A Nonhomothetic Price Index and Inflation Heterogeneity

Philipp Hochmuth[*] Markus Pettersson[†] Christoffer Weissert[‡]

May 20, 2022

Abstract

We derive a microfounded, nonhomothetic generalization of all known superlative price indices, including the Fisher, the Törnqvist, and the Sato-Vartia indices. The index largely avoids the need for estimation, aggregates consistently across heterogeneous households, and admits different index weights across the expenditure distribution. The latter property rationalizes the methods used in most previous measurements of inflation inequality. In an empirical application to the United States using CEX-CPI data for the period 1995–2020, we find: (i) poor and rich households experience on average the same inflation rate; but (ii) inflation for the poorest decile is more than 2.5 times as volatile as that of the richest decile; and (iii) this higher volatility primarily stems from a larger exposure to price changes in food, gas and utilities. In these findings, substitution between goods as prices change plays only a second-order role. Instead, almost all differences come from mechanical changes in the cost of different base-period reference baskets.

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

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

[*]Stockholm University. Email: philipp.hochmuth@iies.su.se.
[†]Stockholm School of Economics. Email: markus.pettersson@phdstudent.hhs.se.
[‡]University of Copenhagen and Danmarks Nationalbank. Email: cjw@econ.ku.dk.

1 Introduction

Does inflation vary with income? Conventional price indices that are used for the measurement of inflation cannot answer this question because these rely fully on the assumption of homothetic preferences. That is, consumers are assumed to make identical consumption allocations, regardless of income level. Yet, one of the oldest empirical economic facts, dating back to at least Engel (1857), is that consumption patterns differ systematically between rich and poor consumers. In other words, preferences are *not* homothetic. Differences in consumption bundles raises the possibility for inflation inequality, with implications for any area where inflation matters, not least monetary policy and the measurement of real incomes. By now there exists an abundance of empirical research investigating this issue by computing standard (homothetic) price indices for separate income groups. The question of how to consistently incorporate nonhomothetic preferences into conventional price index formulas, however, remains unsolved.

The goal of this paper is to tackle this problem head-on. In doing so, we make three main contributions. First, we derive a cost-of-living index that is consistent with nonhomothetic consumer demand theory. 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) and Vartia (1976) indices. Second, we outline a feasible strategy to compute these price indices without being at the mercy of estimating entire demand systems. Instead, under a relatively mild assumption, estimation reduces to two parameters which are identified from a single equation. Third, we implement this approach using consumption and CPI data to investigate US inflation heterogeneity over the last quarter century.

Our framework allows for a characterization of the cost of living for the full expenditure distribution as well as at the aggregate level. We achieve this by deriving cost-of-living indices from a specification of Muellbauer's (1975, 1976) "price independent generalized linearity" (PIGL) preferences that has recently gained popularity in the structural change literature.[1] These preferences are nonhomothetic but maintain tractable aggregation properties that allow us to account for consumer heterogeneity. Like many conventional price indices, we show that the PIGL preferences induce cost-of-living indices that are weighted geometric means of individual price changes. Unlike their homothetic counterparts, however, the weights on these price changes vary systematically across the expenditure distribution. Specifically, richer households allocate higher weights to price changes of luxury goods. Changes in the cost of living are consequently allowed to differ with the expenditure level.

Due to the nonhomothetic nature of the underlying PIGL preferences, the cost-of-living index in its most general form is not directly computable without estimating a complete consumer demand system. We overcome this hurdle by imposing a key assumption: that preferences are weakly separable into necessities and luxuries. Under the weak separability assumption, the cost-of-living index reduces to observed price changes and expenditure shares and only two unknown parameters, which are readily estimated by linear as well as nonlinear regression methods. As in for instance Wachter and Yogo (2010) and Orchard (2021), classification of individual goods as "necessity" or "luxury" is straightforwardly done by investigating Engel curves. The cost-of-living index still nests homothetic price indices as special cases, so the assumption of weakly separable

[1] See for instance Boppart (2014), Alder, Boppart and Müller (2022), and Cravino, Levchenko and Rojas (forthcoming).

preferences is not a hard restriction when compared to conventional price indices. In our empirical application, we also show that it is well justified in the data.

We illustrate this implementation method by an empirical analysis of US inflation heterogeneity over the years 1995 to 2020. In this exercise, we consider twenty-one consumption good categories from the Consumer Expenditure Survey (CEX) that we match with corresponding CPI sub-indices. We obtain three main empirical results. First, households in the first expenditure decile ("the poor") experienced similar inflation as households in the tenth expenditure decile ("the rich") between 1995 and 2020, although substantial differences arise in the period around the Great Recession. In particular, poor households faced on average a 0.37 percentage points higher annual inflation rate between 2004 and 2015. Second, while the average inflation *rate* is relatively similar between the poor and the rich, inflation *volatility* is more than 2.5 times higher for the poor. Third, we find that this higher volatility primarily stems from a larger exposure for the poor to price changes in food, gas and utilities.

Decomposing the price index, we find that the overall development is almost entirely driven by mechanical price changes on the base-period consumption basket. Substitution behavior, as relative prices change, plays only a minor role. Yet, both differential base-period reference baskets and differential substitution are significant drivers in explaining the *differences* between groups, as poor households substitute away from expensive goods to a larger degree than the rich.

Furthermore, we exploit the fact that the price index can be used to retrieve inflation for a sequence of varying real expenditures and document the average inflation experienced over the life-cycle. The results from this analysis indicate that the inflation volatility of the young and poor is somewhat dampened by their life-cycle path of expenditures, while this is not the case for the old and poor. Overall, aging effects do not substantially change the fact that households who are initially poor experience a 2.5 times higher expected volatility of inflation compared to households who are initially rich.

This paper falls within an old literature on the economic approach to price index theory following Konüs (1939), Samuelson and Swamy (1974), Diewert (1976, 1978), Feenstra (1994), Redding and Weinstein (2020), and many others, whereby cost-of-living indices are derived from consumer theory via the expenditure function.[2] Central to this line of research is Diewert's notion of a *superlative* price index, which includes indices that are exact for some homothetic expenditure function and can approximate other homothetic indices 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 these and all other currently known indices within this class.

Nonhomothetic preferences have generally received little attention in the price index literature. Feenstra and Reinsdorf (2000) derive an index for Deaton and Muellbauer's (1980) almost ideal demand system (AIDS) and Oulton (2012) proposes a numerical algorithm to calculate indices of Banks, Blundell and Lewbel's (1997) generalization of the AIDS. The AIDS is a special case of the PIGLOG class of preferences, which is itself a limit case of the PIGL preferences used here. Unlike the PIGL class, the AIDS does not provide a straightforward generalization of

[2] Diewert (1993) surveys the early stages of this literature, which is far too large for us to do justice to here.

conventional price indices. Redding and Weinstein (2020) also derive a theoretical price index for the nonhomothetic CES specification of Hanoch (1975) and Sato (1975). In contrast to our index, however, the nonhomothetic CES specification does not consistently aggregate across heterogeneous consumers and (to the best of our knowledge) provides no easy implementation empirically without being forced to estimate all parameters in the utility function.

We also add to an empirical strand of literature concerned with inflation inequality which dates back at least to the 1950s (see for instance Muellbauer, 1974, and references therein), with recent advances surveyed by Jaravel (2021). The bulk of this literature approximates nonhomothetic cost-of-living indices by computing conventional homothetic price indices separately for different income groups.[3] This "group-specific" approach posits that differences in inflation are driven by differences in *ex ante* tastes and rests on a theoretical foundation where deep preference parameters jump discontinuously between groups.[4] These discrete jumps set aside straightforward comparisons between groups as well as to an aggregate inflation rate, the latter being completely lost. By contrast, inflation heterogeneity here stems endogenously from differences in expenditures. This allows us to characterize the full inflation distribution as well as aggregate inflation, with clear-cut comparisons between subgroups. Nevertheless, the group-specific approach provides an easy way to obtain index weights that vary with income, a feature that lies at the core of our framework. We therefore do not necessarily see these approaches as direct substitutes. Instead, our framework rationalizes previous empirical methodologies within a consistent theory for nonhomothetic consumer demand.

The paper proceeds as follows. Section 2 covers the theoretical framework, derives the nonhomothetic price index and shows how all superlative price indices can be generalized to a nonhomothetic setting. Section 3 outlines the strategy of our empirical implementation and discusses the assumption we make to render the demand system estimation feasible. Section 4 explains the data we employ, classifies the twenty-one goods into necessities and luxuries and reports estimates from the tractable demand system estimation. Section 5 reports the main empirical results, while Section 6 decomposes the price index and inflation further. Section 7 compares the main results to traditional demand system estimation and Section 8 concludes.

2 Theoretical Framework

The framework we consider is one where consumers maximize utility over a set of goods J with a corresponding price vector $\boldsymbol{p}$ and where we wish to investigate the change in the cost of living between a period t and some base period s. In what follows, we drop time subscripts whenever possible to simplify notation, as long as this causes no confusion. The minimum expenditure e required to obtain some utility level u when faced by the price vector $\boldsymbol{p}$ is given by the expenditure function $e = c(u, \boldsymbol{p})$. 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

[3] Recent papers employing this approach include Hobijn and Lagakos (2005), McGranahan and Paulson (2005), Broda and Romalis (2009), Kaplan and Schulhofer-Wohl (2017), Jaravel (2019), Orchard (2020, 2021), Argente and Lee (2021), and Lauper and Mangiante (2021).

[4] The implication being, for instance, that poor consumers allocate a larger share of expenditures to, say, rented housing due to some innate preference for rental homes obtained at birth.

a constant utility level:

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

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

2.1 The Homothetic Case

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,

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

where $B(\boldsymbol{p})$ 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 $c(u, \boldsymbol{p}) = (1 + \varepsilon u)^{1/\varepsilon} B(\boldsymbol{p})$ and using Equation (1), we get

$$P_t = \frac{B(\boldsymbol{p}_t)}{B(\boldsymbol{p}_s)}, \tag{3}$$

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),

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

where $B(\boldsymbol{p})$ and $D(\boldsymbol{p})$ are linearly homogeneous functions of prices and the parameters satisfy $\varepsilon, \gamma \in (0, 1)$ 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). Despite being nonhomothetic, these preferences consistently aggregate across individual-level expenditures. Aggregate expenditure shares in this case correspond to a representative expenditure level which is independent of prices and given by the average expenditure level multiplied by a simple inequality measure.

To gain understanding and intuition of Equation (4), it is convenient to think of $B(\boldsymbol{p})$ and $D(\boldsymbol{p})$ as the expenditure functions of some homothetic subutility functions. We refer to these subutility functions as "goods" or "baskets". The parameter ε controls the degree of nonhomotheticity between the D and B baskets: 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 $\varepsilon \to 0$, 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 $B(\boldsymbol{p}) = D(\boldsymbol{p})$ or in the limit case $\nu \to 0$. 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.

In general, there is nothing restricting an individual good j 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 goods are not directly observable and we obtain what Blundell and Robin (2000) call "latent separability". Latent separability is equivalent to weak separability, but in the latent goods B and D rather than in purchased goods, with weak separability as 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 where 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 shares w_D and w_B allocated to the D and B baskets in the first stage are therefore given by

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

and

$$w_B = 1 - \nu \left(\frac{B(\boldsymbol{p})}{e}\right)^{\varepsilon} \left(\frac{D(\boldsymbol{p})}{B(\boldsymbol{p})}\right)^{\gamma}. \tag{6}$$

Similarly, the shares w_j^D and w_j^B of total D and B expenditures allocated to individual good j are given by

$$w_j^D = p_j \frac{D_j(\boldsymbol{p})}{D(\boldsymbol{p})} \qquad \text{and} \qquad w_j^B = p_j \frac{B_j(\boldsymbol{p})}{B(\boldsymbol{p})}, \tag{7}$$

where D_j and B_j denote, respectively, the partial derivatives of D and B with respect to p_j. Equations (5) to (7) imply an expenditure share w_j of good j in total expenditures of the form

$$w_j = p_j \left[\frac{B_j(\boldsymbol{p})}{B(\boldsymbol{p})} + \left(\frac{D_j(\boldsymbol{p})}{D(\boldsymbol{p})} - \frac{B_j(\boldsymbol{p})}{B(\boldsymbol{p})} \right) \nu \left(\frac{B(\boldsymbol{p})}{e}\right)^{\varepsilon} \left(\frac{D(\boldsymbol{p})}{B(\boldsymbol{p})}\right)^{\gamma} \right]. \tag{8}$$

Therefore, nonhomotheticity between B and D also creates nonhomothetic behavior across individual goods, with a good j being a necessity if $D_j/D > B_j/B$ and a luxury vice versa. Aggregating over any N number of consumers indexed by h, the aggregate expenditure share $\overline{w}_j$ of good j across these consumers is

$$\overline{w}_j = p_j \left[\frac{B_j(\boldsymbol{p})}{B(\boldsymbol{p})} + \left(\frac{D_j(\boldsymbol{p})}{D(\boldsymbol{p})} - \frac{B_j(\boldsymbol{p})}{B(\boldsymbol{p})} \right) \nu \left(\frac{B(\boldsymbol{p})}{\bar{e}}\right)^{\varepsilon} \left(\frac{D(\boldsymbol{p})}{B(\boldsymbol{p})}\right)^{\gamma} \kappa \right], \tag{9}$$

where $\bar{e} \equiv \frac{1}{N} \sum_{h=1}^{N} e_h$ is the average expenditure level and κ is an inequality measure defined by

$$\kappa \equiv \frac{1}{N} \sum_{h=1}^{N} \left(\frac{e_h}{\bar{e}}\right)^{1-\varepsilon}. \tag{10}$$

Aggregate shares $\overline{w}_B$ and $\overline{w}_D$ of the B and D baskets are defined similarly. See Alder, Boppart and Müller (2022, Proposition 2) for a derivation of Equations (9) and (10). The representative agent in Muellbauer's (1975, 1976) sense (henceforth the PIGL RA) is the expenditure level e^{RA} that induces the aggregate expenditure share. By Equation (9), this expenditure level is given by $e^{RA} \equiv \bar{e}\kappa^{-1/\varepsilon}$.

2.3 The Nonhomothetic Case: Price Index

The indirect utility function (4) allows us to extend the homothetic cost-of-living index (3). Unlike the homothetic case, the index now depends on a base-period standard of living, represented by the utility level u in the Konüs definition (1). Because preferences are nonhomothetic over the B and D baskets, with an expenditure elasticity of demand for the D basket always less than 1, this utility level is fully captured by w_{Ds}, the base-period expenditure share of the D good. For the remainder of the paper, let

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

denote the logarithmic mean (Carlson, 1972). Moreover, recalling from Equation (3) that the cost-of-living index of homothetic preferences between some period t and base period s is the ratio of the corresponding expenditure functions over the same periods, we denote the price indices of the B and D baskets by $P_{Bt} \equiv B(\boldsymbol{p}_t)/B(\boldsymbol{p}_s)$ and $P_{Dt} \equiv D(\boldsymbol{p}_t)/D(\boldsymbol{p}_s)$. The following result then shows that the cost-of-living index corresponding to Equation (4) is a function of w_{Ds}, P_{Bt} and P_{Dt} and the two parameters ε and γ.

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*

$$P_t^{PIGL} = P_{Dt}^{\frac{\gamma\phi_t}{\varepsilon}} P_{Bt}^{1-\frac{\gamma\phi_t}{\varepsilon}} \qquad \textit{with} \qquad \phi_t \equiv \frac{L(\psi_{Dt},\ \psi_{Ds})}{L(\psi_{Dt},\ \psi_{Ds})\ +\ L(\psi_{Bt},\ \psi_{Bs})}, \tag{11}$$

where

$$\psi_{Bt} \equiv \left(1 - \frac{\varepsilon w_{Ds}}{\gamma}\right)\left(\frac{P_{Bt}}{\widetilde{P}_t}\right)^{\gamma} \qquad \textit{and} \qquad \psi_{Dt} \equiv \frac{\varepsilon w_{Ds}}{\gamma}\left(\frac{P_{Dt}}{\widetilde{P}_t}\right)^{\gamma} \tag{12}$$

are shares of a CES-type aggregator $\widetilde{P}_t$ defined by

$$\widetilde{P}_t \equiv \left[\left(1 - \frac{\varepsilon\, w_{Ds}}{\gamma}\right) P_{Bt}^{\gamma} + \frac{\varepsilon\, w_{Ds}}{\gamma} P_{Dt}^{\gamma}\right]^{\frac{1}{\gamma}}. \tag{13}$$

The aggregate cost-of-living index over any N number of consumers is given identically using their average expenditure share $\overline{w}_{Ds}$ in ϕ_t.

