

A Nonhomothetic Price Index and Cost-of-Living Inequality*

Philipp Hochmuth[†] Markus Pettersson[‡] Christoffer Jessen Weissert[§]

September 21, 2023

Abstract

We derive a price index based on nonhomothetic preferences and use it to document cost-of-living inequality in the United States. Our framework generalizes all known superlative price indices and admits heterogeneous indices across the distribution of household consumption expenditures which aggregate consistently into welfare-relevant group indices. When necessities and luxuries are separable in the expenditure function, this generalization avoids the need to estimate a complete demand system. Using CEX-CPI data for the period 1995–2020, we find no differences in average inflation rates across the expenditure distribution, but 2.5 times higher inflation volatility for the bottom decile compared to the top decile, stemming from a larger exposure to food, gasoline, and utilities. These results contrast with the inequality found by the typical approach of constructing separate homothetic price indices for different consumer groups, and our analysis suggests that the difference follows from an income-effect bias in the homothetic group-specific indices.

Keywords: cost of living, inequality, nonhomotheticity, superlative price index.

JEL Codes: C43, D11, D12, E31, I30.

* For many helpful suggestions and discussions, we thank Per Krusell, Timo Boppart, Kurt Mitman, John Hassler, Kjetil Storesletten, Vincent Sterk, Morten Ravn, Lars Ljungqvist, Jeppe Druedahl, Søren Hove Ravn, Jessie Handbury, Thomas Jørgensen, Mikkel Plagborg-Møller, Tobias Renkin, Kieran Larkin, and seminar participants at the 14th Nordic Summer Symposium in Macroeconomics, the 2022 Annual Meeting of the Austrian Economic Association, the 7th Swedish Conference in Economics, the IIES Macro Group, the 2021 Stockholm-Uppsala Doctoral Student Workshop in Economics, and the 9th New Developments in Business Cycle Analysis Conference (*The Return of Inflation*) at Norges Bank.

[†] Institute for International Economic Studies, Stockholm University. Email: philipp.hochmuth@iies.su.se.

[‡] Department of Economics, Stockholm University. Email: markus.pettersson@su.se.

[§] Department of Economics and Monetary Policy, Danmarks Nationalbank. Email: cjw@nationalbanken.dk.

1 Introduction

Do changes in the cost of living vary with income and if so, how much? These questions are important for areas in which price changes and inequality matter, not least for the measurement of real incomes and monetary policy. One of the oldest empirical economic facts, dating back to Engel (1857), answers the first question in the affirmative: consumption patterns, and consequently cost-of-living changes, differ systematically between rich and poor individuals. More uncertainty surrounds the magnitude and its drivers, partly due to a lack of tools: inflation measures based on conventional cost-of-living indices cannot provide an answer because they assume identical consumption behavior across individuals, contrary to reality, while data limitations typically prohibit the estimation of flexible demand models. Current workarounds to these shortcomings resort to computing standard price indices separately for different income groups or rely on nonparametric algorithms. How to combine the practical simplicity of conventional cost-of-living index formulas with realistic consumer behavior in a theoretically consistent way, however, remains an unresolved issue.

The goal of this paper is to tackle this problem head-on. In doing so, we make two key contributions. The first one is theoretical and methodological: we derive a cost-of-living index that is consistent with nonhomothetic demand theory and outline a feasible strategy to compute this index without having to estimate an entire consumer demand system. In its most general form, this index nests all known superlative price indices as special cases, including the Fisher (1922), the Törnqvist (1936), and the Sato (1976)-Vartia (1976) indices. We show that when preferences are such that commodities can be divided into a necessity bundle and a luxury bundle, only two parameters need to be estimated: the expenditure elasticity of demand for necessities and a parameter that controls the elasticity of substitution between the two bundles. The second contribution is empirical: we apply our framework to US data and find that while consumption-poor households experienced sharper increases in the cost of living between 1995 and 2014, the gap to the rich disappears between 2015 and 2020, leaving households equally affected on average.

Our framework allows for a theoretically consistent identification of the true cost of living across the full distribution of individual consumption expenditures. We achieve this by deriving cost-of-living indices from a specification of Muellbauer’s (1975, 1976) “price independent generalized linearity” (PIGL) preferences.¹ These preferences are nonhomothetic, so that richer households allocate higher weights to price changes of luxury goods. Changes in the cost of living are consequently allowed to differ with the consumption expenditure level. We show that the corresponding cost-of-living indices can be written as weighted geometric means of price changes of individual commodities, which enables straightforward decompositions to identify the drivers of the overall change in the cost of living. PIGL preferences also maintain tractable aggregation properties that eliminate the conflict between heterogeneous indices at the individual level and the existence of welfare-relevant group cost-of-living indices with an equivalent functional form. This is important because grouping of households is typically required in practical applications.

¹ PIGL preferences have gained recent popularity in the structural change literature, see for instance Boppart (2014), León-Ledesma and Moro (2020), Alder, Boppart and Müller (2022), Cravino, Levchenko and Rojas (2022), and Fan, Peters and Zilibotti (2023).

To implement the cost-of-living index empirically, we present an approach that avoids the curse of dimensionality associated with flexible nonhomothetic demand models, for which the number of parameters to estimate typically explodes with the number of goods considered. Specifically, we consider preferences that are quasi-separable between two bundles of commodities: necessities and luxuries.² This structure is straightforward to implement by investigating Engel curve slopes and classifying goods as necessities or luxuries accordingly. Under quasi-separability, nonhomotheticity runs between the bundles whereas behavior within them remains homothetic, so price indices for each bundle can be obtained with standard homothetic cost-of-living indices. The PIGL cost-of-living index consequently reduces to observed expenditure shares and prices together with the two elasticity parameters above. We still nest homothetic cost-of-living indices as special cases, so quasi-separability is not a strong restriction from a price index theoretical point of view. Comparing our approach to the estimation of an equivalent demand model that does not invoke quasi-separability, we also find that this preference structure is justified in the data.³

In an empirical analysis of US inflation inequality, we apply our methodology to matched Consumer Price Index and Consumer Expenditure Survey data for 21 nondurable commodity groups and obtain three key results. First, between 1995 and 2020, we find no economically significant differences in average annual inflation rates between consumption-poor and consumption-rich households. Second, while average inflation *rates* are similar, inflation *volatility* is much higher for consumption-poor households: the bottom expenditure decile faces a 2.5 times more volatile inflation rate than the top decile, as measured by the standard deviation of annual inflation rates. This volatility implies that substantial differences arise in the short run. For instance, between 2004 and 2015, the time around the Great Recession, the inflation rate for the bottom expenditure decile exceeds that of the top decile by on average 0.37 percentage points per year. Lastly, decomposing the cost-of-living changes, we find that these disparities primarily stem from a larger exposure to price changes in food, gasoline, and utilities for poorer households.

Our methodology contrasts with a strand of literature on the measurement of inflation inequality which dates back at least to the 1950s (see Muellbauer, 1974, and references therein), with recent advances surveyed by Jaravel (2021). This literature approximates nonhomotheticity by computing homothetic cost-of-living indices separately for different income groups based on each group’s observed consumption behavior over time.⁴ Such an approach, however, conflicts with Konüs’s (1939) definition of a cost-of-living index as the percentage change in expenditures needed to maintain a *fixed* utility level when prices change. By holding utility constant, a true cost-of-living index only reflects price changes and subsequent substitution effects. Yet, it is reasonable to expect that *observed* consumption changes include not only substitution effects, but also income effects related to changes in real consumption expenditures (or rather,

² While direct separability groups quantities of commodities in the utility function, quasi-separability groups prices of commodities in the expenditure function. This concept that was first introduced by Gorman (1996) in an unpublished manuscript from the 1970s.

³ Our approach also applies to other popular demand systems such as the quasi-homothetic or the nonhomothetic CES considered by Matsuyama (2019) and Comin, Lashkari and Mestieri (2021). Our PIGL parametrization nests quasi-homothetic preferences and a particular case of nonhomothetic CES preferences as special cases. In general, neither individual expenditure shares nor individual cost-of-living indices aggregate consistently into group-level equivalents for other nonhomothetic CES parametrizations, thus motivating our PIGL choice.

⁴ Recent papers employing this method include Hobijn and Lagakos (2005), McGranahan and Paulson (2005), Broda and Romalis (2009), Kaplan and Schulhofer-Wohl (2017), Jaravel (2019), Argente and Lee (2021), Klick and Stockburger (2021), Orchard (2022), and Lauper and Mangiante (2023).

utility). These income effects drive a potential wedge between the true cost-of-living index and its constructed approximation (see Oulton, 2008, for further discussion). Our theoretically consistent approach avoids this bias.

We evaluate this income effect bias by computing group-specific homothetic price indices and compare those to our baseline results, and this exercise points to a considerable bias in the inflation inequality obtained with group-specific indices. Over the full sample period, the group-specific indices exhibit an average inflation rate difference between the top and bottom consumption deciles that is 0.3 percentage points larger than in our PIGL framework, thus suggesting a bias of a similar size. Interestingly, recent papers on US inflation inequality that compute group-specific price indices suggest that annual inflation rates for low-income households exceed those of high-income households by 0.3 to 0.4 percentage points in recent years, comparable to the bias that we find here.

Besides the empirical contribution, this paper adds to a long-standing literature on the economic approach to price index theory following Konüs (1939), Samuelson and Swamy (1974), Diewert (1976, 1978), Feenstra (1994), Feenstra and Reinsdorf (2000), Oulton (2008, 2012), Redding and Weinstein (2020), and many others, whereby cost-of-living indices are derived from consumer theory via the expenditure function.⁵ Nonhomothetic behavior has received limited attention within this line of research, which rather centers around Diewert’s notion of a *superlative* price index. That is, a price index that corresponds to a homothetic expenditure function that approximates any other homothetic expenditure function to the second order. The Fisher, the Törnqvist and the Sato-Vartia indices are all known to satisfy this property; see Diewert (1976) for the former two and Barnett and Choi (2008) for the latter. Our paper provides a nonhomothetic generalization of all currently known indices within this class.

Rather than measuring welfare changes with exact nonhomothetic cost-of-living index formulas, Baqaee, Burstein and Koike-Mori (forthcoming) and Jaravel and Lashkari (forthcoming) recently proposed numerical approximations of arbitrary preferences that exploit the cross-sectional variation in consumer expenditure survey data. Relative to these papers, the parametric PIGL cost-of-living index provides three advantages. First, our index features a natural decomposition that identifies the contributions of individual commodities to the overall change in the cost of living. Second, while we also use cross-sectional microdata in this paper, we are by no means restricted to this choice: our index can measure aggregate welfare changes using macro data alone, and the PIGL preference parameters can be estimated without aggregation bias from the same data. Third, we only need cross-sectional data for a single reference period to construct a distribution of cost-of-living indices, which allows for timely measures of inflation inequality when cross-sectional microdata only become available with a time lag. Taken together, the PIGL index is therefore applicable when users and producers of inflation statistics (i) want to know where total price changes come from and (ii) have access to aggregate time series but only a few years of distributional data, and this is arguably the relevant case for most countries.⁶

For the remainder of the paper, we proceed as follows. Section 2 covers the theoretical frame-

⁵ For a survey of the early stages of this literature, see Diewert (1993).

⁶ We fully exploit these features in a follow-up paper that constructs a distribution of nonhomothetic Personal Consumption Expenditures price indices and analyzes US inflation inequality in the recent cost-of-living crisis (see Hochmuth, Pettersson and Weissert, 2023).

work, derives the nonhomothetic price index, and shows that it generalizes all superlative price indices. [Section 3](#) outlines our empirical strategy and discusses the assumption we make to simplify the estimation procedure. [Section 4](#) explains the data we employ, classifies goods into necessities and luxuries, and reports estimates for the two preference parameters. [Section 5](#) reports the main empirical results, while [Sections 6](#) and [7](#) compare the price index with previous methods to measure heterogeneous changes in the cost of living. [Section 8](#) concludes.

2 A Nonhomothetic Cost-of-Living Index

The framework we consider is one in which consumers maximize utility over a set of goods J with a corresponding price vector \mathbf{p} and in which we wish to investigate the change in the cost of living between a period t and some base period s . To familiarize readers with the setting, we first review known price index definitions and results and outline our specification of consumer preferences before presenting the cost-of-living index that we derive. We drop time subscripts throughout whenever possible to simplify notation, as long as this causes no confusion.

