Measuring cost of living inequality during an inflation surge

Tao Chen, Peter Levell and Martin O'Connell*

February 2026

Abstract

We study the welfare cost of inflation 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 inflation-driven welfare changes into exposure, substitution, and income effects. We document pronounced "cheapflation": within narrowly defined categories, prices rose fastest for lower-quality necessities disproportionately consumed by poorer households. This generates sharply regressive inflation exposure that is only slightly mitigated by behavioural responses. As purchasing power fell, cheapflation interacted with non-homothetic demand movements down the quality ladder (captured by our income-effects term) to amplify welfare losses and reshape the distribution of exposure to future necessity-driven inflation.

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 has re-emerged as a central economic concern in the aftermath of the COVID-19 pandemic. Many advanced economies have experienced 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 amplifying 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 has affected inequality in both inflation exposure and welfare. 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 necessity products ("cheapflation"), disproportionately consumed by poorer households, and households—especially those whose real expenditure fell most—shifted down the quality ladder toward these same products. To quantify the welfare consequences of these price movements and household responses, we develop and implement an index-number-based welfare decomposition under general non-homothetic preferences. This decomposition maps welfare changes into exposure, substitution, and income effects 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 and that, as purchasing power fell, the interaction of cheapflation with non-homothetic demand adjustments amplified the welfare cost of inflation relative to homothetic benchmarks and reshaped the distribution of exposure to future necessity-driven inflation.[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 ladder: products in the bottom two rungs exhibit average price increases of around one-third,

[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. (2023), and Del Canto et al. (2024).

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-quality (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 low-quality goods. Cheapflation therefore raises inflation exposure disproportionately for worse-off households. We show that much of the resulting inflation inequality is missed by household inflation 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 whose deflated expenditure falls more trade down more along the quality ladder. Together, these facts highlight three central forces for inflation inequality within 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 the welfare effects of inflation that is suited to environments with pronounced within-category relative price dispersion and non-homothetic demand. Our theoretical object of interest is an equivalent-variation quality-of-living index, which expresses the welfare effect of the episode in initial-period prices: it is the proportional change in spending at initial prices required to attain the realised final-period utility level rather than the initial one.[2] We show that, up to a second-order approximation in nominal expenditure growth and price changes, this welfare change equals nominal expenditure growth minus the welfare cost of inflation, defined as the sum of exposure, substitution, and income effects.

The first, *exposure*, term is the log Laspeyres cost-of-living index based on the initial consumption bundle, capturing mechanical exposure to inflation under Leontief (no-substitution) preferences.[3] The second, *substitution*, term splits into a Cobb–Douglas component (the gap between the Laspeyres and geometric-Laspeyres indexes) and a second-order term involving Hicksian responses of budget shares to prices; together these capture how substitution away from goods with strong *relative* price increases mitigates cost-of-living

[2]Equivalently, prior to the inflation surge, the index answers what fraction of initial-period expenditure a household would be willing to forgo to avoid the realised changes in prices and budgets over the episode.

[3]The cost-of-living index is the proportional change in expenditure required to attain a fixed utility level at final-period prices rather than initial-period prices. Under Leontief preferences, this object coincides with the Laspeyres index based on the initial bundle.

growth. The final *income-effects* term arises from non-homothetic demand and is equal to minus the product of a "real expenditure" term (negative when purchasing power falls) and a covariance between relative price changes and expenditure elasticities. In a cheapflation episode where goods with low expenditure elasticities (often called income elasticities) become relatively more expensive, falling real expenditure pushes households along Engel curves toward precisely those goods whose prices are rising, generating additional welfare losses beyond those captured by homothetic benchmarks.

A key advantage of our framework is that each component can be expressed as differences between observable household-level price indexes. Building on Jaravel and Lashkari (2024), we implement a non-homothetic money-metric index that provides a second-order approximation to the cost-of-living index at a given utility level. We show how the terms for exposure, Cobb–Douglas substitution, second-order substitution, and income effects correspond to log differences between a Laspeyres index, a geometric-Laspeyres index, and non-homothetic cost-of-living indexes evaluated at the initial and final utility levels. This delivers an index-number-based implementation of the welfare decomposition that is valid under general non-homothetic preferences and 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 income effects to inflation 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 virtually all 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 25.2%. The Cobb–Douglas term alone would lower it by approximately 1.9 percentage points relative to the Laspeyres benchmark, but richer substitution patterns—captured by the second-order term—raise it by around 0.9 percentage points. This suggests that products experiencing the strongest price growth are precisely those that households are least willing or able to substitute away from, limiting the scope for substitution in mitigating inflation-driven welfare losses. Across the expenditure distribution, substitution modestly dampens the cost-of-living gradient (the bottom–top inter–decile range falls from about 7.5 to 7.1 percentage points), but the gradient remains large and regressive.

We then measure the non-homothetic (income-effects) component—i.e., the difference between non-homothetic cost-of-living indexes evaluated at the initial (2021Q3) and final (2023Q3) utility levels. On average, the cost-of-living index evaluated at final utility is about 0.9 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 low-quality necessity goods whose relative prices are increasing, making it more expensive to attain the realised final utility level than the initial one.

Overall, richer substitution patterns and income effects undo much of the headline cost savings implied by a simple Cobb–Douglas benchmark. As a result, the sharp inflation exposure gradient across both the expenditure and income distributions translates into similarly regressive gradients in the welfare costs of the inflation episode.

Finally, we ask whether cheapflation and trading down have also reshaped households' exposure to *future* inflation episodes. We construct a "repeat-cheapflation exposure" index that applies the observed pattern of relative price changes to each household's final-period consumption bundle—a forward-looking Laspeyres index capturing exposure to a future bout of cheapflation. We find that households with the most negative income effects during 2021Q3–2023Q3 end the episode with consumption baskets that are more exposed to a repeat cheapflation shock than under their initial baskets. At the same time, these households were not initially the most exposed to cheapflation; rather, their Engel-curve adjustments under cheapflation drew them toward goods where future necessity-driven inflation would hit harder. This intertemporal reconfiguration of exposure suggests that the distributional consequences of cheapflation may persist 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, 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 2023, 2025; Becker 2025), by showing how cheapflation in lower-quality necessities and falling purchasing power interact to shape both realised welfare losses and future exposure. Third, we contribute to the literature on non-homothetic demand and measurement of living standards and welfare approximations (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 to an equivalent-variation quality-of-living index, in the spirit of Baqaee and Burstein (2023), and show how the three price-based components of the welfare cost of inflation—exposure, substitution, and income effects—can be implemented as differences between household-level price indexes under general non-homothetic preferences. Our contribution is to provide a non-homothetic, index-number-based decomposition of inflation-driven welfare changes that can be implemented in scanner data, and to show empirically that within-category quality choices, together with trading down the quality ladder as purchasing power falls (the adjustment margin captured by the income-effects term in our framework), are central both to the regressive welfare impact of the 2021–2023 surge and to shaping households' exposure to future cheapflation episodes.

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 quality-of-living decomposition and maps its components to observable price indexes. Section 4 presents our empirical decomposition of inflation-driven welfare changes across the expenditure and income distributions. Section 5 studies how cheapflation and trading down reshape households' exposure to future inflation episodes. 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 foods—brought into the home by a sample of households living in Great Britain (i.e., the UK excluding Northern Ireland). It includes purchases made both in brick-and-mortar stores and online.

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 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—together with alcohol and household goods, key components of fast-moving consumer goods—rose from 103.3 to 133.5 after a prolonged period of stability.[4] Our main analysis sample includes 19,030 households that recorded their purchases in every year-quarter over this period.[5] We compare inflation inequality over 2021Q3–2023Q3 to four earlier nine-quarter periods: 2012Q1–2014Q1, 2014Q1–2016Q1, 2016Q1–2018Q1, and 2018Q1–2020Q1.[6]

Our dataset offers several advantages for measuring inflation inequality compared to 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 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.[7]

Our data are more detailed than the budget surveys and price microdata often used to study household-level inflation inequality (see Office for National Statistics (2022) for the UK and Klick and Stockburger (2021, 2024) for the US). These survey-based measures cover a broader share of total spending but aggregate prices to around 100–200 item groups. By contrast, our scanner data 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. Since some UPCs are occasionally replaced with nearly identical ones, we define products based on

[4]In January 2019, the index was 102.6, and by April 2024 it had reached 135.6. 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).

[5]We exclude households that are not continuously present across all nine quarters, or whose quarterly expenditure falls below the 5^{th} percentile (on average £114) of the expenditure distribution. Our results are not sensitive to these restrictions.

[6]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.

[7]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 for studying the distribution of inflation across households.

the more aggregated combination of brand and package size. Numerator provides highly disaggregated brand information, so this approach involves minimal aggregation over meaningfully distinct products. Over this period, there are approximately 90,000 brand–size pairs, which we refer to as products throughout.[8]

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,[9]

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

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 minor relative to the overall inflation inequality we document below, and acts, albeit modestly, to reinforce 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 (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. 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. However, since some price variation arises from nonlinear pricing across different pack sizes of the same brand (e.g., see Griffith et al. 2009), we adjust for this form of price variation before constructing the quality ladder. We define rungs using adjusted unit-price levels from early in the episode and then hold rung assignments fixed when measuring inflation over 2021Q3–2023Q3, so "cheapflation" reflects differential price growth across a predetermined within-category ranking.

For each product category, using data from 2021Q3–2023Q3 to estimate the pack-size adjustment and to assign rungs to products introduced later in the episode, we estimate the

[8] 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.

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

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, $\tau_{c(i)t}$ are year-quarter fixed effects, $\xi_{b(i)}$ are brand effects, and the α-terms represent a third-order polynomial in demeaned pack size, with coefficients that vary across years, indexed by y.

To assign products to quality-ladder rungs, we first consider the set of products available in the four quarters 2021Q3–2022Q2 (which we index $t = 1, \ldots, 4$) and compute their adjusted prices in each of these quarters, netting out the size polynomial: $\tilde{p}_{it} = p_{it} - \sum_{l=1}^{3} \alpha_{c(i)y}^{(l)} \mathbb{1}\{t \in y\} \times \text{size}_i^{(l)}, \quad t \leq 4$. We then average adjusted product-level prices across these quarters to form a baseline adjusted price for each product and use the expenditure-weighted distribution within each product category to define decile boundaries. For all products, we then use equation (2.2) to predict their adjusted price, averaging over the first four quarters: $\tilde{p}_i = \xi_{b(i)} + \frac{1}{4} \sum_{t=1}^{4} \tau_{c(i)t}$. Using these predicted prices, we assign products to quality-ladder rungs based on the decile boundaries. This procedure ensures that products introduced after 2022Q2 are assigned to a quality rung based on an 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).[10]

In each nine-quarter period, we use equivalised expenditure from the first calendar year to assign households to expenditure percentiles (100 bins) or deciles (10 bins), and we refer to these as expenditure percentiles (or deciles). We also group households into deciles of equivalised income, using the midpoint of each income band 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.

[10]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.

Descriptive evidence

Our objective is to measure inflation inequality during an inflationary surge, and to characterise and quantify the importance of behavioural responses. Figure 2.1 provides preliminary evidence that motivates our approach. It shows that lower-quality products experience more rapid price growth during the inflation surge, that these products are more popular among households with fewer resources, and that households with the largest declines in spending shift more sharply down the quality ladder. This points toward the joint role of price dynamics across the quality ladder and spending adjustments driven by changes in purchasing power 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 lower-quality products than for those at the top of the ladder. For instance, the average price increase for products on the bottom two rungs (1 and 2) is 36.2% and 32.3%, respectively, while for those on the top two rungs (9 and 10), it is 20.7% and 15.8%. Over this period, the within-category price distribution compresses, and 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 growth is relatively flat 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-quality 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 goods that became relatively more expensive are those disproportionately consumed by worse-off households.

Panels (c) and (d) focus on 2021Q3–2023Q3. Panel (c) shows that the share of first-quarter expenditure allocated to bottom-quality-rung products decreases sharply across expenditure percentiles, whereas the spending share on top-quality-rung products increases across percentiles. This pattern is consistent with economic intuition: low-quality products are necessities, while high-quality products are luxuries.

Panel (d) illustrates the relationship between *within*-household changes in expenditure and changes in spending shares on top- and bottom-quality-rung products between 2021Q3 and 2023Q3. On average, households increase their spending shares on both top- and bottom-quality-rung products, but panel (d) shows that households with the largest declines in deflated expenditure shift most strongly toward low-quality products and least strongly toward high-quality products. This pattern is suggestive of households moving along quality-ladder Engel curves in response to changes in living standards.