Sketch proof (full proof in Appendix A.1). Set the reference utility to that of the base period expenditure level, $u \equiv V(e_s, \boldsymbol{p}_s)$. It is then possible to write the period-t expenditure function corresponding to Equation (4) as $c(u, \boldsymbol{p}_t) = c(u, \boldsymbol{p}_s)\widetilde{P}_t^{\frac{\gamma}{\varepsilon}} P_{Bt}^{1-\frac{\gamma}{\varepsilon}}$. Since $\widetilde{P}_t$ is of a CES form, it can be rewritten as a Sato-Vartia index with weights $1 - \phi_t$ and ϕ_t on P_{Bt} and P_{Dt}, respectively. The result then follows from the Konüs definition (1). □

Proposition 1 shows that the PIGL cost-of-living index can be written as something akin to a Sato-Vartia index over the B and D baskets. Unlike the homothetic case, however, the ϕ_t in the weights of the two subindices varies across the expenditure distribution. Richer consumers spend a smaller share w_D on the D basket, which reduces the weights ψ_{Dt} and, subsequently, ϕ_t. In other words, because richer consumers allocate a smaller share to the D basket, the corresponding price index P_{Dt} is weighted less heavily when determining the overall change in the cost of living. The weights ψ_{Dt} and ψ_{Bt} are not directly observable but are readily computed given price indices P_{Bt} and P_{Dt}, an expenditure share w_{Ds}, and parameter values for ε and γ. In Appendix A.1, we show that $w_{Ds}(P_{Dt} \,/\, \widetilde{P}_t)^{\gamma}$ is the expenditure share of the D basket at period-t prices that prevails at the same utility level as w_{Ds}. Therefore, the weights ψ_{Dt} and ψ_{Bt} ensure that the consumer remains on the same indifference curve as in the base period.

Two potential caveats to Proposition 1 are that the underlying preferences are identical across consumers with the same expenditure level and that expenditure shares change monotonically in the level of expenditure. Redding and Weinstein (2020) emphasize accounting for taste heterogeneity in cost-of-living indices while Banks, Blundell and Lewbel (1997) highlight the importance of allowing for hump-shaped expenditure shares to match microeconomic data. In Appendix A.4 we show that it is straightforward to incorporate time- and household-specific tastes between the B and D baskets into the indirect utility function (4) and that this leaves Proposition 1 unaffected. In Appendix A.5 we discuss a generalization that allows for hump-shaped expenditure shares. This generalization works well for household-level indices but requires more stringent conditions for aggregate price indices to have the same form.

2.4 Generalized Superlative Indices

With Proposition 1 at hand, it is straightforward to generalize standard homothetic indices to the nonhomothetic PIGL case: just plug in two homothetic indices for P_{Bt} and P_{Dt} in Equation (11). To emphasize the importance of Proposition 1, we present generalizations of two classes of indices: Diewert's (1976) *quadratic-mean-of-order-r* class, which consists of all indices of the form

$$P_t = \sqrt{\left\{\sum_{j\in J} w_{js}\left(\frac{p_{jt}}{p_{js}}\right)^{\frac{r}{2}}\right\}^{\frac{2}{r}}\left\{\sum_{j\in J} w_{jt}\left(\frac{p_{jt}}{p_{js}}\right)^{-\frac{r}{2}}\right\}^{-\frac{2}{r}}}, \qquad r > 0, \tag{14}$$

and Barnett and Choi's (2008) *Theil-Sato* class, which is defined as

$$P_t = \prod_{j\in J}\left(\frac{p_{jt}}{p_{js}}\right)^{\delta_{jt}}, \qquad \delta_{jt} \equiv \frac{m(w_{jt}, w_{js})}{\sum_{i\in J} m(w_{it}, w_{is})}, \tag{15}$$

where $m(x, y)$ is a *symmetric mean* of two variables, a function class that includes all linearly homogenous functions satisfying $\min\{x, y\} \leq m(x, y) = m(y, x) \leq \max\{x, y\}$. These index classes include several of the most well-known price index formulas. Equation (14) incorporates Fisher's (1922) ideal index ($r = 2$), the arithmetic Walsh (1901) index ($r = 1$) and, as a limit case, the Törnqvist (1936) index ($r \to 0$). Equation (15) nests the Törnqvist index (arithmetic mean, $m(x, y) = (x + y)/2$), the geometric Walsh (1901) index (geometric mean, $m(x, y) = \sqrt{xy}$), the Sato (1976)-Vartia (1976) index (logarithmic mean, $m(x, y) = (x - y)/(\ln x - \ln y)$), 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 consists of all currently known *superlative* price indices (Diewert, 1976). That is, they are exact cost-of-living indices for some homothetic expenditure functions and are second-order approximations of any other homothetic price index.[5] Therefore, even if the indices corresponding to the true expenditure functions $B(\boldsymbol{p})$ and $D(\boldsymbol{p})$ have some other forms than those in Equations (14) and (15), we should still be able to reasonably approximate their corresponding price indices under this specific parameterization. The nonhomothetic generalization of these indices under Proposition 1 is presented below.

Corollary 1 (Generalized superlative indices)**.** *If preferences are of the PIGL form (4), the base-period expenditure share w_{Ds} allocated to the D basket is given, and $B(\boldsymbol{p})$ and $D(\boldsymbol{p})$ are expenditure functions with price indices of the form (14) or (15), the Konüs cost-of-living index is*

$$P_t^{G\text{-}S} = \prod_{j \in J} \left(\frac{p_{jt}}{p_{js}} \right)^{\chi_{jt}}, \tag{16}$$

where

$$\chi_{jt} \equiv \frac{\gamma \phi_t}{\varepsilon} \, \delta_{jt}^D + \left(1 - \frac{\gamma \phi_t}{\varepsilon} \right) \delta_{jt}^B, \tag{17}$$

with ϕ_t as in Proposition 1. The weights δ_{jt}^C, $j \in J$, $C \in \{B, D\}$, are given by

$$\delta_{jt}^C = \frac{1}{2} \left[\frac{\widetilde{w}_{Ljt}^C}{\sum_i \widetilde{w}_{Lit}^C} + \frac{\widetilde{w}_{Pjt}^C}{\sum_i \widetilde{w}_{Pit}^C} \right] \qquad \text{or} \qquad \delta_{jt}^C = \frac{m(w_{jt}^C, w_{js}^C)}{\sum_i m(w_{it}^C, w_{is}^C)} \tag{18}$$

if P_{Ct} is as in Equation (14) or (15), respectively. In the former case,

$$\widetilde{w}_{Ljt}^C \equiv w_{js}^C \, L\left(\left(\tfrac{p_{jt}}{p_{js}} \right)^{\frac{r}{2}}, \left(P_{Lt}^C \right)^{\frac{r}{2}} \right) \qquad \text{and} \qquad \widetilde{w}_{Pjt}^C \equiv w_{jt}^C \, L\left(\left(\tfrac{p_{jt}}{p_{js}} \right)^{-\frac{r}{2}}, \left(P_{Pt}^C \right)^{-\frac{r}{2}} \right),$$

where $P_{Lt}^C = \left[\sum_j w_{js}^C \left(\frac{p_{jt}}{p_{js}} \right)^{\frac{r}{2}} \right]^{\frac{2}{r}}$ and $P_{Pt}^C = \left[\sum_j w_{jt}^C \left(\frac{p_{jt}}{p_{js}} \right)^{-\frac{r}{2}} \right]^{-\frac{2}{r}}$. The aggregate cost-of-living index over any N number of consumers is given identically using their average expenditure shares in χ_{jt}.

Proof. In Appendix A.2. □

[5] This definition differs from Diewert's original definition but is shown by Barnett and Choi (2008, Theorem 1) to be equivalent.

In Corollary 1, we have rewritten the *quardratic-mean-of-order-r* class on a geometric-mean form following Balk (2004) to highlight the intuitive generalization of the homothetic superlative indices that we obtain. In doing so, we denote the weights by L and P to capture the fact that these reduce to standard Laspeyres and Paasche weights when $r = 2$. The resulting cost-of-living index is a weighted geometric average of individual price changes with the index weights (17) of the following structure:

$$\begin{array}{c}\text{Weight}\\ \text{on } j\end{array} = \begin{array}{c}\text{Weight}\\ \text{on } D\end{array} \times \begin{array}{c}\text{Weight}\\ \text{on } j\\ \text{within } D\end{array} + \begin{array}{c}\text{Weight}\\ \text{on } B\end{array} \times \begin{array}{c}\text{Weight}\\ \text{on } j\\ \text{within } B\end{array}.$$

The weights on j within B and D, given by Equation (18), are standard homothetic weights and affect all consumers similarly. The weights on D and B are the same as in Proposition 1. Therefore, the overall weights χ_{jt} vary across the base-period expenditure distribution in a similar way as before. If $B(\boldsymbol{p}) = D(\boldsymbol{p})$, we get that $\delta_{jt}^{B} = \delta_{jt}^{D}$ for all j and the generalized superlative indices immediately collapse to the homothetic indices in Equations (14) and (15).

Another feature of Corollary 1 is that it rationalizes the methodology used in much of the literature concerned with inflation inequality, whereby homothetic price indices are computed for different income groups separately. In particular, papers like Broda and Romalis (2009), Jaravel (2019), and Argente and Lee (2021) construct homothetic price indices of the geometric-mean form $\ln P_t = \sum_j \delta_{jt} \ln(p_{jt}/p_{js})$, where δ_{jt} are weights computed separately for each income group considered. This generates heterogeneous weights across the income distribution. The method therefore mimics an overall geometric-mean price index with income specific weights, which is exactly what we also have in Corollary 1. In contrast to the group-specific approach, however, Corollary 1 allows for a full characterization of the inflation distribution rather than a discontinuous, discrete distribution.

3 Empirical Implementation

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 the B and the D baskets, these across and within expenditure shares are unobserved in the data. The only feasible approach then is to parameterize $B(\boldsymbol{p})$ and $D(\boldsymbol{p})$, estimate the demand system associated with the expenditure share equations (8) via GMM, and infer these shares from the estimated model. This methodology, however, suffers from the usual drawbacks 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.[6] The nonlinear nature of the demand system also implies that there is no guarantee that the GMM estimator converges to the actual global minimum of the GMM objective function. The latter could in principle be solved by a grid search, but this only exacerbates the curse of dimensionality further. Estimating more than a few goods is therefore generally infeasible. These issues, however, are fully circumvented within our framework when a simple assumption on the structure of the demand system is met.

[6] As an illustration, suppose we have n goods and parameterize $B(\boldsymbol{p})$ and $D(\boldsymbol{p})$ by the linearly homogeneous translog expenditure function of Christensen, Jorgenson and Lau (1975), for which the Törnqvist index is an exact cost-of-living index (Diewert, 1976). The PIGL demand system considered here then requires the estimation of $n(n+1) + 3$ independent parameters.

Assumption 1. Preferences are weakly separable into the B and D baskets. ◁

Under Assumption 1, an individual good occurs in either the B basket or the D basket, but not in both. Since the D basket captures necessity needs and B basket luxury needs, it follows that preferences are also weakly separable into necessities and luxuries. The 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 price indices P_{Bt} and P_{Dt} for the B and D baskets using Equation (14) or (15). The only additional information needed to compute the generalized superlative indices are, per Proposition 1, the two parameters ε and γ. Using Equations (3) and (5), we may write the period-t expenditure share on the D good as

$$w_{Dt} = \widetilde{\nu} \left(\frac{P_{Bt}}{e_t} \right)^{\varepsilon} \left(\frac{P_{Dt}}{P_{Bt}} \right)^{\gamma}, \tag{19}$$

where $\widetilde{\nu} \equiv \nu B(\boldsymbol{p}_s)^{\varepsilon-\gamma} D(\boldsymbol{p}_s)^{\gamma}$ is a scale parameter. Since w_{Dt}, e_t, P_{Bt} and P_{Dt} are all known, estimating ε and γ from (19) is easily carried out using either linear (by taking logs of (19)) or nonlinear estimation methods. We summarize this empirical approach by the following proposition:

Proposition 2 (Tractable demand system estimation)**.** *Under Assumption 1, across and within group expenditure shares are observable in the data and computing the generalized superlative indices (16) only requires estimation of two parameters, ε and γ, from the single expenditure share equation (19).*

At first sight, Assumption 1 may seem to be at odds with the nonhomothetic generalization of the superlative indices: Corollary 1 reduces to the standard homothetic case when $B(\boldsymbol{p}) = D(\boldsymbol{p})$, which Assumption 1 excludes by construction. However, the weak separability assumption still nests homothetic preferences. In particular, as $\varepsilon \to 0$ and $\gamma \to 0$, we obtain Cobb-Douglas preferences with $V(e, \boldsymbol{p}) = \ln \left[\frac{e}{B(\boldsymbol{p})^{1-\nu} D(\boldsymbol{p})^{\nu}} \right]$ and a corresponding price index $P_t = P_{Bt}^{1-\nu} P_{Dt}^{\nu}$, where $\nu = w_D$ is the homothetic and time-invariant expenditure share on D. Thus, if preferences truly are homothetic, we still expect Proposition 2 to yield a homothetic price index. This index is approximately equal to the corresponding superlative index when $B(\boldsymbol{p}) = D(\boldsymbol{p})$, by virtue of superlative indices being second-order approximations of any other homothetic index. 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 $B(\boldsymbol{p}) = D(\boldsymbol{p})$: the Törnqvist index.

Proposition 3 (Homothetic Törnqvist index under weak separability)**.** *Suppose that preferences are of the homothetic Cobb-Douglas form,* $V(e, \boldsymbol{p}) = \ln\left[\frac{e}{B(\boldsymbol{p})^{1-\nu} D(\boldsymbol{p})^{\nu}}\right]$*, and that* $B(\boldsymbol{p})$ *and* $D(\boldsymbol{p})$ *are such that their corresponding price indices* P_{Bt} *and* P_{Dt} *are Törnqvist indices. The cost-of-living index under Assumption 1 is then the standard Törnqvist index:*

$$P_t = \prod_{j \in J} \left(\frac{p_{jt}}{p_{js}}\right)^{\delta_{jt}}, \qquad \textit{where} \qquad \delta_{jt} = \frac{w_{js} + w_{jt}}{2}.$$

Proof. In Appendix A.3. □

4 Data and Estimation

We implement the tractable demand system estimation described in the previous section 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 a rather coarse set of consumption goods categories as this allows us to compare the empirical approach in Proposition 2 with a full demand system estimation. All in all, we consider twenty-one expenditure categories using the hierarchical groupings defined by the BLS. We broadly follow Hobijn and Lagakos (2005) and construct prices for these categories by matching them with individual CPI series. Table B.1 in Appendix B lists the CEX categories and shows their mapping to the CPI item codes.

4.1 Classification of Goods Into Luxuries and Necessities

In order to utilize the tractable demand system estimation in Proposition 2, we impose Assumption 1 by allocating luxuries to B and necessities to D. The classification into B and D

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:

$$w_{jh} = \alpha_j + \beta_j \mathsf{d}_h + \epsilon_{jh}.$$

If $\beta_j > 0$, we allocate the good to the B basket, otherwise to the D basket. Figure 1 shows the Engel curves by expenditure decile together with the resulting classification from the regressions and Table B.2 in Appendix B lists the β_j coefficient estimates. The resulting classification is intuitive and all estimates are significantly different from zero. For comparable product groups, our necessity/luxury split is highly similar to those constructed in similar analyses using CEX data (see for instance Wachter and Yogo, 2010, and Orchard, 2021), thus suggesting that this simple classification regression works well on our coarse set of goods.

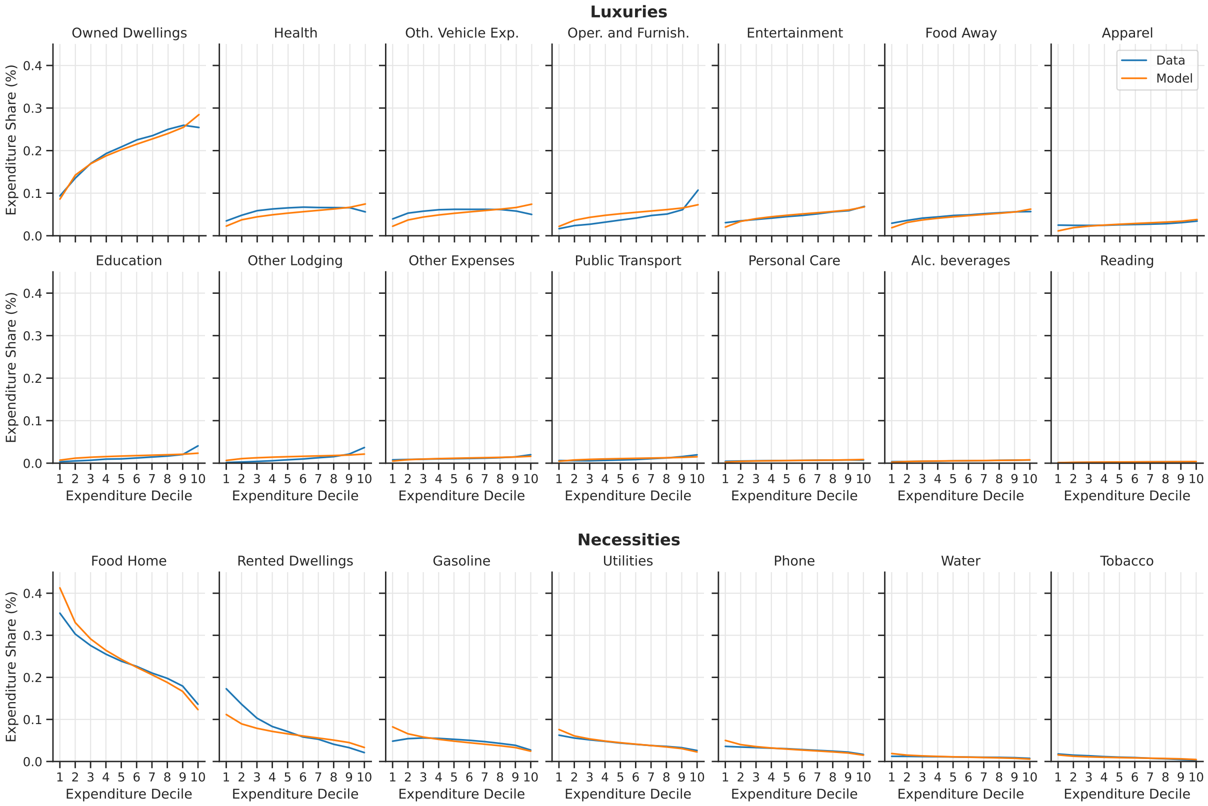

FIGURE 1. Empirical and model implied Engel curves.