2.1 The Cost-of-Living Index and its Homothetic Case

The minimum consumption expenditure e required to obtain some utility level u when faced by the price vector \mathbf{p} is given by the expenditure function $e = c(u, \mathbf{p})$. Following Konüs (1939), we define a cost-of-living index in period t relative to base period s to be the ratio of minimum expenditures required to maintain a constant utility level:

$$P(u, \mathbf{p}_t, \mathbf{p}_s) \equiv \frac{c(u, \mathbf{p}_t)}{c(u, \mathbf{p}_s)}. \quad (1)$$

Hereinafter we typically leave the arguments of the cost-of-living index implicit and simply write $P_t = c(u, \mathbf{p}_t)/c(u, \mathbf{p}_s)$.

In general, the Konüs cost-of-living index (1) depends on the reference standard of living u as well as the prices in the two periods. Samuelson and Swamy (1974) show that independence of u occurs if and only if we consider the special case of homothetic preferences. Suppose for instance that consumer preferences are characterized by an indirect utility function of the standard homothetic CRRA form,

$$V(e, \mathbf{p}) = \frac{1}{\varepsilon} \left[\left(\frac{e}{B(\mathbf{p})} \right)^\varepsilon - 1 \right], \quad (2)$$

where $B(\mathbf{p})$ is a linearly homogenous function of prices and ε is the coefficient of relative risk aversion. Inverting the utility function to obtain the expenditure function $c(u, \mathbf{p}) = (1 + \varepsilon u)^{1/\varepsilon} B(\mathbf{p})$ and using [Equation \(1\)](#), we get

$$P_t = \frac{B(\mathbf{p}_t)}{B(\mathbf{p}_s)}, \quad (3)$$

which is evidently independent of the utility level. All conventional price indices that can be derived from economic theory satisfy this property.

2.2 The Nonhomothetic Case: Preferences

Our framework extends the indirect utility (2) to allow for nonhomothetic behavior. To this end, we characterize preferences by an indirect utility function as in Boppart (2014),

$$V(e, \mathbf{p}) = \frac{1}{\varepsilon} \left[\left(\frac{e}{B(\mathbf{p})} \right)^\varepsilon - 1 \right] - \frac{\nu}{\gamma} \left[\left(\frac{D(\mathbf{p})}{B(\mathbf{p})} \right)^\gamma - 1 \right], \quad (4)$$

where $B(\mathbf{p})$ and $D(\mathbf{p})$ are linearly homogeneous functions of prices and the parameters satisfy $0 < \varepsilon \leq 1$, $0 < \gamma \leq 1$, and $\nu > 0$. This utility function belongs to the class of PIGL preferences defined by Muellbauer (1975, 1976) and more generally to the class of “intertemporally aggregable” preferences defined by Alder, Boppart and Müller (2022), and nests several common preference structures: Equation (4) becomes nonhomothetic CES utility over $B(\mathbf{p})$ and $D(\mathbf{p})$ for $\varepsilon = \gamma$, with an elasticity of substitution equal to $1 - \varepsilon$, and includes homothetic Cobb-Douglas and quasi-homothetic preferences as the special cases where $(\varepsilon, \gamma) \rightarrow (0, 0)$ and $\varepsilon = \gamma = 1$, respectively.

It is convenient to think of $B(\mathbf{p})$ and $D(\mathbf{p})$ as homothetic expenditure functions of some intermediate consumption bundles. The parameter ε controls the degree of nonhomotheticity between these two: the expenditure elasticity of demand for the D basket is $1 - \varepsilon$, which is less than 1 under the restrictions on ε . The D basket therefore covers necessity needs and B conversely covers luxury needs. In the limit case $\varepsilon \rightarrow 0$, the expenditure elasticity is 1 and we obtain homothetic preferences. Comparing Equations (2) and (4), we also obtain homothetic preferences for $\varepsilon \neq 0$ whenever $B(\mathbf{p}) = D(\mathbf{p})$ or in the limit case $\nu \rightarrow 0$. The parameter ν is a scale parameter that controls the level of demand for the D basket and γ controls the nonconstant elasticity of substitution between the B and D baskets.⁷

In general, there is nothing restricting an individual good from occurring in both the B and the D baskets. If there is overlap between the sets of goods within B and D , the allocations to the B and D bundles are not directly observable and we obtain what Blundell and Robin (2000) call “latent separability”. Latent separability is a generalization of standard separability assumptions, with the latter being the special case when there is no overlap between B and D . Two-stage budgeting is still valid under latent separability, meaning that the consumer’s allocation problem can be viewed in two stages in which consumers first allocate expenditures between the B and D baskets and then, conditional on this first-stage decision, allocate expenditures across individual goods within B and D . Applying Roy’s identity, the expenditure share w_D allocated to the D basket in the first stage is therefore given by

$$w_D = \nu \left(\frac{B(\mathbf{p})}{e} \right)^\varepsilon \left(\frac{D(\mathbf{p})}{B(\mathbf{p})} \right)^\gamma, \quad (5)$$

and the corresponding expenditure share for the B basket is the residual $w_B = 1 - w_D$. Similarly, the shares w_j^D and w_j^B of total D and B expenditures allocated to individual good j are given by

$$w_j^D = p_j \frac{D_j(\mathbf{p})}{D(\mathbf{p})} \quad \text{and} \quad w_j^B = p_j \frac{B_j(\mathbf{p})}{B(\mathbf{p})}, \quad (6)$$

⁷ By Boppart (2014, Lemma 3), this elasticity of substitution is given by $1 - \gamma - \frac{w_D}{w_B}(\gamma - \varepsilon)$, where w_D and w_B are the expenditure shares allocated to the D and B baskets.

where D_j and B_j denote the partial derivatives of D and B with respect to p_j . Equations (5) and (6) imply an expenditure share w_j of good j in total expenditures of the form

$$w_j = p_j \left[\frac{B_j(\mathbf{p})}{B(\mathbf{p})} + \left(\frac{D_j(\mathbf{p})}{D(\mathbf{p})} - \frac{B_j(\mathbf{p})}{B(\mathbf{p})} \right) \nu \left(\frac{B(\mathbf{p})}{e} \right)^\varepsilon \left(\frac{D(\mathbf{p})}{B(\mathbf{p})} \right)^\gamma \right]. \quad (7)$$

Therefore, nonhomotheticity between B and D also generates nonhomothetic behavior across individual goods, with a good j being a necessity if $D_j/D > B_j/B$ and a luxury vice versa.

Despite being nonhomothetic, the individual budget shares (7) are easily aggregated into some representative average of our choice. In general, the weighted average of good j over any measure N of consumers, indexed by h and with individual weights μ_h , is given by

$$\bar{w}_j = \int_0^N \mu_h w_{jh} dh = p_j \left[\frac{B_j(\mathbf{p})}{B(\mathbf{p})} + \left(\frac{D_j(\mathbf{p})}{D(\mathbf{p})} - \frac{B_j(\mathbf{p})}{B(\mathbf{p})} \right) \nu \left(\frac{B(\mathbf{p})}{\bar{e}\kappa} \right)^\varepsilon \left(\frac{D(\mathbf{p})}{B(\mathbf{p})} \right)^\gamma \right], \quad (8)$$

where $\bar{e} \equiv \frac{1}{N} \int_0^N e_h dh$ is the average expenditure level and κ is a scale-invariant inequality measure defined by

$$\kappa = \left[\int_0^N \mu_h \left(\frac{e_h}{\bar{e}} \right)^{-\varepsilon} dh \right]^{-\frac{1}{\varepsilon}}. \quad (9)$$

Average shares \bar{w}_B and \bar{w}_D of the B and D baskets are defined similarly. Any aggregated expenditure share can therefore be thought of as resulting from a representative consumer with expenditure level $\bar{e}\kappa$. We obtain the market expenditure share $\frac{1}{N} \int_0^N \frac{e_h}{\bar{e}} w_{jh} dh$ when consumers are weighted by their shares of total expenditures: $\mu_h = e_h / (\bar{e}N)$. The corresponding expenditure level $\bar{e}\kappa$ is what Muellbauer (1975, 1976) specifies as *the* representative agent. But aggregation over any other set of weights μ_h , for example some marginal social value of consumers determined by ethical judgments, is also possible. Muellbauer refers to this case as the *socially* representative agent, a terminology which reflects the fact that the representative agent (market as well as social) embodies a Bergson-Samuelson social welfare function with distributional weights μ_h .⁸

2.3 The Nonhomothetic Case: Price Index

Now to our contribution, which uses the indirect utility function (4) to generalize the homothetic cost-of-living index (3). Unlike the homothetic case, the PIGL cost-of-living index depends on a fixed standard of living, represented by the utility level u in the Konüs definition (1). Because both utility and budget shares vary monotonically with the expenditure level, it is possible to fully capture the reference standard of living with the expenditure share on the D basket in the base period s . This expenditure share together with knowledge about the price indices of the B and D baskets and the two parameters ε and γ turns out to be sufficient to obtain the cost-of-living index corresponding to Equation (4). Since $B(\mathbf{p})$ and $D(\mathbf{p})$ are linearly homogenous, we write their price indices like any homothetic index: $P_{Bt} \equiv B(\mathbf{p}_t)/B(\mathbf{p}_s)$ and $P_{Dt} \equiv D(\mathbf{p}_t)/D(\mathbf{p}_s)$. We then have:

⁸ Specifically, the representative expenditure level $\bar{e}\kappa$ is a product of average expenditures and an Atkinson (1970) inequality index subtracted from unity. As shown by Blackorby and Donaldson (1978), this corresponds to a social welfare function given by a weighted generalized mean of order $-\varepsilon$ of individual expenditures with weights μ_h .

Proposition 1 (PIGL cost-of-living index). *If preferences are of the PIGL form (4) and the base-period expenditure share w_{Ds} allocated to the D basket is given, the Konüs cost-of-living index is*

$$P_t = \tilde{P}_t^{\frac{\gamma}{\varepsilon}} P_{Bt}^{1-\frac{\gamma}{\varepsilon}}, \quad (10)$$

where \tilde{P}_t is a nonhomothetic CES composite defined by

$$\tilde{P}_t = \left[\left(1 - \frac{\varepsilon w_{Ds}}{\gamma}\right) P_{Bt}^{\gamma} + \frac{\varepsilon w_{Ds}}{\gamma} P_{Dt}^{\gamma} \right]^{\frac{1}{\gamma}}. \quad (11)$$

The aggregate cost-of-living index over any measure N of consumers is given identically using their weighted average expenditure share $\bar{w}_{Ds} = \int_0^N \mu_{hs} w_{Dhs} dh$ as weight in \tilde{P}_t .

Sketch proof (details in Online Appendix A.1). Set the reference utility to that of the base period expenditure level: $u \equiv V(e_s, \mathbf{p}_s)$. Substituting the period- s indirect utility function into the period- t expenditure function and using Equation (5) then yields $c(u, \mathbf{p}_t) = c(u, \mathbf{p}_s) \tilde{P}_t^{\frac{\gamma}{\varepsilon}} P_{Bt}^{1-\frac{\gamma}{\varepsilon}}$. The result then follows from the Konüs definition (1). \square

Proposition 1 states that the cost-of-living index associated with the indirect utility function (4) is a Cobb-Douglas combination of the price index of the B basket and a CES composite of the price indices of B and D . The functional form therefore looks much like a standard homothetic index with CES and Cobb-Douglas preferences. Unlike the homothetic case, however, the weights in the CES composite vary across the consumption expenditure distribution. Richer consumers spend a lower share w_D on the D basket, which reduces the weight $\varepsilon w_{Ds}/\gamma$. Consequently, the corresponding price index P_{Dt} is weighted less heavily when determining the overall change in the cost of living. In the limit as the expenditure level approaches infinity, the expenditure share w_D goes to zero, thus reducing the cost-of-living index (10) to the price index of the B basket. Note also that Proposition 1, like the indirect utility function (4), nests a nonhomothetic CES index when $\varepsilon = \gamma$, a homothetic Cobb-Douglas index when $(\varepsilon, \gamma) \rightarrow (0, 0)$, and a quasi-homothetic Laspeyres index when $\varepsilon = \gamma = 1$.

The aggregate cost-of-living index has an identical form as Equations (10) and (11) with the average expenditure share \bar{w}_D as weight in \tilde{P}_t , but it also consistently aggregates individual-level indices. Denote the aggregate index by P_t^{RA} (for *representative agent*) and consumer h 's index by P_t^h . Plugging the right-hand side of $\bar{w}_D = \int_0^N \mu_h w_{Dh} dh$ into P_t^{RA} and using Proposition 1 then yields

$$P_t^{RA} = \left[\int_0^N \mu_{hs} \left(P_t^h\right)^{\varepsilon} dh \right]^{\frac{1}{\varepsilon}}. \quad (12)$$

Equation (12) is a generalized mean of order ε of individual indices, with weights μ_h reflecting each consumer's importance in the aggregated expenditure share \bar{w}_D . The aggregate index therefore generalizes the social cost-of-living indices defined by Prais (1959), Pollak (1980, 1981), and Diewert (1984), who only consider weighted arithmetic and geometric averages of individual price indices (the special cases where $\varepsilon = 1$ and $\varepsilon \rightarrow 0$).

