real consumption.

A.2.2 Refined Second-Order Algorithm

Algorithm A.4 (Refined Second-Order Algorithm)**.** *Consider a sequence of power functions* $\{f_k(z) \equiv z^k\}_{k=0}^{K_N}$ *for some* K_N*, where* N *is the number of consumers in the cross-section. Let* $\widehat{c}_b^n \equiv y_b^n$ *and* $\widehat{\chi}_b(c) = \chi_b(c) = c$*. For each* $t \geq b$*, successively apply the following two steps.*

1. *Nonparametrically fit the true price index between periods* t *and* $t+1$*:*

   Let $\widehat{\pi}_t^{\pm}(y) \equiv \sum_{k=0}^{K_N} \alpha_{k,t}^{\pm} f_k(\log y)$*, based on the estimates of the coefficients* $\left(\widehat{\alpha}_{k,t}^{\pm}\right)_{k=0}^{K_N}$ *that solve the following problem:*

$$\min_{\left(\alpha_{k,t}^{\pm}\right)_{k=0}^{K}} \sum_{n=1}^{N} \left(\pi_t^{\pm,n} - \sum_{k=0}^{K_N} \alpha_{k,t}^{\pm} f_k\left(\log y_t^n\right) \right)^2, \tag{A.11}$$

   where $\left\{\pi_t^{\pm,n}\right\}_n$ *are household-specific price index formulas at time* t *defined by Equation (26) and* $\pi_t^{-,n} \equiv \log \mathbb{P}_G\left(\boldsymbol{p}_t, \boldsymbol{s}_{t+1}^n; \boldsymbol{p}_{t+1}, \boldsymbol{s}_t^n\right)$*.*

2. *Initialize the values of the real consumption in period* $t+1$ *using the refined first-order algorithm:*

   Evaluate $\widehat{c}_{t+1}^{n,(0)}$ *using Algorithm A.1.*

3. *Iteratively find the real consumption in period* $t+1$*:*

   Iterate over the following steps over $\ell \in \{0, 1, \cdots\}$ *until convergence for some tolerance* $\epsilon \ll 1$*.*

   (a) *Update the values of real consumption for consumers in period* $t+1$ *:*

   Apply the rule

$$\log \widehat{c}_{t+1}^{n,(\ell)} = \log \widehat{c}_{t+1}^{n,(\ell-1)} + \frac{1}{1+\widehat{\Lambda}_{t+1}^{n,(\ell-1)}} \left[\log\left(\frac{y_{t+1}^n}{y_t^{n,(\ell-1)}} \right) - \frac{1}{2}\left(\widehat{\pi}_t^{+}\left(y_t^{n,(\ell-1)}\right) + \widehat{\pi}_t^{-}\left(y_{t+1}\right) \right) \right], \tag{A.12}$$

   where we have defined $y_t^{n,(\ell-1)} \equiv \chi_t\left(\widehat{c}_{t+1}^{n,(\ell-1)}\right)$ *and*

$$\widehat{\Lambda}_{t+1}^{n,(\ell-1)} \equiv \frac{\partial \log \widehat{\chi}_t\left(c_{t+1}^{n,(\ell-1)}\right)}{\partial \log c_{t+1}^{n,(\ell-1)}} \left(1 + \frac{1}{2} \frac{\partial \widehat{\pi}_t^{+}(y)}{\partial \log y}\bigg|_{y=y_t^{n,(\ell-1)}} \right) - 1,$$

   (b) *Stopping criterion: if* $\max_n \left| \widehat{c}_{t+1}^{n,(\ell)} - \widehat{c}_{t+1}^{n,(\ell-1)} \right| < \epsilon$ *and set* $\widehat{c}_{t+1}^n \equiv \widehat{c}_{t+1}^{n,(\ell)}$*.*

4. *Estimate the values of real consumption for consumers in period* $t+1$*:*

Let $\widehat{\chi}_{t+1}(c) \equiv \sum_{k=0}^{K_N} \beta_{k,t+1} f_k(\log c)$*, based on the estimates of the coefficients* $\left(\widehat{\beta}_{k,t+1}\right)_{k=0}^{K_N}$ *that solve the following problem:*

$$\min_{\left(\beta_{k,t+1}\right)_{k=0}^{K}} \sum_{n=1}^{N} \left(y_{t+1}^{n} - \sum_{k=0}^{K_N} \beta_{k,t+1} f_k\left(\log c_{t+1}^{n}\right) \right)^2, \tag{A.13}$$

To apply the algorithm backward in time for $t < b$*, simply re-label all the time indices in the data preceding the base period* b *such that* $t - \tau \to t + \tau$ *for all* $1 \leq \tau \leq b$ *and perform the same steps as above.*

This algorithm follows along similar lines as those of the first-order refined Algorithm A.1, but relies on Equation (19) instead of Equation (18) to achieve a second-order approximation.

A.2.3 Algorithm Based on the Real Consumption Function

Algorithm A.5 (Second-Order Algorithm based on the Real Consumption Function (RCF))**.** *Consider a sequence of power functions* $\{f_k(z) \equiv z^k\}_{k=0}^{K_N}$ *for some* K_N*, where* N *is the number of consumers in the cross-section. Let* $\widehat{c}_b^{\,n} \equiv y_b^n$*,* $\widetilde{\chi}_b^b(y) \equiv y$*, and for each* $t \geq b$*, successively apply the following two steps.*

1. *Initialize the values of real consumption in period* $t+1$ *using the first-order algorithm:*

   Evaluate $\widehat{c}_{t+1}^{n,(0)}$ *using Algorithm A.2.*

2. *Iteratively find real consumption in period* $t+1$*:*

   Iterate over the following steps over $\ell \in \{0, 1, \cdots\}$ *until convergence for some tolerance* $\epsilon \ll 1$*.*

   (a) *Nonparametrically fit a first-order approximation of the real consumption function at time* $t+1$*:*

   Estimate the coefficients $\left(\widehat{\alpha}_{k,t}^{(\ell)}\right)_{k=0}^{K_N}$ *solving the following problem:*

$$\min_{\left(\alpha_{k,t+1}^{(\ell)}\right)_{k=0}^{K}} \sum_{n=1}^{N} \left(\log \widehat{c}_{t+1}^{n,(\ell)} - \sum_{k=0}^{K_N} \alpha_{k,t+1}^{(\ell)} f_k\left(\log y_t^n\right) \right)^2, \tag{A.14}$$

   where $\left\{\widehat{c}_{t+1}^{n,(\ell)}\right\}_n$ *are the next period's household-specific values of real consumption in iteration* ℓ*.*

   (b) *Estimate the nonhomotheticity correction and the values of real consumption for consumers in period* $t+1$*:*

Update the real consumption in the next period for each consumer

$$\log \widehat{c}_{t+1}^{n,(\ell+1)} = \log \widehat{c}_t^n + \frac{1}{1+\frac{1}{2}\left(\widehat{\Lambda}_t^n + \widehat{\Lambda}_{t+1}^{n+1,(\ell)}\right)}\left(\log\left(\frac{y_{t+1}^n}{y_t^n}\right) - \pi_t^{*,n}\right), \tag{A.15}$$

where we use the Törnqvist index formula $\pi_t^{*,n} \equiv \log \mathbb{P}_T\left(\boldsymbol{p}_t, \boldsymbol{s}_t^n; \boldsymbol{p}_{t+1}, \boldsymbol{s}_{t+1}^n\right)$*, and where the estimates of the nonhomotheticity correction for household* n *at times* t *and* $t+1$ *are given by*

$$\widehat{\Lambda}_t^n \equiv \left[\sum_{k=0}^{K_N} \widehat{\alpha}_{k,t} f_k'(\log y_t^n)\right]^{-1} - 1, \tag{A.16}$$

$$\widehat{\Lambda}_{t+1}^{n,(\ell)} \equiv \left[\sum_{k=0}^{K_N} \widehat{\alpha}_{k,t+1}^{(\ell)} f_k'\left(\log y_{t+1}^n\right)\right]^{-1} - 1. \tag{A.17}$$

(c) Stopping criterion: *if* $\max_n \left|\widehat{c}_{t+1}^{n,(\ell+1)} - \widehat{c}_{t+1}^{n,(\ell)}\right| < \epsilon$*, set* $\widehat{c}_{t+1}^n \equiv \widehat{c}_{t+1}^{n,(\ell+1)}$ *and* $\widehat{\alpha}_{k,t+1} \equiv \widehat{\alpha}_{k,t+1}^{(\ell+1)}$.

We can motivate the above algorithm based on the results of Lemma B.1 in Appendix B.4.2 below, and in particular the combination of Equations (B.18) and (B.20).

A.3 Algorithms with Observed Heterogeneity

In this section, we provide the algorithms that allow us to approximate changes in real consumption growth to the first and second orders of approximation in the presence of changes in household preferences that relate to observable household characteristics.

First-Order Algorithm

Algorithm A.6. *Let* $\widehat{c}_b^n \equiv c_b^n \equiv c_b^n$ *and consider a sequence* $\{f_k(c, \boldsymbol{x})\}_{k=0}^{K_N}$ *of log-power functions of* c *and* $\boldsymbol{x}$ *where* K_N *depends on* N*, the number of consumers in the cross-section. For each* $t \geq b$*, apply the following steps:*

1. Nonparametrically fit the true price index between periods t and $t+1$:

   Find the coefficients $\left(\widehat{\alpha}_{k,t}\right)_{k=0}^{K_N}$ *that solve the following problem:*

$$\min_{\left(\alpha_{k,t}\right)_{k=0}^{K}} \sum_{n=1}^{N}\left(\pi_t^{+,n} - \sum_{k=0}^{K_N} \alpha_{k,t} f_k\left(\widehat{c}_t^n, \boldsymbol{x}_t^n\right)\right)^2, \tag{A.18}$$

   where $\pi_t^{+,n}$ *is defined by Equation (26).*

2. *Estimate the values of real consumption for consumers in period $t+1$:*

   Compute the real consumption in the next period for each household:

$$\log \widehat{c}_{t+1}^{n} = \log \widehat{c}_{t}^{n} + \frac{1}{1+\widehat{\Lambda}_{t+1}(\widehat{c}_{t}^{n};\boldsymbol{x}_{t}^{n})}\left[\log\left(y_{t+1}^{n}/y_{t}^{n}\right) - \pi_{t}^{+,n} - \sum_{d=1}^{D}\widehat{\Gamma}_{d,t+1}(c_{t}^{n};\boldsymbol{x}_{t}^{n})\cdot\log\left(\frac{x_{d,t+1}}{x_{d,t}}\right)\right] \tag{A.19}$$

   where we have defined the approximate nonhomotheticity correction function as:

$$\widehat{\Lambda}_{t+1}(c;\boldsymbol{x}) = \sum_{k=0}^{K_N}\left(\sum_{\tau=b+1}^{t+1}\widehat{\alpha}_{k,\tau}\right)\frac{\partial f_k(c,\boldsymbol{x})}{\partial \log c}, \tag{A.20}$$

   and the following approximation for the characteristic-d correction function:

$$\widehat{\Gamma}_{d,t+1}(c;\boldsymbol{x}) = \sum_{k=0}^{K_N}\left(\sum_{\tau=b+1}^{t+1}\widehat{\alpha}_{k,\tau}\right)\frac{\partial f_k(c,\boldsymbol{x})}{\partial \log x_d}.$$

Second-Order Algorithm Next, we provide a second-order approximation holding under arbitrary observed heterogeneity across households. Algorithm A.7 thus provides a generalization of Algorithm A.3 to the cases involving observed heterogeneity.

Algorithm A.7. *Let $\widehat{c}_{b}^{n} \equiv c_{b}^{n} \equiv y_{b}^{n}$ and consider a sequence $\{g_k(c,\boldsymbol{x})\}_{k=0}^{K_N}$ of log-power functions of c and $\boldsymbol{x}$ where N is the number of households in the cross-section. For each $t \geq b$, apply the following steps:*

1. *Initialize the values of real consumption in period $t+1$ using the first-order algorithm:*

   Initialize the values of the real consumption $\widehat{c}_{t+1}^{n,(0)}$ for each household at $t+1$ using the results of Algorithm A.6.

2. *Apply the loop to find real consumption in period $t+1$:*

   Iterate over the following steps over $\ell \in \{0,1,\cdots\}$ until convergence for some tolerance $\epsilon \ll 1$.

   (a) *Nonparametrically fit a first-order term needed for finding the true price index between periods t and $t+1$:*

   Solve for the coefficients $\left(\widehat{\alpha}_{k,t}^{\dagger,(\ell)}\right)_{k=0}^{K}$ in the following problem:

$$\min_{\left(\alpha_{k,t}^{\dagger}\right)_{k=0}^{K}} \sum_{n=1}^{N}\left(\pi_{t}^{+,n} - \sum_{k=0}^{K}\alpha_{k,t}^{\dagger} f_k\left(\widehat{c}_{t+1}^{n,(\ell)},\boldsymbol{x}_{t+1}^{n}\right)\right)^2, \tag{A.21}$$

where $\pi_t^{+,n}$ is defined by Equation (26).

(b) Nonparametrically fit the true price index between periods t and $t+1$:
Find the coefficients $\left(\widehat{\beta}_{k,t}^{(\ell)}\right)_{k=0}^{K}$ that solve the following problem:

$$\min_{(\beta_{k,t})_{k=0}^{K}} \sum_{n=1}^{N} \left(\pi_t^{*,n} + \rho_t^{n,(\ell)} - \sum_{k=0}^{K} \beta_{k,t} f_k(\widehat{c}_t^n, \boldsymbol{x}_t^n) \right)^2, \tag{A.22}$$

where $\pi_t^{*,n} \equiv \mathbb{P}_T\left(\boldsymbol{p}_t, \boldsymbol{s}_t^n; \boldsymbol{p}_{t+1}, \boldsymbol{s}_{t+1}^n\right)$ and where $\rho_t^{n,(\ell)}$ is defined as:

$$\rho_t^{n,(\ell)} \equiv \frac{1}{4} \sum_{k=0}^{K} \widehat{\alpha}_{k,t}^{\dagger,(\ell)} \left[\frac{\partial f_k(\widehat{c}_t^n, \boldsymbol{x}_t^n)}{\partial \log \widehat{c}_t^n} + \frac{\partial f_k\left(\widehat{c}_{t+1}^{n,(\ell)}, \boldsymbol{x}_{t+1}^n\right)}{\partial \log \widehat{c}_{t+1}^{n,(\ell)}} \right] \log\left(\frac{\widehat{c}_{t+1}^{n,(\ell)}}{\widehat{c}_t^n} \right) \tag{A.23}$$

$$+ \frac{1}{4} \sum_{d=1}^{D} \sum_{k=0}^{K} \widehat{\alpha}_{k,t}^{\dagger,(\ell)} \left[\frac{\partial f_k(\widehat{c}_t^n, \boldsymbol{x}_t^n)}{\partial \log x_{d,t}^n} + \frac{\partial f_k\left(\widehat{c}_{t+1}^{n,(\ell)}, \boldsymbol{x}_{t+1}^n\right)}{\partial \log x_{d,t+1}^n} \right] \log\left(\frac{x_{d,t+1}}{x_{d,t}} \right).$$

(A.24)

(c) Estimate the values of real consumption for consumers in period $t+1$:
Update real consumption in the next period for each household:

$$\log \widehat{c}_{t+1}^{n,(\ell+1)} = \log \widehat{c}_t^{n,(\ell)} + \frac{1}{1 + \frac{1}{2}\left[\widehat{\Lambda}_t(\widehat{c}_t^n; \boldsymbol{x}_t^n) + \widehat{\Lambda}_{t+1}^{(\ell)}\left(\widehat{c}_{t+1}^{n,(\ell)}; \boldsymbol{x}_{t+1}^n\right) \right]} \tag{A.25}$$

$$\times \left[\log\left(\frac{y_{t+1}^n}{y_t^n} \right) - \frac{1}{2} \sum_{d=1}^{D} \left[\widehat{\Gamma}_{d,t}(c_t^n; \boldsymbol{x}_t^n) + \widehat{\Gamma}_{d,t+1}^{(\ell)}\left(c_{t+1}^n; \boldsymbol{x}_{t+1}^n\right) \right] \cdot \log\left(\frac{x_{d,t+1}}{x_{d,t}} \right) \right], \tag{A.26}$$

where we have defined the approximate correction functions as:

$$\widehat{\Lambda}_{t+1}^{(\ell)}(c; \boldsymbol{x}) = \widehat{\Lambda}_t(c; \boldsymbol{x}) + \sum_{k=0}^{K} \widehat{\beta}_{k,t}^{(\ell)} \frac{\partial f_k(c, \boldsymbol{x})}{\partial \log c}, \tag{A.27}$$

$$\widehat{\Gamma}_{d,t+1}^{(\ell)}(c; \boldsymbol{x}) = \widehat{\Gamma}_{d,t}(c; \boldsymbol{x}) + \sum_{k=0}^{K} \widehat{\beta}_{k,t}^{(\ell)} \frac{\partial f_k(c, \boldsymbol{x})}{\partial \log x_d}, \tag{A.28}$$

(d) Stopping criterion: if $\max_n \left| \widehat{c}_{t+1}^{n,(\ell+1)} - \widehat{c}_{t+1}^{n,(\ell)} \right| < \epsilon$ and set $\widehat{c}_{t+1}^n \equiv \widehat{c}_{t+1}^{n,(\ell+1)}$, $\widehat{\Lambda}_{t+1}(c; \boldsymbol{x}) \equiv \widehat{\Lambda}_{t+1}^{(\ell)}(c; \boldsymbol{x})$, and $\widehat{\Gamma}_{d,t+1}(c; \boldsymbol{x}) \equiv \widehat{\Gamma}_{d,t+1}^{(\ell)}(c; \boldsymbol{x})$.

B Theory Appendix

B.1 Proposition 1 and the Integrability Problem

Proposition 1 has a tight link to the classical integrability problem, which has long been studied in economics (e.g., Antonelli, 1886; Samuelson, 1950). Consider a demand system $\boldsymbol{\omega}(y;\boldsymbol{p})$ that gives the vector of expenditure shares of the household as a function of total expenditure and the vector of prices, satisfying $\sum_i \omega_i(y;\boldsymbol{p}) = 1$ and $\boldsymbol{\omega}(\alpha y;\alpha \boldsymbol{p}) = \boldsymbol{\omega}(y;\boldsymbol{p})$ for all y, $\boldsymbol{p}$, and $\alpha > 0$. The integrability problem asks whether we can uncover the corresponding expenditure function $E(u;\boldsymbol{p})$ from the knowledge of such a demand system. The answer is that the demand system is integrable if and only if the corresponding Slutsky matrix is negative semidefinite and symmetric. Under this assumption, the expenditure function is given as the solution to the following system of differential equations

$$\frac{\partial \log E(u;\boldsymbol{p})}{\partial \log p_i} = \omega_i(E(u;\boldsymbol{p});\boldsymbol{p}), \qquad \text{for all } i, \tag{B.1}$$

with the initial condition $E(u;\boldsymbol{p}_0) = y_0$ for a fixed vector of prices $\boldsymbol{p}_0$ and different choices of y_0.