Notes. The figure shows the empirical and model implied expenditure shares by expenditure group and expenditure decile averaged over all years. The model implied expenditure shares correspond to those under the assumption of weak separability and the Sato-Vartia specification. They are calculated by first taking the model implied expenditure shares on the B and D goods and then use the empirical expenditures shares of all households to obtain shares within these groups.

4.2 Tractable Demand System Estimation

We estimate the preference parameters ε and γ from Equation (19) using a nonlinear GMM estimator on household-level 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, 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 the literature (Blundell, Pashardes and Weber, 1993; Banks, Blundell and Lewbel, 1997) and control for this endogeneity bias by instrumenting expenditures with household income.

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 a necessary and sufficient condition for household h to satisfy the Slutsky restrictions in period t is $\widetilde{\nu}\left(\frac{P_{Bt}}{e_{ht}}\right)^{\varepsilon}\left(\frac{P_{Dt}}{P_{Bt}}\right)^{\gamma} \leq (1-\gamma)/(1-\varepsilon)$. We enforce this constraint by augmenting the GMM estimation with a classic penalty method. 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 arithmetic), the Theil, and the Fisher. 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 statistically different from zero at conventional significance levels and the fact that $\varepsilon > 0$ and $\widetilde{\nu} > 0$ 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.[7] Moreover, Alder, Boppart and Müller (2022, Proposition 3) show that a sufficient condition for expenditure shares to remain globally nonnegative 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 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. We construct the model expenditure shares as the product of model-implied across- and within-expenditure shares w_C and w_j^C for goods

[7] Further note that the parameter estimates are also robust to not using the CEX household sampling weights.

in $C \in \{B, D\}$. The former is computed from Equation (19) at the representative level of expenditures within each expenditure decile. Since $B(\boldsymbol{p})$ and $D(\boldsymbol{p})$ are homothetic, the latter is given empirically by the average within-shares $\overline{w}_j^B$, $\overline{w}_j^D$ across *all* households. Figure 1 shows that the model fits the data quite well. In particular, the model does a much better job at matching the empirical Engel curves than the constant Engel curves resulting from homothetic preferences would. This underscores that the assumption of weak separability between the B and D baskets is not a strong restriction in our sample.

TABLE 1. GMM estimates under weak separability.

	Sato-Vartia	Törnqvist	Geom. Walsh	Theil	Fisher	Arith. Walsh
ε	0.677	0.677	0.677	0.677	0.677	0.677
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
γ	0.211	0.211	0.211	0.211	0.211	0.211
	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)
$\widetilde{\nu}$	327.271	327.437	327.173	327.273	324.273	324.600
	(13.358)	(13.365)	(13.354)	(13.358)	(13.217)	(13.233)
N	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 squared error of the expenditure share on the D good: $\sqrt{\sum_h (w_{Dh} - \widehat{w}_{Dh})^2 / N}$. Each household-level observation is weighted by the CEX sampling weight.

5 Results

The estimation in the previous section suggests that the results are insensitive to the choice of generalized superlative index.[8] For the remainder of the paper, we therefore focus on the nonhomothetic generalization of the Sato-Vartia index (henceforth G-SV). Since the Sato-Vartia index is the corresponding Konüs index 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 therefore makes this a case of particular interest. The main reason for selecting the G-SV, however, is that a parameterization of $B(\boldsymbol{p})$ and $D(\boldsymbol{p})$ 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 G-SV index to a relatively parsimonious estimation of the full PIGL demand system.

5.1 Cost-of-living and inflation inequality

Figure 2 shows the evolution of the G-SV 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}.[9] Even though the generalized superlative price indices allow for characterizations of the entire distribution of indices, here we focus on expenditure deciles for ease of exposition. Figures B.3 and B.4 in Appendix B show the full price index and inflation distributions.

[8] Figures B.7 and B.8 in Appendix B show the main results in this section for all choices considered in Section 4.2, which confirms that this is indeed the case.

[9] Figure B.1 in Appendix B shows that the choice of base period does not affect our results.

Figure 2 corroborates two findings from the literature: inflation rates vary across households and poorer households 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. The changes in the cost of living over the 26 years under study therefore do not diverge dramatically between groups.

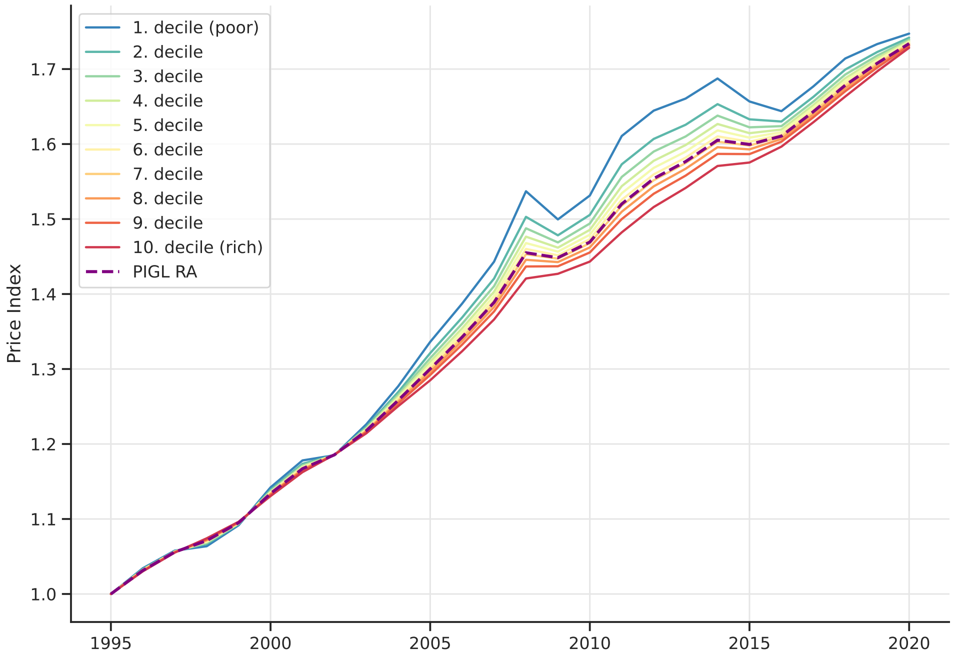

FIGURE 2. G-SV price index under weak separability by expenditure decile.

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. "PIGL RA" stands for the PIGL representative agent over *all* 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 across subperiods is also visible from the implied annual inflation rates. Figure 3 plots these rates, which highlights that there are periods where poorer households experience substantially higher *or* lower inflation. In several years, the range of the inflation rate exceeds 2 percentage points. There is another key finding

Table 2. Inflation heterogeneity in numbers.

	2004 - 2015				1996 - 2020			
Exp. dec.	mean	Δmean	std.	rel. std.	mean	Δmean	std.	rel. std.
1	2.57	0.37	2.71	2.38	2.28	0.06	2.14	2.51
2	2.46	0.26	2.25	1.98	2.26	0.04	1.76	2.06
3	2.41	0.21	2.05	1.79	2.25	0.03	1.59	1.86
4	2.38	0.18	1.89	1.66	2.24	0.03	1.46	1.71
5	2.35	0.15	1.78	1.56	2.24	0.02	1.36	1.60
6	2.33	0.13	1.67	1.46	2.24	0.02	1.28	1.50
7	2.30	0.10	1.57	1.38	2.23	0.02	1.20	1.40
8	2.28	0.08	1.47	1.29	2.23	0.01	1.12	1.31
9	2.25	0.05	1.35	1.18	2.22	0.01	1.02	1.20
10	2.20	0.00	1.14	1.00	2.22	0.00	0.85	1.00

Notes. Arithmetic mean and standard deviation of annual inflation over the respective years. "Δ mean" denotes the difference in the mean annual inflation to the tenth expenditure decile. Rel. std. denotes the relative standard deviation to the standard deviation of the tenth expenditure decile.

that stands out from Figure 3 however: the inflation rate of the poor is much more volatile than that of the rich. More precisely, Table 2 shows that the standard deviation of inflation has been 2.14 and 0.85 for the poorest and richest, respectively, and the poor have thus experienced a 2.5 times *more volatile* inflation rate than the rich.

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 difference in volatility of inflation rates warrants that it is important in future work to understand the economic implications it may have for instance in terms of welfare on individual household levels and economic policies that aim at stabilizing inflation rates.

5.2 Changes in Living Standards

Most empirical evidence on inflation inequality, including the results above, compare the change in inflation rates between two or more groups whose real expenditures stay fixed at some base-period level. In other words, living standards remain constant. Yet, it is a well-known that households experience substantial changes in expenditures over the life-cycle. It is therefore not clear whether the evidence on inflation inequality is of much relevance to the inflation experience of an average household.

Our framework distinguishes itself from previous attempts to study inflation inequality in that it is straightforward to compute inflation rates for households even when real expenditures change over time. The group-specific price index approach poses some challenges in allowing for this due to the discontinuous change underlying preferences across the expenditure distribution. Since our framework expresses the price index as a continuous function of expenditures we do not face that challenge.

Consider an individual whose expenditures, $e_t = c(u_t, \boldsymbol{p}_t)$, change over time. The change in

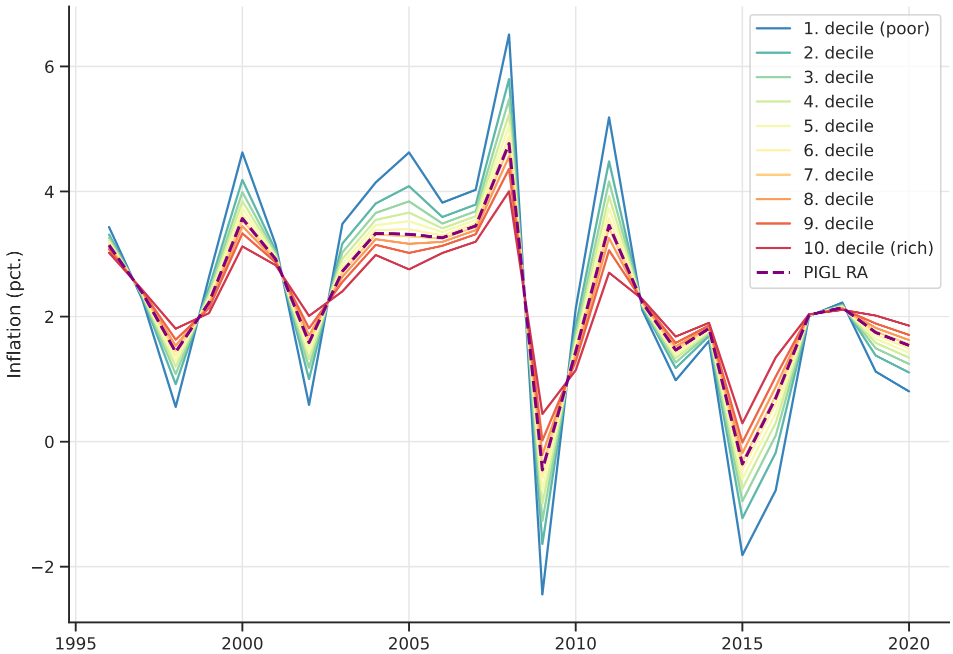

FIGURE 3. G-SV inflation under weak separability by expenditure decile.

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.

expenditures can arise because prices, $\boldsymbol{p}_t$, change or living standards, u_t, change. The relative change in expenditures between period s and period T may be written as

$$\frac{c(u_T, \boldsymbol{p}_T)}{c(u_s, \boldsymbol{p}_s)} = \left[\prod_{t=s+1}^{T} Q(u_{t-1}, u_t, \boldsymbol{p}_t)\right]\left[\prod_{t=s+1}^{T} \frac{P(u_{t-1}, \boldsymbol{p}_t, \boldsymbol{p}_s)}{P(u_{t-1}, \boldsymbol{p}_{t-1}, \boldsymbol{p}_s)}\right], \tag{20}$$

which is made up of two components: (i) a quantity cost, $Q(u_i, u_k, \boldsymbol{p}_t) \equiv c(u_k, \boldsymbol{p}_t)/c(u_i, \boldsymbol{p}_t)$, that tells how much a household will have to pay to go from living standard u_i to u_k at prices $\boldsymbol{p}_t$, and (ii) a per-period price cost, $P(u_{t-1}, \boldsymbol{p}_t, \boldsymbol{p}_s)/P(u_{t-1}, \boldsymbol{p}_{t-1}, \boldsymbol{p}_s)$, that tells how much the cost of living has changed period by period.[10] This latter part is just a chained price index where the base period is always $t-1$.[11]

From Equation (20), we can also compute the inflation cost from the per-period inflation rate

[10] The quantity-cost in Equation (20) measures the cost of obtaining living standard u_t relative to maintaining living standard u_{t-1} in present-period prices. That is, it measures the relative cost of obtaining living standard u_t compared to u_{t-1} in prices when the change takes place. A different definition, the quantity-cost in past-period prices defined as $Q_{t-1}(u_{t-1}, u_t)$, gives a similar expression as the one in Equation (20) but is less intuitive.

[11] When living standards do not change throughout period s to period T, we have that

$$\frac{c(u_T, \boldsymbol{p}_T)}{c(u_s, \boldsymbol{p}_s)} = \prod_{t=s+1}^{T} \frac{P(u_s, \boldsymbol{p}_t, \boldsymbol{p}_s)}{P(u_s, \boldsymbol{p}_{t-1}, \boldsymbol{p}_s)} = \frac{P_T(u_s, \boldsymbol{p}_T, \boldsymbol{p}_s)}{P(u_s, \boldsymbol{p}_s, \boldsymbol{p}_s)} = P(u_s, \boldsymbol{p}_T, \boldsymbol{p}_s),$$

which is simply the PIGL price index at time T with base period s.

$1 + \pi_t(u_{t-1}) \equiv P(u_{t-1}, \boldsymbol{p}_t, \boldsymbol{p}_s)/P(u_{t-1}, \boldsymbol{p}_{t-1}, \boldsymbol{p}_s)$. It forms the basis of the following analysis where we decompose the change in expenditures into quantity and price costs when allowing for changes in expenditures over time.

5.2.1 Life-Cycle Model

Changes in household expenditures over time is a well-documented empirical fact. Fernández-Villaverde and Krueger (2007) for example find that US household expenditures follow a deterministic, hump-shaped pattern over the life-cycle.

Based on our data, we estimate the following deterministic life-cycle model:

$$\ln e_{it} = \alpha_i + \beta_t + \widetilde{\gamma}_{it}, \tag{21}$$

where i denotes age, t time period, α_i are age-fixed effects, β_t year-month-fixed effects, and $\widetilde{\gamma}_{it}$ is the residual. The results of the estimated model in Equation (21) are shown in Figure B.5 of Appendix B. It shows that expenditures increase on average by 30 % for an individual between 25 and 60 years before they revert to the level of a 25-year-old at the age of 90. This is in line with what Fernández-Villaverde and Krueger (2007) find.

We use the estimated deterministic life-cycle model to simulate a 26-year nominal expenditure path for 30 different types. Each type has a distinct age-by-initial-expenditure-level combination, where the age is such that in 1995 the type is either young (25 years old), middle-aged (45 years old) or old (65 years old), and the initial expenditure level is equal to the level of one of the expenditure deciles in the CEX data as in the main results. For each type we then compute the price index, inflation series and standard deviation of inflation according to Equation (20).

In panel (a) of Figure 4 we show the simulated path of nominal expenditures for six types: young-poor, young-rich, middle-aged-poor, middle-aged-rich, old-poor and old-rich. For each age type, poor refers to an individual whose 1995 expenditure level was equal to that of expenditure decile 1 and rich refers to an individual whose 1995 expenditure level was equal to that of expenditure decile 10. The figure shows that nominal expenditures varies greatly across age groups and by initial expenditure levels.