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

(a) Price changes across quality rungs (b) Average quality rung, by expenditure

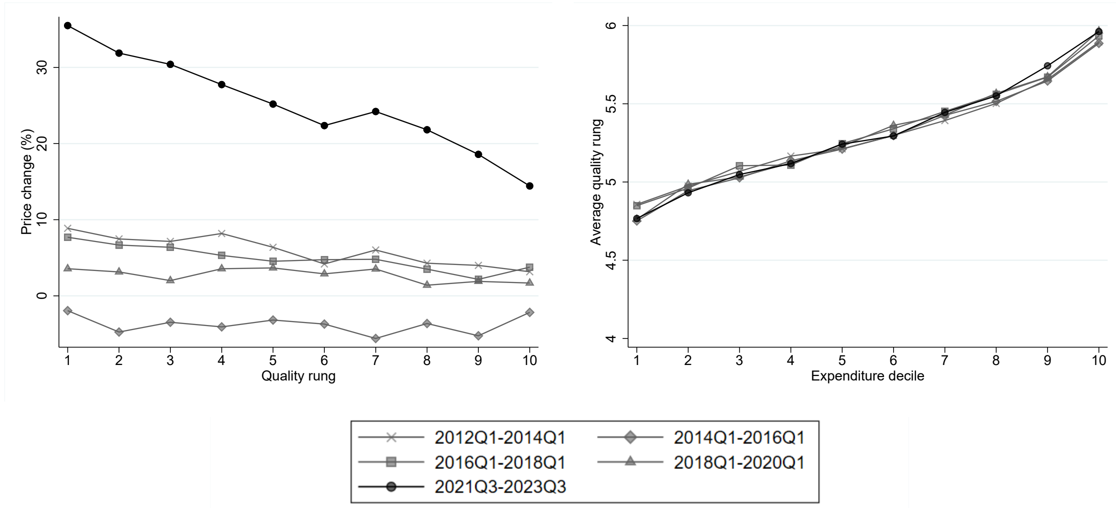

Bottom and top quality rung Engel curve:

(c) cross-sectional (d) within-household

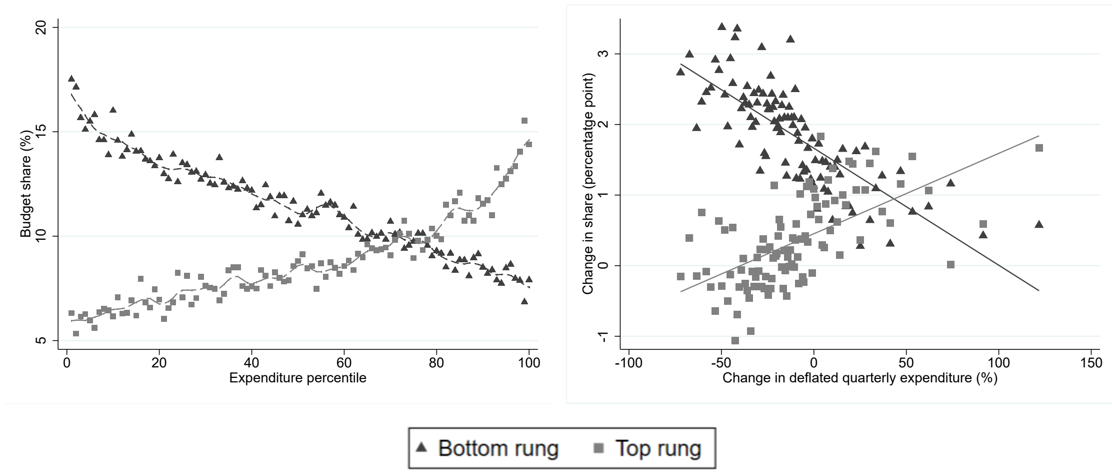

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2012–2023). Panel (a) reports change in average price over the nine-quarter period for products on each rung of the quality ladder, weighted by initial-period aggregate spending shares. Panel (b) reports the average quality rung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

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 reflect genuine Engel-curve movements, 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 2.1 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 purge 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 affect this instrument. The 2SLS estimate in column (2) remains highly significant and only modestly smaller in magnitude than the OLS coefficient, indicating that the negative association between deflated-expenditure declines and quality downgrading primarily reflects genuine behavioural responses along Engel curves rather than an artefact of how deflated expenditure is constructed.

Table 2.1: *Within segment trading-down*

Dependent Variable = log change in segment-level average quality		
	OLS	2SLS
log change in deflated expenditure	0.063***	0.047***
	(0.004)	(0.006)
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 volume-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. A positive coefficient implies that households with larger declines in deflated expenditure (more negative* $\log(x_{hT}/x_{h1}) - \log \Pi^L_{h(1,T)}$*) experience larger quality downgrades (more negative* $\log(Q_{hgT}/Q_{hg1})$*).* $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Discussion. The descriptive evidence in Figure 2.1 and Table 2.1 highlights three features of the 2021Q3–2023Q3 inflation episode that are central to our welfare analysis. First, price growth is sharply tilted toward the bottom of the within-category quality ladder: cheap varieties experienced systematically higher inflation, in line with the "cheapflation" documented using price data by Cavallo and Kryvtsov (2024). Second, households with fewer resources are systematically more exposed to these low-quality products that see the strongest price growth even before the episode begins, 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 trade down along the quality ladder *toward* those same low-quality, necessity-type goods whose prices rise most.

Taken together, these patterns imply that the distributional consequences of the inflation surge cannot be understood solely through household-specific Laspeyres indexes or through homothetic price indexes. In a cheapflation environment with falling purchasing power, homothetic cost-of-living indexes conflate two distinct responses: substitution away from goods whose relative prices are rising most quickly, and movements down the quality ladder toward necessities bearing the brunt of price increases. The interaction of cheapflation with non-homothetic demand—manifested in exposure differences across households and in within-household quality downgrading during the episode—creates an additional, potentially regressive channel through which price changes affect welfare. This combination of within-category price changes, cross-sectional exposure, and within-household trading down motivates the non-homothetic welfare decomposition developed in the next section, together with an index-number implementation that disentangles exposure, substitution, and income-effect components of the welfare costs of inflation.

3 Theory and measurement

In this section, we develop a framework for measuring inflation-driven welfare changes under general non-homothetic preferences. We show that these welfare changes can be decomposed into three components—initial exposure, substitution, and income effects—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- and quality-of-living indexes

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 increasing, homogeneous of degree one and concave in $\mathbf{p}$, and strictly increasing in u.

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

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

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

The quality-of-living index measures the change in the cost of attaining realised utility when prices are held fixed at a reference vector $\mathbf{p}$:

$$Q(u_1, u_T; \mathbf{p}) = \frac{e(\mathbf{p}, u_T)}{e(\mathbf{p}, u_1)},$$

where $u_b = v(\mathbf{p}_b, x_b)$ is the realised utility in period $b \in \{1, T\}$. Cost-of-living and quality-of-living indexes are linked to nominal expenditure growth by

$$P(\mathbf{p}_1, \mathbf{p}_T; u_b) \times Q(u_1, u_T; \mathbf{p}_{b'}) = \frac{x_T}{x_1}, \qquad b, b' \in \{1, T\},\ b \neq b'.$$

3.2 A decomposition of inflation-driven welfare changes

We focus on the quality-of-living index between the initial period ($t = 1$) and the final period ($t = T$), evaluated at the initial price vector. This index measures the proportional change in expenditure needed at the initial prices to attain the realised utility in the final period relative to the initial period:

$$\log Q(u_1, u_T; \mathbf{p}_1) = \log e(\mathbf{p}_1, u_T) - \log e(\mathbf{p}_1, u_1).$$

We show that this welfare change can be decomposed into nominal expenditure growth adjusted for three forces: (i) mechanical exposure to inflation given the initial consumption bundle; (ii) substitution responses to relative price changes; and (iii) non-homothetic (income-effects) adjustments in spending that change the cost of maintaining a fixed living standard.

Proposition 1 (Decomposition of the quality-of-living index)**.** *Let $t = 1, \dots, T$ index time and $i = 1, \dots, I$ goods. Let $\mathbf{p}_t = (p_{1t}, \dots, p_{It})'$ 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 $\Delta p_i = p_{iT}/p_{i1}$ and $\Delta x = x_T/x_1$, 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 $\eta_i = \big(\partial \log q_i(\mathbf{p}, x)/\partial \log x\big)\big|_{(\mathbf{p}_1, x_1)}$ 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 on $U(\cdot)$, a second-order Taylor expansion of $\log Q(u_1, u_T; \mathbf{p}_1)$ around $(\mathbf{p}_1, x_1)$ yields

$$\log Q(u_1, u_T; \mathbf{p}_1) \;\approx\; \log \Delta x \;-\; (\text{Exposure} \;-\; \text{Substitution} \;-\; \text{Income effects})\,, \tag{3.1}$$

where

$$\text{Exposure} = \log\Big(\sum_i s_{i1}\Delta p_i\Big),$$

$$\text{Substitution} = \underbrace{\left(\log\Big(\sum_i s_{i1}\Delta p_i\Big) - \sum_i s_{i1}\log \Delta p_i\right)}_{\text{Cobb–Douglas substitution term}} + \underbrace{\left(-\frac{1}{2}\sum_{i,j}\frac{\partial s_i}{\partial \log p_j}\Big|_{(\mathbf{p}_1, u_1)} \log \Delta p_i \, \log \Delta p_j\right)}_{\text{Second-order substitution term}},$$

$$\text{Income effects} = -\underbrace{\left(\log \Delta x - \sum_j s_{j1}\log \Delta p_j\right)}_{\text{Real expenditure term}} \underbrace{\mathrm{Cov}_s\big(\log \Delta p_i, \eta_i\big)}_{\text{Inflation–expenditure elasticity covariance}}.$$

The approximation neglects terms of order higher than two in $(\log \Delta p_i, \log \Delta x)$.

Proof. See Appendix B.1. ☐

Proposition 1 provides a second-order decomposition of the quality-of-living change into nominal expenditure growth minus a price-driven term that can be split into exposure, a two-part substitution term, and income effects driven by 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 and demand system.

Exposure. The exposure term coincides with the log of the cost-of-living index under Leontief preferences, where the consumption bundle is fixed in the initial-period proportions. It measures how much the cost of the period-1 bundle changes when prices move from $\mathbf{p}_1$ to $\mathbf{p}_T$, holding quantities fixed. Equivalently, it is the log of the household-specific Laspeyres cost-of-living index, and captures differential *exposure* to inflation across households given their initial consumption baskets, before any behavioural substitution or non-homothetic demand adjustments.

Cobb–Douglas substitution. The Cobb–Douglas substitution component captures the substitution gain that would arise under Cobb–Douglas preferences with base-period budget shares s_{i1}. Under Cobb–Douglas preferences, budget shares are fixed and the expenditure function is log-linear in prices, so the log cost-of-living index is exactly $\sum_i s_{i1} \log \Delta p_i$, corresponding to a log geometric-Laspeyres index. The Cobb–Douglas substitution term therefore measures how much allowing for Cobb–Douglas substitution between goods reduces the cost of attaining a given utility level, relative to the fixed-bundle (Leontief/Laspeyres) benchmark.

Second-order substitution. The second-order substitution component captures deviations from the Cobb–Douglas benchmark arising from curvature of the expenditure function in 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 aggregates how cross-price substitution and departures of own-price elasticities from the Cobb–Douglas benchmark shape the welfare impact of a given pattern of price changes.[11] Its sign is a priori ambiguous: when goods whose relative prices increase are also goods that the household is willing to substitute away from strongly, the second-order term tends to make the true cost-of-living index lower than under the Cobb–Douglas (geometric-Laspeyres) benchmark, and conversely when substitution possibilities are weaker.

Income effects. The income-effects term captures how movements along Engel curves interact with the pattern of relative price changes to affect the welfare cost of inflation. It

[11] Using the Hicksian elasticity matrix ϵ^H_{ij}, the second-order substitution term can be written as: $-\frac{1}{2} \operatorname{Cov}_s \Big(\log \Delta p_i,\ \sum_j \left(\epsilon^H_{ij} + \mathbb{1}_{i=j} - s_{j1} \right) \log \Delta p_j \Big)$; that is, as a share-weighted covariance between relative price changes and an elasticity-weighted index of relative price changes.

depends on the product of two objects. The first is the bracketed term $\log \Delta x - \sum_j s_{j1} \log \Delta p_j$, which measures log real-expenditure growth (i.e., nominal expenditure growth deflated by the geometric-Laspeyres price index). When this object is negative, purchasing power falls; in a non-homothetic setting, this tends to shift spending toward lower–η_i goods (trading down); while a positive value corresponds to rising purchasing power and tends to shift spending toward higher–η_i goods (trading up).

The second object is the inflation–expenditure elasticity covariance, $\text{Cov}_s\big(\log \Delta p_i, \eta_i\big)$, which summarises how expenditure elasticities line up with the pattern of price changes. If goods with high expenditure elasticities (those whose quantities adjust most with total expenditure) experience systematically higher price growth, this covariance is positive; if those goods become relatively cheaper, it is negative. The product of the real-expenditure term and this covariance determines how movements along Engel curves, induced by changes in real expenditure, affect the quality of living. Under homothetic preferences, $\eta_i = 1$ for all i, the covariance is zero, the income-effects term vanishes, and the quality-of-living change reduces to nominal expenditure growth minus the cost-of-living index.

For example, suppose real expenditure falls and low–η_i necessity goods become relatively more expensive ("cheapflation"). Households then move along Engel curves, shifting spending toward necessities at the same time as those goods' relative prices are rising. In our decomposition this shows up as a negative income-effects term: the interaction of falling purchasing power with cheapflation raises the welfare cost of inflation and lowers the quality of living.

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. Consider the household-specific Laspeyres index

$$\Pi^L_{1,T} = \sum_i s_{i1} \Delta p_i, \tag{3.2}$$

which measures how much the cost of the initial-period consumption bundle changes between periods 1 and T.[12] Taking logs, the exposure term in Proposition 1 is simply

$$\text{Exposure} = \log \Pi^L_{1,T}.$$

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

A closely related index is the geometric-Laspeyres index,

$$\Pi_{1,T}^{GL} = \prod_i (\Delta p_i)^{s_{i1}} , \tag{3.3}$$

which corresponds to the Cobb–Douglas cost-of-living index with base shares s_{i1}. The Cobb–Douglas term in Proposition 1,

$$\log\Big(\sum_i s_{i1}\Delta p_i\Big) - \sum_i s_{i1}\log\Delta p_i,$$

is therefore equal to

$$\text{Cobb–Douglas substitution} = \log\Pi_{1,T}^{L} - \log\Pi_{1,T}^{GL},$$

i.e., the difference between the log Laspeyres and log geometric-Laspeyres indexes.

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, it captures a mixture of substitution and income effects.[13]

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 $\omega_{it}(u)$

[13]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 (e.g., the initial period), or for chaining the index across multiple periods.

generally differs from observed shares except in the base period, the Törnqvist index no longer provides a second-order approximation to the cost-of-living index 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 the sequence of money-metric utilities $(\tilde{u}_1, \ldots, \tilde{u}_T)$ from observed $(\mathbf{s}_t, \mathbf{p}_t, x_t)$, allowing for flexible non-homothetic preferences while restricting unobserved preference heterogeneity. The key insight is to approximate the log quality-of-living index between adjacent periods and express it as a function of money-metric utility; a "non-homothetic correction" is then identified from nonparametric cross-household regressions of a homothetic price index on $\log \tilde{u}_t$. 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 quality-of-living change $\log Q(u_1, u_T; \mathbf{p}_1) = \log e(\mathbf{p}_1, u_T) - \log e(\mathbf{p}_1, u_1)$. Using the identities $x_t = e(\mathbf{p}_t, u_t)$ and $P(\mathbf{p}_1, \mathbf{p}_T; u) = e(\mathbf{p}_T, u)/e(\mathbf{p}_1, u)$, we can rewrite this exactly as follows:

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

The first price term on the right-hand side, $\log P(\mathbf{p}_1, \mathbf{p}_T; u_1)$, is the log cost-of-living index evaluated at the initial utility level, while the bracketed difference is an exact measure of how the cost-of-living index varies with the reference utility level. Proposition 1 shows that, up to second order in $(\log \Delta p_i, \log \Delta x)$, the first term is given by the exposure and substitution components, while the bracketed term corresponds to the income-effects component: changes in real expenditure interacting with the alignment of relative price changes and expenditure elasticities. Under homothetic preferences, $\eta_i = 1$ for all i, the income-effects term vanishes and the quality-of-living change reduces to nominal expenditure growth minus the cost-of-living index.