We can think of Proposition 1 as an application of the above result along a specific path of prices. In particular, we can write

$$\begin{aligned}
\frac{\partial \log \chi_t(c)}{\partial t} = \frac{\partial \log E\left(M_b^{-1}(c);\boldsymbol{p}_t\right)}{\partial t} &= \sum_i \frac{\partial \log E\left(M_b^{-1}(c);\boldsymbol{p}_t\right)}{\partial \log p_{it}} \frac{d \log p_{it}}{dt}, \\
&= \sum_i \omega_{i,t}(\chi_t(c)) \frac{d \log p_{it}}{dt}, \\
&= \log D_t(\chi_t(c)),
\end{aligned}$$

where in the second equality we have substituted from Equation (B.1), using $\omega_{i,t}(\chi_t(c)) \equiv \omega_i(\chi_t(c);\boldsymbol{p}_t)$. Thus, since we assume that the conditions for integrability hold, the differential equation in Proposition 1 uniquely characterizes the path of the mapping from real consumption to expenditure.

The key distinction between the classical integrability problem and our results are as follows. First, the classical problem aims to recover the full structure of preferences (in terms of an expenditure function or an indirect utility function), given the knowledge of the full demand function. In contrast, we only aim to recover the value of the indirect utility (in money metric terms) for specific combinations of expenditures and prices observed in the data. Second, the classical problem aims to characterize the set of demand functions that may correspond to a valid expen-

diture (or indirect utility) function. In contrast, we already begin with the assumption that the observed data corresponds to demand that is rationalizable by a utility function that corresponds to a smooth expenditure function. Finally, our goal is to find the expenditure function expressed in terms of the particular cardinalization given by the money metric utility. In sum, our paper does not provide new results about the rationalizability of observed demand. Instead, it studies how to uncover money metric utility values corresponding to any combination of prices and expenditures observed in the data, assuming that they are rationalized by a smooth expenditure function.

B.2 Characterization of the Real Consumption Function

We can characterize the real consumption function $\widetilde{\chi}_t^b(y)$, the inverse mapping from expenditure to real consumption in terms of prices in the base period b. This mapping is the indirect utility function, in money-metric terms corresponding to the base period b, and is defined such that

$$\widetilde{\chi}_t^b\left(\chi_t^b(c)\right) \equiv c, \qquad \forall t, c. \tag{B.2}$$

Since this equation holds for any c, we can substitute for c from Equation (B.2) in Equation (11) to find

$$\log \chi_t^b\left(\widetilde{\chi}_t^b(y)\right) = \log \widetilde{\chi}_t^b(y) + \int_b^t \log D_\tau\left(\chi_\tau^b\left(\widetilde{\chi}_t^b(y)\right)\right) d\tau,$$

which leads to the following integral representation for this function:[A2]

$$\log \widetilde{\chi}_t^b(y) = \log y - \int_b^t \log D_\tau\left(\chi_\tau^b\left(\widetilde{\chi}_t^b(y)\right)\right) d\tau. \tag{B.3}$$

For completeness, we state a corollary (and dual) of Proposition 1 to directly characterize the mapping $\widetilde{\chi}_t^b(\cdot)$ as a differential equation in terms of the Divisia index function $\log D_t(\cdot)$.

Corollary 1. *The mapping $\widetilde{\chi}_t^b(\cdot)$ from expenditure to real consumption (under prices in base period b) satisfies the following partial differential equation*

$$\frac{\partial \log \widetilde{\chi}_t^b(y)}{\partial t} + \frac{\partial \log \widetilde{\chi}_t^b(y)}{\partial \log y} \cdot \log D_t(y) = 0, \qquad \textit{for all } t \neq b, \tag{B.4}$$

with the boundary condition $\widetilde{\chi}_b^b(y) = y$. In addition, for any path of total nominal expenditure

[A2] We note that Equation (B.3) first appeared in Baqaee et al. (2022) who, subsequent to our paper, used this integral representation of the solution of the differential equation (9) to construct an alternative algorithm for approximating real consumption, using a logic similar to ours. We study the properties of their approach in relation to ours in Appendix B.3.

y_t over the interval, the growth in real consumption, defined under period-b constant prices, at any point in time satisfies

$$\frac{d\log c_t^b}{dt} = \frac{\partial \log \widetilde{\chi}_t^b(y_t)}{\partial \log y_t} \times \left(\frac{d\log y_t}{dt} - \log D_t(y_t)\right). \tag{B.5}$$

Proof. Taking the full time derivative of $\log \widetilde{\chi}_t^b\left(\chi_t^b(c)\right) \equiv \log c$ we find

$$\begin{aligned} 0 &= \left.\frac{\partial \log \widetilde{\chi}_t^b(y)}{\partial t}\right|_{y=\chi_t(c)} + \left.\frac{\partial \log \widetilde{\chi}_t^b(y)}{\partial \log y}\right|_{y=\chi_t(c)} \cdot \frac{\partial \log \chi_t^b(c)}{\partial \log t}, \\ &= \left.\frac{\partial \log \widetilde{\chi}_t^b(y)}{\partial t}\right|_{y=\chi_t(c)} + \left.\frac{\partial \log \widetilde{\chi}_t^b(y)}{\partial \log y}\right|_{y=\chi_t(c)} \cdot \log D_t\left(\chi_t^b(c)\right), \end{aligned}$$

where in the second equality, we have used Equation (9) to substitute for $\frac{\partial \log \chi_t^b(c)}{\partial \log t}$. Substituting for $y = \chi_t^b(c)$ leads to Equation (B.4). Taking the derivative of $\chi_t^b\left(\widetilde{\chi}_t^b(y)\right) \equiv y$ with respect to y, we find that $\frac{\partial \log \widetilde{\chi}_t^b(y)}{\partial \log y_t} = \left(\frac{\partial \log \chi_t^b(c)}{\partial \log c_t^b}\right)^{-1}$ for $y = \chi_t^b(c)$, which, when combined with Equation (10), leads to Equation (B.5). □

In this paper, we rely on Equation (9), which simplifies the first-order hyperbolic partial differential equation (B.4) into an ordinary differential equation. The differential equation (9) achieves this simplification by constructing the characteristic curves of the function $\widetilde{\chi}_t^b(y)$, that is, the curves along which the value of the mapping remains constant $c = \widetilde{\chi}_t^b(y)$.

More generally, we may combine numerical solutions of the partial differential equation (B.4) with approximations of the Divisa index function to find tighter error bounds for the approximation errors. Thus, there are many alternative approaches to approximately solve the partial differential equation (B.4) based on cross-sectional data. The algorithms presented in Appendix A offer a number of such alternatives.

B.3 Comparison with Baqaee et al. (2022)

In this section of the appendix, we offer a comparison of the approach that Baqaee et al. (2022) (henceforth, BBK) have proposed subsequent to our paper as an alternative to our algorithm. Their approach relies on the discretization of the integral in Equation (B.3) that characterizes the real consumption function as a solution to the differential equation in Proposition 1, which can be approximated as

$$\log \widetilde{\chi}_t(y) \approx \log y - \sum_{\tau=0}^{t-1} \pi_\tau^+\left(\chi_\tau\left(\widetilde{\chi}_t(y)\right)\right),$$

where $\pi_t^+(\cdot)$ stands for the Laspeyres geometric index function.

Overview of BBK's Algorithms BBK provide two alternative algorithms to approximately solve for Equation (B.3) and characterize the real consumption function, always taking the first period as the base $b = 0$. As a result, to simplify the exposition in this section, we drop the superscript $b = 0$. The BBK algorithm relies on estimating I expenditure share functions $\widehat{\omega}_{i,t}(y)$ in terms of total expenditure y, based on observed total expenditure y_t^n and expenditure shares s_{it}^n across households at each time t. They find each function $\widehat{\omega}_{i,t}(\cdot)$ by estimating coefficients $\widehat{\beta}_{k,i,t}$ such that

$$\widehat{\omega}_{i,t}(y) = \sum_k \widehat{\beta}_{k,i,t} f_k(y), \tag{B.6}$$

fits the observed collections of expenditure shares s_{it}^n and expenditures y_t^n across households n. They use the I estimated functions to construct an approximation for the integral in Equation (B.3) between periods t and $t+1$:

$$\widehat{\pi}_t^{BBK}(y) \equiv \sum_i \widehat{\omega}_{i,t}(y) \log\left(\frac{p_{it+1}}{p_{it}}\right) \approx \pi_t^+(y), \tag{B.7}$$

which, as we saw in Equation (18), approximates the true price index as $\widehat{\pi}_t^{BBK}(y) \approx \mathscr{P}_{t,t+1}(\widetilde{\chi}_t(y))$, where $\mathscr{P}_{t,t+1}(\cdot)$ is defined as in Equation (3).

1. *"Iterative" BBK Algorithm.* Start with $\widetilde{\chi}_0(y) = y$. Approximate Equation (B.9) by letting

$$\log \widetilde{\chi}_t(y) \approx \log y - \sum_{\tau=0}^{t-1} \widehat{\pi}_\tau^{BBK}\left(\chi_\tau\left(\widetilde{\chi}_{t-1}(y)\right)\right). \tag{B.8}$$

   In this case, since $\widetilde{\chi}_{t-1}(y)$ known at each period t, there is no need to solve a fixed-point problem.

2. *"Recursive" BBK Algorithm.* Start with $\widetilde{\chi}_0(y) = y$. A first-order approximation of Equation (B.3) is given by

$$\log \widetilde{\chi}_t(y) \approx \log y - \sum_{\tau=0}^{t-1} \widehat{\pi}_\tau^{BBK}(\chi_\tau(\widetilde{\chi}_t(y))). \tag{B.9}$$

   Notice that we need to know $\chi_\tau(\cdot)$ for $\tau < t$, which we can find by inverting $\widetilde{\chi}_\tau(y)$ at each period. Thus, we can solve the fixed-point problem in Equation (B.9) for each y to recover $\widetilde{\chi}_t(y)$.

The Connection between Our Algorithm and the BBK Algorithms Since both our algorithm and the BBK algorithms offer approximate solutions to the same partial differential equation, it should not come as a surprise that they lead to similar results. However, the BBK approach has three distinct features compared with our approach. First, they choose to estimate and approximate the integral equation characterizing the real consumption function, which is the inverse of the mapping from real consumption to expenditure. Second, they choose to estimate Laspeyres geometric function in Equation (B.7) by estimating the Engel curves product by product, as in Equation (B.6). Finally, they specifically choose to limit the domain of this estimated function $\widehat{\pi}_t^{BBK}(\cdot)$ in each period t to the interval $\left[\underline{y}_t, \overline{y}_t\right]$, which is the range of the lowest to the highest levels of nominal expenditure observed in the data in that period. In contrast, we do not impose such a restriction and allow extrapolating the functions outside the range of the observed values of nominal expenditure.

Among the above differences, the last may constitute the most substantive one. In the first period, their method relies on the following equality to recover the estimated value of real consumption: $\log \widetilde{\chi}_1(y) = \log y - \widehat{\pi}_0^{BBK}(\widetilde{\chi}_1(y))$. Based on their restriction, we can solve for the real consumption function $\widetilde{\chi}_1(y)$ only for values of nominal expenditure $y \in \left[\underline{y}_1^*, \overline{y}_1^*\right]$, such that $\underline{y}_1^* \equiv \chi_1\left(\underline{y}_0\right)$ and $\overline{y}_1^* \equiv \chi_1(\overline{y}_0)$. Iterating on this logic using Equation (B.9), we find that we can solve for the real consumption function $\widetilde{\chi}_t(y)$ only for values of nominal expenditure $y \in \left[\underline{y}_t^*, \overline{y}_t^*\right]$, such that

$$\underline{y}_t^* \equiv \max_{\tau<t}\left\{\chi_t\left(\widetilde{\chi}_\tau\left(\underline{y}_\tau\right)\right), \chi_t\left(\underline{y}_\tau^*\right)\right\}, \qquad \overline{y}_t^* \equiv \min_{\tau<t}\left\{\chi_t\left(\widetilde{\chi}_\tau\left(\overline{y}_\tau\right)\right), \chi_t\left(\overline{y}_\tau^*\right)\right\},$$

where we have assumed $\underline{y}_0^* \equiv \underline{y}_0$ and $\overline{y}_0^* \equiv \overline{y}_0$. For instance, in environments featuring growth in real consumption, the BBK algorithm fails to recover real consumption for the richest households due to the restriction that the highest level of real consumption estimated should remain below the highest level of nominal expenditure observed in the initial period, i.e., $\overline{y}_t^* \leq \widetilde{\chi}_t(\overline{y}_0)$.

In contrast, our approach does not suffer from this limitation, since we allow the extrapolation of the mapping from real consumption to expenditure outside the range of observed values, relying on the smoothness of the expenditure function. To see the logic, consider the indirect utility function $V(y;\boldsymbol{p}_t)$, the inverse of the expenditure function, which is smooth whenever the expenditure function is smooth. Therefore, the real consumption function $\widetilde{\chi}_t(\cdot) \equiv E(V(y;\boldsymbol{p}_t);\boldsymbol{p}_b)$ also inherits the smoothness property, implying that the behavior of the function $\widetilde{\chi}_t(y)$ near the upper bound of $\overline{y}_t^*$ can be extrapolated to the neighboring region $y > \overline{y}_t^*$. For this reason, we allow the algorithm to extrapolate the fitted relationship $\widetilde{\chi}_t(\cdot)$ to outside the domain of values $\left[\underline{y}_t^*, \overline{y}_t^*\right]$. As we saw in our simulation exercise Section 2.4, our method indeed recovers accurate estimates of real consumption in an environment featuring sizable growth in real consumption.

We expand on the results of our simulation in Appendix C.1 below, where we compare the results of our algorithm with those of the BBK algorithms across households. In this example, we show how our method effectively yields an extension of the BBK algorithm that allows for extrapolation outside the range they consider, while maintaining low levels of error. Furthermore, in Appendix Figure E.9, we also document that our baseline first-order algorithm and the BBK algorithm yield similar results when applied to the U.S. data.

B.4 Proofs and Derivations

Section B.4.1 presents the proofs of all the results in the main text. Some of these proofs in turn rely on additional lemmas that are presented and proved in Section B.4.2 below.

B.4.1 Proofs of the Main Lemmas and Propositions

Proof of Lemma 1. First, using $\chi_t^b(c) \equiv E\left(M_b^{-1}(c);\boldsymbol{p}_t\right)$, note that

$$\begin{aligned}\frac{\partial \log \chi_t^b(c)}{\partial \log c} &= \frac{\partial \log E(u;\boldsymbol{p}_t)}{\partial \log u}\Big|_{u=M_b^{-1}(c)} \cdot \frac{\partial \log M_b^{-1}(c)}{\partial \log c}, \\ &= \frac{\partial \log E(u;\boldsymbol{p}_t)/\partial \log u}{\partial \log E(u;\boldsymbol{p}_b)/\partial \log u}\Big|_{u=M_b^{-1}(c)},\end{aligned}$$

where in the second equality, we have used the fact that $\frac{\partial \log M_b^{-1}(c)}{\partial \log c} = \frac{1}{\partial \log E(u;\boldsymbol{p}_b)/\partial \log u}\Big|_{u=M_b^{-1}(c)}$. Then, using Proposition 1 and $\Lambda_t^b(c) \equiv \frac{\partial \log \mathscr{P}_{b,t}^b(c)}{\partial \log c} = \frac{\partial \log \chi_t^b(c)}{\partial \log c} - 1$,we have:

$$\begin{aligned}\frac{d\log c_t^{b_2}}{d\log c_t^{b_1}} &= \frac{1+\Lambda_t^{b_1}\left(c_t^{b_1}\right)}{1+\Lambda_t^{b_2}\left(c_t^{b_2}\right)} = \frac{\frac{\partial \log \chi_t^{b_1}(c)}{\partial \log c}\Big|_{c=c_t^{b_1}}}{\frac{\partial \log \chi_t^{b_2}(c)}{\partial \log c}\Big|_{c=c_t^{b_2}}}, \\ &= \frac{\frac{\partial \log E(u;\boldsymbol{p}_t)/\partial \log u}{\partial \log E\left(u;\boldsymbol{p}_{b_1}\right)/\partial \log u}\Big|_{u=u_t}}{\frac{\partial \log E(u;\boldsymbol{p}_t)/\partial \log u}{\partial \log E\left(u;\boldsymbol{p}_{b_2}\right)/\partial \log u}\Big|_{u=u_t}} = \frac{\partial \log E\left(u;\boldsymbol{p}_{b_2}\right)/\partial \log u}{\partial \log E\left(u;\boldsymbol{p}_{b_1}\right)/\partial \log u}\Big|_{u=u_t}, \\ &= \frac{\partial \log \chi_{b_2}^{b_1}(c)}{\partial \log c}\Big|_{c=c_t^{b_1}} = 1+\Lambda_{b_2}^{b_1}\left(c_t^{b_1}\right) = 1+\frac{\partial \log \mathscr{P}_{b_1,b_2}^{b_1}(c)}{\partial \log c}\Big|_{c=c_t^{b_1}}.\end{aligned} \tag{B.10}$$

□

Proof of Lemma 3. From the definition of the true price index in Equation (3), we have $\log \mathscr{P}^{b}_{t,t+1}(c) = \log \chi^{b}_{t+1}(c) - \log \chi^{b}_{t}(c)$. Following Lemma B.2 in Appendix B.4.2 and using a first-order Taylor series expansion of the expenditure function $\chi^{b}_{t+1}(c) \equiv E\left(M_b^{-1}(c); \boldsymbol{p}_{t+1}\right)$ around the vector of prices $\boldsymbol{p}_t$, we find:

$$\log \mathscr{P}^{b}_{t,t+1}(c) = \sum_i \omega_{i,t}\left(\chi^{b}_{t}(c)\right) \log\left(\frac{p_{i,t+1}}{p_{i,t}}\right) + O\left(\Delta_p^2\right), \tag{B.11}$$

where we have used Shephard's lemma in the second step to write the price elasticity of the expenditure function as the expenditure share of the good, i.e.,

$$\frac{\partial \log E\left(M_b^{-1}(c); \boldsymbol{p}_t\right)}{\partial \log p_{i,t}} \equiv \omega_{i,t}\left(\chi^{b}_{t}(c)\right). \tag{B.12}$$

If the preferences are homothetic, we have $s_{it} = \omega_{i,t}\left(\chi^{b}_{t}(c_t)\right) = \omega_{i,t}\left(\chi^{b}_{t}(c)\right)$ for all c and the desired result follows. Otherwise, using Lemma B.2 and performing a first-order Taylor series expansion of the share function, as a function of real consumption c_τ around real consumption c_t, we find:

$$\omega_{i,t}\left(\chi^{b}_{t}(c)\right) = s_{it} + \frac{\partial \omega_{i,t}\left(\chi^{b}_{t}(c)\right)}{\partial \log c} \cdot \log\left(\frac{c}{c_t}\right) + O\left(\left|\log\left(\tfrac{c}{c_t}\right)\right|^2\right),$$

where we have substituted $s_{it} = \omega_{i,t}\left(\chi^{b}_{t}(c_t)\right)$ on the right hand side. Substituting the above equation in Equation (B.12) and using the definition of the geometric index in Equation (7), we find

$$\log \mathscr{P}^{b}_{t,t+1}(c) = \log \mathbb{P}_G\left(\boldsymbol{p}_t, \boldsymbol{s}_t; \boldsymbol{p}_{t+1}, \boldsymbol{s}_{t+1}\right) + \log\left(\frac{c}{c_t}\right) \cdot \sum_{i=1}^{I} \frac{\partial \omega_{i,t}\left(\chi^{b}_{t}(c_t)\right)}{\partial \log c_t} \log\left(\frac{p_{i,t+1}}{p_{i,t}}\right) + O\left(\left|\log\left(\tfrac{c}{c_t}\right)\right|^2\right).$$

If $c = c_t$, then we immediately find the desired result. If $c = c_{t+1}$, we first use Lemma B.3 in Appendix B.4.2 below to let $\left(\log\left(\frac{c_{t+1}}{c_t}\right)\right)^2 = O\left(\Delta^2\right)$, where $\Delta \equiv \max\left\{\Delta_p, \Delta_y\right\}$. Then, since the expenditure function is second order continuously differentiable, we can use the fact that $\frac{\partial \omega_{i,t}\left(\chi^{b}_{t}(c_t)\right)}{\partial \log c_t}$ is bounded, and thus the second term on the right-hand side of the equation above is of the order $O\left(\Delta^2\right)$, which yields Equation (16).