Panel (b) of Figure 4 shows the corresponding inflation for each of the six types whose nominal expenditure levels are those in panel (a). The figure shows that different deterministic life-cycle patterns give rise to different inflation levels for individuals who start out with the same initial standard of living. For example, the young-poor experience a somewhat more dampened volatility of inflation over time compared to the old-poor. This is because the young experience changes in expenditures that exceed the change in prices and therefore allows them to increase their living standards. As our main results shows, households with higher living standards have experienced a much lower volatility of inflation and the young-poor transit towards this over time. The broad picture, however, is that the deterministic life-cycle components are far from being able to mitigate the results from the previous section: columns 3–6 in Table 3 show that the relative standard deviation of inflation remains 2.41 times higher for the young-poor compared to the young-rich, 2.46 for the middle-aged-poor and 2.55 times for the old-poor. Although not shown, the mean and standard deviation of inflation are also largely unchanged.

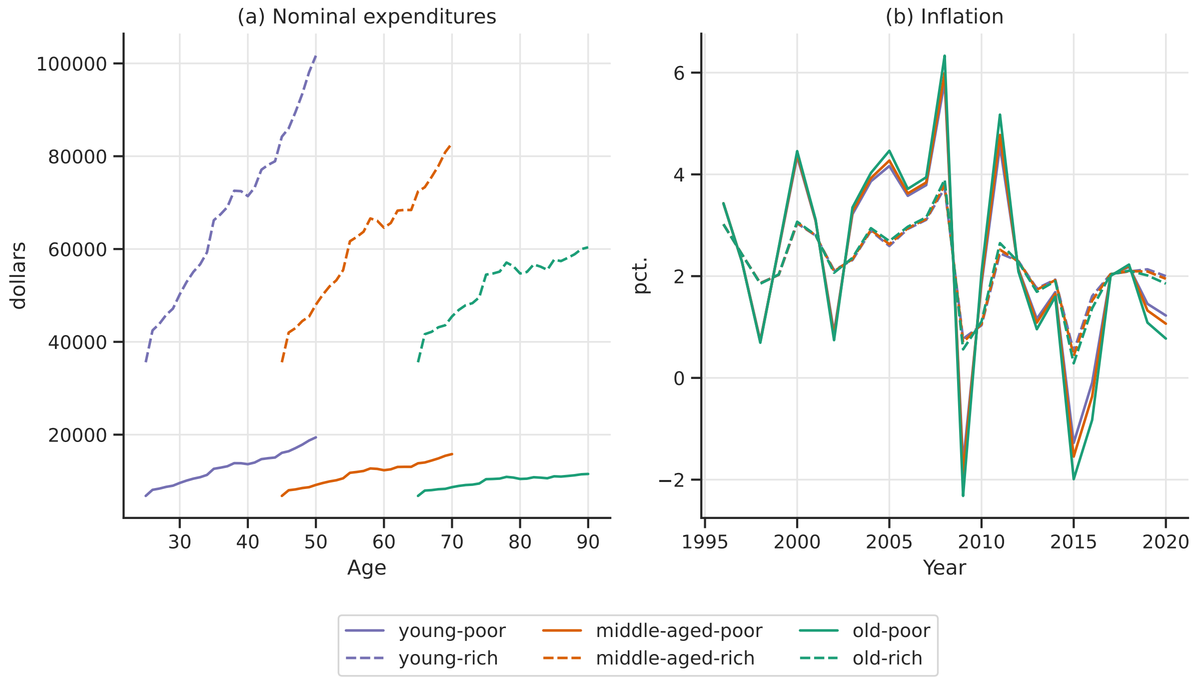

FIGURE 4. Expenditures and inflation in the deterministic life-cycle model.

Notes. Panel (a) shows the nominal expenditure levels for six types simulated from the life-cycle model in Equation (21). Panel (b) shows the corresponding inflation levels for each type based on Equation (20). Young refers to an individual who was 25 years in 1995, middle-aged to an individual who was 45 years and old to an individual who was 65 years. Poor refers to an individual whose 1995 expenditure level was equal to that of expenditure decile 1 in the CEX data. Rich refers to an individual whose 1995 expenditure level was equal to that of expenditure decile 10 in the CEX data.

In addition to the deterministic evolution in the life-cycle expenditures, other empirical studies such as Blundell, Pistaferri and Preston (2008) also find that US households experience significant stochastic shocks to expenditures.

The stochastic component comes from the unexplained change in log expenditures between two ages which is given by

$$\gamma_{it} \equiv \Delta\widetilde{\gamma}_{it} = \ln e_{it} - \ln e_{i-1,t-1} - (\alpha_i - \alpha_{i-1}) - (\beta_t - \beta_{t-1}), \tag{22}$$

and as in Blundell, Pistaferri and Preston (2008), we assume that γ is independently and normally distributed and we use their estimated variance of γ. In order to preserve the expected life-cycle expenditures from the estimated model, we impose the restriction that $\exp\{\gamma\}$ has a mean of one.

We once again consider the same 30 types as in the deterministic life-cycle model. For each type, we draw 26 γ-shocks and simulate a nominal expenditure path according to the stochastic life-cycle model. We repeat the simulation 10,000 times for each type and for each expenditure path we compute the corresponding price index, inflation series and standard deviation of inflation. We then average the 10,000 simulations to get the expected inflation for each type.

TABLE 3. Relative standard deviation.

Exp. dec.	Constant utility	Life-cycle young	Life-cycle middle-aged	Life-cycle old	Stochastic life-cycle young	Stochastic life-cycle middle-aged	Stochastic life-cycle old
1	2.51	2.41	2.46	2.55	2.50	2.55	2.63
2	2.06	1.97	2.01	2.08	2.05	2.08	2.13
3	1.86	1.78	1.81	1.87	1.84	1.88	1.91
4	1.71	1.64	1.67	1.72	1.70	1.72	1.76
5	1.60	1.54	1.56	1.60	1.59	1.61	1.64
6	1.50	1.44	1.46	1.50	1.47	1.50	1.53
7	1.40	1.36	1.37	1.40	1.39	1.40	1.43
8	1.31	1.27	1.29	1.31	1.30	1.31	1.33
9	1.20	1.17	1.18	1.19	1.19	1.19	1.21
10	1.00	1.00	1.00	1.00	1.00	1.00	1.00

Notes. The table reports the relative arithmetic standard deviations of inflation from 1995 to 2020 between each type whose 1995 expenditure level was equal to expenditure decile $\mathsf{d} \in \{1, \dots, 10\}$ and that of those whose 1995 expenditure level was equal to expenditure decile 10 in the CEX. "Constant utility" refers to the baseline results also presented in Table 2 where real expenditures are fixed. "Life-cycle" refers to the deterministic life-cycle model in Equation (21) where the determinsitic component is shut off (i.e. equal to 0). "Stochastic life-cycle" refers to the life cycle model in Equation (21) with the stochastic component activated. Young refers to an individual who was 25 years in 1995, middle-aged to an individual who was 45 years and old to an individual who was 65 years.

Columns 6–8 in Table 3 show the relative standard deviation of expected inflation for each type in the stochastic life-cycle model. We see that the stochastic component makes the poorer types worse off in terms of expected inflation volatility. The relative inflation volatility of the young-poor compared to the young-rich, for example, goes from 2.41 in the deterministic life-cycle model to 2.5 in the stochastic model. The broad picture, however, still remains: even when controlling for deterministic and stochastic changes in expenditures, the expected inflation volatility is around 2.5 times larger for the poorest ten percent compared to the richest ten percent.

These findings show that, although households' real expenditures change over time, the average change over a time period of 26 years from ageing and chance is not enough to alter the conclusions drawn in regards to inflation inequality between households whose real expenditures stay fixed.[12] Instead, this suggests that initial expenditure levels are an important determinant for households' expected inflation rates.

6 What Drives the Inflation Differences?

The generalized superlative price index (16) lends itself to a simple decomposition of inflation into its components. Taking first differences of the log of the price index allows us to compute the contribution of inflation coming from each expenditure category j. Figure 5 plots this decomposition for the rich and poor's inflation. The figure shows the primary expenditure

[12] The results do not reject that other factors such as education could induce changes in expenditures that are strong enough to make a difference.

categories that have driven inflation since 1996.

Panel (a) in Figure 5 shows that the two primary expenditure categories that drive inflation for the poor are "food at home" and "gas and utilities" (energy). In contrast, panel (b) shows that these categories play a minor role in driving inflation for the rich. Moreover, the key point in regards to what drives inflation for the rich is that no expenditure category is a major driver. In panel (b) we plotted the contributions to the rich's inflation from "food at home" and "gas and utilities" to illustrate that these play a minor role. Additionally, we also show how the most important driver for the rich's inflation, "owned dwellings", and the fourth most important driver, "other vehicle expenditures", contribute. While "owned dwelling" indeed drives a considerable amount of total inflation of the rich, it is still minor and "other vehicle expenditures" is almost invisible in some periods. Figure B.6 in Appendix B shows the inflation contribution from all expenditure categories to the rich and poor's inflation, respectively.

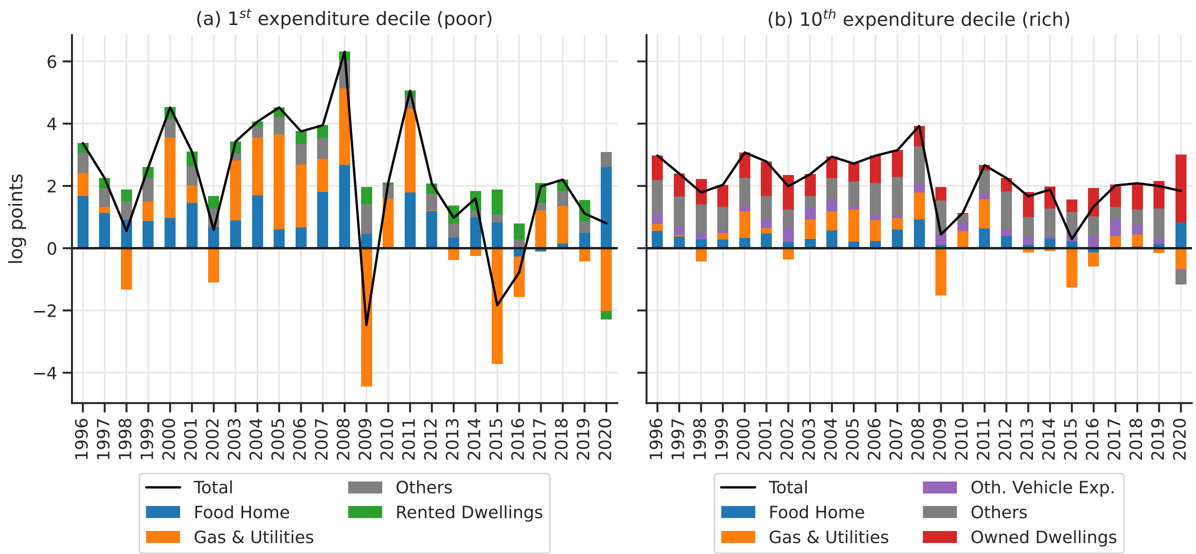

FIGURE 5. Inflation decomposition by expenditure categories.

6.1 Decomposition Into Reference Basket and Product Substitution

In order to shed more light on the inflation heterogeneity between poor and rich households, we decompose the price index into one component that reflects pure price changes of a base-period reference basket and another component that reflects product substitution away from this basket as prices change. The former is simply the Laspeyres price index $P_{Lt} = \sum_j w_{js}(p_{jt}/p_{js})$, which can be written on a weighted geometric-mean form as[13]

$$P_{Lt} = \prod_{j \in J} \left(\frac{p_{jt}}{p_{js}} \right)^{\delta_{jt}^L}, \qquad \text{where} \qquad \delta_{jt}^L \equiv \frac{w_{js} \, L\left(\frac{p_{jt}}{p_{js}}, \, P_{Lt}\right)}{\sum_{i \in J} w_{is} \, L\left(\frac{p_{jt}}{p_{js}}, \, P_{Lt}\right)}. \tag{23}$$

The latter component can then be backed out as the residual between the computed price index and the corresponding Laspeyres index. Taking logs of the generalized superlative index (16) and

[13] See Appendix A.2 or Balk (2004).

adding and subtracting the log of the Laspeyres index (23), we obtain the decomposition

$$\ln P_t = \underbrace{\sum_{j\in J} \delta^L_{jt} \ln\left(\frac{p_{jt}}{p_{js}}\right)}_{\text{Laspeyres price index}} + \underbrace{\sum_{j\in J} \left(\chi_{jt} - \delta^L_{jt}\right) \ln\left(\frac{p_{jt}}{p_{js}}\right)}_{\text{Product substitution}}. \tag{24}$$

Figure 6 shows the decomposition of the G-SV price index for the first and the tenth expenditure decile. It clearly highlights that the biggest share of the increase in the cost of living is driven by the price changes of the reference baskets and substitution effects only marginally reduce the G-SV price index. Interestingly, however, the substitution effect among the households in the first expenditure decile is considerably bigger than the substitution among the households in the tenth expenditure decile.

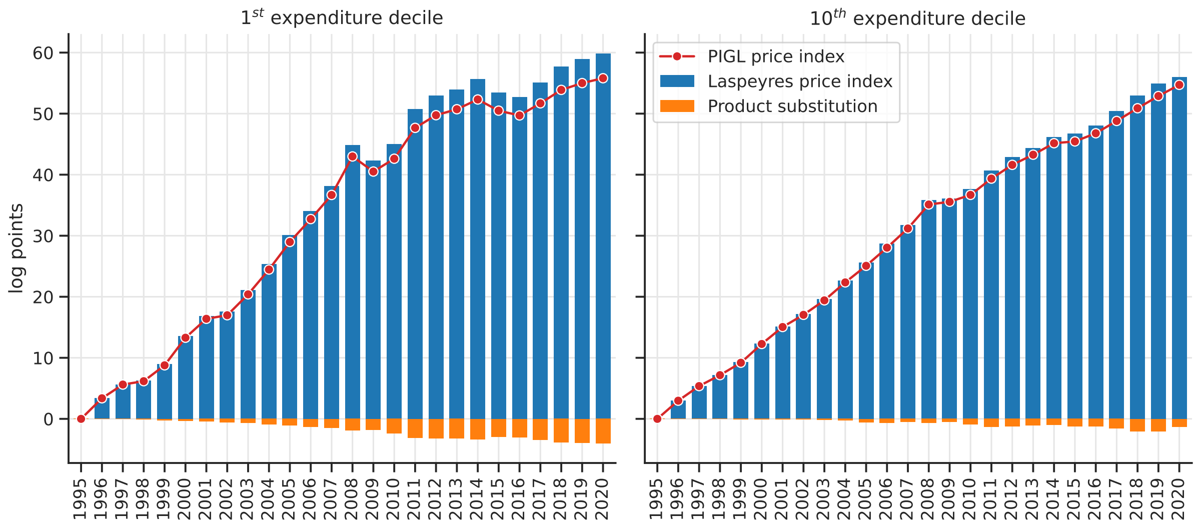

FIGURE 6. Decomposition of the G-SV price index into the Laspeyres price index and product substitution.

Notes. Decomposition is performed as shown in Equation (24). The G-SV price index for each expenditure decile denotes the price index of the representative agent within each respective decile.

Focusing on the difference between the two deciles in Figure 7 shows that the composition of their reference baskets plays an important role but that the difference in product substitution becomes relatively more important.

This section shows that in order to uncover the fundamental drivers of inflation heterogeneity across households it is primarily necessary to understand inflation heterogeneity across products. Therefore, we deem it important for future research to investigate this further. In particular, whether it is a coincidence that the highly volatile product groups are purchased relatively more by poor households, or if there is a deeper and more structural link. In terms of policy, the results show that stabilizing prices of selected product groups, and not just an aggregate price index, are important when inequality is taken into account. Specifically, the results show that the highly volatile inflation of "food at home" and "gas and utilities" affects poorer households disproportionately more.

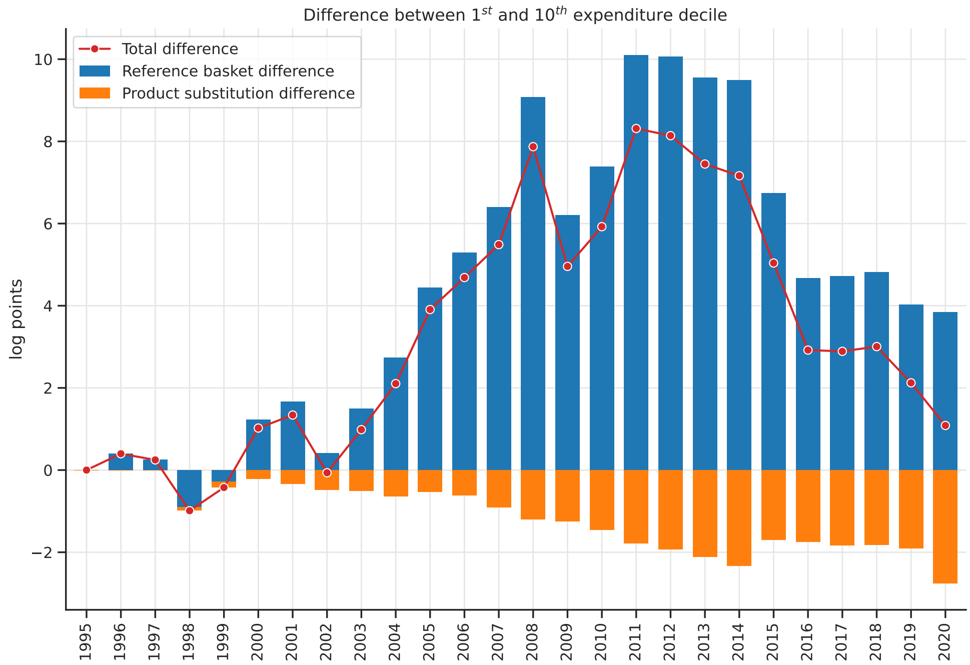

FIGURE 7. Difference of the decomposed G-SV price index between the first and tenth expenditure decile.