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:

$$\begin{aligned}
\text{Exposure} &= \log \Big(\sum_i s_{i1} \Delta p_i \Big) = \log \Pi^{L}_{1,T}, \\
\text{Cobb–Douglas substitution} &= \log \Big(\sum_i s_{i1} \Delta p_i \Big) - \sum_i s_{i1} \log \Delta p_i = \log \Pi^{L}_{1,T} - \log \Pi^{GL}_{1,T}, \\
\text{Second-order substitution} &\approx \log \Pi^{GL}_{1,T} - \log \Pi^{JL}_{1,T}(u_1), \\
\text{Income effects} &\approx \log \Pi^{JL}_{1,T}(u_1) - \log \Pi^{JL}_{1,T}(u_T).
\end{aligned}$$

The exposure and Cobb–Douglas substitution components are exact functions of observed prices and expenditure shares. The second-order substitution and income-effects components use $\Pi^{JL}_{1,T}(\cdot)$, a second-order approximation to the corresponding cost-of-living indexes evaluated at a fixed reference utility level (see Appendix B.1). 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 $(\log \Delta p_i, \log \Delta x)$.

Discussion. Our decomposition is related to the second-order welfare expansions in Baqaee and Burstein (2023) and their application to inflation measurement in Auer et al. (2024). In particular, their formulas show that money-metric welfare differs from chain-weighted real consumption through expenditure-switching terms driven by substitution, income effects, and taste shocks, which can be summarised as covariances between price changes, income elasticities, and taste shifters. Proposition 1 applies similar second-order logic to a two-period quality-of-living index, and makes explicit how the corresponding exposure, substitution, and income-effects components can be written in terms of observable household-level price indexes. Whereas Auer et al. (2024) work with a compensating-variation-based quality-of-living index $Q(u_1, u_T; \mathbf{p}_T)$, evaluated at final-period prices (the proportional expenditure change needed at $\mathbf{p}_T$ to reach u_T rather than u_1), we instead use an equivalent-variation-based index $Q(u_1, u_T; \mathbf{p}_1)$, evaluated at initial prices (the proportional change in expenditure at $\mathbf{p}_1$ required to reach u_T rather than u_1). This choice implies that the way the cost-of-living index varies with the reference utility level shows up as a distinct income-effects component in Proposition 1, rather than being absorbed into a single cost-of-living term.

Our focus on $Q(u_1, u_T; \mathbf{p}_1)$ addresses the empirical question of how the 2021Q3–2023Q3 episode changed households' money-metric welfare relative to their pre-episode situation, measured in a common 2021 price environment. In a non-homothetic setting with cheapflation, the difference between equivalent and compensating variation is driven by the interaction between falling real expenditure and the pattern of relative price changes: movements along Engel curves toward necessity goods whose prices are rising. In our decomposition this appears as the income-effects term, which captures an additional, potentially regressive, component of the welfare cost of inflation that would be missed under a homothetic benchmark or a compensating variation-based treatment.

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 that is valid under general non-homothetic preferences, without imposing a parametric demand system at the product level. Each term in the expansion can be measured as log differences between indexes, which can be computed for every household in scanner data. Under standard revealed-preference assumptions (stable preferences and utility maximisation), this implementation uses only observed prices and expenditure shares; it does not rely on quasi-experimental income or price shocks or on estimating a structural demand system. In this sense, our analysis is complementary to recent structural work that uses non-homothetic demand and oligopolistic pricing to show how quality downgrading can endogenously generate cheapflation through countercyclical markups on low-quality goods (see Becker 2025). Second, we specialise the income-effects component to an environment with within-category quality ladders and "cheapflation", and interpret it as capturing how movements along Engel curves during a necessity-driven inflation episode amplify the distributional consequences of price changes. The framework decomposes the welfare impact of the *observed* price and expenditure changes, but does not, on its own, deliver arbitrary policy or price counterfactuals. However, as we show in Section 5, we can use the same objects to study exposure to a forward-looking "repeat-cheapflation" shock. This provides a tractable bridge between the general welfare formulas in Baqaee and Burstein (2023) and the empirical analysis of inflation inequality and quality downgrading that follows.

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 the household-specific Laspeyres index in equation (3.2). This captures how the cost of the 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 simply have $s_{hi1} = 0$, so they do not enter that household's index. By focusing on products that are available in both 2021Q3 and 2023Q3, this baseline Laspeyres index abstracts from product entry and exit. In Appendix C.2 we show that alternative treatments of product churn (e.g., chaining the index across adjacent quarters, and using a Feenstra–type variety correction; see Feenstra 1994; Broda and Weinstein 2010) have negligible impact on the pattern of inflation inequality.

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, earlier nine-quarter periods saw average cumulative inflation ranging from -3.9% (2014Q1–2016Q1) to 6.1% (2012Q1–2014Q1). Panel (b) illustrates the *distribution* of household-level cumulative inflation across these periods, showing that the elevated average inflation in 2021Q3–2023Q3 was accompanied by much greater dispersion across households.

Figure 4.1: *Household-level inflation*

(a) Average inflation (b) Inflation dispersion

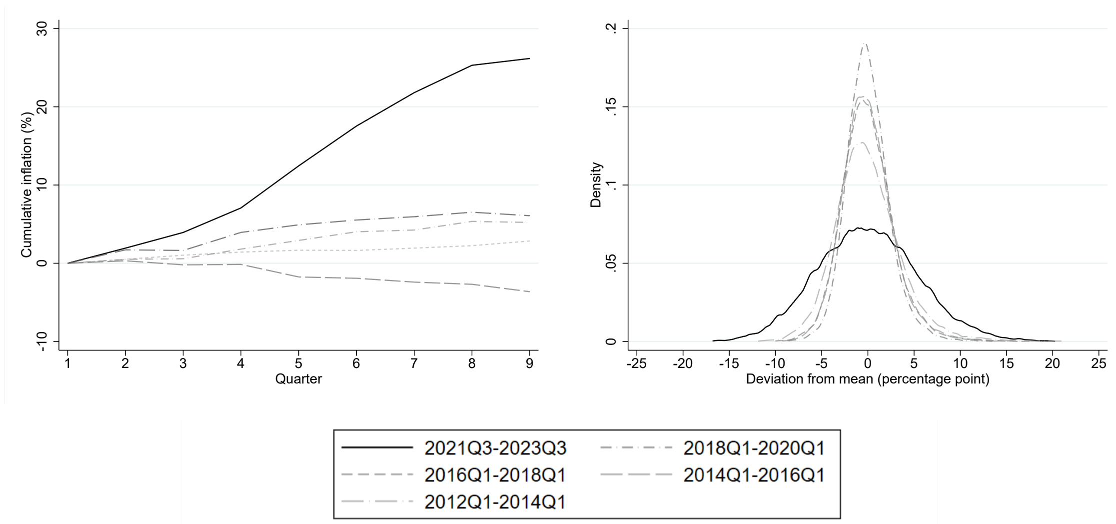

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2012–2023). Panel (a) presents the average of cumulative inflation across households for each nine-quarter period. Panel (b) displays kernel density estimates of the distribution of cumulative inflation in the ninth quarter.

Inflation 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 inflation gradients in 2021Q3–2023Q3. Households in the bottom quartile of the equivalised 2021 expenditure distribution experienced cumulative inflation exposure approximately 4.5 percentage points higher than those in the top quartile, while the difference between bottom and top deciles is 7.5 percentage points. Households in the bottom equivalised income quartile had exposure around 3.0 percentage points higher than households in the top income quartile, while the inter–decile range is 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 worse-off households lead to systematically higher inflation exposure. Quantitatively, the inflation 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 introduces noise into the ranking of households by equivalised income and tends to attenuate differences in average inflation 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.

Figure 4.2: *Inflation inequality*

(a) By total expenditure

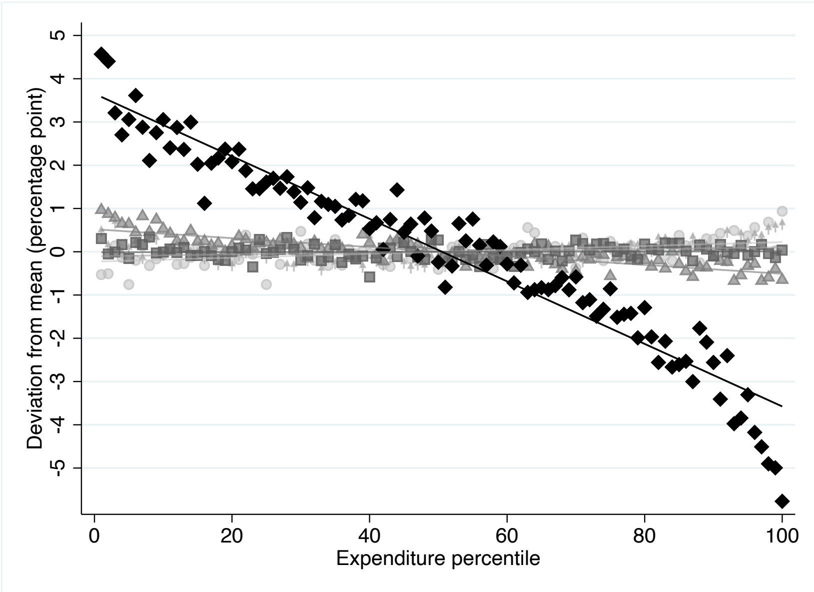

(b) By household income

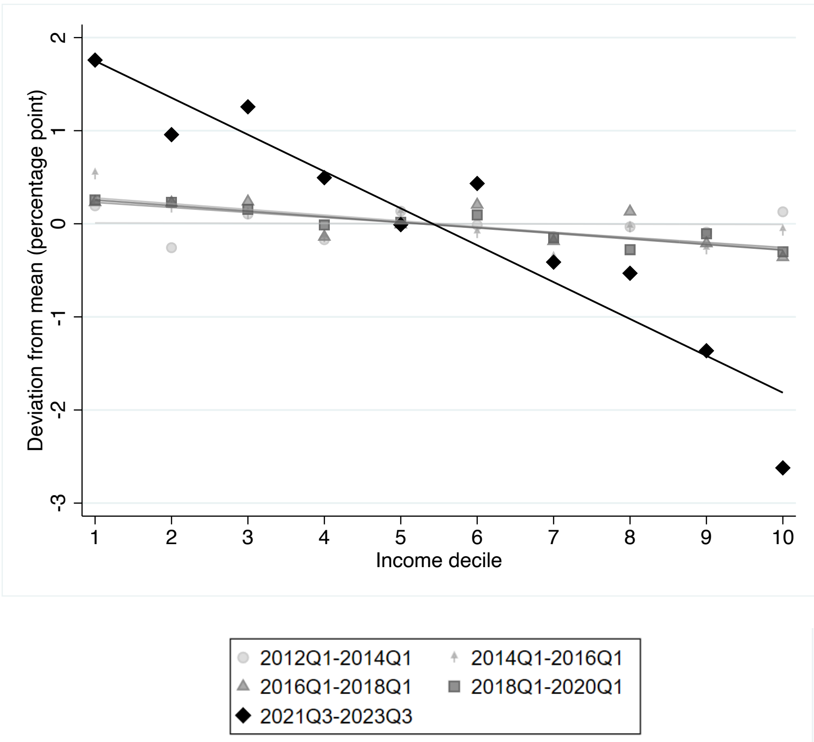

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2012–2023). Figure plots the relationship between ninth-quarter cumulative inflation and a household's percentile in the equivalised expenditure distribution (panel (a)) and income deciles based on a household's equivalised income computed using the midpoint of each income band (panel (b)), with a marker for each percentile/decile and a line of best fit. Households are assigned to expenditure percentiles and income deciles based on the initial calendar year of the relevant nine-quarter period. Cumulative inflation is measured using a Laspeyres index.

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

Concretely, let Ω^c denote the set of products that belong to product category c and Γ^g denote the set of product categories that belong to segment g. Define the household-specific within-category product share, the within-segment category share, and the segment share as:

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

respectively. 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 $\Pi^{HL}_{h(1,T)} = \Pi^L_{h(1,T)}$. We use this representation to decompose household-specific exposure to inflation by sequentially "switching off" heterogeneity at each level—replacing household-specific spending shares at that level and below with population-average shares—while leaving higher-level shares household-specific. For example, to isolate the importance of heterogeneity in segment shares, we replace category- and product-shares with their population averages, leaving only segment-level heterogeneity.

In Figure 4.3 we plot the distribution of household-level cumulative inflation over 2021Q3–2023Q3 (panel (a)) and the across-expenditure inflation gradient (panel (b)) using black lines and markers, replicating information from Figures 4.1 and 4.2. We also show patterns when we switch off heterogeneity in spending across products within categories (red lines and markers) and, additionally, when we switch 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 1.5 percentage points when only segment shares vary across households (all within-segment and within-category shares set to population averages), and rises to 5.4 percentage points in the full household-specific index. Panel (b) shows that heterogeneity in spending across products within product categories is essentially fully responsible for the inflation gradient: once within-category product shares are homogenised across households, the inflation gradient is close to zero. These findings emphasise the importance of using detailed household spending and price information to study inflation inequality.

Figure 4.3: *Household-level hierarchical inflation*

(a) Distribution (b) Inequality

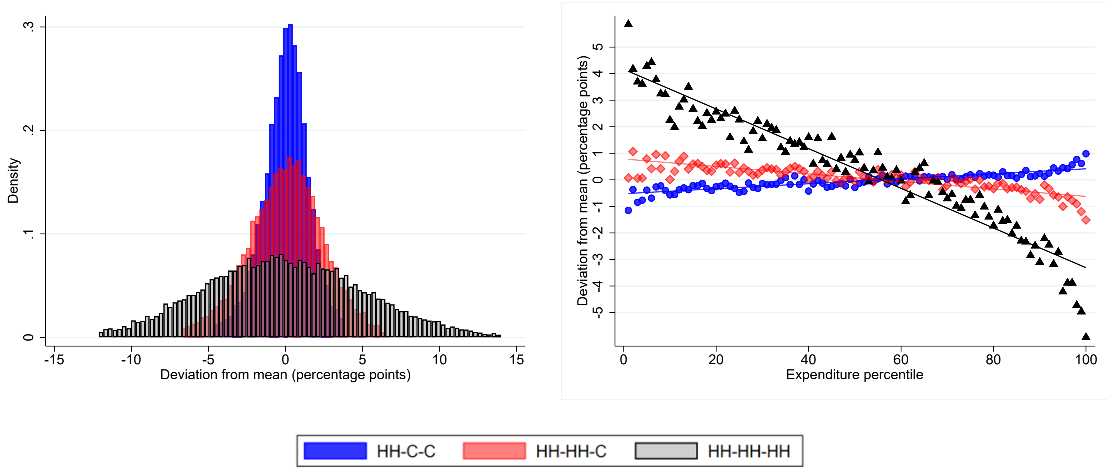

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2021–2023). Panel (a) shows a histogram of the distribution of cumulative inflation from 2021Q3 to 2023Q3. Panel (b) plots the relationship between cumulative inflation and the equivalised expenditure-distribution percentile to which a household belongs, with a marker for each percentile and a line of best fit. HH-C-C uses household-specific segment shares and common (average) within-segment shares, HH-HH-C uses household-specific segment and category shares, and common within-category shares, while 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 is measured using a Laspeyres index.

