Measuring cost of living inequality during an inflation surge

Tao Chen, Peter Levell and Martin O'Connell*

August 2026

Abstract

We study cost-of-living inequality during the 2021–2023 UK inflation surge using household scanner data on fast-moving consumer goods. We develop and implement a non-homothetic, index-number-based decomposition of cost-of-living changes into exposure, substitution, and a non-homothetic adjustment. We document pronounced "cheapflation": within narrowly defined categories, prices rose fastest for lower-quality products disproportionately consumed by lower-resource households. This pattern generates sharply regressive inflation exposure. Behavioural responses compressed, but did not eliminate, the resulting inequality in cost-of-living growth. Substitution attenuated the exposure gradient, while the alignment between cheapflation and household shifts down the quality ladder raised average equivalent-variation losses. The episode also reshaped exposure to a repeat cheapflation shock: households with the largest negative non-homothetic adjustments ended the period with baskets more exposed than their initial baskets.

Keywords: inflation, cost-of-living, inequality, heterogeneity, non-homotheticity
JEL classification: D12, D30, E31, I30
Acknowledgements: The authors would like to gratefully acknowledge financial support from the ESRC under grant number CPP: ES/T014334/1. Chen also gratefully acknowledges financial support from the UBEL DTP Research Training Support Grant. We also would like to thank Richard Blundell, Ian Crawford, Rachel Griffith, Xavier Jaravel, Rasmus Lentz, Lars Nesheim, Cormac O'Dea, Áureo de Paula and Imran Rasul for helpful comments. The contents of this publication are IFS's own analysis and findings. The analysis has been carried out by IFS using data from Worldpanel by Numerator's Take Home data (all use is subject to Worldpanel by Numerator's terms and conditions). Worldpanel by Numerator does not endorse, and is not responsible for, the efficacy or accuracy of IFS's analysis and findings. The use of Worldpanel by Numerator's data does not imply any endorsement by Worldpanel by Numerator's data of the interpretation or analysis of the data. Worldpanel by Numerator cannot independently verify the findings, nor can it endorse the views or conclusions set out in this article. All errors and omissions remain the responsibility of the authors of this publication.
Correspondence: tao.chen.16@ucl.ac.uk, peter_l@ifs.org.uk, moconnell9@wisc.edu

*Chen is at University College London, Levell is at the Institute for Fiscal Studies, and O'Connell is at the University of Wisconsin-Madison and the Institute for Fiscal Studies.

1 Introduction

Inflation re-emerged as a central economic concern in the aftermath of the COVID-19 pandemic, with many advanced economies experiencing the sharpest and most persistent price increases in several decades. Beyond its aggregate level, a key question is how the burden of inflation is distributed across households. Since consumption baskets differ systematically across households with different resources, the same vector of price changes can translate into very different changes in living standards, potentially reinforcing existing inequalities. These distributional effects depend not only on which prices move, but also on how households adjust their spending patterns in response.

In this paper, we study how the recent UK inflation surge affected cost-of-living inequality. We focus on fast-moving consumer goods (food, beverages, household goods, and personal care products) using household-level scanner data covering the period 2012–2023. The 2021Q3–2023Q3 episode we analyse features two salient characteristics: (i) unusually large and heterogeneous relative price movements within narrowly defined product categories, and (ii) falling real expenditure for many households. These features interacted: prices rose fastest for cheaper, lower-quality products ("cheapflation"), disproportionately consumed by lower-resource households, and households—especially those whose real expenditure fell most—shifted down the quality ladder toward the same part of the product space.

To quantify the welfare consequences of these price movements and household responses, we develop and implement an index-number-based decomposition of cost-of-living changes under general non-homothetic preferences. This decomposition separates cost-of-living changes into exposure, substitution, and a non-homothetic adjustment, and is implementable in scanner data without estimating a parametric demand system. We use this framework to show that cheapflation-driven inflation exposure is strongly regressive; that behavioural responses compressed, but did not eliminate, this regressive gradient; and that, as purchasing power fell, the alignment between cheapflation and household shifts down the quality ladder raised average equivalent-variation losses and reshaped households' exposure to a repeat cheapflation shock.[1]

We begin by providing descriptive evidence on prices and spending along a product-level quality ladder. We construct quality rungs within each narrow product category by ranking brand–size combinations by unit price, netting out nonlinear pricing over pack sizes. Over 2021Q3–2023Q3, price growth is sharply tilted toward the bottom of this within-category

[1] We focus on the consumption-basket channel through which heterogeneous spending patterns translate a common price vector into different welfare changes across households. Other channels operate through balance sheets, wage rigidities, and asset prices; see, for example, Doepke and Schneider (2006), Ferreira et al. (2026), and Del Canto et al. (2025).

ladder: products in the bottom two rungs exhibit average price increases of around one-third, compared to roughly one-fifth for products in the top two rungs. In earlier nine-quarter periods, price growth is essentially flat across rungs. This compression of the within-category price distribution is consistent with the cheapflation documented using price data in other countries by Cavallo and Kryvtsov (2024).

We also show that households with fewer resources are systematically more exposed to this pattern of price changes. Throughout our sample, lower-equivalised-expenditure and lower-equivalised-income households purchase lower-rung, cheaper products on average, and the strength of this relationship is remarkably stable over time; the key change in 2021Q3–2023Q3 is that price growth becomes sharply tilted toward these lower-rung products. Cheapflation therefore raises inflation exposure disproportionately for lower-resource households. We show that much of the resulting inflation exposure inequality is missed by measures that do not capture within-category differences in product choice along the quality ladder—for example, distributional inflation statistics that rely only on category-level spending and price indexes. Using within-household variation, we show that households with larger declines in deflated expenditure are also those that shift more toward lower-rung products. Together, these facts highlight three central forces for inflation inequality within fast-moving consumer goods: cheapflation, systematic cross-sectional differences in quality choices, and non-homothetic within-household adjustments along the quality ladder.

To organise these forces, we develop a decomposition of cost-of-living changes that is suited to environments with pronounced within-category relative price dispersion and non-homothetic demand. We focus on the cost-of-living index evaluated at final-period utility, which, for a given path of nominal expenditure growth, determines the proportional equivalent-variation loss from inflation. We show that, up to a second-order approximation, the index decomposes into exposure, substitution, and a non-homothetic adjustment.

The *exposure* term measures inflation exposure at the initial consumption bundle. The *substitution* term captures how substitution away from goods with strong relative price increases mitigates cost-of-living growth. We decompose substitution into a fixed-budget-share adjustment and a deviation from this benchmark that depends on Hicksian responses of budget shares to prices. The *non-homothetic adjustment* captures how the cost-of-living index changes when evaluated at final rather than initial utility. It is governed by the product of a real expenditure term and a covariance between relative price changes and expenditure elasticities. In a cheapflation episode where low-expenditure-elasticity goods become relatively more expensive, falling real expenditure pushes households along Engel curves toward precisely those goods whose prices are rising. This alignment of relative price growth and movements along Engel curves raises the cost-of-living index evaluated at final utility.

A key advantage of our framework is that each component can be implemented using observable household-level price indexes. Exposure is measured by a Laspeyres index, substitution is measured by comparing Laspeyres, geometric-Laspeyres, and non-homothetic cost-of-living indexes evaluated at initial utility, and the non-homothetic adjustment is measured by comparing non-homothetic indexes evaluated at initial and final utility. We approximate these non-homothetic indexes using the algorithm in Jaravel and Lashkari (2024). This delivers an index-number-based implementation of the decomposition that can be computed for every household in scanner data without specifying a parametric demand system at the product level.

We use this framework to quantify the contribution of exposure, substitution, and the non-homothetic adjustment to cost-of-living inequality over 2021Q3–2023Q3. Using a household-specific Laspeyres index, we find that average cumulative inflation exposure over this nine-quarter period is about 26%, with a standard deviation of 5.4 percentage points. Inequality in exposure is striking: the bottom decile of the 2021 equivalised expenditure distribution faces about 7.5 percentage points higher cumulative exposure than the top decile; the corresponding gap across income deciles is around 4.4 percentage points. Using a hierarchical decomposition of the Laspeyres index across segments (e.g., alcohol and dairy), categories, and products, we show that within-category differences in product choice along the quality ladder account for most of the dispersion in cumulative inflation exposure and the majority of the regressive gradient. These patterns do not appear in earlier nine-quarter periods.

We next turn to the cost-of-living index evaluated at the initial utility level, which reflects inflation exposure and substitution responses. On average, allowing for substitution reduces the cumulative cost-of-living increase from 26.2% under the no-substitution benchmark to about 24.5%. The fixed-budget-share Cobb–Douglas benchmark would lower the cost-of-living increase by approximately 1.9 percentage points, but deviations from this benchmark offset around 0.3 percentage points of that gain. This indicates that products experiencing the strongest price growth are precisely those from which households substitute away least strongly. Across the expenditure distribution, substitution modestly dampens the cost-of-living gradient: the bottom–top decile gap falls from about 7.5 to 6.7 percentage points, but remains large and regressive.

We then measure the non-homothetic adjustment. On average, the cost-of-living index evaluated at final utility is about 0.5 percentage points higher than at initial utility. This reflects the interaction of cheapflation and declining real expenditure: as purchasing power falls, households move along Engel curves toward lower-expenditure-elasticity products whose relative prices are increasing, making it more expensive to attain the realised final utility level than the initial one.

Overall, behavioural responses compress, but do not overturn, the regressive distributional pattern implied by initial-basket exposure. The sharp inflation exposure gradient across both the expenditure and income distributions therefore translates into similarly regressive gradients in inflation-driven welfare losses.

Finally, we ask whether cheapflation and trading down reshaped households' exposure to a repeat cheapflation shock. We construct a repeat-exposure Laspeyres index that applies the realised 2021Q3–2023Q3 pattern of relative price changes to each household's final-period consumption bundle. We find that households with the largest negative non-homothetic adjustments during 2021Q3–2023Q3 ended the episode with baskets more exposed to a repeat cheapflation shock than their initial baskets. These households were not initially the most exposed to cheapflation; instead, their movements along Engel curves drew them toward products whose relative prices had risen most. This intertemporal reallocation of exposure suggests that the distributional consequences of cheapflation may persist through household basket composition even after aggregate inflation recedes.

Our paper contributes to three strands of literature. First, we add to work on inflation inequality, including studies using scanner data (e.g., Kaplan and Schulhofer-Wohl 2017; Jaravel 2019; Argente and Lee 2021) and survey-based measures that replicate official CPI methods (e.g., Jaravel 2024). We show that, in a recent high-inflation episode, inflation inequality in the UK is large and regressive across both expenditure and income groups, and that within-category differences in product choice along the quality ladder account for the bulk of this inequality. Second, we connect to the emerging literature that documents recent patterns of cheapflation (Cavallo and Kryvtsov 2024) and links them to firm pricing responses to cost shocks and demand shifts (Sangani 2025, 2026; Becker 2025), by showing how cheapflation in lower-quality products and falling purchasing power interact to shape both realised equivalent-variation losses and repeat exposure. Third, we contribute to the literature on non-homothetic demand and measurement of living standards and welfare (e.g., Atkin et al. 2024; Auer et al. 2024; Baqaee and Burstein 2023; Baqaee et al. 2024; Balk 1990; Comin et al. 2021; Jaravel and Lashkari 2024). We use a flexible second-order approximation in the spirit of Baqaee and Burstein (2023), and show how the resulting decomposition can be implemented as differences between household-level price indexes under general non-homothetic preferences.

The remainder of the paper is structured as follows. Section 2 introduces the scanner data, product classification, and quality ladder, and presents descriptive evidence on cheapflation, cross-sectional exposure, and trading down. Section 3 develops the cost-of-living decomposition and maps its components to observable price indexes, and Section 4 presents the results. Section 5 studies how cheapflation and trading down reshape households' exposure to repeat cheapflation shocks. Section 6 concludes.

2 Data and motivating evidence

Scanner data

We use household-level scanner data from *Worldpanel by Numerator's Take Home Purchase Panel* (henceforth Numerator data). The dataset tracks purchases of fast-moving consumer goods—including food, alcoholic and non-alcoholic beverages, toiletries, cleaning products, and pet food—brought into the home by households living in Great Britain (i.e., the UK excluding Northern Ireland). It includes both brick-and-mortar and online purchases.

Households typically remain in the dataset for several years. Each participating household records all purchased UPCs (barcodes) using a handheld scanner or mobile app and submits receipts electronically or by post. For each transaction, we observe quantity, expenditure, price paid, and UPC characteristics. We also observe socio-demographic information, including household structure and banded household income.

We focus on the nine calendar quarters from 2021Q3 to 2023Q3, a period of elevated inflation. Over this period, the Consumer Price Index (CPI) for food and non-alcoholic beverages, a central component of fast-moving consumer goods, rose from 103.3 to 133.5 after a prolonged period of stability.[2] Our main analysis sample includes 19,030 households that recorded their purchases in every quarter over this period.[3] We compare inflation inequality over 2021Q3–2023Q3 to four earlier nine-quarter periods: 2012Q1–2014Q1, 2014Q1–2016Q1, 2016Q1–2018Q1, and 2018Q1–2020Q1.[4]

Our dataset offers several advantages for measuring inflation inequality compared with commonly used alternatives. First, expenditures are recorded continuously and available with minimal lag, allowing us to construct inflation measures based on contemporaneous spending patterns. By contrast, the expenditure information used in official distributional inflation statistics typically comes from annual budget surveys that are released with a substantial delay relative to the underlying price microdata. Second, the scanner panel tracks

[2] In January 2019, the index was 102.6, and by December 2024 it had reached 137.8. One exception to the prior stability was a spike and subsequent reduction in inflation at the start of the COVID-19 pandemic, driven by a decline and recovery in promotional activity (Jaravel and O'Connell 2020a,b).

[3] We exclude households that are not continuously present across all nine quarters, or whose quarterly expenditure falls below the quarter-specific 5^{th} percentile, which averages £114 across quarters. Our results are not sensitive to these restrictions.

[4] Each of these periods includes 19,000–20,000 households. We exclude 2020Q2–2021Q2, as purchasing patterns during this time are likely atypical due to lockdowns and social distancing measures. Including this period does not materially affect our results.

households over time, enabling us to construct household-specific price indexes. Third, expenditure is recorded at the UPC level, so our inflation measures reflect spending patterns across narrowly defined products.[5]

These advantages matter for our application. Budget surveys and price microdata commonly used to study household-level inflation inequality typically aggregate prices to around 100–200 item groups (see Office for National Statistics 2022 for the UK and Klick and Stockburger 2021, 2024 for the US). Our scanner data cover a narrower set of fast-moving consumer goods, but allow us to study within-category substitution and within-category price changes across narrowly defined products, which is central to our focus on cheapflation and quality downgrading.

Product classification

Our data cover approximately 200,000 unique UPCs over 2021Q3–2023Q3. Because some UPCs are occasionally replaced by nearly identical ones, we define products using brand–package-size pairs rather than individual UPCs. Numerator provides highly disaggregated brand information, so this approach involves minimal aggregation over meaningfully distinct products. Over this period, there are approximately 100,000 brand–size pairs, which we refer to as products throughout.[6]

We denote household h's year-quarter t expenditure on product i as x_{hit} and their total year-quarter expenditure as $x_{ht} = \sum_i x_{hit}$. Household h's year-quarter t quantity of product i is q_{hit}. We measure the period-t price of product i as its unit value,[7]

$$p_{it} = \frac{\sum_h x_{hit}}{\sum_h q_{hit}}. \tag{2.1}$$

This definition assigns a common price to each product-quarter. In practice, lower-income households pay slightly lower prices for identical products—consistent with greater search effort, as in Aguiar and Hurst (2007)—but the price gap between low- and high-income households is modest and relatively stable, narrowing by 0.5 percentage points from 2021Q3 to 2023Q3 (see Appendix A.1). This change is small relative to the inflation inequality we document below and, if anything, modestly reinforces the patterns we find.

We use a hierarchical product classification provided by Numerator, designed to allocate products into well-defined consumer markets. Products are grouped into 10 segments—

[5] Store-level scanner data, which are now used in some countries for CPI construction, also provide disaggregated and timely expenditure information. However, since they are recorded at the store rather than the household level, they are not well suited, in isolation, for studying the distribution of inflation across households.

[6] For some fresh produce, such as fruit, vegetables, and meat, UPCs do not have a well-defined brand. In these cases, we define products using UPCs. Our results are materially unchanged if we define all products based on UPCs.

[7] Because unit values are constructed from transaction expenditures and quantities, they reflect the prices households actually pay, including temporary discounts and multi-buy offers.

bakery, dairy, fresh fruit and vegetables, meat and fish, prepared food, cupboard ingredients, confectionery, non-alcoholic drinks, alcohol, and household goods—and 238 categories. For example, the product *Coca Cola 2 litre bottle* belongs to the category *colas* within the *non-alcoholic drinks* segment.

The quality ladder. We further segment the product space by defining a quality ladder within each product category. Previous work (e.g., Jaravel 2019) classifies products within narrowly defined categories into price deciles, which serve as a proxy for quality. We adopt a similar approach, but first adjust prices for nonlinear pricing across different pack sizes of the same brand (e.g., see Griffith et al. 2009). We define quality rungs using adjusted prices early in the episode and then hold rung assignments fixed when measuring inflation over 2021Q3–2023Q3. Thus, "cheapflation" reflects differential price growth across a fixed within-category ranking.

To remove nonlinear pricing across pack sizes, we estimate the expenditure-weighted regression:

$$p_{it} = \xi_{b(i)} + \tau_{c(i)t} + \sum_{y} \sum_{l=1}^{3} \alpha_{c(i)y}^{(l)} \mathbb{1}\{t \in y\} \times \text{size}_i^{(l)} + \epsilon_{it}, \tag{2.2}$$

where p_{it} is the price of product i at time t, $\xi_{b(i)}$ are brand effects, $\tau_{c(i)t}$ are category-by-year-quarter fixed effects, and the α-terms form a third-order polynomial in demeaned pack size, with coefficients that vary by category and year.

We then construct the quality ladder as follows. For products available in the first four quarters of the episode, 2021Q3–2022Q2, we compute adjusted prices by removing the estimated size component:

$$\tilde{p}_{it} = p_{it} - \sum_{y} \sum_{l=1}^{3} \alpha_{c(i)y}^{(l)} \mathbb{1}\{t \in y\} \times \text{size}_i^{(l)}, \quad t \leq 4.$$

We average these adjusted prices across the first four quarters and use the expenditure-weighted distribution within each product category to define decile boundaries. For all products, including those introduced after 2022Q2, we then use equation (2.2) to predict an adjusted baseline price,

$$\tilde{p}_i = \xi_{b(i)} + \frac{1}{4} \sum_{t=1}^{4} \tau_{c(i)t}$$

and assign products to quality-ladder rungs using the category-specific decile boundaries. This procedure assigns new products to quality rungs based on a baseline-equivalent adjusted price that accounts for category-specific price growth. We construct the quality ladder analogously for the four earlier comparison periods.

Income measure

Our data provide two measures of household economic resources: (i) total household fast-moving consumer goods expenditure and (ii) household income, reported in £10,000 bands and top-coded at £70,000. For both variables, we construct equivalised measures using the OECD-modified scale (Hagenaars et al. 1994).[8]

In each nine-quarter period, we use equivalised expenditure from the calendar year in which the period begins to assign households to expenditure percentiles or deciles. We refer to these groups as expenditure percentiles and expenditure deciles. We also group households into deciles of equivalised income, using assigned band midpoints to construct the equivalised income measure. These two rankings are strongly, though not perfectly, correlated (see Appendix A.2). In Section 4.1, we discuss the conceptual and measurement differences between expenditure- and income-based rankings, and in the results below we systematically report patterns by both.

Descriptive evidence

Our objective is to measure inflation inequality during an inflationary surge, and to characterise the role of household spending patterns and behavioural adjustments. Figure 2.1 provides preliminary evidence that motivates our approach. It shows that lower-rung products experienced more rapid price growth during the inflation surge, that these products are more popular among households with fewer resources, and that households with larger declines in observed spending shifted more strongly down the quality ladder. These patterns point toward the joint role of price dynamics across the quality ladder, cross-sectional differences in consumption baskets, and changes in spending patterns in shaping household exposure and responses to inflation.

Panel (a) reports the change in average price over 2021Q3–2023Q3, weighted by initial-period aggregate spending shares, for products on different rungs of the quality ladder (black line). Proportional price increases are substantially higher for products on lower rungs than for those at the top of the ladder. For instance, the average price increases for products on the bottom two rungs, 1 and 2, are 35.5% and 31.9%, respectively, while those for products on the top two rungs, 9 and 10, are 18.5% and 14.4%. Over this period, the within-category price distribution compresses. This pattern is also evident in the microdata underlying the UK CPI—including for fast-moving consumer goods, clothing, and leisure goods (see Appendix A.3). Cavallo and Kryvtsov (2024) term this phenomenon *cheapflation*. By contrast, price

[8] This involves constructing an equivalised household size, where the first adult counts as 1, additional individuals aged 14 or over count as 0.5, and each child under 14 counts as 0.3.

growth is much flatter across the quality ladder in the preceding nine-quarter periods (grey lines).

Panel (b) illustrates how the average quality rung of products purchased in the initial quarter of each nine-quarter period varies across expenditure deciles. Higher-expenditure households consistently purchase higher-rung products on average, and this relationship is remarkably stable across all nine-quarter periods. A similar pattern holds across deciles of the equivalised income distribution (Appendix A.4). Taken together, panels (a) and (b) show that the 2021Q3–2023Q3 episode combines strong cheapflation with a persistent cross-sectional relationship between resources and average quality rung: the products that became relatively more expensive are disproportionately consumed by lower-resource households.

Panels (c) and (d) focus on 2021Q3–2023Q3. Panel (c) shows that the share of first-quarter expenditure allocated to bottom-rung products decreases sharply across expenditure percentiles, whereas the spending share on top-rung products increases across percentiles. This pattern is consistent with lower-rung products having lower expenditure elasticities than higher-rung products.