Since P_t^{RA} corresponds to a representative agent that embodies a social welfare function with distributional weights μ_h , [Proposition 1](#) provides a simple way to evaluate aggregate welfare changes: aggregate individual base-period expenditure shares using some socially desirable distributional weights μ_h and plug into [Equations \(10\)](#) and [\(11\)](#). For instance, if we follow most statistical agencies and use aggregate market shares in the construction of P_t^{RA} , such that $\mu_h = e_h/(\bar{e}N)$, then [Equation \(12\)](#) becomes what Prais (1959) originally defined as a *plutocratic* index; each dollar spent is assigned equal weight, giving more importance to richer individuals. A more egalitarian option assigns equal weight to each consumer, $\mu_h = 1/N$, which Prais labels a *democratic* index. Or we could use some entirely different set of welfare weights. The point is that any such choice is easily incorporated in the PIGL framework.

Lastly, two potential caveats to [Proposition 1](#) are that preferences are identical across consumers with the same expenditure level and that expenditure shares change monotonically in the level of consumption expenditure. In [Online Appendix B](#) we show that it is straightforward to incorporate taste heterogeneity between the B and D baskets and that this leaves [Proposition 1](#) unaffected. In the same appendix we also discuss a generalization that allows for hump-shaped expenditure shares. This generalization works well at the individual level but requires more stringent conditions for aggregate cost-of-living indices across groups of people.

2.4 Generalized Superlative Indices

[Proposition 1](#) is sufficient to compute total changes in the cost of living across the expenditure distribution. In applications, however, it is often of interest to know also how individual commodities or groups of commodities contribute to the overall price change. An issue with [Proposition 1](#) is that it does not provide a natural decomposition in this regard. We now parameterize the two subindices P_{Bt} , P_{Dt} and show that it is possible to construct such a decomposition by rewriting the PIGL cost-of-living index as a weighted geometric mean of individual price changes.⁹

As an intermediate step, we first decompose the change in the PIGL cost-of-living index into changes in the price indices of the B and D baskets, with weights capturing changes in their expenditure shares over time. Because the standard of living is held fixed, however, these index weights reflect not the observed *Marshallian* expenditure shares but the hypothetical *Hicksian* expenditure shares that prevail at observed prices if utility remains constant. This distinction is important because the nonhomotheticity between B and D implies that changes in observed expenditure shares also include income effects related to changes in the standard of living. By contrast, observed shares equal Hicksian shares whenever preferences are homothetic, so distinguishing between the two does not matter within the subindices P_{Bt} and P_{Dt} (or within any other homothetic price index).

Although not directly observable, the Hicksian allocations for the D and B baskets are readily constructed for the standard of living associated with a base-period allocation. Applying Shephard's lemma on the expenditure function corresponding to the indirect utility function [\(4\)](#) and using [Equations \(5\)](#) and [\(10\)](#) yields a period- t Hicksian expenditure share on the D basket of

⁹ Alternatively, we could consider a decomposition that writes the index as a weighted *arithmetic* mean of individual price changes. Choosing between the two is inconsequential since any multiplicative decomposition can be converted into an additive decomposition and vice versa; see Balk (2008, ch. 4.2).

the form

$$w_{Dt}^h = w_{Ds} \left(\frac{P_{Dt}}{\tilde{P}_t} \right)^\gamma. \quad (13)$$

Given the price indices P_{Bt} and P_{Dt} , a base-period allocation, and values for the parameters ε and γ , Equation (13) becomes straightforward to use. Applying Equation (13) together with the logarithmic mean, which is defined as

$$L(x, y) = \begin{cases} \frac{x - y}{\ln x - \ln y} & \text{if } x \neq y, \\ x & \text{if } x = y, \end{cases} \quad (14)$$

we obtain the following decomposition.

Lemma 1 (Decomposition of the PIGL cost-of-living index). *The rate of change of the Konüs index (10) can be decomposed into price changes of the B and D baskets according to*

$$\frac{P_t}{P_{t-1}} = \left(\frac{P_{Dt}}{P_{Dt-1}} \right)^{\frac{\gamma\phi_t}{\varepsilon}} \left(\frac{P_{Bt}}{P_{Bt-1}} \right)^{1 - \frac{\gamma\phi_t}{\varepsilon}}. \quad (15a)$$

The index weight ϕ_t is given by

$$\phi_t = \frac{L(w_{Dt}^h, w_{Dt-1}^h)}{L(w_{Dt}^h, w_{Dt-1}^h) + L\left(\frac{\gamma}{\varepsilon} - w_{Dt}^h, \frac{\gamma}{\varepsilon} - w_{Dt-1}^h\right)}, \quad (15b)$$

in which $L(\cdot, \cdot)$ is the logarithmic mean (14) and $w_{D\tau}^h$ is the period- τ Hicksian expenditure share (13) on the D basket for the standard of living associated with the base-period expenditure share w_{Ds} .

Sketch proof (details in Online Appendix A.2). Since \tilde{P}_t has a CES form, the change $\tilde{P}_t/\tilde{P}_{t-1}$ can be written as a Sato-Vartia index. Plugging this Sato-Vartia index into the rate of change of (10) gives the result. \square

Lemma 1 makes it straightforward to obtain a multiplicative decomposition of the PIGL cost-of-living index: just plug in two homothetic indices P_{Bt} and P_{Dt} , written on a geometric mean form, into Equation (15). For the sake of generality, we consider two parameterizations of P_{Bt} and P_{Dt} . The first is Diewert's (1976) *quadratic-mean-of-order-r* class, which consists of all indices of the form

$$\frac{P_t}{P_{t-1}} = \sqrt{\left\{ \sum_{j \in J} w_{jt-1} \left(\frac{p_{jt}}{p_{jt-1}} \right)^{\frac{r}{2}} \right\}^{\frac{2}{r}} \left\{ \sum_{j \in J} w_{jt} \left(\frac{p_{jt}}{p_{jt-1}} \right)^{-\frac{r}{2}} \right\}^{-\frac{2}{r}}}, \quad (16)$$

where $r > 0$. The second is Barnett and Choi's (2008) *Theil-Sato* class, which is defined as

$$\frac{P_t}{P_{t-1}} = \prod_{j \in J} \left(\frac{p_{jt}}{p_{jt-1}} \right)^{\delta_{jt}}, \quad \delta_{jt} = \frac{m(w_{jt}, w_{jt-1})}{\sum_{i \in J} m(w_{it}, w_{it-1})}, \quad (17)$$

where $m(x, y)$ is a *symmetric mean* of two variables, a function class that includes all linearly homogenous functions satisfying $\min\{x, y\} \leq m(x, y) = m(y, x) \leq \max\{x, y\}$.

Equations (16) and (17) nest several of the best known price index formulas, including Fisher's (1922) ideal index ($r = 2$), Walsh's (1901) implicit index ($r = 1$) and geometric index (geometric mean, $m(x, y) = \sqrt{xy}$), the Törnqvist (1936) index (arithmetic mean, $r \rightarrow 0$ or $m(x, y) = (x + y)/2$), the Sato (1976)-Vartia (1976) index (logarithmic mean, $m(x, y) = L(x, y)$), and the Theil (1973) index ($m(x, y) = \sqrt[3]{xy(x+y)/2}$). While we could choose any underlying homothetic price indices P_{Dt} and P_{Bt} , these two classes encompass all currently known *superlative* price indices. That is, they correspond to some homothetic expenditure functions that are second-order approximations of any other homothetic expenditure function. Therefore, even if the price indices corresponding to the true expenditure functions $B(\mathbf{p})$ and $D(\mathbf{p})$ have some other forms than those in Equations (16) and (17), we should still hope to reasonably approximate them under this specific parameterization.

While the *Theil-Sato* class splits the overall change into individual commodities by definition, no such natural decomposition exists for the *quadratic-mean-of-order-r* indices. Balk (2004), however, shows how to construct such a decomposition for the special case of the Fisher index ($r = 2$). The generalization to any $r > 0$ is a trivial expansion of Balk's approach, but since we are unaware of any paper that shows this decomposition, we state it as a formal result.

Lemma 2 (Decomposition of the *quadratic-mean-of-order-r* indices). *The quadratic-mean-of-order-r indices (16) can be decomposed into price changes of individual commodities according to*

$$\frac{P_t}{P_{t-1}} = \prod_{j \in J} \left(\frac{p_{jt}}{p_{jt-1}} \right)^{\delta_{jt}}, \quad \delta_{jt} = \frac{1}{2} \left[\frac{\psi_{jt}^L}{\sum_i \psi_{it}^L} + \frac{\psi_{jt}^P}{\sum_i \psi_{it}^P} \right], \quad (18a)$$

where

$$\psi_{jt}^L = w_{jt-1} L \left(\left(\frac{p_{jt}}{p_{jt-1}} \right)^{\frac{r}{2}}, \sum_{i \in J} w_{it-1} \left(\frac{p_{it}}{p_{it-1}} \right)^{\frac{r}{2}} \right), \quad (18b)$$

$$\psi_{jt}^P = w_{jt} L \left(\left(\frac{p_{jt}}{p_{jt-1}} \right)^{-\frac{r}{2}}, \sum_{i \in J} w_{it} \left(\frac{p_{it}}{p_{it-1}} \right)^{-\frac{r}{2}} \right), \quad (18c)$$

in which $L(\cdot, \cdot)$ is the logarithmic mean (14).

Sketch proof (details in Online Appendix A.3). Analogous to Balk (2004). \square

The L and P notation in Lemma 2 refers to the fact that Equation (16) becomes a geometric mean of the Laspeyres and Paasche indices between periods t and $t - 1$ for $r = 2$, and ψ_{jt}^L , ψ_{jt}^P correspond to the weights of good j within these two. Given Lemmas 1 and 2 and the Theil-Sato indices in Equation (17), the decomposition of the PIGL cost-of-living index into individual commodities is immediate.

Proposition 2 (Generalized superlative indices). *If $B(\mathbf{p})$, $D(\mathbf{p})$ are expenditure functions with superlative price indices of the form (16) or (17), the rate of change of the Konüs index (10) can be decomposed into price changes of individual commodities according to*

$$\frac{P_t}{P_{t-1}} = \prod_{j \in J} \left(\frac{p_{jt}}{p_{jt-1}} \right)^{\chi_{jt}}. \quad (19)$$

For each good j , the index weight χ_{jt} is given by

$$\chi_{jt} = \frac{\gamma \phi_t}{\varepsilon} \delta_{jt}^D + \left(1 - \frac{\gamma \phi_t}{\varepsilon} \right) \delta_{jt}^B, \quad (20)$$

where ϕ_t is defined as in Lemma 1 and where δ_{jt}^C , $C \in \{B, D\}$, is constructed according to Equation (17) or Equation (18) using within-basket expenditure shares w_j^C .

Proof. Plug Equations (17) and (18) into Equation (15). \square

Proposition 2 writes the cost-of-living index as a weighted geometric mean of individual price changes with index weights (20) of the following intuitive structure:

$$\text{Weight on } j = \text{Weight on } D \times \text{Weight on } j \text{ within } D + \text{Weight on } B \times \text{Weight on } j \text{ within } B.$$

Similarly to Proposition 1, these weights vary across the base-period expenditure distribution via the weights on B and D due to the nonhomotheticity between these two bundles. The weights on good j within B and D are standard homothetic weights and affect all consumers similarly. If $B(\mathbf{p}) = D(\mathbf{p})$, we get that $\delta_{jt}^B = \delta_{jt}^D$ for all j and the cost-of-living index immediately collapses to its homothetic counterpart in Equations (16) and (17). Proposition 2 therefore provides a natural nonhomothetic generalization of all known superlative price indices.

At first sight, Proposition 2 also looks similar to the methodology used in the literature concerned with inflation inequality, whereby homothetic price indices are computed separately for different income groups. Argente and Lee (2021), Broda and Romalis (2009), Jaravel (2019), and Klick and Stockburger (2021) all construct homothetic price indices of the form $\ln(P_t/P_{t-1}) = \sum_j \delta_{jt} \ln(p_{jt}/p_{jt-1})$, where δ_{jt} are weights computed separately for each income group considered. This method generates heterogeneous weights across the income distribution

and therefore mimics a cost-of-living index with income-specific weights like the one in [Proposition 2](#). An important distinction between the two, however, is that [Proposition 2](#) ensures a fixed standard of living over time. Previous papers, by contrast, construct index weights based on *observed* changes in consumer behavior and consequently also capture changes in the standard of living, thereby generating a potential bias in these cost-of-living indices.