For the second order approximation, we apply the second-order expansion in Lemma B.2 in Appendix B.4.2 to the expenditure function $E\left(M_b^{-1}(c); \boldsymbol{p}\right)$, which yields:

$$\log \mathscr{P}^{b}_{t,t+1}(c) = \log \frac{E\left(M_b^{-1}(c); \boldsymbol{p}_{t+1}\right)}{E\left(M_b^{-1}(c); \boldsymbol{p}_t\right)},$$

$$= \frac{1}{2}\sum_{i=1}^{I}\left[\omega_{i,t+1}\left(\chi_{t+1}^{b}(c)\right)+\omega_{i,t}\left(\chi_{t}^{b}(c)\right)\right]\log\left(\frac{p_{i,t+1}}{p_{i,t}}\right)+O\left(\Delta_{p}^{3}\right), \tag{B.13}$$

where we have again used Equation (B.12). Assuming homotheticity, we have that $s_{it}=\omega_{i,t}\left(\chi_{t}^{b}(c)\right)$ for all c and the desired result follows. Otherwise, using Lemma B.2 in Appendix B.4.2, applied to the Hicksian expenditure share function, we find:

$$\omega_{i,t}\left(\chi_{t}^{b}(c)\right)=s_{it}+\frac{1}{2}\left[\frac{\partial\omega_{i,t}\left(\chi_{t}^{b}(c_t)\right)}{\partial\log c_t}+\frac{\partial\omega_{i,t}\left(\chi_{t}^{b}(c)\right)}{\partial\log c}\right]\cdot\log\left(\frac{c}{c_t}\right)+O\left(\left|\log\left(\tfrac{c}{c_t}\right)\right|^{3}\right),$$
$$\omega_{i,t+1}\left(\chi_{t+1}^{b}(c)\right)=s_{it+1}+\frac{1}{2}\left[\frac{\partial\omega_{i,t+1}\left(\chi_{t+1}^{b}(c_{t+1})\right)}{\partial\log c_{t+1}}+\frac{\partial\omega_{i,t+1}\left(\chi_{t+1}^{b}(c)\right)}{\partial\log c}\right]\cdot\log\left(\frac{c}{c_{t+1}}\right)+O\left(\left|\log\left(\tfrac{c}{c_t}\right)\right|^{3}\right),$$

Substituting this expression in Equation (B.13), using Lemma B.3 in Appendix B.4.2 to write $\left|\log\left(\frac{c}{c_t}\right)\right|=O(\Delta)$, and using the definition of the Törnqvist index in Equation (7), we find:

$$\begin{aligned}\log\mathscr{P}_{t,t+1}^{b}(c)=&\log\mathbb{P}_{T}\left(\boldsymbol{p}_t,\boldsymbol{s}_t;\boldsymbol{p}_{t+1},\boldsymbol{s}_{t+1}\right)\\&+\tfrac{1}{2}\log\left(\frac{c}{c_t}\right)\cdot\sum_{i=1}^{I}\left[\frac{\partial\omega_{i,t}\left(\chi_{t}^{b}(c_t)\right)}{\partial\log c_t}+\frac{\partial\omega_{i,t}\left(\chi_{t}^{b}(c)\right)}{\partial\log c}\right]\log\left(\frac{c}{c_{t_0}}\right)\log\left(\frac{p_{i,t+1}}{p_{i,t}}\right)\\&+\tfrac{1}{2}\log\left(\frac{c}{c_{t+1}}\right)\cdot\sum_{i=1}^{I}\left[\frac{\partial\omega_{i,t+1}\left(\chi_{t+1}^{b}(c_{t+1})\right)}{\partial\log c_{t+1}}+\frac{\partial\omega_{i,t+1}\left(\chi_{t+1}^{b}(c)\right)}{\partial\log c}\right]\log\left(\frac{c}{c_t}\right)\log\left(\frac{p_{i,t+1}}{p_{i,t}}\right)+O\left(\Delta^{3}\right),\end{aligned} \tag{B.14}$$

where $\Delta\equiv\max\left\{\Delta_p,\Delta_y\right\}$. Now, we use the third-order continuously differentiable property of the expenditure function to find

$$\frac{\partial\omega_{i,t''}\left(\chi_{t''}^{b}(c_{t'})\right)}{\partial\log c_{t'}}=\frac{\partial\omega_{i,t}\left(\chi_{t}^{b}(c_t)\right)}{\partial\log c_t}+O(\Delta),\qquad t',t''\in\{t,t+1\},$$

which we use to substitute for the expressions within the square brackets in Equation (B.14). This leads to the following result:

$$\begin{aligned}\log\mathscr{P}_{t,t+1}^{b}(c)=&\log\mathbb{P}_{T}\left(\boldsymbol{p}_t,\boldsymbol{s}_t;\boldsymbol{p}_{t+1},\boldsymbol{s}_{t+1}\right)\\&+\log\left[\frac{(c)^2}{c_t\cdot c_{t+1}}\right]\cdot\sum_{i=1}^{I}\frac{\partial\omega_{i,t}\left(\chi_{t}^{b}(c_t)\right)}{\partial\log c_t}\log\left(\frac{p_{i,t+1}}{p_{i,t}}\right)+O\left(\Delta^{3}\right).\end{aligned}$$

Thus, letting $c=\sqrt{c_t\cdot c_{t+1}}$, the second term on the right hand side vanishes and we obtain Equation (17), as desired. □

Proof of Lemma 4. The proof follows the steps outlined in proving Equations (B.11) and (B.13) in the proof of Lemma 3 above. □

Proof of Proposition 2. First, note that since the expenditure function is second-order continuously differentiable, using Lemma B.2 in Appendix B.4.2 for function $\omega_i\left(\chi_t^b(c);\boldsymbol{p}\right)$ around $(\boldsymbol{p}_t, c_t)'$, we have

$$\log\left(\frac{s_{it+1}}{s_{it}}\right)=\sum_i \frac{\partial \log \omega_i\left(\chi_t^b(c_t);\boldsymbol{p}_t\right)}{\partial \log p_{it}}\log\left(\frac{p_{it+1}}{p_{it}}\right)+\frac{\partial \log \omega_i\left(\chi_t^b(c_t);\boldsymbol{p}_t\right)}{\partial \log c_t}\log\left(\frac{c_{t+1}}{c_t}\right)+O\left(\Delta^2\right), \tag{B.15}$$

where $\Delta \equiv \Delta_p$ if preferences are homothetic, and $\Delta \equiv \left\{\Delta_p, \Delta_y\right\}$. Using the second order continuously differentiable property of the expenditure function, we conclude that $\log\left(\frac{s_{it+1}}{s_{it}}\right)=O(\Delta)$.

For the Laspeyres price index formula, we have:

$$\begin{aligned}
\log P_L &= \log\left(\sum_i s_{it}\frac{p_{it+1}}{p_{it}}\right),\\
&= \log\left(1+\sum_i s_{it}\log\left(\frac{p_{it+1}}{p_{it}}\right)+\frac{1}{2}\sum_i s_{it}\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2+O\left(\Delta^3\right)\right),\\
&= \sum_i s_{it}\log\left(\frac{p_{it+1}}{p_{it}}\right)+\frac{1}{2}\sum_i s_{it}\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2-\frac{1}{2}\left(\sum_i s_{it}\log\left(\frac{p_{it+1}}{p_{it}}\right)+\frac{1}{2}\sum_i s_{it}\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2\right)^2+O\left(\Delta^2\right),\\
&= \log P_G+O\left(\Delta^2\right),
\end{aligned}$$

where in the second equality we use the Taylor series expansion of $\exp(x)$ for $x \equiv \log\left(\frac{p_{it+1}}{p_{it}}\right)$, and in the third equality we use the Taylor series expansion of $\log(1+x)$ for $x \equiv \sum_i s_{it}\log\left(\frac{p_{it+1}}{p_{it}}\right)+\frac{1}{2}\sum_i s_{it}\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2+O\left(\Delta^3\right)$.

For the Paasche price index formula, we find:

$$\begin{aligned}
\log P_P &= -\log\left(\sum_i s_{it+1}\frac{p_{it}}{p_{it+1}}\right),\\
&= -\log\left(1-\sum_i s_{it+1}\log\left(\frac{p_{it+1}}{p_{it}}\right)+\frac{1}{2}\sum_i s_{it+1}\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2+O\left(\Delta^3\right)\right),\\
&= \sum_i s_{it+1}\log\left(\frac{p_{it+1}}{p_{it}}\right)-\frac{1}{2}\sum_i s_{it}\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2+\frac{1}{2}\left(\sum_i s_{it}\log\left(\frac{p_{it+1}}{p_{it}}\right)+\frac{1}{2}\sum_i s_{it}\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2\right)^2+O\left(\Delta^2\right),\\
&= \sum_i s_{it}\times\frac{s_{it+1}}{s_{it}}\times\log\left(\frac{p_{it+1}}{p_{it}}\right)+O\left(\Delta^2\right),\\
&= \log P_G+O\left(\Delta^2\right),
\end{aligned}$$

where in the second equality we use the Taylor series expansion of $\exp(-x)$ for $x \equiv \log\left(\frac{p_{it+1}}{p_{it}}\right)$, and in the third equality we used the Taylor series expansion of $\log(1-x)$ for $x \equiv \sum_i s_{it+1} \log\left(\frac{p_{it+1}}{p_{it}}\right) - \frac{1}{2}\sum_i s_{it+1}\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2 + O\left(\Delta^3\right)$. In the last equality, we use the fact that $\log\left(\frac{s_{it+1}}{s_{it}}\right) = O(\Delta)$ from Equation (B.15) above.

For the Fisher price index formula, we repeat the same steps used in the arguments above for the Laspeyres and Paasche indices:

$$\begin{aligned}
\log \mathbb{P}_F &= \frac{1}{2}\log \mathbb{P}_L + \frac{1}{2}\log \mathbb{P}_P, \\
&= \frac{1}{2}\log\left(s_{it}\left(\frac{p_{it+1}}{p_{it}}\right)\right) - \frac{1}{2}\log\left(s_{it+1}\left(\frac{p_{it}}{p_{it+1}}\right)\right), \\
&= \frac{1}{2}\log\left(1 + \sum_i s_{it}\log\left(\frac{p_{it+1}}{p_{it}}\right) + \frac{1}{2}\sum_i s_{it}\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2\right) \\
&\quad - \frac{1}{2}\log\left(1 - \sum_i s_{it+1}\log\left(\frac{p_{it+1}}{p_{it}}\right) + \frac{1}{2}\sum_i s_{it+1}\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2\right) + O\left(\Delta^3\right), \\
&= \frac{1}{2}\sum_i \left(s_{it} + s_{it+1}\right)\log\left(\frac{p_{it+1}}{p_{it}}\right) + \frac{1}{4}\sum_i s_{it}\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2 - \frac{1}{4}\sum_i s_{it+1}\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2 \\
&\quad + \frac{1}{4}\left(\sum_i s_{it+1}\log\left(\frac{p_{it+1}}{p_{it}}\right) - \frac{1}{2}\sum_i s_{it+1}\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2\right)^2 \\
&\quad - \frac{1}{4}\left(\sum_i s_{it}\log\left(\frac{p_{it+1}}{p_{it}}\right) + \frac{1}{2}\sum_i s_{it}\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2\right)^2 + O\left(\Delta^3\right).
\end{aligned} \tag{B.16}$$

To simplify this expression further, note that using Equation (B.15), we have

$$\begin{aligned}
\tfrac{s_{it+1}-s_{it}}{s_{it}} &= \exp\left(\log\left(\tfrac{s_{it+1}}{s_{it}}\right)\right) - 1 = O(\Delta), \\
\sum_i \bar{s}_{it}\left(\log\left(\tfrac{p_{it+1}}{p_{it}}\right)\right)^m &= O\left(\Delta^m\right), \qquad 1 \leq m,
\end{aligned}$$

where we have let $\bar{s}_{it} \equiv \frac{1}{2}\left(s_{it} + s_{it+1}\right)$. Using this result, we can rewrite Equation (B.16) as

$$\begin{aligned}
\log \mathbb{P}_F &= \log \mathbb{P}_T - \frac{1}{4}\sum_i \left(s_{it+1} - s_{it}\right)\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2 + O\left(\Delta^3\right) \\
&\quad + \frac{1}{4}\left[\sum_i s_{it+1}\log\left(\frac{p_{it+1}}{p_{it}}\right) - \sum_i s_{it}\log\left(\frac{p_{it+1}}{p_{it}}\right) - \frac{1}{2}\sum_i s_{it+1}\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2 - \frac{1}{2}\sum_i s_{it}\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2\right] \\
&\quad \times \left[\sum_i s_{it+1}\log\left(\frac{p_{it+1}}{p_{it}}\right) - \frac{1}{2}\sum_i s_{it+1}\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2 + \sum_i s_{it}\log\left(\frac{p_{it+1}}{p_{it}}\right) + \frac{1}{2}\sum_i s_{it}\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2\right], \\
&= \log \mathbb{P}_T - \frac{1}{4}\sum_i s_{it}\left(\tfrac{s_{it+1}-s_{it}}{s_{it}}\right)\left(\log\left(\frac{p_{it+1}}{p_{it}}\right)\right)^2
\end{aligned}$$

$$
\begin{aligned}
&+\frac{1}{2}\left[\sum_i s_{it}\left(\tfrac{s_{it+1}-s_{it}}{s_{it}}\right)\log\left(\tfrac{p_{it+1}}{p_{it}}\right)-\sum_i \bar{s}_{it}\left(\log\left(\tfrac{p_{it+1}}{p_{it}}\right)\right)^2\right]\\
&\times\left[\sum_i \bar{s}_{it}\log\left(\tfrac{p_{it+1}}{p_{it}}\right)-\frac{1}{4}\sum_i s_{it}\left(\tfrac{s_{it+1}-s_{it}}{s_{it}}\right)\left(\log\left(\tfrac{p_{it+1}}{p_{it}}\right)\right)^2\right]+O\left(\Delta^3\right),\\
&=\log\mathbb{P}_T+O\left(\Delta^3\right).
\end{aligned}
$$

Finally, for Sato-Vartia, we begin with the following approximation:

$$
\begin{aligned}
\frac{s_{it+1}-s_{it}}{\log\left(\frac{s_{it+1}}{s_{it}}\right)}&=\frac{\frac{1}{2}s_{it}\left(\frac{s_{it+1}}{s_{it}}-1\right)+\frac{1}{2}s_{it+1}\left(1-\frac{s_{it}}{s_{it+1}}\right)}{\log\left(\frac{s_{it+1}}{s_{it}}\right)},\\
&=\frac{1}{2}s_{it}\left(1+\frac{1}{2}\log\left(\tfrac{s_{it+1}}{s_{it}}\right)+\frac{1}{6}\left(\log\left(\tfrac{s_{it+1}}{s_{it}}\right)\right)^2\right)\\
&\quad+\frac{1}{2}s_{it+1}\left(1-\frac{1}{2}\log\left(\tfrac{s_{it+1}}{s_{it}}\right)+\frac{1}{6}\left(\log\left(\tfrac{s_{it+1}}{s_{it}}\right)\right)^2\right)+O\left(\Delta^3\right),\\
&=\bar{s}_{it}\left(1+\tfrac{1}{6}\left(\log\left(\tfrac{s_{it+1}}{s_{it}}\right)\right)^2\right)-\frac{1}{4}\left(s_{it+1}-s_{it}\right)\log\left(\tfrac{s_{it+1}}{s_{it}}\right),\\
&=\bar{s}_{it}\left(1+\tfrac{1}{6}\left(\log\left(\tfrac{s_{it+1}}{s_{it}}\right)\right)^2\right)-\frac{s_{it}}{4}\left(\tfrac{s_{it+1}}{s_{it}}-1\right)\log\left(\tfrac{s_{it+1}}{s_{it}}\right)+O\left(\Delta^3\right),\\
&=\bar{s}_{it}\left(1+\tfrac{1}{6}\left(\log\left(\tfrac{s_{it+1}}{s_{it}}\right)\right)^2\right)-\frac{s_{it}}{4}\left(\log\left(\tfrac{s_{it+1}}{s_{it}}\right)\right)^2+O\left(\Delta^3\right),
\end{aligned}
$$

where in the second equality, we use a Taylor expansion of $\exp(x)-1$ for $x\equiv\log\left(\frac{s_{it+1}}{s_{it}}\right)$ to simplify $\frac{s_{it+1}}{s_{it}}-1$, and a Taylor expansion of $1-\exp(-x)$ for $x\equiv\log\left(\frac{s_{it+1}}{s_{it}}\right)$ to simplify $1-\frac{s_{i1}}{s_{it+1}}$. We use the former approximation again in the fourth equality, as well as Equation (B.15). Substituting this result in the definition of the Sato-Vartia price index formula, we find

$$
\begin{aligned}
\log\mathbb{P}_S&\equiv\frac{\sum_i\left(\frac{s_{it+1}-s_{it}}{\log\left(\frac{s_{it+1}}{s_{it}}\right)}\right)\log\left(\frac{p_{it+1}}{p_{it}}\right)}{\sum_j\left(\frac{s_{j,t+1}-s_{jt}}{\log\left(\frac{s_{j,t+1}}{s_{jt}}\right)}\right)},\\
&=\frac{\log\mathbb{P}_T+O\left(\Delta^3\right)}{1+O\left(\Delta^2\right)},\\
&=\log\mathbb{P}_T+O\left(\Delta^3\right),
\end{aligned}
$$

where we use the fact that $\sum_i\bar{s}_{it}=1$. □

Proof of Proposition 3. We can write the growth in consumer expenditures as

$$\frac{d\log E\left(M_b^{-1}(c_t);\boldsymbol{p}_t,\boldsymbol{x}_t\right)}{dt}=\sum_i\frac{\partial\log E\left(M_b^{-1}(c_t);\boldsymbol{p}_t,\boldsymbol{x}_t\right)}{\partial\log p_{it}}\frac{d\log p_{it}}{dt}+\sum_d\frac{\partial\log E\left(M_b^{-1}(c_t);\boldsymbol{p}_t,\boldsymbol{x}_t\right)}{\partial\log x_{dt}}\frac{d\log x_{dt}}{dt}$$
$$+\frac{\partial\log E\left(M_b^{-1}(c_t);\boldsymbol{p}_t,\boldsymbol{x}_t\right)}{\partial\log c_t}\frac{d\log c_t}{dt},$$

where the left hand side equals $\frac{d\log y_{it}}{dt}$ and where we omit the base period superscripts b to simplify the expression. The desired result follows from the observation that $c=E\left(M_b^{-1}(c);\boldsymbol{p}_b,\boldsymbol{x}\right)$ for all $\boldsymbol{x}$. □

B.4.2 Additional Lemmas

In this section, we derive the additional lemmas and propositions used in some of the steps of the main proofs in Section B.4.1.

