IFS

Institute for Fiscal Studies

Tao Chen
Peter Levell
Martin O'Connell

Working paper

25/21

Measuring cost of living inequality during an inflation surge

UKRI Economic and Social Research Council

Measuring cost of living inequality during an inflation surge

Tao Chen, Peter Levell and Martin O'Connell*

May, 2025

Abstract

We provide new evidence that inflation inequality surged during the 2021–2023 cost-of-living crisis, driven by systematically higher price growth for lower-quality goods disproportionately consumed by poorer households. While substitution in response to relative price changes helped mitigate cost-of-living increases, it did not reverse historically high cost-of-living inequality. Declining living standards drove many households to trade down to lower-quality goods, further exposing them to the strongest price increases. Our findings have important implications for cost-of-living measurement and policymaking in an inflationary environment and underscore rising political discontent, as lower-income households face the steepest rise in their living costs.

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, Áureo de Paula and Imran Rasul for helpful comments. The contents of this publication are the authors' own analysis and findings. The analysis and findings have been undertaken by the authors on the basis of their own work using data supplied by the Kantar Take Home Purchase Panel (all use is subject to Kantar UK Limited's terms and conditions) and Kantar does not represent any endorsement of the efficacy or accuracy of authors' analysis and findings. 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 emerged as a central economic concern following the COVID-19 pandemic, as many countries have experienced persistently high price growth driven by supply disruptions, geopolitical instability, and shifts in global trade policy. However, the burden of inflation is not equally shared across households, with differences in consumption patterns meaning inflationary pressures can disproportionately affect certain income groups, potentially exacerbating economic inequality.

In this paper, we provide novel evidence on how the return of high and volatile aggregate inflation has widened inflation-driven inequality. We show that systematically larger price increases for lower-quality, cheaper goods—disproportionately consumed by poorer households—drove an unprecedented divergence in inflation exposure. A key feature of this environment is the interaction between rapid relative price changes, which induce substitution toward goods with weaker price growth, and declining purchasing power, which drives households to trade down from higher- to lower-quality necessities. We show how these behavioural adjustments interact, shaping disparities in cost-of-living increases across the income distribution.

We leverage household panel data on UPC-level purchases of fast-moving consumer goods during the 2021–2023 cost-of-living crisis—a period marked by the most rapid inflation in wealthy countries in over three decades, alongside stagnant incomes. Our data allow us to track household consumption baskets over time at a granular level, distinguishing quality differences within narrowly defined product categories. This enables us to measure household-level cost-of-living changes across a key segment of the economy that strongly influences household inflation expectations (D'Acunto et al., 2021), while capturing substitution and trading-down effects that aggregate data, including official inflation statistics, often obscure.

We begin by documenting inflation using a Laspeyres index, the measure underlying the official Consumer Price Index (CPI), which captures households' exposure to price changes based on initial consumption baskets. Over the nine quarters beginning in 2021Q3, average cumulative inflation was 26.2%. However, this aggregate figure masks substantial heterogeneity, with household-level inflation rates exhibiting a standard deviation of 5.4 percentage points. We show that variation in inflation is systematically related to household income measures. For example, households in the top decile of the 2021 equivalised expenditure distribution experienced inflation that was 7.4 percentage points lower than those in the bottom decile. This inflation-income gradient is unprecedented in the UK. Moreover, the magnitude of inflation inequality we document exceeds that observed in the US during periods of relative aggregate price stability (see Jaravel, 2021).

Using a decomposition of inflation inequality into contributions from across and within product group heterogeneity in consumption baskets, we show that differences within narrowly defined product categories account for 54.5% of inflation dispersion across households and are entirely responsible for the inflation-income gradient. This finding suggests that the income-group-specific inflation rates published by statistical agencies—based solely on differences in spending across product categories—likely understate inflation inequality, especially during periods of high inflation.

To understand why poorer households face steeper inflation, we segment the product space into a 10-rung *quality ladder*. Specifically, we caterogise products based on within-category price differences after accounting for nonlinear pricing across different package sizes of the same brand. We show price growth is stronger for products on lower rungs of the quality ladder. For instance, over just nine quarters, products on the bottom two rungs exhibit average price increases of 34.2%, compared to 18.3% for those on the top two rungs. This patterns leads to a compression of the price distribution, a phenomenon recently labelled cheapflation (Cavallo and Kryvtsov, 2024). Worse-off households allocate a disproportionate share of their spending to lower-quality products, making them more exposed to rising prices. Moreover, by comparing the recent inflation spike to earlier periods of price stability, we show that the rise in inflation inequality is driven by this new pattern of price dynamics across the quality ladder, rather than shifts in consumption basket composition across households with different levels of resources.

When faced with significant relative price changes, households have an incentive to substitute toward goods with lower price growth to mitigate increases in their cost-of-living. When aggregate inflation outpaces nominal income growth, reducing purchasing power, households can also limit reductions in their living standards by trading down from luxuries to necessities. During inflationary periods characterised by falling real incomes and the most rapid price growth occurring in low-quality products, these two forces—substitution and trading down—can act in opposing directions.

The extent to which substitution effects lower cost-of-living increases is captured by the difference between a Konüs (1939) cost-of-living index—which measures changes over time in the cost of maintaining fixed living standards—and the Laspeyres index, which captures cost-of-living changes under no-substitution (Leontief) preferences. We construct an approximation to the cost-of-living index that is robust to non-homothetic preferences using the method recently developed by Jaravel and Lashkari (2024). This approach adjusts a standard homothetic price index to remove the influence of spending reallocation driven by changes in living standards.

We show that household substitution responses reduced the cost of maintaining pre-inflation surge living standards by an average of 1.5 percentage points, meaning, on average, the cost of living rose by 24.7% over the nine quarters beginning in 2021Q3. The magnitude of substitution responses is similar across the equivalised expenditure distribution. However, because worse-off households face greater inflation exposure, the percentage reduction in their cost-of-living increase due to substitution is smaller than that of better-off households—for instance, 5.5% in the bottom expenditure quartile versus 6.9% in the top quartile. Using a theory-consistent decomposition of substitution responses into within-segment and between-segment components, we find that substitution effects were quantitatively significant across all product segments. For worse-off households, substitution responses mitigated cost-of-living increases more than for better-off households in some segments—such as dairy, prepared food, and confectionery. However, this effect was offset by better-off households allocating a lower share of spending to segments where substitution effects are weaker and their larger between-segment substitution effect.

Despite the rise in the relative price of lower-quality goods, we show that many households switch toward them. Leveraging within-household variation, we provide evidence that this trading-down behaviour results from an income effect driven by reduced purchasing power. Specifically, households experiencing the largest declines in living standards reallocated their spending most strongly down the quality ladder. This generates a quantitatively significant non-homotheticity bias—up to two-thirds the magnitude of the substitution effect—on superlative price indexes typically used to measure cost-of-living changes. Standard cost-of-living indexes assume households adjust spending only through substitution, meaning any income-effect-induced trading down can bias these measures. While income-effect-driven trading down does not directly enter the cost-of-living index, it serves as a mechanism for households to lower the effective price level of their consumption basket. We find that the reduction in the average within-category final-period price paid due to adjustments to consumption baskets is 3.3% on average and 5.1% for the tenth of households experiencing the largest decline in living standards. However, the extent to which households can exploit this adjustment is constrained by the rising relative price of lower-quality goods. Additionally, it also leaves household consumption baskets more exposed to future inflationary surges driven by rising prices of necessities.

Our work contributes to the growing literature on how differences in consumption baskets shape inflation inequality.[1] One strand of this literature uses granular household scanner data (see Kaplan and Schulhofer-Wohl, 2017; Jaravel, 2019; Argente and Lee, 2021) and shows that, in the US, during periods of relative price stability, lower-income households experience higher inflation rates. A key strength of these studies—and ours—is their ability to measure spending at the UPC level, which is crucial for uncovering inflation inequality in our context. A complementary strand of research, dating back to Michael (1979), uses household expenditure surveys to measure inflation inequality across the full consumption basket. Recent work by Jaravel (2024) replicates official CPI methodology and provides evidence that differences in inflation rates across income groups have contributed to rising long-run inequality. A second, nascent strand of literature characterises recent patterns of price dynamics. Cavallo and Kryvtsov (2024) were the first to document faster price growth among cheaper products, a phenomenon observed across ten countries, including Canada, France, Germany, and the US. Sangani (2023) provides evidence that this pattern is, in part, driven by cost shock pass-through. A third strand highlights the importance of accounting for preference non-homotheticities when measuring long-run changes in living standards (Comin et al., 2021; Atkin et al., 2024; Baqaee et al., 2024; Jaravel and Lashkari, 2024). We contribute to these literatures by showing that high aggregate inflation was fuelled by stronger price growth down the product quality ladder, which drove a significant widening of inflation inequality—a trend largely hidden in official expenditure surveys. Furthermore, we show that households responded through a combination of substitution to relative price changes and income-effect-driven trading down as purchasing power declined. Crucially, we demonstrate that separating these two channels is essential for distinguishing adjustments that mitigate cost-of-living increases from those that reduce the price level of consumption baskets.

The return of high inflation has significant implications for economic policymaking and social stability. We show that rising prices for necessities have been a key driver of cost-of-living increases, which has important implications for central banks seeking to stabalise inflation and the output gap (Olivi et al., 2024). This also helps explain the widespread perception that inflation disproportionately burdens lower-income households (Binetti et al., 2024). While households responded to

[1] Beyond consumption patterns, inflation can also affect inequality through differences in nominal asset holdings (Doepke and Schneider, 2006), the stickiness of nominal income sources (Ferreira et al., 2023), and exposure to asset price changes (Del Canto et al., 2024).

eroding purchasing power by trading down to lower-quality goods, thereby reducing the price level of their consumption basket, this adjustment increases their vulnerability to future necessity-driven inflation surges and may weaken aggregate labor demand, as lower-quality goods tend to be produced with less labour (Jaimovich et al., 2019). Additionally, the asymmetric impact of inflation may have contributed to rising political discontent, fuelling increased support for populist parties (Federle et al., 2024) and heightening concerns about inflation in the 2024 US presidential election (Gillespie, 2024). As the global economy transitions to a more protectionist international trade system, with concerns that this will further fuel inflation (Federal Reserve, 2025), our findings underscore the risk that inflationary environments disproportionately burden less well-off households—a key consideration shaping the trade-offs policymakers face.

2 Data

Scanner data

We use household-level scanner data that is collected by Kantar's Take Home Purchase Panel, a market research firm. 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 on average. Each participating household records all purchased UPCs (or barcodes) using a handheld scanner or mobile phone app and submits receipts electronically or by post. For each transaction, we observe quantity, expenditure, price paid, and UPCs characteristics. Additionally, we have access to 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. During this time, the Consumer Price Index (CPI) for food and non-alcoholic beverages—which, along with alcohol and household goods, comprises fast-moving consumer goods—rose from 103.3 to 133.5 after a prolonged period of stability.[2] Our analysis sample includes 19,030 households that recorded their

[2] 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).

purchases in every year-quarter over this period.[3] We compare inflation inequality over 2021Q3-2023Q3 to four earlier nine-quarter periods: 2012Q1-2014Q1, 2014Q1-2016Q1, 2016Q1-2018Q1 and 2018Q1-2020Q1.[4]

Our dataset offers several advantages for measuring inflation inequality compared to commonly used alternatives. First, it contains up-to-date expenditures, allowing us to construct inflation measures based on up-to-date spending patterns. Second, it tracks households through time, enabling us to construct household-specific inflation rates. Third, expenditure is recorded at the UPC level, meaning our inflation measures reflect spending patterns across narrowly defined products.[5] Our data are more detailed than the budget surveys and price microdata often used to study household-level inflation inequality. The UK Office for National Statistics publishes inflation estimates across household income quintiles using spending data from the Living Costs and Food Survey and prices from approximately 100 CPI 'classes' (e.g., bread and cereals, meat) across all spending categories (Office for National Statistics, 2022). Similarly, household-group-specific inflation estimates from the US based on the Consumer Expenditure Survey use budget shares and price series for 234 'basic items' (Klick and Stockburger, 2021, 2024).[6] While these data sources cover a broader share of households' total spending, they are not well suited for measuring the role played by differential price increases across similar products.