3 Tractable Demand System Estimation

To compute the generalized superlative indices in practice, we need total expenditure shares between the B and D baskets and the expenditure shares within each basket. Yet, if individual goods occur in both baskets, these across and within expenditure shares are not observed in the data. The only feasible approach then is to parameterize $B(\mathbf{p})$ and $D(\mathbf{p})$, estimate the demand system associated with the expenditure share equations (7), and infer these shares from the estimated model. This methodology, however, suffers from a common drawback of nonlinear demand system estimation. In particular, for standard parameterizations the number of parameters to estimate quickly grows out of proportion as we increase the number of goods considered, and this puts strong requirements on the amount of data and computational power a researcher needs to have available.¹⁰ Consequently, demand system estimation is typically restricted to models with only a handful of goods in practice. This issue, however, is fully circumvented within our framework when a simple assumption on the structure of the demand system is met.

Assumption 1. Preferences are quasi-separable between $B(\mathbf{p})$ and $D(\mathbf{p})$. ◁

Under [Assumption 1](#), the price of an individual good occurs in either $B(\mathbf{p})$ or $D(\mathbf{p})$, but not in both. Since the D basket captures necessity needs and B basket luxury needs, it follows that preferences are also quasi-separable between necessities and luxuries. This assumption is therefore easily implemented empirically by allocating luxuries to B and necessities to D .

The immediate consequence of [Assumption 1](#) is that across and within expenditure shares become observable in the data. Summing the total expenditure shares w_j (which are always observable) over goods in D gives the across share w_D . Within shares are then obtained as $w_j^D = w_j/w_D$. The same applies for the B basket. This knowledge is enough to compute basket price indices P_{Bt} and P_{Dt} using [Equation \(16\)](#) or (17). By [Proposition 2](#), the only additional components needed to compute the generalized superlative indices are then the two parameters ε and γ . Using [Equations \(3\)](#) and (5), we may write the period- t expenditure share on the D good as

$$w_{Dt} = \tilde{\nu} \left(\frac{P_{Bt}}{e_t} \right)^\varepsilon \left(\frac{P_{Dt}}{P_{Bt}} \right)^\gamma, \quad (21)$$

where $\tilde{\nu} \equiv \nu B(\mathbf{p}_s)^{\varepsilon-\gamma} D(\mathbf{p}_s)^\gamma$ is a scale parameter. Since w_{Dt} , e_t , P_{Bt} and P_{Dt} are all known, estimating ε and γ from (21) is easily carried out using either linear (by taking logs of (21)) or nonlinear estimation methods. We summarize this empirical approach by the following proposition:

¹⁰ As an illustration, suppose we have n goods and parameterize $B(\mathbf{p})$ and $D(\mathbf{p})$ by the linearly homogeneous translog expenditure function, for which the Törnqvist index is an exact cost-of-living index as shown by Diewert (1976). The PIGL demand system considered here then requires the estimation of $n(n+1)+3$ independent parameters, a number which grows quadratically with n .

Proposition 3 (Tractable demand system estimation). *Under [Assumption 1](#), across- and within-basket expenditure shares are observable in the data and computing the generalized superlative indices [\(19\)](#) only requires estimation of two parameters, ε and γ , from the single expenditure share equation [\(21\)](#).*

At first sight, [Assumption 1](#) may seem to be at odds with the nonhomothetic generalization of the superlative indices: [Proposition 2](#) reduces to the standard homothetic case when $B(\mathbf{p}) = D(\mathbf{p})$, which [Assumption 1](#) excludes by construction. However, the quasi-separability assumption still nests homothetic preferences. In particular, as $\varepsilon \rightarrow 0$ and $\gamma \rightarrow 0$, we obtain Cobb-Douglas preferences with an indirect utility function given by $V(e, \mathbf{p}) = \ln e - \ln [B(\mathbf{p})^{1-\nu} D(\mathbf{p})^\nu]$. The corresponding price index is $P_t = P_{Bt}^{1-\nu} P_{Dt}^\nu$, where $\nu = w_D$ is the homothetic and time-invariant expenditure share on D . Thus, if preferences truly are homothetic, we still expect [Proposition 3](#) to yield a homothetic price index. This index is approximately equal to the corresponding superlative index when $B(\mathbf{p}) = D(\mathbf{p})$, by virtue of superlative indices being second-order approximations of any other homothetic expenditure function. This highlights that using the generalized superlative indices under [Assumption 1](#) should at least (approximately) be weakly better than using the standard homothetic indices. Since nonhomothetic preferences is the empirically relevant case, we do not expect the “approximately” part to matter much, and the empirical application below confirms this. For cases where it nevertheless might be of importance, it turns out that a special case exists where [Assumption 1](#) exactly nests the corresponding homothetic index when $B(\mathbf{p}) = D(\mathbf{p})$: the Törnqvist index.

Proposition 4 (Homothetic Törnqvist index under quasi-separability). *If $B(\mathbf{p})$, $D(\mathbf{p})$ are expenditure functions with price indices of the Törnqvist form and preferences are homothetic Cobb-Douglas according to $V(e, \mathbf{p}) = \ln e - \ln [B(\mathbf{p})^{1-\nu} D(\mathbf{p})^\nu]$, then the cost-of-living index under [Assumption 1](#) is the standard Törnqvist index:*

$$\frac{P_t}{P_{t-1}} = \prod_{j \in J} \left(\frac{p_{jt}}{p_{jt-1}} \right)^{\delta_{jt}}, \quad \text{where} \quad \delta_{jt} = \frac{w_{jt} + w_{jt-1}}{2}.$$

Proof. In Online Appendix A.4. □

4 Data and Estimation

We implement the tractable demand system estimation using consumption and price data from two sources. Household consumption is taken from the interview component of the Consumer Expenditure Survey (CEX) and price data are taken from the product-level Consumer Price Index (CPI) series for all urban consumers. Both are provided by the US Bureau of Labor Statistics (BLS). The CEX interview survey is a quarterly rotating panel of households who are representative of the US population. New households are sampled every month and each household is tracked for up to four consecutive quarters. The survey covers around 95 percent of total household consumption and contains additional information on annual income, socioeconomic

characteristics and other background characteristics like ownership of a car. The survey has been continuously conducted since 1980, though we focus on the years 1995 to 2020 to ensure consistency across waves and to match the availability of the CPI subindices.

As is standard in the literature, we select a sample of respondents between the ages of 25 and 65 who report strictly positive income. To avoid issues with seasonality, we aggregate expenditures to annual levels and, consequently, drop households that do not respond to all four quarterly interviews. To account for differences in household size, we also divide household income and expenditures by the number of adult equivalents in the household using the equivalence scale of the US Census Bureau (see Fox and Burns, 2021). The final dataset on expenditures consists of approximately 3,000 households per year.

We aggregate nondurable consumption expenditures into 21 commodity groups using the hierarchical groupings defined by the BLS and construct prices for these categories by matching them with individual CPI series in a similar manner as Hobijn and Lagakos (2005).¹¹ While Jaravel (2021) stresses that one should consider the most disaggregated product level possible to avoid aggregation biases that cancel out various nonhomotheticities, we consider a rather coarse set of consumption goods categories as this allows us to compare the empirical approach in Proposition 3 with a full demand system estimation that does not invoke quasi-separability. Note that this is not a restriction of our approach as Proposition 3 applies to arbitrarily many goods; the estimation requirement remains at two parameters. Despite our level of product aggregation, we obtain a similar degree of inflation inequality as for instance Jaravel (2019), Argente and Lee (2021), and Klick and Stockburger (2021) when we use the same methodology as these papers (see Section 6), which objects to aggregation bias being a first-order issue in our analysis.

4.1 Classification of Goods Into Luxuries and Necessities

To utilize the tractable demand system estimation, we impose quasi-separability by allocating luxuries to B and necessities to D . The classification of goods is implemented by investigating slopes of the budget-share Engel curves: if the Engel curve of a good decreases as expenditures increase, it is a necessity. Conversely, a good is a luxury if its Engel curve increases with increasing expenditures. We split households into expenditure deciles and, for each good j , run a household-level regression of the expenditure share w_{jh} on the expenditure decile d_h of household h :

$$w_{jh} = \alpha_j + \beta_j d_h + \epsilon_{jh}.$$

If $\beta_j > 0$, we allocate the good to the B basket, otherwise we allocate it to the D basket. Figure 1 shows the Engel curves by expenditure decile together with the resulting classification from the regressions. The results are intuitive and for comparable groups the necessity/luxury split is highly similar to those constructed in similar analyses using CEX data (see for instance Wachter and Yogo, 2010, Lauper and Mangiante, 2023, and Orchard, 2022), thus suggesting that this approach works well on our set of goods. Table C.2 in Online Appendix C lists the estimates of the β_j coefficients along with robustness specifications including regressions on the level and log of expenditures instead of the decile group and additional controls. The estimates of the β_j coefficients are always significantly different from zero and their signs never change across robustness specifications.

¹¹ Table C.1 in Online Appendix C lists the CEX categories and their mapping to the CPI item codes.




FIGURE 1. Empirical and model-implied Engel curves.

Notes. The figure shows the empirical and model-implied expenditure shares by commodity group and expenditure decile averaged over all years. The model-implied expenditure shares correspond to those under the assumption of quasi-separability and a Sato-Vartia price index parameterization of P_{Dt} and P_{Bt} . They are calculated by first taking the model-implied expenditure shares on the B and D baskets and then use the empirical average expenditures shares of all households to obtain shares within these groups.

4.2 Estimation Results

We estimate the preference parameters ε and γ from Equation (21) using a nonlinear GMM estimator on the CEX data. In doing so, we make explicit corrections for two potential issues.

First, it is well known that infrequently bought items, like clothing and transportation goods and services, create a measurement error in the observed level of expenditures. Although we alleviate much of this concern by excluding durable goods and by aggregating expenditures to an annual level, we follow for instance Blundell, Pashardes and Weber (1993) and Banks, Blundell and Lewbel (1997) and control for this endogeneity bias by using household income as an instrument for household consumption expenditure.

Second, an indirect utility specification like the PIGL requires additional attention with respect to the regularity conditions for utility maximization. Specifically, we need to certify that the parameter estimates yield a symmetric and negative semidefinite Slutsky matrix. Under the indirect utility function (4) and Assumption 1, it follows from Boppart (2014, Lemma 1) that

TABLE 1. GMM estimates of the PIGL parameters under quasi-separability.

	Sato-Vartia	Törnqvist	Geometric Walsh	Theil	Fisher's Ideal	Implicit Walsh
ε	0.677 (0.004)	0.677 (0.004)	0.677 (0.004)	0.677 (0.004)	0.677 (0.004)	0.677 (0.004)
γ	0.211 (0.023)	0.211 (0.023)	0.211 (0.023)	0.211 (0.023)	0.211 (0.023)	0.211 (0.023)
$\tilde{\nu}$	327.271 (13.358)	327.437 (13.365)	327.173 (13.354)	327.273 (13.358)	324.273 (13.217)	324.600 (13.233)
Observations	74,372	74,372	74,372	74,372	74,372	74,372
RMSE	0.1487	0.1486	0.1487	0.1487	0.1486	0.1486

Notes. Robust standard errors in parentheses. “RMSE” refers to the root-mean-square error of the expenditure share on the D good: $\sqrt{\sum_h (w_{Dh} - \widehat{w}_{Dh})^2 / N}$. Observations are weighted by their CEX sampling weights.

a necessary and sufficient condition for household h to satisfy the Slutsky restrictions in period t is $\tilde{\nu} \left(\frac{P_{Bt}}{e_{ht}}\right)^\varepsilon \left(\frac{P_{Dt}}{P_{Bt}}\right)^\gamma \leq (1 - \gamma)/(1 - \varepsilon)$. We enforce this constraint by augmenting the GMM estimation with a standard penalty method, in which we iteratively add prohibitive penalties to the GMM objective function outside of the feasible region of parameter values. Consequently, the reported parameter estimates below satisfy the Slutsky restrictions for all observations in the sample.

To gauge the sensitivity to different choices of underlying superlative price indices, we estimate ε and γ for six different choices of P_{Bt} and P_{Dt} . These choices correspond to the indices listed in [Section 2.4](#): the Sato-Vartia, the Törnqvist, the Walsh (geometric and implicit), the Theil, and the Fisher indices. This robustness check is instructive since there is generally no guarantee that superlative indices are numerically similar, despite being second-order approximations of each other (see for instance Hill, 2006). For the estimation exercise, we compute these indices on a monthly frequency using aggregate expenditure shares of all households in the CEX. Since individual household’s expenditures are aggregated to a twelve-month period, we also aggregate the price levels they face by taking the expenditure weighted average over the months each household is in the sample.