Lemma B.1. *Assume that the expenditure function $E(\cdot;\cdot)$ is third-order continuously differentiable. Then the growth in real consumption between periods t and $t+1$ satisfies*[A3]

$$\log\left(\frac{c_{t+1}^b}{c_t^b}\right)=\frac{1}{1+\Lambda_{t+1}^{b,(1)}}\log\left(\frac{y_{t+1}/y_t}{\mathbb{P}_G\left(\boldsymbol{p}_t,\boldsymbol{s}_t;\boldsymbol{p}_{t+1},\boldsymbol{s}_{t+1}\right)}\right)+O\left(\Delta^2\right),\tag{B.17}$$
$$=\frac{1}{1+\Lambda_{t+1}^{b,(2)}}\log\left(\frac{y_{t+1}/y_t}{\mathbb{P}_T\left(\boldsymbol{p}_t,\boldsymbol{s}_t;\boldsymbol{p}_{t+1},\boldsymbol{s}_{t+1}\right)}\right)+O\left(\Delta^3\right),\tag{B.18}$$

where the first-order $\Lambda_{t+1}^{b,(1)}$ and second-order $\Lambda_{t+1}^{b,(2)}$ nonhomotheticity correction are defined by

$$1+\Lambda_{t+1}^{b,(1)}\equiv\frac{\partial\log\chi_t^b\left(c_t^b\right)}{\partial\log c_t^b}=\left(\frac{\partial\log\tilde{\chi}_t^b\left(y_t\right)}{\partial\log y_t}\right)^{-1},\tag{B.19}$$

$$1+\Lambda_{t+1}^{b,(2)}\equiv\frac{1}{2}\left[\frac{\partial\log\chi_t^b\left(c_t^b\right)}{\partial\log c_t^b}+\frac{\partial\log\chi_{t+1}^b\left(c_{t+1}^b\right)}{\partial\log c_{t+1}^b}\right]=\frac{1}{2}\left[\left(\frac{\partial\log\tilde{\chi}_t^b\left(y_t\right)}{\partial\log y_t}\right)^{-1}+\left(\frac{\partial\log\tilde{\chi}_{t+1}^b\left(y_{t+1}\right)}{\partial\log y_{t+1}}\right)^{-1}\right],\tag{B.20}$$

[A3] We can also show the following second-order approximation holds:

$$\log\left(\frac{c_{t+1}^b}{c_t^b}\right)=\frac{1}{2}\left[\frac{\partial\log\tilde{\chi}_t^b\left(y_t\right)}{\partial\log y_t}+\frac{\partial\log\tilde{\chi}_{t+1}^b\left(y_{t+1}\right)}{\partial\log y_{t+1}}\right]\log\left(\frac{y_{t+1}}{y_t}\right)$$
$$-\frac{1}{2}\left(\frac{\partial\log\tilde{\chi}_t^b\left(y_t\right)}{\partial\log y_t}\right)\log\mathbb{P}_G\left(\boldsymbol{p}_t,\boldsymbol{s}_t;\boldsymbol{p}_{t+1},\boldsymbol{s}_{t+1}\right)+\frac{1}{2}\left(\frac{\partial\log\tilde{\chi}_{t+1}^b\left(y_{t+1}\right)}{\partial\log y_{t+1}}\right)\log\mathbb{P}_G\left(\boldsymbol{p}_{t+1},\boldsymbol{s}_{t+1};\boldsymbol{p}_t,\boldsymbol{s}_t\right)+O\left(\Delta^3\right).$$

Using this result leads to another alternative to the second-order Algorithm A.5.

and where $\Delta \equiv \max\left\{\Delta_p, \Delta_y\right\}$*, with* Δ_y *and* Δ_p *defined as in Equation (15).*

Proof. Note that we have

$$\log\left(\frac{y_{t+1}}{y_t}\right) = \log\left(\frac{E\left(M_b^{-1}\left(c_{t+1}^b\right);\boldsymbol{p}_{t+1}\right)}{E\left(M_b^{-1}\left(c_t^b\right);\boldsymbol{p}_t\right)}\right). \tag{B.21}$$

Using a first-order Taylor expansion of the left-hand-side of the equation above in terms of c_t^b, as well as Lemma B.3, we obtain:

$$\log\left(\frac{y_{t+1}}{y_t}\right) = \log\frac{E\left(M_b^{-1}\left(c_1^b\right);\boldsymbol{p}_{t+1}\right)}{E\left(M_b^{-1}\left(c_t^b\right);\boldsymbol{p}_t\right)} + \left.\frac{\partial \log E\left(M_b^{-1}\left(c^b\right);\boldsymbol{p}_t\right)}{\partial \log c}\right|_{c\equiv c_t^b} \cdot \log\left(\frac{c_{t+1}^b}{c_t^b}\right) + O\left(\Delta^2\right),$$

which gives us Equation (B.17) for the first equality in Equation (B.19), since $E\left(M_b^{-1}(c);\boldsymbol{p}_t\right) = \chi_t^b(c)$. Since $\widetilde{\chi}_t\left(\chi_t(c)\right) = c$, the second equality in that equation then follows.

Next, in Equation (B.21), we use Lemma B.2 for a vector of variables $(\boldsymbol{p}, c)'$ to find:

$$\begin{aligned}
\log\left(\frac{y_{t+1}^n}{y_t^n}\right) &= \frac{1}{2}\sum_{i=1}^{I}\left[\left.\frac{\partial \log E\left(M_b^{-1}(c);\boldsymbol{p}\right)}{\partial \log p_i}\right|_{(c;\boldsymbol{p})\equiv(c_t;\boldsymbol{p}_t)} + \left.\frac{\partial \log E\left(M_b^{-1}(c);\boldsymbol{p}\right)}{\partial \log p_i}\right|_{(c;\boldsymbol{p})\equiv\left(c_{t+1};\boldsymbol{p}_{t+1}\right)}\right]\cdot\log\left(\frac{p_{i,t+1}}{p_{i,t}}\right) \\
&\quad + \frac{1}{2}\left[\left.\frac{\partial \log E\left(M_b^{-1}(c);\boldsymbol{p}\right)}{\partial \log c}\right|_{(c;\boldsymbol{p})\equiv(c_t;\boldsymbol{p}_t)} + \left.\frac{\partial \log E\left(M_b^{-1}(c);\boldsymbol{p}\right)}{\partial \log c}\right|_{(c;\boldsymbol{p})\equiv\left(c_{t+1};\boldsymbol{p}_{t+1}\right)}\right]\cdot\log\left(\frac{c_{t+1}}{c_t}\right) \\
&\quad + O\left(\Delta^3\right), \\
&= \frac{1}{2}\sum_{i=1}^{I}\left[\omega_{i,t}\left(\chi_t^b(c_t)\right) + \omega_{i,t+1}\left(\chi_{t+1}^b\left(c_{t+1}\right)\right)\right]\log\left(\frac{p_{i,t+1}}{p_{i,t}}\right) \\
&\quad + \left(\frac{\partial \log \chi_t^b\left(c_t^b\right)}{\partial \log c_t^b} + \frac{\partial \log \chi_{t+1}^b\left(c_{t+1}^b\right)}{\partial \log c_{t+1}^b}\right)\cdot\log\left(\frac{c_t}{c_{t_0}}\right) + O\left(\Delta^3\right),
\end{aligned}$$

where in the second equality we use Shephard's lemma, as well as the definition of the mapping $\chi_t(c) \equiv E\left(M_b^{-1}(c);\boldsymbol{p}\right)$. Noting again that $\widetilde{\chi}_t\left(\chi_t(c)\right) = c$, Equation (B.18) follows. □

Lemma B.2. *Consider a function* $f(\boldsymbol{x})$ *defined in the space of* $\boldsymbol{x} \in \mathbb{R}^I$*. If* $f(\cdot)$ *is second order continuously differentiable, we have:*

$$f(\boldsymbol{y}) - f(\boldsymbol{x}) = \sum_{i=1}^{I}\frac{\partial f(\boldsymbol{x})}{\partial x_i}(y_i - x_i) + O\left(\delta^2\right), \tag{B.22}$$

and if it is third order continuously differentiable, we have:

$$f(\boldsymbol{y})-f(\boldsymbol{x})=\frac{1}{2}\sum_{i=1}^{I}\left[\frac{\partial f(\boldsymbol{y})}{\partial y_i}+\frac{\partial f(\boldsymbol{x})}{\partial x_i}\right](y_i-x_i)+O\left(\delta^3\right), \tag{B.23}$$

where we have defined $\delta \equiv \max_i |y_i - x_i|$.

Proof. Using Taylor's series expansion, we have

$$f(\boldsymbol{y})=f(\boldsymbol{x})+\sum_{i=1}^{I}\frac{\partial f(\boldsymbol{x})}{\partial x_i}(y_i-x_i)+\sum_{i,j}R_2(\boldsymbol{y})(y_i-x_i)\left(y_j-x_j\right).$$

From Taylor's theorem, we have the bound$|R_2(\boldsymbol{y})| \leq \frac{1}{2}B_2$ where B_2 is the upper bound of the value of the second order derivatives within the ball of radius $|\boldsymbol{y}-\boldsymbol{x}|$ around $\boldsymbol{x}$. This implies that the absolute value of the residual can be bounded above by $\frac{I^2}{2}B_2\delta^2 = O\left(\delta^2\right)$, which leads to Equation (B.22).

Following similar steps, we can show that if function f is third order continuously differentiable, we have:

$$\begin{aligned}
f(\boldsymbol{y})&=f(\boldsymbol{x})+\sum_{i=1}^{I}\frac{\partial f(\boldsymbol{x})}{\partial x_i}(y_i-x_i)+\frac{1}{2}\sum_{i,j=1}^{I}\frac{\partial^2 f(\boldsymbol{x})}{\partial x_i\partial x_j}(y_i-x_i)\left(y_j-x_j\right)+O\left(\delta^3\right),\\
f(\boldsymbol{x})&=f(\boldsymbol{y})+\sum_{i=1}^{I}\frac{\partial f(\boldsymbol{y})}{\partial x_i}(x_i-y_i)+\frac{1}{2}\sum_{i,j=1}^{I}\frac{\partial^2 f(\boldsymbol{y})}{\partial x_i\partial x_j}(y_i-x_i)\left(y_j-x_j\right)+O\left(\delta^3\right).
\end{aligned}$$

Together, the two equations imply:

$$\begin{aligned}
f(\boldsymbol{y})-f(\boldsymbol{x})=&\frac{1}{2}\sum_{i=1}^{I}\left[\frac{\partial f(\boldsymbol{y})}{\partial y_i}+\frac{\partial f(\boldsymbol{x})}{\partial x_i}\right](y_i-x_i)\\
&+\frac{1}{4}\sum_{i,j=1}^{I}\left[\frac{\partial^2 f(\boldsymbol{x})}{\partial x_i\partial x_j}-\frac{\partial^2 f(\boldsymbol{y})}{\partial x_i\partial x_j}\right](y_i-x_i)\left(y_j-x_j\right)+O\left(\delta^3\right).
\end{aligned}$$

This gives us the desired result in Equation (B.23), since:

$$\frac{\partial^2 f(\boldsymbol{y})}{\partial x_i\partial x_j}-\frac{\partial^2 f(\boldsymbol{x})}{\partial x_i\partial x_j}=\sum_k\frac{\partial^3 f(\boldsymbol{x})}{\partial x_k\partial x_i\partial x_j}(y_k-x_k).$$

□

Lemma B.3. *Assume* $\Delta_y \equiv \left|\log\left(y_{t+1}/y_t\right)\right|$ *and* $\Delta_p \equiv \max_i \left|\log\left(p_{i,t+1}/p_{i,t}\right)\right|$. *If the expenditure function is second order continuously differentiable, we have* $\left|\log\left(c^b_{t+1}/c^b_t\right)\right| = O(\Delta)$ *where* $\Delta \equiv$

$\max\left\{\Delta_p, \Delta_y\right\}$.

Proof. We use a Taylor expansion of the expenditure function $E\left(M_b^{-1}(c);\boldsymbol{p}\right)$ as a function of real consumption c around $c = c_{t_0}^b$ and $\boldsymbol{p}_{t_0}$, noting that the expenditure function $E\left(M_b^{-1}(c);\boldsymbol{p}\right)$ is continuously differentiable in all its arguments, to obtain:

$$\begin{aligned}\log\left(\frac{y_{t+1}}{y_t}\right) &= \log\left(\frac{E\left(M_b^{-1}\left(c_{t+1}^b\right);\boldsymbol{p}_{t+1}\right)}{E\left(M_b^{-1}\left(c_t^b\right);\boldsymbol{p}_t\right)}\right),\\ &= \sum_i h_i^p\left(c_{t+1}^b,\boldsymbol{p}_{t+1}\right)\log\left(\frac{p_{i,t+1}}{p_{i,t}}\right) + h^c\left(c_{t+1}^b,\boldsymbol{p}_{t+1}\right)\log\left(\frac{c_{t+1}^b}{c_t^b}\right),\end{aligned}$$

where the values of $h_i^p\left(c_{t+1}^b,\boldsymbol{p}_{t+1}\right)$ and $h^c\left(c_{t+1}^b,\boldsymbol{p}_{t+1}\right)$ are bounded by the maximum value of the gradient of $E\left(M_b^{-1}(c);\boldsymbol{p}\right)$ in the ball around $\left(\boldsymbol{p}_t, c_t^b\right)$ with radius $\left|\left(\boldsymbol{p}_{t+1}, c_{t+1}^b\right)' - \left(\boldsymbol{p}_t, c_t^b\right)'\right|$. We thus have:

$$\left|\log\left(\frac{c_{t+1}^b}{c_t^b}\right)\right| < M\,\Delta,$$

for some $M > 0$.

□

C Additional Simulation Results

In this appendix, we report additional results from our illustrative simulation exercise in Section 2.4 in the main text. We first show how the mapping between real consumption and total expenditure $\chi_t^b(\cdot)$ changes over time depending on the covariance between income elasticities and inflation across products. We then use the simulation to assess the accuracy of our algorithm in estimating changes in real consumption over time.

C.1 Mapping between Real Consumption and Expenditure

We document how the mapping between real consumption and expenditure $\chi_t^b(\cdot)$, defined in Equation (2), evolves over time, depending on the sign of the covariance between income elasticity and inflation. We first consider the case with a positive covariance, which is illustrated in Figures C.1a and C.1b. These figures compare the mapping in terms of real consumption between the nonhomothetic and homothetic specifications, with the initial ($b = 0$) and the last ($b = 70$) periods as the base, respectively.[A4] The figures depict how the expenditure functions change over time in each case.

[A4] Specifically, we compare the nonhomothetic specification against a homothetic CES specification with $(\sigma, \varepsilon_a, \varepsilon_m, \varepsilon_s) = (0.26, 1, 1, 1)$.