Notes. The figure shows the decomposition of the first expenditure decile minus the corresponding numbers for the tenth expenditure decile in Figure 6.

7 Comparison with Full Demand System Estimation

To assess the validity of the weak separability assumption and the robustness of our results, we perform a full demand system estimation. Full demand system estimation can be carried out via GMM by using the expenditure share equations (8) and the parameterization for $B(\boldsymbol{p})$ and $D(\boldsymbol{p})$. The parameterization of the Sato-Vartia specification used in our empirical analysis is given by

$$B(\boldsymbol{p}) = \left(\sum_{j \in J} \omega_j p_j^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \qquad \text{and} \qquad D(\boldsymbol{p}) = \left(\sum_{j \in J} \theta_j p_j^{1-\varphi} \right)^{\frac{1}{1-\varphi}}, \tag{25}$$

where $\sigma, \varphi > 0$ and the taste parameters satisfy $\sum_{j \in J} \omega_j = \sum_{j \in J} \theta_j = 1$ with $\omega_j, \theta_j \geq 0$ for all $j \in J$.[14]

Since the GMM estimation under the full demand system specification might exhibit several local minima, it generally must not hold that a particular local solution of the GMM estimation is indeed also the global minimum. Therefore, the following results should be interpreted with this in mind. However, by choosing an appropriate guess for the preference parameters in the full demand system that corresponds to the preferences under weak separability, we obtain a local solution for the full demand system which is in the vicinity of the weakly separable case.

[14] This parameterization is also considered in Alder, Boppart and Müller (2022).

In particular, we set the values for ε, γ and ν to the values in Table 1 and equally distribute ω_j and θ_j among the goods classified as B and D, respectively.[15]

Table 4 and Figure 8 show the estimates for the entire set of preference parameters in the full demand system.[16] We can note that ε is practically unchanged, but γ substantially higher. ν is not important for determining the price index, but still lies fairly close to the estimates under weak separability. Turning to σ and φ, we can note that they both lie below one, which indicates that goods within the B and D basket are complements. Figure 8 shows the point estimates of the complete set of taste parameters ω_j and θ_j for the B and the D basket respectively. Each category has a blue and an orange bar for ω_j and θ_j, but with the minor exception of "tobacco", every single good only has one strictly positive taste parameter. Thus, the estimated parameters from the full demand system and the implied classification into two baskets is almost perfectly in line with the classification applied in the weakly separable case. In order to gauge the measure of fit, we computed the root mean squared error (RMSE) of the expenditure share on the D good. Both measures are close to each other and the full demand system actually has a slightly higher RMSE. Although it is more flexible in matching the expenditure shares of all individual expenditure categories, it not necessarily needs to perform better in matching the expenditure share on the D good.

Table 4. Estimates of the preference parameters in the full demand system compared to the estimates under weak separability.

	Full demand system	Weak separability
ε	0.685	0.677
	(0.004)	(0.004)
γ	0.505	0.211
	(0.018)	(0.023)
ν	346.736	327.271
	(11.793)	(13.358)
σ	0.050	
	(0.012)	
φ	0.360	
	(0.006)	
N	74,372	74,372
RMSE	0.1494	0.1487

Notes. Standard errors in parentheses. "RMSE" refers to the root mean squared error of the expenditure share on the D good: $\sqrt{\sum_h (w_{Dh} - \widehat{w}_{Dh})^2 / N}$. Each household-level observation is weighted by the CEX sampling weight.

[15] In principle we could perform a grid search over the whole set of feasible parameters. However, depending on the initial guess, convergence of the GMM estimator is slow and can sometimes take more than 24 hours. Thus, we deemed a reliable grid search as infeasible.

[16] While point estimates are still consistent when some parameters lie on the boundary of the parameter space, standard errors obtained from the standard covariance matrix are not (see Andrews, 1999, 2002). This should be kept in mind with the standard errors presented here.

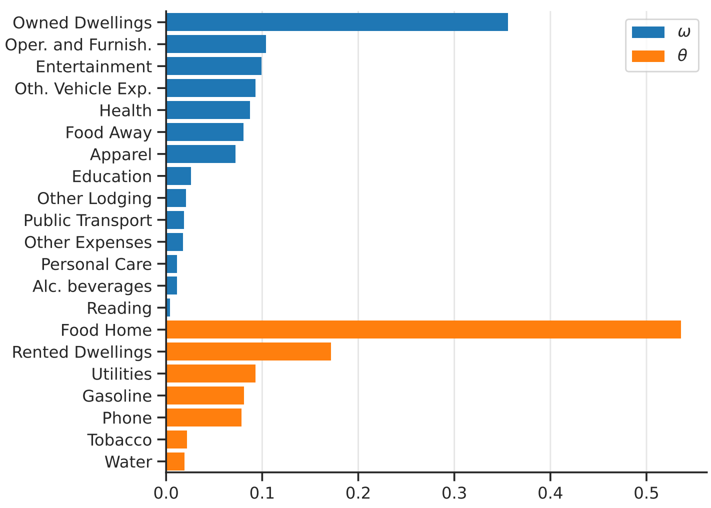

FIGURE 8. Point estimates for the taste parameters ω_j and θ_j of the full demand system.

Notes. The figure shows the estimated taste parameters from the full demand system estimation. Results are for the closest local minimum to the weakly separable case which has been used as initial guess for the parameters.

Figure 9 shows a comparison of the resulting price index from the full demand system (left panel) with the respective price index under weak separability (right panel). The figure shows that the price indices for the individual expenditure deciles are less spread out under the full demand system, but the general difference is relatively minor.

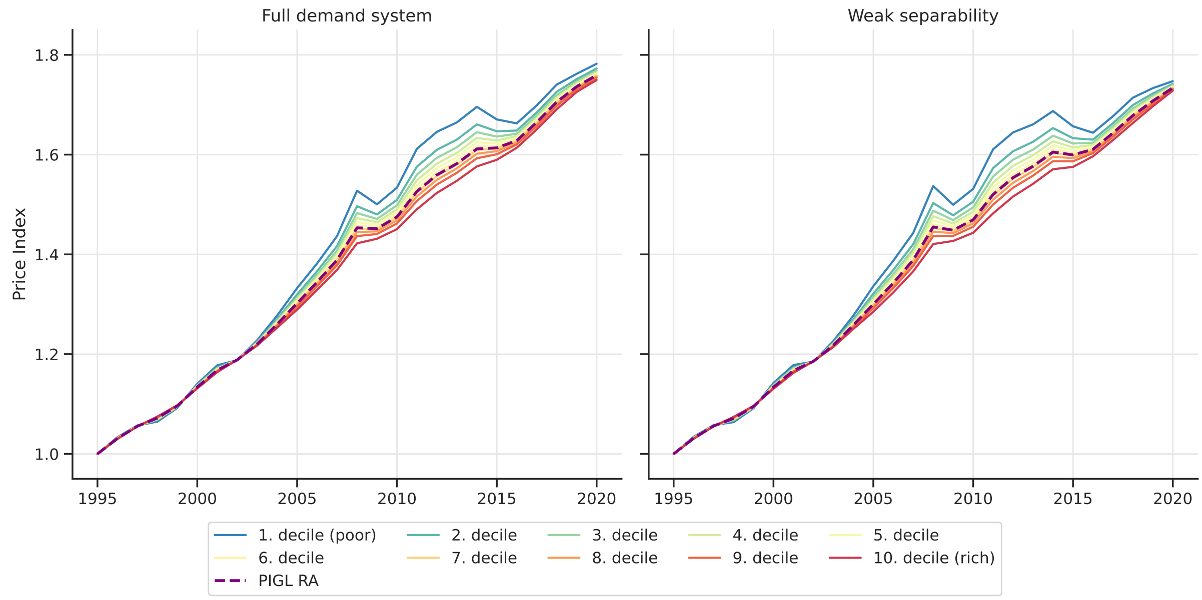

FIGURE 9. Comparison of the G-SV price index for the full demand system and under weak separability.

The differences between the full demand system estimation and estimation under weak separability are even less pronounced when looking at the comparison of inflation rates in Figure 10.

Taking these comparisons at face value indicates that the full demand system estimation does

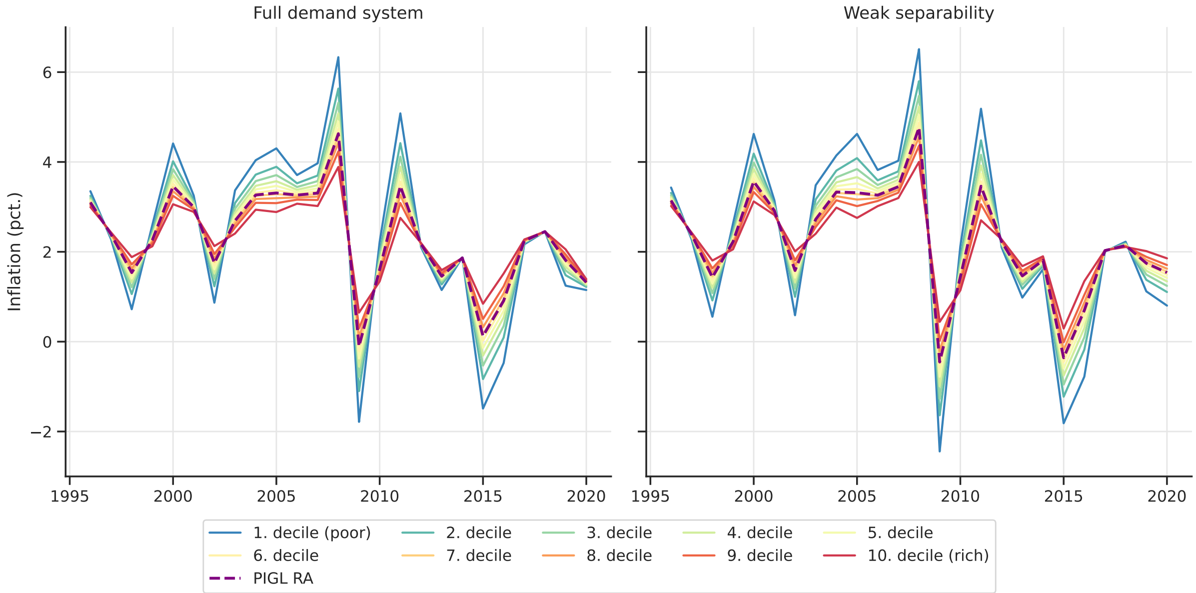

FIGURE 10. Comparison of the G-SV inflation for the full demand system and under weak separability.

not strongly reject the case of weakly separable preferences. Further, the price index along with inflation from the full demand system and from the weakly separable preferences are remarkably close and do not change any of our main conclusions.

8 Conclusion

We derive a nonhomothetic cost-of-living index which allows us to describe inflation heterogeneity along the full expenditure distribution. The cost-of-living index is microfounded with PIGL preferences and we show that it can be computed without a full demand system estimation if weak separability of consumption goods into necessary and luxury goods is imposed. The theoretical rigor and practical simplicity of our index makes it especially appealing compared to other approaches previously taken in the literature on inflation heterogeneity.

The price index generalizes several classes of homothetic price indices, some of them particularly interesting for their superlative property. We present results for generalized Sato-Vartia, Törnqvist, geometric Walsh, Theil, Fisher and arithmetic Walsh price indices. Homothetic indices such as the Sato-Vartia and Törnqvist are used in previous empirical studies to approximate inflation heterogeneity by computing separate price indices for different income groups. We show that this approach can be rationalized through the lens of our framework but the usual caveats of the group-specific approach still apply.

Our empirical results show that from 1996 to 2020 there was a substantial heterogeneity in inflation between poorer and richer households in the US. The particular striking result we find is that while mean inflation is around 2.25 percent for everyone, the standard deviation of inflation has been 2.14 for the poor compared to 0.85 for the rich. Thus, inflation volatility is 2.5 times higher for the poor.

We find that poorer households are much more exposed to the highly volatile inflation rates of food, gas and utilities compared to the rich. We furthermore show that substitution behavior is only of second-order importance. Our findings hence suggest that in order to uncover the fundamental drivers of inflation inequality it is first and foremost important to understand *why* households make the consumption choices they do and what the explanation for inflation heterogeneity across product groups is.

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 pages 1, 4, 6, 13, 23, 35, 37).

Andrews, D. W. K. (1999). "Estimation when a parameter is on a boundary." *Econometrica*, vol. 67(6), pp. 1341–1383 (cited on page 24).

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 page 24).

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 pages 3, 9, 15).

Balk, B. M. (2004). "Decompositions of Fisher indexes." *Economics Letters*, vol. 82(1), pp. 107–113 (cited on pages 9, 21, 32).

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 pages 2, 7, 13, 35).

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 pages 2, 7, 8).

Blundell, R., Pashardes, P. and Weber, G. (1993). "What do we learn about consumer patterns from micro data?" *American Economic Review*, vol. 83(3), pp. 570–597 (cited on page 13).

Blundell, R., Pistaferri, L. and Preston, I. (2008). "Consumption inequality and partial insurance." *American Economic Review*, vol. 98(5), pp. 1887–1921 (cited on page 19).

Blundell, R. and Robin, J.-M. (2000). "Latent separability: Grouping goods without weak separability." *Econometrica*, vol. 68(1), pp. 53–84 (cited on page 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 pages 1, 4, 13).

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*. Washington, DC: Advisory Commission to Study the Consumer Price Index. (Cited on page 15).

Broda, C. and Romalis, J. (2009). *The welfare implications of rising price dispersion* [Unpublished manuscript], Booth School of Business, University of Chicago. (Cited on pages 3, 9).

Carlson, B. C. (1972). "The logarithmic mean." *American Mathematical Monthly*, vol. 79(6), pp. 615–618 (cited on page 6).

Christensen, L. R., Jorgenson, D. W. and Lau, L. J. (1975). "Transcendental logarithmic utility functions." *American Economic Review*, vol. 65(3), pp. 367–383 (cited on page 9).

Cravino, J., Levchenko, A. A. and Rojas, M. (forthcoming). "Population aging and structural transformation." *American Economic Journal: Macroeconomics* (cited on pages 1, 34).

Deaton, A. and Muellbauer, J. (1980). "An almost ideal demand system." *American Economic Review*, vol. 70(3), pp. 312–326 (cited on page 2).

Diewert, W. E. (1976). "Exact and superlative index numbers." *Journal of Econometrics*, vol. 4(2), pp. 115–145 (cited on pages 2, 7–9).

Diewert, W. E. (1978). "Superlative index numbers and consistency in aggregation." *Econometrica*, vol. 46(4), pp. 883–900 (cited on page 2).

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. 33–71). Elsevier Science Publishers B.V. (Cited on page 2).

Engel, E. (1857). "Die vorherrschenden Gewerbszweige in den Gerichtsämtern mit Beziehung auf die Productions- und Consumtionsverhältnisse des Königreichs Sachsen." *Zeitschrift des Statistischen Büreaus des Königlich Sächsischen Ministeriums des Innern*, vol. 8–9, pp. 153–182 (cited on page 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 page 2).
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 page 2).
Fernández-Villaverde, J. and Krueger, D. (2007). "Consumption over the life cycle: Facts from Consumer Expenditure Survey data." *Review of Economics and Statistics*, vol. 89(3), pp. 552–565 (cited on page 18).
Fisher, I. (1922). *The making of index numbers: A study of their varieties, tests, and reliability*. Houghton Mifflin Company. (Cited on pages 1, 8).
Fox, L. E. and Burns, K. (2021). *The supplemental poverty measure: 2020* (Current Population Reports P60–275). Washington, DC: U.S. Census Bureau. (Cited on page 11).
Hanoch, G. (1975). "Production and demand models with direct or indirect implicit additivity." *Econometrica*, vol. 43(3), pp. 395–419 (cited on page 3).
Hill, R. J. (2006). "Superlative index numbers: Not all of them are super." *Journal of Econometrics*, vol. 130(1), pp. 25–43 (cited on page 13).
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 pages 3, 11).
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 pages 3, 9, 15).
Jaravel, X. (2021). "Inflation inequality: Measurement, causes, and policy implications." *Annual Review of Economics*, vol. 13, pp. 599–629 (cited on page 3).
Kaplan, G. and Schulhofer-Wohl, S. (2017). "Inflation at the household level." *Journal of Monetary Economics*, vol. 91, pp. 19–38 (cited on page 3).
Konüs, A. A. (1939). "The problem of the true index of the cost of living." *Econometrica*, vol. 7(1), pp. 10–29 (cited on pages 2, 3).
Lauper, C. and Mangiante, G. (2021). *Monetary policy shocks and inflation heterogeneity* [Unpublished manuscript], University of Lausanne. (Cited on page 3).
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 page 3).
Muellbauer, J. (1974). "Prices and inequality: The United Kingdom experience." *Economic Journal*, vol. 84(333), pp. 32–55 (cited on page 3).
Muellbauer, J. (1975). "Aggregation, income distribution and consumer demand." *Review of Economic Studies*, vol. 42(4), pp. 525–543 (cited on pages 1, 4, 6, 37).
Muellbauer, J. (1976). "Community preferences and the representative consumer." *Econometrica*, vol. 44(5), pp. 979–999 (cited on pages 1, 4, 6, 37).
Orchard, J. (2020). *Household inflation and aggregate inflation* [Unpublished manuscript], University of California, San Diego. (Cited on page 3).
Orchard, J. (2021). *Cyclical demand shifts and cost of living inequality* [Unpublished manuscript], University of California, San Diego. (Cited on pages 1, 3, 12).
Oulton, N. (2012). "How to measure living standards and productivity." *Review of Income and Wealth*, vol. 58(3), pp. 424–456 (cited on page 2).
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 pages 2, 3, 7, 34).
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 pages 2, 4).
Sato, K. (1976). "The ideal log-change index number." *Review of Economics and Statistics*, vol. 58(2), pp. 223–228 (cited on pages 1, 8, 14).
Sato, R. (1975). "The most general class of CES functions." *Econometrica*, vol. 43(5/6), pp. 999–1003 (cited on page 3).
Theil, H. (1973). "A new index number formula." *Review of Economics and Statistics*, vol. 55(4), pp. 498–502 (cited on page 8).