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 the slightly more aggregated combination of brand and package size. Kantar provides highly disaggregated brand information, ensuring that this approach involves minimal aggregation over meaningfully distinct products. Over this pe-

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

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

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

[6]The US Bureau of Labor Statistics does not publish data at this level of granularity. Jaravel (2024) uses publicly available CPI data on 211 'item strata' to construct distributional measures of inflation.

riod, there are approximately 90,000 brand-size pairs, which we refer to as products throughout.[7]

We denote household h's year-quarter t expenditure on good i as x_{hit} and their total period expenditure as $x_{ht} = \sum_i x_{hit}$. We denote household h's year-quarter t quantity purchased of good i quantity by q_{hit}. We measure the period t price of product i as $p_{it} = \frac{\sum_h x_{hit}}{\sum_h q_{hit}}$. In Appendix A.1, we show that lower-income households pay lower prices for identical products (consistent with greater search effort (Aguiar and Hurst, 2007)). However, the price gap between low- and high-income households is modest and relatively stable, narrowing by 0.48 percentage points from 2021Q3 to 2023Q3. This change is minor in context of the overall inflation inequality we document below, and acts, albeit modestly, to reinforce the patterns we document.

We utilise a hierarchical product classification provided by Kantar, 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 liter bottle* belongs to the category *colas* within the *non-alcoholic drinks* segment.

The quality ladder. We further segment the product space by defining a product 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 non-linear 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 measuring the quality ladder.

For each product category, over 2021Q3-2023Q3, we estimate the expenditure-weighted regression:

$$p_{it} = \xi_{b(i)} + \tau_{c(i)t} + \sum_{y} \sum_{l=1}^{3} \alpha_{c(i)y}^{(l)} \mathbb{1}\{t \in y\} \times \text{size}_i^{(l)} + \epsilon_{it}, \tag{2.1}$$

where p_{it} is the price of product i at time t, $\tau_{c(j)t}$ are year-quarter fixed effects, $\xi_{b(j)}$ 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 and compute their adjusted prices in each of these quarters, netting out the size polynomial; $\tilde{p}_{it} = p_{it} - \sum_{l=l} \alpha_{m(i)y}^{(l)} \mathbb{1}\{t \in$

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

$y\} \times \text{size}_j^{(l)}$, $t \leq 4$. We then average adjusted product-level prices across these quarters and use the expenditure-weighted distribution within each product category to define decile boundaries. For all products, we then use equation (2.1) 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 a price that adjusts for category-specific price growth. We construct the quality ladder analogously for the four earlier comparison periods.

Income measure

Our data provide two alternative measures of household economic well-being: (i) total household fast-moving consumer goods expenditure and (ii) current household income, reported in £10,000 bands and top-coded at £70,000. We use total household expenditure as our baseline measure for two reasons. First, current consumption is likely a better proxy for households' lifetime resources than current income (Poterba, 1989; Slesnick, 1993; Meyer and Sullivan, 2023). Second, total expenditure is reported as a continuous measure, whereas current income is both banded and top-coded. In Appendix A.2, we show that total expenditure strongly correlates with banded income. We also show that our main conclusions hold when using the banded household income measure.

For both total expenditure and banded income (using the band midpoint), we apply an equivalisation using the OECD modified scale (Hagenaars et al., 1994).[8] In each nine-quarter period, we use equalized expenditure from the first calendar year to assign households to 100 or 10 fixed bins based on expenditure percentiles or deciles. Throughout we refer to these as expenditure percentiles (or deciles).

Descriptive evidence

Our objective is to measure inflation inequality in an inflationary environment, and to characterise and quantify the significance of behavioural responses. In Figure 2.1, we provide 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 lower-income households, and that households with the largest declines in spending shift more sharply down the quality ladder. This points towards the importance of price dynamics across the

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

quality ladder and spending adjustments driven by declining purchasing power in understanding patterns of 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). It shows that proportional price increases were substantially higher for lower-quality products than those at the top of the quality 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%. In Appendix A.3, we show that over this period, there was compression in the category-level price distribution, and that this pattern is also evident in the micro-data underlying the UK CPI, including for fast-moving consumer goods, clothing and leisure goods. Cavallo and Kryvtsov (2024) term this phenomenon cheapflation. In contrast, price growth was relatively flat across the quality ladder in the set of preceding nine-quarter periods (as shown by the 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. It show that higher-income households consistently purchase higher-quality products on average. This relationship remains remarkably stable across all nine-quarter periods.

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

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

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

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

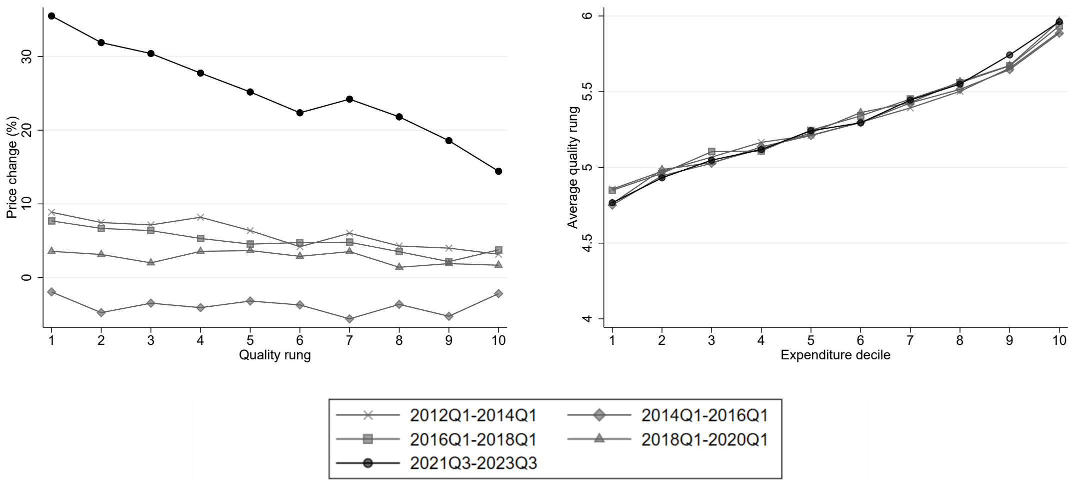

Bottom and top quality rung Engel curve:

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

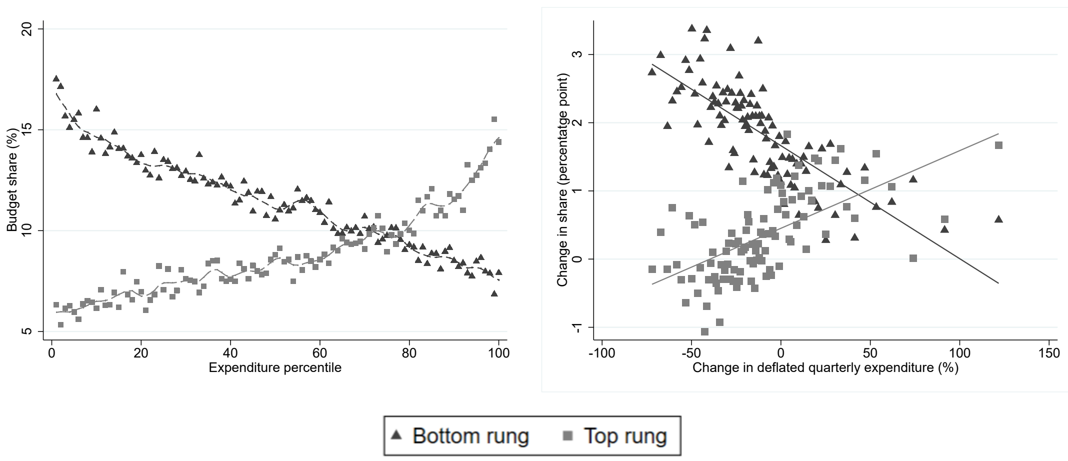

Notes: Authors' calculations using Kantar's Take Home Purchase Panel (2012-2023). Panel (a) reports change in average price over the nine-quarter period for products on each rung of the quality ladder, weighted by initial-period aggregate spending shares. Panel (b) reports the average quality rungs of 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 between 2021Q1 to 2023Q3 for each percentile of the distribution of percent changes in deflated quarterly expenditure over this time. Expenditure changes are deflated using a Laspeyres price index.

3 Measurement

We document inflation inequality using several measures, each corresponding to a cost-of-living index under specific restrictions on household preferences. By doing so, we separate the influence of exposure to inflation from household behavioural adjustments in response to changing living costs. Given that we study a period of sharp declines in purchasing power, with evidence suggesting that this triggered spending reallocation via an income effect (Figure 2.1(d)), we employ methods that account for non-homothetic preferences. We compute all our price indexes at the household level. For notational simplicity, we omit a household index.

The cost-of-living index. Let $t = 1, \dots, T$ index time, $i = 1, \dots, I$ index products, $\mathbf{p}_t = (p_{1t}, \dots, p_{It})'$ denote the vector of prices in period t, with x_t denoting the household's total available expenditure. The price indexes we discuss below each have a precise welfare interpretation under specific restrictions on the following consumer problem.

Each period, the household chooses a consumption bundle (also referred to as a consumption basket) $\mathbf{q} = (q_1, \dots, q_I)'$, according to:

$$v(\mathbf{p}_t, x_t) = \max_{\mathbf{q}} U(\mathbf{q}) \text{ subject to } \mathbf{p}_t'\mathbf{q} = x_t,$$

where $U(.)$ is the household's utility function, which we assume is increasing and quasi-concave, and $u = v(\mathbf{p}, x)$ is the corresponding indirect utility function. The expenditure function, defined as $e(\mathbf{p}, u) = v^{-1}(\mathbf{p}, .)$, gives the minimum cost of attaining utility level u at prices $\mathbf{p}$.

The Konüs (1939) cost-of-living index between periods 1 to T, which measures the change in the cost of maintaining a fixed living standard, is given by:

$$\mathbb{P}(\mathbf{p}_1, \mathbf{p}_T; u) = \frac{e(\mathbf{p}_T, u)}{e(\mathbf{p}_1, u)}. \tag{3.1}$$

In general, this index depends on the living standards at which it is evaluated, represented by the utility level u. The quality-of-living index measures the change in the cost of attaining realised utility at fixed prices:

$$\mathbb{Q}(u_1, u_T; \mathbf{p}) = \frac{e(\mathbf{p}, u_T)}{e(\mathbf{p}, u_1)}. \tag{3.2}$$

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

Households can mitigate rises in the cost of living through their substitution responses to relative price changes. Conversely, they can mitigate declines in quality of living by reallocating spending through income effects. We illustrate this graphically in Figure 3.1. Specifically, suppose there are two goods, q_1 and q_2, with the price of good 1 normalised to 1. The budget constraint in period 1 is AB, and the optimal choice is at point X. In period T, the budget constraint shifts to CD, with the new optimal choice at Y. The left-hand graph illustrates the construction of the cost-of-living index, evaluated at initial period's living standards, while the right-hand graph depicts the construction of the quality-of-living index, evaluated at final-period prices. Both indexes depend on observed initial and final-period spending, as well the unobserved level of spending required at final-period prices to reach initial-period living standards.

The cost-of-living index compares initial-period spending, $x_1 = e(p_1, u_1)$, with the expenditure required to maintain the same living standards, u_1, at final-period prices, $e(p_T, u_1)$, corresponding to the consumption bundle Z on the budget constraint $A'B'$. If the household did not adjust their consumption bundle in response to changes in relative prices, their expenditure would be given by $x^{NS} > e(p_T, u_1)$. Thus, the cost-of-living index increases by less than it would in the absence of substitution responses.

The quality-of-living index compares final-period spending, $x_T = e(p_T, u_T)$, with the expenditure necessary, at the same prices, to attain initial-period living standards, $e(p_T, u_1)$. If the household maintained the same *spending share* on each good as in the final period–represented by the ray from the origin to the consumption bundle Y', which passes through the final period optimal choice–then the expenditure required at final-period prices to attain initial-period living standards would be given by $x^H > e(p_T, u_1)$. Thus, the quality-of-living declines by less than it would in the absence of an income effect response.