Figure C.1: nhCES Example: Mapping from Real Consumption to Expenditure $\chi_t^b(\cdot)$

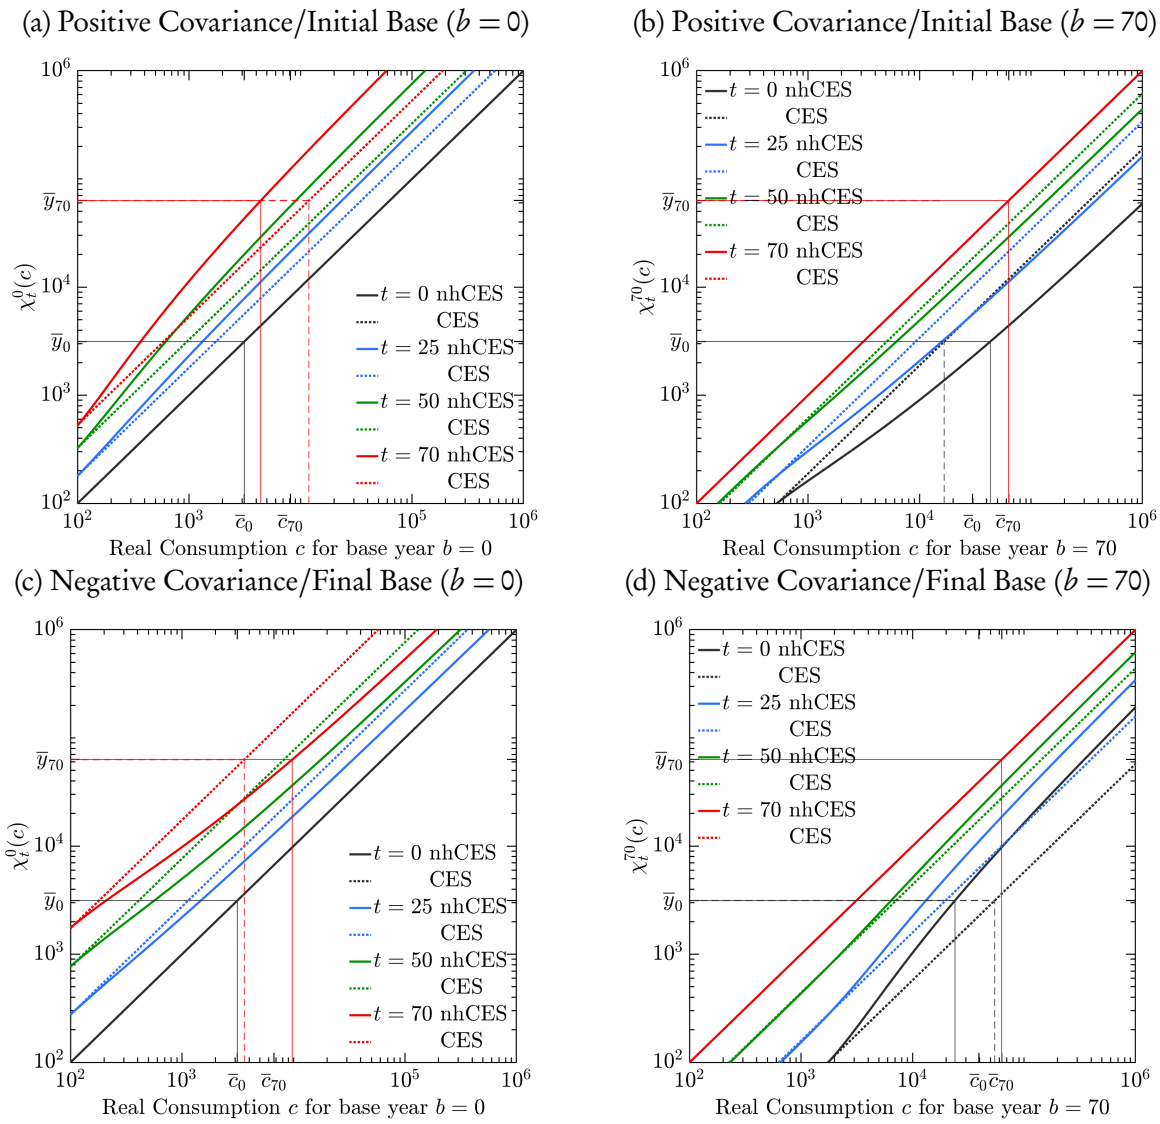

Note: The figure shows the change over time in the mapping between real consumption and expenditure, for the preferences defined in Equation (29) with parameters corresponding to a nonhomothetic CES $(\sigma, \varepsilon_a, \varepsilon_m, \varepsilon_s) = (0.26, 0.2, 1, 1.65)$ (nhCES) and homothetic CES $(\sigma, \varepsilon_a, \varepsilon_m, \varepsilon_s) = (0.26, 1, 1, 1)$ functions. Panels (a) and (b) show the results for initial and final periods as the base for the case with positive income elasticity-inflation covariance, respectively. Panels (c) and (d) show the same results for the case with negative income elasticity-inflation covariance.

In the homothetic case, the expenditure function always has a linear form. Due to the overall inflation in prices, the expenditure function uniformly shifts upward over time for the homothetic CES preferences.

In the nonhomothetic case, let us first consider the initial period as the base in Figure C.1a. By definition, the mapping begins as the identity function in the initial base period. As time passes, the costs of achieving higher levels of real consumption rises faster, since to achieve these higher levels households need to shift their consumption toward goods that have featured higher cumu-

lative inflation between the base and current periods. Thus, the mapping $\chi_t^b(\cdot)$, which characterizers the expenditure function in terms of real consumption, increasingly deviates from linearity and becomes more convex as time passes. The figure shows that, compared to the homothetic case, the upward shift in the expenditure function is larger for higher levels of real consumption.

Next, consider the final period as base as in Figure C.1b. By definition, in this case the mapping is the identity function in the final period. As we move backward in time, the costs of achieving higher levels of real consumption fall faster, since to achieve these levels of welfare households shift their consumption toward necessities, whose prices were relatively higher in the past. Thus, the mapping increasingly deviates from linearity and becomes more concave as we move toward the initial period. The simulation thus illustrates that, regardless of the choice of the base period, the expenditure function is more convex in later periods under nonhomothetic preferences with a positive income elasticity-inflation covariance.

Figures C.1c and C.1d examine the same patterns in the case with a negative covariance between price inflation and income elasticities across goods. In this case, the mapping becomes more concave over time, since now consumers shift the composition of their expenditures toward goods that have lower inflation. With the initial period as base, the mapping begins with a linear form and becomes more concave as we move forward in time. With the final period as base, the mapping ends with the identity function in the last period and becomes more convex as we move backward in time.

C.2 Accuracy of the Approximation for Different Algorithms

Figures C.2a-C.2d document the average error in the measures of real consumption approximated across households over time by each of the different algorithms presented in Section A with that found based on the conventional uncorrected approach. As previously, the results are reported for different base periods and income elasticity-inflation covariance combinations. In all cases, we have used third-order polynomials for the nonparametric approximations. As in the main text, we use the underlying preference parameters to compute the correct value of real consumption $c_t^{b,n}$ for each household n at each point in time t, and compare that value with the approximate value $\widehat{c}_t^{b,n}$ found with each approach.

As we saw before, the conventional approach leads to sizable levels of error, whereas all our algorithms achieve substantially more precise approximations. The three first-order methods typically lead to similar errors, with the benchmark and refined algorithms being very close and the algorithm based on the real consumption function leading to slightly higher (lower) error in the positive (negative) covariance case. As expected, the second-order algorithms achieve lower levels of error. The refined algorithm and that based on real consumption function lead to par-

Figure C.2: nhCES Example: Nonparametric Approximation of Real Consumption

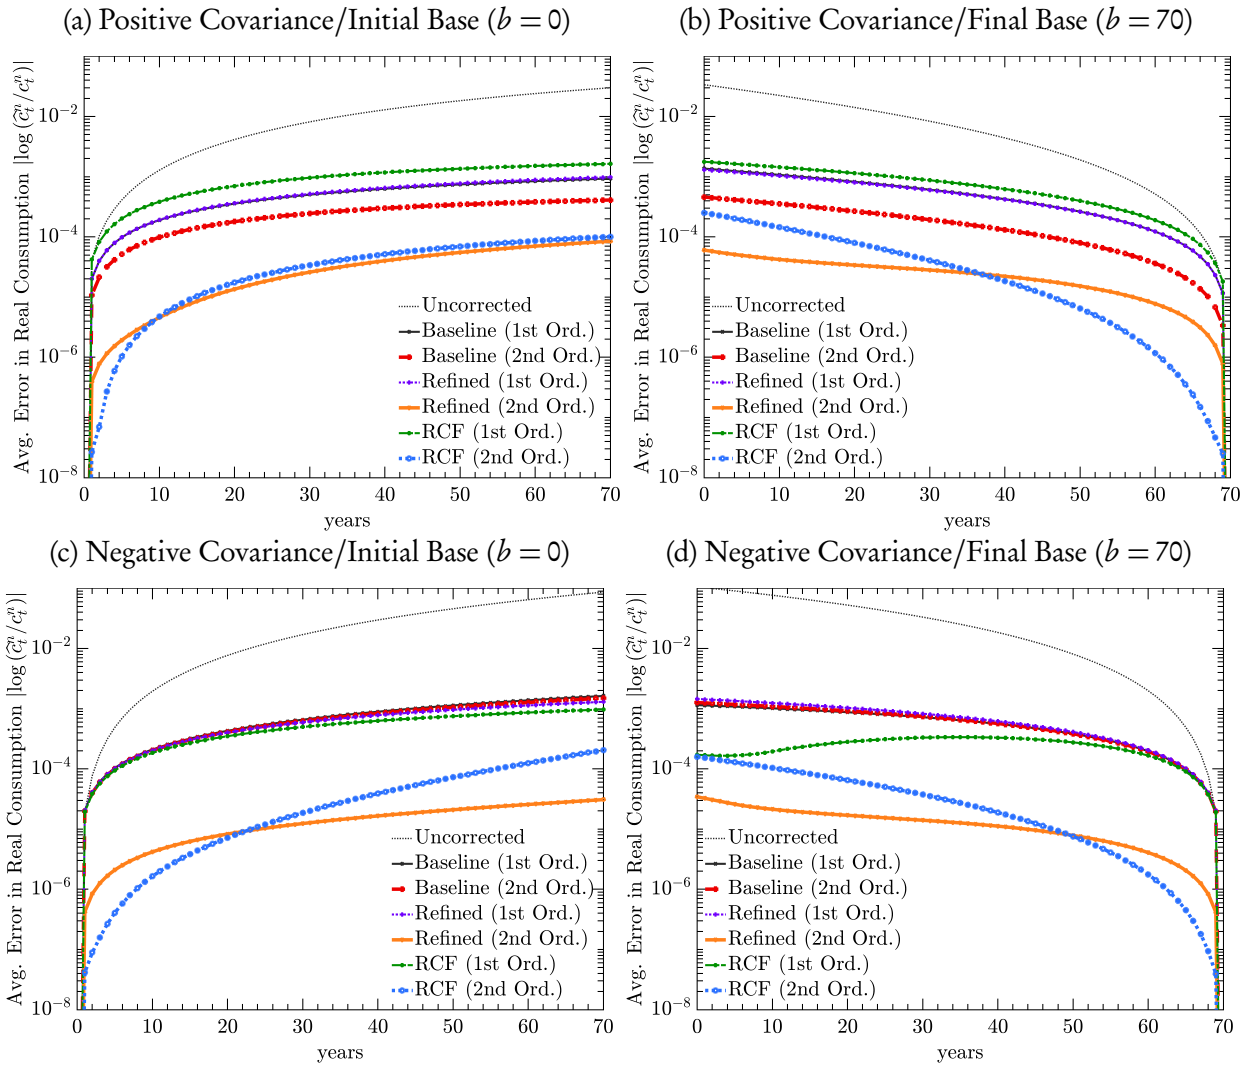

Note: The figures compare the average error in the approximated value of real consumption across households across different algorithms presented in Section A. The correct value of real consumption is calculated based on the underlying parameters of the nhCES preferences. The panels show the error for the choices of base period (a) $b = 0$ and (b) $b = 70$ with the positive income elasticity-inflation covariance and (c) $b = 0$ and (d) $b = 70$ with the negative covariance.

ticularly improved approximations, with the former outperforming the latter over longer time horizons.

Comparison with Baqaee et al. (2022) Figures C.3a-C.3b compare the sizes of the approximation error between our three first-order algorithms and the two algorithms suggested by Baqaee et al. (2022) (BBK) across different households in the terminal period. Since Baqaee et al. (2022) only present their algorithms in the case of the intial period as base, we only focus on this case. As the figure shows, the recursive BBK algorithm leads to levels of error that are fairly close to those of our first-order algorithms. However, the main distinction between the methods lies in

Figure C.3: nhCES Example: Terminal-Period Error Across Different Households

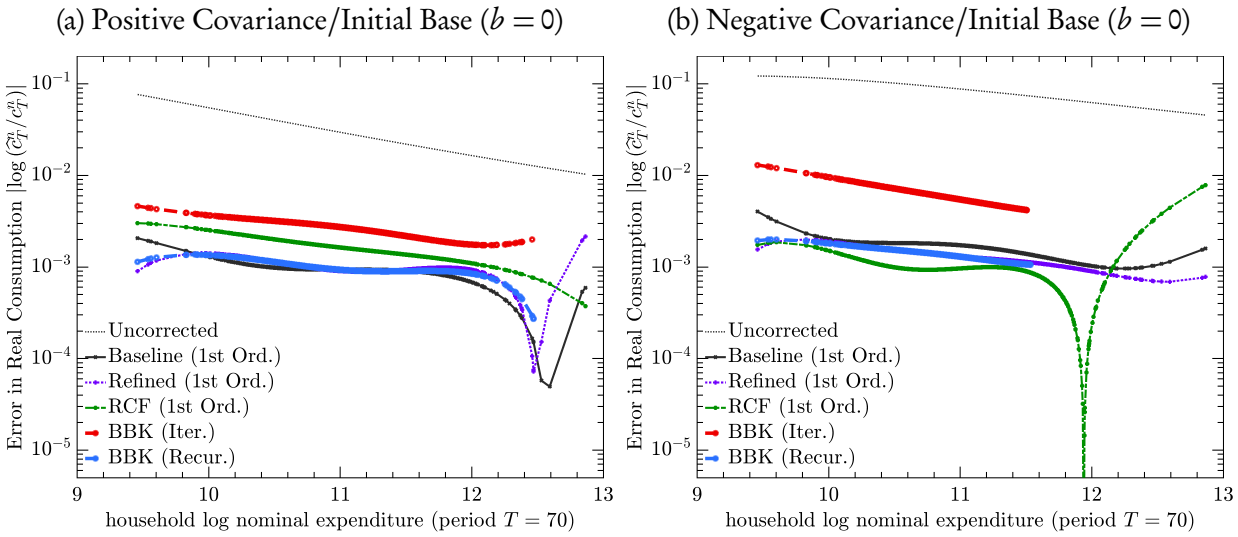

Note: The figures compare the error in the approximate value of real consumption between our first-order algorithm and those of Baqaee et al. (2022) (BBK) in the terminal period, across different households as a function of their total expenditure. The correct value of real consumption is calculated based on the underlying parameters of the nhCES preferences. The panels show the error for the choices of base period (a) the positive income elasticity–inflation covariance and (b) the negative covariance.

the no-extrapolation restriction that BBK imposes on their nonparametric fits to the expenditure share functions $\widehat{\omega}_{it}(\cdot)$. As a result, the BBK algorithms fail to recover the real consumption for the richest households in the final period. In particular, in the case with the negative covariance, which features faster growth in real consumption, the BBK method fails to produce any estimates for around 16% of the richest households in the data (the corresponding share in the positive covariance case is 0.4%). In contrast, all our algorithms recover approximate values of real consumption with levels of error far lower than those in the uncorrected approach. In particular, our benchmark and refined first-order algorithms yield estimates that are fairly close to the iterative BBK method over the range for which it produces estimates, but removes the restriction imposed by the BBK algorithm.

C.3 Extension to Other Values of Inflation-Income Elasticity Covariance

To show how the results extend to other ranges of the values of covariance between price inflations and expenditure elasticities, we perform one last exercise with our illustrative simulation. We consider alternative trends in prices, varying the deviations between inflation in services and agriculture from that in manufacturing (fixed to the average level of 3.19%) symmetrically from -2% to +2%. As previously, we compare the chained measures of deflated nominal consumption growth with and without our correction. Figure C.4 reports the error in the approximated val-

Figure C.4: Example: Real Consumption Error and Income Elasticity-Inflation Covariance

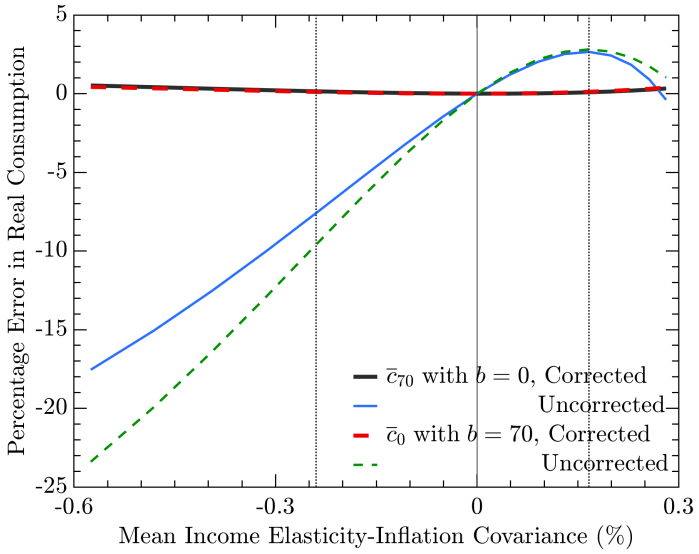

Note: The figure compares the error in the corrected and uncorrected approximations of the average final and initial real consumption for the initial and final periods as base, respectively, as a function of the mean covariance between price inflations and expenditure elasticities over the period.

ues of average real consumption, depending on the choice of the base period.[A5] As previously, the figure considers two cases, with either positive or negative income elasticity–inflation covariances.

The figure shows that when income elasticities are uncorrelated with the level of inflation across goods, the uncorrected measures approximate the correct values with negligible errors. However, as the covariance deviates from zero, the bias in the uncorrected measures grows. As the covariance falls to around -0.6% per year, the error in the uncorrected measure grows to around 20% of average real consumption.[A6] In contrast, the error in the approximation achieved with our nonhomotheticity correction remains close to zero over the entire range of values of the covariance, which highlights the accuracy of our algorithm.

D Data Appendix

In this appendix, we describe the data construction steps for the dataset in our main analysis, as well as for robustness checks.

[A5] We focus on the period that is most distant from the base period so that the error can potentially cumulate. Thus, we report the error in the final period when the initial period is taken as base. Symmetrically, we report the error in the initial period when the final period is taken as base.

[A6] As the covariance grows above zero, the error initially rises but ultimately begins to fall for large and positive values of covariance. This is because those scenarios lead to negligible growth in average household real consumption, which mechanically reduces the size of the bias in the reduced-form indices.

D.1 Dataset for the Main Analysis

Our main analysis dataset covers the period from 1955 to 2019, combining price series from the Consumer Price Index (CPI) to household expenditure data from the consumer expenditure survey (CEX).

Consumer Price Index Dataset The Consumer Price Index (CPI) data series contain monthly or quarterly price indices for over 200 detailed product categories. The price series are available over various time frames.[A7] To obtain a balanced panel of inflation series derived from the CPI, whenever a category is missing we use a more aggregate series in the product hierarchy as a proxy, since higher-level series usually have longer time coverage.[A8] The category-level inflation rate is obtained by averaging these price series at the desired frequency (annual or quarterly).