Our descriptive evidence in Figure 2.1 shows that price growth is significantly stronger down the quality ladder, and that the consumption baskets of worse-off households make them more exposed to these price dynamics. Figures 4.2 and 4.3 demonstrate that these patterns translate into wide inflation-exposure inequality, and that this inequality is dominated by within-category differences in cumulative price growth across rungs of the quality ladder, together with heterogeneity in product choice across households. In the welfare decomposition below, the same cheapflation–necessity alignment matters again through the income-effects term, via the covariance between relative price changes and expenditure elasticities: goods that are both necessity-type (low η_i) and subject to strong relative price increases are key drivers of the distribution of welfare losses.

Replacing the common prices in the Laspeyres index with household-specific prices has minimal impact on the inflation inequality gradient (see Appendix C.1). This aligns with our evidence that poorer households pay only modestly lower prices for identical products, and that this pattern remains relatively stable over time (see Appendix A.1). A further concern is that the Laspeyres index abstracts from product entry and exit. Neither chaining the Laspeyres index nor using the Feenstra–Broda–Weinstein CES index, which incorporates variety effects under CES preferences, has a 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, 20.7% of total household expenditure (and 36.3% of total spending on goods) in the LCFS 2021, with budget shares declining from roughly 23.3% in the bottom equivalised-income decile to 17.5% in the top decile. Consistent with Engel's law, food and basic consumer goods therefore absorb a much larger fraction of the budgets of poorer households. Combined with the stronger cheapflation we document for FMCG, this suggests that regressive exposure is amplified once we account for higher FMCG budget shares among worse-off households.[14]

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-group-specific CPI indexes under two scenarios: one in which FMCG categories have common inflation across all groups, and one in which they inherit the group-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 inter–decile gap in CPI inflation by about 0.9 percentage points, a 40.3% 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 suggest that the inequality we document for FMCG is likely reinforced, rather than offset, when considering the broader consumption basket. First, the LCFS–CPI analysis shows that, even with common inflation within CPI categories, 2021Q3–2023Q3 is a period of substantial inequality in inflation exposure, driven by very large increases in energy prices combined with higher energy expenditure shares for poorer households (see Levell et al. 2025); when residential energy is excluded from the CPI basket, the income gradient in the common-price index becomes much smaller. 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. Taken together, these facts imply that our scanner-data-based results capture one important margin of inflation inequality in the 2021Q3–2023Q3 episode, and that cheapflation is likely to reinforce regressive patterns of exposure across wider consumption baskets.

[14] 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).

Discussion. Taken together, these results show that the 2021Q3–2023Q3 UK inflation episode was unusual not only in terms of its aggregate level, but also in the extent of dispersion and inequality in household-specific inflation exposure. Earlier work documents 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 shows that the degree of measured inflation inequality can be attenuated when using data aggregated across products (see Jaravel 2021).

A key empirical finding of this section is that, in a period of sharp and highly uneven relative price changes within the fast-moving consumer goods sector, inflation inequality across households is large and strongly regressive across both income and expenditure distributions, and that 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) and recent evidence linking within-category cheapflation to variation in inflation inequality in US data (Sangani 2025), indicating that our UK-based results on cheapflation and inflation inequality are relevant beyond this setting. It also complements evidence in Jaravel (2024) on recent inflation heterogeneity across income groups based on expenditure differences across more aggregated COICOP spending categories. The price-growth differentials we estimate between low- and high-quality rungs 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, which, using our decomposition, quantifies how substitution and non-homothetic demand responses shape the distributional consequences of such "cheapflation" episodes.

4.2 Substitution responses

Proposition 1 shows that household reoptimisation through both substitution and income effects shapes the proportional change in resources required, at initial (2021Q3) prices, to attain realised utility between 2021Q3 and 2023Q3. In the absence of such behavioural responses, the quality-of-living index would simply equal nominal expenditure growth deflated by exposure to inflation at the initial-period consumption basket (captured by the Laspeyres price index). The substitution terms in the decomposition capture the improvement in the quality of living that arises when households adjust their consumption bundles in response to relative price changes, holding the utility level at its initial value. For a given set of relative price changes, the sum of the exposure and substitution terms is therefore the (log) change in the cost of living: the proportional change in resources needed to attain initial-period utility as prices move from 2021Q3 to 2023Q3.

Index construction. To measure substitution responses we compare the Laspeyres exposure index to household-specific cost-of-living indexes that allow for substitution. We construct three such indexes. First, the Cobb–Douglas (geometric-Laspeyres) index uses the same 2021Q3 base shares (s_{hi1}) as the Laspeyres index and its log is given by $\log \Pi^{GL}_{h(1,T)} \equiv \sum_{i \in \Omega} s_{hi1} \log \frac{p_{iT}}{p_{i1}}$. 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 $\log \Pi^{T}_{h(1,T)} \equiv \sum_{t=1}^{T-1} \sum_{i \in \Omega_t} \frac{1}{2} (s_{hit} + s_{hit+1}) \log \frac{p_{it+1}}{p_{it}}$ where s_{hit} is the expenditure share of product i in period t for household h. 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. We compute non-homothetic cost-of-living indexes evaluated at the initial and final utility levels (see Appendix B.2 for details). We obtain the final-utility index by running the algorithm forward from 2021Q3 to 2023Q3, and the initial-utility index by running the analogous procedure in reverse.

Average cost of living. Column (1) of Table 4.1 summarises how substitution responses affect both the average cumulative change in the cost of living over 2021Q3–2023Q3 under different restrictions on household preferences.[15]

Row 1 reports numbers under the Leontief (no-substitution) benchmark. Under Leontief preferences households do not switch across products in response to relative price changes, so their cost-of-living index coincides with the Laspeyres index. Row 2 reports cost-of-living increases based on the geometric-Laspeyres index in equation (3.3), which coincides with the cost-of-living index under Cobb–Douglas preferences with base-period shares. The difference between the Cobb–Douglas index (24.31% on average) and the Laspeyres index (26.21% on average) corresponds to the Cobb–Douglas substitution term in Proposition 1: allowing for Cobb–Douglas substitution reduces the average cost-of-living increase by 1.90 percentage points relative to the Leontief benchmark.

Under Cobb–Douglas preferences, household budget shares are fixed, so deviations from constant shares—reflecting richer substitution patterns—show up in the second-order substitution term in Proposition 1. Under homothetic preferences, this second-order substitution term is approximated by the difference between a Törnqvist index (row 3) and the

[15] 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).

geometric-Laspeyres/Cobb–Douglas index. Under non-homothetic preferences, the second-order term is approximated by the difference between the Jaravel–Lashkari non-homothetic index, evaluated at the initial utility level (row 4) and the geometric-Laspeyres index.

Allowing for non-homothetic preferences, the average cost-of-living increase over 2021Q3–2023Q3 is 25.16%, so that substitution responses lower the average cost of living by 1.05 percentage points. This net effect is smaller in magnitude than the Cobb–Douglas substitution term alone (1.90 percentage points), implying the second-order substitution term acts to *raise* the cost of living by 0.85 percentage points. This pattern is consistent with products experiencing the strongest price growth being those that households are least willing or able to substitute away from, so that richer substitution patterns attenuate some of the gains implied by the Cobb–Douglas benchmark.

The Törnqvist index records an average cost-of-living increase that is 0.15 percentage points higher than the non-homothetic Jaravel–Lashkari index (25.31% versus 25.16%). This difference arises from the non-homothetic correction in the Jaravel–Lashkari index and is therefore a direct manifestation of non-homothetic demand. The Törnqvist index is a homothetic superlative index constructed from observed (Marshallian) shares, which reflect both substitution responses to relative prices and income-effect movements along Engel curves. When households respond to reduced purchasing power by shifting spending toward necessity goods that experienced relatively strong price growth, a homothetic index such as Törnqvist misinterprets part of this income-effect-based reallocation as limited scope for substitution and therefore reports a higher cost-of-living index. In this sense, the Törnqvist index is subject to a non-homotheticity bias: it overstates the cost-of-living increase relative to the non-homothetic index that holds utility fixed at its initial level.

Cost-of-living inequality. Columns (2)–(5) of Table 4.1 summarise how substitution responses affect the inter–quartile and inter–decile ranges in cost-of-living increases across the equivalised expenditure and equivalised income distributions.

Under general non-homothetic preferences, the cost-of-living gradient is slightly dampened relative to the inflation-exposure gradient. Across the expenditure distribution, the inter–quartile and inter–decile ranges fall by 0.27 percentage points (from 5.50 to 5.23 in absolute value) and 0.35 percentage points (from 7.48 to 7.13), respectively. Across the income distribution, they fall by 0.48 percentage points (from 2.98 to 2.50) and 0.62 percentage points (from 4.44 to 3.82). This pattern indicates that substitution responses were somewhat stronger among worse-off households. Nonetheless, the cost-of-living gradient remains large and regressive: households with lower income and lower total expenditure levels continue to experience substantially higher cost-of-living increases than better-off households.

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

	Average change (%) (%)	Inter–quartile range (p.p)		Inter–decile range (p.p)	
		Expenditure	Income	Expenditure	Income
	(1)	(2)	(3)	(4)	(5)
Leontief preferences ($\Pi^{L}_{1,T}$)	26.21	-5.50	-2.98	-7.48	-4.44
Cobb–Douglas preferences ($\Pi^{GL}_{1,T}$)	24.31	-5.27	-2.83	-7.17	-4.22
General homothetic preferences ($\Pi^{T}_{1,T}$)	25.31	-4.75	-2.48	-6.48	-3.81
General non-homothetic preferences ($\Pi^{JL}_{1,T}(u_1)$)	25.16	-5.23	-2.50	-7.13	-3.82

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2021–2023). Average change is the mean cumulative cost-of-living increase (in percent). Inter–quartile (inter–decile) ranges are measured in percentage points and are defined as top minus bottom quartile (decile), so that negative values indicate higher cost-of-living changes at the bottom of the distribution.

4.3 Income effects

The final component in Proposition 1 is the income-effects term. It reduces the quality of living when purchasing power falls and households trade down toward necessity-type goods whose relative prices rise most. This component can be understood through how the cost-of-living index varies with the reference living standard. Since the consumption bundle required to attain a lower utility level is more concentrated in necessities, in a cheapflation environment the same pattern of price changes raises the cost of living more at low utility than at high utility. When a household's living standard falls from the initial to the final period, the cost of living evaluated at the realised final-period utility therefore exceeds that evaluated at the realised initial-period utility. The gap between the initial- and final-utility cost-of-living indexes is exactly the utility-dependence term in Section 3.3 and, in Proposition 1, corresponds to the income-effects term.

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 income effects. On average, the cost of living rises by 25.16% when evaluated at realised initial-period utility u_1, but by 26.04% when evaluated at final-period utility u_T. This 0.88 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. In Proposition 1, this means that income effects raise the welfare cost of inflation (and lower the quality of living).

Income effects and inequality. Columns (2)–(5) of Table 4.2 compare inter–quartile and inter–decile differences in cost-of-living changes across the expenditure and income distributions when the Jaravel–Lashkari index is evaluated at initial-period utility u_1 versus at realised final-period utility u_T. The differences between these two evaluations capture how the cost-of-living index varies with the reference utility level; in our decomposition, this corresponds to the income-effects component. Across the expenditure distribution, the inter–quartile and inter–decile ranges shrink in absolute value from 5.23 to 4.36 percentage points and from 7.13 to 5.95, respectively, when moving from the initial-utility to the final-utility index. Evaluating the cost of living at u_T rather than u_1 therefore attenuates the expenditure gradient: the inter–quartile gap narrows by 0.87 percentage points and the inter–decile gap by 1.18 percentage points. By contrast, the gradient across the income distribution is essentially unchanged: the inter–quartile range remains 2.50 percentage points in absolute value, and the inter–decile range changes only slightly, from 3.82 to 3.87 percentage points in absolute value.

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

	Average change (%) (%)	Inter–quartile range (p.p)		Inter–decile range (p.p)	
		Expenditure	Income	Expenditure	Income
	(1)	(2)	(3)	(4)	(5)
Initial-period cost-of-living index ($\Pi^{JL}_{1,T}(u_1)$)	25.16	-5.23	-2.50	-7.13	-3.82
Final-period cost-of-living index ($\Pi^{JL}_{1,T}(u_T)$)	26.04	-4.36	-2.50	-5.95	-3.87

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2021–2023). Average change is the mean cumulative cost-of-living increase (in percent). Inter–quartile (inter–decile) ranges are measured in percentage points and are defined as top minus bottom quartile (decile), so that negative 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 combine the exposure, substitution, and income-effects components of Proposition 1 to summarise the price-driven welfare consequences of the 2021Q3–2023Q3 inflation episode. Figure 4.4 brings these components together, showing how inflation exposure and behavioural adjustments shape inflation-driven welfare losses across the equivalised expenditure and income distributions.

We focus on the price-driven component of welfare. For each household we define the welfare cost of inflation as

$$W^{\pi} \equiv \text{Exposure} - \text{Cobb–Douglas substitution} - \text{Second-order substitution} - \text{Income effects}.$$

By Proposition 1, the total welfare change is then $\log Q(u_1, u_T; \mathbf{p}_1) \approx \log \Delta x - W^{\pi}$. We focus on W^{π} because it isolates the contribution of price changes and behavioural adjustments to welfare losses.

Panels (a) and (b) plot, by decile, average log Laspeyres inflation exposure and the welfare cost of inflation W^{π}. By the mapping in Section 3.3, W^{π} is numerically equal (up to approximation error) to the Jaravel–Lashkari cost-of-living index evaluated at final-period utility, $\log \Pi^{JL}_{1,T}(u_T)$. Relative to the exposure-only benchmark, behavioural adjustments modestly compress the inflation gradient: for households in the bottom expenditure decile, W^{π} is 0.88 log points ($\times$ 100) below inflation exposure based on the initial basket, whereas for the top expenditure decile behavioural adjustments raise welfare losses by 0.31 log points relative to the exposure-only benchmark. Across income deciles, the corresponding figures are 0.40 and 0.16 log points.