The final-period-denominated quality-of-living index is analogous to the household's compensating variation for the change in their budget set, expressed in proportional terms rather than as a difference. The index is equal to nominal expenditure growth $\frac{x_T}{x_1}$, divided by the cost-of-living index evaluated at initial-period living standards, which serves as the proportional analog to compensating variation from the price change alone (holding nominal expenditure at its initial-period value). In Appendix B.1, we include figures illustrating the construction of the cost-of-living index, evaluated at final-period living standards, and the quality-of-living index, evaluated at initial-period prices, both of which are closely related to equivalent variation.

The central challenge in measuring quality- and cost-of-living changes is that they depend on the expenditure function evaluated at prices and living standards from different periods (e.g., $e(p_T, u_1)$ in the example in Figure 3.1), which is not directly observable. Index numbers address this challenge by providing first- or second-order approximations to cost-of-living changes. While they entail placing economic restrictions on the structure of preferences or preference heterogeneity, they avoid the need for parametric functional forms. Although our primary focus is on cost-of-living measurement, we emphasise that, given the cost-of-living index, the corresponding quality-of-living index follows directly.

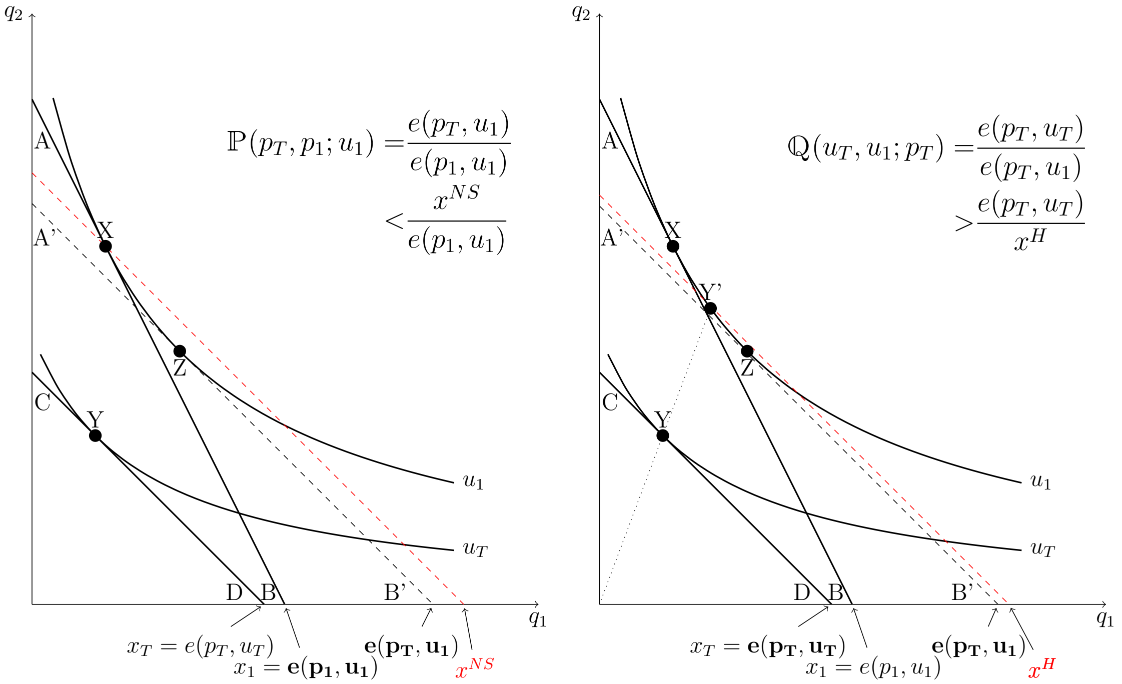

Figure 3.1: *Cost- and quality-of-living indexes and behavioural response*

Leontief preferences. If household preferences are Leontief, their choice behaviour exhibits zero substitution effects, meaning they do not switch between products in response to relative price changes. Additionally, they do not reallocate expenditure in response to changes in living standards. In this case, the cost-of-living takes the form a Laspeyres index:

$$\Pi^L_{1,T} = \sum_i s_{i1} \left(\frac{p_{iT}}{p_{i1}} \right), \tag{3.3}$$

where $s_{i1} = \frac{x_{i1}}{\sum_{i'} x_{i'1}}$ is the household's spending share on product i in period 1.

Although the Laspeyres index corresponds to the cost-of-living index under restrictive assumptions about preferences, it is useful for several reasons. First, it provides a measure of household exposure to inflation, capturing how much the cost of their initial consumption basket changes over time. Second, by comparing it with indexes that allow for substitution effects, we can quantify the significance of these adjustments (corresponding to the difference in cost-of-living index and

x^{NS}/x_1 in Figure 3.1). Third, it serves as the basis for official CPI inflation measurement.[9]

Another useful property of the Laspeyres index is that it be can re-expressed hierarchically. Let Ω^c denote the set of products that belong to product category c and define the household's within-product category spending share of product i and the spending share of category c as

$$s_{it}^c = \frac{x_{it}}{\sum_{i' \in \Omega^c} x_{i't}} \qquad s_{ct} = \frac{\sum_{i \in \Omega^c} x_{it}}{\sum_{c'} \sum_{i \in \Omega^{c'}} x_{it}},$$

respectively. The hierarchical Laspeyres index is given by:

$$\Pi_{1,T}^{HL} = \sum_c s_{c1} \mathbb{P}_{1,T}^c \quad \text{where} \quad \mathbb{P}_{1,T}^c = \sum_{j \in \Omega^c} s_{i1}^c \left(\frac{p_{iT}}{p_{i1}} \right). \tag{3.4}$$

Since $s_{it} = s_{it}^c s_{ct}$ we have $\Pi_{1,T}^{HL} = \Pi_{1,T}^{L}$. This decomposition naturally extends to allow for additional layers. We use the hierarchical index to decompose household-specific exposure to inflation into contributions of differences in spending shares at the segment, category, and product levels, sequentially shutting down heterogeneity at each level.

Homothetic preferences. If preferences are homothetic but otherwise general, substitution responses are unrestricted, but spending reallocation in response to changing living standards is ruled out. Under these conditions, the cost-of-living index is independent of living standards. If price and quantity data are available continuously, the log of the cost-of-living index is given by a Divisia index (Divisia, 1926):

$$\log \Pi_{1,T}^D = \int_1^T \sum_i s_i(t) \mathrm{d} \log p_i(t). \tag{3.5}$$

With discrete data, this can be approximated by chaining a superlative index, where (i) chaining breaks the Divisia integral into discrete summations over time, and (ii) a superlative index provides a second-order approximation to the cost-of-living change over each discrete period (see Diewert, 1976).[10] Using the superlative Törnqvist

[9] In practice, CPI calculations typically allow for limited substitution behaviour by updating weights annually using lagged spending shares and chaining the resulting index.

[10] While, in theory, chaining across the shortest feasible time intervals improves the accuracy of this approximation, in practice, high-frequency chaining can introduce severe chain drift bias (see Ivancic et al., 2011; Fox et al., 2023). This bias arises when expenditure weights in one period are correlated with price changes in another (Diewert, 2022). The issue stems from the fact that the underlying consumer problem, which underpins the cost-of-living interpretation of the index, assumes optimisation of a static and stable utility function—an assumption that becomes less reasonable over very short time intervals, as product storage means purchases observed in the

index, the cost-of-living index is approximated by:

$$\log \Pi_{1,T}^{\mathcal{T}} = \sum_{t=1}^{T-1} \left(\frac{1}{2} \sum_i (s_{it} + s_{it+1}) (\log p_{it+1} - \log p_{it}) \right). \tag{3.6}$$

Non-homothetic preferences. If preferences are non-homothetic, households adjust their spending due to both substitution and income effects, reallocating expenditures due to changes in purchasing power. In this case, with continuous data, the cost-of-living index resembles the Divisia index (equation (3.5)), but with the key difference that observed spending shares are replaced with utility-constant compensated shares. Likewise, the Törnqvist approximation (equation (3.6)) is modified by replacing the observed spending shares, s_{it}, with compensated shares, $\omega_{it}(u)$. Because compensated shares are unobserved, this approximation is not directly implementable

Under non-homothetic preferences, homothetic price indexes provide a biased measure of cost-of-living changes, as they fail to isolate substitution from income effects. For instance, suppose (i) household purchasing power declines over time, causing a reallocation of spending toward necessities (where $s_{it} > \omega_{it}(u_1)$ for $t > 1$) and away from luxuries (where $s_{it} < \omega_{it}(u_1)$ for $t > 1$), and (ii) the relative price of necessities is rising. In this case, by overweighting the rapid price growth of necessities while underweighting the more moderate price growth of luxuries, equation (3.6) overstates the cost-of-living index when evaluated at the initial period's utility level.[11]

Jaravel and Lashkari (2024) have recently developed an algorithm for approximating an arbitrary cost-of-living index. Their approach does not restrict the form of the expenditure function, allowing for non-homotheticities, but it does rule out preference heterogeneity beyond that captured by observables. We briefly outline the algorithm here and provide further details in Appendix B.2.

Let utility be cardinalised in period 1 money-metric terms: the attained utility in period t, u_t, equals the expenditure necessary at period 1 prices to achieve the same living standard as in period t.[12] Log expenditure growth between period 1 and

data and consumption may not align. In practice, we chain across calendar quarters. For example, when constructing a 2021Q3-2023Q3 index, the time indices are $t = \{1, 2, 3\}$, corresponding to 2021Q3, 2022Q3 and 2023Q3. This ensures our results are robust to seasonality in preferences.

[11] Conversely, it understates the cost-of-living index when evaluated at the final period's utility. Diewert (1976) shows that the Törnqvist index computed between period 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. However, this result is not useful when constructing a cost-of-living index at a pre-specified utility level (e.g., initial period), or when comparing more than two periods.

[12] Suppose utility is cardinalised in arbitrary utils, with indirect utility and expenditure functions given by, $\mathbb{u} = \mathbb{v}(\mathbf{p}, x)$ and $x = \mathbb{e}(\mathbf{p}, \mathbb{u})$. The period 1 denominated money-metric utility

2 can be expressed as the sum of a log cost-of-living index and a log quality-of-living index:

$$\begin{aligned}\log x_2 - \log x_1 =& [\log e(\mathbf{p}_2, u_1) - \log e(\mathbf{p}_1, u_1)] + [\log e(\mathbf{p}_2, u_2) - \log e(\mathbf{p}_2, u_1)] \\ \approx& \sum_i s_{i1}(\log p_{i2} - \log p_{i1}) + \frac{\partial \log e(\mathbf{p}_2, u_1)}{\partial \log u}(\log u_2 - \log u_1), \end{aligned} \quad (3.7)$$

where the second line follows from a first-order approximation. The first term on the right-hand side corresponds to the log Geometric-Laspeyres price index, $\log \Pi^{GL}_{2,1} = \sum_i s_{i1}(\log p_{i2} - \log p_{i1})$, which is a function of observables. Since the log cost-of-living index between periods 1 and 2, evaluated at initial utility, satisfies $\log \mathbb{P}(\mathbf{p}_1, \mathbf{p}_2, u_1) = \log e(\mathbf{p}_2, u_1) - \log e(\mathbf{p}_1, u_1)$ and, given that the utility cardinalisation implies $e(\mathbf{p}_1, u_1) = u_1$, the derivative of the expenditure function in equation (3.7) satisfies: $\frac{\partial e(\mathbf{p}_2, u_1)}{\partial \log u} \equiv \Lambda_2(u_1) = \frac{\partial \log \mathbb{P}(\mathbf{p}_1, \mathbf{p}_2, u_1)}{\partial \log u} + 1$. Jaravel and Lashkari (2024) propose recovering this via a non-parametric cross-household regression of a log price index on period 1 log expenditure, controlling for observable preference shifters.[13] Given this, u_2 can then be recovered from equation (3.7). By applying this process recursively, the sequence $(u_1, \ldots, u_T)$ is identified. The log cost-of-living index evaluated at final period utility is then $\log \mathbb{P}(\mathbf{P}_1, \mathbf{P}_T; u_T) = \log x_T - \log u_T$.[14]

As we discuss in the appendix, this approach readily extends to a second-order approximation, which involves adjusting a superlative index for non-homotheticities. Applying the algorithm in reverse produces a cost-of-living index evaluated at the initial period's utility.