Consumer Expenditure Survey Datasets from 1984 to 2019 We obtain household expenditures from the Consumer Expenditure Survey (CEX) public-use microdata.[A9] Specifically, we use the interview survey data, which covers the full consumption basket from 1990 to 2019. Sampled households are interviewed at a quarterly frequency for four to five consecutive rounds, and report monthly expenditures at the universal classification code (UCC) level for the three months prior to the interview month in each round. Households also provide socio-demographic characteristics in each quarter of the survey, such as annual income and age of all household members. We use self-reported before-tax annual income prior to 2004 and imputed annual income in or after 2004 to classify households into income groups (e.g., deciles, quintiles, or percentiles) in each quarter.[A10]

We restrict the expenditure data to only include the UCCs that appear in the annual hierarchical grouping auxiliary files provided by the BLS. Indeed, these auxiliary files define the set of relevant UCCs that the BLS uses to produce the CE summary tables of household expenditures by socio-demographic characteristics.[A11] Furthermore, we exclude the UCCs belonging to

[A7] The data is available at https://download.bls.gov/pub/time.series/cu.

[A8] For example, series CUUR0000SEME (Health insurance) is a level 2 series beginning in 2005; filling it back to 1955 requires using level 1 series CUSR0000SAM2 (Medical care services) for 1957 to 2005 and level 0 series CUSR0000SAM (Medical care) for 1955 to 1956.

[A9] The data is available at https://www.bls.gov/cex/pumd_data.htm.

[A10] The main dataset is restricted to households with strictly positive before-tax income. In a robustness check, we keep households with zero imputed income from 2004 onwards. The results are similar (unreported).

[A11] The hierarchical grouping auxiliary files are only available back to 1997, so we apply the UCC restriction as specified in the 1997 file to earlier years with minor adjustments that come from comparing our estimates of average annual expenditure by product and income quintile with the CE tables from 1990 to 1996. The CE summary tables can be found at: https://www.bls.gov/cex/tables/calendar-year/mean-item-share-average-standard-error.htm (2012 onward); https://www.bls.gov/cex/csxstnd.htm (prior to 2012).

the categories "pensions & social security," "life and other personal insurance," and "education," which are long-run investments.[A12] We thus obtain a dataset containing 598 UCC product codes.

We benchmark our data against official estimates provided by the BLS in CE summary tables. Using the expenditure microdata for the relevant product UCCs, we calculate average annual expenditure for 32 product categories[A13] by income quintiles. Our results closely approximate the values reported in the CE summary tables, but we do not match them exactly because CE summary tables source expenditure data from both interview and diary surveys, while we only use interview data. To be exactly consistent with the annual consumption patterns published by BLS, we compute a scaling factor to adjust the expenditure microdata by the ratio between CE table values and our estimates, such that all average annual expenditures match the CE summary tables exactly, for each of the 32 product categories and income quintile.[A14]

The BLS provides monthly expenditure microdata by UCC and households starting in 1990 only; data prior to this date require special treatment. From 1980 to 1989, CEX microdata files are not suitable for our analyses. Indeed, for the period 1982-1989, BLS does not provide expenditure microdata at the UCC level. Moreover, in 1980 and 1981, expenditure microdata contain many legacy UCCs that were no longer in use in 1997, which is the earliest year for which the hierarchical grouping auxiliary files are available; therefore, we cannot reliably define the universe of relevant UCCs for these two years. However, the BLS provides CE summary tables from 1984 onward,[A15] which we combine with the 1990 microdata to obtain expenditure patterns from 1984 to 1989. Specifically, we assume that the expenditure shares for any given income group *within* each of the 32 product categories remain the same as in the 1990 microdata, and we use the CE summary table to adjust expenditure shares for each of the 32 categories from 1984 to 1989. The scaling factors are computed at the level of before-tax income quintile and the 32 product categories from CE summary tables in each year. We then aggregate the microdata and calculate average annual expenditures for the desired income groups and product categories. For the main analysis dataset, we compute the average annual expenditures at the "before-tax income percentile

[A12] For these categories, changes in returns to investment – and therefore the effective inflation rate for these categories – are difficult to measure accurately. Building a nonhomothetic price index accounting for savings and investment behavior is an important direction for future research, which is outside of the scope of this paper.

[A13] The 32 product categories are: food at home; food away from home; alcoholic beverages; shelter; utilities, fuels, and public services; household operations; household furnishings and equipment; clothes for men and boys; clothes for women and girls; clothes for children under 2; footwear; other apparel products and services; vehicle purchases (net outlay); gasoline, other fuels, and motor oil; other vehicle expenses; public and other transportation; health insurance; medical services; prescription drugs; medical supplies; fees and admissions; audio and visual equipment and services; pets, toys, hobbies, and playground equipment; other entertainment supplies, equipment, and services; personal care products and services; reading; education; tobacco products and smoking supplies; miscellaneous; cash contributions; life and other personal insurance; pensions and social security.

[A14] The scaling factor is applied to each of the "product categories by income quintile" cells.

[A15] CE summary tables from 2012 to 2020 can be found here. Historical summary tables from 1984 to 2011 can be found here.

by UCC" level, using household project weights provided by the CEX.[A16]

In all analysis and robustness datasets, we include a set of seven household characteristics that can serve as controls in regression specifications: (1) the raw number of household members; (2) family size with adjustment based on the OECD-modified equivalence scale;[A17] (3) family size after restricting to members aged 18 and over; (4) the average age of all household members; (5) the average age of all household members aged 18 and over; (6) household race[A18]; and (7) the highest level of education among all household members.[A19]

Consumer Expenditure Survey Datasets from 1955 to 1983 We also build a dataset tracking households' expenditure patterns back to 1955, using the expenditures shares at the level of 32 product categories in 1984, 1972 and 1960 documented in available CE summary tables.[A20]

The 1972 table provides annual average expenditures by income decile for 42 product categories. We harmonize these items with the 32 product categories available in summary tables available from 1984 onward. Using the average annual expenditure levels by product categories and income decile in 1972 and 1984,[A21] we interpolate expenditure shares in each of the intervening years, assuming constant increments in expenditure shares for each product category and income decile.[A22] As previously, we keep expenditure shares for any given income group *within* each of the 32 product categories at the level observed in the 1990 microdata.

We follow analogous steps using the 1960 CE summary table, which provides annual average

[A16] Since we use calendar year as the time unit, and households that are interviewed in February and March report expenditures across two calendar years, we apply an adjustment to the survey weights as instructed by Section 6 of the "Consumer Expenditure Surveys Public Use Microdata Getting Started Guide," which can be found here.

[A17] According to the OECD-modified equivalence scale, the first adult in a household has an equivalence value of 1; any additional adult or child aged 14 and over has an equivalence value of 0.5; any child aged 13 and under has an equivalence value of 0.3.

[A18] The majority race is chosen to represent the household. In the event of a tie, the household race is randomly determined.

[A19] When aggregating the data to the level of pre-tax income percentiles, for household race and highest level of education, we convert each factor variable into multiple variables capturing the percentage of households corresponding to each distinct value. Therefore, we have five variables expressed in percentages for race (Asian or Pacific Islander, Black, White, Native American, Multi-race or Other), and eleven variables for highest level of education (Never attended, Some or completed elementary school, Some or completed middle school, Some high school (no diploma), High school graduates, Some college (no diploma), Associate or professional degree, Bachelor's degree, Some graduate school (no diploma), Master's degree, Doctorate degree).

[A20] Prior to 1990, we do not have reliable microdata at the household or UCC level but annual CE summary tables on household expenditures are available by socio-demographic characteristics back to 1984. Prior to 1984, we do not have CE summary tables except for 1972 and 1960, which can be downloaded here and here.

[A21] The BLS only provides summary expenditure table by income decile in 1972, and by income quintile in 1984. To harmonize the income class and allow for direct comparisons, we first compute scaling factors at the level of income quintiles using the 1984 table as the benchmark. The scaling factor is applied to households depending on the income quintile to which they belong, and we then aggregate the household-level data to the level of income percentiles.

[A22] Results with alternative interpolation methods are similar (unreported).

expenditures for nine income brackets and 19 product categories, which we link to the 42 product categories observed in 1972 by building a one-to-many crosswalk. To create meaningfully comparable income groups between 1960 and 1972, we first convert the data structure from the nine income brackets to income deciles.[A23] We then interpolate expenditure shares between 1960 and 1972. We thus obtain a dataset matching CE summary tables exactly back to 1960. Given that there is no CE table prior to 1960, we assume that expenditure shares remain constant for the period 1955-1960.

Finally, as with the dataset in our main analysis from 1984 to 2019, after making adjustments by the scaling factors from the historical CE summary tables, we aggregate the microdata to "before-tax income percentile by UCC" cells in each year.

Data on Consumption Expenditures by Income and Age Following the same data construction steps as for the main dataset on consumption expenditures by income groups and products, we build an alternative dataset aggregating households into "income decile by age decile" cells. Specifically, households are first assigned into before-tax income deciles, then further divided into age deciles within each income decile based on the average age of all adults in the household.[A24] Just like the main dataset, the microdata is adjusted so that we exactly match the CE summary tables by income quintile from 1984 to 2020, as well as in 1960 and 1972. As previously, we use interpolation to obtain expenditure shares in intervening years. As a robustness check, we calculate alternative scaling factors using CE summary tables by household head (reference person) age bracket instead, while keeping all other data treatment unchanged, which yields similar results (unreported).

Since the size of the bias from the household aging correction is governed by changes in average age over time, it is important to check the accuracy of the age data. We check that average age in our household survey data matches the benchmark series of the UN World Population Prospects. Average age in our data is close to this external benchmark. To guarantee an exact match, we apply a year-specific scaling factor to the age variable in our data; this scaling factor is the same for all households in a given year. For all years prior to 1984 in which the CE summary tables are not available, we use the benchmark series of the UN World Population Prospects to impute average household age.

Linking Consumption and Price Datasets To link the CPI price series to household expenditures from the CEX, we manually build a crosswalk, starting from the UCC to CPI concordance

[A23] We translate the boundaries of income brackets into percentiles using the 1960 before-tax income distribution in the U.S.; we then assign income brackets to income deciles to maximize overlap.

[A24] The average age of all adults in the household is calculated by averaging the age of all household members at or above the age of 18.

provided by the BLS[A25] and extending coverage back in time. All expenditure categories are mapped to at least one inflation series from the CPI price data.[A26] Our main dataset is thus at the UCC level and includes 598 unique product codes that map to 159 CPI inflation series.

Year-specific Scaling Factor to Match BEA's Aggregate Personal Consumption Expenditure For all datasets, we ensure that we match the BEA's aggregate personal consumption expenditures. We apply a year-specific scaling factor to the household consumption data so that we match the BEA's nominal personal consumption expenditure per household in each year. This step is useful for our purposes since the bias from the nonhomotheticity correction depends on consumption growth over time, and since household expenditure surveys are known to miss some expenditures. Our approach allows us to compute inflation inequality in an empirical setting that is fully in line with the average nominal consumption growth observed in the U.S. national accounts. This scaling step follows the spirit of distributional national accounts of Piketty et al. (2018), ensuring that our analysis is consistent with macroeconomic aggregates.

D.2 Datasets for Sensitivity Analysis

We build four alternative datasets to assess the robustness of our findings to data construction choices.

Sensitivity to Aggregation Level: Robustness Datasets #1, #2 and #3 To assess whether our results are sensitive to aggregation choices, we build two alternative datasets that closely follow our main dataset but use different levels of aggregation, grouping UCCs into broader categories. First, we create a version of the dataset using the 32 product categories from the CE summary tables. The crosswalk between UCCs and the 32 CE table product categories is provided in the hierarchical grouping auxiliary files. Second, we manually group the 598 UCCs into 119 mutually exclusive product categories that are continuously available from 1984 to 2019.[A27]

In addition, we use Nielsen scanner data for consumer packaged goods to implement Algorithm 1 on highly disaggregated data. The main product categories covered in the Nielsen data are food and drinks at home, housekeeping supplies, household cleaning products, as well as personal care products, smoking products, tableware, tools, nonelectric cookware, and apparel.

[A25] The up-to-date UCC-ELI concordance can be found here: https://www.bls.gov/cpi/additional-resources/ce-cpi-concordance.htm.

[A26] While most UCCs are mapped to a single CPI category, when there is more than one relevant CPI series, we take the simple average of all relevant series to obtain the price change for that UCC.

[A27] These two robustness datasets allow us to compute additional price indices that require observing the same set of product categories between consecutive periods, e.g. a Törnqvist price index. In contrast, there is substantial churn for UCC items across years.

These product categories account for 13.39% of overall household spending, which corresponds to close to 40% of expenditures on goods. We conduct the analysis at the level of "product modules by price decile" cells, as in Jaravel (2019).

Sensitivity to Official Aggregate Expenditure Weights in CPI: Robustness Datasets #4 The fourth alternative dataset for robustness is based on the official consumption weights used by the Bureau of Labor Statistics when calculating the CPI.[A28] We use the official consumption weights for eight product categories that are available every year back to 1955. The eight broad product categories included in this dataset are: food and beverages, housing, apparel, transportation, medical care, recreation, education and communication, other goods and services. Due to the evolution of product categories and product hierarchy over the years, some sub-categories are reassigned by BLS from one broad category to another over time. For example, BLS places "Telephone services" under housing until 1997, then under "Education and communication." To address this issue, we adjust the placement of certain sub-categories and their allocated weights so that the composition of broad categories remains consistent from 1955 to 2019.

In addition to the aggregate consumption weights, our linked dataset uses expenditure shares by income quintiles from the CE summary tables published by the BLS, which are available from 1984 onward, as in the main dataset. Prior to 1984, we assume that the expenditure shares remain identical to 1984. We use the expenditure shares of each income quintile to distribute aggregate consumption across income groups, so that we obtain a linked dataset with consumption patterns that vary across income groups while keeping aggregate, category-level consumption weights identical to the official weights of the BLS for its eight product categories that can be tracked back to 1955.

E Additional Figures

[A28] The official consumption weights are available at https://www.bls.gov/cpi/tables/relative-importance/home.htm. They can differ from the expenditure patterns reported in the CE summary tables.

Figure E.1: Additional Evidence on Inflation Inequality over Time

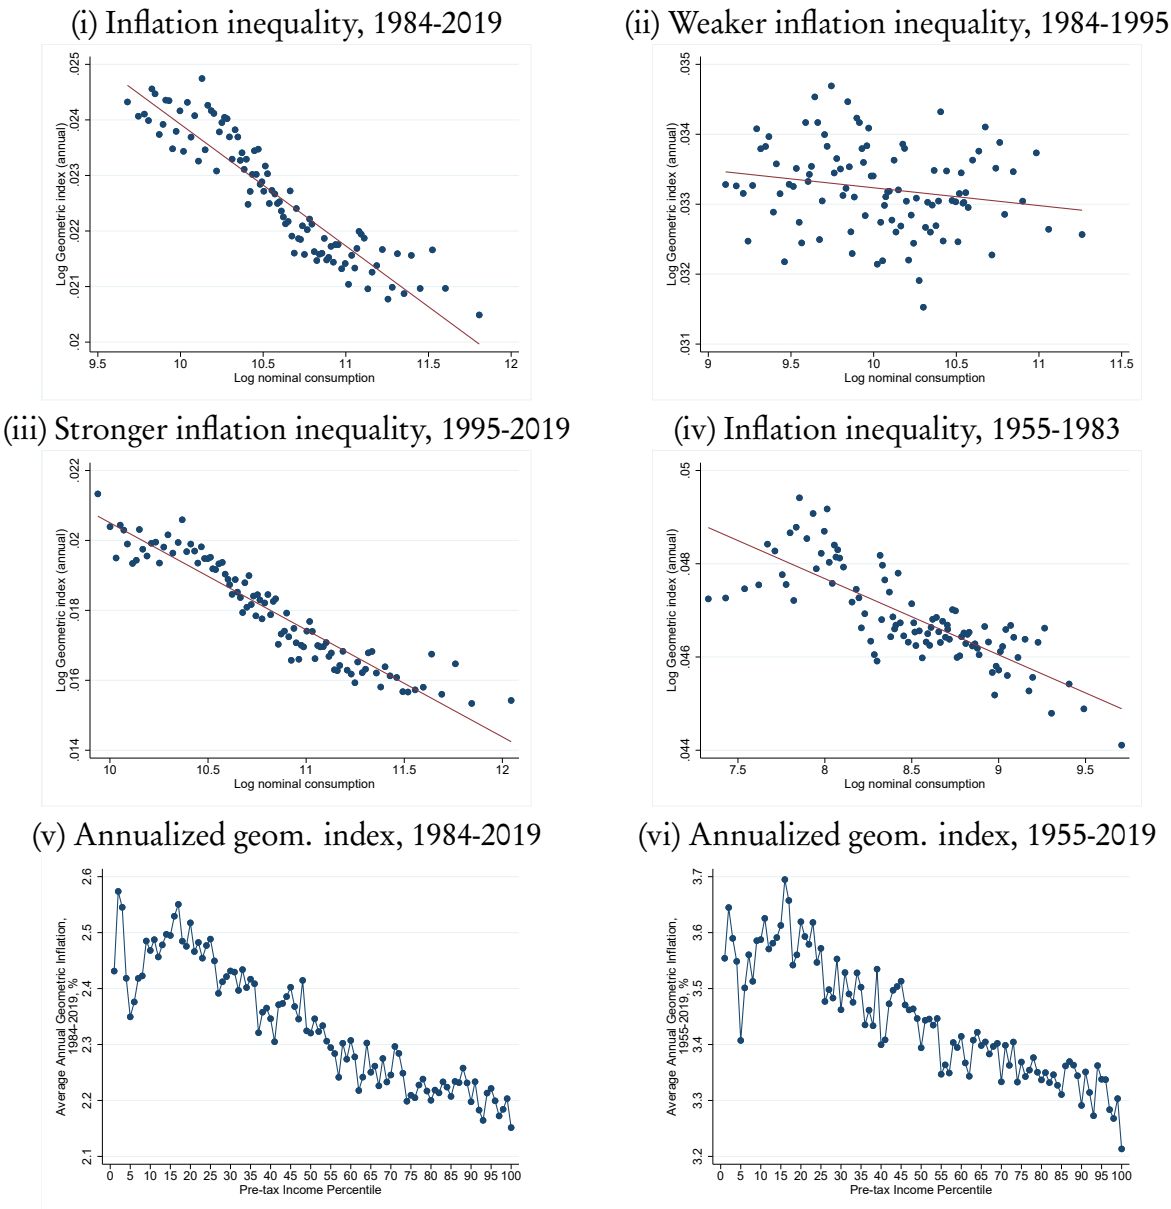

Note: This figure reports descriptive patterns on inflation inequality. In panels (i) through (iv), households are grouped by pre-tax income percentile in each year. These panels report binned scatter plots depicting the relationship between the annual inflation rate and log nominal consumption, absorbing time fixed effects. Each dot represents 1% of the data and all panels use the geometric price index. Panels (v) and (vi) report the annualized inflation rate, for the periods 1984-2019 and 1955-2019, respectively, using the chained geometric index.