The welfare cost of inflation remains sharply regressive. The bottom decile of the equivalised expenditure distribution faces a welfare loss about 5.95 percentage points larger than the top decile, and the corresponding bottom–top gap across income deciles is 3.87 percentage points. Thus, even after accounting for substitution and income effects, behavioural adjustments compress—but by no means eliminate—the regressive pattern in welfare losses implied by initial-basket inflation exposure.

Panels (c) and (d) decompose the gap Exposure $-$ W^{π} into Cobb–Douglas substitution, second-order substitution, and income effects. Cobb–Douglas substitution lowers the welfare cost of inflation for all groups, with somewhat larger effects at the bottom of the expenditure and income distributions. Second-order substitution partially offsets this first-order term across all deciles, although the total substitution effect (Cobb–Douglas plus second-order) still reduces the welfare cost of inflation throughout the distribution.

Finally, the income-effects component 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 shift spending toward necessity-type goods whose relative prices rose most over the period. These Engel-curve adjustments are optimal given prices and budgets, but because they reallocate expenditure toward goods with stronger relative price growth, they raise the welfare cost of inflation in Proposition 1 (and therefore enter panels (c) and (d) as a negative contribution to Exposure $-$ W^{π}). The magnitude of this income-effects term is somewhat smaller in the lower part of the expenditure distribution, acting against but not fully undoing the regressive pattern in inflation-driven welfare changes.[16]

[16]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 as a non-homothetic income-effects component that can add to welfare losses even when they lower the average price level of the consumption basket.

Figure 4.4: *Welfare costs of inflation*

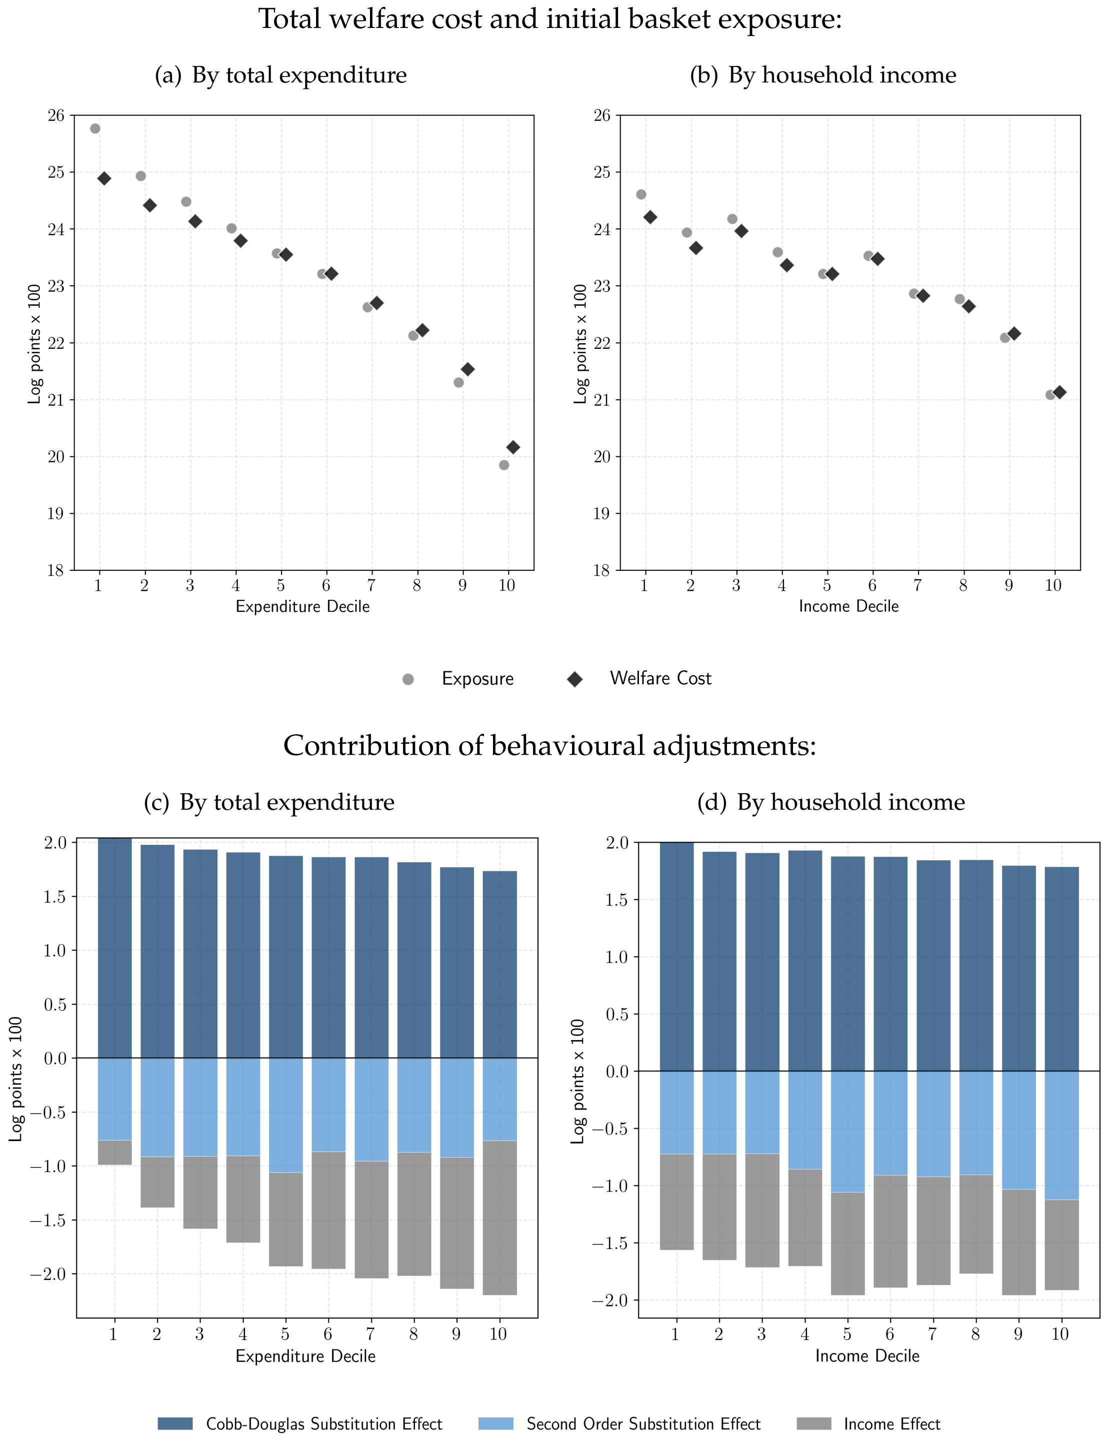

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2021–2023). Panels (a) and (c) group households into deciles of equivalised 2021 total fast-moving consumer good expenditure; panels (b) and (d) into deciles of equivalised 2021 income. Panels (a) and (b) plot, by decile, average log Laspeyres inflation exposure and the welfare cost of inflation W^{π}, measured by the log Jaravel–Lashkari cost-of-living index evaluated at the final-period utility level. Panels (c) and (d) decompose Exposure $-$ W^{π} into terms for Cobb–Douglas substitution, second-order substitution, and income effects. All values are expressed in log points ($\times$100) and averaged within each decile.

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 in exposure 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 the second-order substitution component offsets a substantial part of the Cobb–Douglas benchmark gain. This pattern is consistent with the goods experiencing the strongest relative price increases being those from which households are least willing or able to substitute away.

Third, income effects add a separate welfare channel: when purchasing power falls, households shift toward necessity-type 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. This non-homothetic channel raises the inflation-driven welfare cost over and above what is captured by exposure and substitution.

5 Dynamic exposure to future cheapflation

The analysis so far focuses on the welfare costs of the price changes during the 2021Q3–2023Q3 inflation episode. Similar within-category "cheapflation" patterns can, however, arise from other shocks, including trade-policy changes. Recent evidence suggests that tariffs can generate within-category cheapflation, with cheaper varieties exhibiting larger price increases than more expensive ones (Cavallo et al. 2025). In this final section, we ask whether the combination of cheapflation and falling purchasing power has reshaped households' *future* vulnerability to similar episodes. If households respond to cheapflation by trading down along the quality ladder toward necessity goods whose relative prices rose most, their *final* consumption baskets may be more exposed to a similar pattern of relative price changes than their initial baskets.

A repeat-cheapflation exposure index. To quantify this channel, we construct a "repeat-cheapflation exposure" index that applies the 2021Q3–2023Q3 relative price changes to each household's final-period consumption bundle. 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

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

and use $\log \Pi^{L,\text{rep}}_{h(1,T)}$ as a measure of how exposed the final-period consumption basket would be if the observed pattern of relative price changes were to repeat.[17] We also consider the change in exposure implied by the episode,

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

which captures how movements in consumption bundles over 2021Q3–2023Q3 have reconfigured households' vulnerability to a repeat of the episode's relative-price shock.

Income effects and future exposure. The income-effects term captures the extent to which falling purchasing power shifts households' spending toward necessity-type goods whose relative prices rose most during the episode. If these Engel-curve adjustments tilt the final-period basket toward cheapflation-intensive products, households with more negative income effects should exhibit higher repeat-cheapflation exposure, $\log \Pi^{L,\text{rep}}_{h(1,T)}$, and larger increases in exposure, $\Delta^{\exp}_h$. Figure 5.1 documents this relationship by plotting both indexes against percentiles of the income-effects term.

Panel (a) reveals a strong negative association between income effects and repeat-cheapflation exposure: households with more negative income effects (i.e., larger welfare losses through the income-effects channel) tend to have *higher* values of $\log \Pi^{L,\text{rep}}_{h(1,T)}$. Moving from the top decile of the income-effects distribution (relatively small or near-zero income effects) to the bottom decile (the most negative income effects) is associated with an increase of around 0.8 log points in repeat-cheapflation exposure. Panel (b) shows an even stronger negative relationship between the income-effects term and the change in exposure, $\Delta^{\exp}_h$: households that experienced the most negative income effects also experienced the largest increases in exposure between the beginning and end of the episode.

[17] 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 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.

Figure 5.1: *Vulnerability to future inflation*

(a) Repeat-exposure index

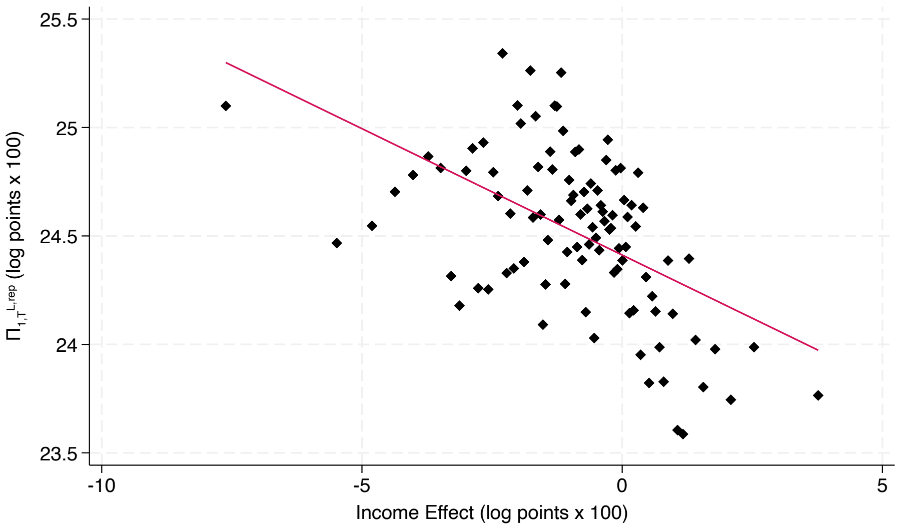

(b) Change in exposure

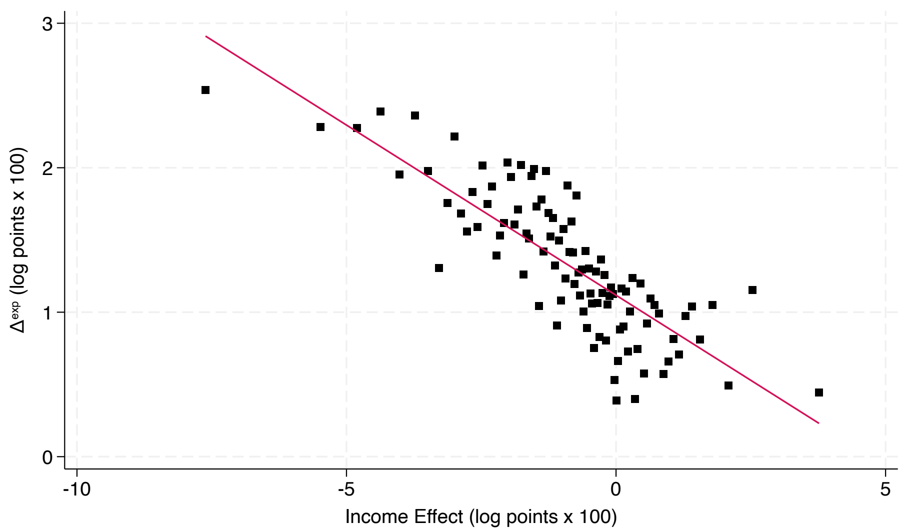

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2021–2023). Households are assigned to 100 bins based on percentiles of the income-effects component of the welfare cost of inflation. Panel (a) plots the repeat-cheapflation exposure index by bin. Panel (b) plots, by the same bins, the difference between the repeat-cheapflation exposure index and the exposure component of the welfare cost of inflation. All values are expressed in log points ($\times$100*).*

Taken together, these patterns imply that households most affected by the interaction of cheapflation and falling purchasing power ended the episode with consumption baskets that are *more* exposed to a repeat cheapflation shock than their initial baskets. At the same time, initial exposure during 2021Q3–2023Q3 is higher among households with less negative income effects. Thus, the households with the most negative income effects did not start out as the most exposed to cheapflation; rather, their Engel-curve adjustments under cheapflation drew them into parts of the product space where a future cheapflation episode would hit harder.