4 Results

4.1 Inflation exposure

We begin by describing heterogeneity in inflation using the household-specific Laspeyres index (equation (3.3)), which measures household exposure to cost-of-living rises based on their pre-inflation surge consumption basket.

Figure 4.1(a) presents average household-level cumulative inflation across the nine quarters from 2021Q3 to 2023Q3, alongside earlier nine-quarter periods. Inflation during this recent period was historically high, with a cumulative increase

function is then defined $u = v(\mathbf{p}, x) = \mathbb{e}(\mathbf{p}_1, \mathbb{v}(\mathbf{p}, x))$, and the corresponding expenditure function is $x = e(\mathbf{p}, u) = \mathbb{e}(\mathbf{p}, \mathbb{v}(\mathbf{p}_1, u))$. This is simply a convenient re-labelling of utility levels.

[13] We control for household composition, exploiting variation in inflation within households of the same composition.

[14] This follows from the convenient re-labelling of utility (see previous footnote): in terms of the util-utility cardinalisation: $\log \mathbb{P}(\mathbf{p}_1, \mathbf{p}_T; u_T) = \log x_T - \log u_T = \log \mathbb{e}(\mathbf{p}_T, \mathbb{u}_T) - \log \mathbb{e}(\mathbf{p}_1, \mathbb{u}_T)$.

exceeding 26%. 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 rates across these periods, showing the elevated average inflation in 2021Q3-2023Q3 was accompanied by greater dispersion across households.

Figure 4.1: *Household-level inflation*

(a) Average inflation (b) Inflation dispersion

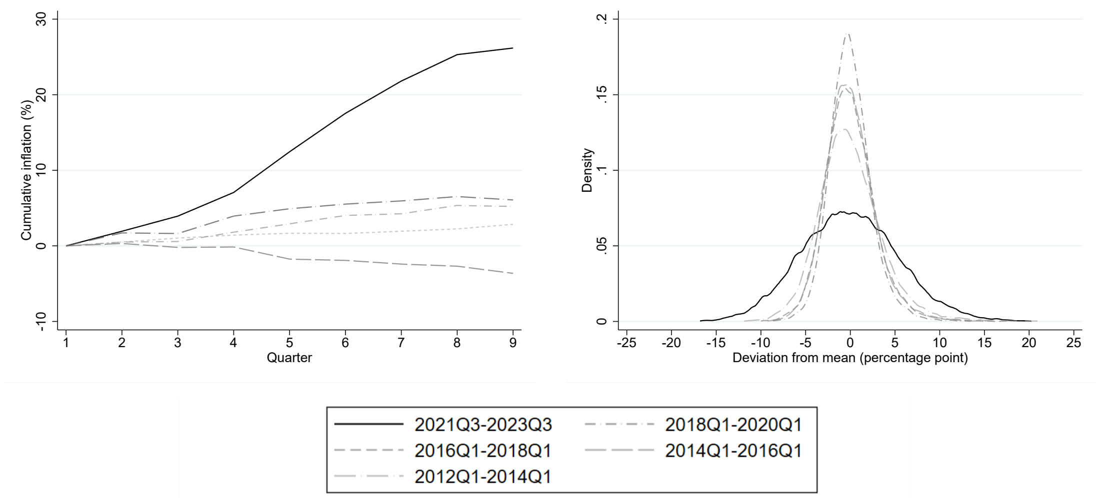

(c) Inflation inequality

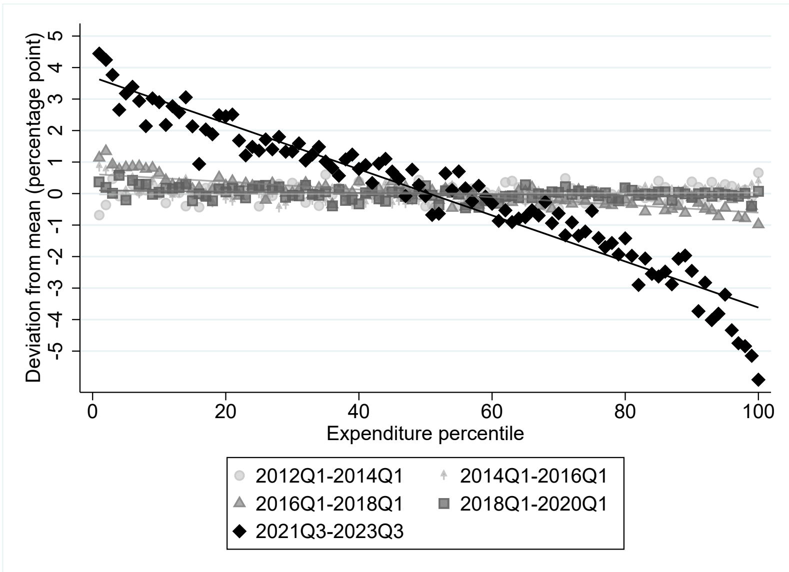

Notes: Authors' calculations using Kantar's Take Home Purchase Panel (2012–2023). Panel (a) presents the average of cumulative inflation across households for each of the nine quarters. Panel (b) display kernel density estimates of the distribution of cumulative inflation in the ninth quarter. Panel (c) 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 measured using a Laspeyres index.

Panel (c) shows the relationship between household-level cumulative inflation and expenditure percentiles, revealing historically large differences in inflation rates across expenditure levels in 2021Q3-2023Q3. Households in the 1^{st} quartile of the expenditure distribution, on average, experienced inflation rates 5.5 percentage points higher than those in the 4^{th} quartile. In contrast, none of the nine-quarter periods between 2012 and 2020 exhibit a significant gradient, Moreover, previous research has found only minor disparities between income groups in the UK over 1976-2014 (Crawford and Oldfield, 2002; Leicester et al., 2008; Adams and Levell, 2014).

We use a three-level hierarchical index (see equation (3.4)) to isolate the contributions of differences in spending shares across households at the segment, category and product level in accounting for variation in household-specific inflation. We do this by switching household-specific spending shares with population-average spending shares at different levels of the index. For example, to isolate the importance of heterogeneity in segment shares, we replace the household-specific category and product shares with their population averages.

In Figure 4.2 we plot the distribution of household-level cumulative inflation over 2021Q3-2023Q3 (panel (a)) and the inflation gradient (panel (b)) using black lines and markers, replicating information from Figure 4.1. 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 across household variation in spending across categories within segments (blue lines and markers). Panel (a) shows that variation in spending across products *within* product categories contributes substantially to the dispersion in household inflation rates, increasing the standard deviation from 1.5 (based only on heterogeneity in segment shares) to 5.4 (after accounting for product-level variation). Panel (b) shows that heterogeneity in spending across products within product categories is *solely* responsible for the gradient in inflation. These findings emphasise the importance of using detailed household spending and price information to study inflation inequality.

Our descriptive evidence in Figure 2.1 shows that price growth is significantly stronger down the quality ladder, with the consumption baskets of worse-off households making them more exposed to these price dynamics. Figures 4.1 and 4.2 demonstrate that these patterns translate into wide inflation inequality, and that they are the sole cause of this inequality.

Figure 4.2: *Household-level hierarchical inflation*

(a) Distribution (b) Inequality

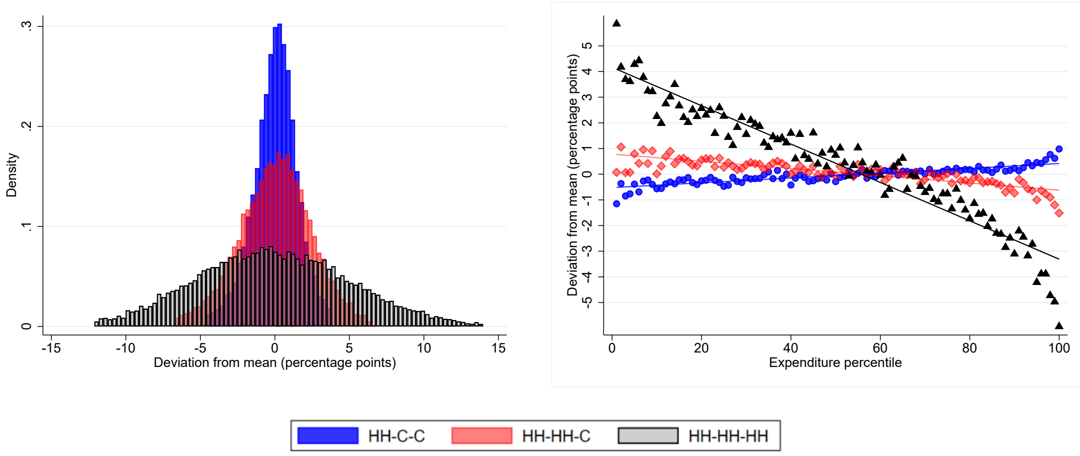

Notes: Authors' calculations using Kantar's Take Home Purchase Panel (2021–2023). Panel (a) shows a histogram of the distribution of cumulative inflation from 2021Q3 to 2023Q3. Panel (b) plots the relationship between cumulative inflation and the expenditure distribution percentile to which a household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

In Appendix C.1, we show that replacing the common prices in the Laspeyres index with prices that vary across households has minimal impact on the inflation inequality gradient. This aligns with our evidence that poorer households pay only modestly lower prices for identical products, and this pattern remains relatively stable over time (see Appendix A.1). We also show that when we instead use banded household income as our measure of household well-being, we continue to find that the most recent nine-quarter period exhibits a pronounced gradient in inflation inequality that is absent in earlier periods.

4.2 Substitution responses

To what extent do household behavioural adjustments reduce cost-of-living increases? To answer this, we estimate a cost-of-living index (equation (3.1)) that accounts for non-homothetic preferences and compare it to the Laspeyres inflation exposure index. With non-homothetic preferences, the cost-of-living index depends on the utility level, indexing living standards, at which it is evaluated. We construct an index that holds living standards fixed at their pre-inflation surge level (2021Q3). The index therefore measures the change in expenditure required to maintain 2021Q3 living standards. Our estimation follows a version of the algorithm outlined in Section 3 that adjusts a Törnqvist index for non-homothetic preferences (see Appendix B.2 for further details). The resulting cost-of-living index differs from the Laspeyres index, by accounting for household substitution responses to relative price changes, and it differs form a homothetic cost-of-living index approximation (for example, equation (3.6)) by adjusting for income effects that arises from changes in living standards over time.

Figure 4.3(a) illustrates how the difference between our estimate of the cost-of-living index and the Laspeyres index varies across expenditure percentiles. On average, this difference is 1.5 percentage points, implying that household substitution responses reduced the average cost-of-living increase from 26.3% (under no substitution) to 24.7%—a 5.9% reduction. There is no systematic relationship between the difference in inflation exposure and cost-of-living increase across the expenditure percentiles. However, since worse-off households face greater inflation exposure, the percentage reduction in their cost-of-living rise due to substitution responses is smaller compared to better-off households. For instance, among households in the bottom expenditure quartile, substitution responses reduced the cost-of-living increase by 5.5%, compared with 6.9% for those in the top quartile.

In Figure 4.3(b), we decompose substitution effects by reporting the contribution of each product segment. We do this by constructing a cost-of-living subindex for each segment. For each expenditure quartile, we report both average inflation exposure and the extent to which substitution responses mitigate cost-of-living increases, measured by the difference in the segment cost-of-living and Laspeyres subindexes. Additionally, we report a between-segment component. By construction, this is zero for inflation exposure (see equation (3.4)). For substitution effects, it equals the difference between our cost-of-living index and the initial share-weighted average of the cost-of-living subindexes. This decomposition not only provides a purely statistical within-between breakdown of substitution effects, but also offers a theory-consistent decomposition under the additional preference restriction of weak separability across product segments.