Panel (d) illustrates the relationship between *within*-household changes in expenditure and changes in spending shares on top- and bottom-rung products between 2021Q3 and 2023Q3. On average, households increase their spending shares on both top- and bottom-rung products. However, households with larger declines in deflated expenditure are also those that increase their spending shares on bottom-rung products more, and increase their spending shares on top-rung products less. This descriptive evidence is consistent with households moving along quality-ladder Engel curves as their real expenditure changes.

One possibility is that the relationship in Figure 2.1(d) partly reflects a mechanical link between quality downgrading and measured expenditure growth: trading down to cheaper products reduces nominal expenditure, which may mechanically lower deflated expenditure. To assess whether this is the case, we use a household–segment regression that relates within-segment quality changes to changes in deflated total expenditure. Instrumenting deflated-expenditure growth with a leave-one-out measure that excludes the segment whose quality change is being studied yields a similar, though slightly smaller, coefficient (see Appendix A.5). This suggests that the pattern in Figure 2.1(d) is not primarily driven by the mechanical construction of deflated expenditure. We nevertheless interpret the evidence as descriptive: it is consistent with households moving along quality-ladder Engel curves as real expenditure changes, but does not by itself isolate exogenous changes in purchasing power.

Figure 2.1: *Prices and spending across the quality ladder*

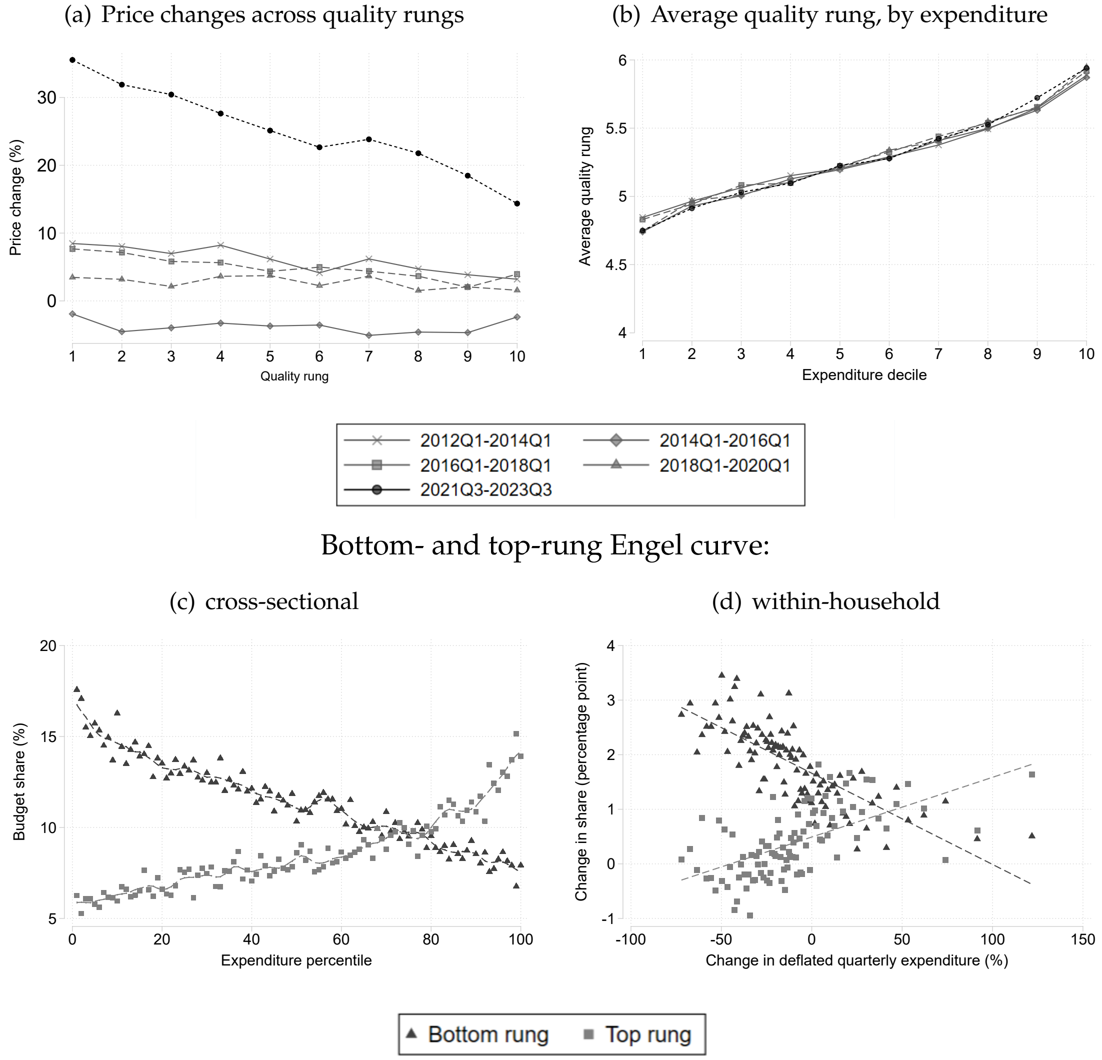

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2012–2023). Panel (a) reports change in average price over the nine-quarter period for products on each rung of the quality ladder, weighted by initial-period aggregate spending shares. Panel (b) reports the average quality rung of households' purchases by expenditure deciles. Panel (c) reports the average household spending share allocated to top- and bottom-rung products in 2021Q3, by expenditure percentile. The dashed lines are local polynomial-smoothed regressions. Panel (d) shows the average percentage point change in spending share allocated to bottom- and top-rung products from 2021Q3 to 2023Q3 for each percentile of the distribution of percentage changes in deflated quarterly expenditure over this period. Expenditure changes are deflated using a household-specific Laspeyres price index.

Discussion. The descriptive evidence in Figure 2.1 highlights three features of the 2021Q3–2023Q3 inflation episode that are central to our cost-of-living analysis. First, price growth is sharply tilted toward the bottom of the within-category quality ladder: lower-rung products experienced systematically higher inflation, in line with the cheapflation documented using price data by Cavallo and Kryvtsov (2024). Second, households with fewer resources were already more exposed to these lower-rung products before the episode began, echoing evidence in Argente and Lee (2021) and Jaravel (2024) that lower-income households buy lower-priced varieties and are more concentrated in necessity goods. Third, households with larger declines in deflated expenditure are observed to shift down the quality ladder *toward* products whose prices rose more quickly.

Taken together, these patterns imply that the distributional consequences of the inflation surge cannot be understood solely through household-specific Laspeyres indexes or by treating observed changes in expenditure shares as substitution responses to relative price changes. Laspeyres indexes capture differences in initial exposure, but do not allow for behavioural adjustment. Indexes that use changing expenditure shares can capture basket adjustment, but a homothetic cost-of-living interpretation treats those changes as substitution responses to relative prices. With non-homothetic preferences, changing shares may also reflect movements along Engel curves as real expenditure changes. In a cheapflation environment with falling purchasing power, this distinction matters: households may substitute away from products whose relative prices are rising, while movements along Engel curves may shift them toward lower-rung products whose prices are rising more quickly. This combination of within-category price changes, cross-sectional exposure differences, and within-household shifts down the quality ladder motivates the non-homothetic cost-of-living decomposition developed in the next section, together with an index-number implementation that disentangles exposure, substitution, and the non-homothetic adjustment.

3 Theory and measurement

In this section, we develop a framework for measuring cost-of-living changes under general non-homothetic preferences. We show that these changes can be decomposed into three components—initial-basket exposure, substitution, and a non-homothetic adjustment—which map into household-specific price indexes. All indexes are computed at the household level; for notational simplicity, we suppress the household index throughout this section.

3.1 Cost-of-living indexes and welfare

Let $t = 1, \ldots, T$ index time, $i = 1, \ldots, I$ index products, and let $\mathbf{p}_t = (p_{1t}, \ldots, p_{It})'$ denote the vector of prices in period t. The household's total expenditure in period t is x_t. In each period, the household chooses a consumption bundle $\mathbf{q}_t = (q_{1t}, \ldots, q_{It})'$ by solving

$$v(\mathbf{p}_t, x_t) = \max_{\mathbf{q}_t} \; U(\mathbf{q}_t) \quad \text{s.t. } \mathbf{p}_t'\mathbf{q}_t = x_t,$$

where $U(\cdot)$ is increasing and strictly quasi-concave. The indirect utility function is $v(\mathbf{p}, x)$, and the corresponding expenditure function $e(\mathbf{p}, u)$ is defined by

$$e(\mathbf{p}, u) = \min_{\mathbf{q}} \ \mathbf{p}'\mathbf{q} \quad \text{s.t. } U(\mathbf{q}) \geq u,$$

where $e(\mathbf{p}, u)$ is homogeneous of degree one, increasing and concave in $\mathbf{p}$, and strictly increasing in u.

The cost-of-living index between periods 1 and T measures the change in the cost of maintaining a fixed living standard (Konüs 1939):

$$P(\mathbf{p}_1, \mathbf{p}_T; u) = \frac{e(\mathbf{p}_T, u)}{e(\mathbf{p}_1, u)}.$$

In general, $P(\mathbf{p}_1, \mathbf{p}_T; u)$ depends on the reference living standard u.

3.2 A decomposition of cost-of-living changes

We focus on the cost-of-living index between the initial period ($t = 1$) and the final period ($t = T$), evaluated at the final utility level. This index measures the proportional change in expenditure required to attain realised final-period utility at final-period rather than initial-period prices:

$$\log P(\mathbf{p}_1, \mathbf{p}_T; u_T) = \log e(\mathbf{p}_T, u_T) - \log e(\mathbf{p}_1, u_T).$$

We decompose this cost-of-living index into an initial-basket exposure term, a substitution adjustment evaluated at the initial utility level, and a non-homothetic adjustment that captures how movements along Engel curves change the cost-of-living index when the index is evaluated at u_T rather than u_1.

To derive the decomposition, define the period-1-price-denominated money-metric indirect utility function

$$\mu(\mathbf{p}, x; \mathbf{p}_1) \equiv e(\mathbf{p}_1, v(\mathbf{p}, x)).$$

This is the expenditure at initial prices $\mathbf{p}_1$ required to attain the utility generated by expenditure x at prices $\mathbf{p}$. The following proposition uses this money-metric representation to obtain a second-order approximation to the final-utility cost-of-living index.

Proposition 1 (Decomposition of the cost-of-living index)**.** *Let $t = 1, \ldots, T$ index time and $i = 1, \ldots, I$ goods. Let $\mathbf{p}_t = (p_{1t}, \ldots, p_{It})'$ denote the price vector in period t, x_t total expenditure, and $s_{i1} = p_{i1}q_{i1}/x_1$ the initial-period spending share on good i. Define*

$$g_i \equiv \log p_{iT} - \log p_{i1}, \qquad g_x \equiv \log x_T - \log x_1,$$

and define the initial-period-share-weighted covariance

$$\mathrm{Cov}_s(a_i, b_i) = \sum_i s_{i1} a_i b_i - \Big(\sum_i s_{i1} a_i\Big)\Big(\sum_i s_{i1} b_i\Big).$$

Let $\eta_i = \left(\partial \log q_i(\mathbf{p}, x)/\partial \log x\right)\big|_{(\mathbf{p}_1, x_1)}$ *denote the expenditure elasticity of good* i*, and let* $\partial s_i/\partial \log p_j\big|_{(\mathbf{p}_1, u_1)}$ *denote the compensated (Hicksian) response of budget shares to prices, where* $u_1 = v(\mathbf{p}_1, x_1)$*.*

Under standard regularity conditions ensuring that the relevant derivatives exist, a second-order Taylor expansion of the money-metric function $\log \mu(\mathbf{p}, x; \mathbf{p}_1)$ *in log prices and log expenditure around* $(\log \mathbf{p}_1, \log x_1)$*, combined with the identity*

$$\log P(\mathbf{p}_1, \mathbf{p}_T; u_T) = \log x_T - \log \mu(\mathbf{p}_T, x_T; \mathbf{p}_1),$$

implies

$$\log P(\mathbf{p}_1, \mathbf{p}_T; u_T) \;\approx\; \text{Exposure} \;-\; \text{Substitution} \;-\; \text{Non-homothetic adjustment}, \tag{3.1}$$

where

$$\text{Exposure} = \log\Big(\sum_i s_{i1} \exp(g_i)\Big),$$

$$\text{Substitution} = \underbrace{\left(\log\Big(\sum_i s_{i1}\exp(g_i)\Big) - \sum_i s_{i1} g_i\right)}_{\text{Fixed-budget-share adjustment}} + \underbrace{\left(-\frac{1}{2}\sum_{i,j} \frac{\partial s_i}{\partial \log p_j}\bigg|_{(\mathbf{p}_1, u_1)} g_i\, g_j\right)}_{\text{Deviation from fixed-budget-share benchmark}},$$

$$\text{Non-homothetic adjustment} = -\underbrace{\left(g_x - \sum_j s_{j1} g_j\right)}_{\text{Real expenditure term}} \underbrace{\mathrm{Cov}_s(g_i, \eta_i)}_{\substack{\text{Inflation–expenditure}\\ \text{elasticity covariance}}}.$$

The approximation neglects terms of order higher than two in $(g_1, \ldots, g_I, g_x)$*.*

Proof. See Appendix B.1. □

Proposition 1 provides a second-order decomposition of the cost-of-living change into exposure, substitution, and a non-homothetic adjustment. The substitution term captures the total adjustment from the fixed-basket exposure benchmark to the initial-utility cost-of-living index. We decompose it around a fixed-budget-share benchmark, separating the adjustment implied by fixed budget shares from the deviation from this benchmark. The non-homothetic adjustment is governed by the interaction between real-expenditure growth and the covariance between relative price growth and expenditure elasticities. Combined with the mapping in Section 3.3, this yields a one-to-one link between these theoretically defined components and differences between household-level price indexes. The empirical implementation is therefore index-number-based and can be carried out for every household in scanner data without committing to a particular parametric preference structure or demand system.

Exposure. The exposure term is the log cost of the initial consumption bundle at final-period prices relative to initial-period prices. It coincides with the log cost-of-living index under Leontief preferences, where the consumption bundle is fixed in initial-period proportions. Equivalently, it is the log household-specific Laspeyres index. Differences in this term across households capture differential *exposure* to inflation arising from differences in initial consumption baskets, before allowing for substitution or non-homothetic adjustments.

Substitution. The substitution term is the total adjustment from the fixed-basket exposure term to the initial-utility cost-of-living index. Equivalently, it captures the difference between the Laspeyres benchmark, which holds the initial bundle fixed, and the cost of attaining the initial utility level when households can substitute in response to relative price changes. We decompose this total adjustment around a fixed-budget-share benchmark.

The first component is the log difference between the Laspeyres and geometric-Laspeyres indexes. The geometric-Laspeyres index corresponds to a first-order log-price approximation to the initial-utility cost-of-living index. Equivalently, it is the exact cost-of-living index for Cobb–Douglas preferences whose expenditure shares equal the period-1 shares. The Laspeyres–geometric-Laspeyres gap therefore measures the reduction in the cost-of-living index implied by replacing the fixed-basket benchmark with the fixed-budget-share geometric-Laspeyres benchmark.

The second component is the deviation from the fixed-budget-share benchmark. It captures curvature of the expenditure function in log prices. It is proportional to the Hicksian responses of budget shares to relative price changes, $\partial s_i / \partial \log p_j \big|_{(\mathbf{p}_1, u_1)}$, and therefore vanishes under Cobb–Douglas preferences, for which $\log e(\mathbf{p}, u)$ is linear in $(\log p_i)$. More generally, this term reflects whether households substitute away from products with rising relative prices by more or less than implied by the fixed-budget-share benchmark.[9] Its sign is a priori ambiguous. If households substitute away from goods with rising relative prices more strongly than is implied by Cobb–Douglas preferences, the deviation term lowers the

[9] Using the Hicksian elasticity matrix ϵ^H_{ij}, the deviation term can be written as: $-\frac{1}{2}\,\mathrm{Cov}_s\Big(g_i,\ \sum_j \big(\epsilon^H_{ij} + \mathbb{1}_{i=j} - s_{j1}\big) g_j\Big)$; that is, as a share-weighted covariance between relative price changes and an elasticity-weighted index of relative price changes.

initial-utility cost-of-living index relative to the fixed-budget-share benchmark; if they do so less strongly, it raises it.

Non-homothetic adjustment. The non-homothetic adjustment captures the difference between the cost-of-living index evaluated at the initial utility level, u_1, and the same index evaluated at the final utility level, u_T. In the second-order expansion, this difference depends on the product of two objects.

The first is the real-expenditure term $g_x - \sum_j s_{j1} g_j$, which measures nominal expenditure growth deflated by the geometric-Laspeyres price index. When this term is negative, purchasing power falls; under non-homothetic preferences, this moves households along Engel curves toward lower–η_i goods, i.e., away from luxuries and toward necessities. When it is positive, purchasing power rises and households shift spending toward higher–η_i goods.

The second is the inflation–expenditure elasticity covariance, $\mathrm{Cov}_s(g_i, \eta_i)$, which summarises how expenditure elasticities line up with relative price growth. If higher-η_i goods experience systematically higher price growth, this covariance is positive; if lower-η_i goods experience systematically higher price growth, it is negative. Under homothetic preferences, $\eta_i = 1$ for all i, the covariance is zero, and the non-homothetic adjustment vanishes.

For example, suppose real expenditure falls and necessities become relatively more expensive. Households then move along Engel curves, shifting spending toward goods whose relative prices are rising. The same price changes are therefore more costly at the household's realised final living standard than at its initial living standard. Equivalently, the final-utility cost-of-living index exceeds the initial-utility cost-of-living index: the non-homothetic adjustment offsets part of the substitution adjustment.

3.3 Mapping to observable price indexes

The decomposition in Proposition 1 maps directly into differences between observable household-level price indexes. This subsection introduces the indexes we use and shows how they correspond to the components of the decomposition.

Laspeyres and geometric-Laspeyres indexes. The household-specific Laspeyres index,

$$\Pi^L_{1,T} = \sum_i s_{i1} \exp(g_i), \tag{3.2}$$

measures how much the cost of the initial-period consumption bundle changes between periods 1 and T.[10] Its log is the exposure term in Proposition 1.

[10] Let $\mathbf{q}_1 = (q_{11}, \ldots, q_{I1})'$ denote the initial consumption bundle. The Laspeyres index is $\Pi^L_{1,T} = \frac{\sum_i p_{iT} q_{i1}}{\sum_i p_{i1} q_{i1}} = \sum_i s_{i1} \exp(g_i)$, where $s_{i1} \equiv \frac{p_{i1} q_{i1}}{\sum_j p_{j1} q_{j1}}$ and $g_i \equiv \log p_{iT} - \log p_{i1}$.

The geometric-Laspeyres index,

$$\Pi_{1,T}^{GL} = \exp\left(\sum_i s_{i1} g_i\right), \tag{3.3}$$

is the first-order log-price approximation to the cost-of-living index and is exact for Cobb–Douglas preferences whose expenditure shares equal the period-1 shares. The difference $\log \Pi_{1,T}^{L} - \log \Pi_{1,T}^{GL}$ is the fixed-budget-share adjustment: it is the adjustment from the fixed-basket Laspeyres benchmark to the fixed-budget-share geometric-Laspeyres benchmark.

Törnqvist index. In continuous time and with homothetic preferences, the log cost-of-living index corresponds to a Divisia index,

$$\log \Pi_{1,T}^{D} = \int_1^T \sum_i s_i(\tau)\, d \log p_i(\tau), \tag{3.4}$$

where $s_i(\tau)$ are expenditure shares along the path $\tau \in [1, T]$ (Divisia 1926). With discrete data, superlative indexes provide a second-order approximation to the cost-of-living index over each interval, and chaining them across time approximates the Divisia integral (Diewert 1976). We use the superlative Törnqvist index

$$\log \Pi_{1,T}^{T} = \sum_{t=1}^{T-1} \sum_i \frac{s_{it} + s_{it+1}}{2} \log \frac{p_{it+1}}{p_{it}}, \tag{3.5}$$

where $s_{it} = p_{it} q_{it} / x_t$ are observed period-t Marshallian shares. Under homothetic preferences, this index allows for substitution between goods and provides a second-order approximation to the cost-of-living index. With non-homothetic preferences, however, changes in observed Marshallian shares reflect both substitution responses to relative price changes and movements along Engel curves. As a result, the Törnqvist index does not isolate substitution at a fixed reference utility level.[11]

Non-homothetic preferences and the Jaravel–Lashkari index. If preferences are non-homothetic, then in continuous time the cost-of-living index can still be represented by a Divisia-type expression, but with Hicksian shares $\omega_i(\tau; u)$ evaluated at reference utility u in place of observed shares $s_i(\tau)$. With discrete data, a superlative approximation would likewise replace the observed shares in equation (3.5) with Hicksian shares $\omega_{it}(u)$. Because $\omega_{it}(u)$ generally differs from observed shares except in the base period, the Törnqvist index no

[11] Diewert (1976) shows that the Törnqvist index computed between two periods t and $t + 1$ is a second-order approximation to the cost-of-living index evaluated at the geometric mean of utilities across the two periods. This result is not directly useful when constructing a cost-of-living index at a pre-specified reference utility level, such as the initial or final utility level, or for chaining the index across multiple periods.

longer provides a second-order approximation to the cost-of-living index at a fixed reference utility level when preferences are non-homothetic.

We therefore implement the approach of Jaravel and Lashkari (2024), which provides a second-order approximation to a non-homothetic cost-of-living index. This entails a recursive algorithm that recovers a sequence of money-metric utility levels from observed $(\mathbf{s}_t, \mathbf{p}_t, x_t)$. The algorithm can then be used to construct household-level cost-of-living indexes evaluated at specified reference utility levels. The approach allows for flexible non-homothetic preferences, but assumes that, conditional on observable household controls, households share a common function linking money-metric utility to the non-homothetic correction. Under this assumption, cross-sectional variation in money-metric utility is informative about how household price indexes change as households move along Engel curves. This relationship provides the non-homothetic correction needed to evaluate the cost-of-living index at a fixed reference utility level. We use the resulting household-level index, denoted $\Pi^{JL}_{1,T}(u)$, as a second-order approximation to the cost-of-living index $P(\mathbf{p}_1, \mathbf{p}_T; u)$ evaluated at utility level u. In particular, we use $\Pi^{JL}_{1,T}(u_1)$ and $\Pi^{JL}_{1,T}(u_T)$ as approximations to $P(\mathbf{p}_1, \mathbf{p}_T; u_1)$ and $P(\mathbf{p}_1, \mathbf{p}_T; u_T)$, respectively. See Appendix B.2 for full details.