Törnqvist, L. (1936). "The Bank of Finland's consumption price index." *Bank of Finland Monthly Bulletin*, vol. 16(10), pp. 27–34 (cited on pages 1, 8).
Vartia, Y. O. (1976). "Ideal log-change index numbers." *Scandinavian Journal of Statistics*, vol. 3(3), pp. 121–126 (cited on pages 1, 8).
Wachter, J. A. and Yogo, M. (2010). "Why do household portfolio shares rise with wealth?" *Review of Financial Studies*, vol. 23(11), pp. 3929–3965 (cited on pages 1, 12).
Walsh, C. M. (1901). *The measurement of general exchange-value*. The Macmillan Company. (Cited on page 8).

Appendix A Proofs and Extensions

A.1 Proof of Proposition 1

Proof. Inverting the indirect utility function (4) gives the expenditure function

$$c(u,\boldsymbol{p})=\left[1+\varepsilon\left(u+\frac{\nu}{\gamma}\left\{\left(\frac{D(\boldsymbol{p})}{B(\boldsymbol{p})}\right)^{\gamma}-1\right\}\right)\right]^{\frac{1}{\varepsilon}}B(\boldsymbol{p}).$$

Suppose that the reference utility u corresponds to the expenditure level in some base period s such that $c(u,\boldsymbol{p}_s)=e_s$ and $u\equiv V(e_s,\boldsymbol{p}_s)$. Using the indirect utility function (4) to substitute this into some period-t expenditure function and rearranging terms yields

$$\begin{aligned}
c(u,\boldsymbol{p}_t)&=e_s\left[1+\frac{\varepsilon\,\nu}{\gamma}\left(\frac{B(\boldsymbol{p}_s)}{e_s}\right)^{\varepsilon}\left(\frac{D(\boldsymbol{p}_s)}{B(\boldsymbol{p}_s)}\right)^{\gamma}\left\{\left(\frac{D(\boldsymbol{p}_t)}{D(\boldsymbol{p}_s)}\right)^{\gamma}\left(\frac{B(\boldsymbol{p}_t)}{B(\boldsymbol{p}_s)}\right)^{-\gamma}-1\right\}\right]^{\frac{1}{\varepsilon}}\frac{B(\boldsymbol{p}_t)}{B(\boldsymbol{p}_s)}\\
&=e_s\left[1+\frac{\varepsilon\,w_{Ds}}{\gamma}\left\{\left(\frac{P_{Dt}}{P_{Bt}}\right)^{\gamma}-1\right\}\right]^{\frac{1}{\varepsilon}}P_{Bt}\\
&=e_s\left[\left(1-\frac{\varepsilon\,w_{Ds}}{\gamma}\right)P_{Bt}^{\gamma}+\frac{\varepsilon\,w_{Ds}}{\gamma}P_{Dt}^{\gamma}\right]^{\frac{1}{\gamma}\cdot\frac{\gamma}{\varepsilon}}P_{Bt}^{1-\frac{\gamma}{\varepsilon}},
\end{aligned}$$

where the second equality uses $P_{Bt}\equiv B(\boldsymbol{p}_t)/B(\boldsymbol{p}_s)$, $P_{Dt}\equiv D(\boldsymbol{p}_t)/D(\boldsymbol{p}_s)$ and the expenditure share (5). By the Konüs definition (1), the cost-of-living index is then

$$P_t^{PIGL}=\left[\left(1-\frac{\varepsilon\,w_{Ds}}{\gamma}\right)P_{Bt}^{\gamma}+\frac{\varepsilon\,w_{Ds}}{\gamma}P_{Dt}^{\gamma}\right]^{\frac{1}{\gamma}\cdot\frac{\gamma}{\varepsilon}}P_{Bt}^{1-\frac{\gamma}{\varepsilon}}=\widetilde{P}_t^{\frac{\gamma}{\varepsilon}}P_{Bt}^{1-\frac{\gamma}{\varepsilon}}. \tag{A.1}$$

Since $\widetilde{P}_t$ has a CES form, we may define hypothetical budget shares corresponding to this price function by

$$\psi_{Bt}\equiv\left(1-\frac{\varepsilon w_{Ds}}{\gamma}\right)\left(\frac{P_{Bt}}{\widetilde{P}_t}\right)^{\gamma}\qquad\text{and}\qquad\psi_{Dt}\equiv\frac{\varepsilon w_{Ds}}{\gamma}\left(\frac{P_{Dt}}{\widetilde{P}_t}\right)^{\gamma}, \tag{A.2}$$

with $\psi_{Bs}=1-\varepsilon w_{Ds}/\gamma$ and $\psi_{Ds}=\varepsilon w_{Ds}/\gamma$. These shares ensure that the price index remains at the same utility level; $w_{Ds}(P_{Dt}\,/\,\widetilde{P}_t)^{\gamma}$ is the expenditure share of the D basket at period-t prices that prevails at the same utility level as w_{Ds}. To see this, use the expenditure share (5) to get

$$\begin{aligned}
w_{Dt}&=\nu\left(\frac{B(\boldsymbol{p}_t)}{e_t}\right)^{\varepsilon}\left(\frac{D(\boldsymbol{p}_t)}{B(\boldsymbol{p}_t)}\right)^{\gamma} && \text{(by (5))}\\
&=\nu\left(\frac{B(\boldsymbol{p}_s)}{e_s}\right)^{\varepsilon}\left(\frac{D(\boldsymbol{p}_s)}{B(\boldsymbol{p}_s)}\right)^{\gamma}\left(\frac{e_sP_{Dt}^{\frac{\gamma}{\varepsilon}}P_{Bt}^{1-\frac{\gamma}{\varepsilon}}}{e_t}\right)^{\varepsilon} && \text{(by (3))}
\end{aligned}$$

$$= w_{Ds} \left(\frac{P_{Dt}^{\frac{\gamma}{\varepsilon}} P_{Bt}^{1-\frac{\gamma}{\varepsilon}}}{P_t Q_t} \right)^{\varepsilon} \qquad \text{(by (5) and } e_t/e_s = P_t Q_t)$$

$$= w_{Ds} \left(\frac{P_{Dt}}{\widetilde{P}_t} \right)^{\gamma} Q_t^{-\varepsilon} \qquad \text{(by (A.1))}.$$

The third equality uses the decomposition $e_t/e_s = P_t Q_t$, where P_t is the Konüs price index and Q_t the corresponding quantity index. Along the same indifference curve as w_{Ds}, we necessarily have $Q_t = 1$ for all t, and the result then immediately follows. Equation (A.2) allows us to write $\widetilde{P}_t$ as a Sato-Vartia index. The procedure is the same as in the standard case: solve for $\widetilde{P}_t$ from the shares in Equation (A.2), take logs, multiply by the difference in shares over time, sum over both B and D and solve for $\widetilde{P}_t$. For $C \in \{B, D\}$, the first two steps yields

$$\ln \widetilde{P}_t = \ln P_{Ct} - \frac{1}{\gamma} \ln \left(\frac{\psi_{Ct}}{\psi_{Cs}} \right) \qquad \Longleftrightarrow \qquad -\frac{1}{\gamma} = \frac{\ln \widetilde{P}_t - \ln P_{Ct}}{\ln \psi_{Ct} - \ln \psi_{Cs}}.$$

Multiplying both sides by $\psi_{Ct} - \psi_{Cs}$, summing over both $C \in \{B, D\}$, and rearranging terms results in

$$\ln \widetilde{P}_t \sum_{C \in \{B,D\}} \frac{\psi_{Ct} - \psi_{Cs}}{\ln \psi_{Ct} - \ln \psi_{Cs}} = \sum_{C \in \{B,D\}} \frac{\psi_{Ct} - \psi_{Cs}}{\ln \psi_{Ct} - \ln \psi_{Cs}} \ln P_{Ct}.$$

Then solving for $\widetilde{P}_t$ yields

$$\widetilde{P}_t = P_{Dt}^{\phi_t} P_{Bt}^{1-\phi_t}, \qquad \text{where} \qquad \phi_t \equiv \frac{L(\psi_{Dt}, \psi_{Ds})}{L(\psi_{Dt}, \psi_{Ds}) \; + \; L(\psi_{Bt}, \psi_{Bs})}. \tag{A.3}$$

Plugging Equation (A.3) into Equation (A.1) gives the household-level price indices.

Because a representative level of expenditures e^{RA} exists over any group of households, group-level behavior is characterized by the same indirect utility function and expenditure function as household-level behavior.[17] Aggregate-level cost-of-living indices are therefore derived identically to above, with the only difference that group-level expenditure shares $\overline{w}_{Ds}$ and representative levels of expenditure e^{RA} are used instead of household-level ones. □

A.2 Proof of Corollary 1

Proof. The result is immediate by setting both P_{Dt} and P_{Dt} to either the *quadratic-mean-of-order-r* index (14) or the *Theil-Sato* index (15) and substituting these into the general PIGL index (11). We only need to rewrite (14) into a geometric-mean form. Balk (2004) does this for the Fisher ideal index ($r = 2$), and the generalization to any $r > 0$ is analogous. As in Corollary 1, define

$$P_{Lt} = \left[\sum_{j \in J} w_{js} \left(\frac{p_{jt}}{p_{js}} \right)^{\frac{r}{2}} \right]^{\frac{2}{r}} \qquad \text{and} \qquad P_{Pt} = \left[\sum_{j \in J} w_{jt} \left(\frac{p_{jt}}{p_{js}} \right)^{-\frac{r}{2}} \right]^{-\frac{2}{r}},$$

[17] This can be shown by substituting for $e^{RA} = \bar{e}\kappa^{-\frac{1}{\varepsilon}}$ in the aggregate expenditure share (9) and integrating to obtain the group-level indirect utility function.

such that the *quadratic-mean-of-order-r* index (14) can be written $P_t = \sqrt{P_{Lt} P_{Pt}}$. P_{Lt} weighs price changes by base-period expenditure shares while P_{Pt} uses current-period expenditure shares, and the definition nests the Laspeyres and Paasche indices as the special case where $r = 2$, thus motivating the L and P notation. By the definition of P_{Lt} and the logarithmic mean, it holds that

$$\begin{aligned} 0 &= \sum_{j \in J} w_{js} \left(\frac{p_{jt}}{p_{js}} \right)^{\frac{r}{2}} - P_{Lt}^{\frac{r}{2}} = \sum_{j \in J} w_{js} \left[\left(\frac{p_{jt}}{p_{js}} \right)^{\frac{r}{2}} - P_{Lt}^{\frac{r}{2}} \right] \\ &= \sum_{j \in J} w_{js} \, L \left((\tfrac{p_{jt}}{p_{js}})^{\frac{r}{2}}, \, P_{Lt}^{\frac{r}{2}} \right) \, \ln \left(\frac{p_{jt}/p_{js}}{P_{Lt}} \right)^{\frac{r}{2}} \\ &= \frac{r}{2} \sum_{j \in J} \widetilde{w}_{Ljt} \left[\ln \left(\frac{p_{jt}}{p_{js}} \right) - \ln P_{Lt} \right], \end{aligned}$$

with $\widetilde{w}_{Ljt}$ defined as in Corollary 1. Solving for $\ln P_{Lt}$, we get

$$\ln P_{Lt} = \sum_{j \in J} \frac{\widetilde{w}_{Ljt}}{\sum_{i \in J} \widetilde{w}_{Lit}} \ln \left(\frac{p_{jt}}{p_{js}} \right).$$

Identical steps for P_{Pt} yields

$$\ln P_{Pt} = \sum_{j \in J} \frac{\widetilde{w}_{Pjt}}{\sum_{i \in J} \widetilde{w}_{Pit}} \ln \left(\frac{p_{jt}}{p_{js}} \right),$$

with $\widetilde{w}_{Pjt}$ defined as in Corollary 1. Substituting these into the overall index P_t yields

$$P_t = \prod_{j \in J} \left(\frac{p_{jt}}{p_{js}} \right)^{\delta_{jt}}, \qquad \text{where} \qquad \delta_{jt} \equiv \frac{1}{2} \left[\frac{\widetilde{w}_{Ljt}}{\sum_i \widetilde{w}_{Lit}} + \frac{\widetilde{w}_{Pjt}}{\sum_i \widetilde{w}_{Pit}} \right],$$

and we are done. □

A.3 Proof of Proposition 3

Proof. If preferences are of the Cobb-Douglas form $V(e, \boldsymbol{p}) = \left[\frac{e}{B(\boldsymbol{p})^{1-\nu} D(\boldsymbol{p})^{\nu}} \right]$, then the cost-of-living index is $P_t = P_{Bt}^{1-\nu} P_{Dt}^{\nu}$ by Proposition 1. If P_{Bt} and P_{Dt} are Törnqvist indices, Corollary 1 and Assumption 1 allow us to write this index as

$$P_t = \prod_{j \in J_B} \left(\frac{p_{jt}}{p_{js}} \right)^{(1-\nu)\delta_{jt}^B} \prod_{j \in J_D} \left(\frac{p_{jt}}{p_{js}} \right)^{\nu \delta_{jt}^D}, \qquad \delta_{jt}^C = \frac{w_{js}^C + w_{jt}^C}{2}, \quad C \in \{B, D\}, \tag{A.4}$$

where J_B and J_D denote the sets of goods in B and D, respectively. (Under the weak separability assumption, it holds that $J_B \cup J_D = J$ and $J_B \cap J_D = \varnothing$.) Meanwhile, the standard Törnqvist index reads

$$P_t = \prod_{j \in J} \left(\frac{p_{jt}}{p_{js}} \right)^{\delta_{jt}}, \qquad \delta_{jt} = \frac{w_{js} + w_{jt}}{2}. \tag{A.5}$$

Equations (A.4) and (A.5) are equal if $(1-\nu)\delta_{jt}^B = \delta_{jt}$ for $j \in J_B$ and $\nu\delta_{jt}^D = \delta_{jt}$ for $j \in J_D$. Under Cobb-Douglas preferences, $\nu = w_D$ is the homothetic and time-invariant expenditure share on D. Under the weak separability assumption, the total expenditure share on good $j \in J_C$, $C \in \{B, D\}$, is given by $w_j = w_C \, w_j^C$. Thus,

$$\begin{aligned}
(1-\nu)\delta_{jt}^B &= w_B \frac{w_{js}^B + w_{jt}^B}{2} = \frac{w_{js} + w_{jt}}{2} = \delta_{jt}, \qquad j \in J_B, \\
\nu\delta_{jt}^D &= w_D \frac{w_{js}^D + w_{jt}^D}{2} = \frac{w_{js} + w_{jt}}{2} = \delta_{jt}, \qquad j \in J_D,
\end{aligned}$$

and it follows that the Törnqvist index under weak separability (A.4) is the same as the standard Törnqvist index (A.5). □

A.4 Allowing for Heterogeneity in Tastes

Redding and Weinstein (2020) stress the importance of accounting for heterogeneity in tastes for the cost of living and it is possible to extend the baseline framework to allow for this. Following Cravino, Levchenko and Rojas (forthcoming), let the preferences of household h be characterized by an indirect utility function of the form

$$V_h(e_h, \boldsymbol{p}) = \frac{1}{\varepsilon} \left[\left(\frac{e_h}{B(\boldsymbol{p})} \right)^{\varepsilon} - 1 \right] - \frac{\nu_h}{\gamma} \left[\left(\frac{D(\boldsymbol{p})}{B(\boldsymbol{p})} \right)^{\gamma} - 1 \right], \tag{A.6}$$

where the only difference to the PIGL specification in Equation (4) is that we allow the taste parameter ν_h to vary across households (and time). As before, the expenditure share of the latent good with price function $D(\cdot)$ is given by Roy's identity as

$$w_{Dh} = \nu_h \left(\frac{B(\boldsymbol{p})}{e_h} \right)^{\varepsilon} \left(\frac{D(\boldsymbol{p})}{B(\boldsymbol{p})} \right)^{\gamma},$$

and the corresponding aggregate expenditure share over any N households is now

$$\overline{w}_D = \left(\frac{B(\boldsymbol{p})}{\bar{e}} \right)^{\varepsilon} \left(\frac{D(\boldsymbol{p})}{B(\boldsymbol{p})} \right)^{\gamma} \kappa, \qquad \text{where} \qquad \kappa \equiv \frac{1}{N} \sum_{h=1}^{N} \nu_h \left(\frac{e_h}{\bar{e}} \right)^{1-\varepsilon}.$$

A representative expenditure level $e^{RA} = \bar{e}\kappa^{-\frac{1}{\varepsilon}}$ therefore exists and incorporates any heterogeneity in tastes. Substituting back into the aggregate expenditure share $\overline{w}_D$ and integrating back yields aggregate-level behavior characterized by the indirect utility function

$$V(e^{RA}, \boldsymbol{p}) = \frac{1}{\varepsilon} \left[\left(\frac{e^{RA}}{B(\boldsymbol{p})} \right)^{\varepsilon} - 1 \right] - \frac{1}{\gamma} \left[\left(\frac{D(\boldsymbol{p})}{B(\boldsymbol{p})} \right)^{\gamma} - 1 \right],$$

with corresponding expenditure function

$$c(u^{RA}, \boldsymbol{p}) = \left[1 + \varepsilon\left(u^{RA} + \frac{1}{\gamma}\left\{\left(\frac{D(\boldsymbol{p})}{B(\boldsymbol{p})}\right)^{\gamma} - 1\right\}\right)\right]^{\frac{1}{\varepsilon}} B(\boldsymbol{p}).$$

This expenditure function is independent of the taste parameters ν_h. We can therefore follow the same steps as in Appendix A.1 to derive an identical price index as in Proposition 1. Again, this index is only a function of the base-period expenditure share for the D basket, price indices P_{Dt} and P_{Bt}, and the parameters ε and γ. Heterogeneity in the taste parameters ν_h only affect the price index indirectly to the extent that they affect expenditure shares. Whenever expenditure shares are observed in the data, we therefore do not need to know these individual tastes to compute the price index. We obtain the same result for household-level cost-of-living indices as the special case where $N = 1$.