Dynamic vulnerability to necessity-driven inflation. These results highlight an intertemporal dimension to inflation inequality. Earlier sections show that cheapflation, combined with non-homothetic demand, generates large and regressive differences in *realised* welfare losses across the expenditure and income distributions. Figure 5.1 suggests that the same forces also reshape the cross-sectional distribution of *future* vulnerability to similar shocks. Households that adjust most strongly along the quality ladder toward low-quality necessity goods end the episode with consumption baskets that are more heavily tilted toward cheapflation-intensive products. If the pattern of relative price changes observed over 2021Q3–2023Q3 were to recur, these households' post-episode baskets would make them more exposed than their initial baskets.

From a policy perspective, this dynamic exposure margin suggests that the welfare consequences of cheapflation episodes can persist through the composition of households' baskets even after the episode ends, so long as those baskets remain tilted toward necessities. In this sense, the income-effects term in our decomposition not only captures an additional source of welfare loss during the cheapflation episode itself, but also signals how that episode has reallocated exposure across households in a way that would matter if similar necessity-driven relative-price shocks recur.

6 Conclusion

We study the 2021Q3–2023Q3 UK inflation episode and show that cheapflation—faster price growth at the lower end of the quality ladder—interacted with non-homothetic demand to generate large, regressive welfare losses. Using a second-order decomposition of an equivalent-variation quality-of-living index implemented with household price indexes, we disentangle exposure, substitution, and Engel-curve (income-effect) adjustments under general non-homothetic preferences. A central mechanism is that when real purchasing power falls, households optimally reallocate spending toward low-expenditure-elasticity necessities, and in a cheapflation episode these are precisely the goods with the strongest relative price increases, amplifying welfare losses.

A further implication is that these Engel-curve adjustments can also reshape households' exposure to future necessity-driven inflation: households that trade down more end the episode with baskets more tilted toward cheapflation-intensive goods, and would therefore be more exposed 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 necessity-driven inflation 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*.

Cavallo, A., P. Llamas, and F. M. Vazquez (2025). Tracking the short-run price impact of us tariffs.

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*. London: The Stationary 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, and E. Qian (2024). Are Inflationary Shocks Regressive? A Feasible Set Approach. *mimeo*.

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. Technical Report 1.

Ferreira, C., J. M. Leiva, G. Nuño, Á. Ortiz, T. Rodrigo, and S. Vazquez (2023). The Heterogeneous Impact of Inflation on Households' Balance Sheets. *BIS Working Papers No 1152*.

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*.

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

Klick, J. and A. Stockburger (2021). Experimental CPI for lower and higher income households. Technical Report Working Paper 537, U.S Bureau of Labor Statistics.

Klick, J. and A. Stockburger (2024). Examining U.S. inflation across households grouped by equivalized income. Technical Report Monthly Labor Review, Bureau of Labor Statistics.

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. *The American Economic Review 79*(2), 325–330.

Sangani, K. (2023). Pass-Through in Levels and the Unequal Incidence of Commodity Shocks. *SSRN Electronic Journal*.

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

APPENDIX:
FOR ONLINE PUBLICATION

Measuring cost of living inequality during an inflation surge

Tao Chen, Peter Levell and Martin O'Connell
February 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. Therefore our baseline analysis incorporates substitution toward the main low-cost seller 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 price 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 $p_{hit} = x_{hit}/q_{hit}$.[18] 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:

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

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

[18]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*

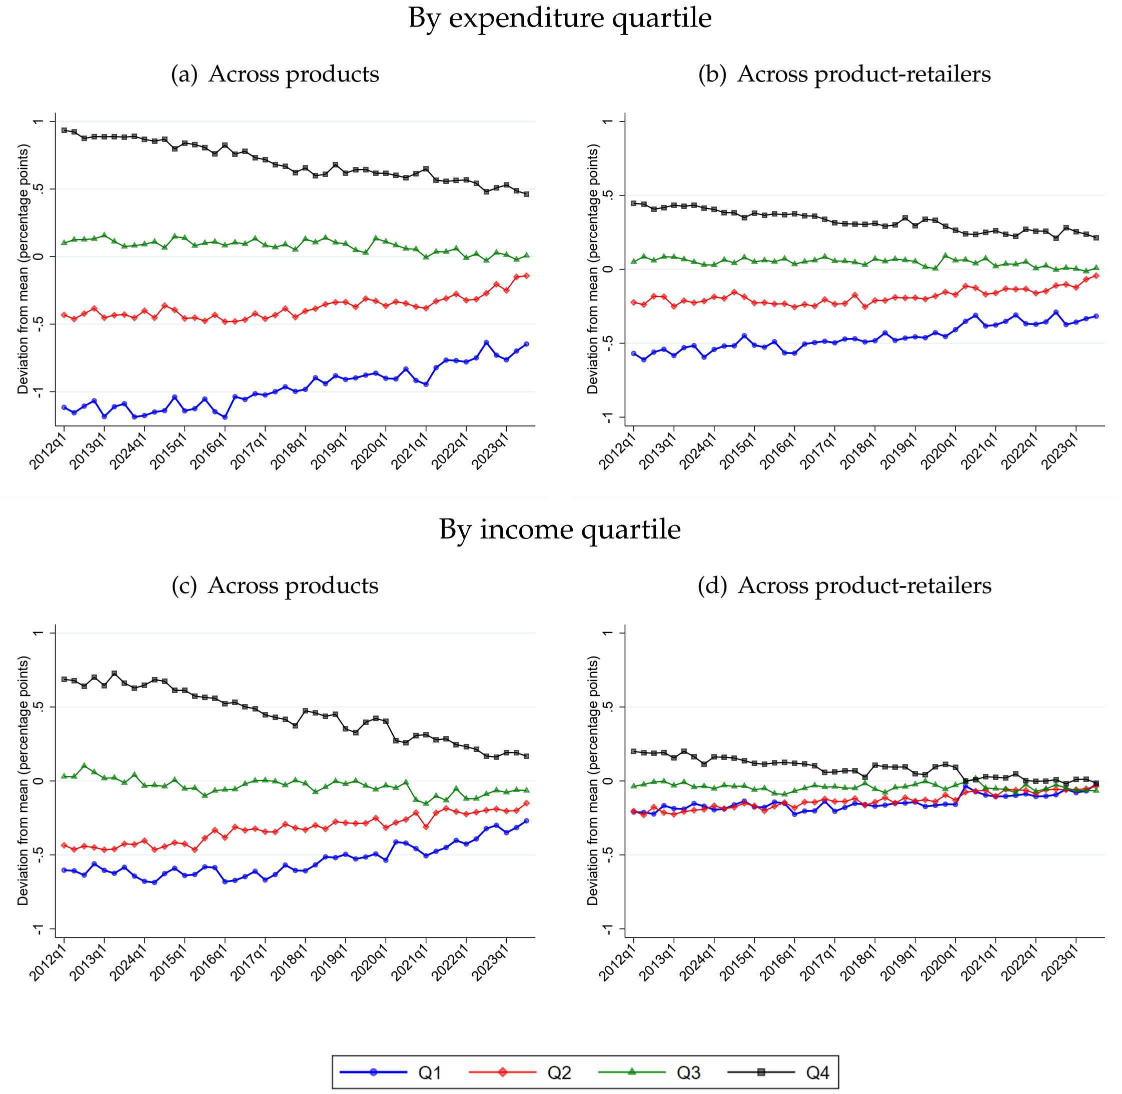

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2012–2023). Panel (a) and (c) define products at the brand–pack-size level; panel (b) and (d) further conditions on retailer, defining products at the brand-pack-size-retailer level. Panels (a) and (b) plot, for each year–quarter, the mean AH index within equivalised expenditure quartiles (across households), expressed as deviations from the quarterly mean. Panels (c) and (d) repeat (a) and (b), instead splitting households based on their 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 poorer 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 poorer households. 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 (reported in bands) in 2021. 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 within product categories in the third quarter of 2021, 2022 and 2023, using the Numerator data. When lower-priced products within a category experience faster price growth—a phenomenon known as “cheapflation”—we expect average price dispersion to fall, causing these distributions to shift left. We indeed observe a leftward shift in 2023 compared to 2022 and 2021.

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

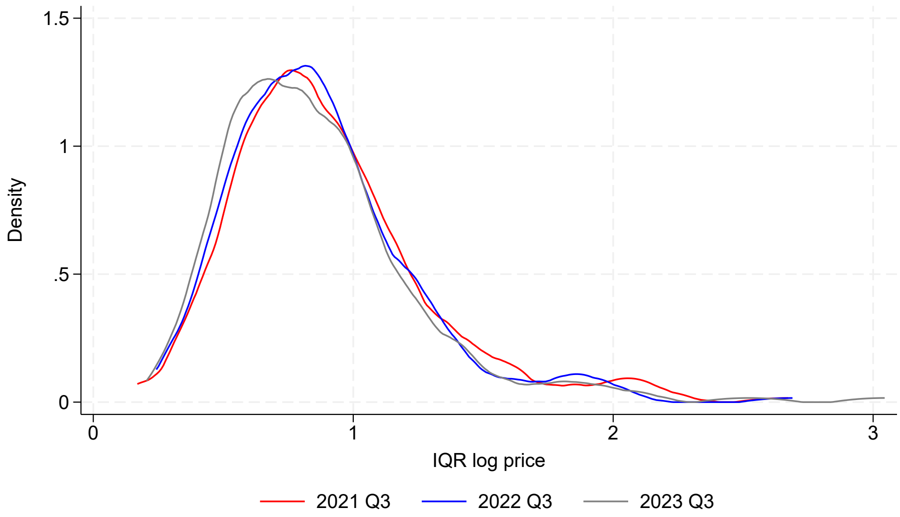

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2012–2023). The graph shows the distribution of interquartile ranges of log prices across products within product categories in 2021Q3, 2022Q3 and 2023Q3.

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

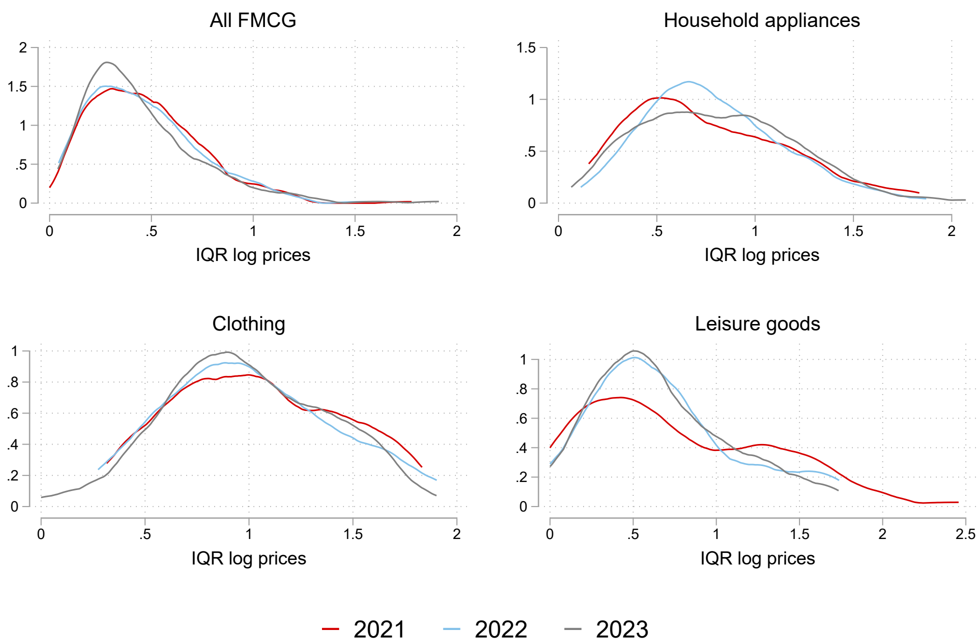

Notes: Authors' calculations using CPI microdata produced by the Office for National Statistics. The graph shows the distribution of interquartile ranges of log prices across sampled prices within CPI 'items' in December 2021, 2022 and 2023. FMCG denotes fast-moving consumer goods.

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 for December in years 2021, 2022 and 2023. Items are disaggregated product groups (e.g., "large white loaf unsliced, 800g") which are more granular than the product categories we use in Figure A.3.

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.

A.4 Spending across 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, showing a similar relationship to variation across expenditure deciles in Figure 2.1(b).

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

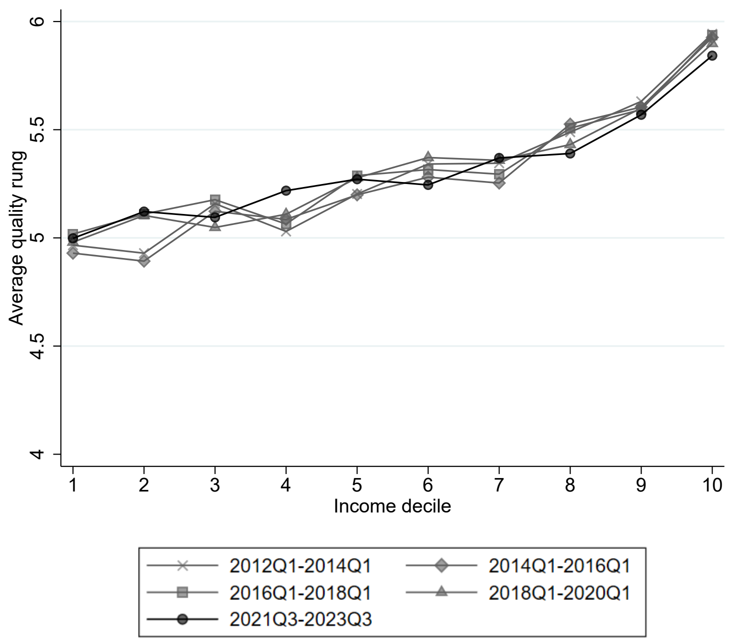

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2012–2023). Figure reports the average quality rung of households' purchases by deciles of the equivalised income distribution.

B Measurement appendix

B.1 Proof of Proposition 1

Lemma 1. *Let preferences be such that the indirect utility function $v(\mathbf{p}, x)$ and expenditure function $e(\mathbf{p}, u)$ 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*

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

Let

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

denote the Marshallian budget share 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. By the chain rule,

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

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

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

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

Roy's identity implies the Marshallian budget share

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

Combining (B.1), (B.5), and (B.6) yields

$$\left.\frac{\partial \log \mu(\mathbf{p}, x)}{\partial \log p_i}\right|_{(\mathbf{p}_1, x_1)} = -\, s_{i1}, \qquad \left.\frac{\partial \log \mu(\mathbf{p}, x)}{\partial \log x}\right|_{(\mathbf{p}_1, x_1)} = 1.$$

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

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

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

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

and evaluating (B.7) at $(\mathbf{p}_1, x)$ gives

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

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

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

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

$$\begin{aligned} \left.\frac{\partial^2 \log \mu(\mathbf{p}, x)}{\partial \log x \, \partial \log p_i}\right|_{(\mathbf{p}_1, x_1)} &= \left.\frac{\partial^2 v(\mathbf{p}, x)}{\partial \log x \, \partial \log p_i}\right|_{(\mathbf{p}_1, x_1)} \Bigg/ \left.\frac{\partial v(\mathbf{p}, x)}{\partial \log x}\right|_{(\mathbf{p}_1, x_1)} \\ &\quad + s_{i1} \left.\frac{\partial^2 v(\mathbf{p}, x)}{\partial (\log x)^2}\right|_{(\mathbf{p}_1, x_1)} \Bigg/ \left.\frac{\partial v(\mathbf{p}, x)}{\partial \log x}\right|_{(\mathbf{p}_1, x_1)}. \end{aligned}$$

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

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

Rearranging and evaluating at $(\mathbf{p}_1, x_1)$ yields

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

Moreover, since $\log \mu(\mathbf{p}_1, x) = \log x$, we also have

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

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

$$\frac{\partial^2 \log \mu(\mathbf{p},x)}{\partial \log p_i\, \partial \log p_j} = \frac{\partial \log \mu(\mathbf{p},x)}{\partial v}\frac{\partial^2 v(\mathbf{p},x)}{\partial \log p_i\, \partial \log p_j} + \frac{\partial^2 \log \mu(\mathbf{p},x)}{\partial v^2}\frac{\partial v(\mathbf{p},x)}{\partial \log p_i}\frac{\partial v(\mathbf{p},x)}{\partial \log p_j}. \quad \text{(B.10)}$$

Evaluating (B.10) at $(\mathbf{p}_1, x_1)$ and using (B.5), (B.8), and $\frac{\partial v}{\partial \log p_i} = -s_i \frac{\partial v}{\partial \log x}$ yields

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

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

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

Using the expression for $\partial s_j/\partial \log x$ derived above, one obtains

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

Evaluating at $(\mathbf{p}_1, x_1)$ gives

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

□

Proof of Proposition 1. Define

$$\begin{aligned}\log Q(u_1, u_T; \mathbf{p}_1) &\equiv \log e(\mathbf{p}_1, u_T) - \log e(\mathbf{p}_1, u_1)\\ &= \log e(\mathbf{p}_1, v(\mathbf{p}_T, x_T)) - \log x_1\end{aligned}$$

Define $\mu(\mathbf{p}, x) \equiv e(\mathbf{p}_1, v(\mathbf{p}, x))$ and let $\Delta p_i = p_{iT}/p_{i1}$ and $\Delta x = x_T/x_1$. Since $\log Q(u_1, u_T; \mathbf{p}_1) = \log \mu(\mathbf{p}_T, x_T) - \log \mu(\mathbf{p}_1, x_1)$ and $\log \mu(\mathbf{p}_1, x_1) = \log x_1$, a second-order Taylor expansion of $\log \mu(\mathbf{p}_T, x_T)$ around $(\mathbf{p}_1, x_1)$ and subtraction of $\log \mu(\mathbf{p}_1, x_1)$ yields:

$$\begin{aligned}\log Q(u_1, u_T; \mathbf{p}_1) = &\sum_i \frac{\partial \log \mu(\mathbf{p}_1, x_1)}{\partial \log p_i} \log \Delta p_i + \frac{\partial \log \mu(\mathbf{p}_1, x_1)}{\partial \log x} \log \Delta x\\ &+ \frac{1}{2} \sum_{i,j} \frac{\partial^2 \log \mu(\mathbf{p}_1, x_1)}{\partial \log p_i \log p_j} \log \Delta p_i \log \Delta p_j + \sum_i \frac{\partial^2 \log \mu(\mathbf{p}_1, x_1)}{\partial \log p_i \log x} \log \Delta x \log \Delta p_i\\ &+ \frac{1}{2} \frac{\partial^2 \log \mu(\mathbf{p}_1, x_1)}{\partial \log x^2} (\log \Delta x)^2 + R^{(2)}\end{aligned}$$

where $R^{(2)} = O(|| \log \Delta \mathbf{p}, \log \Delta x||^3)$. By Lemma 1, we have:

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

Hence,

$$\begin{aligned}\log Q(u_1, u_T; \mathbf{p}_1) = &\log \Delta x - \sum_i s_{i1} \log \Delta p_i\\ &- \frac{1}{2} \sum_{i,j} \left(\left.\frac{\partial s_i}{\partial \log p_j}\right|_{(\mathbf{p}_1, x_1)} - s_{i1} \left.\frac{\partial s_j}{\partial \log x}\right|_{(\mathbf{p}_1, x_1)} \right) \log \Delta p_i \log \Delta p_j\\ &- \sum_i \left.\frac{\partial s_i}{\partial \log x}\right|_{(\mathbf{p}_1, x_1)} \log \Delta p_i \log \Delta x + R^{(2)}\end{aligned}$$

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,

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

and, at fixed u_1,

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

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

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

Note also that

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

Hence,

$$\begin{aligned}
\log Q(u_1, u_T; \mathbf{p}_1) =& \log \Delta x - \sum_i s_{i1} \log \Delta p_i \\
& - \Bigg(\frac{1}{2}\sum_{i,j} \frac{\partial s_i}{\partial \log p_j}\Big|_{(\mathbf{p}_1,u_1)} \log \Delta p_i \log \Delta p_j - \frac{1}{2}\sum_{i,j} s_{i1}s_{j1}(\eta_i - 1)\log \Delta p_i \log \Delta p_j \\
& - \frac{1}{2}\sum_{i,j} s_{i1}s_{j1}(\eta_j - 1)\log \Delta p_i \log \Delta p_j + \sum_i s_{i1}(\eta_i - 1)\log \Delta p_i \log \Delta x \Bigg) + R^{(2)} \\
=& \log \Delta x - \sum_i s_{i1} \log \Delta p_i \\
& - \frac{1}{2}\sum_{i,j} \frac{\partial s_i}{\partial \log p_j}\Big|_{(\mathbf{p}_1,u_1)} \log \Delta p_i \log \Delta p_j \\
& - \Big(\log \Delta x - \sum_j s_{j1} \log \Delta p_j \Big) \sum_i s_{i1}(\eta_i - 1)\log \Delta p_i + R^{(2)}.
\end{aligned}$$

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

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

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

$$\log\Big(\sum_i s_{i1}\Delta p_i\Big)$$

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

$$\log Q(u_1, u_T; \mathbf{p}_1) \approx \log \Delta x - \underbrace{\log \Big(\sum_i s_{i1} \Delta p_i \Big)}_{\text{Exposure}}$$
$$+ \underbrace{\Big[\log \Big(\sum_i s_{i1} \Delta p_i \Big) - \sum_i s_{i1} \log \Delta p_i - \tfrac{1}{2} \sum_{i,j} \tfrac{\partial s_i}{\partial \log p_j} \big|_{(\mathbf{p}_1, u_1)} \log \Delta p_i \log \Delta p_j \Big]}_{\text{Substitution}}$$
$$+ \underbrace{\Big[- \Big(\log \Delta x - \sum_j s_{j1} \log \Delta p_j \Big) \mathrm{Cov}_s \big(\log \Delta p_i, \eta_i \big) \Big]}_{\text{Income effects}}.$$

This is precisely the decomposition stated in Proposition 1. □

Remark 1. *Using the identity*

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

and comparing with the decomposition above, we see that

$$\log P(\mathbf{p}_1, \mathbf{p}_T; u_1) \approx \textit{Exposure} - \textit{Substitution}, \qquad \log P(\mathbf{p}_1, \mathbf{p}_T; u_1) - \log P(\mathbf{p}_1, \mathbf{p}_T; u_T) \approx \textit{Income effects}.$$

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

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

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

$$\log P(\mathbf{p}_1, \mathbf{p}_T; u_1) - \log P(\mathbf{p}_1, \mathbf{p}_T; u_T) \approx - \Big(\log \Delta x - \sum_j s_{j1} \log \Delta p_j \Big) \mathrm{Cov}_s \big(\log \Delta p_i, \eta_i \big).$$

B.2 Jaravel–Lashkari algorithm

Here we outline the algorithm developed in Jaravel and Lashkari (2024) for approximating a cost-of-living index under arbitrary non-homothetic preferences.

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 function for an arbitrary cardinalisation of utility. Define period-1-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 nonparametrically 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),

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

and use the derivative of the fitted function to approximate

$$\Lambda_t(u_{t-1}) \approx \sum_{\tau<t} g_\tau'(\log u_{\tau-1}).$$

This first-order algorithm therefore adjusts a geometric-Laspeyres index for non-homotheticities.

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

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

Second,

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

where the first-order approximation is taken (i) in prices around $(\mathbf{p}_t, u_t)$ and (ii) in utility around $(\mathbf{p}_{t-1}, u_{t-1})$.

Taking the arithmetic average of these two expressions gives

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

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

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

As in the first-order case, the functions $\Lambda_{t-1}(\cdot)$ and $\Lambda_t(\cdot)$ are obtained from flexible cross-sectional regressions of household-specific price indexes on lagged money-metric utility and household characteristics; see Jaravel and Lashkari (2024) for implementation details.

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, to construct the cost-of-living index evaluated at initial-period utility u_1.

Implementation with scanner data. We implement the 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-index objects (geometric Laspeyres/Törnqvist) are computed over the set of products with observed prices in both periods (the overlap set $\Omega_{t-1,t}$), and households that do not purchase a product simply have zero share for that item. We implement $g_t(\cdot)$ using polynomial regressions, estimated separately by period from cross-household variation in lagged money-metric utility and household composition controls. Identification of $\Lambda_t(\cdot)$ is based on the maintained assumption that, conditional on $\mathbf{X}_h$, the non-homothetic correction is common across household; cross-sectional variation in $\log u_{ht-1}$ then disciplines the mapping from lagged money-metric utility to non-homotheticity.

C Results appendix

C.1 Inflation exposure with household-specific prices

In Appendix A.1 we show that the gap in prices paid for identical products between less-well-off and better-off households is modest and changes little over time. This suggests differences in price gaps for identical products are modest and fairly stable over time, suggesting that within-product price dispersion is unlikely to be an important driver of inflation inequality over 2021Q3 to 2023Q3.

Figure C.1: *Inflation inequality: household-specific prices*

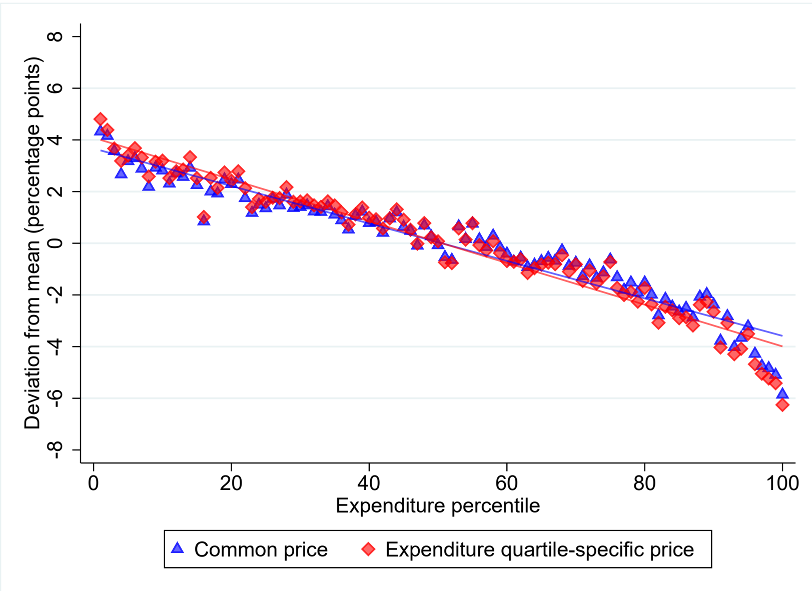

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2021–2023). Figure illustrates the relationship between ninth-quarter cumulative inflation exposure (2021Q3–2023Q3) and a household's percentile in the 2021 expenditure distribution, with a marker for each percentile and a line of best fit. Cumulative inflation is measured using a Laspeyres index. The blue series uses common prices across households; the red series uses expenditure–quartile-specific prices as defined in the text.

Figure C.1 confirms this by comparing the relationship between cumulative household inflation over 2021Q3–2023Q3 based on common prices (the blue line and markers; replicating Figure 4.2(a)) with an alternative series that recomputes prices separately within expenditure quartiles (the red line and markers).[19] (the red line and markers). The gradient is slightly steeper with quartile-specific prices, implying that differential prices paid for identical goods slightly reinforce, rather than attenuate, the substantial inflation inequality documented in the main text.

C.2 Product churn

The Laspeyres index in equation (3.2) fixes each household's consumption basket at its initial composition, so it does not reflect the effects of product entry and exit. In Figure C.2 we show inflation inequality computed using two chained variants of the Laspeyres index. Panel (a) shows an index chained year-on-year (comparing 2021Q3, 2022Q3 and 2023Q3), while panel (b) presents an index chained quarter-on-quarter over nine quarters from 2021Q3 to 2023Q3. Because these indexes compare price growth across consecutive years or quarters, they are less affected by product churn. The figure shows that they exhibit essentially the same pattern of inflation inequality as in Figure 4.2(a).

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

(a) Calendar-quarter-to-calendar-quarter (b) Quarter-to-quarter

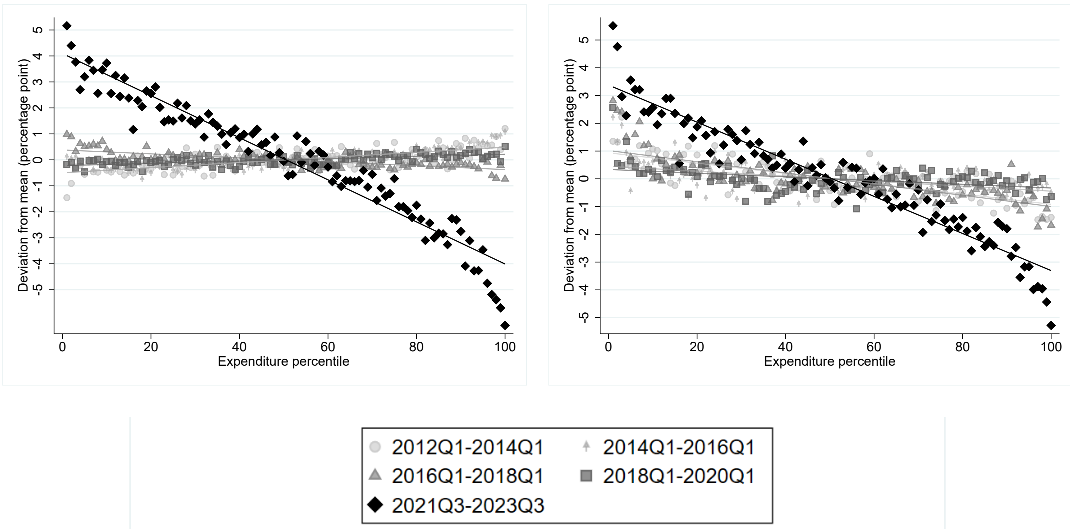

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2012–2023). The figure illustrates the relationship between ninth-quarter cumulative inflation and a household's percentile in the expenditure distribution, with a marker for each percentile and a line of best fit. Households are assigned to expenditure percentiles based on their equivalised spending over the initial calendar year of the relevant nine-quarter period. Cumulative inflation is based on a Laspeyres index chained across 2021Q3, 2022Q3 and 2023Q3 in panel (a) and on a Laspeyres index chained quarter-to-quarter in panel (b).