Mapping the decomposition to observables. Proposition 1 provides a second-order decomposition of the cost-of-living change $\log P(\mathbf{p}_1, \mathbf{p}_T; u_T) = \log e(\mathbf{p}_T, u_T) - \log e(\mathbf{p}_1, u_T)$. It is useful to rewrite this as:

$$\log P(\mathbf{p}_1, \mathbf{p}_T; u_T) = \log P(\mathbf{p}_1, \mathbf{p}_T; u_1) - \left[\log P(\mathbf{p}_1, \mathbf{p}_T; u_1) - \log P(\mathbf{p}_1, \mathbf{p}_T; u_T)\right].$$

The first term on the right-hand side is the log cost-of-living index evaluated at the initial utility level. Proposition 1 approximates this term by the exposure term minus the substitution adjustment. The bracketed term captures how the cost-of-living index changes when the reference utility level moves from u_1 to u_T; this is the non-homothetic adjustment. Under homothetic preferences, $\eta_i = 1$ for all i, this adjustment vanishes.

Using the price indexes defined above, and noting that $\log \Pi^{JL}_{1,T}(u)$ is a second-order approximation to $\log P(\mathbf{p}_1, \mathbf{p}_T; u)$, we implement the components in Proposition 1 as:

$$\begin{aligned}
\text{Exposure} &= \log \Pi^{L}_{1,T}, \\
\text{Substitution} &\approx \log \Pi^{L}_{1,T} - \log \Pi^{JL}_{1,T}(u_1) \\
&= \underbrace{\log \Pi^{L}_{1,T} - \log \Pi^{GL}_{1,T}}_{\substack{\text{Fixed-budget-share} \\ \text{adjustment}}} + \underbrace{\log \Pi^{GL}_{1,T} - \log \Pi^{JL}_{1,T}(u_1)}_{\substack{\text{Deviation from} \\ \text{fixed-budget-share benchmark}}}, \\
\text{Non-homothetic adjustment} &\approx \log \Pi^{JL}_{1,T}(u_1) - \log \Pi^{JL}_{1,T}(u_T).
\end{aligned}$$

The exposure and fixed-budget-share adjustment are exact functions of observed prices and expenditure shares. The deviation from the fixed-budget-share benchmark and the non-homothetic adjustment use $\Pi^{JL}_{1,T}(\cdot)$, a second-order approximation to the corresponding cost-of-living indexes evaluated at a fixed reference utility level. Substituting these expressions into equation (3.1) yields an index-number-based implementation of each component, up to terms of order higher than two in $(g_1, \ldots, g_I, g_x)$.

Discussion. We formulate Proposition 1 in terms of the final-utility cost-of-living index, but this object has an immediate money-metric welfare interpretation. Since $x_T = e(\mathbf{p}_T, u_T)$ and $x_1 = e(\mathbf{p}_1, u_1)$, the proportional equivalent-variation change in living standards, measured at initial prices, is

$$\log e(\mathbf{p}_1, u_T) - \log e(\mathbf{p}_1, u_1) = \log \frac{x_T}{x_1} - \log P(\mathbf{p}_1, \mathbf{p}_T; u_T).$$

Thus, given nominal expenditure growth, measuring $\log P(\mathbf{p}_1, \mathbf{p}_T; u_T)$ is equivalent to measuring the money-metric welfare change at initial prices. Analogously, $\log P(\mathbf{p}_1, \mathbf{p}_T; u_1)$ pins down compensating variation, the money-metric welfare change measured at final prices.

This welfare interpretation connects our decomposition to the second-order welfare expansions in Baqaee and Burstein (2023), which show that money-metric welfare can differ from chain-weighted real consumption through expenditure-switching terms driven by substitution, income effects, and taste shocks. Our decomposition applies similar second-order logic to a two-period cost-of-living index and makes explicit how the components can be implemented as differences between price indexes. As in Auer et al. (2024)'s analysis of unequal expenditure switching in response to foreign-price changes, the decomposition contains terms reflecting inflation exposure and substitution. Because we use an equivalent-variation welfare measure, our decomposition also includes a non-homothetic adjustment, which captures how the cost-of-living index changes when it is evaluated at the final rather than the initial utility level.

Our use of the final-utility cost-of-living index is tied to the equivalent-variation welfare question: how did the episode change households' living standards relative to their pre-episode situation, measured in the common price environment of 2021Q3? Under homothetic preferences, proportional equivalent- and compensating-variation welfare changes coincide. Under non-homothetic preferences, they differ because the cost of a given price change depends on the reference utility level. In a cheapflation episode with falling real expenditure, this difference is economically meaningful: movements along Engel curves can push households toward goods whose relative prices are rising, raising the final-utility cost-of-living index relative to the initial-utility index.

The contribution of our approach is twofold. First, by combining this decomposition with the non-homothetic money-metric index of Jaravel and Lashkari (2024), we obtain an *index-number-based* implementation of the welfare components under general non-homothetic preferences. Each term can be measured as a log difference between household-level price indexes, without imposing a parametric demand system. Under stable preferences and utility maximisation, the implementation uses observed prices, expenditure shares, and total expenditure.[12]

Second, we apply and interpret the non-homothetic adjustment in an environment with within-category quality ladders and cheapflation. In such an environment, movements along Engel curves can push households toward goods whose relative prices are rising, thereby amplifying inflation-driven welfare losses. The framework decomposes the welfare impact of the *observed* price and expenditure changes; it does not, on its own, deliver arbitrary policy or price counterfactuals. However, the same household-level price-index objects can be used to study how realised basket reallocation changes exposure to a repeat price shock (see Section 5).

4 Inflation inequality and welfare effects

In this section, we present our empirical results, documenting how exposure to inflation, substitution responses to relative price changes, and non-homothetic demand responses shape the welfare effects of the inflationary surge.

4.1 Inflation exposure

We measure exposure to inflation using household-specific Laspeyres indexes, following equation (3.2). These indexes capture how the cost of each household's initial-period consumption bundle changes over time.

Index construction. Let $T = $ 2023Q3 and let Ω denote the set of products observed in both periods 1 and T (2021Q3 and 2023Q3). For each household h, let x_{hi1} denote period-1 expenditure on product i, and define the corresponding expenditure share as $s_{hi1} \equiv \frac{x_{hi1}}{\sum_{j\in\Omega} x_{hj1}}$. Let p_{it} denote the price of product i in period $t \in \{1, T\}$ (see equation (2.1)). The household's Laspeyres index between 2021Q3 and 2023Q3 is then $\Pi^L_{h(1,T)} \equiv \sum_{i\in\Omega} s_{hi1} \frac{p_{iT}}{p_{i1}}$. Products that household h does not buy in 2021Q3 are assigned $s_{hi1} = 0$, so they do not enter that

[12]An alternative is to fully specify a non-homothetic demand system. This requires stronger assumptions, but allows researchers to compute counterfactuals and endogenise firms' pricing behaviour. For example, Becker (2025) uses a model with non-homothetic CES preferences and oligopolistic pricing to show how quality downgrading can endogenously generate cheapflation through countercyclical markups on low-quality goods.

household's index. By focusing on products available in both 2021Q3 and 2023Q3, this baseline Laspeyres index abstracts from product entry and exit.

Household-level exposure. Figure 4.1 summarises household-level cumulative inflation exposure across the nine quarters from 2021Q3 to 2023Q3, alongside earlier nine-quarter periods. Panel (a) shows that *average* inflation exposure over 2021Q3–2023Q3 was historically high, with a cumulative increase of 26.2%. By comparison, average cumulative inflation exposure in earlier nine-quarter periods ranged from -3.6% (2014Q1–2016Q1) to 6.1% (2012Q1–2014Q1). Panel (b) illustrates the *distribution* of household-level exposure across these periods: the elevated average inflation exposure in 2021Q3–2023Q3 was accompanied by much greater dispersion across households.

Figure 4.1: *Household-level inflation exposure*

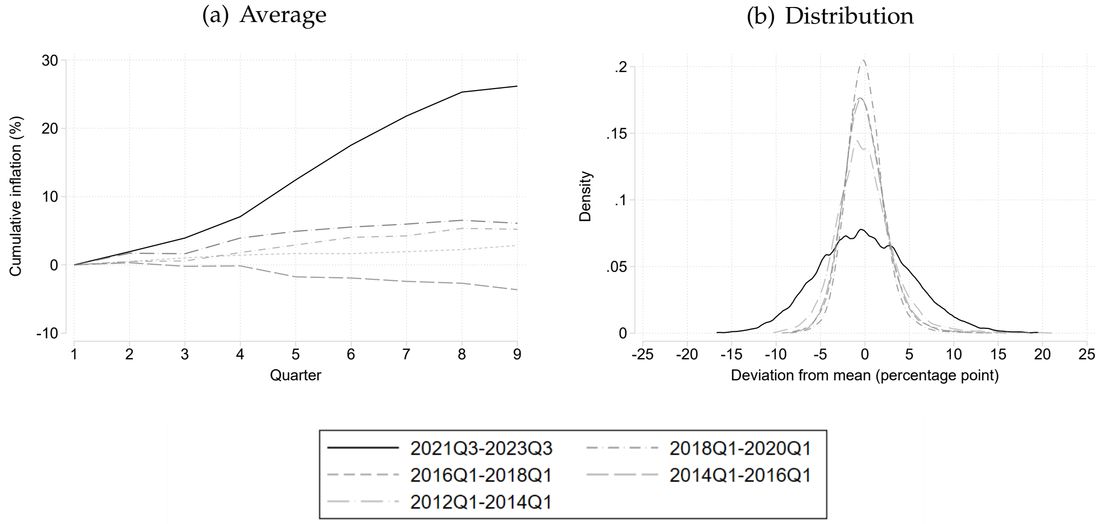

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2012–2023). Panel (a) presents mean cumulative inflation exposure for each nine-quarter period. Panel (b) displays kernel density estimates of the distribution of household-level cumulative inflation exposure over each period. Cumulative inflation is measured using a household-specific Laspeyres index.

Inequality. Figure 4.2 summarises the relationship between inflation exposure and equivalised fast-moving consumer goods expenditure (panel (a)) and equivalised income (panel (b)). Both panels reveal historically large gradients in 2021Q3–2023Q3. Households in the bottom quartile of the equivalised 2021 expenditure distribution experienced cumulative inflation exposure approximately 5.5 percentage points higher than those in the top quartile, while the difference between bottom and top deciles was 7.5 percentage points. Households in the bottom equivalised income quartile had exposure around 3.0 percentage points higher than households in the top quartile, while the corresponding decile difference was 4.4 percentage points. By contrast, none of the nine-quarter periods between 2012 and 2020 exhibits an economically meaningful gradient. This absence of a systematic gradient in low-inflation years is consistent with survey-based evidence for the UK over 1976–2014, which finds only modest long-run inflation differences across income groups (Crawford and Oldfield 2002; Leicester et al. 2008; Adams and Levell 2014), but stands in contrast to the secular inflation inequality documented for the US over 2004–2015 by Jaravel (2019).

Across both expenditure and income distributions, the pattern of inflation inequality is stark: the initial consumption baskets of lower-resource households lead to systematically higher inflation exposure. Quantitatively, the gradient is steeper across the expenditure distribution than across the income distribution. This likely reflects two factors. First, income is reported in £10,000 bands and is top-coded at £70,000, whereas expenditure is observed on a continuous scale. This banding and top-coding introduce noise into the ranking of households by equivalised income and tend to attenuate differences in average inflation exposure across income groups. Second, total expenditure and current income are conceptually distinct measures of household resources. Expenditure decisions are typically shaped by longer-run or "permanent" income and therefore reflect broader economic well-being (e.g., Poterba 1989; Meyer and Sullivan 2023), whereas current income can be influenced by more transitory short-run factors.

Hierarchical decomposition. The inflation inequality documented in Figure 4.2 could in principle be driven by heterogeneity in consumption bundles across broad segments, across product categories within segments, across products and quality rungs within categories, or by some combination of these margins. To isolate the contribution of each margin, we exploit the fact that the Laspeyres index can be re-expressed hierarchically.

Let Ω^c denote the set of products in category c, and let Γ^g denote the set of categories in segment g. Define the household-specific within-category product share, the within-segment category share, and segment share as

$$s^c_{hit} = \frac{x_{hit}}{\sum_{i' \in \Omega^c} x_{hi't}}, \quad s^g_{hct} = \frac{\sum_{i \in \Omega^c} x_{hit}}{\sum_{c' \in \Gamma^g} \sum_{i' \in \Omega^{c'}} x_{hi't}}, \quad s_{hgt} = \frac{\sum_{c \in \Gamma^g} \sum_{i \in \Omega^c} x_{hit}}{\sum_{g'} \sum_{c' \in \Gamma^{g'}} \sum_{i \in \Omega^{c'}} x_{hi't}}.$$

The hierarchical Laspeyres index can then be written as

$$\Pi^{HL}_{h(1,T)} = \sum_g s_{hg1} \mathbb{P}^g_{h(1,T)}, \quad \text{where}$$
$$\mathbb{P}^g_{h(1,T)} = \sum_{c \in \Gamma^g} s^g_{hc1} \mathbb{P}^c_{h(1,T)} \quad \text{and} \quad \mathbb{P}^c_{h(1,T)} = \sum_{i \in \Omega^c} s^c_{hi1} \left(\frac{p_{iT}}{p_{i1}} \right).$$

Since $s_{hi1} = s^c_{hi1} \, s^g_{hc1} \, s_{hg1}$, it follows that $\Pi^{HL}_{h(1,T)} = \Pi^L_{h(1,T)}$.

Figure 4.2: *Inflation inequality*

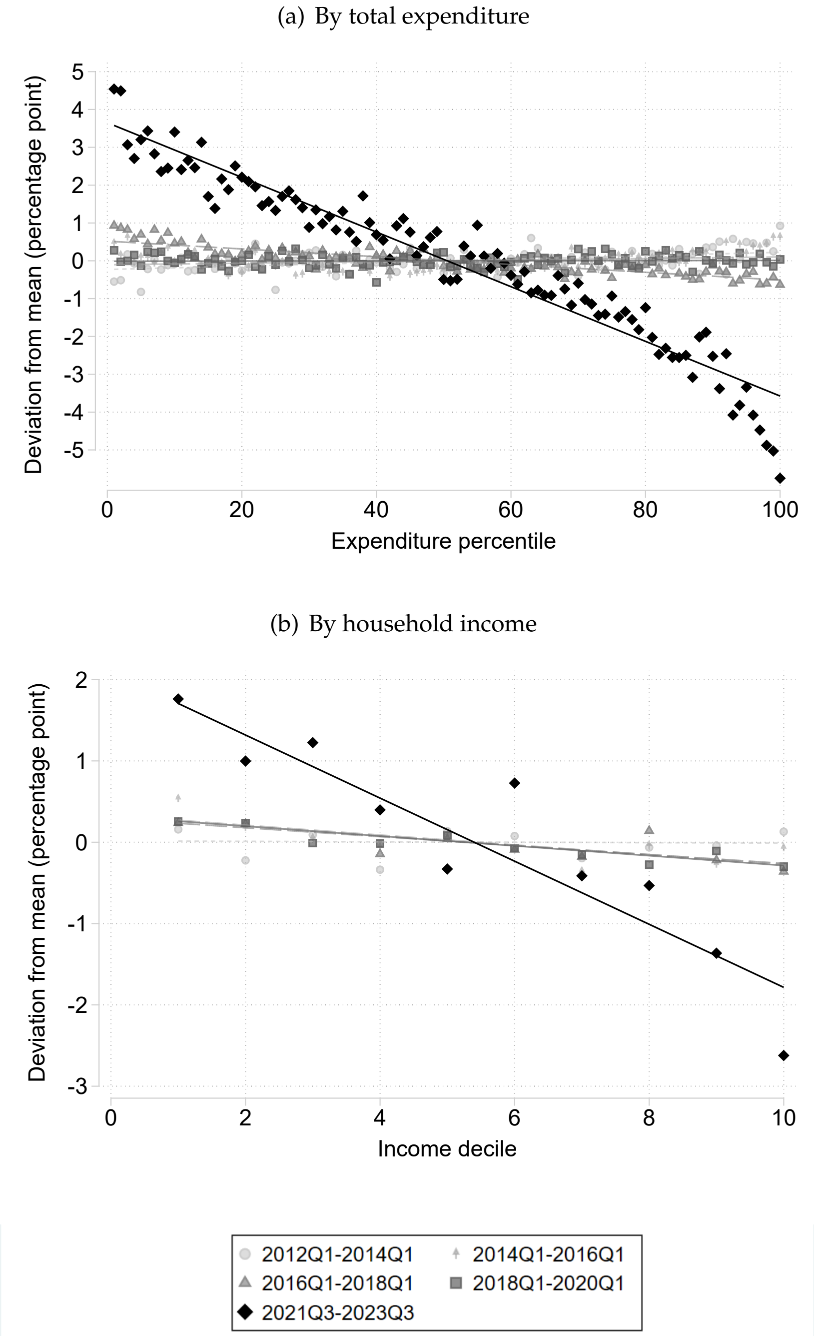

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2012–2023). The figure plots the relationship between cumulative inflation exposure over each nine-quarter period and a household's percentile in the equivalised expenditure distribution (panel (a)) or decile in the equivalised income distribution (panel (b)). Each marker corresponds to one expenditure percentile or income decile; lines show fitted linear relationships. Households are assigned to expenditure percentiles and income deciles based on expenditure or income in the calendar year in which the relevant nine-quarter period begins. Cumulative inflation exposure is measured using a household-specific Laspeyres index.

We use this representation to decompose household-specific exposure to inflation by sequentially replacing household-specific spending shares with population-average shares at different levels of the hierarchy. Starting from a benchmark that preserves only segment-level heterogeneity, we then allow households to differ in category shares within segments, and finally in product shares within categories. Comparing the resulting gradients shows how much of inflation inequality is attributable to broad segment-level differences, category-level differences within segments, and product-level differences within categories.

In Figure 4.3, we plot the distribution of household-level cumulative inflation exposure over 2021Q3–2023Q3 (panel (a)) and the across-expenditure gradient (panel (b)) using black lines and markers, replicating information from Figures 4.1 and 4.2. We also show patterns when switching off heterogeneity in spending across products within categories (red lines and markers) and, additionally, when switching off heterogeneity across categories within segments (blue lines and markers). Panel (a) shows that within-category product choice is central in driving dispersion: the standard deviation is about 2.4 percentage points when only segment shares vary across households, rising to 5.4 percentage points in the full household-specific index. Panel (b) shows that heterogeneity in spending across products within categories accounts for the majority of the inflation-exposure gradient. These findings emphasise the importance of using detailed household spending and price information to study inflation inequality.

Figure 4.3: *Hierarchical decomposition of inflation exposure*

(a) Distribution

(b) Inequality

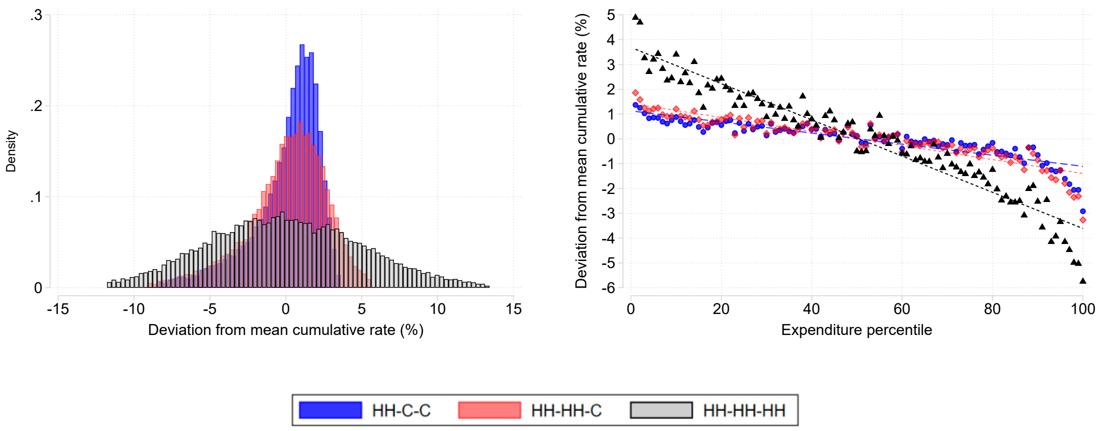

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2021–2023). Panel (a) shows a histogram of the distribution of cumulative inflation exposure from 2021Q3 to 2023Q3. Panel (b) plots the relationship between cumulative inflation exposure and households' equivalised expenditure percentiles, with a marker for each percentile and a fitted linear relationship. HH-C-C uses household-specific segment shares and common within-segment category shares and within-category product shares. HH-HH-C uses household-specific segment and category shares, and common within-category product shares. HH-HH-HH uses household-specific shares at all levels. Households are allocated to expenditure percentiles based on their equivalised spending in 2021. Cumulative inflation exposure is measured using a Laspeyres index.

The descriptive evidence in Figure 2.1 shows that price growth is stronger lower down the quality ladder, and that the consumption baskets of lower-resource households make them more exposed to these price dynamics. Figures 4.2 and 4.3 show that these patterns translate into substantial inequality in inflation exposure, driven primarily by the interaction of within-category differences in cumulative price growth and heterogeneity in product choice across households. In the welfare decomposition below, this alignment between relative price growth and expenditure elasticities matters through two channels. First, because lower-resource households initially spend more on lower-rung, lower-expenditure-elasticity products, cheapflation raises their exposure to inflation. Second, as real expenditure falls, the same alignment enters the non-homothetic adjustment: movements along Engel curves shift households toward products with lower expenditure elasticities, which are also the products experiencing stronger relative price increases.