The estimated parameters for the six cases are reported in [Table 1](#). All parameters are significantly different from zero at conventional significance levels and the fact that $\varepsilon > 0$ and $\tilde{\nu} > 0$ directly rejects homotheticity. Reassuringly, the choice of price indices for the B and D baskets turns out to be completely inconsequential as all specifications yield close to identical estimates.¹² Moreover, Alder, Boppart and Müller (2022, Proposition 3) show that a sufficient condition for predicted expenditure shares to remain nonnegative as the expenditure level approaches infinity is $0 < \gamma \leq \varepsilon < 1$. This condition is also met in our estimation, though we do not impose the constraint explicitly. Other preference specifications, like the Almost Ideal Demand System, typically violate expenditure share nonnegativity for sufficiently large

¹² The estimates in [Table 1](#) weigh observations by their CEX sampling weights but are robust to not using these weights.

expenditures levels.

To get an idea of how well the estimated model matches the data, we compute budget-share Engel curves using the parameter estimates for the Sato-Vartia specification and plot these against their empirical counterparts in [Figure 1](#). For a good j in basket $C \in \{B, D\}$, we construct model expenditure shares as the product of the model-implied across-share w_C and observed within-shares w_j^C . The across-share is computed from [Equation \(21\)](#) at the representative level of expenditures within each expenditure decile. Since $B(\mathbf{p})$ and $D(\mathbf{p})$ are homothetic, the within-shares are given empirically by the average within-shares \bar{w}_j^B and \bar{w}_j^D across *all* households. Although we of course do not capture *all* nonhomothetic consumption patterns,¹³ the resulting model-implied Engel curves are nonconstant and exhibit reasonably similar patterns as in the data. This contrasts to the constant Engel curves induced by homothetic preferences, which is the underlying preference structure assumed in all other conventional price indices. Using this as our point of departure, [Figure 1](#) underscores that the assumption of quasi-separability between the B and D baskets is not a strong restriction for our price index purposes.

5 Cost-of-Living and Inflation Inequality

The estimation in the previous section suggests that the results are insensitive to the choice of generalized superlative index.¹⁴ For the remainder of the paper, we therefore focus on the nonhomothetic generalization of the Sato-Vartia index (which we refer to as the Generalized Sato-Vartia price index from now on). Since the Sato-Vartia index is the Konüs index corresponding to the canonical CES expenditure function (see Sato, 1976), this choice implies that we are investigating a generalization of homothetic CES preferences. The prominence of CES preferences within macroeconomics and international trade makes this a case of particular interest. The main reason for selecting the generalization of the Sato-Vartia index, however, is that a parameterization of $B(\mathbf{p})$ and $D(\mathbf{p})$ as CES aggregates contains many fewer parameters than the parameterizations that induce, for instance, the Fisher or the Törnqvist indices. This allows us to compare the Generalized Sato-Vartia price index later on to a parsimonious demand system estimation of these preferences when we do not invoke the quasi-separability assumption.

5.1 Main Empirical Results

[Figure 2](#) shows the evolution of the Generalized Sato-Vartia price index from 1995 to 2020. We set the base period to 1995 and, in contrast to [Section 4.2](#), use annual indices for P_{Bt} and P_{Dt} .¹⁵ Even though the generalized superlative price indices allow for characterizations of the entire

¹³ [Figure C.1](#) in [Online Appendix C](#) shows how expenditure shares in the data vary *within* the luxury and necessity baskets across expenditure deciles once we impose [Assumption 1](#), which highlights the nonhomotheticity ignored under quasi-separability.

¹⁴ [Figures C.8](#) and [C.9](#) in [Online Appendix C](#) show the main results for all choices considered in [Section 4.2](#), which confirms that this is indeed the case.

¹⁵ We define a base period as the year from which we collect the reference utility level for the cost-of-living index [\(1\)](#). This differs from that of Jaravel and Lashkari ([forthcoming](#)), who refers to a base period as the year from which they collect a reference price vector for a corresponding quantity index. Just as Jaravel and Lashkari's results change with the choice of reference prices, so do our results change with the choice of reference utility. For a given point in the expenditure distribution, however, the utility level (as captured by the necessity budget share) barely moves over time, so changing the base period does not affect our results; see [Figures C.2](#) to [C.4](#) in [Online Appendix C](#).

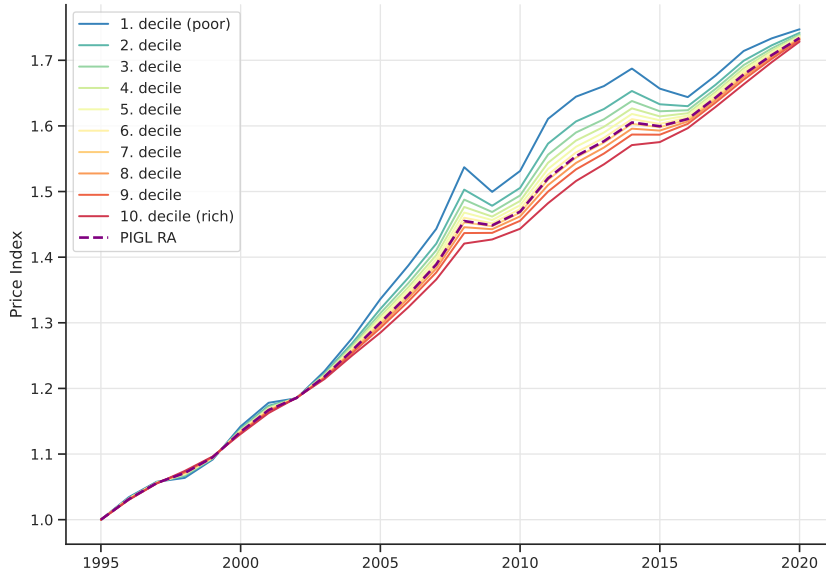


FIGURE 2. Generalized Sato-Vartia price index under quasi-separability by expenditure decile.

Notes. The price index for each expenditure decile is calculated as the price index of the PIGL representative agent over households within each respective decile in the base year. “PIGL RA” stands for the PIGL representative agent over *all* households.

distribution of indices, here we focus on expenditure deciles for ease of exposition. Figures C.5 and C.6 in Online Appendix C show the full price index and inflation distributions.

Figure 2 corroborates two findings from the literature: inflation rates vary across households and poorer households have experienced a larger increase in the cost of living than richer households over the last quarter century. It is noteworthy, though, that the cumulative differences are small. Table 2 makes this point clear: the mean annual inflation rate of the poorest ten percent is only 0.06 percentage points higher than that of the richest ten percent over the full sample period. Consequently, the changes in the cost of living over the 26 years under study do not diverge dramatically between households.

The small differences in the cost of living by 2020 is striking given the substantial heterogeneity observed in subperiods of the sample. For instance, if we zoom in on the years 2004 to 2015, the annual change in the cost of living for the poorest ten percent are on average 0.37 percentage points higher than the change for the richest ten percent. Jaravel (2019) and Argente and Lee (2021) focus on the same years, the former using a CEX-CPI dataset similar to ours and the latter using scanner data for the retail sector, and both find results close to ours. To put this difference in perspective, the Boskin Commission Report estimated the total bias in the aggregate US CPI to be 1.1 percentage points (Boskin *et al.*, 1996). Of these, substitution biases alone account for 0.4 percentage points. The difference we find here is therefore substantial when compared to previously estimated biases in aggregate price indices.

That the differences in the change in the cost of living varies substantially across subperiods is also visible from the implied annual inflation rates, which we plot in Figure 3. In several years, the range of inflation rates across the expenditure distribution exceeds 2 percentage points.

TABLE 2. Inflation rate levels and dispersion over time across the expenditure distribution.

Decile	1996–2020				2004–2015			
	Level		Dispersion		Level		Dispersion	
	Mean (%)	Relative to top decile (pp. diff.)	Standard deviation (%)	Relative to top decile (x Std Dev)	Mean (%)	Relative to top decile (pp. diff.)	Standard deviation (%)	Relative to top decile (x Std Dev)
1	2.28	0.06	2.14	2.51	2.57	0.37	2.71	2.38
2	2.26	0.04	1.76	2.06	2.46	0.26	2.25	1.98
3	2.25	0.03	1.59	1.86	2.41	0.21	2.05	1.79
4	2.24	0.03	1.46	1.71	2.38	0.18	1.89	1.66
5	2.24	0.02	1.36	1.60	2.35	0.15	1.78	1.56
6	2.24	0.02	1.28	1.50	2.33	0.13	1.67	1.46
7	2.23	0.02	1.20	1.40	2.30	0.10	1.57	1.38
8	2.23	0.01	1.12	1.31	2.28	0.08	1.47	1.29
9	2.22	0.01	1.02	1.20	2.25	0.05	1.35	1.18
10	2.22	0.00	0.85	1.00	2.20	0.00	1.14	1.00

Notes. Arithmetic mean and standard deviation of annual inflation over the time periods 2004–2015 and 1996–2020. Under inflation rate levels, the “Relative to top decile” columns show the percentage point difference in the average annual inflation rate to that of the tenth expenditure decile. Under inflation rate dispersion, the same columns show the standard deviation of annual inflation as a multiple of that of the tenth expenditure decile.

Moreover, poor households experience higher inflation than rich households in periods of high inflation but also experience *lower* inflation than the rich in periods of *low* inflation. [Table 2](#) outlines these differences in inflation rate volatility and shows that the standard deviation of inflation is 2.14 for the poorest and 0.85 for the richest. The poorest households therefore experience a 2.5 times *more volatile* inflation rate than the richest households.

In sum, despite the fact that the overall change in the cost of living has not diverged dramatically between groups, there is a considerable difference in the volatility of inflation rates. This in turn generates significant differences in cost-of-living changes during subperiods of our sample. These findings immediately raise the question of *why* poorer households face this higher cost-of-living volatility, and this is what we turn to next.

5.2 What Drives Inflation and the Cost of Living for Different Households?

The nonhomothetic nature of the underlying preferences that we rely on suggests two main channels at work in [Figures 2](#) and [3](#). On the one hand, poorer households allocate a larger share of consumption expenditures on necessity items like food and energy, as highlighted by the Engel curves in [Figure 1](#). This implies that poorer households are more exposed to price changes of these product categories than richer households. Conversely, rich households put relatively more weight on price changes of luxury goods. Therefore, if prices of necessities fluctuate more than those of luxuries, we should expect inflation for poor households to also fluctuate more. On the other hand, the direct impact of a price change on the overall cost of living may be offset by substitution towards relatively cheaper expenditure categories. If this substitution effect varies across the expenditure distribution, this too contributes to the different changes in the cost of




FIGURE 3. Generalized Sato-Vartia inflation under quasi-separability by expenditure decile.

Notes. Inflation for each expenditure decile is calculated as the first difference in the price index of the PIGL representative agent over households within each respective decile. “PIGL RA” stands for the PIGL representative agent over *all* households.

living that we observe. We investigate these two channels in turn.

5.2.1 Different Goods Price Exposure

The geometric-mean form of the generalized superlative price indices in [Proposition 2](#) makes a multiplicative decomposition of these indices straightforward: taking first differences of the log of the price index [\(19\)](#) immediately yields the contribution of each individual expenditure category to overall inflation. [Figure 4](#) plots this decomposition for the poor and the rich and highlights the main contributing expenditure categories to inflation since the start of our sample.¹⁶

Panel (a) in [Figure 4](#) shows that the two primary drivers of inflation for the poor are the expenditure categories “food at home” and “gasoline and utilities”, which together dwarf all other consumption categories combined. The decomposition also reveals that the high inflation volatility of the poor stems from their large exposure to price changes in “gasoline and utilities”, which fluctuates greatly from year to year. By contrast, panel (b) instead identifies housing as the main inflation driver for the rich, as seen by the contribution of “owned dwellings”. The contribution of any individual category for rich households is much less pronounced, however. This is particularly the case for “food at home” and “gasoline and utilities”, thus resulting in a more stable inflation rate for the rich.

¹⁶ [Figure C.7](#) in [Online Appendix C](#) shows the contribution from all expenditure categories to the inflation of the rich and poor.