However, chained indexes do not fully address product churn, because even across successive quarters there can be a non-trivial amount of product entry and exit. This can give rise to a new-product-variety bias, since the indexes fail to account for welfare losses from product exits and gains from entry. In our data, 7.1% of 2021Q3 aggregate spending is

[19] Specifically, we compute the expenditure-quartile-specific price for (i,t) as $p^r_{it} = \frac{\sum_{h\in\mathcal{R}_r} x_{hit}}{\sum_{h\in\mathcal{R}_r} q_{hit}}$ where $\mathcal{R}_r$ denotes the set of households that belong to quartile r of the calendar-year-specific annual equivalised expenditure distribution.

on products not available in 2022Q3, and 8.6% of 2022Q3 aggregate spending is on products that were not available in 2021Q3.[20] The analogous comparison between 2022Q3 and 2023Q3 results in 5.1% of 2022Q3 spending on exiting products and 6.5% of 2023Q3 spending on entering products.

The Feenstra-corrected CES price index (Feenstra 1994) provides a convenient way of quantifying the importance of product entry and exit. If the household has CES preferences, their cost-of-living index over period s to t takes the form $P^{CES}_{h,(s,t)} = \left[\prod_i \left(\frac{p_{it}}{p_{is}}\right)^{\phi_{hi,(s,t)}}\right] \times \mathcal{F}_{h,(s,t)}$, where $\phi_{hi,(s,t)}$ is a weight that depends on the household's period-s and period-t spending shares over products available in both comparison periods.[21], and $\mathcal{F}_{h,(s,t)} = \left(\frac{1-s^N_{ht}}{1-s^X_{hs}}\right)^{\frac{1}{\sigma-1}}$ is the correction for the influence of entering and exiting products. This correction depends positively on the share of initial period spending the household allocates to exiting goods, s^X_{hs} (the cost of living rises if products the household consumes disappear), and negatively on the share of final-period spending allocated to entering goods, s^N_{ht} (since consumption of new products lowers the cost of living). The sensitivity of the cost-of-living index to net product entry depends on the elasticity of substitution $\sigma > 0$: when σ is low, entering and exiting products do not have close substitutes, so net entry has a relatively large effect on the cost of living.

In Figure C.3(a) we plot the inflation-inequality gradient computed using a CES price index with no Feenstra correction. The pattern of inflation inequality is very similar to that shown in Figure 4.2(a). Panel (b) plots the difference between entry and exit shares, $s^N_{ht} - s^X_{hs}$, while panel (c) shows the adjustment to the baseline CES index implied by net product entry by reporting the difference between the Feenstra-corrected CES and non-corrected CES index (shown in panel (a)) for a range of values of σ. In each case we report results for each percentile of the expenditure distribution. The figure shows that households allocate, on average, a higher share of their spending to entering than exiting products, which acts to lower cost-of-living increases. However, the magnitude of the effect is very similar across the expenditure distribution.

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

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

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

(a) Not adjusted for net entry

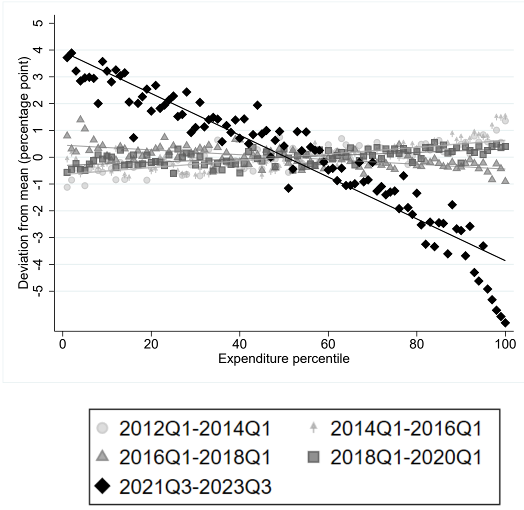

(b) Average net entry

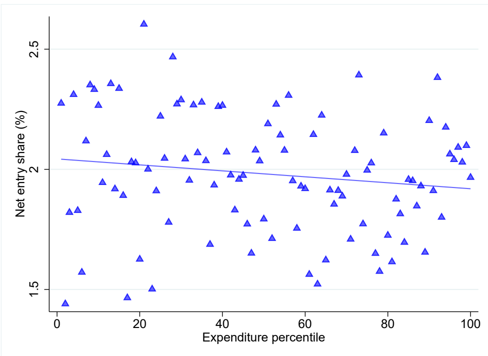

(c) Net entry adjustment

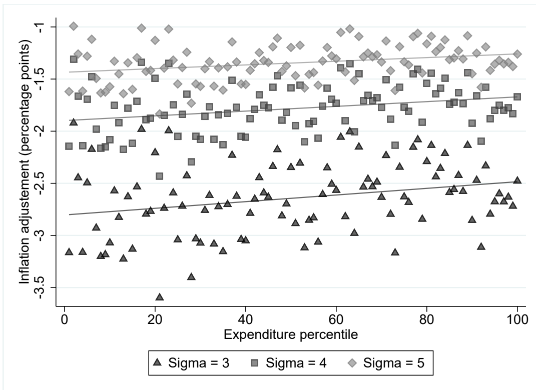

Notes: Authors' calculations using Numerator's Take Home Purchase Panel (2021–2023). Panel (a) shows the relationship between nine-quarter cumulative inflation, based on a CES price index with no Feenstra-correction, and expenditure percentiles. Panel (b) plots the average difference in entry and exit shares, $s^N_{ht} - s^X_{hs}$*, across expenditure percentiles, averaging over the periods 2021Q3–2022Q3 and 2022Q3–2023Q3. Panel (c) plots, across expenditure percentiles, the difference between the Feenstra-corrected CES index and uncorrected CES index, for three different values of CES elasticity of substitution* σ*. Households are allocated to expenditure percentiles based on their equivalised spending in 2021. The index is chained across 2021Q3, 2022Q3 and 2023Q3.*

C.3 Inflation inequality in the broader consumption basket

This appendix describes how we construct income-group-specific Consumer Price Indices (CPI) that combine LCFS expenditure data, official CPI price indices, and our scanner-data-based FMCG inflation measures. We construct two alternative indices: (i) a baseline index that uses *common* FMCG price relatives across all income groups, and (ii) an alternative index that incorporates *group-specific* FMCG price relatives by income decile, derived from the Numerator scanner data.

Household expenditure data come from the Living Costs and Food Survey (LCFS), with spending categories classified using the COICOP system.[22] Official price indices P_{ct}^{CPI} and national CPI weights w_{ct}^{CPI} are taken from ONS CPI data at the COICOP class level. For FMCG categories, we supplement these with FMCG price relatives from the Numerator data, constructed either at the aggregate level (common across income groups) or separately by income decile.

We treat as FMCG the COICOP 3-digit classes listed in Table C.1.

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

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

$$X_{gct} \equiv \sum_{i \in G_g} x_{ict},$$

where x_{ict} is household i's expenditure on category c in year t and G_g is the set of households in decile g. The corresponding LCFS budget share is

$$w_{gct}^{\text{LCFS}} \equiv \frac{X_{gct}}{\sum_{c'} X_{gc't}}.$$

[22]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).

These LCFS shares do not exactly match the national CPI weights w_{ct}^{CPI} because of survey under-reporting and differences between LCFS definitions and the national accounts. To construct group-specific weights that are consistent with the CPI, we re-scale LCFS expenditure within each COICOP class. Let

$$X_{ct} \equiv \sum_{g} X_{gct}, \qquad X_t \equiv \sum_{c} X_{ct}.$$

For each (g, c, t) we define adjusted expenditure

$$x_{gct}^{*} \equiv w_{ct}^{\text{CPI}} \times \frac{X_{gct}}{X_{ct}} \times X_t.$$

By construction, $\sum_g x_{gct}^{*} = w_{ct}^{\text{CPI}} X_t$ for every c, so the re-scaled group-level expenditures aggregate to the official CPI weights within each category while preserving the LCFS distribution across income groups.

We then define group-specific CPI-consistent budget shares as

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

For later use, we also define the group-specific FMCG share in year t as

$$w_{g\text{FMCG}t}^{*} \equiv \sum_{c \in \text{FMCG}} w_{gct}^{*},$$

Baseline CPI by income group. Let P_{ct}^{CPI} denote the official CPI index for category c in period t, and let $t \in \{1, 2, 3\}$ where $t = 1$, $t = 2$ and $t = 3$ correspond to 2021Q3, 2022Q3 and 2023Q3, respectively. For non-FMCG categories $c \notin$ FMCG and $t \in \{1, 2\}$ we define the category-level CPI price change

$$p_{ct} \equiv \frac{P_{ct+1}^{\text{CPI}}}{P_{ct}^{\text{CPI}}}.$$

From the Numerator scanner data, we construct a *common* FMCG Laspeyres price relative p_t^{FMCG} between 2021Q3 and 2022Q3 ($t = 1$) and between 2022Q3 and 2023Q3 ($t = 2$), 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 Laspeyres inflation for group g over 2021Q3-2023Q3 is then

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

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

$$\tilde{\Pi}_d \equiv \Pi_d - \Pi_1.$$

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

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

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

The corresponding decile d gap with the bottom decile is

$$\tilde{\Pi}_d^{\text{alt}} \equiv \Pi_d^{\text{alt}} - \Pi_1^{\text{alt}}.$$

Results. Figure C.4 summarises the contribution of FMCG cheapflation to overall inflation inequality across the full CPI basket. Panel (a) shows that FMCG budget shares decline strongly with income: in LCFS 2021, the bottom income decile devotes around 23.3% of total expenditure to FMCG, compared to about 17.5% in the top decile, a gap of 5.7 percentage points.

Panel (b) compares the decile gap, relative to the bottom income decile, in nine-quarter cumulative inflation under the two index constructions. In the *common-FMCG benchmark*, FMCG categories share the same inflation across all income groups, given by the aggregate FMCG Laspeyres price relatives p_t^{FMCG}, while non-FMCG categories follow class-level CPI inflation. In this case, the cumulative inflation of the bottom decile exceeds that of the top decile by 2.16 percentage points. In the *FMCG-adjusted index*, we replace the common FMCG inflation with group-specific FMCG price relatives $p_{g\text{FMCG}t}$ constructed from the Numerator data. Under this specification, the corresponding bottom–top gap rises to 3.03 percentage points. Thus, differential FMCG inflation accounts for about 0.87 percentage points of the total 3.03 percentage point bottom–top gap, increasing the inter–decile gradient by 40.3% relative to the common-FMCG benchmark.

Panel (b) also shows that, even in the common-FMCG benchmark (which ignores within-FMCG heterogeneity across income groups), there is substantial inflation inequality over 2021Q3–2023Q3. Panel (c) repeats the exercise excluding residential energy (gas and electricity) from the CPI basket. Once energy is removed, the inter–decile gradient in the common-FMCG benchmark becomes much smaller, indicating that most of the inflation inequality captured by the standard CPI over this period reflects the large spike in household energy prices during the European energy crisis.

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

(a) FMCG budget share

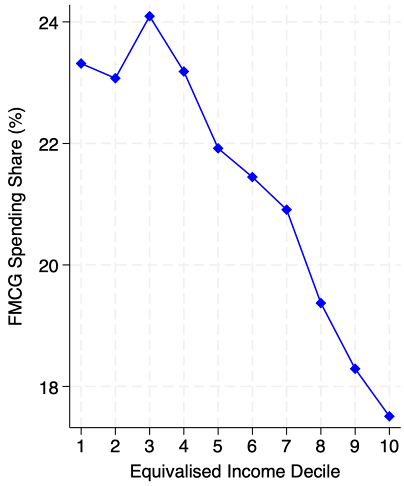

(b) Inter–decile inflation differences

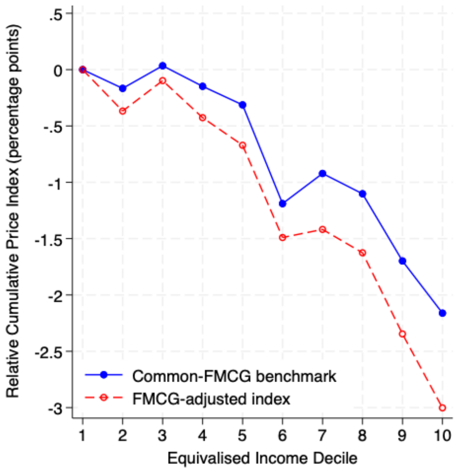

(c) Inter–decile inflation differences (no energy)

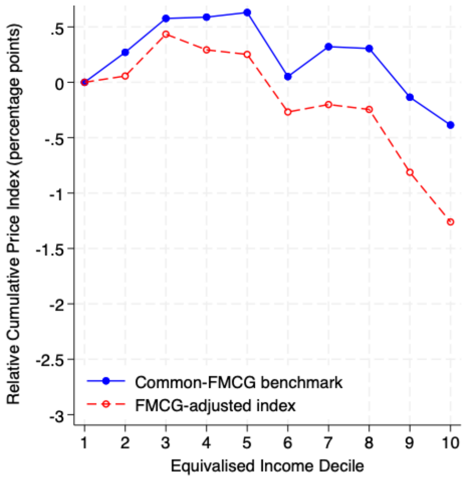

Notes: Authors' calculations using the Living Costs and Food Survey (LCFS) 2021, ONS CPI data, and Numerator's Take Home Purchase Panel (2021–2023). Panel (a) plots the share of total expenditure allocated to FMCG goods by equivalised income decile. Panel (b) reports cumulative inflation over 2021Q3–2023Q3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