Replacing the common prices in the Laspeyres index with expenditure-quartile-specific prices has minimal impact on the inflation inequality gradient (see Appendix C.1). This aligns with our evidence that lower-resource households pay only modestly lower prices for identical products, and that this pattern remains relatively stable over time (see Appendix A.1). The baseline Laspeyres index also abstracts from product entry and exit. Alternative treatments of product churn—including chaining the Laspeyres index across adjacent periods and using the Feenstra-corrected CES index, which incorporates variety effects under CES preferences, following Feenstra (1994) and Broda and Weinstein (2010)—have no discernible impact on the pattern of inflation inequality (see Appendix C.2).

Broader implications. Fast-moving consumer goods (FMCG) account for only part of the overall consumption basket, so it is useful to gauge how our findings translate into differences in *overall* inflation. Using the Living Costs and Food Survey (LCFS)—the UK's official household expenditure survey—we find that the COICOP categories corresponding to our FMCG coverage represent, on average, 19.3% of total household expenditure (and 33.8% of total spending on goods) in the 2021 LCFS, with budget shares declining from roughly 23.1% in the bottom equivalised-income decile to 16.6% in the top decile. Consistent with Engel's law, food and basic consumer goods absorb a larger fraction of the budgets of lower-income households. Combined with the stronger cheapflation we document for FMCG, these expenditure patterns amplify the contribution of within-FMCG inflation inequality to the bottom–top gap in overall inflation: lower-income households have baskets that place greater weight on precisely the categories in which inflation is most regressive.[13]

To connect our scanner-data-based measures to the full CPI basket, we combine LCFS expenditure shares with CPI category-level price indexes in a simple back-of-the-envelope calculation (see Appendix C.3). We construct income-decile-specific price indexes under two scenarios: one in which FMCG categories have common inflation across all groups, and

[13] Fast-moving consumer goods are also highly salient and play a disproportionate role in shaping households' inflation perceptions and expectations; see, for example, D'Acunto et al. (2021).

one in which they inherit the decile-specific FMCG Laspeyres inflation we estimate from scanner data. Under the realised pattern of non-FMCG CPI inflation over 2021Q3–2023Q3, allowing FMCG inflation to differ across income groups raises the bottom–top decile gap in CPI inflation by about 0.9 percentage points, a 28% increase. Although FMCG categories account for only about one-fifth of total expenditure, cheapflation in these goods nonetheless makes a quantitatively meaningful contribution to overall inflation inequality.

Finally, two additional pieces of evidence help place the within-FMCG cheapflation gradient in the context of the broader consumption basket. First, the calculation combining LCFS shares with CPI price indexes shows that, even when inflation is common across households within CPI categories, 2021Q3–2023Q3 features substantial inequality in inflation exposure. This common-price gradient is driven largely by the exceptionally large increase in residential energy prices combined with higher energy expenditure shares among lower-resource households (see Levell et al. 2025); when residential energy is excluded from the CPI basket, the gradient becomes much smaller. Thus, absent the unusual energy-price shock, within-category cheapflation would account for an even larger share of overall inflation inequality. Second, using CPI microdata, we find evidence that within-category price distributions compress in several non-FMCG goods categories as well (see Appendix A.3), indicating that cheapflation is not confined to the segment we study in detail. Our scanner-data-based estimates of FMCG inflation inequality may therefore understate the contribution of cheapflation across the broader consumption basket.

Discussion. Taken together, these results show that the 2021Q3–2023Q3 UK inflation episode was unusual not only in its aggregate magnitude, but also in the dispersion and regressivity of household-specific inflation exposure. Earlier work documents substantial dispersion in household-level price indexes (Kaplan and Schulhofer-Wohl 2017), a widening of the inflation gap between high- and low-income households in the US during the Great Recession (Argente and Lee 2021), and attenuation in measured inflation inequality when using data aggregated across products (see Jaravel 2021).

A key empirical finding of this section is that, in a period of sharp and uneven relative price changes within fast-moving consumer goods, inflation exposure is substantially higher for lower-resource households across both the income and expenditure distributions. Within-category differences in product choice along the quality ladder account for the bulk of this inequality. This complements recent work showing that cheapflation over 2021–2023 is a widespread international phenomenon (Cavallo and Kryvtsov 2024), that within-category cheapflation is linked to inflation inequality in US data (Sangani 2025), and that recent inflation heterogeneity across income groups also arises from expenditure differences across more aggregated COICOP categories (Jaravel 2024). The price-growth differentials we estimate between low and high rungs of the quality ladder in UK fast-moving consumer goods are of similar magnitude to the inflation gap between the first and fourth quality quartiles reported for the UK by Cavallo and Kryvtsov (2024). These findings motivate the welfare analysis that follows, where we use our decomposition to quantify how substitution and non-homothetic demand responses shape the distributional consequences of cheapflation episodes.

4.2 Substitution responses

We next quantify the substitution adjustment, holding the reference utility level fixed at its initial value. Proposition 1 shows that the final-utility cost-of-living index can be written as exposure minus substitution minus the non-homothetic adjustment. The exposure term is the no-substitution benchmark: it is the change in the cost of the initial-period consumption basket at final- compared to initial-period prices, captured by the household-specific Laspeyres index. The substitution term captures the reduction in the cost-of-living index that arises when households can reoptimise in response to relative price changes, while remaining on the initial indifference curve. Thus, exposure minus substitution approximates the initial-utility cost-of-living index: the proportional change in expenditure needed to attain 2021Q3 utility at 2023Q3 rather than 2021Q3 prices.

Index construction. To measure substitution responses, we compare the Laspeyres exposure index with household-specific price indexes that allow for basket adjustment. We construct three such indexes. First, the geometric-Laspeyres index uses the same 2021Q3 base shares (s_{hi1}) as the Laspeyres index: $\log \Pi^{GL}_{h(1,T)} \equiv \sum_{i\in\Omega} s_{hi1} \log \frac{p_{iT}}{p_{i1}}$. Second, we construct a household-level chained Törnqvist index. Let $t \in \{1, 2, 3\}$ correspond to 2021Q3, 2022Q3, and 2023Q3, and let $T = 3$. Let Ω_t denote the set of products available in both t and $t + 1$. The log chained Törnqvist index is $\log \Pi^{T}_{h(1,T)} \equiv \sum_{t=1}^{T-1} \sum_{i\in\Omega_t} \frac{1}{2} \left(s_{hit} + s_{hi,t+1}\right) \log \frac{p_{i,t+1}}{p_{it}}$ where s_{hit} is household h's expenditure share on product i in period t. Third, we construct non-homothetic cost-of-living indexes using the algorithm in Jaravel and Lashkari (2024), which adjusts the household-specific chained Törnqvist for non-homothetic demand responses. This procedure starts from household-specific chained price indexes and estimates the non-homothetic correction from how price indexes vary with money-metric utility, conditional on household composition. We compute non-homothetic cost-of-living indexes evaluated at both the initial and final utility levels. We obtain the final-utility index by running the algorithm forward from 2021Q3 to 2023Q3, and the initial-utility index by applying the analogous procedure in reverse. See Appendix B.2 for details.

Average. Column (1) of Table 4.1 summarises average cumulative cost-of-living changes over 2021Q3–2023Q3 under different restrictions on household preferences.[14]

Rows 1 and 2 provide the Leontief and Cobb–Douglas benchmarks. Under Leontief preferences, households do not substitute across products in response to relative price changes, so the cost-of-living index coincides with the Laspeyres index and rises by 26.18%. Under Cobb–Douglas preferences with expenditure shares equal to initial-period values, the geometric-Laspeyres index is the exact cost-of-living index. This fixed-budget-share adjustment reduces the average cost-of-living increase to 24.26%, 1.92 percentage points below the Leontief benchmark.

Rows 3 and 4 allow for more general substitution patterns. Relative to the geometric-Laspeyres benchmark, the deviation is approximated by $\log \Pi^{GL}_{h(1,T)} - \log \Pi^{T}_{h(1,T)}$ under homothetic preferences and by $\log \Pi^{GL}_{h(1,T)} - \log \Pi^{JL}_{h(1,T)}(u_1)$ under non-homothetic preferences. In the latter case, the average initial-utility cost-of-living increase is 24.54%, so substitution responses lower the average cost of living increase by 1.64 percentage points. The deviation from the fixed-budget-share benchmark therefore offsets 0.28 percentage points of the 1.92 percentage point Cobb–Douglas adjustment. This is consistent with households substituting away least strongly from the products experiencing the strongest price growth.

The Törnqvist index records an average cost-of-living increase that is 0.33 percentage points higher than the non-homothetic Jaravel–Lashkari index evaluated at the initial utility level (24.87% versus 24.54%). This difference reflects the non-homothetic correction in the Jaravel–Lashkari index. The Törnqvist index is a homothetic superlative index constructed from observed expenditure shares, which reflect both substitution responses to relative price changes and movements along Engel curves. When households respond to reduced purchasing power by shifting spending toward lower-expenditure-elasticity products that experienced relatively strong price growth, a homothetic cost-of-living interpretation attributes part of this Engel-curve movement to substitution. In this setting, the Törnqvist index is therefore subject to a non-homotheticity bias: it overstates the initial-utility cost-of-living increase relative to the non-homothetic index that holds utility fixed at its initial level.

Inequality. Columns (2)–(5) of Table 4.1 summarise differences in cost-of-living increases across the equivalised expenditure and income distributions under different restrictions on household preferences. We report bottom–top quartile and decile differences, so positive values indicate higher cost-of-living increases for lower-resource households.

[14] For interpretability, we report the cost-of-living indexes in Tables 4.1 and 4.2 as changes rather than log changes (which coincide up to a standard first-order approximation error). In Figure 4.4 below, when we present the results of the decomposition in Proposition 1, we report results in log points ($\times$ 100).

Under general non-homothetic preferences, the cost-of-living gradient is slightly attenuated relative to the inflation-exposure gradient. Across the expenditure distribution, the bottom–top quartile difference narrows by 0.65 percentage points, from 5.49 to 4.84, while the bottom–top decile difference narrows by 0.84 percentage points, from 7.50 to 6.66. Across the income distribution, the corresponding differences narrow by 0.54 percentage points, from 3.01 to 2.47, and by 0.69 percentage points, from 4.45 to 3.76. This pattern indicates that substitution responses somewhat offset the exposure gradient, consistent with stronger substitution among lower-resource households. Nonetheless, the cost-of-living gradient remains large and regressive: households with lower income and lower total expenditure continue to experience substantially higher cost-of-living increases than higher-resource households.

Table 4.1: *Cost-of-living changes over 2021Q3–2023Q3*

	Average change (%)	Bottom–top quartile difference (p.p.)		Bottom–top decile difference (p.p.)	
		Expenditure	Income	Expenditure	Income
	(1)	(2)	(3)	(4)	(5)
Leontief preferences ($\Pi^{L}_{1,T}$)	26.18	5.49	3.01	7.50	4.45
Cobb–Douglas preferences ($\Pi^{GL}_{1,T}$)	24.26	5.26	2.86	7.19	4.23
General homothetic preferences ($\Pi^{T}_{1,T}$)	24.87	4.58	2.44	6.29	3.75
General non-homothetic preferences ($\Pi^{JL}_{1,T}(u_1)$)	24.54	4.84	2.47	6.66	3.76

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2021–2023). Average change is the mean cumulative cost-of-living increase (in percent). Bottom–top differences are measured in percentage points; positive values indicate higher cost-of-living changes at the bottom of the distribution.

4.3 Non-homothetic adjustment

The final component in Proposition 1 is the non-homothetic adjustment. This adjustment captures how the cost-of-living index changes when it is evaluated at the realised final-period utility level rather than the initial-period utility level. In a cheapflation environment, this distinction matters because lower utility levels are associated with consumption bundles that place greater weight on lower-expenditure-elasticity goods. If these goods also experience stronger relative price increases, then the same pattern of price changes raises the cost of living more at lower utility levels than at higher utility levels. When a household's living standard falls between the initial and final periods, the cost-of-living index evaluated at realised final-period utility can therefore exceed the index evaluated at initial-period utility. This difference is the non-homothetic adjustment in Proposition 1.

Table 4.2 compares non-homothetic cost-of-living indexes evaluated at the realised initial- and final-period utility levels, using the Jaravel–Lashkari approach.

Average. On average, the cost of living rises by 24.54% when evaluated at realised initial-period utility u_1, but by 25.00% when evaluated at final-period utility u_T. This 0.46 percentage point gap implies that the 2021Q3–2023Q3 pattern of relative price changes raises the cost of living more at the lower living standards households reached by 2023Q3 than at their pre-episode living standards. This means that, in this episode, the non-homothetic adjustment raises the equivalent-variation loss from inflation.

Inequality. Columns (2)–(5) of Table 4.2 compare bottom–top quartile and decile differences in cost-of-living changes across the expenditure and income distributions when the Jaravel–Lashkari index is evaluated at initial- and final-period utility. Across the expenditure distribution, the bottom–top quartile difference narrows from 4.84 to 4.24 percentage points, while the bottom–top decile difference narrows from 6.66 to 5.86 percentage points. Evaluating the cost of living at u_T rather than u_1 therefore attenuates the expenditure gradient: the quartile gap narrows by 0.60 percentage points and the decile gap by 0.80 percentage points. By contrast, the gradient across the income distribution is essentially unchanged: the bottom–top quartile difference changes only from 2.47 to 2.48 percentage points, while the bottom–top decile difference changes from 3.76 to 3.79 percentage points.

Table 4.2: *Non-homothetic cost-of-living indexes over 2021Q3–2023Q3*

	Average change (%)	Bottom–top quartile difference (p.p.)		Bottom–top decile difference (p.p.)	
		Expenditure	Income	Expenditure	Income
	(1)	(2)	(3)	(4)	(5)
Initial-utility cost-of-living index ($\Pi^{JL}_{1,T}(u_1)$)	24.54	4.84	2.47	6.66	3.76
Final-utility cost-of-living index ($\Pi^{JL}_{1,T}(u_T)$)	25.00	4.24	2.48	5.86	3.79

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2021–2023). Average change is the mean cumulative cost-of-living increase (in percent). Bottom–top differences are measured in percentage points; positive values indicate higher cost-of-living growth at the bottom of the distribution. Row (1) repeats row (4) of Table 4.1.

4.4 Inflation-driven welfare changes

We now bring together the components of Proposition 1, applying the decomposition to the final-utility cost-of-living index, $\log P(\mathbf{p}_1, \mathbf{p}_T; u_T)$. As discussed in Section 3.3, this index has a direct equivalent-variation interpretation: given nominal expenditure growth, a higher final-utility cost-of-living index implies a larger inflation-driven loss in living standards, measured at initial-period prices.

Figure 4.4 shows how inflation exposure and behavioural adjustments shape the final-utility cost-of-living index over 2021Q3–2023Q3 across the equivalised expenditure and income distributions. Panels (a) and (b) plot, by decile, average log Laspeyres inflation exposure and the final-utility cost-of-living index, approximated by the final-utility Jaravel–Lashkari index, $\log \Pi^{JL}_{1,T}(u_T)$.

Relative to the exposure-only benchmark, behavioural adjustments modestly compress the inflation gradient. For households in the bottom expenditure decile, $\log \Pi^{JL}_{1,T}(u_T)$ is 1.72 log points below inflation exposure based on the initial basket, whereas for the top expenditure decile the corresponding difference is 0.47 log points. Across income deciles, the corresponding figures are 1.22 and 0.73 log points.

Panels (c) and (d) decompose the gap Exposure $- \log P(\mathbf{p}_1, \mathbf{p}_T; u_T)$ into substitution and the non-homothetic adjustment. Positive components reduce the final-utility cost-of-living index relative to the Laspeyres exposure benchmark; negative components raise it. The substitution component is split into the fixed-budget-share adjustment and the deviation from the fixed-budget-share benchmark. The fixed-budget-share adjustment lowers the cost-of-living index for all groups, with somewhat larger effects at the bottom of the expenditure and income distributions. The deviation from the fixed-budget-share benchmark partially offsets this effect across all deciles, although total substitution still reduces the cost-of-living index throughout the distribution.

The non-homothetic adjustment is negative for all deciles. This reflects the interaction of non-homothetic demand with the pattern of relative price changes: as purchasing power falls, households move along Engel curves toward lower-expenditure-elasticity products whose relative prices rose most over the period. These adjustments are optimal given prices and budgets, but because they shift the relevant reference bundle toward products with stronger relative price growth, they raise the final-utility cost-of-living index. In panels (c) and (d), they therefore enter as a negative contribution to Exposure $- \log \Pi^{JL}_{1,T}(u_T)$. Thus, the non-homothetic adjustment raises the average equivalent-variation loss from inflation, but its distributional effect is less regressive than initial-basket exposure; in particular, across expenditure deciles it partly offsets, without undoing, the regressive gradient in cost-of-living changes.[15]

[15] This treatment differs from homothetic CES exact price-index decompositions, where trading down in quality is absorbed into a "product substitution" term (see, e.g., Argente and Lee 2021). In our framework, movements down the quality ladder in a cheapflation environment are captured explicitly through the non-homothetic adjustment, which can raise the equivalent-variation loss from inflation even when households shift toward lower-priced products.

Figure 4.4: *Inflation-driven welfare changes*

Final-utility cost-of-living index and initial-basket exposure:

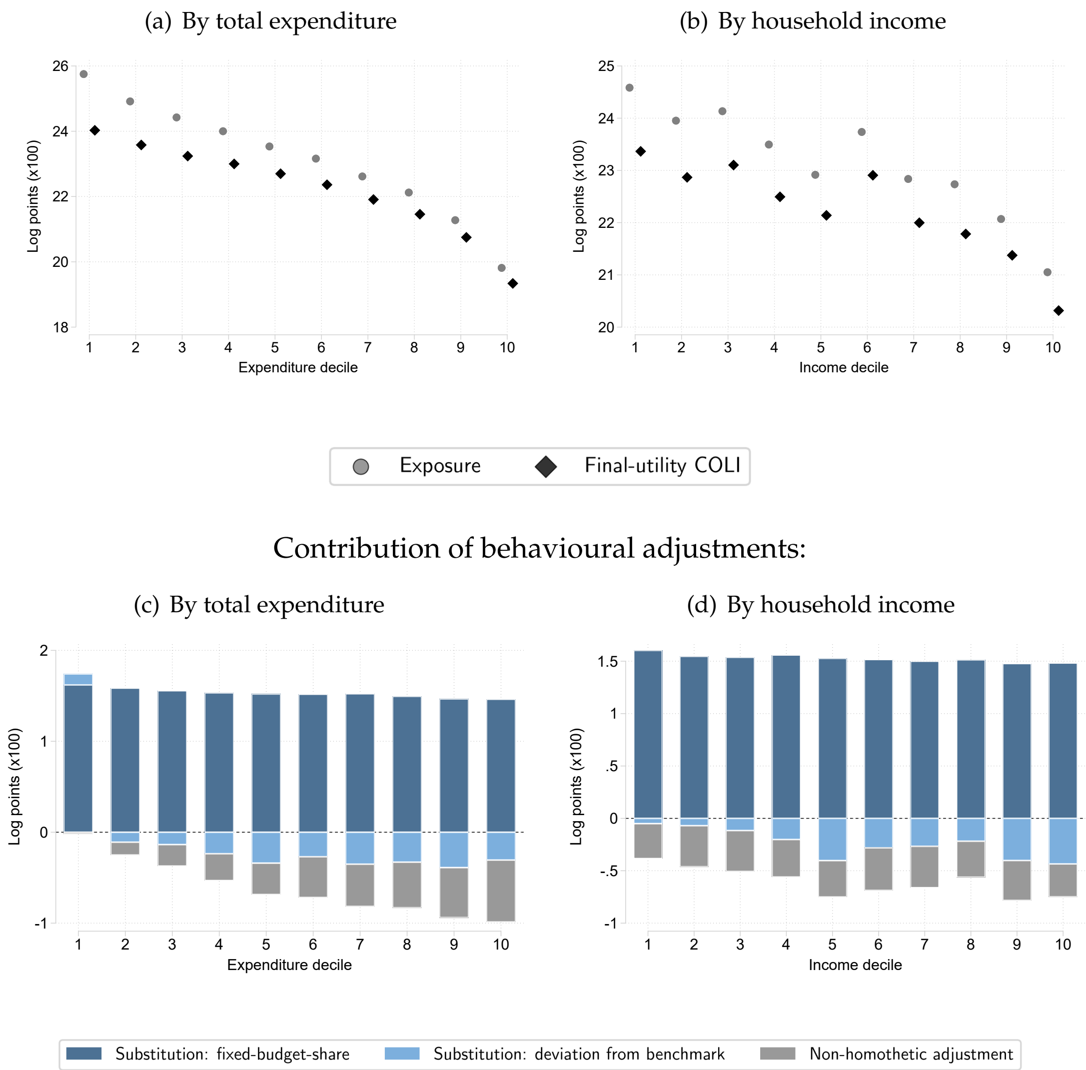

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2021–2023). Panels (a) and (c) group households into deciles of equivalised 2021 total fast-moving consumer goods expenditure; panels (b) and (d) group households into deciles of equivalised 2021 income. Panels (a) and (b) plot, by decile, average log Laspeyres inflation exposure and the final-utility cost-of-living index, approximated by the log Jaravel–Lashkari index evaluated at the final-period utility level, $\log \Pi^{JL}_{1,T}(u_T)$*. Panels (c) and (d) decompose* Exposure $- \log \Pi^{JL}_{1,T}(u_T)$ *into substitution and the non-homothetic adjustment, splitting substitution into the fixed-budget-share adjustment and the deviation from this benchmark. Positive values reduce the final-utility cost-of-living index relative to Laspeyres exposure; negative values raise it. All values are expressed in log points (*$\times 100$*) and averaged within each decile. Table 4.3 reports the corresponding inequality accounting in percentage points.*

Table 4.3 provides the corresponding term-by-term inequality accounting in percentage points, decomposing a single bottom–top gap into exposure, substitution, and the non-homothetic adjustment. The bottom–top expenditure-decile gap falls from 7.50 percentage points under Laspeyres exposure to 5.86 percentage points in the final-utility cost-of-living index: substitution attenuates the gap by 0.84 percentage points, and the non-homothetic adjustment attenuates it by a further 0.80 percentage points. Across income deciles, the gap falls from 4.45 to 3.79 percentage points: substitution attenuates the gap by 0.69 percentage points, while the non-homothetic adjustment slightly increases it by 0.04 percentage points. Thus, behavioural adjustments compress—but by no means eliminate—the regressive pattern in inflation-driven welfare losses implied by initial-basket exposure.