Taste heterogeneity also poses no challenge with respect to estimating ε and γ. Since PIGL preferences aggregate consistently, it is possible to estimate these parameters from a aggregate data without any aggregation bias. Therefore, taking an aggregate time series and estimating

$$\overline{w}_{Dt} = \left(\frac{B(\boldsymbol{p}_t)}{\bar{e}_t}\right)^{\varepsilon}\left(\frac{D(\boldsymbol{p}_t)}{B(\boldsymbol{p}_t)}\right)^{\gamma}\kappa_t,$$

where κ_t is just a standard time fixed effect, is sufficient and we therefore avoid the need to estimate all the household-level effects ν_h.

A.5 Allowing for Hump-Shaped Expenditure Shares

Banks, Blundell and Lewbel (1997) stress the importance of allowing for hump-shaped expenditure shares to match the microeconomic data and it is possible to extend the baseline framework to allow for this at the household level. Following Alder, Boppart and Müller (2022), let preferences be characterized by an indirect utility function of the form

$$V(e, \boldsymbol{p}) = \frac{1}{\varepsilon}\left[\left(\frac{e - A(\boldsymbol{p})}{B(\boldsymbol{p})}\right)^{\varepsilon} - 1\right] - \frac{\nu}{\gamma}\left[\left(\frac{D(\boldsymbol{p})}{B(\boldsymbol{p})}\right)^{\gamma} - 1\right], \tag{A.7}$$

where the only difference to the PIGL specification in Equation (4) is the addition of a linearly homogeneous function $A(\boldsymbol{p})$ of prices. The expenditure shares of the three latent goods with price functions $A(\cdot)$, $B(\cdot)$ and $D(\cdot)$ are given by Roy's identity as

$$w_A = \frac{A(\boldsymbol{p})}{e}, \tag{A.8}$$

$$w_B = \left(1 - \frac{A(\boldsymbol{p})}{e}\right)\left[1 - \nu\left(\frac{B(\boldsymbol{p})}{e - A(\boldsymbol{p})}\right)^{\varepsilon}\left(\frac{D(\boldsymbol{p})}{B(\boldsymbol{p})}\right)^{\gamma}\right], \tag{A.9}$$

$$w_D = \left(1 - \frac{A(\boldsymbol{p})}{e}\right)\nu\left(\frac{B(\boldsymbol{p})}{e - A(\boldsymbol{p})}\right)^{\varepsilon}\left(\frac{D(\boldsymbol{p})}{B(\boldsymbol{p})}\right)^{\gamma}. \tag{A.10}$$

The shares w_j^A, w_j^B and w_j^D of total A, B and D expenditures allocated to an individual good j are given as before by $w_j^C = p_j C_j(\boldsymbol{p})/C(\boldsymbol{p})$, $C \in \{A, B, D\}$. Together with Equations (A.8) to (A.10), this implies an expenditure share w_j of good j in total expenditures of the form

$$w_j = p_j \left\{ \frac{A(\boldsymbol{p})}{e} \frac{A_j(\boldsymbol{p})}{A(\boldsymbol{p})} + \left(1 - \frac{A(\boldsymbol{p})}{e}\right) \left[\frac{B_j(\boldsymbol{p})}{B(\boldsymbol{p})} + \left(\frac{D_j(\boldsymbol{p})}{D(\boldsymbol{p})} - \frac{B_j(\boldsymbol{p})}{B(\boldsymbol{p})} \right) \nu \left(\frac{B(\boldsymbol{p})}{e - A(\boldsymbol{p})} \right)^{\varepsilon} \left(\frac{D(\boldsymbol{p})}{B(\boldsymbol{p})} \right)^{\gamma} \right] \right\}. \tag{A.11}$$

Since the first term on the right-hand side of (A.11) is decreasing in e while the second term can be either increasing or decreasing in e, this allows for expenditure shares that are non-monotonic in expenditures. The derivation of the exact price index of (A.7) is virtually identical to the PIGL case in Appendix A.1. The corresponding expenditure function of (A.7) is

$$c(u, \boldsymbol{p}) = \left[1 + \varepsilon \left(u + \frac{\nu}{\gamma} \left\{ \left(\frac{D(\boldsymbol{p})}{B(\boldsymbol{p})} \right)^{\gamma} - 1 \right\} \right) \right]^{\frac{1}{\varepsilon}} B(\boldsymbol{p}) + A(\boldsymbol{p}).$$

Suppose again that the reference utility u is that corresponding to the expenditure level in some base period s such that $c(u, \boldsymbol{p}_s) = e_s$ and $u \equiv V(e_s, \boldsymbol{p}_s)$. Using the indirect utility function (A.7) to substitute this into some period-t expenditure function, rearranging terms, and using $P_{Ct} = C(\boldsymbol{p}_t)/C(\boldsymbol{p}_s)$ together with Equations (A.8) to (A.10) yields

$$c(u, \boldsymbol{p}_t) = e_s \left\{ (1 - w_{As}) \left[\left(1 - \frac{\varepsilon}{\gamma} \frac{w_{Ds}}{1 - w_{As}} \right) P_{Bt}^{\gamma} + \frac{\varepsilon}{\gamma} \frac{w_{Ds}}{1 - w_{As}} P_{Dt}^{\gamma} \right]^{\frac{1}{\gamma} \cdot \frac{\gamma}{\varepsilon}} P_{Bt}^{1 - \frac{\gamma}{\varepsilon}} + w_{As} P_{At} \right\},$$

and it follows that the price index is

$$P_t^{IA} = (1 - w_{As}) \widetilde{P}_t^{\frac{\gamma}{\varepsilon}} P_{Bt}^{1 - \frac{\gamma}{\varepsilon}} + w_{As} P_{At}$$

where

$$\widetilde{P}_t \equiv \left[\left(1 - \frac{\varepsilon}{\gamma} \frac{w_{Ds}}{1 - w_{As}} \right) P_{Bt}^{\gamma} + \frac{\varepsilon}{\gamma} \frac{w_{Ds}}{1 - w_{As}} P_{Dt}^{\gamma} \right]^{\frac{1}{\gamma}}.$$

Writing $\widetilde{P}_t$ as a Sato-Vartia index, we finally obtain the household-level cost-of-living index

$$P_t^{IA} = (1 - w_{As}) P_{Dt}^{\frac{\gamma \phi_t}{\varepsilon}} P_{Bt}^{1 - \frac{\gamma \phi_t}{\varepsilon}} + w_{As} P_{At},$$

where ϕ_t is a Sato-Vartia weight as in Proposition 1 with

$$\psi_{Bt} \equiv \left(1 - \frac{\varepsilon}{\gamma} \frac{w_{Ds}}{1 - w_{As}} \right) \left(\frac{P_{Bt}}{\widetilde{P}_t} \right)^{\gamma} \qquad \text{and} \qquad \psi_{Dt} \equiv \frac{\varepsilon}{\gamma} \frac{w_{Ds}}{1 - w_{As}} \left(\frac{P_{Dt}}{\widetilde{P}_t} \right)^{\gamma}.$$

This index is a direct generalization of Proposition 1 and, as before, is computable given expenditure shares w_{As}, w_{Ds}, price indices P_{At}, P_{Bt}, P_{Dt} and parameter values for ε and γ.

Similarly to Proposition 2, under Assumption 1 and appropriate choices for P_{At}, P_{Bt} and P_{Dt}, estimation reduces to only the two parameters ε and γ which are readily obtained from Equations (A.8) and (A.10).

Unlike the baseline framework, however, these preferences do not aggregate as easily. As shown in Alder, Boppart and Müller (2022, Proposition 2), aggregate expenditure shares over N households are now

$$\overline{w}_A = \frac{A(\boldsymbol{p})}{\bar{e}},$$

$$\overline{w}_B = \left(1 - \frac{A(\boldsymbol{p})}{\bar{e}}\right)\left[1 - \nu\left(\frac{B(\boldsymbol{p})}{\bar{e} - A(\boldsymbol{p})}\right)^{\varepsilon}\left(\frac{D(\boldsymbol{p})}{B(\boldsymbol{p})}\right)^{\gamma}\right]\kappa,$$

$$\overline{w}_D = \left(1 - \frac{A(\boldsymbol{p})}{\bar{e}}\right)\nu\left(\frac{B(\boldsymbol{p})}{\bar{e} - A(\boldsymbol{p})}\right)^{\varepsilon}\left(\frac{D(\boldsymbol{p})}{B(\boldsymbol{p})}\right)^{\gamma}\kappa.$$

where

$$\kappa \equiv \frac{1}{N}\sum_{h=1}^{N}\left(\frac{e_h - A(\boldsymbol{p})}{\bar{e} - A(\boldsymbol{p})}\right)^{1-\varepsilon}.$$

Unlike the PIGL case, there is no representative level of expenditure in Muellbauer's (1975, 1976) sense, even though a representative agent exists.[18] Therefore, it is not possible to bake in the parameter κ into some representative level of expenditure and proceed as for an individual household. Instead, the expenditure function of the representative agent is now

$$c(u^{RA}, \boldsymbol{p}) = \left[1 + \varepsilon\left(u^{RA} + \frac{\nu\,\kappa}{\gamma}\left\{\left(\frac{D(\boldsymbol{p})}{B(\boldsymbol{p})}\right)^{\gamma} - 1\right\}\right)\right]^{\frac{1}{\varepsilon}} B(\boldsymbol{p}) + A(\boldsymbol{p}).$$

Similar steps as before gives an aggregate price index of the same form as above, $P_t^{IA} = (1 - \overline{w}_{As})P_{Dt}^{\frac{\gamma\phi_t}{\varepsilon}} P_{Bt}^{1-\frac{\gamma\phi_t}{\varepsilon}} + \overline{w}_{As}P_{At}$, but with weights given by

$$\psi_{Bt} \equiv \left(1 - \frac{\varepsilon}{\gamma}\frac{\kappa_t}{\kappa_s}\frac{\overline{w}_{Ds}}{1 - \overline{w}_{As}}\right)\left(\frac{P_{Bt}}{\widetilde{P}_t}\right)^{\gamma} \qquad \text{and} \qquad \psi_{Dt} \equiv \frac{\varepsilon}{\gamma}\frac{\kappa_t}{\kappa_s}\frac{\overline{w}_{Ds}}{1 - \overline{w}_{As}}\left(\frac{P_{Dt}}{\widetilde{P}_t}\right)^{\gamma}.$$

where

$$\widetilde{P}_t \equiv \left[\left(1 - \frac{\varepsilon}{\gamma}\frac{\kappa_t}{\kappa_s}\frac{\overline{w}_{Ds}}{1 - \overline{w}_{As}}\right)P_{Bt}^{\gamma} + \frac{\varepsilon}{\gamma}\frac{\kappa_t}{\kappa_s}\frac{\overline{w}_{Ds}}{1 - \overline{w}_{As}}P_{Dt}^{\gamma}\right]^{\frac{1}{\gamma}}.$$

Thus, to compute aggregate price indices, we now either need to know the inequality measures κ in the time periods considered, or we need to impose the rather strong assumption that these measures remain constant over time for *all* groups considered.

[18] The expenditure level e^{RA} that induces the average expenditure shares for A and D are given by $e^{RA} = \bar{e}$ and $\left(1 - \frac{A(\boldsymbol{p})}{e^{RA}}\right)\left(e^{RA} - A(\boldsymbol{p})\right)^{-\varepsilon} = \left(1 - \frac{A(\boldsymbol{p})}{\bar{e}}\right)\left(\bar{e} - A(\boldsymbol{p})\right)^{-\varepsilon}\kappa$, respectively, and these generally differ.

Appendix B Additional Figures and Tables

Table B.1. CEX-CPI crosswalk.

	CEX category	**CPI name**	**CPI code**
1	Food at home	Food at home	SAF11
2	Food away from home	Food away from home	SEFV
3	Alcoholic beverages	Alcoholic beverages	SAF116
4	Rented dwellings	Rent of primary residence	SEHA
5	Owned dwellings[a]	Owners' equivalent rent of primary residence	SEHC
6	Other lodging	Lodging while out of town[b]	MUUR0000SE2102
		Lodging away from home[b]	SEHB
7	Utilities	Household energy	SAH21
8	Water	Water and sewerage maintenance	SEHG01
9	Phone	Communication	SAE2
10	Household F&O[c]	Household furnishings and operations	SAH3
11	Apparel	Apparel	SAA
12	Gasoline	Motor fuel	SETB
13	Other vehicle expenses	Motor vehicle maintenance and repair	SETD
		Motor vehicle insurance	SETE
		Motor vehicle fees	SETF
14	Public transportation	Public transportation	SETG
15	Health	Medical care	SAM
16	Entertainment	Recreation	SAR
17	Personal care	Personal care	SAG1
18	Reading	Recreational reading materials	SERG
19	Education	Education and communcation	SAE
20	Tobacco	Tobacco and smoking products	SEGA
21	Other expenses	Miscellaneous personal services	SEGD

Notes. The CEX categories follow the hierachical groupings defined by the BLS. CPI are non-seasonally adjusted nationwide data for urban consumers. *a.* Rental equivalence value of owned dwellings as reported by the households. *b.* "Lodging away from home" from 1995–1997 and "lodging while out of town" afterwards. *c.* Furnishing and operations, includes "household operations", "housekeeping supplies" and "household furnishings and equipment".

Table B.2. Marginal effect of a change in expenditure decile on expenditure share.

Classification	Goods category	Marginal effect	SE
Luxuries	Owned dwellings	1.757	(0.019)
	Household operations and furnishing	0.690	(0.010)
	Entertainment	0.386	(0.006)
	Other lodging	0.303	(0.005)
	Food away from home	0.299	(0.005)
	Education	0.280	(0.006)
	Health	0.215	(0.007)
	Public transport	0.137	(0.003)
	Other vehicle expenses	0.124	(0.006)
	Other expenses	0.105	(0.002)
	Apparel	0.098	(0.004)
	Alcoholic beverages	0.044	(0.001)
	Personal care	0.031	(0.001)
	Reading	0.028	(0.001)
Necessities	Food at home	−2.031	(0.012)
	Rented dwellings	−1.496	(0.017)
	Utilities	−0.362	(0.004)
	Gasoline	−0.219	(0.004)
	Phone	−0.186	(0.003)
	Tobacco	−0.158	(0.003)
	Water	−0.045	(0.002)
Expenditure category dummies		Yes	
Observations		1,562,211	
Adjusted R^2		0.579	

Notes. The table shows the coefficient estimates (clustered standard errors in parentheses) from a weighted least square regression of expenditure shares on the expenditure decile interacted with individual good dummies using the CEX household sampling weights. We include expenditure category dummies and cluster standard errors on household level.