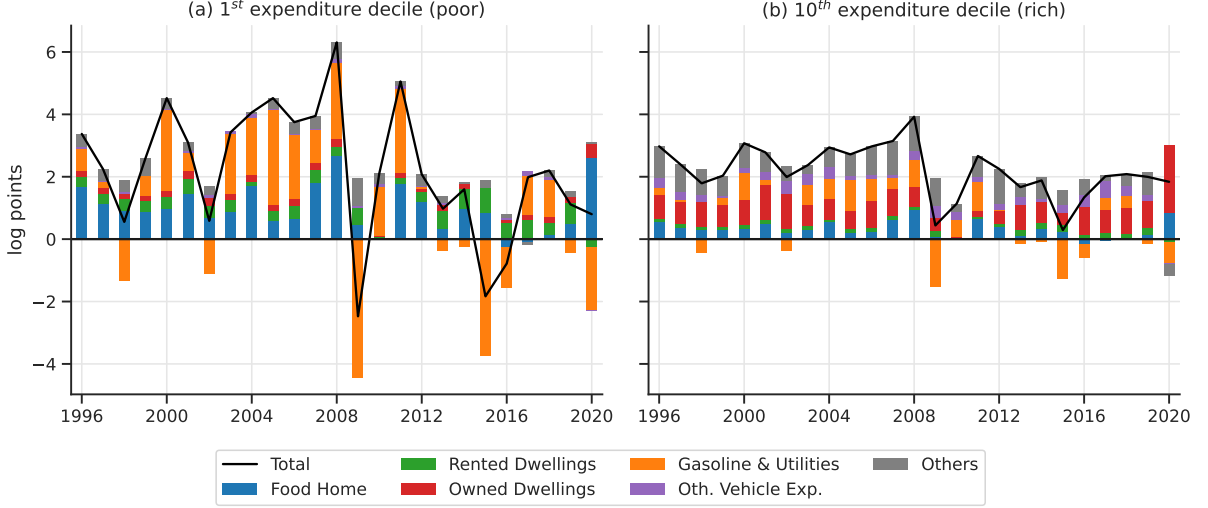


FIGURE 4. Inflation decomposition by expenditure categories.

5.2.2 Different Substitution Behavior

To investigate how substitution behavior affects changes in the cost of living, it is useful to consider a simple behavioral decomposition of the cost-of-living index. Specifically, if households do not substitute at all, consumed quantities of individual goods remain fixed over time. Households' cost-of-living indices should behave exactly like standard Laspeyres indices in that case, with the only differences across households coming from relative quantity differences in their base-period consumption baskets. The log difference between a generalized superlative index and a corresponding Laspeyres index therefore gives a measure of product substitution away from the base-period consumption basket as prices change. The Laspeyres price index is readily obtained as $P_t^L = \sum_j w_{js} (p_{jt}/p_{js})$, and its rate of change between two periods can be decomposed into individual commodities according to

$$\frac{P_t^L}{P_{t-1}^L} = \prod_{j \in J} \left(\frac{p_{jt}}{p_{jt-1}} \right)^{\delta_{jt}^L}, \quad \delta_{jt}^L = \frac{w_{js} \frac{p_{jt-1}}{p_{js}} L\left(\frac{p_{jt}}{p_{jt-1}}, \frac{P_t^L}{P_{t-1}^L}\right)}{\sum_{i \in J} w_{is} \frac{p_{it-1}}{p_{is}} L\left(\frac{p_{it}}{p_{it-1}}, \frac{P_t^L}{P_{t-1}^L}\right)}, \quad (22)$$

where $L(\cdot, \cdot)$ again denotes the logarithmic mean (14).¹⁷ Paired with the generalized superlative index (19), we then obtain the decomposition

$$\ln\left(\frac{P_t}{P_{t-1}}\right) = \underbrace{\sum_{j \in J} \delta_{jt}^L \ln\left(\frac{p_{jt}}{p_{jt-1}}\right)}_{\text{Laspeyres price index}} + \underbrace{\sum_{j \in J} (\chi_{jt} - \delta_{jt}^L) \ln\left(\frac{p_{jt}}{p_{jt-1}}\right)}_{\text{Product substitution}}. \quad (23)$$

¹⁷ The Laspeyres index for periods t and $t-1$ together with the logarithmic mean implies that $0 = \sum_j w_{js} \frac{p_{jt}}{p_{js}} - \frac{P_t^L}{P_{t-1}^L} \sum_j w_{js} \frac{p_{jt-1}}{p_{js}} = \sum_j w_{js} \frac{p_{jt-1}}{p_{js}} \left[\frac{p_{jt}}{p_{jt-1}} - \frac{P_t^L}{P_{t-1}^L} \right] = \sum_j w_{js} \frac{p_{jt-1}}{p_{js}} L\left(\frac{p_{jt}}{p_{jt-1}}, \frac{P_t^L}{P_{t-1}^L}\right) \left[\ln\left(\frac{p_{jt}}{p_{jt-1}}\right) - \ln\left(\frac{P_t^L}{P_{t-1}^L}\right) \right]$, and solving for $\ln(P_t^L/P_{t-1}^L)$ gives the decomposition.

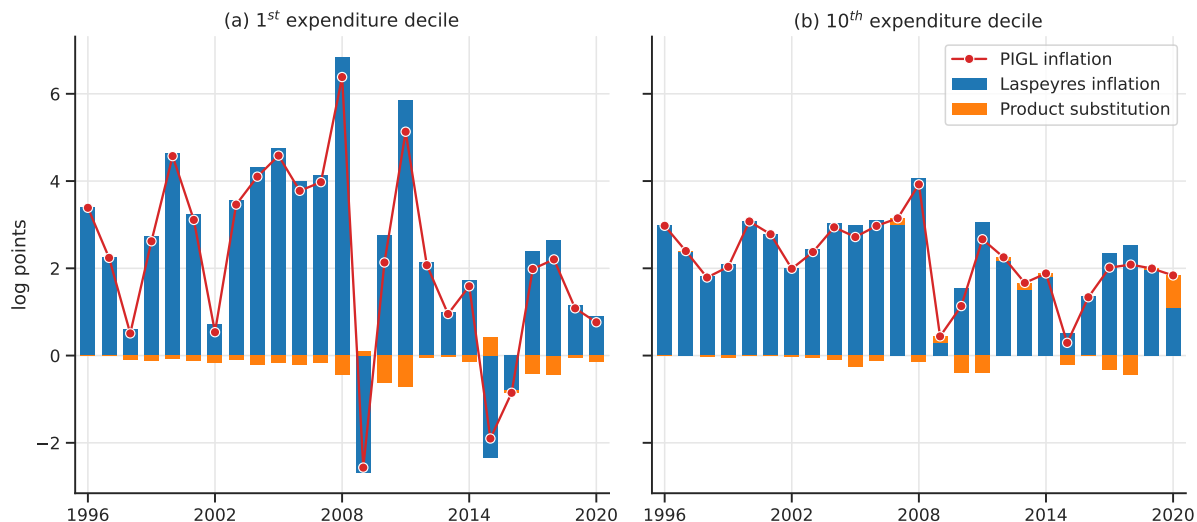


FIGURE 5. Decomposition of the Generalized Sato-Vartia price inflation into the Laspeyres price inflation and product substitution.

Notes. Decomposition according to Equation (23). The Generalized Sato-Vartia price inflation for each expenditure decile denotes the index of the representative agent within each decile.

Equation (23) separates the overall change in the cost of living into mechanical price changes on the base-period consumption basket in the first sum and the overall substitution effect in the second sum.

Figure 5 plots the decomposition (23) of the Generalized Sato-Vartia price inflation for the first and tenth expenditure deciles. The overwhelming majority of the change in the cost of living is caused by pure price changes on the base-period baskets. This is unsurprising given that Laspeyres indices are first-order approximations of true cost-of-living indices (see for instance Deaton and Muellbauer, 1980, ch. 7.1). Despite the small role for product substitution for inflation, Figure 5 nevertheless highlights a more prominent substitution margin among the poor than among the rich. This contrasts with the findings of for instance Argente and Lee (2021), who argue that low-income households experience higher inflation than high-income households, partly because the latter group has more options to substitute away from high-quality goods (luxury goods, presumably).¹⁸

To what extent do these differences affect the disparities in inflation and the cost of living? Figure 6 plots the log difference between the decompositions for the first and tenth expenditure deciles. Over the last 25 years, Laspeyres inflation for the poor is approximately 0.18 percentage point larger than that of the rich. Behavioral differences limit this gap somewhat: more substitution among the poor reduces the difference by around 0.12 percentage points throughout the last quarter century. This effect is seemingly stable, and most swings from year to year come from pure price changes in the different goods households consume.

¹⁸ A key difference to Argente and Lee (2021) is that we evaluate substitution in *quantities consumed* whereas Argente and Lee define product substitution to be *changes in expenditure shares*. Another distinction is that Argente and Lee use scanner data from the retail sector and therefore primarily consider food and other groceries, whereas our CEX data includes virtually all household consumption.




FIGURE 6. Difference of the decomposed Generalized Sato-Vartia price inflation between the first and tenth expenditure decile.

5.2.3 Taking Stock

The two decompositions above suggest that the higher inflation volatility for consumption-poor households is due to differences in types of goods that people consume across the expenditure distribution. Poorer households allocate a larger share to more price-volatile necessity goods, especially “food at home” and “gasoline and utilities”. Behavioral differences help to dampen inflation level differences, as poorer households substitute more than the rich, but does not seem to affect the differences in inflation volatility.

6 Income Effects and Comparison With Previous Methods

A natural question arises at this point. If all we do in our main results is to construct and decompose cost-of-living indices across the expenditure distribution, then why do we bother with developing a new methodology for it? Previous papers studying inflation inequality simply split households into different groups and compute homothetic indices for each group separately based on observed expenditure shares within these groups (examples include Broda and Romalis, 2009, Jaravel, 2019, Argente and Lee, 2021, and Klick and Stockburger, 2021). These group-specific indices do not require any estimation or classification of goods, and nothing stops us from conducting all the analyses in Section 5 with this method instead. So why do we jump through all the hoops needed to get our approach to work?

One reason is that our methodology avoids a particular bias which is present in the group-specific index approach. Recall that a cost-of-living index is defined as the percentage change in consumption expenditures needed to maintain a *constant* standard of living when prices change. Consequently, a true cost-of-living index only reflects price changes and subsequent substitution effects, while ignoring income effects on consumer behavior. The corresponding index weights are therefore based not on actual observed expenditure shares, but on hypothetical *Hicksian* expenditure shares (that is, the shares that prevail at observed prices if utility remains fixed at its base-period level). This distinction does not matter under homothetic preferences, because




FIGURE 7. Comparison of group-specific Sato-Vartia cost-of-living index with PIGL cost-of-living index.

Notes. The group-specific Sato-Vartia index is computed based on a grouping of households into expenditure deciles of the respective year. The PIGL price index for each expenditure decile is calculated as the price index of the PIGL representative agent over households within each respective decile.

Hicksian shares equal observed shares in this case. If preferences are nonhomothetic, however, then observed expenditure shares also capture income effects related to changes in the standard of living over time. This drives a wedge between observed and Hicksian shares, thereby causing a bias in price indices that rely on the former.¹⁹ Our approach overcomes this issue by relying on Hicksian rather than observed expenditure shares.

We examine the magnitude of this bias by comparing the Generalized Sato-Vartia index to the group-specific price index approach. To that end, we partition the households in the CEX data into year-specific consumption deciles and use the observed average expenditure shares w_{jt} within these groups to compute decile-specific Sato-Vartia indices $\ln P_t = \sum_j \delta_{jt} \ln(p_{jt}/p_{js})$, where $\delta_{jt} = L(w_{jt}, w_{js}) / \sum_i L(w_{it}, w_{is})$. Figure 7 plots the price indices and the inflation rates for the top and bottom consumption deciles using this approach against those of the Generalized Sato-Vartia index. Table 3 summarizes average inflation rates and inflation volatility over the full sample period across the consumption distribution for both approaches.

Compared to the Generalized Sato-Vartia index, the group-specific indices over-estimate inflation at the lower end of the distribution and under-estimate it at the upper end. The size of the bias ranges from a 0.07 percentage point higher average annual inflation rate for the lowest decile to a 0.23 percentage point lower inflation rate for the top decile. These differences generate a larger inflation gap across the distribution: the average annual inflation rate in the bottom decile is 2.35 while the corresponding number for the top decile is 1.99 percent, a 0.36 percentage point difference. The group-specific Sato-Vartia indices therefore produce a divergence in the cost of living across the expenditure distribution, as shown in panel (a) of Figure 7.

¹⁹ Oulton (2008) labels this a “path-dependence bias” and provides a more detailed discussion on the issue.

TABLE 3. Inflation across the consumption distribution for two different approaches, 1996–2020.