Table 4.3: *Accounting for inequality in inflation-driven welfare changes*

	Bottom–top quartile gap		Bottom–top decile gap	
	Expenditure (p.p.)	Income (p.p.)	Expenditure (p.p.)	Income (p.p.)
Exposure contribution	5.49	3.01	7.50	4.45
Substitution contribution	-0.65	-0.54	-0.84	-0.69
Fixed-budget-share adjustment	*-0.23*	*-0.15*	*-0.31*	*-0.22*
Deviation from fixed-budget-share benchmark	*-0.41*	*-0.39*	*-0.53*	*-0.47*
Non-homothetic adjustment contribution	-0.60	0.01	-0.80	0.04
Final-utility cost-of-living index	4.24	2.48	5.86	3.79

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2021–2023). The table decomposes bottom–top differences in cumulative cost-of-living changes, measured in percentage points. Positive values increase the bottom–top gap; negative values attenuate it. The exposure row is the Laspeyres gap. The substitution row is the change in the gap when moving from the Laspeyres index to the Jaravel–Lashkari index evaluated at initial-period utility. The italicised rows decompose this substitution contribution into the fixed-budget-share adjustment and the deviation from that benchmark. The non-homothetic adjustment row is the change in the gap when moving from the initial-utility to the final-utility Jaravel–Lashkari index. The final row is the bottom–top gap in the final-utility cost-of-living index.

Summary. Our results highlight three broad lessons about inflation episodes with cheapflation and non-homothetic demand. First, inflation exposure is highly unequal across households and strongly regressive along both the expenditure and income distributions. This regressive pattern is driven largely by within-category differences in product choice along the quality ladder.

Second, substitution responses reduce average cost-of-living growth only moderately, and by less than implied by the fixed-budget-share Cobb–Douglas benchmark. This pattern is consistent with the products experiencing the strongest relative price increases being those from which households substitute away least strongly.

Third, the non-homothetic demand adjustment adds a separate welfare channel. When purchasing power falls, households move along Engel curves toward lower-expenditure-elasticity goods, which in a cheapflation episode are precisely the goods whose relative prices rise most. As a result, the cost of living is higher at the realised final-period living standard than at the initial one. On average, this non-homothetic channel raises the equivalent-variation loss from inflation over and above what is captured by exposure and substitution.

5 Dynamic exposure to repeat cheapflation

The analysis so far measures the welfare consequences of the price changes that occurred during the 2021Q3–2023Q3 inflation episode. In this final section, we ask whether the basket adjustments made during the episode change households' exposure to a repeat of the same pattern of relative price changes.

This exercise holds fixed the realised 2021Q3–2023Q3 relative-price vector and applies it to households' initial- and final-period consumption baskets. The goal is to quantify whether the combination of cheapflation and falling purchasing power left some households with post-episode baskets that are more exposed to a similar cheapflation shock than the baskets they held before the episode began. This matters because similar within-category cheapflation patterns can arise from other shocks, including trade-policy changes: recent evidence suggests that tariffs can generate larger price increases for cheaper varieties than for more expensive ones (Cavallo et al. 2025).

Repeat-cheapflation exposure. Let $\Delta p_i = p_{iT}/p_{i1}$ denote the cumulative gross price change for product i over the episode. For each household h, let $s_{hiT} \equiv p_{iT}q_{hiT}/x_{hT}$ denote the final-period spending share, where $x_{hT} = \sum_{j\in\Omega} p_{jT}q_{hjT}$. We define the repeat-cheapflation Laspeyres index

$$\Pi^{L,\text{rep}}_{h(1,T)} = \sum_{i\in\Omega} s_{hiT}\Delta p_i.$$

This index measures how much the cost of household h's final-period basket would rise if the same product-level relative-price changes were applied again.[16] The change

$$\Delta^{\text{exp}}_h \equiv \log \Pi^{L,\text{rep}}_{h(1,T)} - \log \Pi^{L}_{h(1,T)},$$

therefore measures how observed basket reallocation between 2021Q3 and 2023Q3 changed household h's exposure to a repeat cheapflation shock.

Non-homothetic adjustment and repeat exposure. Figure 5.1 relates repeat-cheapflation exposure to the non-homothetic adjustment from the decomposition above. The figure assesses whether households with more negative non-homothetic adjustments ended the episode with baskets more exposed to the realised cheapflation price vector, and whether their exposure increased relative to their initial baskets.

[16] In this index we treat the period-T consumption bundle as the reference bundle, with period-T prices as the initial prices, and consider the effect of a further product-specific price change Δp_i. The repeat-cheapflation Laspeyres index is $\Pi^{L,\text{rep}}_{h(1,T)} = \frac{\sum_{i\in\Omega} p_{iT}\Delta p_i q_{hiT}}{\sum_{i\in\Omega} p_{iT}q_{hiT}} = \sum_{i\in\Omega} s_{hiT}\Delta p_i$. For comparison, the initial exposure index studied above is $\Pi^{L}_{h(1,T)} = \sum_{i\in\Omega} s_{hi1}\Delta p_i$, with s_{hi1} defined at the start of the 2021Q3–2023Q3 episode.

Panel (a) shows that households with more negative non-homothetic adjustments have higher repeat-cheapflation exposure. This is consistent with these households ending the episode with baskets more tilted toward the products whose relative prices rose most. Panel (b) shows that households with the most negative non-homothetic adjustments also experienced the largest increases in exposure between the beginning and end of the episode. Moving from the top decile of the non-homothetic-adjustment distribution to the bottom decile is associated with an increase in repeat-cheapflation exposure of around 0.5 log points, and a substantially larger increase in exposure relative to the household's initial basket.

These results highlight an intertemporal exposure margin. Households with the largest negative non-homothetic adjustments ended the episode with baskets more tilted toward cheapflation-intensive products, and their exposure increased most over the episode. Because the gradient in the change in exposure is larger than the gradient in repeat exposure itself, these households were not initially the most exposed to the 2021Q3–2023Q3 cheapflation vector. The episode therefore reallocated exposure toward households whose baskets moved most strongly along Engel curves toward lower-expenditure-elasticity products. If a similar pattern of relative price changes were to recur, these households would be more exposed than they were at the start of the episode.

Figure 5.1: *Dynamic exposure to repeat cheapflation*

(a) Repeat-exposure index (b) Change in exposure

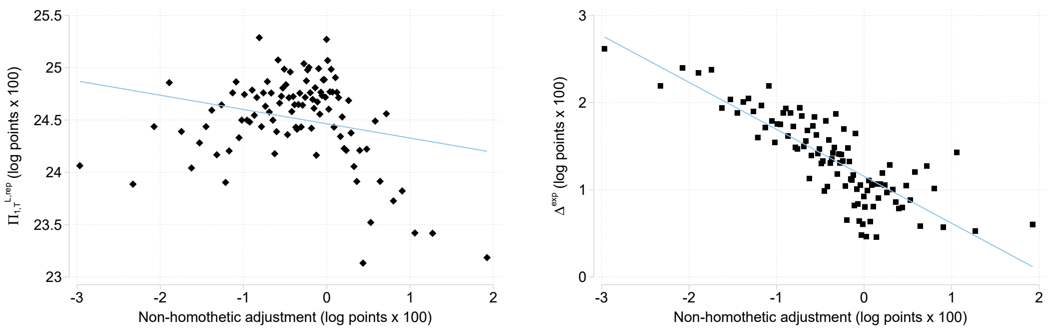

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2021–2023). Households are assigned to 100 bins based on percentiles of the non-homothetic adjustment in Proposition 1. Panel (a) plots the repeat-cheapflation exposure index by bin. Panel (b) plots, by the same bins, the change in exposure, Δ_h^{exp}, defined as the log repeat-cheapflation exposure index minus the log initial-basket exposure index. All values are expressed in log points ($\times 100$).

6 Conclusion

We study cost-of-living inequality during the 2021Q3–2023Q3 UK inflation episode. Using household scanner data and a second-order decomposition of cost-of-living changes implemented with household price indexes, we separate the roles of initial-basket exposure, substitution, and the non-homothetic adjustment. We show that cheapflation—faster price growth at the lower end of the quality ladder—generated large and regressive differences in inflation exposure. Behavioural responses compressed, but did not overturn, the regressive exposure gradient.

A key mechanism in this episode is the interaction between falling real expenditure and cheapflation. When real purchasing power falls, households move along Engel curves toward products lower down the quality ladder. Cheapflation means that these products are also those experiencing the strongest relative price increases. This alignment raises average equivalent-variation losses and explains why the cost-of-living index evaluated at final utility is higher than the corresponding index evaluated at initial utility. It also changes the composition of households' baskets in a way that can increase exposure if similar relative-price patterns recur.

These results speak to broader debates about inflation's macroeconomic and political consequences. On the macroeconomic side, they complement work showing that inflation concentrated in necessities can alter the transmission and optimal design of monetary policy (e.g., Olivi et al. 2024) and that trading down can have non-trivial general-equilibrium effects through the composition of demand and labour use (Jaimovich et al. 2019). On the political economy side, our findings help rationalise why households with fewer resources report feeling particularly squeezed by recent inflation (Binetti et al. 2024), and why inflationary episodes have been linked to rising support for populist parties and political discontent (Federle et al. 2024). Taken together, the evidence underscores that in cheapflation episodes, the relevant policy question is not only how high inflation is, but also which prices move—and how those movements interact with non-homothetic demand to shape both the level and distribution of living standards over time.

References

Adams, A. and P. Levell (2014). Measuring poverty when inflation varies across households. *JRF Report*.

Aguiar, M. and E. Hurst (2007). Life-cycle prices and production. *American Economic Review 97*(5), 1533–1559.

Argente, D. and M. Lee (2021). Cost of Living Inequality During the Great Recession. *Journal of the European Economic Association 19*(2), 913–952.

Atkin, D., B. Faber, T. Fally, and M. Gonzalez-Navarro (2024). Measuring Welfare and Inequality with Incomplete Price Information. *Quarterly Journal of Economics 139*(1), 419–475.

Auer, R., A. Burstein, S. Lein, and J. Vogel (2024). Unequal expenditure switching: Evidence from Switzerland. *Review of Economic Studies 91*(5), 2572–2603.

Balk, B. M. (1990). On Calculating Cost-of-Living Index Numbers for Arbitrary Income Levels. *Econometrica 58*(1), 75–92.

Baqaee, D. R. and A. Burstein (2023). Welfare and Output With Income Effects and Taste Shocks. *Quarterly Journal of Economics 138*(2), 769–834.

Baqaee, D. R., A. T. Burstein, and Y. Koike-Mori (2024). Measuring Welfare by Matching Households Across Time. *Quarterly Journal of Economics 139*(1), 533–573.

Becker, J. (2025). Do poor households pay higher markups in recessions? *mimeo*.

Binetti, A., F. Nuzzi, and S. Stantcheva (2024). People's understanding of inflation. *Journal of Monetary Economics 148*, 103652.

Broda, C. and D. E. Weinstein (2010). Product Creation and Destruction: Evidence and Price Implications. *American Economic Review 100*(3), 691–723.

Cavallo, A. and O. Kryvtsov (2024). Price Discounts and Cheapflation During the Post-Pandemic Inflation Surge. *Journal of Monetary Economics 148*(103644).

Cavallo, A., P. Llamas, and F. M. Vazquez (2025). Tracking the Short-Rrun Price Impact of US Tariffs. *NBER WP 34496*.

Comin, D., D. Lashkari, and M. Mestieri (2021). Structural Change With Long-Run Income and Price Effects. *Econometrica 89*(1), 311–374.

Competition Commission (2008). *The Supply of Groceries in the UK: Market Investigation*. London: The Stationery Office.

Crawford, I. and Z. Oldfield (2002). Distributional aspects of inflation. *IFS Commentary No. 90*.

D'Acunto, F., U. Malmendier, J. Ospina, and M. Weber (2021). Exposure to Grocery Prices and Inflation Expectations. *Journal of Political Economy 129*(5), 1615–1639.

Del Canto, F., J. Grigsby, E. Qian, and C. Walsh (2025). Are inflationary shocks regressive? a feasible set approach. *Quarterly Journal of Economics 140*(4), 2685–2747.

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

Divisia, F. (1926). L'indice Monétaire Et la Théorie de la Monnaie (Suite et fin). *Revue d'économie politique 40*(1), 49–81.

Doepke, M. and M. Schneider (2006). Inflation and the Redistribution of Nominal Wealth. *Journal of Political Economy 114*(6), 1069–1097.

Federle, J., C. Mohr, and M. Schularick (2024). Inflation Surprises and Election Outcomes. *SSRN 5032871*.

Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *The American Economic Review*, 157–177.

Ferreira, C., J. M. Leiva, G. Nuño, Á. Ortiz, T. Rodrigo, and S. Vazquez (2026). The heterogeneous impact of inflation on households' balance sheets. *SERIEs 17*(1), 31–53.

Griffith, R., E. Leibtag, A. Leicester, and A. Nevo (2009). Consumer Shopping Behavior: How Much Do Consumers Save? *Journal of Economic Perspectives 23*(2), 99–120.

Hagenaars, A. J., K. De Vos, M. Asghar Zaidi, et al. (1994). Poverty statistics in the late 1980s: Research based on micro-data.

Jaimovich, N., S. Rebelo, and A. Wong (2019). Trading down and the business cycle. *Journal of Monetary Economics 102*, 96–121.

Jaravel, X. (2019). The Unequal Gains from Product Innovations: Evidence from the U.S. Retail Sector. *Quarterly Journal of Economics 134*(2), 715–783.

Jaravel, X. (2021). Inflation Inequality: Measurement, Causes, and Policy Implications. *Annual Review of Economics 13*(1), 599–629.

Jaravel, X. (2024). Distributional Consumer Price Indices. *mimeo*.

Jaravel, X. and D. Lashkari (2024). Measuring Growth in Consumer Welfare with Income-Dependent Preferences: Nonparametric Methods and Estimates for the United States. *Quarterly Journal of Economics 139*(1), 477–532.

Jaravel, X. and M. O'Connell (2020a). High-Frequency Changes in Shopping Behaviours, Promotions and the Measurement of Inflation: Evidence from the Great Lockdown. *Fiscal Studies 41*(3), 733–755.

Jaravel, X. and M. O'Connell (2020b). Real-time price indices: Inflation spike and falling product variety during the Great Lockdown. *Journal of Public Economics 191*(104270).

Kaplan, G. and S. Schulhofer-Wohl (2017). Inflation at the Household Level. *Journal of Monetary Economics 91*, 19–38.

Klick, J. and A. Stockburger (2021). Experimental CPI for lower and higher income households. *Monthly Labor Review July 2024*.

Klick, J. and A. Stockburger (2024). Examining U.S. inflation across households grouped by equivalized income. *BLS Working Paper No. 537*.

Konüs, A. A. (1939). The Problem of the True Index of the Cost of Living. *Econometrica 7*(1), 10–29.

Leicester, A., C. O'Dea, and Z. Oldfield (2008). The inflation experience of older households. *IFS Commentary No. 106*.

Levell, P., M. O'Connell, and K. Smith (2025). The welfare effects of price shocks and household relief packages: Evidence from an energy crisis. *CEPR Discussion Paper 19939*.

Meyer, B. D. and J. X. Sullivan (2023). Consumption and Income Inequality in the United States since the 1960s. *Journal of Political Economy 131*(2), 247–284.

Office for National Statistics (2022). Inflation and the cost of living for household groups, UK: October 2022. Technical report, Office for National Statistics.

Olivi, A., V. Sterk, and D. Xhani (2024). Optimal Monetary Policy during a Cost-of-Living Crisis. *mimeo*.

Poterba, J. M. (1989). Lifetime Incidence and the Distributional Burden of Excise Taxes. *American Economic Review 79*(2), 325–330.

Sangani, K. (2025). Cheapflation cycles. *Kilts Center at Chicago Booth Marketing Data Center Paper* (1).

Sangani, K. (2026). Complete pass-through in levels. *The Quarterly Journal of Economics 141*(2), 1077–1135.

APPENDIX:
FOR ONLINE PUBLICATION

Measuring cost of living inequality during an inflation surge

Tao Chen, Peter Levell and Martin O'Connell
August 2026

A Data appendix

A.1 Price dispersion for identical products

Our baseline inflation measure uses prices that are common across households. In the UK grocery industry, where retailers have national store coverage and typically follow national pricing policies (see Competition Commission 2008), this assumption is arguably more innocuous than in other settings. Furthermore, the two most prominent low-cost retailers, Aldi and Lidl, primarily sell private-label products unique to them. Our baseline analysis therefore incorporates substitution toward the main low-cost retailers' product lines.

Nonetheless, there is some dispersion in the prices paid for identical products, due, for instance, to temporary promotions and—among manufacturer-branded products available at multiple retailers—to retailer-specific discounts. If households with different resources change their propensity to take advantage of such low prices over time, this could affect measured inflation inequality. To assess this, we use a price index suggested by Aguiar and Hurst (2007), which measures dispersion in the prices households pay for a *fixed* shopping basket.

Let q_{hit} denote the volume of product i purchased by household h in year–quarter t, and define the household-specific price by $p_{hit} = x_{hit}/q_{hit}$.[17] Had household h paid the common (average) prices for its realised basket, its expenditure would have been $\tilde{x}_{ht} = \sum_i q_{hit} p_{it}$. The Aguiar and Hurst (AH) index compares the true cost of the household's basket ($x_{ht} = \sum_i x_{hit}$) with its cost at average prices:

$$\Pi^{AH}_{ht} = \frac{x_{ht}}{\tilde{x}_{ht}}. \tag{A.1}$$

Figure A.1 summarises the evolution of the AH index over time. In panels (a) and (b) we show how the index varies across quartiles of the equivalised expenditure distribution between 2012 and 2023. For each calendar year, we assign households to expenditure quartiles (based on that year) and then, for each year-quarter, report the average AH index across households in each quartile, expressed as a deviation from the quarter mean across all households. Panel (a) defines products at the brand–pack-size level. Panel (b) redefines products at the brand–pack-size–retailer level (which matters for branded products sold across multiple retailers, but leaves the definition of private-label products unchanged). The lines in panel

[17] Common prices are quantity-weighted averages of household-specific prices: $p_{it} = \frac{\sum_h x_{hit}}{\sum_h q_{hit}} = \sum_h \frac{q_{hit}}{\sum_{h'} q_{h'it}} p_{hit}$.

(a) reflect the influence of price dispersion both across and within retailers; those in panel (b) strip out the former and therefore reflect only within-retailer price dispersion.

Figure A.1: *Aguiar and Hurst price dispersion index*

By expenditure quartile

(a) Across products

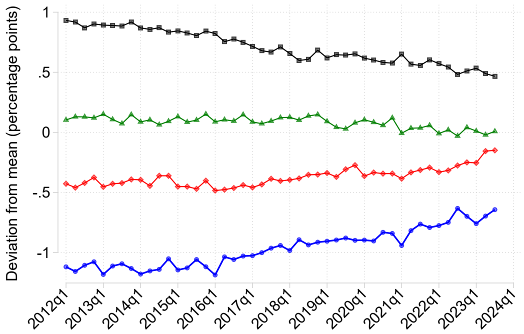

(b) Across product-retailers

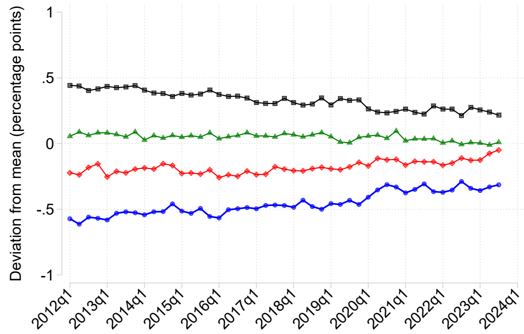

By income quartile

(c) Across products

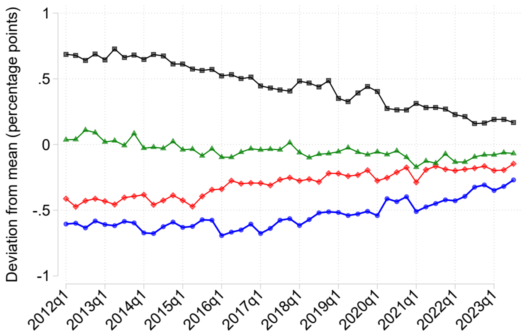

(d) Across product-retailers

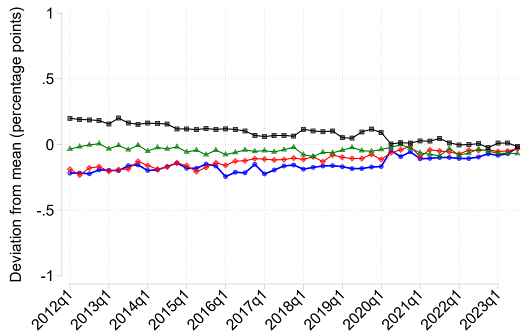

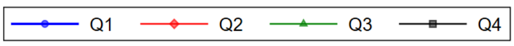

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2012–2023). Panels (a) and (c) define products at the brand–pack-size level; panels (b) and (d) further condition on retailer, defining products at the brand–pack-size–retailer level. Panels (a) and (b) plot, for each year–quarter, the mean AH index within equivalised expenditure quartiles, expressed as deviations from the quarterly mean. Panels (c) and (d) repeat (a) and (b), instead splitting households by equivalised income quartile.

In 2012, households in the top expenditure quartile paid around 2 percentage points more for a fixed basket of goods than those in the bottom quartile; roughly 1 percentage point of this gap reflects cross-retailer variation and 1 percentage point within-retailer variation. This gap narrows over time: by 2023, households in the top expenditure quartile paid only around 1.1 percentage points more than those in the bottom quartile, again split approximately evenly between across- and within-retailer dispersion. Because lower-resource households pay slightly lower prices for identical goods, the narrowing of this price gap tends to reinforce measured inflation inequality by shrinking the price-shopping advantage of lower-resource households. The effect is small: over 2021Q3–2023Q3, the decline in the top–bottom AH gap was 0.48 percentage points when products are defined at the brand–pack-size level, with virtually no change when products are defined at the brand–pack-size–retailer level.