Figure E.2: Nonhomotheticity Correction and the Consumption Deflator

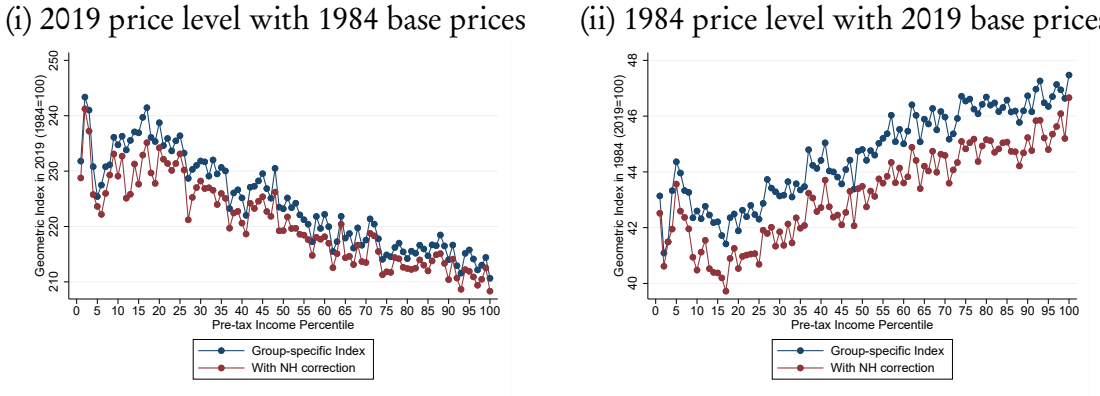

Note: This figure reports the chained index formula, $\Pi_t \pi_t^n$, compared with the corrected non-homothetic deflator, y_t^n / c_t^n.

Figure E.3: Biases in 1984-2019 Cumulative Real Consumption Growth by Income Percentile

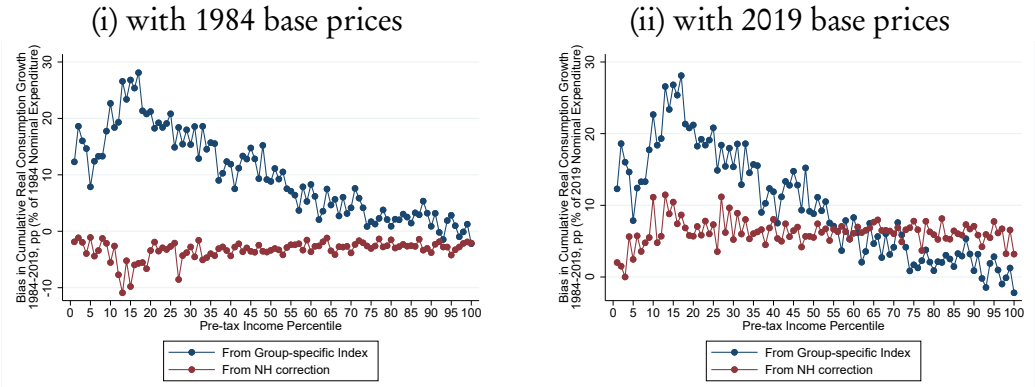

Note: This figure compares the magnitude of biases in the measurement of cumulative consumption growth from 1984 to 2019, reporting the deviation from the aggregate homothetic price index due to (a) percentile-specific homothetic price indices, and (b) due to the nonhomotheticity correction. Panel (i) uses 1984 prices as the base for the nonhomotheticity correction, while panel (ii) uses 2019 prices. The bias from percentile-specific indices is identical in both panels.

Figure E.4: Cumulative and Annualized Growth Rates across Price Indices

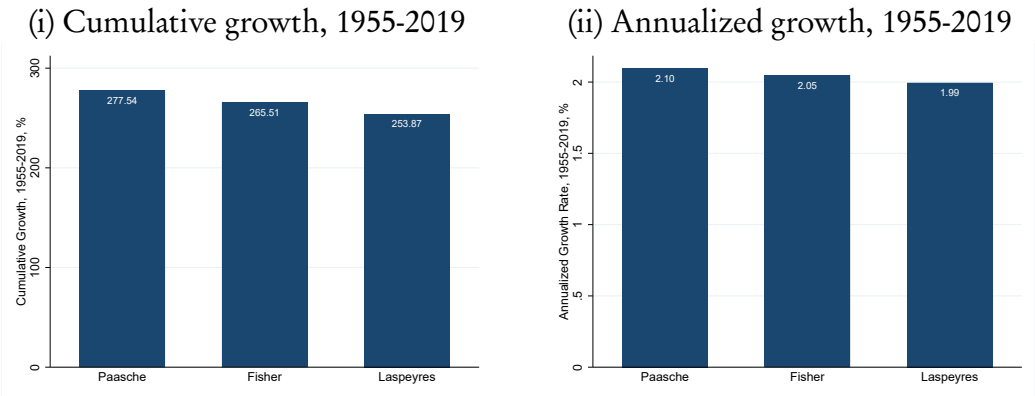

Note: This figure reports cumulative and annualized growth rates from 1955 to 2019 for three price indices, Paasche, Fisher and Laspeyres.

Figure E.5: Sensitivity Analysis for the Annual Bias in Real Consumption Growth

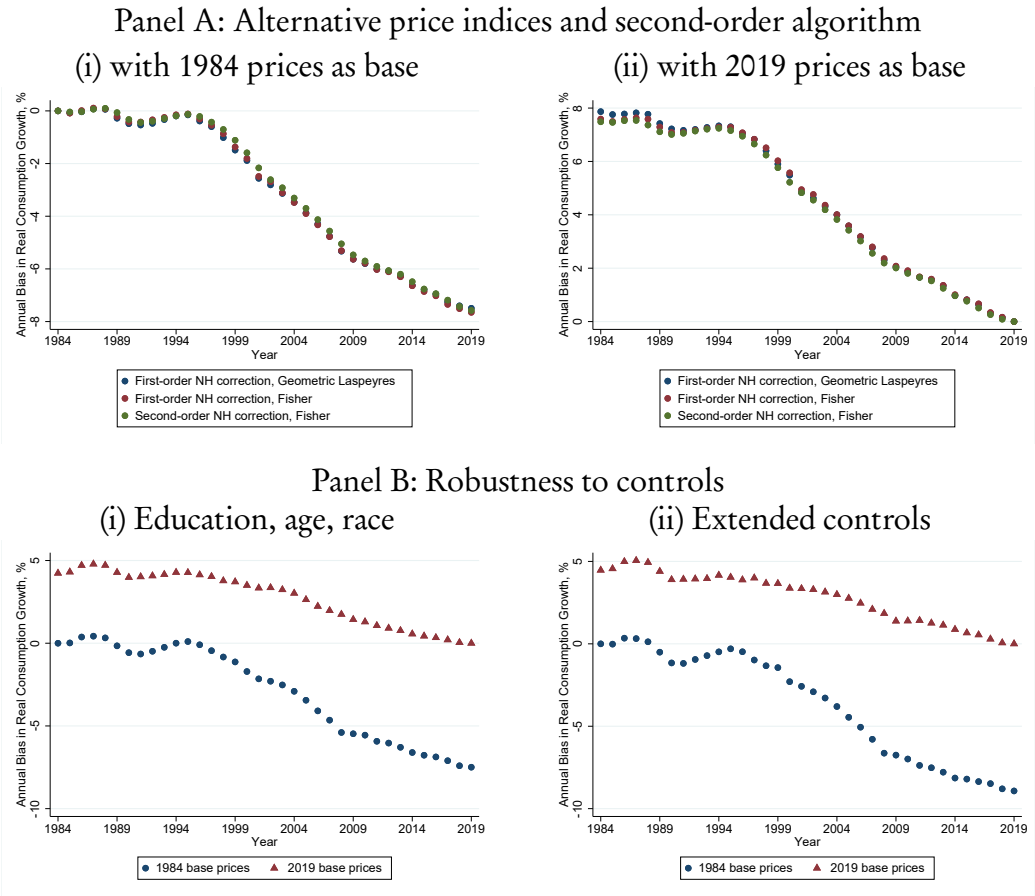

Note: This figure report the biases in annual average real consumption growth per household due to the nonhomotheticity correction under different specifications. Panel A reports the results under alternative price indices, geometric or Fisher, with the first-order algorithm, as well as with the second order algorithm. Panel A(i) uses 1984 prices as the base, while Panel A(ii) uses 2019 prices. Panel B reports the results with controls, using the geometric index and the first order algorithm. Panel B(i) controls for education, age, and race in the estimation of the income elasticity of inflation. Panel B(ii) controls for region, urban vs. rural area, gender, and city population size, in addition to education, age, and race.

Figure E.6: Sensitivity Analysis for the Degree of the Polynomial Used in Algorithm 1

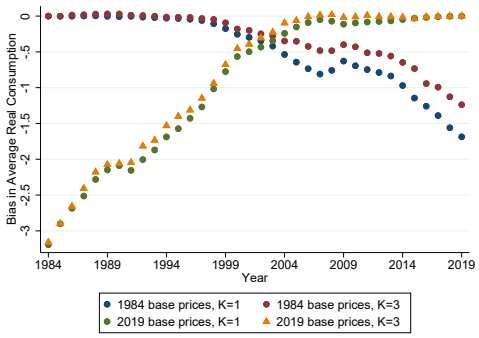

Note: This figure reports the biases in the level of average real consumption per household depending on the degree of the polynomial used in Algorithm 1. We report the results for first-order ($K = 1$) and third-order ($K = 3$) polynomials. The figure is otherwise identical to panel (a) of Figure 3 in the main text, which uses a second-order polynomial.

Figure E.7: Results with Alternative Algorithms

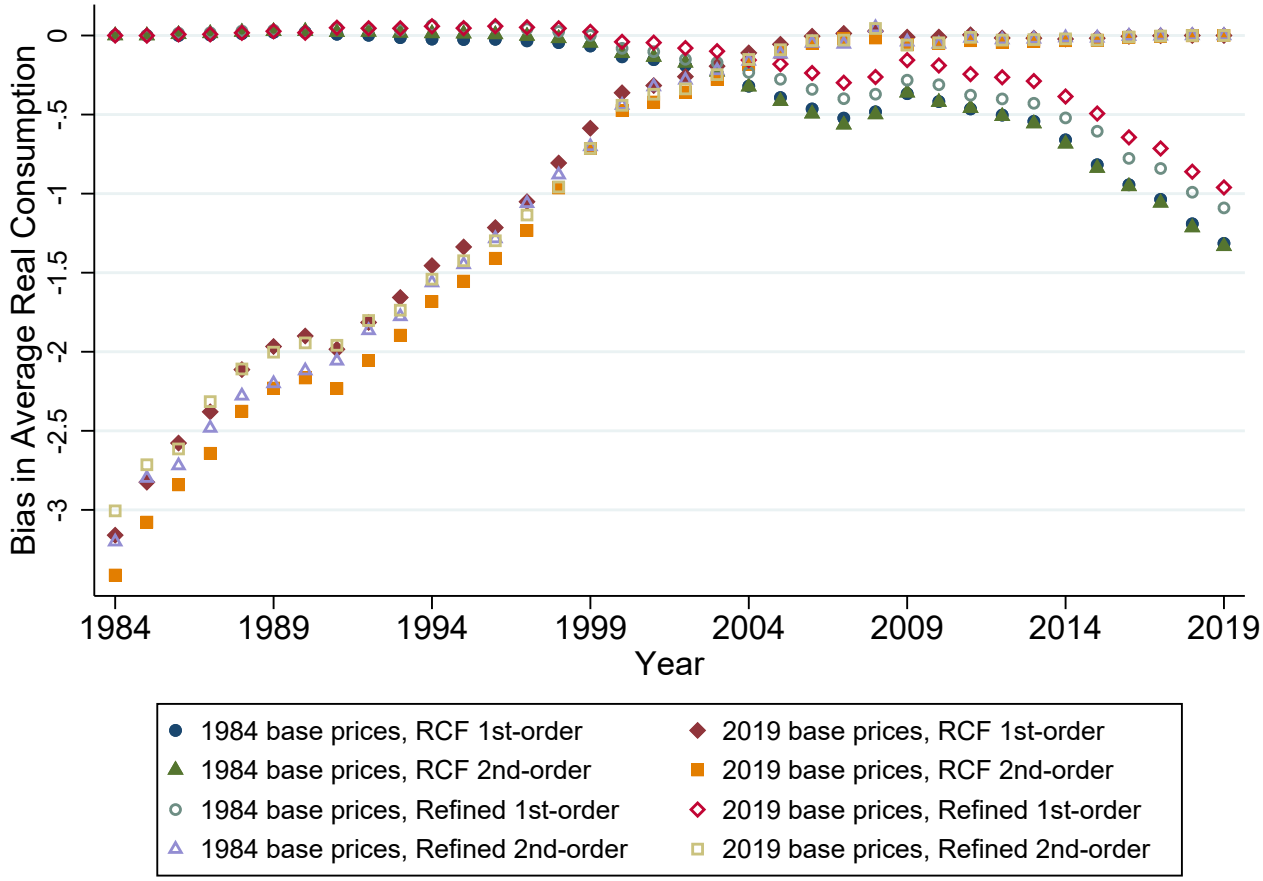

Note: This figure is identical to Panel (a) of Figure 3 in the main text, except that we use the alternative algorithms described in Appendix A. Specifically, we present the results obtained with the first-order refined algorithm (Algorithm A.1), the second-order refined algorithm (Algorithm A.4), the estimation of the real consumption function (RCF) to the first order (Algorithm A.2), and to the second order (Algorithm A.5). The figure reports the biases implied by each algorithm for the level of average real consumption estimated using the conventional uncorrected approach.

Figure E.8: Results with State Fixed Effects

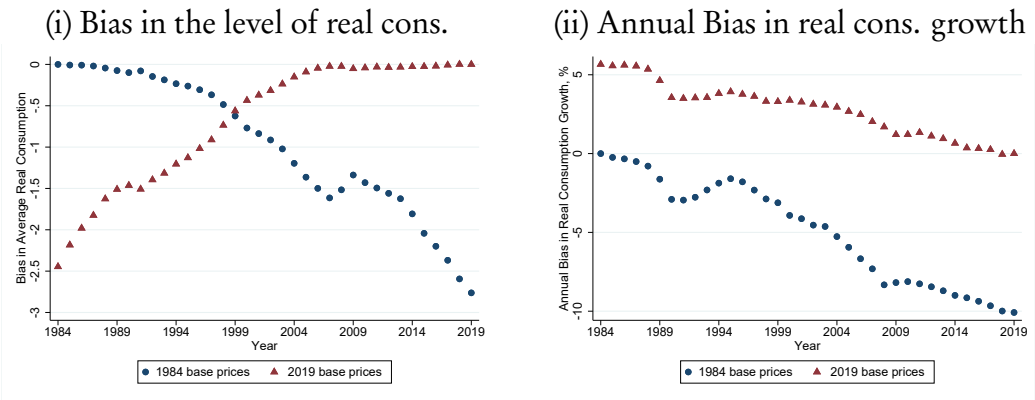

Note: This figure reports the biases when using state fixed effects in the estimation of the income elasticity of inflation. Specifically, we control for state fixed effects, education, age, race, urban vs. rural area, gender, and city population size. The figure is thus identical to Panel B(ii) of Figure 6 in the main text, except that we control for state instead of region.

Figure E.9: Nonhomotheticity Correction and Bias in Average Real Consumption with JL and BBK Algorithms, 1984-2019

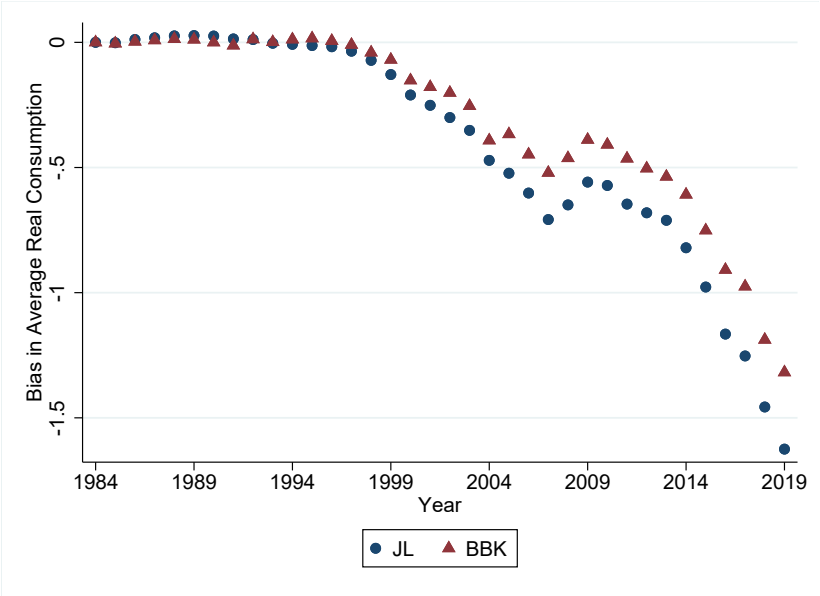

Note: This figure report the nonhomotheticity biases in the level of average real consumption per household. For each year, the bias is expressed in percentage of current average consumption. The bias is computed by applying our baseline first-order algorithm (JL) and the algorithm of Baqaee et al. (2022) (BBK) to obtain the nonhomotheticity correction. We then compare standard measures of average real consumption to corrected measures. The algorithms are applied to our main dataset at the level of pre-tax income percentiles, in the range for which the BBK algorithm can produce estimates (with no extrapolation step). We then average percentile-level results to obtain average real consumption per household. We use 1984 as the base period for our money metric.