Several notable patterns emerge. First, households in lower expenditure quartiles exhibit higher inflation exposure across all product segments, underscoring that higher inflation exposure for worse-off households is a widespread phenomenon. Second, across all expenditure quartiles and product segments, substitution effects materially reduce cost-of-living increases, typically by more than one percentage point. Third, in some segments—such as dairy, prepared food, and confectionery—substitution responses mitigate cost-of-living increases substantially more for bottom-quartile households than for top-quartile households, with differences exceeding 0.5 percentage points. However, when measured across all fast-moving consumer goods, the overall size of substitution effects differs little by expenditure

Figure 4.3: *Substitution effects*

(a) In total

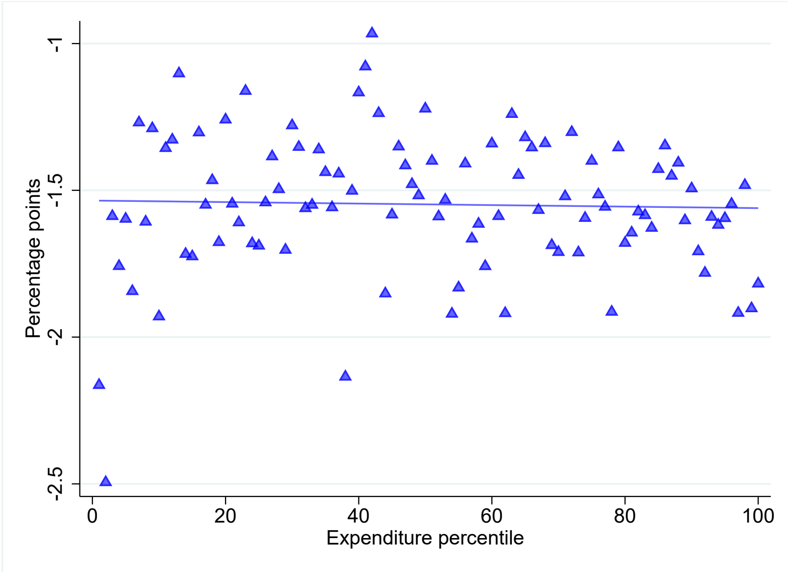

(b) Product segment decomposition

	Expenditure quartile							
	Bottom		2		3		Top	
	Exp. (%)	Sub. (p.p.)	Exp. (%)	Sub. (p.p.)	Exp. (%)	Sub. (p.p.)	Exp. (%)	Sub. (p.p.)
Bakery	32.97	-1.34	31.97	-1.28	31.02	-1.18	29.59	-1.04
Dairy	35.95	-1.14	34.26	-0.62	33.33	-0.58	31.92	-0.53
Fresh fruit and veg	20.64	-1.96	19.31	-2.05	18.07	-2.12	15.73	-2.07
Meat and fish	28.95	-1.28	28.06	-1.17	27.12	-1.40	25.45	-1.38
Prepared	32.54	-1.84	30.42	-1.44	29.10	-1.45	26.52	-1.11
Cupboard ingredients	33.76	-1.94	32.27	-1.63	31.00	-1.70	29.16	-1.65
Confectionery	32.37	-1.39	31.12	-0.94	30.27	-0.92	28.60	-0.85
Drinks	26.55	-1.50	25.62	-1.66	24.49	-1.78	22.61	-1.65
Alcohol	9.58	-1.12	8.84	-1.02	8.19	-1.07	7.61	-0.86
Household goods	23.81	-1.42	23.12	-1.23	22.25	-1.36	21.41	-1.68
Between	0	-0.05	0	-0.13	0	-0.18	0	-0.30
Total	27.19	-1.59	25.74	-1.45	24.19	-1.55	21.76	-1.61

Notes: Authors' calculations using Kantar's Take Home Purchase Panel (2021-2023). Panel (a) plots our estimates for the impact of substitution effects on cost-of-living changes across expenditure percentiles. Panel (b) reports how inflation exposure and substitution effects vary across product segments and expenditure quartiles. All numbers are cumulative effects over 2021Q3-2023Q3.

quartile. This is because better-off households allocate a lower share of spending to categories where substitution effects are weaker (for instance, dairy), combined with their larger between-segment substitution effect. Finally, substitution effects within segments are much stronger than between-segment responses, highlighting the importance of detailed spending data for accurately measuring behavioural adjustments to inflation.

Another limitation of the Laspeyres inflation exposure index (equation (3.3)), aside from not accounting for substitution responses, is that it fails to account for the effect of product entry and exit on the cost of living. In Appendix C.2 we show that there is no evidence that product entry and exit contributed to inequality in the cost of living over the nine-quarter inflation surge. Specifically, we demonstrate that chaining the Laspeyres index has no distinguishable effect on inequality in inflation, despite allowing the consumption baskets to update over time.[15] We also compute the price index developed by Feenstra (1994) and Broda and Weinstein (2010), which, under the assumption of CES preferences, incorporates into the cost-of-living index the value of product entry and exist, showing that doing so has no discernible impact on the pattern of inflation inequality.

4.3 Trading down

The primary adjustment margin through which households can limit increases in their cost of living is by substituting away from products with rising relative prices toward those with falling relative prices. In a high-inflation environment, where price growth is stronger at the lower end of the quality ladder, substitution effects push households toward purchasing *higher* quality goods. This is evident in Figure 2.1(d), where nearly all households increase their spending share on top-quality-rung products.

At the same time, high inflation contributed to declining living standards. If households have non-homothetic preferences, this leads to a shift in consumption from luxuries to necessities. Figure 2.1(c) and (d) provide both cross-sectional and *within*-household evidence that low-quality products are necessities, while high-quality products are luxuries. Households experiencing the largest declines in deflated spending most strongly switch to bottom-rung products, and least strongly

[15] The direct 2021Q3 to 2023Q3 comparison in equation (3.3) implicitly conditions on products available in both periods, representing 86% of expenditure in our data. A chained 2021Q3 to 2023Q3 comparison is less affected by product churn since it is built on a serious of quarter-to-quarter comparison, therefore incorporating the larger set of products available any two consecutive periods.

to top-rung products. While these income-effect-driven reallocations do not directly enter the cost-of-living index—and instead present a challenge in separating substitution and income effects—they are relevant for changes in the quality of living.

Figure 4.4 provides evidence on the extent and consequences of trading down the quality ladder. We group households into 100 bins using the interaction of expenditure deciles (as defined above, based on 2021 equivalised spending) and deciles of living standard changes over 2021Q1-2023Q3. We measure living standard changes using within-household changes in the log of our quality-of-living index. Households in the bottom seven deciles all experience declines in living standards.

Panel (a) illustrates how adjustments along the product quality ladder vary across household groups. For each household-product category, we compute the change in the average quality rung of consumption baskets between 2021Q3 and 2023Q3, then take a spending-share-weighted average across categories. We report how this summary measure of quality ladder adjustments varies across household groups.

While these adjustments reflect a combination of substitution and income effects, comparing households with similar 2021 equivalised expenditure but different changes in living standards allows us to focus on variation driven by income effects. The graph shows that households experiencing the largest declines in living standards exhibit the most pronounced movements down the quality ladder, consistent with an income effect leading them to trade down. Additionally, among those with the steepest declines, higher initial-spending households move most strongly down the quality ladder. This aligns with better-off households—who tend to initially purchase higher quality products—having greater scope for trading down.

In panel (b), we show how the non-homotheticity bias in the homothetic cost-of-living approximation in equation (3.6) varies across household groups. Specifically, we report the difference between our non-homothetic cost-of-living index and the homothetic approximation. For households experiencing a decline in living standards, the homothetic approximation overstates the increase in the cost of maintaining pre-inflation surge living standards. This bias is largest—up to 1 percentage point—for households in the top expenditure decile that experience the largest declines in living standards. The bias arises because income effects from reductions in living standards induce households to reallocate spending toward lower-quality products, which exhibit the strongest price growth—an effect that homothetic indexes fail to account for.

A corollary is that, for households experiencing declines in living standards, the homothetic approximation overstates the final-period-price-denominated magnitude of this reduction. This occurs because, in the initial period—when these households were better off—luxuries were relatively more expensive than at final-period prices. As a result, their initial living standards were lower than inferred by a homothetic index.

Conditional on the size of the living standard reduction, less well-off households—those in lower expenditure deciles—exhibit smaller bias. This is consistent with these households having less scope to trade down in response to living standard falls. For the 30% of households that experience rising living standards during this period, the cost-of-living bias from neglecting non-homothetic preferences operates in the opposite direction.

In panel (c), we illustrate how trading down the quality ladder—despite entailing spending reallocation toward goods with rising relative prices—acts as an adjustment mechanism for households during periods of high inflation and declining living standard. Specifically, for each household-product category, we compute the quantity-weighted average price *level* paid in 2023Q3 and compare it with the hypothetical price level households would have paid in 2023Q3 had there consumption baskets remained as they were in 2021Q3. Similar to panel (a), this index of prices paid reflects both substitution and income effects. However, by leveraging our within-household variation, we can compare households with different sizes of declines in living standards, isolating the variation driven by the income effect associated with falling purchasing power.

The figure shows that the consumption basket adjustments made by households with the largest declines in living standards had the strongest effect in reducing the average within-category price level of their purchases. For households in the bottom 2021 expenditure decile, the fall in the average price paid ranges from 1.9% for those in the top living standard change decile to 4.3% for those experiencing the largest decline in living standards. In comparison, for households in the top initial expenditure decile, the change in price paid ranges from 0.3% and -4.7%. While across all initial spending deciles, larger reductions in living standards are associated with the largest falls in price paid, the difference across the living cost change distribution is most pronounced for the initially best-off households.

In Appendix C.3, we show that the trading-down patterns summarised in Figure 4.4 hold for alternative measures of household well-being and changes in living standards. Specifically, the main findings remain when we measure household well-being using equivalised current income, or average equivalised expenditure from 2021 to 2023. Similarly, measuring changes in living standards based on differences in annual equivalised expenditure between 2021 and 2023 does not alter the key conclusions.

Figure 4.4: *Trading-down effects*

(a) Movements along the quality ladder

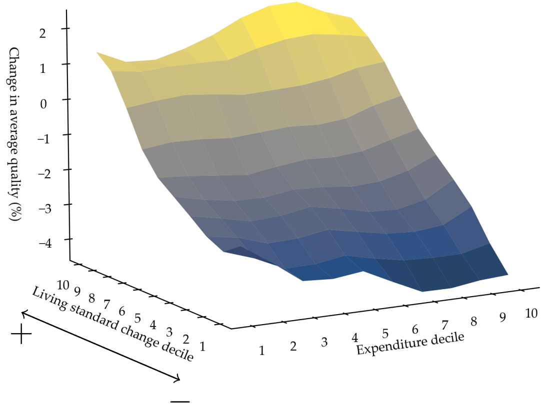

(b) Non-homotheticity cost-of-living bias

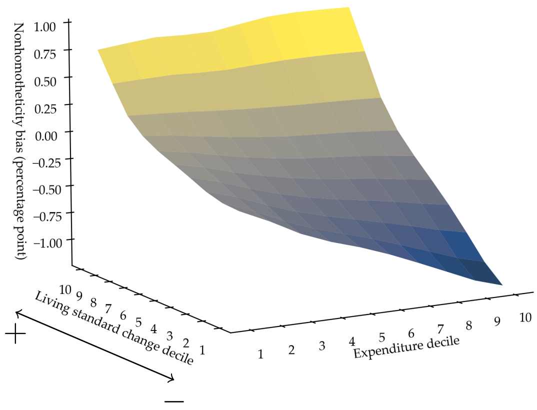

(c) Change in price level of consumption baskets

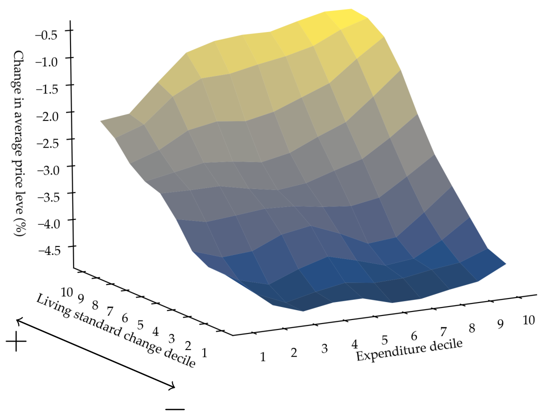

Notes: Authors' calculations using Kantar's Take Home Purchase Panel (2021-2023). For each product category, we compute the percent change in the quality rung of purchases between 2021Q3 to 2023Q3; panel (a) reports the spending-weighted average across categories. Panel (b) shows the difference between our cost-of-living index and a homothetic Törnqvist approximation. For each product category, we compute the percent change in the quantity-weighted average 2023Q3 price, comparing 2021Q1 and 2023Q3 weights; panel (c) reports the spending-weighted average across categories.