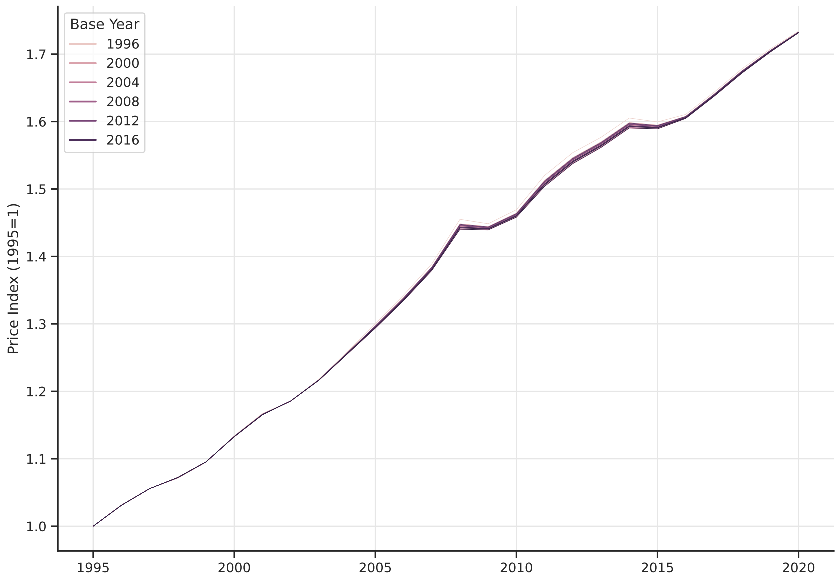

FIGURE B.1. PIGL representative agent G-SV price index for different base years.

Notes. The price index is calculated under weak separability. Each line represents the representative agent price index for a different base year, but normalized to one in 1995.

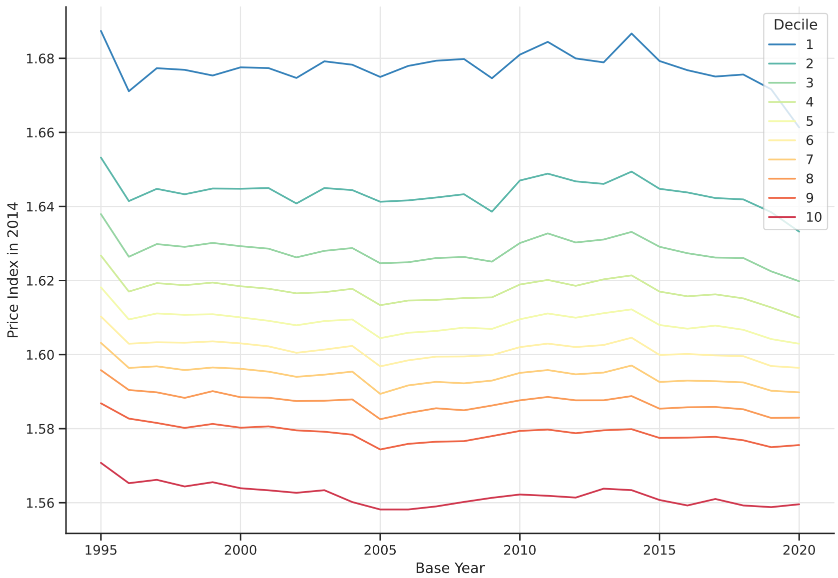

FIGURE B.2. G-SV price index in 2014 by expenditure decile for different base years.

Notes. The price index is calculated under weak separability. The horizontal axis describes the base year of the price index and the vertical axis the respective value of the price index in 2014. Price indices are all normalized to one in 1995. The price index for each expenditure decile is calculated as the price index of the PIGL representative agent over households within each respective decile.

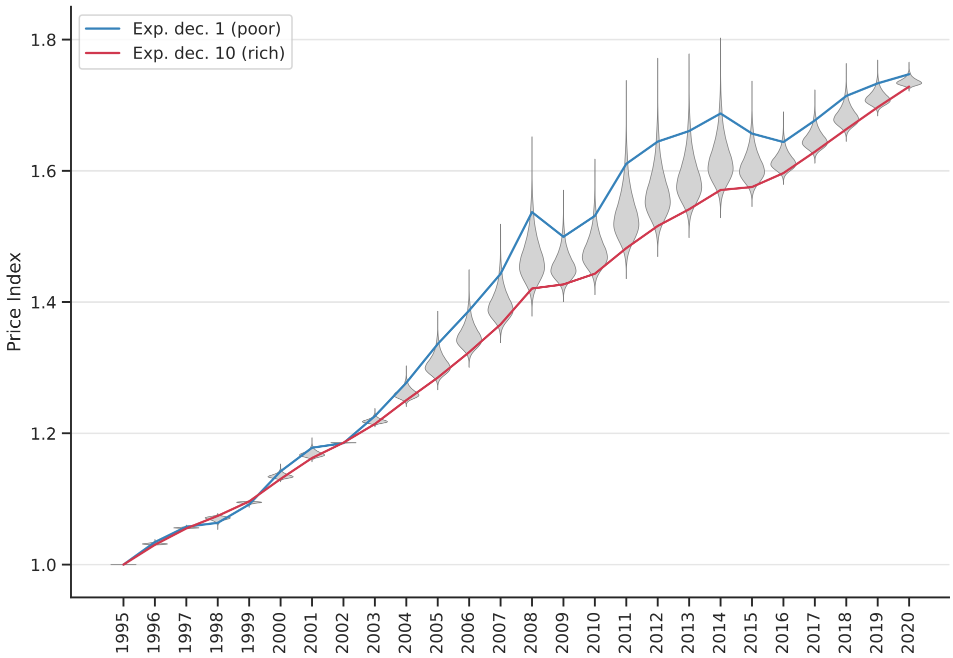

Figure B.3. Distribution of the G-SV price index under weak separability.

Notes. The gray shaded area shows the kernel density estimate of the distribution of G-SV price indices. The index is calculated for each household in the sample of 1995. The blue line shows the PIGL RA price index for the poorest 10 percent. The red line shows the PIGL RA price index for the richest 10 percent.

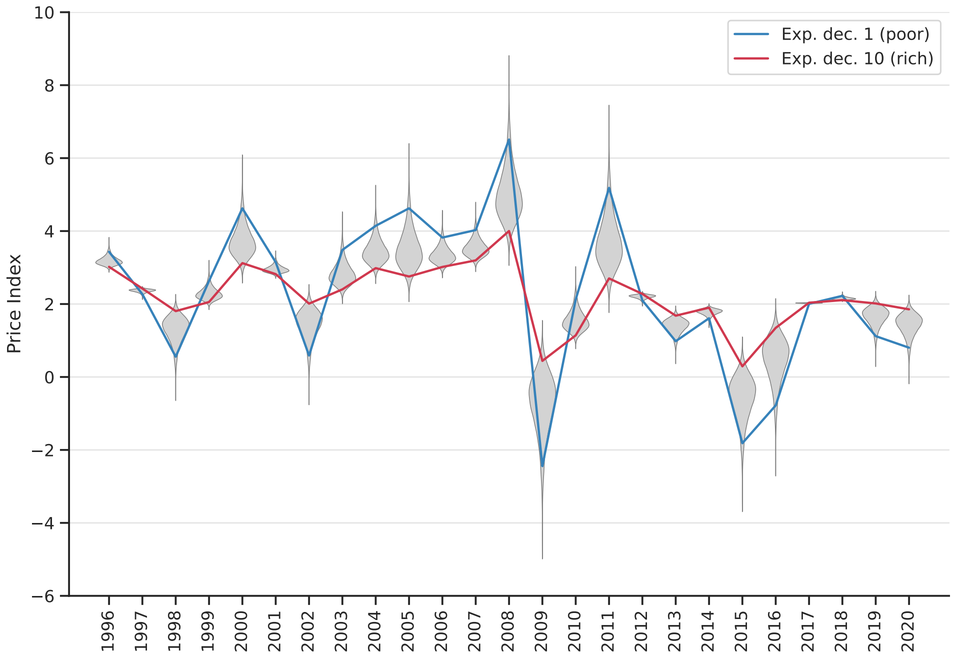

Figure B.4. Distribution of the G-SV inflation under weak separability.

Notes. The gray shaded area shows the kernel density estimate of the distribution of inflation rates. The inflation is calculated for each household in the sample of 1995. The blue line shows the PIGL RA inflation for the poorest 10 percent. The red line shows the PIGL RA inflation for the richest 10 percent.

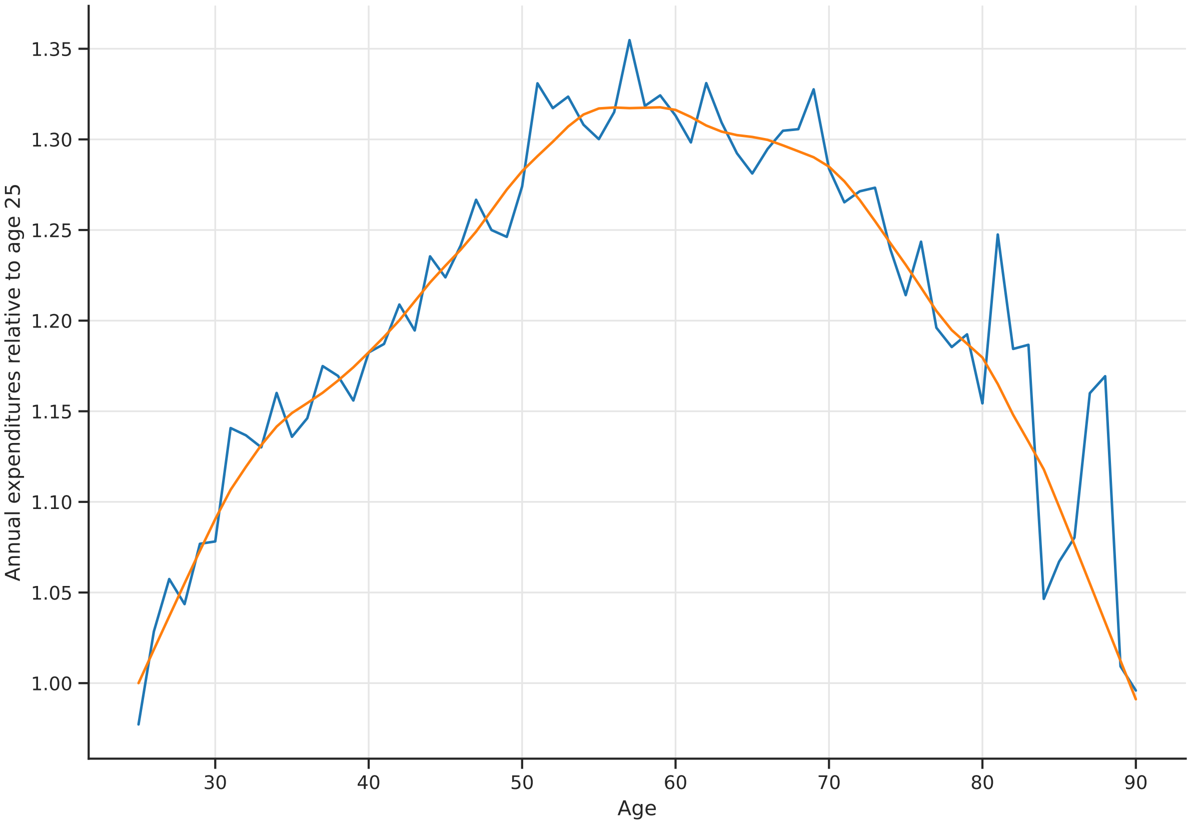

FIGURE B.5. Expenditures over the life cycle.

Notes. The blue line shows the estimated life-cycle expenditures, $\hat{\alpha}_i$'s, from the model in Equation (21). The orange line shows smoothed expenditure levels using a Locally Weighted Scatterplot Smoothing (LOWESS). All measures are relative to the smoothed expenditure level at age 25.

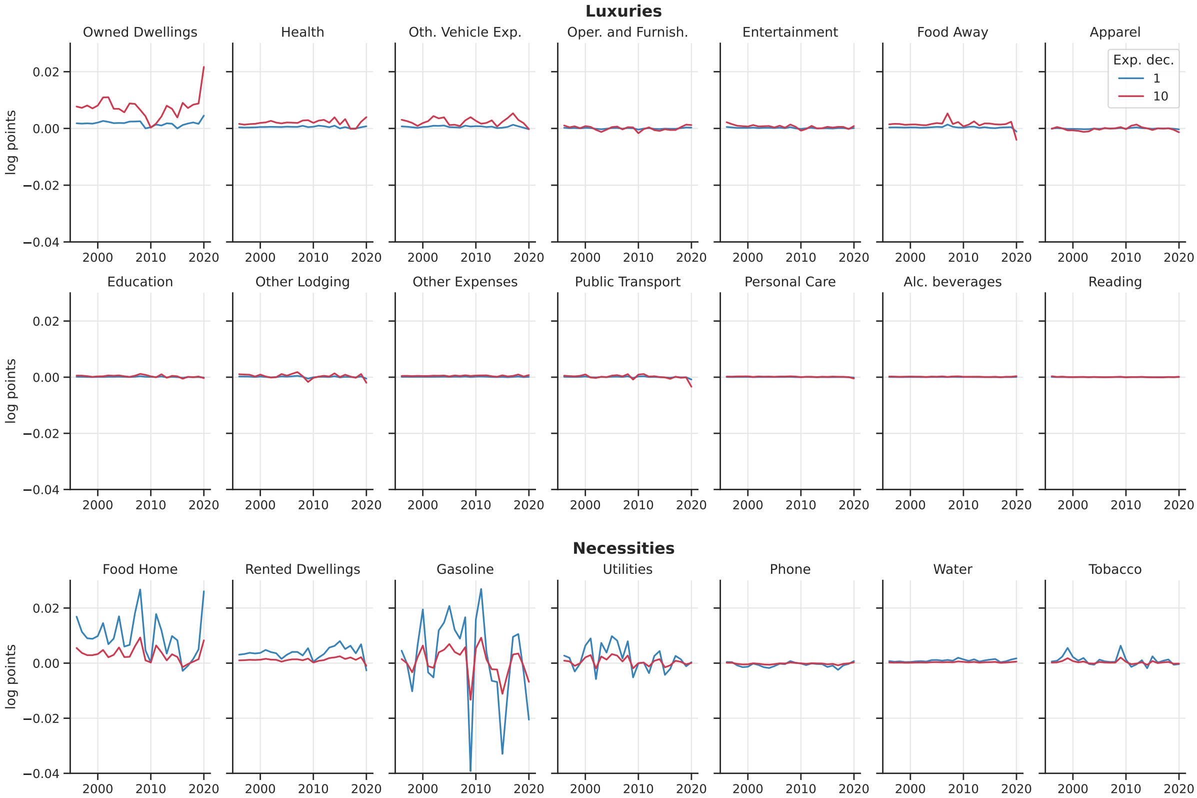

FIGURE B.6. Inflation decomposition by expenditure categories.

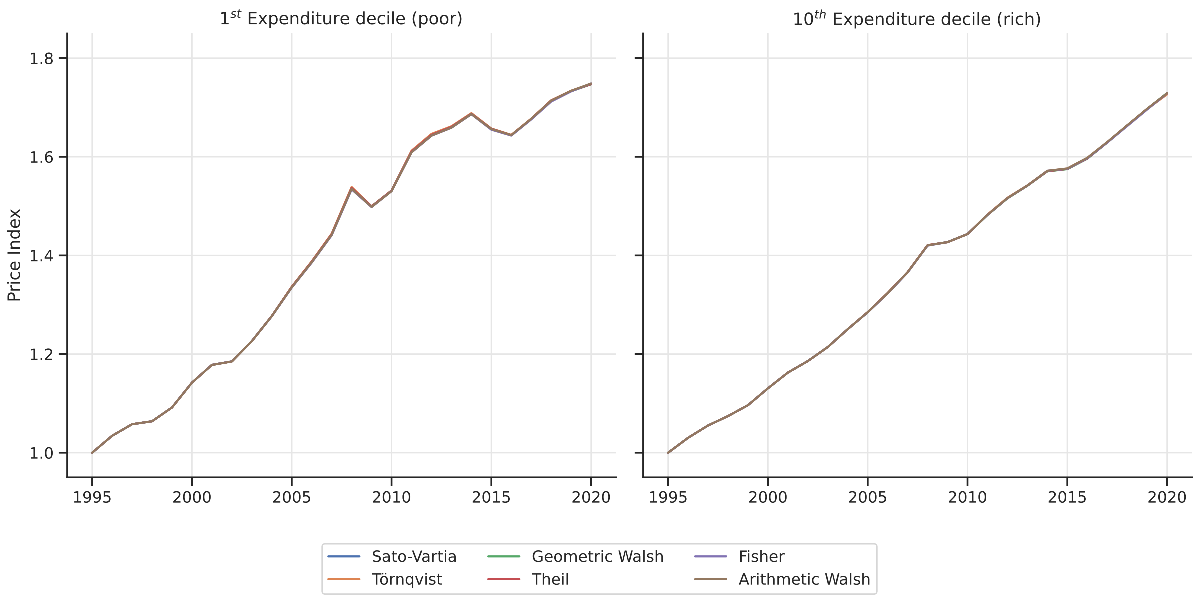

FIGURE B.7. Comparison of different generalized superlative price indices.

Figure B.8. Comparison of inflation across different generalized superlative indices.