Panels (c) and (d) repeat the analysis grouping households by equivalised income quartiles. The results are similar, confirming that modest and slowly changing price dispersion for identical products plays little role in driving inflation inequality.

A.2 Income measures

Figure A.2 illustrates the relationship between annual equivalised fast-moving consumer goods expenditure and equivalised income in 2021, with income reported in bands. The fitted line shows a clear, increasing relationship, confirming that the two measures are strongly, though not perfectly, correlated. Departures from a one-for-one mapping reflect both the banded and top-coded nature of the income data and the conceptual distinction between current income and expenditure as a proxy for longer-run household resources.

Figure A.2: *Annual expenditure and household income*

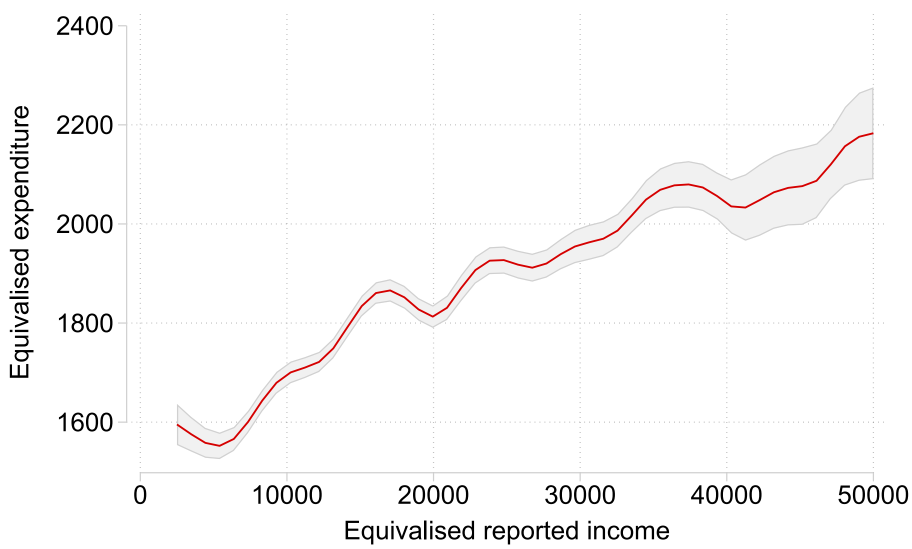

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2012–2023). The figure plots a local-polynomial regression of annual equivalised expenditure on equivalised household income in 2021, with a 95% confidence band. Income is reported in £10,000 bands; we assign each household the midpoint of its reported band.

A.3 Price dispersion across similar products

Figure A.3 shows the distributions of interquartile ranges for log prices across products within Numerator product categories in 2021Q3, 2022Q3, and 2023Q3. When lower-priced products within a category experience faster price growth—a phenomenon known as "cheapflation"—within-category price dispersion should fall, shifting the distribution to the left. Consistent with this pattern, we observe a clear leftward shift in 2023 relative to 2022 and 2021.

We cross-check this result and assess whether cheapflation affects categories beyond fast-moving consumer goods (FMCG), using data from the Office for National Statistics consumer price microdata—the raw price data underlying the UK Consumer Price Index (CPI). Figure A.4 plots the distributions of interquartile ranges for log prices within CPI "item" categories in December 2021, 2022 and 2023. Items are disaggregated product groups, such as "large white loaf unsliced, 800g", which are more granular than the product categories used in Figure A.3.

Figure A.3: *Price dispersion with Numerator categories*

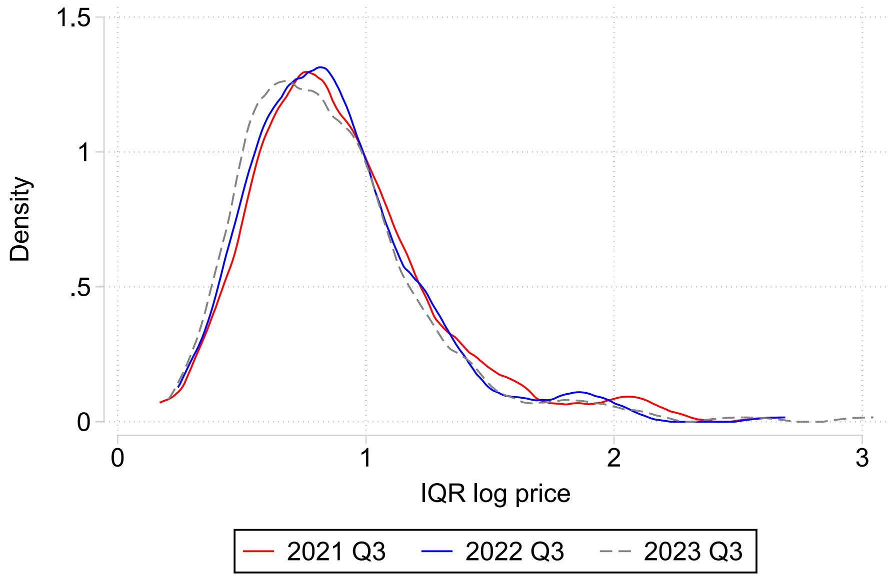

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2012–2023). The graph shows the distribution of interquartile ranges of log prices across products within product categories in 2021Q3, 2022Q3 and 2023Q3.

We plot these distributions for FMCG—constructed by combining food consumed at home, alcoholic and soft drinks consumed at home, pet products and pharmaceutical and personal care products—as well as for other non- and semi-durable goods categories: clothing, household appliances (including tools, household articles and electronics), and leisure goods (including toys, sports equipment and books). As with the Numerator data, we see a clear leftward shift in the distribution for FMCG in 2023. There is also evidence of falling price dispersion for leisure goods (beginning in 2022) and, to a lesser extent, for clothing items, whereas we do not see as clear a pattern of falling price dispersion for household appliances.

Figure A.4: *Price dispersion with CPI microdata categories*

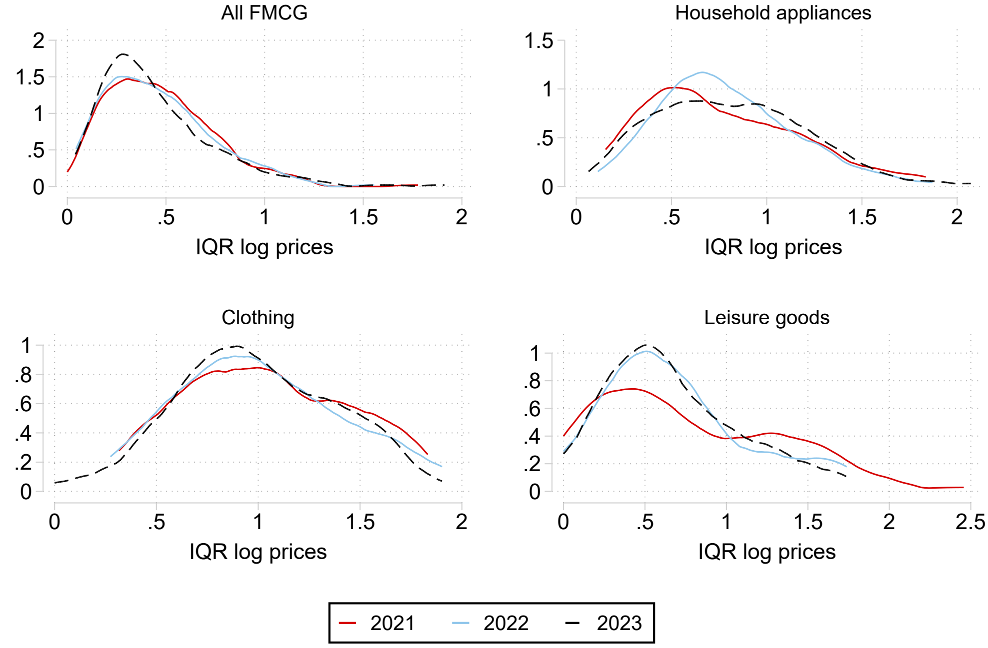

Notes: Authors' calculations using CPI microdata produced by the Office for National Statistics. The graph shows the distribution of interquartile ranges of log prices across sampled prices within CPI "items" in December 2021, 2022 and 2023. FMCG denotes fast-moving consumer goods.

A.4 Spending across the quality ladder

Figure A.5 illustrates how the average quality rung of products purchased in the initial quarter of each nine-quarter period varies across income deciles. The relationship is similar to the pattern across expenditure deciles shown in Figure 2.1(b).

Figure A.5: *Average quality rung, by income*

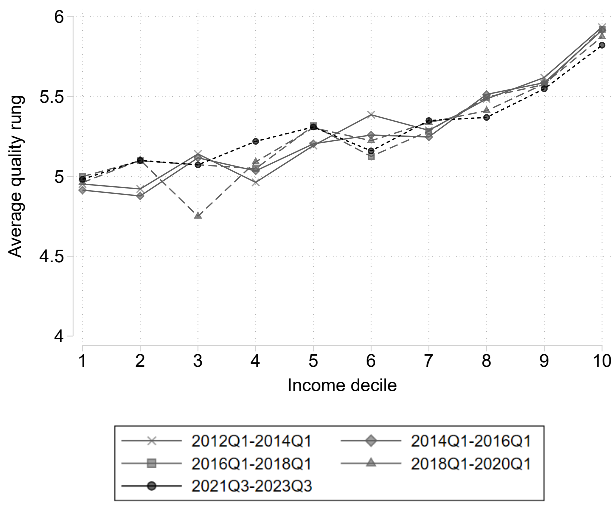

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2012–2023). The figure reports the average quality rung of households' purchases by decile of the equivalised income distribution.

A.5 Quality downgrading and expenditure changes

A potential concern is that the negative relationship shown in Figure 2.1(d) partly reflects a mechanical effect: trading down to lower-quality, cheaper goods reduces nominal expenditure, which could mechanically induce a relationship between quality downgrading and deflated-expenditure changes.

To assess whether the trading-down patterns are mechanically induced by the construction of deflated expenditure, we estimate a simple household–segment regression, first by OLS and then by 2SLS using an instrument for deflated expenditure. For each household–segment (h, g) pair, we compute the log change between 2021Q3 and 2023Q3 in segment-level quality and relate this to the household's log change in deflated total expenditure over the same period, including segment fixed effects. Column (1) of Table A.1 reports the OLS estimate, which shows a strong and statistically precise relationship: households whose deflated expenditure falls more tend to shift more toward lower-quality products within segments, consistent with Figure 2.1(d).

To remove the mechanically induced component of this relationship, we instrument the log change in deflated expenditure with a "leave-one-out" measure that recomputes deflated-expenditure growth for each household excluding segment g from both nominal expenditure and the Laspeyres deflator. By construction, within-segment changes in quality choice do not directly affect this instrument. The 2SLS estimate remains statistically significant and is only modestly smaller than the OLS coefficient. This indicates that the negative association between deflated-expenditure declines and quality downgrading is not primarily an artefact of including segment-g quality choices in the construction of deflated expenditure. However, this exercise does not isolate exogenous changes in purchasing power or rule out correlated changes in quality choice across segments, so we interpret the evidence as descriptive.

Table A.1: *Within-segment trading down*

Dependent variable: log change in segment-level average quality		
	OLS	2SLS
Log change in deflated expenditure	0.063***	0.047***
	(0.006)	(0.004)
Observations	179,682	179,682

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2012–2023). The dependent variable is the log change in household h's segment-level quality in segment g between 2021Q3 and 2023Q3, $\log(Q_{hgT}/Q_{hg1})$*, where* Q_{hgt} *is the expenditure-weighted average quality rung of household h's purchases in segment g at time t. The regressor is the log change in household expenditure deflated with a household-specific Laspeyres price index,* $\log(x_{hT}/x_{h1}) - \log \Pi^L_{h(1,T)}$*. Column (1) reports OLS estimates. Column (2) reports 2SLS estimates in which* $\log(x_{hT}/x_{h1}) - \log \Pi^L_{h(1,T)}$ *is instrumented with a leave-one-out measure constructed by excluding segment g from both nominal expenditure and the Laspeyres deflator. All regressions include segment fixed effects; standard errors, in parentheses, are clustered at the segment level. Because both variables are log changes, a positive coefficient implies that households with larger declines in deflated expenditure experience larger quality downgrades.* $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

B Measurement appendix

B.1 Proof of Proposition 1

Lemma 1. *Suppose that, in a neighbourhood of* $(\mathbf{p}_1, x_1)$*, the indirect utility function, expenditure function, and Marshallian budget-share functions are twice continuously differentiable and satisfy the standard duality identity* $e(\mathbf{p}, v(\mathbf{p}, x)) = x$*. Fix a reference price vector* $\mathbf{p}_1$ *and define, suppressing the dependence on* $\mathbf{p}_1$*,*

$$\mu(\mathbf{p}, x) \equiv e(\mathbf{p}_1, v(\mathbf{p}, x)).$$

Let

$$s_i(\mathbf{p}, x) \equiv \frac{p_i q_i(\mathbf{p}, x)}{x}$$

denote the Marshallian budget-share function of good i at $(\mathbf{p}, x)$*, and write* $s_{i1} \equiv s_i(\mathbf{p}_1, x_1)$*. Then the first and second log-derivatives of* $\log \mu(\mathbf{p}, x)$ *at* $(\mathbf{p}_1, x_1)$ *satisfy:*

$$\left.\frac{\partial \log \mu(\mathbf{p}, x)}{\partial \log p_i}\right|_{(\mathbf{p}_1, x_1)} = -\, s_{i1},$$

$$\left.\frac{\partial \log \mu(\mathbf{p}, x)}{\partial \log x}\right|_{(\mathbf{p}_1, x_1)} = 1,$$

$$\left.\frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial \log p_i\, \partial \log x}\right|_{(\mathbf{p}_1, x_1)} = -\left.\frac{\partial s_i(\mathbf{p}, x)}{\partial \log x}\right|_{(\mathbf{p}_1, x_1)},$$

$$\left.\frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial (\log x)^2}\right|_{(\mathbf{p}_1, x_1)} = 0,$$

$$\left.\frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial \log p_i\, \partial \log p_j}\right|_{(\mathbf{p}_1, x_1)} = -\left.\frac{\partial s_i(\mathbf{p}, x)}{\partial \log p_j}\right|_{(\mathbf{p}_1, x_1)} + s_{i1} \left.\frac{\partial s_j(\mathbf{p}, x)}{\partial \log x}\right|_{(\mathbf{p}_1, x_1)}.$$

Proof of Lemma 1. Fix $\mathbf{p}_1$ and define $\mu(\mathbf{p}, x) \equiv e(\mathbf{p}_1, v(\mathbf{p}, x))$. Throughout, derivatives are taken holding the other arguments fixed, and we use the log-derivative operators $\partial/\partial \log p_i$ and $\partial/\partial \log x$.

Step 1: First derivatives. Let $m(u) \equiv \log e(\mathbf{p}_1, u)$, so that $\log \mu(\mathbf{p}, x) = m(v(\mathbf{p}, x))$. By the chain rule,

$$\frac{\partial \log \mu(\mathbf{p}, x)}{\partial \log p_i} = m'(v(\mathbf{p}, x))\, \frac{\partial v(\mathbf{p}, x)}{\partial \log p_i}, \tag{B.1}$$

$$\frac{\partial \log \mu(\mathbf{p}, x)}{\partial \log x} = m'(v(\mathbf{p}, x))\, \frac{\partial v(\mathbf{p}, x)}{\partial \log x}. \tag{B.2}$$

At $\mathbf{p} = \mathbf{p}_1$ we have, by duality,

$$\mu(\mathbf{p}_1, x) = e(\mathbf{p}_1, v(\mathbf{p}_1, x)) = x, \tag{B.3}$$

hence $\log \mu(\mathbf{p}_1, x) = \log x$ for all x and therefore

$$\left.\frac{\partial \log \mu(\mathbf{p}, x)}{\partial \log x}\right|_{(\mathbf{p}_1, x)} = 1. \tag{B.4}$$

Evaluating (B.2) at $(\mathbf{p}_1, x)$ and using (B.4) gives

$$m'(v(\mathbf{p}_1, x)) = \left(\left.\frac{\partial v(\mathbf{p}, x)}{\partial \log x}\right|_{(\mathbf{p}_1, x)} \right)^{-1}. \tag{B.5}$$

Roy's identity gives

$$s_i(\mathbf{p}, x) = -\frac{\partial v(\mathbf{p}, x)/\partial \log p_i}{\partial v(\mathbf{p}, x)/\partial \log x}. \tag{B.6}$$

Combining (B.1), (B.5), and (B.6), evaluated at $(\mathbf{p}_1, x_1)$, yields

$$\left.\frac{\partial \log \mu(\mathbf{p}, x)}{\partial \log p_i}\right|_{(\mathbf{p}_1, x_1)} = -s_{i1}.$$

Together with (B.4), evaluated at $x = x_1$, this gives

$$\left.\frac{\partial \log \mu(\mathbf{p}, x)}{\partial \log x}\right|_{(\mathbf{p}_1, x_1)} = 1.$$

Step 2: Expenditure–expenditure second derivative and curvature relation. Differentiate (B.2) with respect to $\log x$:

$$\frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial (\log x)^2} = m'(v(\mathbf{p}, x)) \frac{\partial^2 v(\mathbf{p}, x)}{\partial (\log x)^2} + m''(v(\mathbf{p}, x)) \left(\frac{\partial v(\mathbf{p}, x)}{\partial \log x} \right)^2. \tag{B.7}$$

Since $\log \mu(\mathbf{p}_1, x) = \log x$ for all x, we have

$$\left.\frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial (\log x)^2}\right|_{(\mathbf{p}_1, x)} = 0.$$

Evaluating at $x = x_1$ gives

$$\left.\frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial (\log x)^2}\right|_{(\mathbf{p}_1, x_1)} = 0.$$

Moreover, evaluating (B.7) at $(\mathbf{p}_1, x)$ and using (B.5) gives the curvature relation

$$m''(v(\mathbf{p}_1, x)) \left(\left.\frac{\partial v(\mathbf{p}, x)}{\partial \log x}\right|_{(\mathbf{p}_1, x)} \right)^2 = -\frac{\left.\frac{\partial^2 v(\mathbf{p}, x)}{\partial (\log x)^2}\right|_{(\mathbf{p}_1, x)}}{\left.\frac{\partial v(\mathbf{p}, x)}{\partial \log x}\right|_{(\mathbf{p}_1, x)}}. \tag{B.8}$$

Step 3: Cross derivative. Differentiate (B.1) with respect to $\log x$:

$$\frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial \log p_i \, \partial \log x} = m'(v(\mathbf{p}, x)) \frac{\partial^2 v(\mathbf{p}, x)}{\partial \log p_i \, \partial \log x} + m''(v(\mathbf{p}, x)) \frac{\partial v(\mathbf{p}, x)}{\partial \log p_i} \frac{\partial v(\mathbf{p}, x)}{\partial \log x}. \tag{B.9}$$

Evaluate (B.9) at $(\mathbf{p}_1, x_1)$, use (B.5) and (B.8), and substitute $\frac{\partial v}{\partial \log p_i} = -s_i \frac{\partial v}{\partial \log x}$ from (B.6) to obtain

$$\left.\frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial \log p_i \, \partial \log x}\right|_{(\mathbf{p}_1, x_1)} = \left.\frac{\partial^2 v(\mathbf{p}, x)}{\partial \log p_i \, \partial \log x}\right|_{(\mathbf{p}_1, x_1)} \Bigg/ \left.\frac{\partial v(\mathbf{p}, x)}{\partial \log x}\right|_{(\mathbf{p}_1, x_1)} + s_{i1} \left.\frac{\partial^2 v(\mathbf{p}, x)}{\partial (\log x)^2}\right|_{(\mathbf{p}_1, x_1)} \Bigg/ \left.\frac{\partial v(\mathbf{p}, x)}{\partial \log x}\right|_{(\mathbf{p}_1, x_1)}. \tag{B.10}$$

On the other hand, differentiating (B.6) with respect to $\log x$ gives

$$\begin{aligned}\frac{\partial s_i(\mathbf{p}, x)}{\partial \log x} &= -\frac{\partial^2 v / \partial \log x \, \partial \log p_i}{\partial v / \partial \log x} + \frac{\partial v / \partial \log p_i}{(\partial v / \partial \log x)^2} \frac{\partial^2 v}{\partial (\log x)^2} \\ &= -\frac{\partial^2 v / \partial \log x \, \partial \log p_i}{\partial v / \partial \log x} - s_i(\mathbf{p}, x) \frac{\partial^2 v / \partial (\log x)^2}{\partial v / \partial \log x}. \end{aligned} \tag{B.11}$$

Comparing this expression with (B.10) and evaluating at $(\mathbf{p}_1, x_1)$ yields

$$\left.\frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial \log p_i \, \partial \log x}\right|_{(\mathbf{p}_1, x_1)} = -\left.\frac{\partial s_i(\mathbf{p}, x)}{\partial \log x}\right|_{(\mathbf{p}_1, x_1)}.$$

Step 4: Price–price second derivative. Differentiate (B.1) with respect to $\log p_j$:

$$\frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial \log p_i \, \partial \log p_j} = m'(v(\mathbf{p}, x)) \frac{\partial^2 v(\mathbf{p}, x)}{\partial \log p_i \, \partial \log p_j} + m''(v(\mathbf{p}, x)) \frac{\partial v(\mathbf{p}, x)}{\partial \log p_i} \frac{\partial v(\mathbf{p}, x)}{\partial \log p_j}. \tag{B.12}$$

Evaluating (B.12) at $(\mathbf{p}_1, x_1)$ and using (B.5), (B.8), and $\frac{\partial v}{\partial \log p_k} = -s_k \frac{\partial v}{\partial \log x}$ for $k = i, j$ yields

$$\left.\frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial \log p_i \, \partial \log p_j}\right|_{(\mathbf{p}_1, x_1)} = \left.\frac{\partial^2 v(\mathbf{p}, x)}{\partial \log p_i \, \partial \log p_j}\right|_{(\mathbf{p}_1, x_1)} \Bigg/ \left.\frac{\partial v(\mathbf{p}, x)}{\partial \log x}\right|_{(\mathbf{p}_1, x_1)} - s_{i1} s_{j1} \left.\frac{\partial^2 v(\mathbf{p}, x)}{\partial (\log x)^2}\right|_{(\mathbf{p}_1, x_1)} \Bigg/ \left.\frac{\partial v(\mathbf{p}, x)}{\partial \log x}\right|_{(\mathbf{p}_1, x_1)}. \tag{B.13}$$

Now differentiate (B.6) with respect to $\log p_j$:

$$\begin{aligned}\frac{\partial s_i(\mathbf{p}, x)}{\partial \log p_j} &= -\frac{\partial^2 v / \partial \log p_i \, \partial \log p_j}{\partial v / \partial \log x} + \frac{\partial v / \partial \log p_i}{(\partial v / \partial \log x)^2} \frac{\partial^2 v}{\partial \log x \, \partial \log p_j} \\ &= -\frac{\partial^2 v / \partial \log p_i \, \partial \log p_j}{\partial v / \partial \log x} - s_i(\mathbf{p}, x) \frac{\partial^2 v / \partial \log x \, \partial \log p_j}{\partial v / \partial \log x}. \end{aligned}$$

Combining with (B.11), evaluated for good j, gives

$$-\frac{\partial s_i(\mathbf{p}, x)}{\partial \log p_j} + s_i(\mathbf{p}, x) \frac{\partial s_j(\mathbf{p}, x)}{\partial \log x} = \frac{\partial^2 v / \partial \log p_i \, \partial \log p_j}{\partial v / \partial \log x} - s_i(\mathbf{p}, x) s_j(\mathbf{p}, x) \frac{\partial^2 v / \partial (\log x)^2}{\partial v / \partial \log x}.$$

Evaluating at $(\mathbf{p}_1, x_1)$ and comparing with (B.13) gives

$$\left.\frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial \log p_i \, \partial \log p_j}\right|_{(\mathbf{p}_1, x_1)} = -\left.\frac{\partial s_i(\mathbf{p}, x)}{\partial \log p_j}\right|_{(\mathbf{p}_1, x_1)} + s_{i1} \left.\frac{\partial s_j(\mathbf{p}, x)}{\partial \log x}\right|_{(\mathbf{p}_1, x_1)}.$$

$\square$

Proof of Proposition 1. Since $u_T = v(\mathbf{p}_T, x_T)$ and $e(\mathbf{p}_T, u_T) = x_T$,

$$\begin{aligned}\log P(\mathbf{p}_1, \mathbf{p}_T; u_T) &= \log e(\mathbf{p}_T, u_T) - \log e(\mathbf{p}_1, u_T) \\ &= \log x_T - \log e(\mathbf{p}_1, v(\mathbf{p}_T, x_T)).\end{aligned}$$

Define $\mu(\mathbf{p}, x) \equiv e(\mathbf{p}_1, v(\mathbf{p}, x))$ and let $g_i \equiv \log p_{iT} - \log p_{i1}$ and $g_x \equiv \log x_T - \log x_1$. Since $\log P(\mathbf{p}_1, \mathbf{p}_T; u_T) = \log x_T - \log \mu(\mathbf{p}_T, x_T)$ and $\log \mu(\mathbf{p}_1, x_1) = \log x_1$, subtracting from $\log x_T$ a second-order Taylor expansion of $\log \mu(\mathbf{p}_T, x_T)$ in log prices and log expenditure around $(\log \mathbf{p}_1, \log x_1)$, evaluated at $(\mathbf{p}_T, x_T)$, yields:

$$\begin{aligned}\log P(\mathbf{p}_1, \mathbf{p}_T; u_T) =& g_x - \sum_i \frac{\partial \log \mu(\mathbf{p}, x)}{\partial \log p_i}\bigg|_{(\mathbf{p}_1, x_1)} g_i - \frac{\partial \log \mu(\mathbf{p}, x)}{\partial \log x}\bigg|_{(\mathbf{p}_1, x_1)} g_x \\ & - \frac{1}{2}\sum_{i,j} \frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial \log p_i \partial \log p_j}\bigg|_{(\mathbf{p}_1, x_1)} g_i g_j - \sum_i \frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial \log p_i \partial \log x}\bigg|_{(\mathbf{p}_1, x_1)} g_i g_x \\ & - \frac{1}{2} \frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial (\log x)^2}\bigg|_{(\mathbf{p}_1, x_1)} g_x^2 + R^{(2)}\end{aligned}$$

where $R^{(2)}$ collects terms of order higher than two, with $R^{(2)} = o\big(\|(g_1, \ldots, g_I, g_x)\|^2\big)$. By Lemma 1, we have:

$$\begin{aligned}\frac{\partial \log \mu(\mathbf{p}, x)}{\partial \log p_i}\bigg|_{(\mathbf{p}_1, x_1)} &= -\, s_{i1} \\ \frac{\partial \log \mu(\mathbf{p}, x)}{\partial \log x}\bigg|_{(\mathbf{p}_1, x_1)} &= 1 \\ \frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial \log p_i \partial \log p_j}\bigg|_{(\mathbf{p}_1, x_1)} &= -\, \frac{\partial s_i}{\partial \log p_j}\bigg|_{(\mathbf{p}_1, x_1)} + s_{i1}\, \frac{\partial s_j}{\partial \log x}\bigg|_{(\mathbf{p}_1, x_1)} \\ \frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial \log p_i \partial \log x}\bigg|_{(\mathbf{p}_1, x_1)} &= -\, \frac{\partial s_i}{\partial \log x}\bigg|_{(\mathbf{p}_1, x_1)} \\ \frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial (\log x)^2}\bigg|_{(\mathbf{p}_1, x_1)} &= 0.\end{aligned}$$

Hence,

$$\begin{aligned}\log P(\mathbf{p}_1, \mathbf{p}_T; u_T) = \sum_i s_{i1} g_i + \frac{1}{2}\sum_{i,j} \left(\frac{\partial s_i}{\partial \log p_j}\bigg|_{(\mathbf{p}_1, x_1)} - s_{i1}\, \frac{\partial s_j}{\partial \log x}\bigg|_{(\mathbf{p}_1, x_1)} \right) g_i g_j \\ + \sum_i \frac{\partial s_i}{\partial \log x}\bigg|_{(\mathbf{p}_1, x_1)} g_i g_x + R^{(2)}\end{aligned}$$

Let $\epsilon_{ij} \equiv \frac{\partial \log q_i}{\partial \log p_j}\Big|_{(\mathbf{p}_1, x_1)}$ denote the Marshallian price elasticity, $\epsilon^H_{ij} \equiv \frac{\partial \log q_i}{\partial \log p_j}\Big|_{(\mathbf{p}_1, u_1)}$ the Hicksian price elasticity, and $\eta_i \equiv \frac{\partial \log q_i}{\partial \log x}\Big|_{(\mathbf{p}_1, x_1)}$ the expenditure elasticity. Note that, at fixed x,

$$\frac{\partial s_i}{\partial \log p_j}\bigg|_{(\mathbf{p}_1, x_1)} = s_{i1}\left(\epsilon_{ij} + \mathbb{1}_{i=j}\right),$$

and, at fixed u_1,

$$\frac{\partial s_i}{\partial \log p_j}\bigg|_{(\mathbf{p}_1, u_1)} = s_{i1}\left(\epsilon^H_{ij} + \mathbb{1}_{i=j} - s_{j1}\right).$$

Using the Slutsky relation $\epsilon_{ij} = \epsilon^H_{ij} - s_{j1}\eta_i$, simple algebra yields

$$\frac{\partial s_i}{\partial \log p_j}\bigg|_{(\mathbf{p}_1, x_1)} = \frac{\partial s_i}{\partial \log p_j}\bigg|_{(\mathbf{p}_1, u_1)} - s_{i1}s_{j1}(\eta_i - 1).$$

Note also that

$$\frac{\partial s_i}{\partial \log x}\bigg|_{(\mathbf{p}_1, x_1)} = s_{i1}(\eta_i - 1).$$

Hence,

$$\begin{aligned}
\log P(\mathbf{p}_1, \mathbf{p}_T; u_T) =& \sum_i s_{i1}g_i + \left(\frac{1}{2}\sum_{i,j}\frac{\partial s_i}{\partial \log p_j}\bigg|_{(\mathbf{p}_1, u_1)} g_i g_j - \frac{1}{2}\sum_{i,j} s_{i1}s_{j1}(\eta_i - 1)g_i g_j\right. \\
&\left. -\frac{1}{2}\sum_{i,j} s_{i1}s_{j1}(\eta_j - 1)g_i g_j + \sum_i s_{i1}(\eta_i - 1)g_i g_x\right) + R^{(2)} \\
=& \sum_i s_{i1}g_i + \frac{1}{2}\sum_{i,j}\frac{\partial s_i}{\partial \log p_j}\bigg|_{(\mathbf{p}_1, u_1)} g_i g_j + \left(g_x - \sum_j s_{j1}g_j\right)\sum_i s_{i1}(\eta_i - 1)g_i + R^{(2)}.
\end{aligned}$$

By Engel aggregation, $\sum_i s_{i1}\eta_i = 1$, so

$$\begin{aligned}
\sum_i s_{i1}(\eta_i - 1)g_i =& \sum_i s_{i1}\eta_i g_i - \sum_i s_{i1}g_i \\
=& \mathrm{Cov}_s(g_i, \eta_i),
\end{aligned}$$

where $\mathrm{Cov}_s(\cdot, \cdot)$ denotes initial-period-share-weighted covariance. Finally, adding and subtracting

$$\log\left(\sum_i s_{i1}\exp(g_i)\right)$$

and dropping $R^{(2)}$, we obtain

$$\begin{aligned}\log P(\mathbf{p}_1, \mathbf{p}_T; u_T) \approx\ & \underbrace{\log\Big(\sum_i s_{i1}\exp(g_i)\Big)}_{\text{Exposure}} \\ & - \underbrace{\Big[\log\Big(\sum_i s_{i1}\exp(g_i)\Big) - \sum_i s_{i1} g_i - \tfrac{1}{2}\sum_{i,j} \tfrac{\partial s_i}{\partial \log p_j}\Big|_{(\mathbf{p}_1, u_1)} g_i g_j\Big]}_{\text{Substitution}} \\ & - \underbrace{\Big[-\Big(g_x - \sum_j s_{j1} g_j\Big)\mathrm{Cov}_s(g_i, \eta_i)\Big]}_{\text{Non-homothetic adjustment}}.\end{aligned}$$

This is precisely the decomposition stated in Proposition 1. $\square$

Remark 1. *Using the exact identity*

$$\log P(\mathbf{p}_1, \mathbf{p}_T; u_T) = \log P(\mathbf{p}_1, \mathbf{p}_T; u_1) - \big[\log P(\mathbf{p}_1, \mathbf{p}_T; u_1) - \log P(\mathbf{p}_1, \mathbf{p}_T; u_T)\big],$$

and comparing with the decomposition above, we see that, to second order,

$$\begin{aligned}\log P(\mathbf{p}_1, \mathbf{p}_T; u_1) &\approx \text{Exposure} - \text{Substitution}, \\ \log P(\mathbf{p}_1, \mathbf{p}_T; u_1) - \log P(\mathbf{p}_1, \mathbf{p}_T; u_T) &\approx \text{Non-homothetic adjustment}.\end{aligned}$$

Substituting the explicit formulas for these components yields the familiar second-order approximation to the cost-of-living index at u_1:

$$\log P(\mathbf{p}_1, \mathbf{p}_T; u_1) \approx \sum_i s_{i1} g_i + \frac{1}{2}\sum_{i,j} \frac{\partial s_i}{\partial \log p_j}\bigg|_{(\mathbf{p}_1, u_1)} g_i g_j,$$

and the corresponding approximation for the difference between cost-of-living indexes at u_1 and u_T:

$$\log P(\mathbf{p}_1, \mathbf{p}_T; u_1) - \log P(\mathbf{p}_1, \mathbf{p}_T; u_T) \approx -\Big(g_x - \sum_j s_{j1} g_j\Big)\mathrm{Cov}_s(g_i, \eta_i).$$

B.2 Jaravel–Lashkari algorithm

Here we outline the algorithm developed in Jaravel and Lashkari (2024) for approximating cost-of-living indexes under general non-homothetic preferences. The algorithm allows for observable preference heterogeneity through household controls. Conditional on these controls, the function mapping money-metric utility into the non-homothetic correction is assumed to be common across households. For notational simplicity, we suppress household subscripts until the implementation paragraph.

Utility cardinalisation. Let $\mathbb{u} = \mathbb{v}(\mathbf{p}, x)$ and $x = \mathbb{e}(\mathbf{p}, \mathbb{u})$ denote the household's indirect utility and expenditure functions for an arbitrary cardinalisation of utility. Define the period-1-price-denominated money-metric utility function $u = v(\mathbf{p}, x) \equiv \mathbb{e}(\mathbf{p}_1, \mathbb{v}(\mathbf{p}, x))$ so that u is measured in units of period-1 expenditure, and let the corresponding expenditure function be $x = e(\mathbf{p}, u) = \mathbb{e}(\mathbf{p}, \mathbb{v}(\mathbf{p}_1, u))$.

First-order algorithm. The household's log expenditure growth between period t and $t-1$ can be decomposed as:

$$\log x_t - \log x_{t-1} = [\log e(\mathbf{p}_t, u_{t-1}) - \log e(\mathbf{p}_{t-1}, u_{t-1})] + [\log e(\mathbf{p}_t, u_t) - \log e(\mathbf{p}_t, u_{t-1})]$$

Taking a first-order approximation of the two brackets, (i) in prices around $(\mathbf{p}_{t-1}, u_{t-1})$ and (ii) in utility around $(\mathbf{p}_t, u_{t-1})$ yields

$$\log x_t - \log x_{t-1} \approx \sum_i s_{it-1}(\log p_{it} - \log p_{it-1}) + \frac{\partial \log e(\mathbf{p}_t, u_{t-1})}{\partial \log u}(\log u_t - \log u_{t-1})$$

where s_{it-1} is the expenditure share of good i in period $t-1$.

Let $P(\mathbf{p}_1, \mathbf{p}_t; u)$ denote the true cost-of-living index between periods 1 and t evaluated at utility u. Using

$$\log P(\mathbf{p}_1, \mathbf{p}_t; u) = \log e(\mathbf{p}_t, u) - \log e(\mathbf{p}_1, u), \qquad e(\mathbf{p}_1, u) = u,$$

we obtain

$$\frac{\partial \log P(\mathbf{p}_1, \mathbf{p}_t; u)}{\partial \log u} = \frac{\partial \log e(\mathbf{p}_t, u)}{\partial \log u} - 1.$$

Define the non-homotheticity correction

$$\Lambda_t(u) \equiv \frac{\partial \log P(\mathbf{p}_1, \mathbf{p}_t; u)}{\partial \log u} = \frac{\partial \log e(\mathbf{p}_t, u)}{\partial \log u} - 1.$$

Denote the log geometric-Laspeyres index between $t-1$ and t by

$$\log \Pi^{GL}_{(t-1,t)} \equiv \sum_i s_{it-1}(\log p_{it} - \log p_{it-1}).$$

Combining the expressions above, log money-metric utility follows the recursion

$$\log u_t = \log u_{t-1} + \frac{\log x_t - \log x_{t-1} - \log \Pi^{GL}_{(t-1,t)}}{1 + \Lambda_t(u_{t-1})},$$

with initial condition $\log u_1 = \log x_1$.

In practice, $\Lambda_t(\cdot)$ is recovered flexibly from cross-sectional variation. For each t we run a flexible regression of the household-specific geometric-Laspeyres index on lagged money-metric utility (and household composition controls),

$$\log \Pi^{GL}_{h(t-1,t)} = f_t(\log u_{ht-1}) + \mathbf{X}'_h \beta_t + \varepsilon_{ht},$$

and use the cumulative derivatives of the fitted function to approximate

$$\Lambda_t(u_{t-1}) \approx \sum_{\tau=2}^{t} f'_\tau(\log u_{t-1}).$$

This first-order algorithm recovers money-metric utility by adjusting chained geometric-Laspeyres indexes for non-homotheticity.

Second-order algorithm. To obtain a second-order approximation, we symmetrise the decomposition in two ways. First,

$$\begin{aligned}\log x_t - \log x_{t-1} &= \big[\log e(\mathbf{p}_t, u_{t-1}) - \log e(\mathbf{p}_{t-1}, u_{t-1})\big] + \big[\log e(\mathbf{p}_t, u_t) - \log e(\mathbf{p}_t, u_{t-1})\big] \\ &\approx \sum_i s_{it-1}(\log p_{it} - \log p_{it-1}) + \big(1 + \Lambda_t(u_t)\big)(\log u_t - \log u_{t-1}).\end{aligned}$$

where the two first-order approximations are taken, respectively, in prices around $(\mathbf{p}_{t-1}, u_{t-1})$ and in utility around $(\mathbf{p}_t, u_t)$. Second,

$$\begin{aligned}\log x_t - \log x_{t-1} &= \big[\log e(\mathbf{p}_t, u_t) - \log e(\mathbf{p}_{t-1}, u_t)\big] + \big[\log e(\mathbf{p}_{t-1}, u_t) - \log e(\mathbf{p}_{t-1}, u_{t-1})\big] \\ &\approx \sum_i s_{it}(\log p_{it} - \log p_{it-1}) + \big(1 + \Lambda_{t-1}(u_{t-1})\big)(\log u_t - \log u_{t-1}),\end{aligned}$$

where the two first-order approximations are taken, respectively, in prices around $(\mathbf{p}_t, u_t)$ and in utility around $(\mathbf{p}_{t-1}, u_{t-1})$.

Taking the arithmetic average of these two expressions gives

$$\begin{aligned}\log x_t - \log x_{t-1} \approx\ & \frac{1}{2}\sum_i (s_{it-1} + s_{it})(\log p_{it} - \log p_{it-1}) \\ & + \Big(1 + \tfrac{1}{2}\big[\Lambda_{t-1}(u_{t-1}) + \Lambda_t(u_t)\big]\Big)(\log u_t - \log u_{t-1}).\end{aligned}$$

The first term is the log Törnqvist index between $t-1$ and t

$$\log \Pi^T_{(t-1,t)} \equiv \frac{1}{2}\sum_i (s_{it-1} + s_{it})(\log p_{it} - \log p_{it-1}).$$

Solving for $\log u_t$ yields the second-order recursion

$$\log u_t = \log u_{t-1} + \frac{\log x_t - \log x_{t-1} - \log \Pi^T_{(t-1,t)}}{1 + \frac{1}{2}\big[\Lambda_{t-1}(u_{t-1}) + \Lambda_t(u_t)\big]}.$$

Since $\Lambda_t(\cdot)$ is evaluated at the unknown u_t, this recursion defines a fixed-point problem. In the empirical implementation, we follow the iterative second-order procedure in Algorithm A.7 of Jaravel and Lashkari (2024), which uses cross-sectional price-index variation to update the non-homothetic correction at each iteration.

Cost-of-living indexes at u_T and u_1. Given the sequence $\{u_t\}_{t=1}^{T}$, the log cost-of-living index evaluated at final-period utility is

$$\log P(\mathbf{p}_1, \mathbf{p}_T; u_T) = \log \frac{e(\mathbf{p}_T, u_T)}{e(\mathbf{p}_1, u_T)} = \log x_T - \log u_T = \log \mathbb{e}(\mathbf{p}_T, \mathbb{u}_T) - \log \mathbb{e}(\mathbf{p}_1, \mathbb{u}_T).$$

The same logic can be applied in reverse, denominating utility in period-T prices and running the algorithm backward from T to 1. This produces the reverse cost-of-living index, $P(\mathbf{p}_T, \mathbf{p}_1; u_1) = \frac{e(\mathbf{p}_1, u_1)}{e(\mathbf{p}_T, u_1)}$, whose inverse is the cost-of-living index evaluated at initial-period utility, $P(\mathbf{p}_1, \mathbf{p}_T; u_1)$. In the empirical implementation, we use the forward recursion to construct $\Pi_{1,T}^{JL}(u_T)$ and invert the reverse price index produced by the backward recursion to construct $\Pi_{1,T}^{JL}(u_1)$.

Implementation with scanner data. We implement the second-order algorithm using household-level expenditure shares s_{hit}, constructed from UPC expenditures aggregated to our product definition, and product-level prices p_{it} defined as unit values. For each adjacent pair of periods, the price indexes are computed over the set of products with observed prices in both periods, the overlap set $\Omega_{t-1,t}$. Households that do not purchase a product have zero expenditure share for that item. For each link, expenditure shares are renormalised over the overlap set $\Omega_{t-1,t}$. We implement the cross-sectional functions in Algorithm A.7 of Jaravel and Lashkari (2024) using polynomial regressions and solve the resulting fixed-point problem iteratively. We include household composition controls $\mathbf{X}_h$ in these regressions. Estimation of the non-homothetic correction is based on the maintained assumption that, conditional on $\mathbf{X}_h$, households share the same function linking money-metric utility to the non-homothetic correction.

C Results appendix

C.1 Inflation exposure with expenditure-quartile-specific prices

Appendix A.1 shows that differences in the prices paid for identical products across households with different levels of resources are modest and change little over time. Within-product price dispersion is therefore unlikely to be an important driver of inflation inequality over 2021Q3–2023Q3.

Figure C.1 confirms this by comparing cumulative household inflation exposure based on prices that are common across households—the blue line, which replicates Figure 4.2(a)—with an alternative series that recomputes product prices separately within expenditure quartiles.[18] The gradient is slightly steeper when using expenditure-quartile-specific prices, implying that differences in the prices paid for identical goods slightly reinforce, rather than attenuate, the substantial inflation inequality documented in the main text.

[18] Specifically, we compute the expenditure-quartile-specific price for product i in period t as $p_{it}^{r} = \frac{\sum_{h \in \mathcal{R}_r} x_{hit}}{\sum_{h \in \mathcal{R}_r} q_{hit}}$ where $\mathcal{R}_r$ denotes the set of households belonging to quartile r of the calendar-year-specific annual equivalised expenditure distribution.

Figure C.1: *Inflation inequality with expenditure-quartile-specific prices*

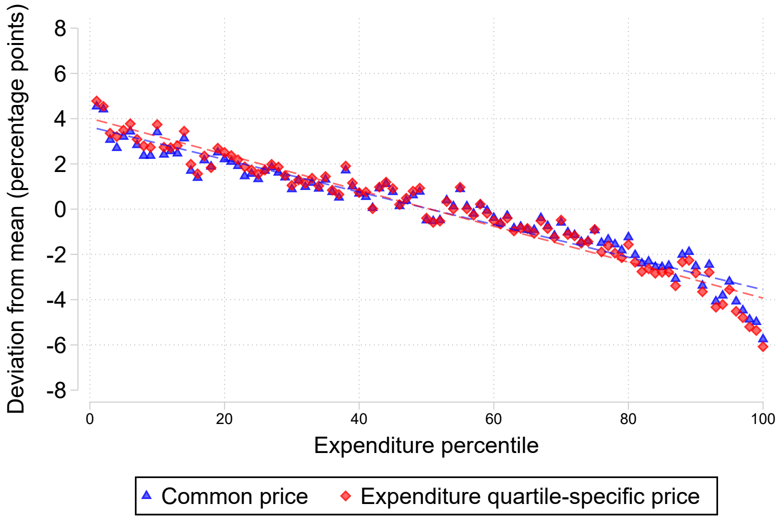

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2021–2023). The figure plots the relationship between cumulative inflation exposure over 2021Q3–2023Q3 and a household's percentile in the 2021 expenditure distribution, with a marker for each percentile and fitted linear relationships. Cumulative inflation exposure is measured using household-specific Laspeyres indexes. The blue series uses prices that are common across households; the red series uses expenditure-quartile-specific prices as defined in the text.

C.2 Product churn

The Laspeyres index in equation (3.2) fixes each household's consumption basket at its initial composition and therefore does not capture the effects of product entry and exit. Figure C.2 shows inflation inequality computed using two chained variants of the Laspeyres index. Panel (a) chains the index annually, comparing 2021Q3 with 2022Q3 and 2022Q3 with 2023Q3. Panel (b) chains the index quarter by quarter over the nine quarters from 2021Q3 to 2023Q3. Because these indexes compare prices over shorter adjacent periods, they reduce the influence of product churn. Both exhibit essentially the same pattern of inflation inequality as in Figure 4.2(a).

Chained indexes do not fully address product churn, because there can be substantial product entry and exit even between adjacent quarters. This can generate new-product bias: conventional indexes do not capture the welfare losses associated with disappearing products or the gains associated with newly available products. In our data, 7.1% of aggregate spending in 2021Q3 is on products that are unavailable in 2022Q3, while 8.6% of aggregate spending in 2022Q3 is on products that are unavailable in 2021Q3.[19] Between 2022Q3 and 2023Q3, the corresponding shares are 5.1% for exiting products and 6.5% for entering products.

[19]We define a product as unavailable in a quarter if no household in our entire sample purchases it.

Figure C.2: *Inflation inequality: chained Laspeyres index*

(a) Annual chaining (b) Quarter-to-quarter chaining

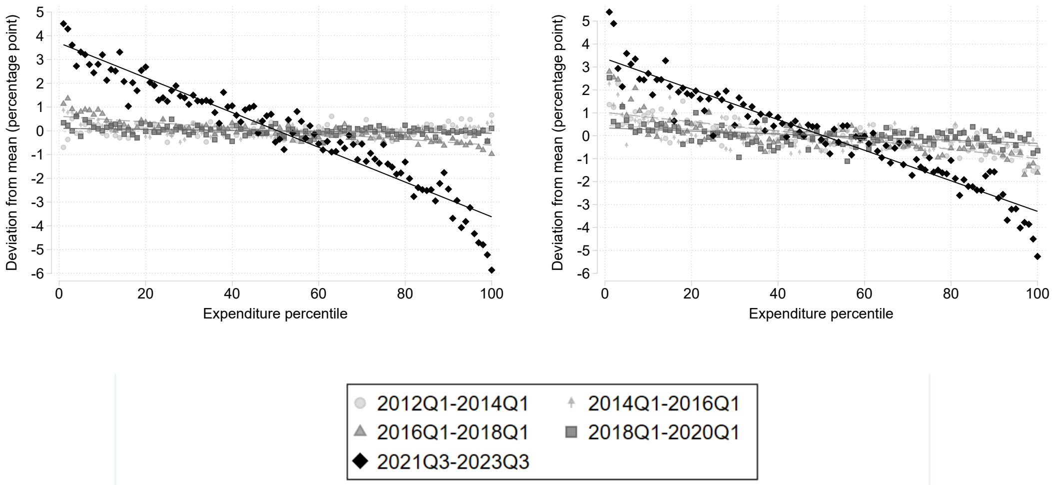

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2012–2023). The figure plots the relationship between cumulative inflation exposure and a household's percentile in the equivalised expenditure distribution, with a marker for each percentile and a fitted linear relationship. Households are assigned to expenditure percentiles based on their equivalised spending in the calendar year in which the relevant nine-quarter period begins. Panel (a) uses a Laspeyres index chained across the initial, fifth, and ninth quarters of each period; for 2021Q3–2023Q3, these are 2021Q3, 2022Q3, and 2023Q3. Panel (b) uses a Laspeyres index chained across adjacent quarters.

The Feenstra-corrected CES price index (Feenstra 1994) provides a convenient way to quantify the importance of product entry and exit. Under CES preferences, the household's cost-of-living index between periods s and t can be written as $P^{CES}_{h,(s,t)} = \left[\prod_{i\in\Omega_{h,(s,t)}} \left(\frac{p_{it}}{p_{is}}\right)^{\phi_{hi,(s,t)}}\right] \times \mathcal{F}_{h,(s,t)}$, where $\Omega_{h,(s,t)}$ is the set of comparison products and $\phi_{hi,(s,t)}$ is a weight determined by the household's period-s and period-t expenditure shares.[20] The variety correction is $\mathcal{F}_{h,(s,t)} = \left(\frac{1-s^N_{ht}}{1-s^X_{hs}}\right)^{\frac{1}{\sigma-1}}$, where s^X_{hs} is the household's initial-period expenditure share on products that subsequently exit, s^N_{ht} is its final-period expenditure share on products that entered, and $\sigma > 1$ is the elasticity of substitution. The correction increases with the share allocated to exiting products, because the disappearance of products raises the cost of living, and decreases with the share allocated to entering products, because access to new products lowers it. These variety effects are larger when σ is closer to one, because entering and exiting products have fewer close substitutes.

Figure C.3(a) plots the inflation-inequality gradient using a CES price index without the Feenstra variety correction. The pattern is very similar to that in Figure 4.2(a). Panel (b) plots net entry, measured as the difference between the household's expenditure shares on entering and exiting products, $s^N_{ht} - s^X_{hs}$. Panel (c) reports the adjustment to the CES index implied by product entry and exit—the difference between the Feenstra-corrected and non-corrected CES indexes—for several values of σ. Each panel reports averages by expenditure percentile. On average, households allocate a larger share of expenditure to

[20]Specifically, $\phi_{hi,(s,t)} = \frac{(s_{hit}-s_{his})/(\log s_{hit} - \log s_{his})}{\sum_{i'}(s_{hi't}-s_{hi's})/(\log s_{hi't} - \log s_{hi's})}$.

entering products than to exiting products, lowering measured cost-of-living growth. The magnitude of the effect, however, varies little across the expenditure distribution.

Figure C.3: *CES cost-of-living index*

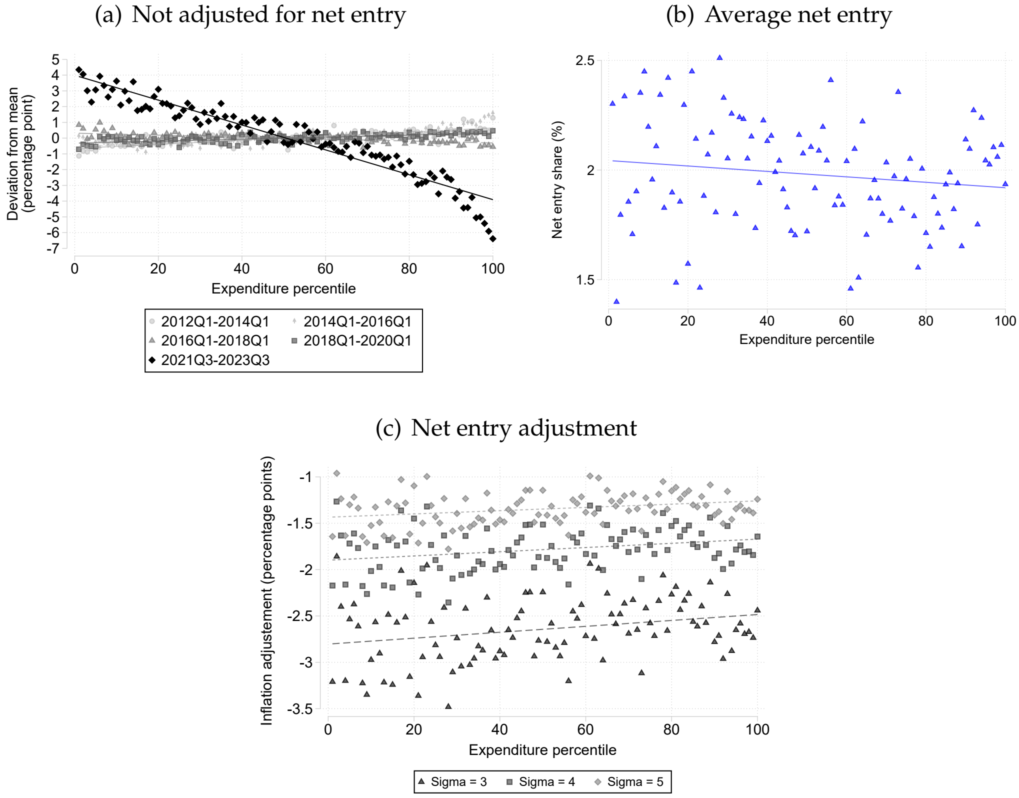

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2021–2023). Panel (a) plots cumulative inflation over 2021Q3–2023Q3, as well as each earlier 9-quarter period, measured using a chained CES price index without the Feenstra variety correction, by expenditure percentile. Panel (b) plots average net entry, $s^N_{ht} - s^X_{hs}$, by expenditure percentile, averaging over the 2021Q3–2022Q3 and 2022Q3–2023Q3 comparisons. Panel (c) plots the difference between the Feenstra-corrected and uncorrected CES indexes for three values of the elasticity of substitution σ. Households are assigned expenditure percentiles based on their equivalised expenditure in 2021. All indexes are chained across 2021Q3, 2022Q3 and 2023Q3.

C.3 Inflation inequality in the broader consumption basket

This appendix describes how we construct income-decile-specific price indexes for the broader consumption basket by combining LCFS expenditure data, official Consumer Prices Index (CPI) category indexes, and our scanner-data-based FMCG inflation measures. We construct two alternative indexes: (i) a baseline index that uses *common* FMCG price relatives across income deciles, and (ii) an alternative index that incorporates *income-decile-specific* FMCG price relatives derived from the Numerator scanner data.

Household expenditure data come from the 2021 Living Costs and Food Survey (LCFS), with spending categories classified using the COICOP system.[21] Let $t \in \{1, 2, 3\}$ index 2021Q3, 2022Q3, and 2023Q3, respectively. For $t \in \{2, 3\}$, we refer to the change from $t-1$ to

[21] Each expenditure item is mapped to a COICOP code with structure $X.Y.Z$, where X is the division (1–12), Y is the group (1–9 within each division), and Z is the class (1–9 within each group).

t as annual link t. Official price indexes P_{ct}^{CPI} and the national CPI weights applicable in each period, w_{ct}^{CPI}, are taken from ONS CPI data at the most disaggregated COICOP level used in the analysis.[22] For FMCG categories, we supplement these with FMCG price relatives from the Numerator data, constructed either at the aggregate level (common across income deciles) or separately by income decile.

We classify as FMCG the COICOP categories listed in Table C.1.

Table C.1: *FMCG classification by COICOP category*

COICOP code	Category name
01.1.1	Bread and cereals
01.1.2	Meat
01.1.3	Fish
01.1.4	Milk, cheese and eggs
01.1.5	Oils and fats
01.1.6	Fruit
01.1.7	Vegetables
01.1.8	Sugar and sweet products
01.1.9	Food products n.e.c.
01.2.1	Coffee, tea and cocoa
01.2.2	Mineral water and soft drinks
02.1.1	Spirits
02.1.2	Wine
02.1.3	Beers
02.2	Tobacco
05.6.1	Non-durable household goods
06.1.1	Pharmaceutical products
06.1.2	Other medical products
09.3.4	Pets and related products
12.1.3	Other products for personal care

Income-decile-specific expenditure weights. Households are grouped into income deciles based on equivalised net household income, indexed by $g \in \{1, \ldots, 10\}$. For each income decile g and COICOP category c, total LCFS expenditure in 2021 is

$$X_{gc,2021} \equiv \sum_{i \in G_{g,2021}} v_i x_{ic,2021}$$

where $x_{ic,2021}$ is household i's expenditure on category c, $G_{g,2021}$ is the set of households in decile g in the 2021 data, and v_i is a grossing weight taken from the survey. The corresponding LCFS budget share is

$$w_{gc,2021}^{\text{LCFS}} \equiv \frac{X_{gc,2021}}{\sum_{c'} X_{gc',2021}}.$$

[22] The CPI weight w_{ct}^{CPI} is the weight applicable to the price change between periods $t-1$ and t. CPI expenditure weights are based largely on expenditure data from approximately two years earlier.

These LCFS shares do not exactly reproduce the national CPI weights because of survey under-reporting and differences in coverage and definitions between the LCFS and the national accounts. We therefore rescale LCFS expenditure within each COICOP category while preserving its distribution across income groups. Let

$$X_{c,2021} \equiv \sum_{g} X_{gc,2021}, \qquad X_{2021} \equiv \sum_{c} X_{c,2021}.$$

For each (g, c, t), $t \in \{2, 3\}$, we define adjusted expenditure for annual link t using the applicable CPI weight, w_{ct}^{CPI}:

$$x_{gct}^{*} \equiv w_{ct}^{\mathrm{CPI}} \times \frac{X_{gc,2021}}{X_{c,2021}} \times X_{2021}.$$

By construction, $\sum_g x_{gct}^{*} = w_{ct}^{\mathrm{CPI}} X_{2021}$ for every c, so the rescaled group-level expenditures aggregate to the official CPI weights within each category while preserving the LCFS distribution of expenditure on each category across income deciles.

We then define income-decile-specific CPI-consistent budget shares as

$$w_{gct}^{*} \equiv \frac{x_{gct}^{*}}{\sum_{c'} x_{gc't}^{*}}.$$

Variation in w_{gct}^{*} across links reflects changes in the official CPI weights, while the distribution of expenditure on each category across income deciles is held fixed at its 2021 LCFS value.

For later use, we also define the income-decile-specific FMCG share in period t as

$$w_{g\mathrm{FMCG}t}^{*} \equiv \sum_{c \in \mathrm{FMCG}} w_{gct}^{*}.$$

Baseline broader-basket index by income decile. Let P_{ct}^{CPI} denote the official CPI index for category c in period t. For non-FMCG categories $c \notin$ FMCG and $t \in \{2, 3\}$ we define the category-level CPI price change

$$p_{ct} \equiv \frac{P_{ct}^{\mathrm{CPI}}}{P_{ct-1}^{\mathrm{CPI}}}.$$

From the Numerator scanner data, we construct a *common* FMCG Laspeyres price relative p_t^{FMCG} between 2021Q3 ($t = 1$) and 2022Q3 ($t = 2$) and between 2022Q3 and 2023Q3 ($t = 3$), using aggregate FMCG expenditure weights across all households. This provides a single FMCG inflation measure that is the same for all income groups in each subperiod.

Baseline cumulative chained inflation for income decile g over 2021Q3–2023Q3 is

$$\Pi_g \equiv \left[\prod_{t=2}^{3} \left(w_{g\mathrm{FMCG}t}^{*}\, p_t^{\mathrm{FMCG}} + \sum_{c \notin \mathrm{FMCG}} w_{gct}^{*}\, p_{ct} \right) \right] - 1.$$

For each income decile $g \in \{1, \ldots, 10\}$, we report cumulative inflation relative to the bottom decile,

$$\tilde{\Pi}_g \equiv \Pi_g - \Pi_1.$$

By construction, $\tilde{\Pi}_1 = 0$. A negative value indicates lower cumulative inflation for decile g than for the bottom decile.

Incorporating income-decile-specific FMCG inflation. To isolate the role of within-FMCG heterogeneity, we replace the common FMCG inflation measure p_t^{FMCG} with income-decile-specific FMCG inflation measures. From the Numerator scanner data, we construct an income-decile-specific FMCG Laspeyres price relative $p_{g\text{FMCG}t}$ for the 2021Q3–2022Q3 and 2022Q3–2023Q3 links. We then form an alternative income-decile-specific index that incorporates these income-decile-specific FMCG price relatives:

$$\Pi_g^{\text{alt}} \equiv \prod_{t=2}^{3} \left(w^*_{g\text{FMCG}t}\, p_{g\text{FMCG}t} + \sum_{c \notin \text{FMCG}} w^*_{gct}\, p_{ct} \right) - 1.$$

The corresponding difference between decile g and the bottom decile is

$$\tilde{\Pi}_g^{\text{alt}} \equiv \Pi_g^{\text{alt}} - \Pi_1^{\text{alt}}.$$

Results. Figure C.4 summarises the contribution of FMCG cheapflation to inflation inequality across the broader consumption basket. Panel (a) shows that FMCG budget shares decline strongly with income, both in the raw LCFS data and after rescaling expenditures to match CPI weights. Using the raw LCFS shares, the bottom income decile allocates around 23.1% of total expenditure to FMCG, compared with approximately 16.6% in the top decile, a gap of 6.5 percentage points.

Panel (b) compares cumulative inflation over 2021Q3–2023Q3 across income deciles, expressed relative to the bottom decile, under the two index constructions. In the *common-FMCG benchmark*, FMCG categories have the same inflation rate across income groups, given by the aggregate FMCG Laspeyres price relatives p_t^{FMCG}, while non-FMCG categories follow class-level CPI inflation. Under this construction, cumulative inflation for the bottom decile exceeds that for the top decile by 3.16 percentage points. In the *FMCG-adjusted index*, we replace common FMCG inflation with income-decile-specific FMCG price relatives $p_{g\text{FMCG}t}$ constructed from the Numerator data. The corresponding bottom–top decile gap rises to 4.04 percentage points. Income-decile-specific FMCG inflation therefore accounts for 0.88 percentage points of the total gap, increasing it by 27.8% relative to the common-FMCG benchmark.

Panel (b) also shows that substantial inflation inequality remains in the common-FMCG benchmark, which abstracts from within-FMCG inflation differences across income deciles. Panel (c) repeats the exercise after excluding residential gas, electricity and other domestic fuels and renormalising the remaining expenditure weights within each income decile. The bottom–top decile gap in the common-FMCG benchmark then becomes much smaller. This indicates that much of the inflation inequality generated by broad category expenditure differences during this period reflects the exceptionally large increase in residential energy prices combined with higher energy budget shares among lower-income households.

Figure C.4: *Inflation inequality in the broader CPI basket*

(a) FMCG budget share

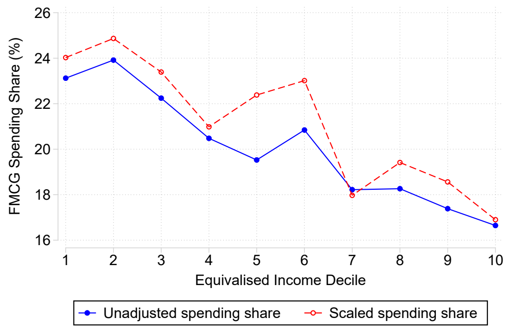

(b) Cumulative inflation relative to bottom decile

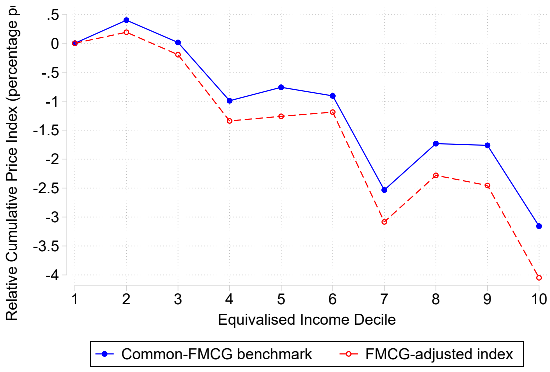

(c) Cumulative inflation relative to bottom decile (excluding energy)

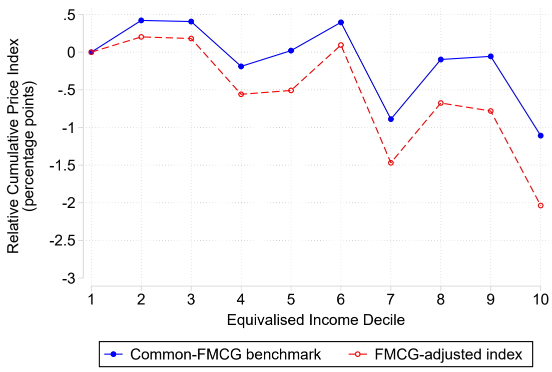

Notes: Authors' calculations using the Living Costs and Food Survey (LCFS) 2021, ONS CPI data, and Numerator's Take Home Purchase Panel (2021–2023). Panel (a) plots the share of total expenditure allocated to FMCG goods by equivalised income decile, before and after scaling spending to match weights used in the ONS CPI. Panel (b) reports cumulative inflation over 2021Q3–2023Q3 for (i) a common-FMCG benchmark that combines class-level CPI inflation for non-FMCG categories with common Numerator-based FMCG inflation across income deciles, and (ii) an FMCG-adjusted index that replaces the common FMCG rate with income-decile-specific Numerator-based FMCG price relatives; values are expressed relative to the bottom decile. Panel (c) repeats the analysis in panel (b) excluding residential gas and electricity and other domestic fuels, and renormalising the remaining expenditure weights within each income decile.