Figure E.10: Results with 32 Product Categories, 1984-2019

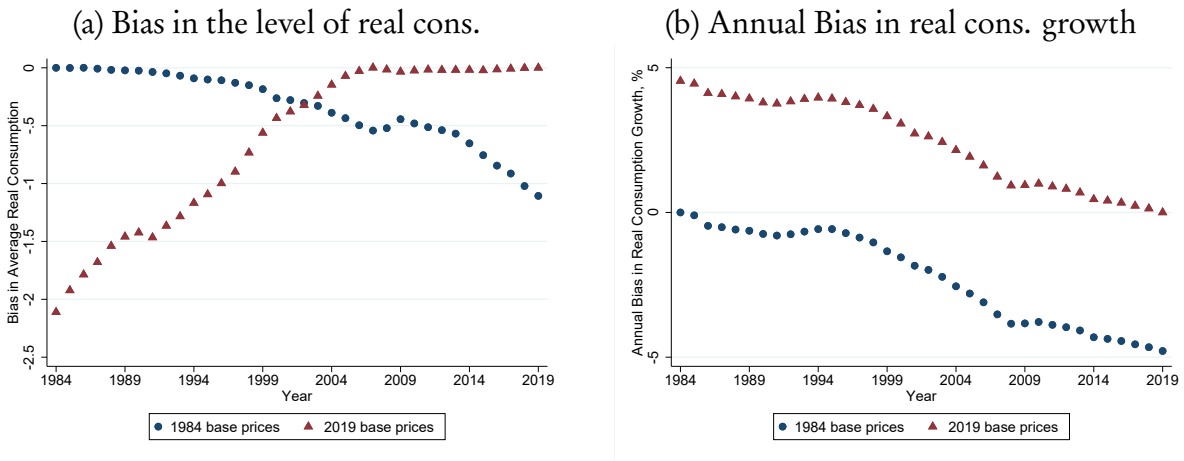

Note: This figure is identical to Figure 3 in the main text, except that we use our robustness dataset #1, i.e., we work with data at the level of 32 product categories from the CE summary tables. This figure reports the biases in the level of average real consumption per household, in panel (a), and in annual growth in real consumption per household, in panel (b). The bias is computed by applying Algorithm 1 to obtain the nonhomotheticity correction. We then compare standard measures of real consumption to corrected measures. In panel (b), the bias is expressed as a percentage of the standard homothetic measure of current-period growth. Algorithm 1 is applied to our robustness dataset #1 at the level of pre-tax income percentiles, using geometric price indices. We then average percentile-level results to obtain average real consumption per household.

Figure E.11: Results with 114 Product Categories, 1984-2019

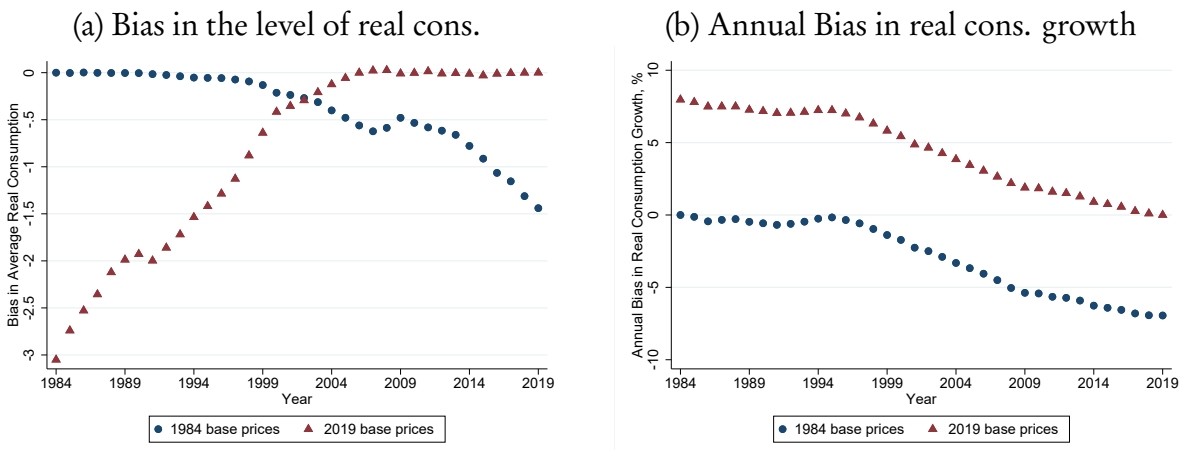

Note: This figure is identical to Figure 3 in the main text, except that we use our robustness dataset #2, i.e., we work with data at the level of 114 product categories that are continuously available between 1984 and 2019. This figure reports the biases in the level of average real consumption per household, in panel (a), and in annual growth in real consumption per household, in panel (b). The bias is computed by applying Algorithm 1 to obtain the nonhomotheticity correction. We then compare standard measures of real consumption to corrected measures. In panel (b), the bias is expressed as a percentage of the standard homothetic measure of current-period growth. Algorithm 1 is applied to our robustness dataset #2 at the level of pre-tax income percentiles, using geometric price indices. We then average percentile-level results to obtain average real consumption per household.

Figure E.12: Results for Fast-Moving Consumer Goods with 9131 Product Categories, 2004-2014

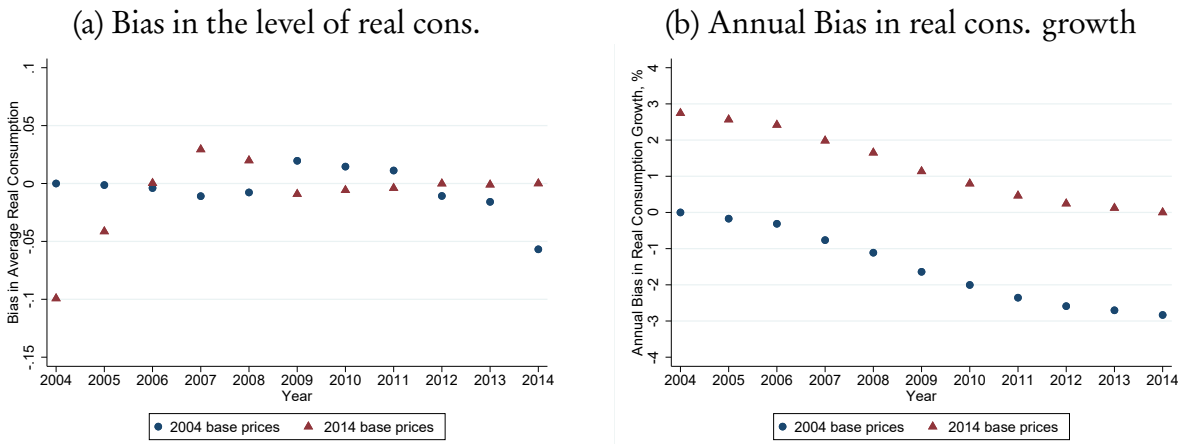

Note: This figure is identical to Figure 3 in the main text, except that we use our robustness dataset #4, i.e., we work with data at the level of 9131 product categories that are available in the Nielsen Homescan Consumer Panel Data between 2004 and 2014. This figure reports the biases in the level of average real consumption per household, in panel (a), and in annual growth in real consumption per household, in panel (b). The bias is computed by applying Algorithm 1 to obtain the nonhomothecticity correction. We then compare standard measures of real consumption to corrected measures. In panel (b), the bias is expressed as a percentage of the standard homothetic measure of current-period growth. Algorithm 1 is applied to our robustness dataset #4 at the level of pre-tax income deciles, using geometric price indices. We then average decile-level results to obtain average real consumption per household.

Figure E.13: Results for CPI-CEX Data Restricted to Fast-Moving Consumer Goods, 2004-2014

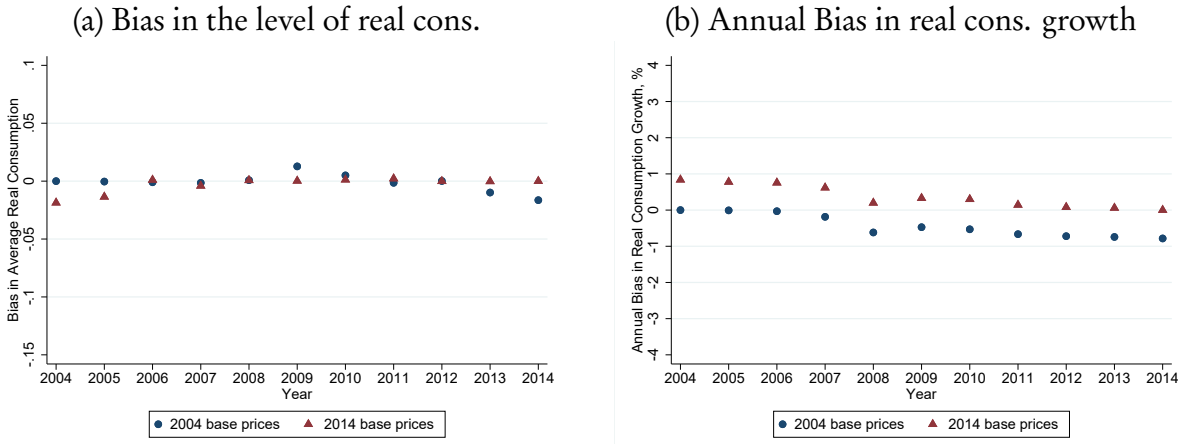

Note: This figure is identical to Figure 3 in the main text, except that we restrict the sample to product categories belonging to fast-moving consumer goods and focus on the period from 2004 to 2014. These sample restrictions allow for a comparison of the results obtained with the CPI-CEX dataset and the Nielsen dataset analyzed in Figure E.12. The restricted sample covers 44 UCC items belonging to the following categories: alcoholic beverage; food at home; personal care products; pets, toys, hobbies, and playground equipment; sewing machines, fabric and supplies; tools, hardware, outdoor equipment and supplies. The figure reports the biases in the level of average real consumption per household, in panel (a), and in annual growth in real consumption per household, in panel (b). The bias is computed by applying Algorithm 1 to obtain the nonhomothecticity correction. We then compare standard measures of real consumption to corrected measures. In panel (b), the bias is expressed as a percentage of the standard homothetic measure of current-period growth. Algorithm 1 is applied at the level of pre-tax income deciles, using geometric price indices. We then average decile-level results to obtain average real consumption per household.

Figure E.14: Results for Fast-Moving Consumer Goods with 9131 Product Categories Accounting for Increasing Product Variety, 2004-2014

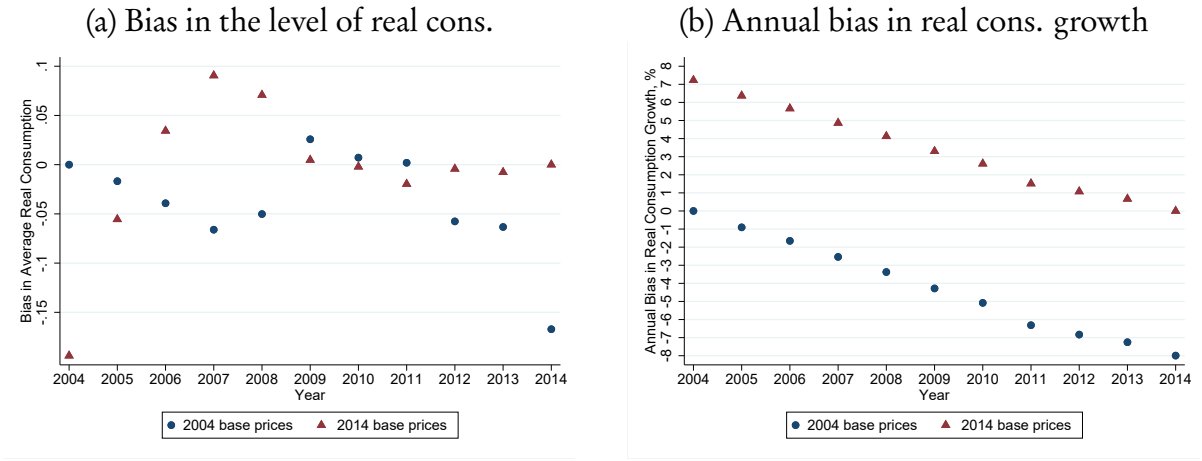

Note: This figure is identical to Appendix Figure E.12, except that for each the 9131 product categories available in the Nielsen data we build a CES price index accounting for changes in product variety over time, using the methodology of Feenstra (1994), which was applied to scanner data in Broda and Weinstein (2010) and Jaravel (2019). The figure reports the biases in the level of average real consumption per household, in panel (a), and in annual growth in real consumption per household, in panel (b), using Algorithm 1. As previously, we average decile-level results to obtain average real consumption per household.

Figure E.15: Results with Official CPI Aggregate Expenditure Weights, 1955-2019

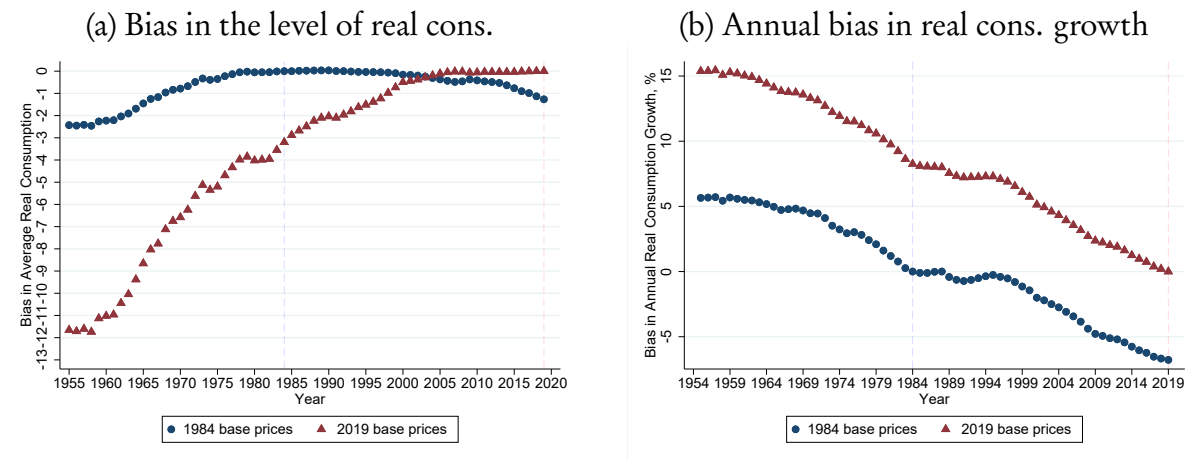

Note: This figure is identical to Figure 3 in the main text, except that we use our robustness dataset #3, i.e., we work with use official CPI aggregate expenditure weights for eight broad expenditure categories to rescale the household-level expenditure patterns, thus ensuring that our data is consistent with aggregate expenditures used by the BLS when computing the CPI. This figure reports the biases in the level of average real consumption per household, in panel (a), and in annual growth in real consumption per household, in panel (b). The bias is computed by applying Algorithm 1 to obtain the nonhomotheticity correction. We then compare standard measures of real consumption to corrected measures. In panel (b), the bias is expressed as a percentage of the standard homothetic measure of current-period growth. Algorithm 1 is applied to our robustness dataset #3 at the level of pre-tax income percentiles, using geometric price indices. We then average percentile-level results to obtain average real consumption per household.

Figure E.16: Inflation across Age Groups and over Time

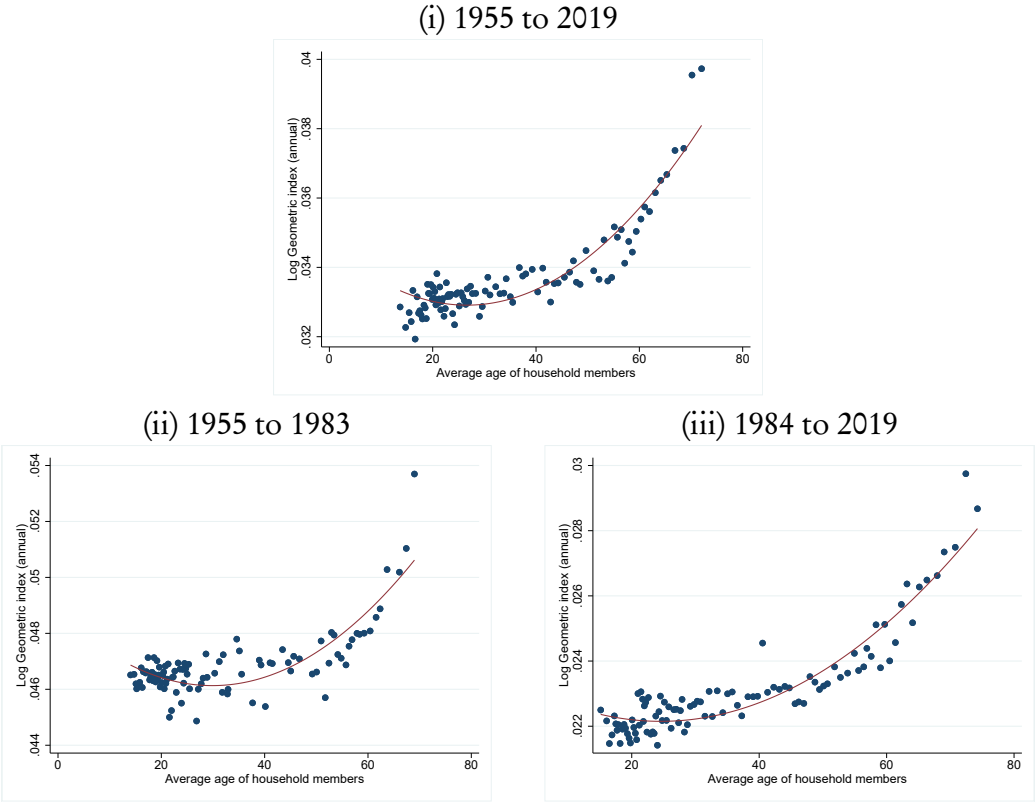

Note: This figure reports binned scatter plots depicting the relationship between the geometric index and the average age of household members. Each panel focuses on a different period. In each panel, each bin represents 1% of households. In each year, the unit of observation is "age decile by income decile" cells. All specifications include year fixed effects.

Figure E.17: Average Household Age Over Time in the United States

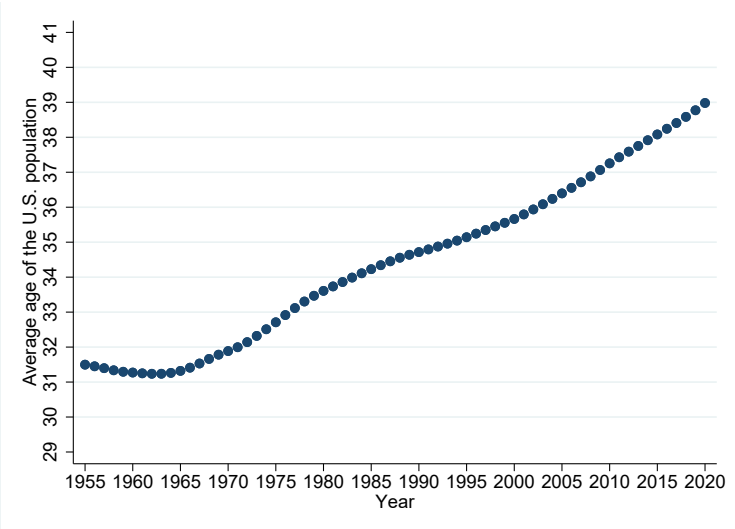

Note: This figure reports the change in average household age over time.