Generalized Sato-Vartia index (PIGL)											
Group-specific homothetic Sato-Vartia indices											
Level				Level				Dispersion			
Decile	Mean (%)	Relative to top decile (pp. diff.)	Standard deviation (%)	Relative to top decile (x Std Dev)	Mean (%)	Relative to top decile (pp. diff.)	Relative to PIGL (pp. diff.)	Standard deviation (%)	Relative to top decile (x Std Dev)	Relative to PIGL (pp. diff.)	Relative to top decile (x Std Dev)
1	2.28	0.06	2.14	2.51	2.35	0.36	0.07	1.48	1.51	0.69	0.69
2	2.26	0.04	1.76	2.06	2.33	0.34	0.07	1.41	1.45	0.80	0.80
3	2.25	0.03	1.59	1.86	2.31	0.32	0.06	1.51	1.54	0.95	0.95
4	2.24	0.03	1.46	1.71	2.30	0.31	0.05	1.47	1.51	1.01	1.01
5	2.24	0.02	1.36	1.60	2.24	0.25	0.00	1.36	1.39	1.00	1.00
6	2.24	0.02	1.28	1.50	2.24	0.25	0.01	1.32	1.35	1.03	1.03
7	2.23	0.02	1.20	1.40	2.22	0.23	-0.02	1.29	1.32	1.08	1.08
8	2.23	0.01	1.12	1.31	2.19	0.20	-0.03	1.20	1.23	1.07	1.07
9	2.22	0.01	1.02	1.20	2.15	0.16	-0.08	1.12	1.14	1.10	1.10
10	2.22	0.00	0.85	1.00	1.99	0.00	-0.23	0.98	1.00	1.15	1.15

Notes. Arithmetic mean and standard deviation of annual inflation over the time period 1996–2020. Under inflation rate levels, the “Relative to top decile” columns show the percentage point difference in the average annual inflation rate to that of the tenth expenditure decile while the “Relative to PIGL” column shows the difference in average inflation between the group-specific Sato-Vartia index and the PIGL index in the same expenditure decile. Under inflation rate dispersion, the same columns show the standard deviation of annual inflation as a multiple of that of the tenth expenditure decile and of the same decile for the PIGL index.

The group-specific approach also generates a more stable inflation rate over both time and consumption groups, as shown in panel (b) of [Figure 7](#). Compared to the Generalized Sato-Vartia index, the standard deviation of inflation over time is now 31 percent lower for the bottom decile and 15 percent higher for the top decile. This evidently reduces the volatility gap: the relative inflation rate volatility between the poorest and richest households is now 1.51, considerably lower than the 2.51 found with the Generalized Sato-Vartia index.

Taken together, these results point to the importance of purging price indices from any income effects on household behavior. Failing to do so in our application increases the inflation rate difference between the top and bottom of the consumption distribution by 0.3 percentage points and generates biased disparities in the cost of living. Interestingly, 0.3 percentage points is close to the range of inflation rates found in papers employing the group-specific approach. Perhaps most comparable are Jaravel ([2019](#)) and Klick and Stockburger ([2021](#)), who construct income-specific Törnqvist indices using similar CEX-CPI data. These authors find a 0.35–0.38 percentage point inflation rate difference between the top and bottom 20–25 percent of the income distribution. Our Generalized Sato-Vartia index exhibits similar inflation differences for the years considered by these authors, but the size of the bias found here nevertheless prompts the question of how much of their findings are driven by the same bias.

7 The (Un)Importance of the Separability Assumption

The results above rely on an implementation method that rests on one key assumption: that goods can be split up into two separate consumption bundles, one that includes all necessity goods and one that includes all luxury goods. How strong is this assumption in our application? To answer this question, we finish the paper with a robustness check in which we compare the baseline results in [Section 5](#) with those of a demand system estimation on a fully parametrized model.

To that end, we parametrize the cost functions $B(\mathbf{p})$ and $D(\mathbf{p})$ and perform a nonlinear GMM estimation of the system of the expenditure shares defined by [Equation \(7\)](#). Since the Sato-Vartia index is the Konüs index corresponding to the CES expenditure function, we parametrize the cost functions of the two baskets as such:

$$B(\mathbf{p}) = \left(\sum_{j \in J} \omega_j p_j^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \quad \text{and} \quad D(\mathbf{p}) = \left(\sum_{j \in J} \theta_j p_j^{1-\varphi} \right)^{\frac{1}{1-\varphi}}, \quad (24)$$

where $\sigma, \varphi > 0$ and where the share parameters satisfy $\sum_{j \in J} \omega_j = \sum_{j \in J} \theta_j = 1$ and $\omega_j, \theta_j \geq 0$ for all $j \in J$.²⁰ The possibility of binding boundary constraints on the parameters generates a nonstandard asymptotic distribution, so neither the bootstrap nor the standard covariance matrix produces consistent standard errors in this case (see Andrews, [1999](#), [2000](#)). Although this complicates statistical inference, the point estimates of the GMM estimator still remains consistent when a true parameter lies on its boundary (Andrews, [2002](#)). Our GMM estimation can therefore be seen as something akin to a calibration rather than a full estimation. As an initial parameter guess, we set ε , γ and $\tilde{\nu}$ to the values in [Table 1](#) and distribute ω_j (θ_j) equally

²⁰ A similar parametrization is also considered by Alder, Boppart and Müller ([2022](#)).

TABLE 4. GMM estimates of the preference parameters.

	Quasi-separability (Sato-Vartia)	Full demand system (CES)
ε	0.677 (0.004)	0.685 (0.004)
γ	0.211 (0.023)	0.505 (0.018)
$\tilde{\nu}$	327.271 (13.358)	346.736 (11.793)
σ		0.050 (0.012)
φ		0.360 (0.006)
Observations	74,372	74,372
RMSE	0.1487	0.1494

Notes. Robust standard errors in parentheses. “RMSE” refers to the root-mean-square error of the expenditure share on the D good: $\sqrt{\sum_h (w_{Dh} - \hat{w}_{Dh})^2 / N}$. Observations are weighted by their CEX sampling weights.

among the goods that are classified as luxuries (necessities) in the quasi-separability case while the remaining parameters are set to zero.

Table 4 and Figure 8 show the point estimates associated with the local minimum obtained from our initial guess.²¹ The key parameter with respect to nonhomotheticity, ε , and the scale parameter $\tilde{\nu}$ are virtually unchanged. The CES weights ω_j and θ_j are also fully in line with the quasi-separability assumption: all necessities from our baseline approach hold zero weight in the luxury basket B , and vice versa for luxuries. The main difference relates to γ . The higher value of γ in the full demand system lowers the elasticity of substitution between B and D for everyone, but more so for the poor than for the rich. For instance, the elasticity of substitution for a household with a 90 percent expenditure share on the D basket is 4.98 in the baseline and 2.12 in the full demand system. The corresponding numbers for a household with a 10 percent expenditure share on D are 0.84 and 0.52. Relative to our baseline approach, the higher value of γ therefore reduces the importance of substitution behavior in explaining any cost-of-living differences between rich and poor households.

Overall, the parameter estimates from the estimated demand system are fully in line with the quasi-separability assumption in the baseline approach. We plot a comparison between the cost-of-living indices resulting from the two approaches in Figure 9 to make sure that the higher

²¹ The usual caveats of nonlinear estimation apply, in particular that the GMM objective function may exhibit multiple local minima, with no guarantee that our solution is the global minimum. In principle, we could get around this issue by a grid search. There are 45 parameters to estimate, however, and the associated curse of dimensionality makes such a grid search infeasible in practice.

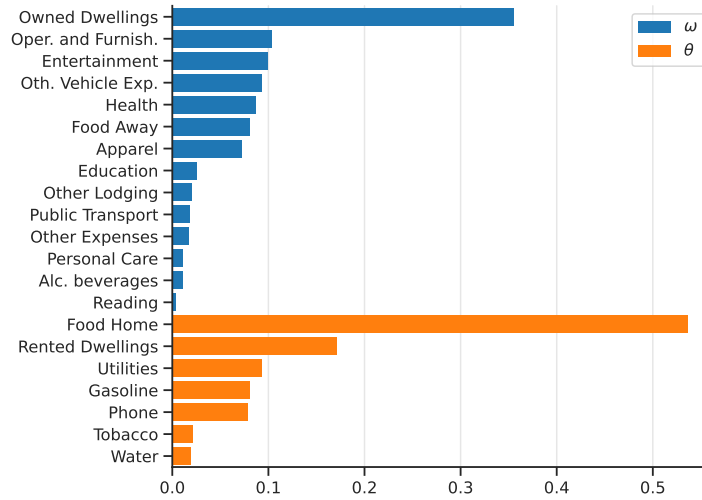


FIGURE 8. Point estimates for the CES share parameters ω_j and θ_j .

Notes. Results are for the closest local minimum to the quasi-separable case which has been used as initial guess for the parameters. Point estimates on the zero boundary are omitted.

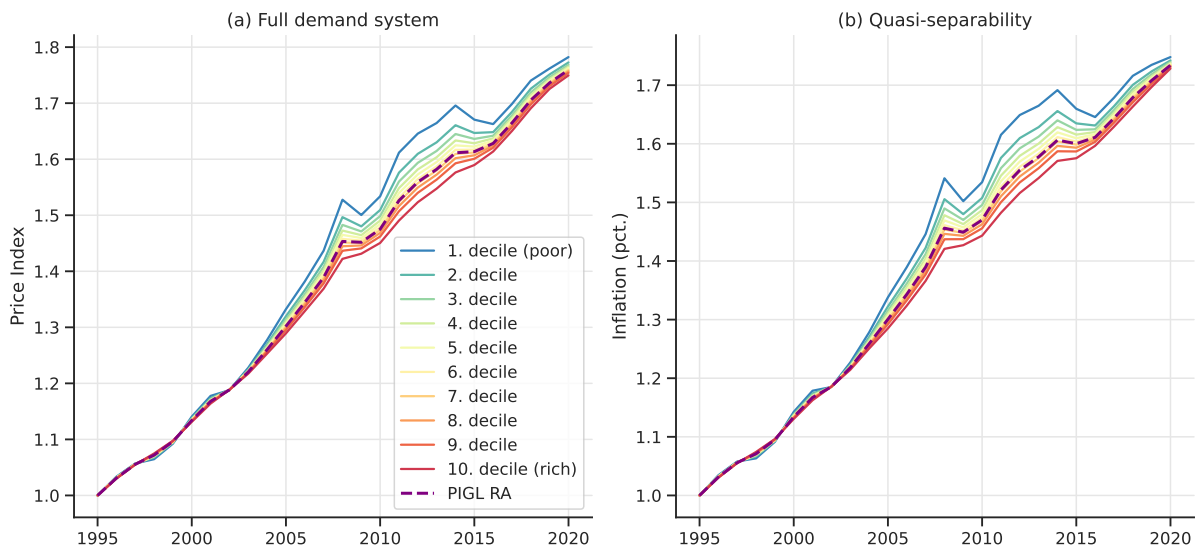


FIGURE 9. Comparison of the Generalized Sato-Vartia price index for the full demand system and under quasi-separability.

value of γ does not radically alter our baseline findings.²² Fortunately, the two cases are barely distinguishable. Thus, taken at face value, the findings in this section seem to fully justify the quasi-separability assumption for the dataset that we use.

²² A similar comparison of inflation rates is shown in Figure C.10 in Online Appendix C.

8 Conclusion

All commonly used cost-of-living indices that can be derived from economic theory rely on the key assumption that consumption behavior is independent of income, despite over 160 years of empirical evidence that suggest otherwise. We abandon this assumption and present a generalization of these indices that accounts for the fact that income groups are affected differently if prices of luxuries and necessities change differently. By considering a nonhomothetic preference structure that divides commodities into a necessity bundle and luxury bundle, we can measure cost-of-living changes using observed prices and consumption and only two estimated parameters. We thereby maintain much of the practical simplicity of standard price indices without compromising on the nonhomothetic foundation of consumer demand theory.

Our empirical analysis suggests that inflation disparities between US households are primarily a transitory phenomenon. Between 1995 and 2020, inflation rates remain similar on average across the expenditure distribution, though substantial differences occur in the short run as consumption-poor households face more volatile inflation rates than the consumption-rich. These findings speak especially to the monetary economics literature: heterogeneous inflation rates throughout the business cycle implies heterogeneous real interest rates, which impacts the consumption responses of households to monetary policy shocks. Inflation inequality therefore has potential implications for the transmission mechanism of these shocks and for optimal monetary policy. The PIGL preferences that we use are shown by Boppart (2014) to be compatible with standard macroeconomic modeling and could be used in monetary policy research to investigate these issues further.

Finally, we want to stress that although our empirical analysis uses microdata for a fixed set of relatively broad commodity groups, the approach that we propose is not limited to this setting. For example, increasing the number of goods considered is trivial; the estimation burden remains at two parameters. Product entry and exit as in Feenstra (1994) and taste shocks as in Redding and Weinstein (2020) are also fully permitted in the CES parametrization that we employ. Meanwhile, the aggregation properties of PIGL preferences means that its parameters can be estimated from macro data without aggregation bias. It is thus possible to obtain a full distribution of nonhomothetic cost-of-living indices from aggregate time series, as long as there exists a cross-sectional consumption distribution for the base period. This property is useful when microdata is not available for most of the time periods considered, for instance if working with national accounts data. Our methodology therefore demonstrates sufficient applicability as well as simplicity to appeal to researchers and practitioners alike.

References

- Alder, S., Boppart, T., and Müller, A. (2022). “A theory of structural change that can fit the data”. *American Economic Journal: Macroeconomics*, vol. 14(2), pp. 160–206. (Cited on pp. 1, 5, 16, 26.)
- Andrews, D. W. K. (1999). “Estimation when a parameter is on a boundary”. *Econometrica*, vol. 67(6), pp. 1341–1383. (Cited on p. 26.)
- Andrews, D. W. K. (2000). “Inconsistency of the bootstrap when a parameter is on the boundary of the parameter space”. *Econometrica*, vol. 68(2), pp. 399–405. (Cited on p. 26.)
- Andrews, D. W. K. (2002). “Generalized method of moments estimation when a parameter is on a boundary”. *Journal of Business and Economic Statistics*, vol. 20(4), pp. 530–544. (Cited on p. 26.)
- Argente, D., and Lee, M. (2021). “Cost of living inequality during the Great Recession”. *Journal of the European Economic Association*, vol. 19(2), pp. 913–952. (Cited on pp. 2, 11, 14, 18, 22, 23.)
- Atkinson, A. B. (1970). “On the measurement of inequality”. *Journal of Economic Theory*, vol. 2(3), pp. 244–263. (Cited on p. 6.)
- Balk, B. M. (2004). “Decompositions of Fisher indexes”. *Economics Letters*, vol. 82(1), pp. 107–113. (Cited on pp. 10, 11.)
- Balk, B. M. (2008). *Price and quantity index numbers: Models for measuring aggregate change and difference*. Cambridge University Press. (Cited on p. 8.)
- Banks, J., Blundell, R., and Lewbel, A. (1997). “Quadratic Engel curves and consumer demand”. *Review of Economics and Statistics*, vol. 79(4), pp. 527–539. (Cited on p. 15.)
- Baqae, D., Burstein, A., and Koike-Mori, Y. (forthcoming). “A fixed point approach to measuring welfare”. *Quarterly Journal of Economics*. (Cited on p. 3.)
- Barnett, W. A., and Choi, K.-H. (2008). “Operational identification of the complete class of superlative index numbers: An application of Galois theory”. *Journal of Mathematical Economics*, vol. 44(7–8), pp. 603–612. (Cited on pp. 3, 10.)
- Blackorby, C., and Donaldson, D. (1978). “Measures of relative equality and their meaning in terms of social welfare”. *Journal of Economic Theory*, vol. 18(1), pp. 59–80. (Cited on p. 6.)
- Blundell, R., Pashardes, P., and Weber, G. (1993). “What do we learn about consumer demand patterns from micro data?” *American Economic Review*, vol. 83(3), pp. 570–597. (Cited on p. 15.)
- Blundell, R., and Robin, J.-M. (2000). “Latent separability: Grouping goods without weak separability”. *Econometrica*, vol. 68(1), pp. 53–84. (Cited on p. 5.)
- Boppart, T. (2014). “Structural change and the Kaldor facts in a growth model with relative price effects and non-Gorman preferences”. *Econometrica*, vol. 82(6), pp. 2167–2196. (Cited on pp. 1, 5, 15, 29.)
- Boskin, M. J., Dulberger, E. R., Gordon, R. J., Griliches, Z., and Jorgenson, D. (1996). *Toward a more accurate measure of the cost of living: Final report to the Senate Finance Committee from the Advisory Commission To Study The Consumer Price Index* (tech. rep.). Washington, DC: Advisory Commission to Study the Consumer Price Index. (Cited on p. 18.)
- Broda, C., and Romalis, J. (2009). *The welfare implications of rising price dispersion*, Booth School of Business, University of Chicago. Unpublished manuscript. (Cited on pp. 2, 11, 23.)
- Comin, D., Lashkari, D., and Mestieri, M. (2021). “Structural change with long-run income and price effects”. *Econometrica*, vol. 89(1), pp. 311–374. (Cited on p. 2.)
- Cravino, J., Levchenko, A. A., and Rojas, M. (2022). “Population aging and structural transformation”. *American Economic Journal: Macroeconomics*, vol. 14(4), pp. 479–498. (Cited on p. 1.)
- Deaton, A., and Muellbauer, J. (1980). *Economics and consumer behavior*. Cambridge University Press. (Cited on p. 22.)
- Diewert, W. E. (1976). “Exact and superlative index numbers”. *Journal of Econometrics*, vol. 4(2), pp. 115–145. (Cited on pp. 3, 9, 12.)
- Diewert, W. E. (1978). “Superlative index numbers and consistency in aggregation”. *Econometrica*, vol. 46(4), pp. 883–900. (Cited on p. 3.)
- Diewert, W. E. (1984). “Group cost of living indexes: Approximations and axiomatics”. *Methods of Operations Research*, vol. 48, pp. 23–45. (Cited on p. 7.)

- Diewert, W. E. (1993). “The early history of price index research”. In W. E. Diewert and A. O. Nakamura (Eds.), *Essays in index number theory* (pp. 33–65, Vol. 1). North-Holland. (Cited on p. 3.)
- Engel, E. (1857). “Die vorherrschenden Gewerbszweige in den Gerichtsämtern mit Beziehung auf die Productions- und Consumtionsverhältnisse des Königreichs Sachsen”. *Zeitschrift des Statistischen Bureaus des Königlich Sächsischen Ministeriums des Innern*, vol. 3(8–9), pp. 153–182. (Cited on p. 1.)
- Fan, T., Peters, M., and Zilibotti, F. (2023). “Growing like India—the unequal effects of service-led growth”. *Econometrica*, vol. 91(4), pp. 1457–1494. (Cited on p. 1.)
- Feenstra, R. C. (1994). “New product varieties and the measurement of international prices”. *American Economic Review*, vol. 84(1), pp. 157–177. (Cited on pp. 3, 29.)
- Feenstra, R. C., and Reinsdorf, M. B. (2000). “An exact price index for the almost ideal demand system”. *Economics Letters*, vol. 66(2), pp. 159–162. (Cited on p. 3.)
- Fisher, I. (1922). *The making of index numbers: A study of their varieties, tests, and reliability*. Houghton Mifflin Company. (Cited on pp. 1, 10.)
- Fox, L. E., and Burns, K. (2021). *The supplemental poverty measure: 2020* (Current Population Reports No. P60–275). Washington, DC: U.S. Census Bureau. (Cited on p. 14.)
- Gorman, W. M. (1996). “Quasi-separable preferences, costs, and technologies”. In C. Blackorby and A. F. Shorrocks (Eds.), *Separability and aggregation: The collected works of W. M. Gorman* (pp. 101–114, Vol. 1). Oxford University Press. (Cited on p. 2.)
- Hill, R. J. (2006). “Superlative index numbers: Not all of them are super”. *Journal of Econometrics*, vol. 130(1), pp. 25–43. (Cited on p. 16.)
- Hobijn, B., and Lagakos, D. (2005). “Inflation inequality in the United States”. *Review of Income and Wealth*, vol. 51(4), pp. 581–606. (Cited on pp. 2, 14.)
- Hochmuth, P., Pettersson, M., and Weissert, C. J. (2023). *A distributional PCE price index from aggregate data*. Unpublished manuscript. (Cited on p. 3.)
- Jaravel, X. (2019). “The unequal gains from product innovations: Evidence from the U.S. retail sector”. *Quarterly Journal of Economics*, vol. 134(2), pp. 715–783. (Cited on pp. 2, 11, 14, 18, 23, 26.)
- Jaravel, X. (2021). “Inflation inequality: Measurement, causes, and policy implications”. *Annual Review of Economics*, vol. 13, pp. 599–629. (Cited on pp. 2, 14.)
- Jaravel, X., and Lashkari, D. (forthcoming). “Measuring growth in consumer welfare with income-dependent preferences”. *Quarterly Journal of Economics*. (Cited on pp. 3, 17.)
- Kaplan, G., and Schulhofer-Wohl, S. (2017). “Inflation at the household level”. *Journal of Monetary Economics*, vol. 91, pp. 19–38. (Cited on p. 2.)
- Klick, J., and Stockburger, A. (2021). *Experimental CPI for lower and higher income households* (BLS Working Papers No. 537). Washington, DC: U.S. Bureau of Labor Statistics. (Cited on pp. 2, 11, 14, 23, 26.)
- Konüs, A. A. (1939). “The problem of the true index of the cost of living” (J. Bronfenbrenner, Trans.). *Econometrica*, vol. 7(1), pp. 10–29. (Original work published 1924). (Cited on pp. 2–4.)
- Lauper, C., and Mangiante, G. (2023). *Monetary policy shocks and inflation heterogeneity*, University of Lausanne. Unpublished manuscript. (Cited on pp. 2, 14.)
- León-Ledesma, M., and Moro, A. (2020). “The rise of services and balanced growth in theory and data”. *American Economic Journal: Macroeconomics*, vol. 12(4), pp. 109–146. (Cited on p. 1.)
- Matsuyama, K. (2019). “Engel’s law in the global economy: Demand-induced patterns of structural change, innovation, and trade”. *Econometrica*, vol. 87(2), pp. 497–528. (Cited on p. 2.)
- McGranahan, L., and Paulson, A. (2005). *Constructing the Chicago Fed income based economic index – consumer price index: Inflation experiences by demographic group: 1983–2005* (Working Paper No. 2005–20). Chicago, IL: Federal Reserve Bank of Chicago. (Cited on p. 2.)
- Muellbauer, J. (1974). “Prices and inequality: The United Kingdom experience”. *Economic Journal*, vol. 84(333), pp. 32–55. (Cited on p. 2.)
- Muellbauer, J. (1975). “Aggregation, income distribution and consumer demand”. *Review of Economic Studies*, vol. 42(4), pp. 525–543. (Cited on pp. 1, 5, 6.)
- Muellbauer, J. (1976). “Community preferences and the representative consumer”. *Econometrica*, vol. 44(5), pp. 979–999. (Cited on pp. 1, 5, 6.)

- Orchard, J. (2022). *Cyclical demand shifts and cost of living inequality*, Federal Reserve Board. Unpublished manuscript. (Cited on pp. 2, 14.)
- Oulton, N. (2008). “Chain indices of the cost-of-living and the path-dependence problem: An empirical solution”. *Journal of Econometrics*, vol. 144(1), pp. 306–324. (Cited on pp. 3, 24.)
- Oulton, N. (2012). “How to measure living standards and productivity”. *Review of Income and Wealth*, vol. 58(3), pp. 424–456. (Cited on p. 3.)
- Pollak, R. A. (1980). “Group cost-of-living indexes”. *American Economic Review*, vol. 70(2), pp. 273–278. (Cited on p. 7.)
- Pollak, R. A. (1981). “The social cost of living index”. *Journal of Public Economics*, vol. 15(3), pp. 311–336. (Cited on p. 7.)
- Prais, S. J. (1959). “Whose cost of living?” *Review of Economic Studies*, vol. 26(2), pp. 126–134. (Cited on pp. 7, 8.)
- Redding, S. J., and Weinstein, D. E. (2020). “Measuring aggregate price indices with taste shocks: Theory and evidence for CES preferences”. *Quarterly Journal of Economics*, vol. 135(1), pp. 503–560. (Cited on pp. 3, 29.)
- Samuelson, P. A., and Swamy, S. (1974). “Invariant economic index numbers and canonical duality: Survey and synthesis”. *American Economic Review*, vol. 64(4), pp. 566–593. (Cited on pp. 3, 4.)
- Sato, K. (1976). “The ideal log-change index number”. *Review of Economics and Statistics*, vol. 58(2), pp. 223–228. (Cited on pp. 1, 10, 17.)
- Theil, H. (1973). “A new index number formula”. *Review of Economics and Statistics*, vol. 55(4), pp. 498–502. (Cited on p. 10.)
- Törnqvist, L. (1936). “The Bank of Finland’s consumption price index”. *Bank of Finland Monthly Bulletin*, vol. 16(10), pp. 27–34. (Cited on pp. 1, 10.)
- Vartia, Y. O. (1976). “Ideal log-change index numbers”. *Scandinavian Journal of Statistics*, vol. 3(3), pp. 121–126. (Cited on pp. 1, 10.)
- Wachter, J. A., and Yogo, M. (2010). “Why do household portfolio shares rise in wealth?” *Review of Financial Studies*, vol. 23(11), pp. 3929–3965. (Cited on p. 14.)
- Walsh, C. M. (1901). *The measurement of general exchange-value*. The Macmillan Company. (Cited on p. 10.)