5 Conclusion

Inflationary surges lead to shifts in relative prices and, if nominal incomes do not keep pace with rising prices, declining purchasing power and living standards. Differences in consumption baskets lead to variation in exposure to price increases, and the simultaneous shifts in relative prices and living standards drive both substitution and income-effect responses.

We show that the 2021-2023 inflationary surge entailed more rapid price growth for low-quality products most popular with poorer households, leaving them disproportionately exposed to cost-of-living increases. Substitution responses mitigated cost-of-living increases, but households nonetheless experienced an increase in their cost of living that is unprecedented in recent decades and highly skewed toward poorer households. Falling living standards drove many households to trade down to lower-quality necessities, despite their more rapid price growth. This income-effect-driven adjustment enables households to lower the price level of their consumption basket, helping to limit declines in living standards. However, it also leaves them more highly exposed to any future necessity-driven inflationary surge.

Our results underscore the disproportionate burden that high inflation places on less well-off households. While the origins of the 2021-2023 inflationary surge are still debated, it is generally accepted that policy decisions are a key determinant of inflation. Recently, sixteen Nobel Laureate economists (2024) warned against the inflationary effects of the current US policy trajectory. It is important for policymakers to carefully consider the implications that surging prices have for inequality, and to design policies that mitigate the regressive impacts of inflation.

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.

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.

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*.

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.

Diewert, W. E. (2022). Scanner Data, Elementary Price Indexes and the Chain Drift Problem. In D. Chotikapanich, A. N. Rambaldi, and N. Rohde (Eds.), *Advances in Economic Measurement*, pp. 445–606. Singapore: Springer Nature Singapore.

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.

Federal Reserve (2025). Minutes of the Federal Open Market Committee, January 28-29, 2025. Technical Report Board of Governonr of the Federal Researve System.

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*.

Fox, K. J., P. Levell, and M. O'Connell (2023). Inflation measurement with high frequency data. *CEPR Discussion Paper 18539*.

Gillespie, L. (2024, October). Survey: 41% Of Americans Say Inflation Is Their Top Economic Issue For The 2024 Election. https://www.bankrate.com/personal-finance/2024-election-and-the-economy-survey/.

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.

Ivancic, L., W. Erwin Diewert, and K. J. Fox (2011). Scanner data, time aggregation and the construction of price indexes. *Journal of Econometrics 161*(1), 24–35.

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*.

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.

Michael, R. T. (1979). Variation Across Households in the Rate of Inflation. *Journal of Money, Credit and Banking 11*(1), 32–46.

Nobel Laureates (2024, February). Letter from 16 Nobel Laureates on Inflation and Monetary Policy. https://www.documentcloud.org/documents/24777566-nobel-letter-final/.

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*.

Slesnick, D. T. (1993). Gaining Ground: Poverty in the Postwar United States. *Journal of Political Economy 101*(1), 1–38.

Appendix:
For online publication

Measuring cost of living inequality during an inflation surge

Tao Chen, Peter Levell and Martin O'Connell

May, 2025

A Data appendix

A.1 Price dispersion for identical products

Our baseline inflation measure uses prices that are common across households. In the context of the UK grocery industry, where retailers have national store coverage and practice national pricing policies (see Competition Commission, 2008), the assumption of common prices is arguably more innocuous than in other settings. Furthermore, the two most prominent low-cost retailers, Aldi and Lidl, primary sell private-label products unique to them, meaning our cost-of-living measures incorporate substitution toward discounter products.

Nonetheless, there is likely to be some price dispersion for identical products, due, for instance, to discounts that are in place for only a limited number of weeks and, in the case of national brands, only in some retailers. If households of different incomes alter their propensity to take advantage of low prices over time this could impact inflation inequality. To assess this we use a price index suggested in Aguiar and Hurst (2007), which measures dispersion in price paid for a *fixed* basket of products.

Let q_{hit} denote the volume of product i purchased by household h in year-quarter t and define the household-specific price by $p_{hit} = x_{hit}/q_{hit}$.[16] Had the household paid average prices for their basket their 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:

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

In Figure A.1 we summarise the evolution of the AH index over time. In panels (a) and (b) we summarise how the AH index varies over 2012-2023, across quartiles of the equivalised expenditure distribution. To do this, in each calendar year (2012-2023), we group households into quartiles of the expenditure distribution and we report the average AH index across households in each expenditure quartile (expressed as a deviation from the quarter mean across all households) in each year-quarter. Panel (a) shows the index when products are defined based on brand-pack size. Panel (b) show the index when we re-define products based on the combination of brand-pack size and retailer (which is relevant for branded products sold by multiple retailers, but leaves the definition of private-label products unaf-

[16] Note that the common price and household-specific prices are related by: $p_{it} = \frac{\sum_h x_{hit}}{\sum_h q_{hit}} = \sum_h \frac{q_{hit}}{\sum_{h'} q_{h'it}} p_{hit}$.

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

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

By expenditure quartile

(a) Across products

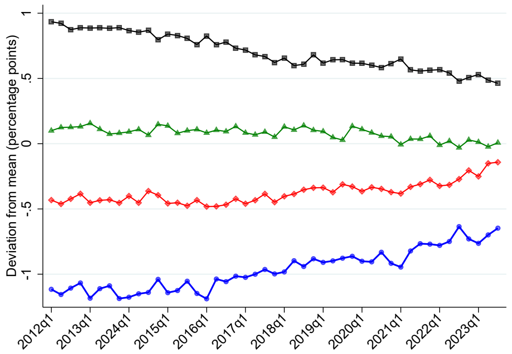

(b) Across product-retailers

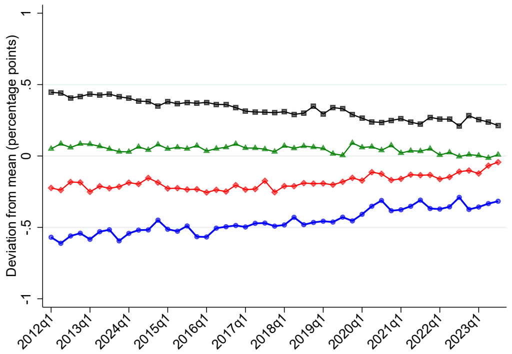

By income quartile

(c) Across products

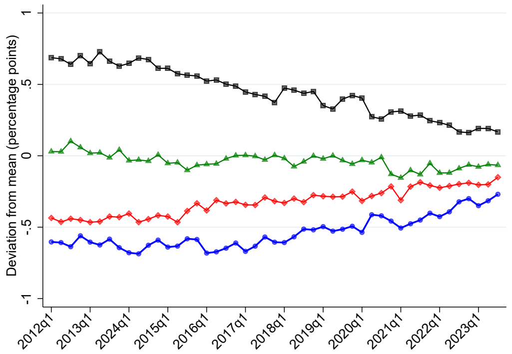

(d) Across product-retailers

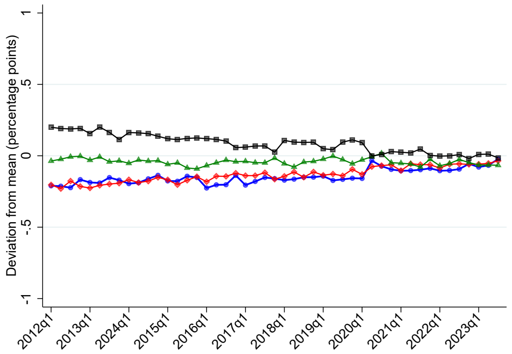

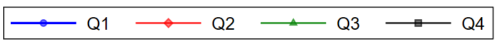

Notes: Authors' calculations using Kantar's Take Home Purchase Panel (2012–2023). Panels (a) and (b) show within equivalised expenditure quartile average (across households) AH index at quarterly frequency over 2012Q1-2023Q3. In panel (a) products are based on brand-size and in panel (b) they are based on brand-size and retailer. 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. Around 1 percentage point was due to cross-retailer variation and 1 percentage point due to within-retailer variation. This gap has closed over time; by 2023 those in the top expenditure quartile paid only around 1.1 percentage point more than those in the bottom quartile (split approximately evenly between the influence of across and within retailer dispersion). This decline, implies differences in price paid for identical goods act to reinforce inflation inequality. Over the nine-quarter period 2021Q3 to 2023Q3 the decline in the price gap was 0.48 percentage points based on the brand-pack size product definition, with virtually no change based on the brand-pack size-retailer definition.

In panels (c) and (d) we repeat the results from panels (a) and (b), instead splitting households based on quartiles of the current equivalised income distribution. The results are similar.

A.2 Income measures

In Figure A.2 we report a local-polynomial smoothed regression of 2021 equivalised expenditure and income.

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

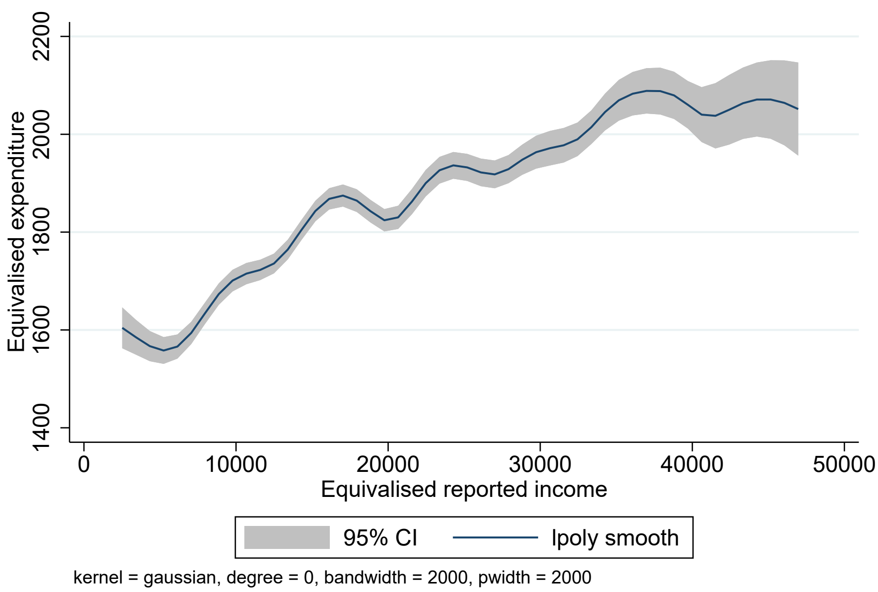

Notes: Authors' calculations using Kantar's Take Home Purchase Panel (2012–2023). Graph shows the fitted line from a local-polynomial smoothed regression of annual equivalised household expenditure and reported income in 2021. Income data is reported in £10,000 bands. We use the band midpoint.

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 Kantar data. When the prices of cheapest products grow the fastest—a phenomenon known as "cheapflation"—we expect average price dispersion to fall, causing these distributions to shift leftwards. The figure shows there was indeed a leftward shift in 2023 compared to 2022 and 2021.

We can 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 shows the distributions of interquartile ranges for log prices within 'item' categories of the UK CPI for the month of December in years 2021, 2022 and 2023. Items represent disaggregated product groups (e.g., "large white loaf unsliced, 800g") which are more disaggregated than product categories we use in Figure A.3.

We plot these for FMCG—constructed by combining food consumed at home, alcoholic and soft drinks consumed at home, and pharmaceutical and personal care products—as well as for other goods categories: clothing, household appliances (including tools, household articles and electronics), and leisure goods (which includes toys, sports equipment and books). As with the Kantar 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. However, we do not see as clear a pattern of falling price dispersion for household appliances.

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

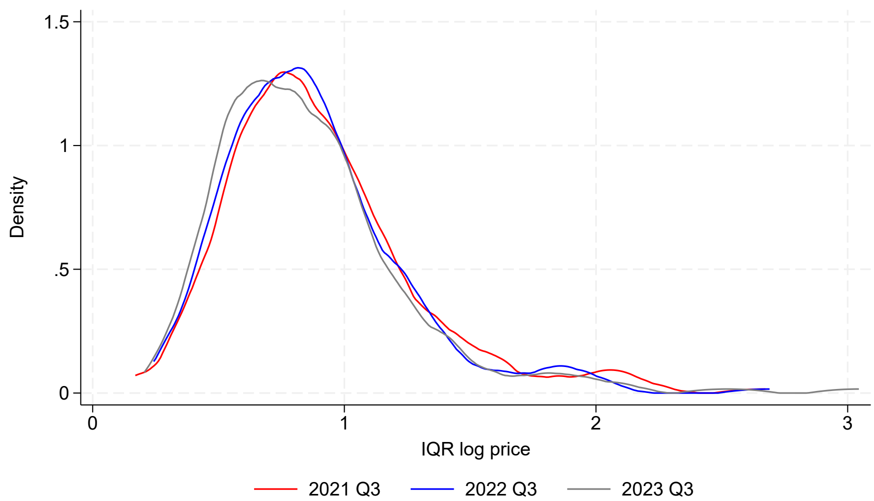

Notes: Authors' calculations using Kantar's Take Home Purchase Panel (2012–2023). Graph shows the distribution of interquartile ranges of log prices computed over products within product categories in 2021Q3, 2022Q2 and 2023Q3.

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

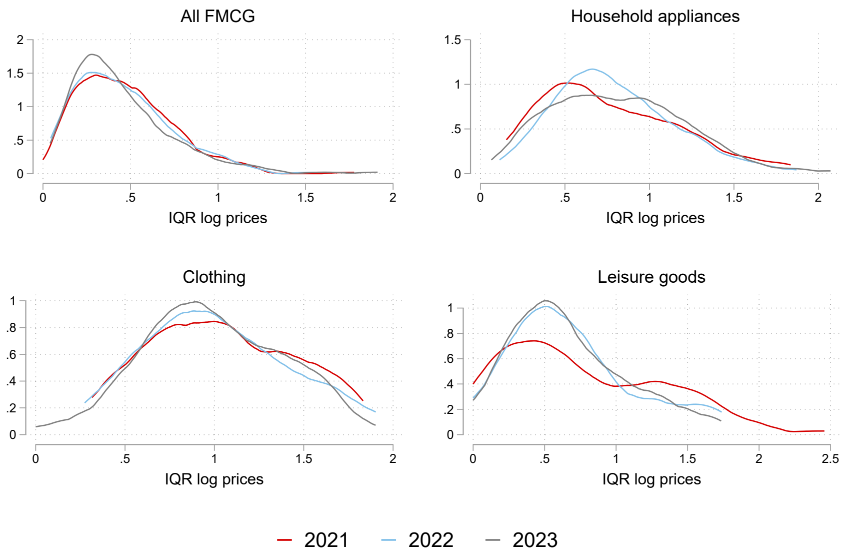

Notes: Authors' calculations using CPI microdata produced by the Office for National Statistics. Graph shows the distribution of interquartile ranges of log prices computed over sampled price within CPI 'items' in December 2021, 2022 and 2023. FMCG are 'fast moving consumer goods'.

B Measurement appendix

B.1 Equivalent variation measures

Figure 3.1 illustrates construction of the cost-of-living index, evaluated at initial period's living standards and the quality-of-living index, evaluated at final-period prices. In Figure B.1 we illustrate construction of the cost-of-living index, evaluated at final period's living standards (left-hand graph) and the quality-of-living index, evaluated at initial-period prices (right-hand graph). Both indexes depend on the unobserved level of spending required at initial-period prices to reach final-period living standards.

Like in Figure 3.1, the budget constraint in period 1 is AB, and the optimal choice is at point X. In period T, the budget constraint shifts to CD, with the new optimal choice at Y.

The cost-of-living index compares final-period spending, $x_T = e(p_T, u_T)$, with the expenditure required to maintain the same living standards, u_T, at initial-period prices, $e(p_1, u_T)$, corresponding to the consumption bundle Z on the budget constraint $C'D'$. If the household did not adjust their consumption bundle in response to changes in relative prices, their expenditure would be given by $x^{NS} < e(p_1, u_T)$.

The quality-of-living index compares initial-period spending, $x_1 = e(p_1, u_1)$, with the expenditure necessary, at the same prices, to attain final-period living standards, $e(p_1, u_T)$. If the household maintained the same *spending share* on each good as in the initial period–represented by the ray from the origin to the initial period optimal choice–then the expenditure required at initial-period prices to attain final-period living standards would be given by $x^H < e(p_1, u_T)$.

In this case, the quality-of-living index is analogous to the household's equivalent variation for the change in their budget set, expressed in proportional terms rather than as a difference. The cost-of-living index, in turn, serves as the proportional analog to equivalent variation from the price change alone, holding nominal expenditure at its final-period value.

The unobserved expenditure function in this context is $e(p_1, u_T)$, which appears in the denominator of the cost-of-living index. As a result, the no-substitution consumption bundle (corresponding to a Paasche index) understates cost-of-living increases. Similarly, since $e(p_1, u_T)$ appears in the numerator of the quality-of-living index, the homothetic bundle overestimates the quality-of-living index.

Figure B.1: *Cost- and quality-of-living indexes at period T utility and 1 prices*

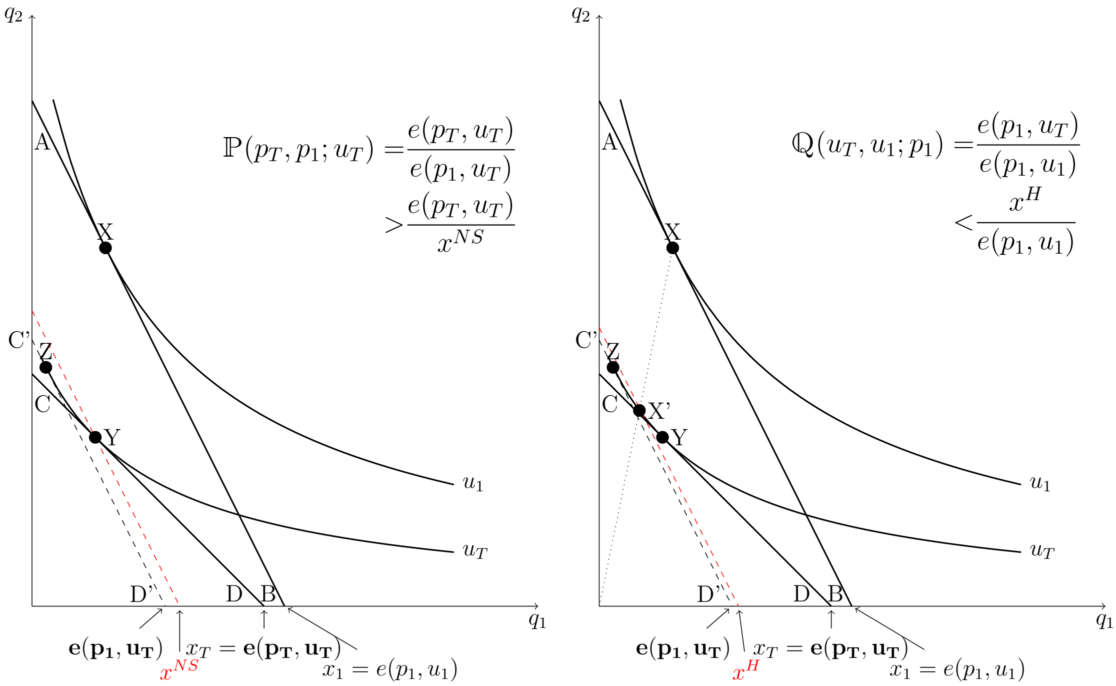

B.2 Second-order Jaravel-Lashkari algorithm

In Section 3 would outline the first-order algorithm suggested in Jaravel and Lashkari (2024). Here we outline the second-order algorithm. This entails adjusting a superlative price index for non-homotheticities.

Utility cardinalisation. Denote the household's indirect utility and expenditure function by $\mathbb{u} = \mathbb{v}(\mathbf{p}, x)$ and $x = \mathbb{e}(\mathbf{p}, \mathbb{u})$ for an arbitrary cardinalisation of utility. Define the period 1 denominated money-metric utility function $u = v(\mathbf{p}, x) = \mathbb{e}(\mathbf{p}_1, \mathbb{v}(\mathbf{p}, x))$, and the corresponding expenditure function is $x = e(\mathbf{p}, u) = \mathbb{e}(\mathbf{p}, \mathbb{v}(\mathbf{p}_1, u))$.

Algorithm. The household's log expenditure growth between period t and $t-1$ can be written:

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

where the second line follows from a first-order approximation.

Note that:

$$\begin{aligned}\log \mathbb{P}(\mathbf{p}_1, \mathbf{p}_t, u_t) =& \log e(\mathbf{p}_t, u_t) - \log e(\mathbf{p}_1, u_t) \implies \\ \frac{\partial \log \mathbb{P}(\mathbf{p}_1, \mathbf{p}_t, u_t)}{\partial \log u} =& \frac{\partial \log e(\mathbf{p}_t, u_t)}{\partial \log u} - 1\end{aligned}$$

where the last line uses $e(\mathbf{p}_1, u_t) = u_t$. Define $\Lambda_t(u_t) = \frac{\partial \log \mathbb{P}(\mathbf{p}_1, \mathbf{p}_t, u_t)}{\partial \log u} = \frac{\partial \log e(\mathbf{p}_t, u_t)}{\partial \log u} - 1$.

The household's log expenditure growth between period t and $t-1$ can alternative by written:

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

where the second line follows from a first-order approximation.

Taking an arithmetic average of the two first-order approximation:

$$\begin{aligned}\log x_t - \log x_{t-1} =& \frac{1}{2}\sum_i (s_{it-1} + s_{it})(\log p_{it} - \log p_{it-1}) \\ &+ \left(1 + \frac{1}{2}(\Lambda_{t-1}(u_{t-1}) + \Lambda_t(u_t))\right)(\log u_t - \log u_{t-1})\end{aligned}$$

Note that the first term is a log-Törnqvist index: $\log \Pi^{\mathcal{T}}_{(t,t-1)} = \frac{1}{2}\sum_i (s_{it-1} + s_{it})(\log p_{it} - \log p_{it-1})$.

This implies the recursive sequence:

$$
\begin{aligned}
\log u_1 &= \log x_1 \\
\log u_2 &= \log u_1 + \frac{1}{1 + \frac{1}{2}\left(\Lambda_{t-1}(u_{t-1}) + \Lambda_t(u_t)\right)}\left(\log x_2 - \log x_1 - \log \Pi^{\mathcal{T}}_{(1,2)}\right) \\
&\dots \\
\log u_T &= \log u_{T-1} + \frac{1}{1 + \frac{1}{2}\left(\Lambda_{T-1}(u_{T-1}) + \Lambda_T(u_T)\right)}\left(\log x_T - \log x_{T-1} - \log \Pi^{\mathcal{T}}_{(T,T-1)}\right)
\end{aligned}
$$

At each stage $\Lambda_{t-1}(u_{t-1})$ can be constructed based on cross-section (across household) regression: $\log \Pi^{\mathcal{T}}_{(t-1,t)} = g_t(\log u_{t-1}) + \epsilon$, with: $\Lambda_{t-1}(u_{t-1}) = \sum_{\tau \leq t-1} g'_\tau(\log u_{1-t})$. The algorithm adjusts a superlative Törnqvist index for non-homotheticities.

The log cost-of-living index evaluated at final period utility is then

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

Reverse algorithm. The algorithm as described constructs a cost-of-living index evaluated at final-period realised utility. The cost-of-living index evaluated at initial period utility can be constructed by denominating utility in final-period prices and running the algorithm in reverse (stating at period T).

C Results appendix

C.1 Inflation exposure

In Appendix A.1 we show there is little change in the price gap paid for identical products between poor and better-off households over time. This suggests changes in price paid for identical goods play little role in driving inflation inequality over 2021Q3 to 2023Q3. Figure C.1 confirms that this is the case, by showing the relationship between household-level cumulative inflation over 2021Q3-2023Q3 based on common prices—the blue line and markers; repeating information in Figure 4.1(c)—and based on prices computed at the expenditure quartile level[17]—the red line and markers. It shows that the inflation gradient is, to a modest degree, larger

[17] 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.

in the latter case, and therefore that variation in price paid slightly reinforces, rather than unwinds, our finding of significant recent inflation inequality.

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

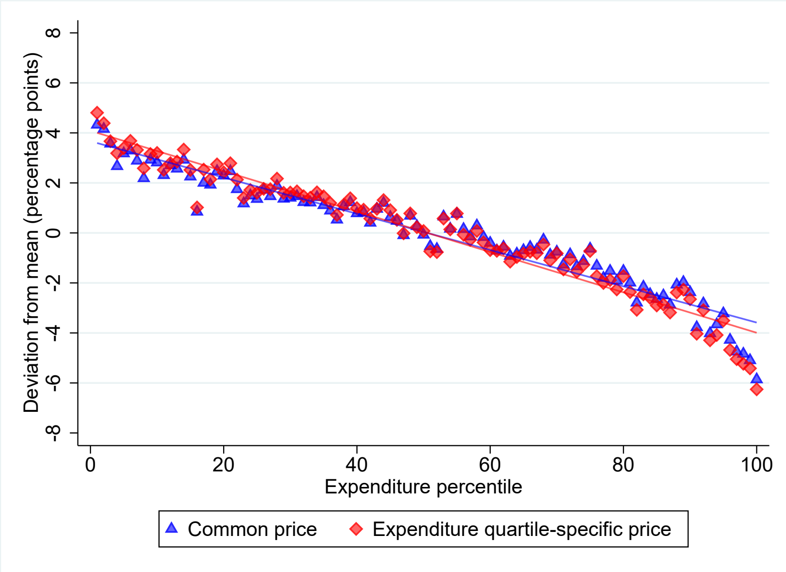

Notes: Authors' calculations using Kantar's Take Home Purchase Panel (2021–2023). Figure illustrates the relationship between ninth-quarter cumulative inflation and a household's percentile in the expenditure distribution, with a marker for each percentile and a line of best fit. Households are assigned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

Figure C.2: *Inflation inequality: across income deciles*

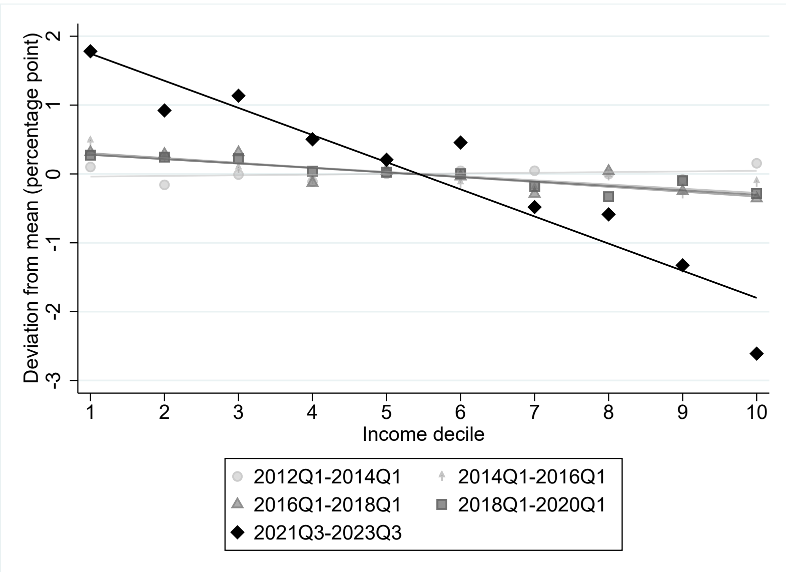

Notes: Authors' calculations using Kantar's Take Home panel (2012-2023). Figure plots the relationship between ninth-quarter cumulative inflation and income deciles based on a household's equivalised income computed based on the band midpoint of banded household income in the initial calendar year of the relevant nine quarter period. Cumulative inflation is computed using a Laspeyres index.

In Figure C.2 we show the relationship between household-level cumulative inflation and what deciles of the equivalised household income distribution a household belongs to. Consistent with Figure 4.1(c), it shows a high degree of inflation inequality over the nine-quarters beginning in 2021Q3 that is absent in any of the earlier nine-quarter periods.

C.2 Product churn

The Laspeyres index in equation (3.3) holds the consumption basket fixed at the household's initial basket, meaning it does not reflect the effects of product entry and exit. In Figure C.3 we show inflation inequality computed using two variants of the Laspeyres index. Panel (a) shows an index chained year-on-year (comparing Q3 in 2021, 2022 and 2023), 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. Nevertheless, the figures shows that they nevertheless exhibit the same pattern of inflation inequality as in Figure 4.1(c).

Figure C.3: *Inflation inequality: chained Laspeyres index*

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

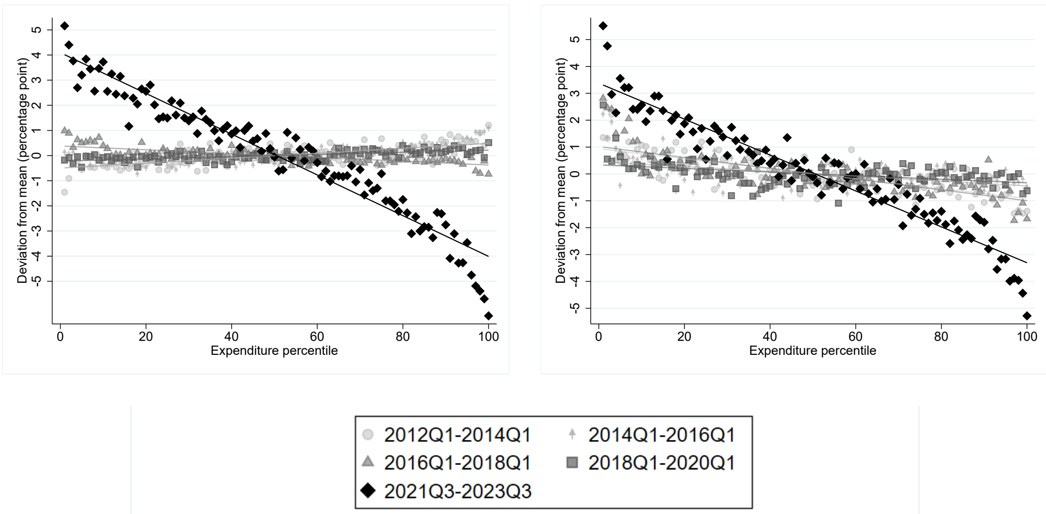

Notes: Authors' calculations Kantar's Take Home Purchase Panel (2012-2023). Figure illustrates the relationship between ninth-quarter cumulative inflation and a household's percentile in the expenditure distribution, with a marker for each percentile and a line of best fit. Households are assigned to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

However, chained indexes do not fully address the issue, as even across successive quarters, there can be a considerable degree 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 on products not available in 2022Q3 and 8.6% of 2022Q3 aggregate spending is on products that were not available in 2021Q3.[18] 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 $\mathbb{P}^{CES}_{h,(s,t)} = \left[\prod_i \left(\frac{p_{it}}{p_{is}}\right)^{\phi_{hi,(s,t)}}\right] \times \mathcal{F}_{h,(s,t)}$, where $\phi_{hi,(s,t)}$ is a weight that depends on the household's period s and t spending shares over products available in both comparison periods.[19] $\mathcal{F}_{h,(s,t)} = \left(\frac{1-s^N_{ht}}{1-s^X_{hs}}\right)^{\frac{1}{\sigma-1}}$ is the correction for the influence of entering and exiting products. It depends positively on the share of initial period spending the household allocates to exiting goods, s^X_{hs} (the cost-of-living grows if products the household consumes disappear) and negatively on the share of final period spending they allocate to entering goods, s^N_{ht} (as the consumption of new products contributes towards a lower 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, exiting and entering products do not have close substitutes and therefore net entry has a relatively large effect of the cost-of-living.

In Figure C.4(a) we plot the inflation inequality gradient computed using a CES price index with no Feentra adjustment. It shows that the inflation inequality pattern is very similar to that shown in Figure 4.1(c). In panel (b) we plot the difference in entry and exit shares, $s^N_{ht} - s^X_{hs}$. In panel (c) we show 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 allocated 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.

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

[19] 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.4: *CES cost-of-living index*

(a) Not adjusted for net entry

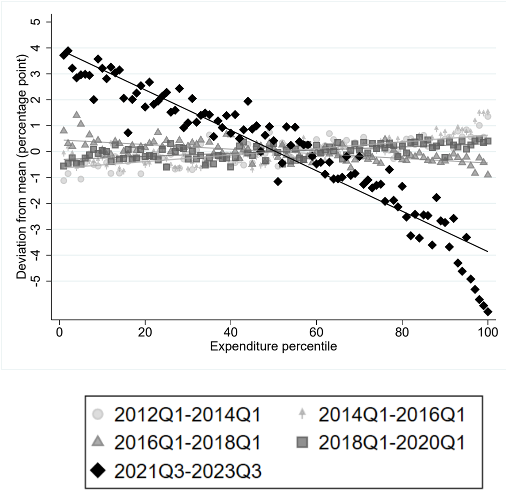

(b) Average net entry

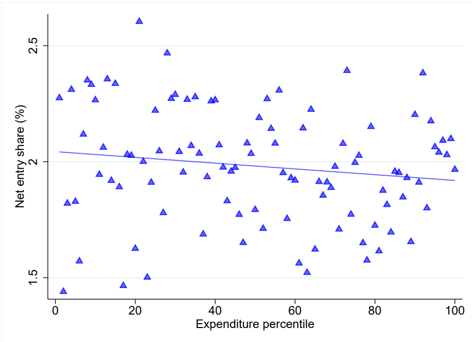

(c) Net entry adjustment

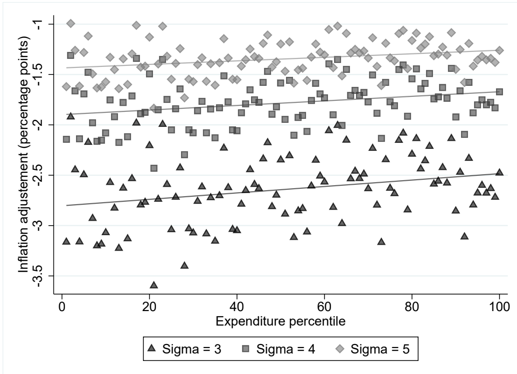

Notes: Authors' calculations Kantar's Take Home Purchase Panel (2021-2023). Panel (a) the relationship between nine-quarter cumulative inflation based on a CES price index (with no Feenstra-correction) and expenditure percentiles. Panel (b) plots the average difference in entry and exit shares $s^N_{ht}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

C.3 Trading-down effects

In Figures C.5-C.7, we reproduce the trading down results (Figure 4.4) under different measures of well-being and living standard changes. Figure C.5 measures well-being based on deciles of the equivalised income distribution. Figure C.6 instead uses average equivalised expenditure across 2021-2023. Figure C.7 uses the 2021 expenditure deciles (in common with Figure 4.4), but measures living standard changes based on deciles of the log change in equivalised annual expenditure between 2021 and 2023. All graph share the same qualitative main takeaways.

Figure C.5: *By income*

(a) