MEASURING GROWTH IN CONSUMER WELFARE WITH INCOME-DEPENDENT PREFERENCES: NONPARAMETRIC METHODS AND ESTIMATES FOR THE UNITED STATES*

XAVIER JARAVEL AND DANIAL LASHKARI

How should we measure changes in consumer welfare given observed data on prices and expenditures? This article proposes a nonparametric approach that holds under arbitrary preferences that may depend on observable consumer characteristics, for example, 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 because of 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. *JEL Codes*: D60, E31, I30.

*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 article 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 grant PLP-2019-069 and United Kingdom Research and Innovation (UKRI) grant EP/X016056/1 for generous financial support. Researcher(s)' own analyses calculated (or derived) based in part on data from Nielsen Consumer LLC and marketing databases provided through the NielsenIQ Data Sets 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 those of the researcher(s) 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.

The Quarterly Journal of Economics (2024), 477–532. https://doi.org/10.1093/qje/qjad039.
Advance Access publication on September 11, 2023.

I. 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 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 (2023) 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 article 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 with the standard setting, the only additional data requirement is access to a cross section of product prices and quantities for consumers with heterogeneous incomes. Such data are 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

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

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 indexing the poverty line and more efficiently targeting welfare benefits. This approach also provides a blueprint for distributional national accounts (Piketty, Saez, and Zucman 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. Because index formulas approximate the Divisia 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, 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 rise in total expenditure translates 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]

2. One way to understand this change in concavity is that it accounts for the cumulative effect of the past inflation inequality. Consumers who were previously poor may not have immediately benefited 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 II.B, 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.

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 recursively apply the same strategy in later periods to approximate the values of the 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 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 second-order approximations.[4]

We demonstrate the accuracy of our different algorithms using a simulation with known preference parameters, relying

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

on the nonhomothetic constant elasticity of substitution (nhCES) preferences of Comin, Lashkari, and Mestieri (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 article, 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 data set providing price changes and expenditure shares at a granular level from 1955 to 2019 across percentiles of the income distribution. This data set 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 data set 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.

Because 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 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 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% a 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, for example, age, family size, or education. 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 the 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

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. Because 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, Decancq, and Ooghe 2018, for the properties of money metric social-welfare functions).

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 III.

that the implied adjustments to real consumption are economically meaningful but much smaller than the nonhomotheticity correction, which justifies our focus on the latter.

I.A. Prior Work

Our article 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 II.C.[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 between income elasticities and price changes (e.g., Fajgelbaum and Khandelwal 2016; Baqaee and Burstein 2023; Atkin et al. 2024). In particular, Baqaee and Burstein (2023) 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

7. Our approach assumes utility maximization, and thus contrasts with the approach of Blundell, Browning, and Crawford (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.

8. Baqaee and Burstein (2023) also study the consequences of the endogeneity of prices in general equilibrium, as well as unobserved heterogeneity, for example, taste shocks. The latter effects have also been considered by Redding and Weinstein (2020). We note that, subsequent to our article, Baqaee, Burstein, and Koike-Mori (2024) proposed an alternative to our algorithms. We discuss the close connections between their approach and ours in Online Appendix B.3.

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. (2024). 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 almost ideal demand system (AIDS) (Deaton and Muellbauer 1980).[9] Atkin et al. (2024) 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. (2024) 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 data set 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 article is organized as follows: Section II presents our theory, approximation algorithms, and simulations. Section III reports the empirical analysis, and Section IV generalizes our approach to settings where preferences vary with observable consumer characteristics. Several proofs and additional results are reported in the Online Appendix.

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

II. 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 II.A introduces the notation and defines the main concepts used for the measurement of welfare, cost of living, and real consumption. Section II.B 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 II.C derives our approximate results in terms of observable data. Finally, in Section II.D we perform a simulation to illustrate and validate the accuracy of our approach.

II.A. Definitions

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 $y \equiv \Sigma_i p_i q_i$ and $s_i \equiv \frac{p_i q_i}{y}$. The function $\boldsymbol{\omega}_t(\cdot)$ characterizes the Marshallian demand for the vector of prices prevailing at time t.[10] Because 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 $0 \leqslant b \leqslant T$), define real consumption under constant time-b (base) prices as a monotonic transformation $M_b(\cdot)$ of utility u given by

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

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

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 often drop the superscript to simplify the expressions whenever it is clear that the base b is fixed.

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 $y_t = \chi_t^b(c_t^b)$. 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 $\frac{c_t^b}{c_{t_0}^b}$, 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 $\mathcal{P}_{t_0,t}^b(c)$ for a consumer with real consumption c (defined under base time period b) between periods t_0 and t ($0 \leqslant t_0$, $t \leqslant T$) as

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

Let us specifically consider the true price index defined between the base period b and the current period t, which satisfies $c \equiv \frac{\chi_t^b(c)}{\mathcal{P}_{b,t}^b(c)}$. 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:

$$\text{(4)} \qquad \frac{c_t^b}{c_{t_0}^b} = \frac{y_t}{y_{t_0}} \frac{\mathcal{P}_{b,t_0}^b(c_{t_0})}{\mathcal{P}_{b,t}^b(c_t)} = \frac{y_t}{y_{t_0}} \frac{1}{\mathcal{P}_{t_0,b}^b(c_{t_0}) \times \mathcal{P}_{b,t}^b(c_t)}.$$

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, $\mathcal{P}_{t_0,b}^b(c_{t_0})$, and the true price index between the base period b and the final period t, $\mathcal{P}_{b,t}^b(c_t)$. 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]

i. *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 $\boldsymbol{E}(u; \boldsymbol{p}) = P(\boldsymbol{p}) \cdot \boldsymbol{F}(u)$, 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 $\mathcal{P}_{t_0,t}^b(c)$ 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]

$$\text{(5)} \qquad \frac{c_t^b}{c_{t_0}^b} = \frac{y_t}{y_{t_0}} \frac{1}{\mathcal{P}_{t_0,t}^b(c)}, \qquad \text{for any } c \text{ and for any } b,$$

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 $EV = \frac{c_t^{t_0}}{c_{t_0}^{t_0}} = \frac{y_t}{y_{t_0}} \frac{1}{\mathcal{P}_{t_0,t}^{t_0}(c_t^{t_0})}$. Alternatively, choosing the final period t as the base leads to the concept of compensating variation (CV), given as $CV = \frac{c_t^t}{c_{t_0}^t} = \frac{y_t}{y_{t_0}} \frac{1}{\mathcal{P}_{t_0,t}^t(c_{t_0}^t)}$.

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

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. *Price Index Formulas.* The indices defined above 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 $\mathbb{P}(\boldsymbol{p}_{t_0}, \boldsymbol{s}_{t_0}; \boldsymbol{p}_t, \boldsymbol{s}_t)$ 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:

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

$$\mathbb{P}_G \equiv \prod_i \left(\frac{p_{i,t}}{p_{i,t_0}}\right)^{s_{i,t_0}}, \tag{6}$$

where we have suppressed the arguments $(\boldsymbol{p}_{t_0}, \boldsymbol{s}_{t_0}; \boldsymbol{p}_t, \boldsymbol{s}_t)$ to avoid repetition. As is well known, these 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 $\mathbb{P}_S$ index formulas defined as

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

$$\mathbb{P}_S = \prod_i \left(\frac{p_{i,t}}{p_{i,t_0}}\right)^{\bar{s}_{S,i}}, \tag{7}$$

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

II.B. 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

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

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, that is,

13. Online 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, Online Appendix B.2 characterizes the inverse mapping from total expenditure to real consumption, which we call the real consumption function.

$y_t = \chi_t^b(c_t^b) = E\left(M_b^{-1}(c_t^b); \boldsymbol{p}_t\right)$. 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

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

which leads to equation (10) after rearranging terms, since the first term on the right 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 and (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, we 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

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

The integral on the right 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

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.

differential equation (9):

(11)
$$\log \mathcal{P}^b_{t_0,t}(c) = \log \chi^b_t(c) - \log c = \int_{t_0}^{t} \log D_\tau \left(\chi^b_\tau(c)\right) d\tau, \qquad \forall b, c.$$

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 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 $\Lambda^b_t(\cdot)$ as the elasticity of the true index to real consumption from the base period to the current period, that is,

$$\Lambda^b_t(c) \equiv \frac{\partial \log \mathcal{P}^b_{b,t}(c)}{\partial \log c} = \frac{\partial \log \chi^b_t(c)}{\partial \log c} - 1, \tag{12}$$

so that the multiplicative factor in equation (10) is given by $(1 + \Lambda^b_t(c_t))^{-1}$. Under homothetic preferences, this nonhomotheticity correction is zero $\Lambda^b_t(c) \equiv 0$ 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).

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 $\mathcal{P}^b_{b,t}(c)$ is small. Second, the dependence of the index on real consumption

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

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]

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

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. 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 II.C, we present a number of other theoretical implications of Proposition 1.

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

16. We note that the importance of the covariance between income elasticities and price changes for measuring welfare change in the presence of nonhomotheticity has been highlighted in prior work (e.g., Fajgelbaum and Khandelwal 2016; Baqaee and Burstein 2023; Atkin et al. 2024). 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.

in period b_1, satisfies

$$(13) \qquad \frac{d\log c_t^{b_2}}{d\log c_t^{b_1}} = 1 + \Lambda_{b_2}^{b_1}\big(c_t^{b_1}\big) = 1 + \frac{\partial \log \mathcal{P}_{b_1,b_2}^{b_1}\big(c_t^{b_1}\big)}{\partial \log c_t^{b_1}}.$$

Proof. See Online 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 of real consumption growth are biased downward because they do not account for income-elastic goods becoming relatively cheaper at the same time that they become relatively more important in 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

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.

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]

2. *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 $\widetilde{\chi}_t^b(y) \equiv (\chi_t^b)^{-1}(y)$, 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. □

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.

Online 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 differential equation (9). The Online 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.

II.C. 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. Here 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.

1. *Setting for the Approximation.* As in Section II.A, we consider continuous paths for prices and total expenditure in some fixed time interval, but now 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 $T+1$:

$$\Delta_p \equiv \max_{i,t} \left\{ \left| \log\left(\frac{p_{i,t+1}}{p_{i,t}}\right) \right| \right\},$$

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

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 mth-order approximation of the sequence $\{g_t\}_{t=0}^{T}$, and denote this by $f_t - g_t = O(\Delta^{m+1})$, 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 $a > 0$.

For the key results presented later, 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. *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

$$\log \mathcal{P}^{b}_{t,t+1}(c) = \log \mathbb{P}_G\left(\boldsymbol{p}_t, \boldsymbol{s}_t; \boldsymbol{p}_{t+1}, \boldsymbol{s}_{t+1}\right) + O\left(\Delta^2\right),$$

$$\text{if } c \in \left\{c^b_t, c^b_{t+1}\right\}, \tag{16}$$

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.

20. As we discuss later, we can generalize this result for broader classes of index formulas defined in Section II.A. 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.

$$\log \mathcal{P}^b_{t,t+1}(c) = \log \mathbb{P}_T(\boldsymbol{p}_t, \boldsymbol{s}_t; \boldsymbol{p}_{t+1}, \boldsymbol{s}_{t+1}) + O\left(\Delta^3\right),$$

$$\text{if } c = \sqrt{c^b_t \cdot c^b_{t+1}}, \tag{17}$$

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 Online 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 $\log \mathcal{P}^b_{t_0,t}(c)$. 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 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.

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 $\mathcal{P}^b_{t,t+1}(c) \equiv \frac{\chi^b_{t+1}(c)}{\chi^b_t(c)}$ satisfies

$$\log \mathcal{P}^b_{t,t+1}(c) = \pi^+_t\left(\chi^b_t(c)\right) + O\left(\Delta^2_p\right), \tag{18}$$

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 $T^{-1} = O(\Delta)$.

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

where $\Delta \equiv \max\{\Delta_p,\ \Delta_y\}$ 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 Online Appendix B.4.

Lemma 4 offers an approximate, discretized restatement 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, we can nonparametrically estimate these functions using observed cross-sectional variations in index formulas across households.

We use Lemma 4 to construct the central contributions of this article, that is, 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.

4. *First-Order Algorithms.* We begin with our simplest algorithm that uses equation (18) to approximate the true index function $\mathcal{P}^{b}_{t,t+1}(\cdot)$ 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

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

where we have used the fact that $\chi^{b}_{t}(c^{b,n}_{t}) = y^{n}_{t}$ and that $\pi^{+}_{t}(y^{n}_{t}) = \pi^{+,n}_{t} \equiv \log \mathbb{P}_G(\boldsymbol{p}_t, \boldsymbol{s}^{n}_{t}; \boldsymbol{p}_{t+1}, \boldsymbol{s}^{n}_{t+1})$ 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 side of equation (22) as

$$\log \mathcal{P}^b_{t,t+1}\left(c^{b,n}_t\right) = \log \chi^b_{t+1}\left(c^{b,n}_t\right) - \log y^n_t$$

$$(23) \qquad = \log\left(\frac{y^n_{t+1}}{y^n_t}\right) - \frac{\partial \log \chi^b_{t+1}\left(c^{b,n}_t\right)}{\partial \log c^{b,n}_t} \log\left(\frac{c^{b,n}_{t+1}}{c^{b,n}_t}\right) + O\left(\Delta^2_y\right),$$

where we have used the definition $\mathcal{P}^b_{t,t+1}(c) \equiv \frac{\chi^b_{t+1}(c)}{\chi^b_t(c)}$ and the fact that $\chi^b_t(c^{b,n}_t) = y^n_t$.

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:

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

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 $\hat{c}^n_t$.

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 $\log \mathcal{P}_{t,t+1}(c^n_t) \approx \pi^{+,n}_t$, to nonparametrically estimate $\log \mathcal{P}_{t,t+1}(\cdot)$ 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 $\hat{c}^n_b = c^n_b = y^n_b$. Thus, we can apply a nonparametric regression of the log geometric index formula $\pi^{+,n}_b$ on real consumption $\hat{c}^n_b$ across households in this period; we recover an estimated true cost-of-living index $\log \widehat{\mathcal{P}}_{b,b+1}(\cdot)$ as a function of real consumption. We can now use the derivative of this function to evaluate the nonhomotheticity correction $\widehat{\Lambda}_{b+1}\left(\hat{c}^n_b\right)$ for each household and apply the update rule in equation (24) to find the value of real consumption for each household $\hat{c}^n_{b+1}$ in the next period. This allows us to nonparametrically estimate the true cost-of-living index $\log \widehat{\mathcal{P}}_{b+1,b+2}(\cdot)$ 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 $\hat{c}^n_{b+1}$ across households.

Applying the two steps successively moving forward from the base period, we can recover the cumulative true cost of living $\log \widehat{\mathcal{P}}_{b,t+1}(c) \equiv \sum_{\tau=b}^{t} \log \widehat{\mathcal{P}}_{\tau,\tau+1}(c)$ as a function of real consumption for each period $t \geqslant b$. 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 $t \geqslant b$, successively apply the following two steps.*

(i) *Nonparametrically fit the true price index between periods t and $t+1$: Estimate the coefficients $(\widehat{\alpha}_{k,t})_{k=0}^{K_N}$ solving the following problem:*

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

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

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

(ii) *Estimate the values of real consumption for consumers in period $t+1$: Use equation (24), where the approximate nonhomotheticity correction function is given by*

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

To apply the algorithm backward in time for $t < b$, simply relabel all the time indices in the data preceding the base period b such that $t - \tau \to t + \tau$ for all $1 \leqslant \tau \leqslant b$ 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

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).

regressions to recover the true cost-of-living index as a function 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).

i. *Sources of Error and Algorithm Refinement.* In addition to the first-order discretization error introduced in equation (22), Algorithm 1 includes two additional sources of error. First, by performing a Taylor expansion of the true cost-of-living function $\mathcal{P}_{t,t+1}(\cdot)$ 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). Although 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 $\chi_{t+1}(c_{t+1}^{n,(\ell)}) \to y_{t+1}^n$, we can write the latter's expansion as

$$\log \mathcal{P}_{t,t+1}\left(c_{t+1}^{n,(\ell)}\right) = \log y_{t+1}^n - \log \chi_t\left(c_{t+1}^{n,(\ell)}\right)$$

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

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

The second source of error stems from the nonparametric estimation step. In particular, Algorithm 1 makes the simplifying choice to directly estimate the true cost-of-living index $\mathcal{P}_{t,t+1}(\cdot)$ 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. Online Appendix Algorithm A.1 combines this strategy with the Taylor series approximation in equation (28) to provide a refinement of our first-order algorithm. Online Appendix A.1.2 further discusses the sources of approximation error in this algorithm.

5. *Second-Order Algorithms.* In parallel to the approach just laid out, 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). Online Appendix Algorithm A.3 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 $\frac{c_{t+1}}{c_t}$. In contrast, our second-order algorithm also 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. Online Appendix Algorithm A.4 provides a refinement of the second-order algorithm, along the same lines as we discussed above.

6. *Other Extensions.* We discuss three additional extensions of our baseline and refined algorithms.

i. *Alternative Algorithms.* 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. An alternative approach is to approximate the nonhomotheticity correction by nonparametrically estimating the elasticity of the inverse mapping, $\widetilde{\chi}_t^b(y)$. Online Appendix Algorithm A.2 and Algorithm A.5 provide first- and second-order schemes based on this alternative approach. In Section II.D below, we provide a comparison of the approximation errors among all of our alternative algorithms.

Finally, as another example, subsequent to our work Baqaee, Burstein, and Koike-Mori (2024) have presented a different alternative to our benchmark first-order algorithm. In Online Appendix B.3, we establish the tight theoretical connection between their approach and ours. We provide evidence using synthetic and real-world data that in practice their approach leads to results that are similar to those produced by Algorithm 1.[23]

ii. *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.

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 II satisfy

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

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

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

Proof. See Online 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, whereas 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

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

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 III where, due to data limitations, the most natural choice for a second-order index is the Fisher index.

iii. *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 (i) 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 discuss in Section IV, such a generalization leads to further corrections in our standard measures of real consumption, beyond the nonhomotheticity correction, to account for the effect of potential changes in consumer characteristics on consumer welfare over time.[24]

7. *Discussion.* As discussed already, 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]

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

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

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.

II.D. Simulation

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, Lashkari, and Mestieri (2021) have shown that the nhCES preferences lead to a demand system compatible with the cross-sectional relationship between household income and the composition of expenditure among three main sectors of the economy: agriculture, manufacturing, and services. Following their

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.

specification, we assume that the expenditure function satisfies:

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

We use the same parameters as in Comin, Lashkari, and Mestieri (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, Blundell, and Lewbel 2009). We consider a horizon of 70 years and assume that over this horizon nominal expenditure grows at the constant rate of 4.48% a year, in line with the U.S. data for 1953–2019. In each case 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 U.S. 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 United States over 1953–2019. We 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 1 percentage point higher in services and 1 percentage point 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.

28. The corresponding shares in the United States based on the Bureau of Labor Statistics data are 0.14, 0.27, and 0.59 for Agriculture, Manufacturing, and Services, respectively.

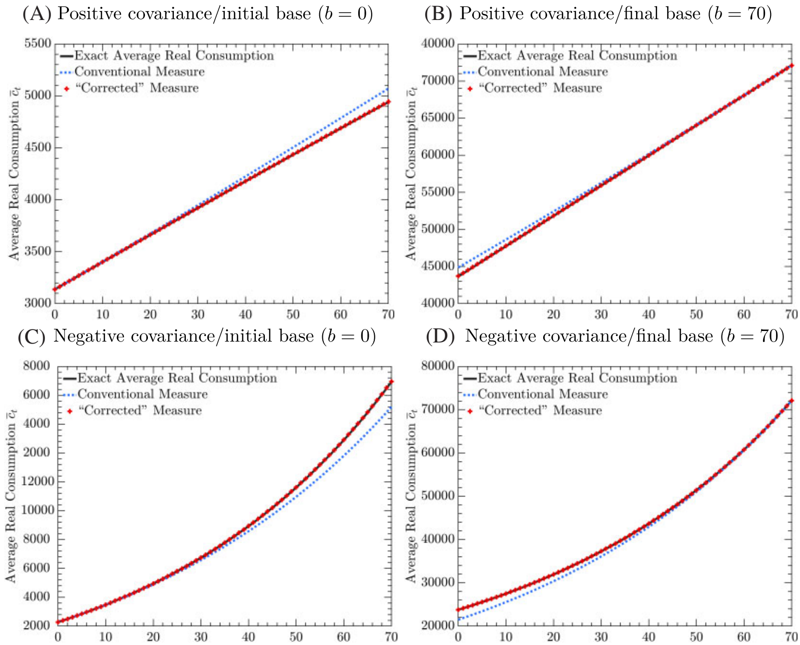

FIGURE I

Illustrative Simulation of the Evolution of Average Real Consumption

The figures compare the evolution of the true value of average real consumption with two different approaches to approximating this value: (i) the average of the uncorrected nominal real consumption growth deflated by household-specific geometric price indices, and (ii) 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.

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.

Figure I reports 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 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 Online 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 uncorrected measures of real consumption growth. The Online 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.

III. Empirics

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

III.A. Data and Descriptive Statistics

1. *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 CEX. These items, called universal classification codes (UCCs), are defined by the Bureau of Labor Statistics (BLS) and cover the entire consumption basket of households in the United States. 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. Online 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 sociodemographic characteristics at the household level. We then aggregate the household-level data to the level of pretax income percentiles. We thus obtain expenditure patterns that vary across income percentiles, which we use to compute the income elasticity of inflation. We also use this data set 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 reweight 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, Saez, and Zucman 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 sociodemographic groups are no longer available, except in 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 nonhomotheticicities to extrapolate expenditure shares back to 1955, and we obtain the growth rate of aggregate consumption expenditures from the BEA.[30] Given the data limitations before 1984, we present two sets of results, first focusing on the period from 1984 to 2019 for which high-quality CEX data are available annually, and then a longer historical analysis going back to 1955.

2. *Descriptive Inflation Statistics.* This new linked data set 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. Figure II, Panels A and B 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

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

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

(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

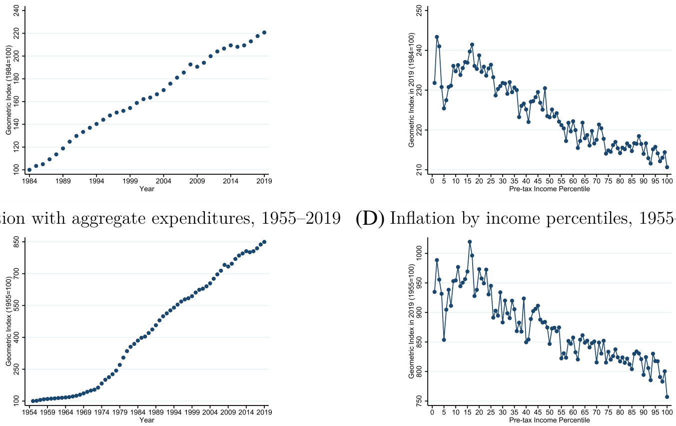

FIGURE II

Descriptive Inflation Statistics

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 pretax income percentiles. In this panel, price indices are built using expenditure shares that are specific to each pretax income percentile. Panels C and D repeat the analysis for a longer period, from 1955 to 2019. All panels use chained geometric price indices.

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). Although this literature focuses on post-2000 patterns, our data show 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 article is the first to build a data set 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]

III.B. Main Estimates

1. *Analysis from 1984 to 2019.* We first implement Algorithm 1 using our main data set 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 II, the negative covariance between household income and price indices shown in Figure II 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 Figure III, Panel A, 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]

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

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

32. Explaining these patterns of inflation inequality falls beyond the scope of this article, 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).

33. Algorithm 1 is implemented using each pretax 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 $K \geqslant 1$. 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 note 5).

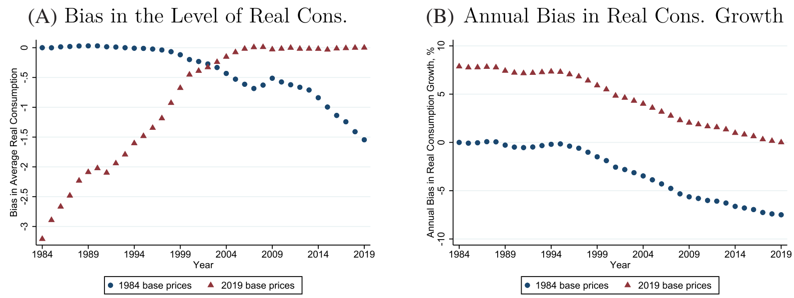

FIGURE III

Nonhomotheticity Correction and Bias in Average Real Consumption, 1984–2019

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 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 data set at the level of pretax income percentiles, using geometric price indices. We then average percentile-level results to obtain average real consumption per household.

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 Figure III, Panel A, 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, because 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 because the level of real consumption is underestimated in all past periods. Panel B of Figure III reports these results, expressing the size of the bias as a share of measured growth.[34] With 1984 prices as the base, the conventional measure underestimates real consumption growth by about 7.5% in 2019. Taking 2019 prices as the base, the conventional measure overestimates real consumption growth by approximately 7.5% in 1984.

It is also instructive to examine the disaggregated patterns for the nonhomotheticity correction across pretax income percentiles. Figure IV plots these results. Panel A reports 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.[35] 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(i) and B(ii) 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

34. 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 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.

35. The biases are expressed as a share of measured growth, as given by λ_t^n defined in note 34 for each percentile n in 2019.

Panel A: Percentile-specific biases for annual growth in real consumption

(i) In 2019 with 1984 prices as base

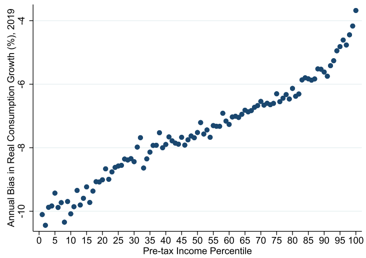

(ii) In 1984 with 2019 prices as base

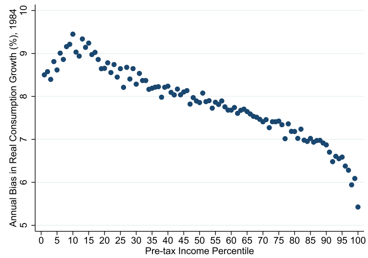

Panel B: Percentile-specific biases for the level of real consumption

(i) In 2019 with 1984 prices as base

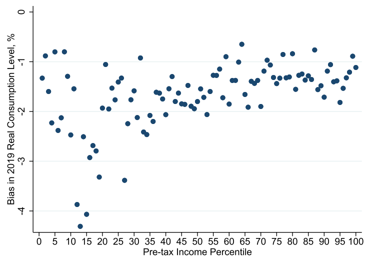

(ii) In 1984 with 2019 prices as base

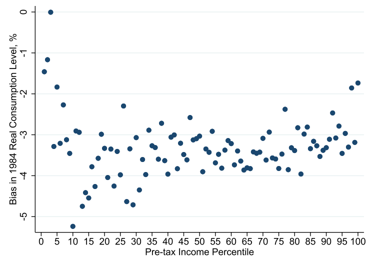

FIGURE IV

Nonhomotheticity Correction and Biases in Real Consumption by Income Percentiles

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.

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, $\frac{y_t^n}{c_t^n}$: the correction is similar in magnitude for all pretax income percentiles. To assess the quantitative relevance of the nonhomotheticity correction, it is instructive to compare its size with 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.

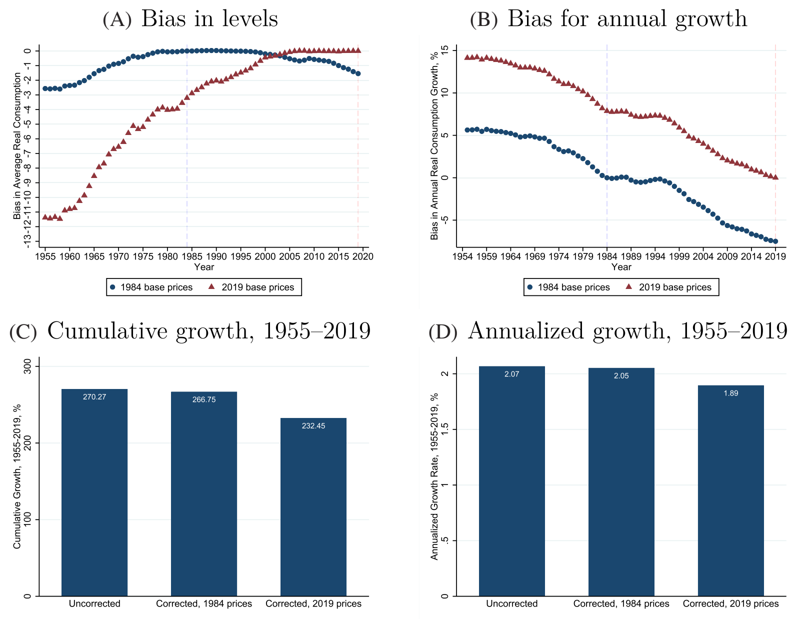

FIGURE V

Nonhomotheticity Correction and Bias in Average Real Consumption, 1984–2019

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 pretax 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.

2. *Analysis from 1955 to 2019*. Next we extend the analysis back to 1955, reporting the results in Figure V.[36] Panel A reports the bias in levels; the patterns are identical to Figure III 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

36. As explained in Online Appendix D, due to data limitations (i) we assume the expenditure shares observed in 1960 remain constant for 1955–1960, and (ii) we interpolate expenditure shares between 1960 and 1972, and between 1972 and 1984.

nonhomotheticity correction, simply because the two biases in levels in 2019 and 1955 turn out to be of the same magnitude.[37]

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, Figure V, Panel B 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, Figure V, Panel C 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%. With 2019 prices as the base, the difference becomes large: cumulative consumption growth falls to 232%, or an annualized growth rate of 1.89% a year. Intuitively, from today's perspective, consumer welfare in the past was higher than

37. 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 V, Panel A, the nonhomotheticity correction changes sign before and after the base period (see Figure V, Panel B), and thus cancels out when the base period is somewhere in the middle of two periods under consideration. However, 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.

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, 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.

III.C. Sensitivity Analysis

We 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 data set as in our baseline specifications. We build alternative data sets to assess the stability of the results depending on data construction choices and the level of aggregation of expenditure data.[38]

1. *Alternative Algorithms, Indices, and Controls*. We implement several sensitivity tests using the same data sets 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 examine whether the results change when we use Algorithm A.3, which implements a second-order approximation. The results are shown in Figure VI, Panel A: the patterns remain unchanged with the Fisher index as well as with the algorithm providing a second-order approximation.

38. 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).

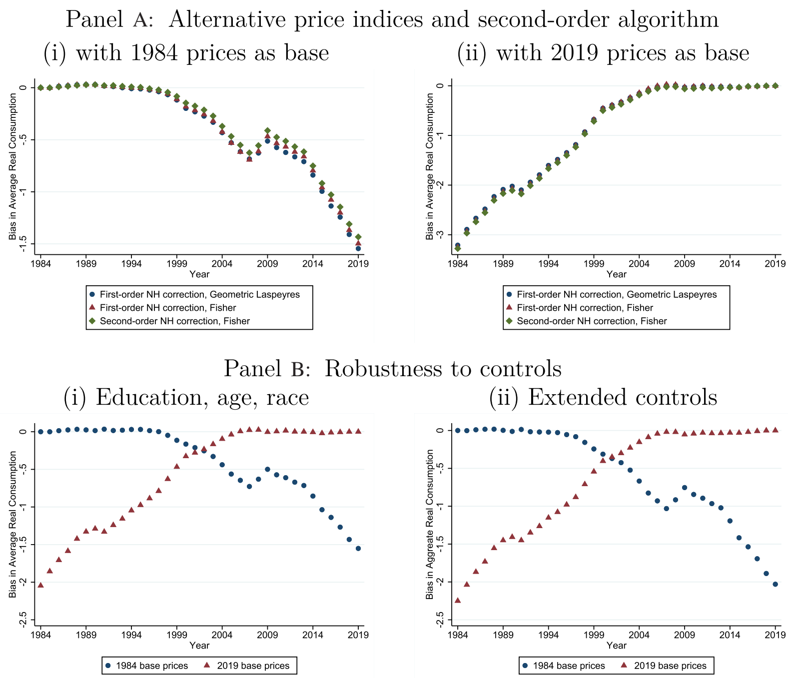

FIGURE VI

Sensitivity Analysis

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, and 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 versus rural area, gender, and city population size, in addition to education, age, and race.

Next we assess whether the patterns remain similar when including controls. We implement Algorithm 1 as in Section III.B, 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 Figure VI, Panel B. We introduce additional controls for region (Midwest, Northeast, West, South), rural versus urban area, gender, and city population size. 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.[39]

2. *Sensitivity Analysis with Alternative Data sets*. To assess the sensitivity of our findings to data construction choices, we build and study four alternative data sets.[40]

To document whether our results are sensitive to aggregation choices, we build two alternative data sets that closely follow our main data set but use different levels of aggregation, grouping UCCs into broader categories. First, we create a version of the data set 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 data set. The results are very similar to those obtained with our main data set, with slightly smaller magnitudes due to the higher level of aggregation.[41]

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

39. 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 $K \geqslant 1$. Online 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 Online Appendix Figure E.7. Third, we consider a specification controlling for state fixed effects and obtain similar results (Online Appendix Figure E.8). Fourth, we apply the algorithm of Baqaee, Burstein, and Koike-Mori (2024) to our data (Online 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.

40. Online Appendix D provides a complete description of the data construction steps.

41. 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).

are almost indistinguishable from those obtained with the data set in our main analysis.

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 are available, using NielsenIQ 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, in 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 NielsenIQ 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.15. The magnitude of the annual bias in real consumption growth increases faster than in our alternative data sets reaching 3% of the uncorrected measure after only a decade.[42]

42. To provide a precise comparison of the magnitude of the biases obtained with the NielsenIQ data, we repeat the analysis with our main CEX-CPI data set restricted to the product categories covered in the NielsenIQ data between 2004 and 2014. The restricted CPI-CEX sample covers 44 UCC items belonging to the following categories: alcoholic beverages; 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 Online 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 NielsenIQ data, and −0.016% with the CEX-CPI sample; the bias in annual real consumption growth in 2014 is −2.83% with the NielsenIQ data, and −0.78% with the CEX-CPI sample. Thus, the biases are about 3.5 times larger with the detailed NielsenIQ 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 NielsenIQ data. We account for the welfare effects of changes in product variety using a CES price index, which we compute for each of the 9,131 NielsenIQ 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

Finally, we implement a robustness test inspired by the distributional national accounts of Piketty, Saez, and Zucman (2018): we discipline our household-level data such that aggregate expenditure shares match exactly the official CPI consumption weights used by the 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.[43] 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.[44] 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 data set, 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 with 11.4% in the baseline specification.

Overall, these robustness checks show that the findings obtained with our baseline data set 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.

IV. Measuring Welfare Changes with Observed Heterogeneity

We extend the results of Section II.B 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

consumption growth in 2014 is −7.99% (Online Appendix Figure E.14). Thus, compared with our baseline NielsenIQ estimates, the biases are about three times larger when accounting for new goods.

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

44. See Online Appendix D for a detailed description of this step.

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.

IV.A. Correction for Changes 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.[45] We assume that consumer preferences are characterized by a well-behaved utility function $u = U(\boldsymbol{q}; \boldsymbol{x})$ that depends on the consumer's 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 $0 \leqslant b \leqslant T$), 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 because they have distinct preferences. 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, that is, $\boldsymbol{x}_t \equiv \boldsymbol{x}_{t_0}$, then Definition 3 reduces to Definition 1, under homogeneous preferences. Second, if prices do not change, that is, $\boldsymbol{p}_t \equiv \boldsymbol{p}_{t_0}$,

45. 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.

the growth in real consumption simply accounts for the growth in nominal expenditure even if consumer characteristics change, $\frac{c_t^b}{c_{t_0}^b} \equiv \frac{y_t}{y_{t_0}}$.

In parallel to the definitions introduced in Section II.A, we denote by $\chi_t^b(c;\boldsymbol{x}) \equiv E(M_b^{-1}(c;\boldsymbol{x});\boldsymbol{x})$ 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 $\boldsymbol{x}$:

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

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

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

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

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

Proof. See Online Appendix B.4.

Proposition 3 extends the 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 $\mathcal{P}^b_{t_0,t}(c;\boldsymbol{x})$ under characteristic-dependent preferences:

$$\mathcal{P}^b_{t_0,t}(c;\boldsymbol{x}) \equiv \frac{\chi^b_t(c;\boldsymbol{x})}{\chi^b_{t_0}(c;\boldsymbol{x})}, \tag{34}$$

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 $\log \mathcal{P}^b_{t_0,t}(c;\boldsymbol{x}) = \int_{t_0}^{t} \log D_\tau(\chi^b_\tau(c;\boldsymbol{x});\boldsymbol{x})\, d\tau$. 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 need the characteristic correction function index $\Gamma^b_{d,t} \equiv \frac{\partial \log \chi^b_t}{\partial \log x} \equiv \frac{\partial \log \mathcal{P}^b_{b,t}}{\partial \log x}$, 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 $\zeta_{i,d,t}(c;\boldsymbol{x}) \equiv \frac{\partial \log \omega_{i,\tau}(\chi^b_\tau(c);\boldsymbol{x})}{\partial \log x_d}$ 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, $\frac{\partial \log \mathcal{P}^{b}_{b,t}(c_t;\boldsymbol{x}_t)}{\partial age_t}\frac{dage_t}{dt}$, 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.

IV.B. 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 IV.A. Online Appendix Algorithms A.6 and A.7 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. Then we use the estimated relationship with total expenditure and with other characteristics to approximate the corresponding correction functions.

IV.C. Application to the Measurement of Real Consumption in the United States with Consumer Aging

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

1. *Data and Summary Statistics*. To study the effect of consumer aging on real consumption growth, we build another version of the data set 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 10 deciles of the (pretax) income distribution and, within each income decile,

(A) Inflation by Age Decile, 1955–2019 (B) Bias in the Level of Real Cons.

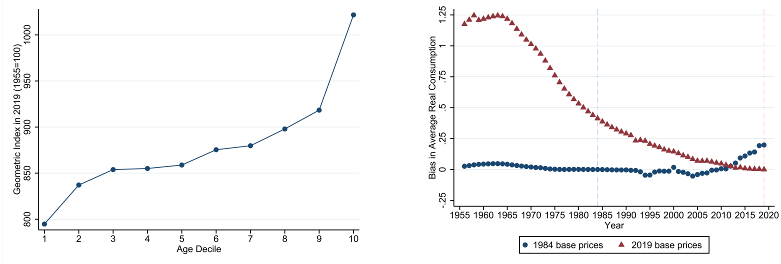

FIGURE VII

Consumer Aging and Real Consumption

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 data set at the level of "age decile by income decile" units, using geometric Laspeyres price indices. We average the results to obtain average real consumption per household with the aging correction.

we compute 10 age deciles. We then compute average age in each of these cells.[46]

Using this data set, we compute inflation rates across age groups and find higher inflation rates for older households, as shown in Figure VII, Panel A. 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 article 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 United States, especially from 1970 onward. Therefore, by the logic of Section IV.A, conventional

46. As in our main data set, 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.

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

2. *Aging Correction for Average Real Consumption*. We apply Algorithm A.6 to quantify the adjustment to average real consumption implied by consumer aging. Figure VII, Panel B reports the results. Specifically, we report the deviation in the level of average real consumption when accounting for aging and nonhomotheticities, relative to the benchmark measure with only the nonhomotheticity correction.[47]

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.[48]

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. Although there is a strong relationship between age and inflation, the correction to average real consumption implied by aging is

47. In the data set 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 III with our baseline data set using income percentiles.

48. To understand the difference in the magnitude of the aging correction depending on the choice of base years, 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.

smaller than the nonhomotheticity correction primarily because the change in average household age over time is relatively slow.

V. Conclusion

In this article, 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 postwar 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 indices across countries with preference heterogeneity, is a promising direction for future research.

Our results may have important implications for how national statistical agencies around the world construct measures of real economic value. The approach suggested here has the potential to be widely adopted for at least three reasons. First, it 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 constructing homothetic price indices. In Section III, we offered a blueprint for how the BLS can use data that are already available to construct improved measures of real consumption growth and inequality in the United States. Second, our approach has a light computational burden. In its first-order renditions, our algorithms simply require one cross-sectional regression 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 are available. Depending on their goals, different data users may opt to rely on data expressed in terms of 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.

London School of Economics, United Kingdom
Federal Reserve Bank of New York, United States

Supplementary Material

An Online Appendix for this article can be found at *The Quarterly Journal of Economics* online.

Data Availability

The data underlying this article are available in the Harvard Dataverse, https://doi.org/10.7910/DVN/IAKREC. (Jaravel and Lashkari 2023).

References

Argente, David, and Munseob Lee, "Cost of Living Inequality during the Great Recession," *Journal of the European Economic Association*, 19 (2021), 913–952. https://doi.org/10.1093/jeea/jvaa018

Atkin, David, Benjamin Faber, Thibault Fally, and Marco Gonzalez-Navarro, "Measuring Welfare and Inequality with Incomplete Price Information," *Quarterly Journal of Economics*, 139 (2024), 419–475. https://doi.org/10.1093/qje/qjad037

Baqaee, David, and Ariel Burstein, "Welfare and Output with Income Effects and Taste Shocks," *Quarterly Journal of Economics*, 138 (2023), 769–834. https://doi.org/10.1093/qje/qjac042

Baqaee, David R., Ariel T. Burstein, and Yasutaka Koike-Mori, "Measuring Welfare by Matching Households across Time," *Quarterly Journal of Economics*, 139 (2024), 533–573. https://doi.org/10.1093/qje/qjad038

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

Blackorby, Charles, and David Donaldson, "Money Metric Utility: A Harmless Normalization?" *Journal of Economic Theory*, 46 (1988), 120–129. https://doi.org/10.1016/0022-0531(88)90152-4.

Blundell, Richard W., Martin Browning, and Ian A. Crawford, "Nonparametric Engel Curves and Revealed Preference," *Econometrica*, 71 (2003), 205–240.

Bosmans, Kristof, Koen Decancq, and Erwin Ooghe, "Who's Afraid of Aggregating Money Metrics?," *Theoretical Economics*, 13 (2018), 467–484.

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

Comin, Diego, Danial Lashkari, and Marti Mestieri, "Structural Change with Long-Run Income and Price Effects," *Econometrica*, 89 (2021), 311–374.

Deaton, Angus, and John Muellbauer, "An Almost Ideal Demand System," *American Economic Review*, 70 (1980), 312–326.

Diewert, W. Erwin, "Exact and Superlative Index Numbers," *Journal of Econometrics*, 4 (1976), 115–145.

———, "Superlative Index Numbers and Consistency in Aggregation," *Econometrica*, 46 (1978), 883–900. https://doi.org/10.2307/1909755.

———, "The Economic Theory of Index Numbers: A Survey." In *Essays in Index Number Theory*, Walter E. Diewert and Alice O. Nakamura, eds. (Bingley, UK: Emerald Group Publishing, (1993), 163–208.

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

Fajgelbaum, Pablo D., and Amit K. Khandelwal, "Measuring the Unequal Gains from Trade," *Quarterly Journal of Economics*, 131 (2016), 1113–1180. https://doi.org/10.1093/qje/qjw013

Feenstra, Robert C., "New Product Varieties and the Measurement of International Prices," *American Economic Review*, 84 (1994), 157–177.

Feenstra, Robert C., Hong Ma, and D. S. Prasada Rao, "Consistent Comparisons of Real Incomes across Time and Space," *Macroeconomic Dynamics*, 13 (2009), 169–193.

Hobijn, Bart, and David Lagakos, "Inflation Inequality in the United States," *Review of Income and Wealth*, 51 (2005), 581–606. https://doi.org/10.1111/j.1475-4991.2005.00170.x

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

Jaravel, Xavier, "The Unequal Gains from Product Innovations: Evidence from the US Retail Sector," *Quarterly Journal of Economics*, 134 (2019), 715–783. https://doi.org/10.1093/qje/qjy031

———, "Inflation Inequality: Measurement, Causes, and Policy Implications," *Annual Review of Economics*, 13 (2021), 599–629. https://doi.org/10.1146/annurev-economics-091520-082042

Jaravel, Xavier, and Danial Lashkari, “Replication Data for: ‘Measuring Growth in Consumer Welfare with Income-Dependent Preferences: Nonparametric Methods and Estimates for the United States’,” (2023). Harvard Dataverse, https://doi.org/10.7910/DVN/IAKREC.

Kaplan, Greg, and Sam Schulhofer-Wohl, “Inflation at the Household Level,” *Journal of Monetary Economics*, 91 (2017), 19–38. https://doi.org/10.1016/j.jmoneco.2017.08.002

Klick, Josh, and Anya Stockburger, “Experimental CPI for Lower and Higher Income Households,” Bureau of Labor Statistics Working Paper no. WP-537, 2021.

McGranahan, Leslie, and Anna Paulson, “Constructing the Chicago Fed Income Based Economic Index–Consumer Price Index: Inflation Experiences by Demographic Group: 1983–2005,” Federal Reserve Bank of Chicago Working Paper no. 2005–20, 2005.

Newey, Whitney K., “Convergence Rates and Asymptotic Normality for Series Estimators,” *Journal of Econometrics*, 79 (1997), 147–168. https://doi.org/10.1016/S0304-4076(97)00011-0

Oberfield, Ezra, “Inequality and Measured Growth,” Princeton University Working Paper, 2022.

Oulton, Nicholas, “Chain Indices of the Cost-of-Living and the Path-Dependence Problem: An Empirical Solution,” *Journal of Econometrics*, 144 (2008), 306–324.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman, “Distributional National Accounts: Methods and Estimates for the United States,” *Quarterly Journal of Economics*, 133 (2018), 553–609. https://doi.org/10.1093/qje/qjx043

Pollak, Robert A., “The Theory of the Cost-of-Living Index,” In *Price Level Measurement*, W. Erwin Diewert, ed, (Amsterdam: Elsevier North-Holland, (1990), 5–77.

Redding, Stephen J., and David E. Weinstein, “Measuring Aggregate Price Indices with Taste Shocks: Theory and Evidence for CES Preferences,” *Quarterly Journal of Economics*, 135 (2020), 503–560. https://doi.org/10.1093/qje/qjz031

Samuelson, Paul A., and S. Swamy, “Invariant Economic Index Numbers and Canonical Duality: Survey and Synthesis,” *American Economic Review*, 64 (1974), 566–593.

Slesnick, Daniel T., “Aggregate Deadweight Loss and Money Metric Social Welfare,” *International Economic Review*, 32 (1991), 123–146.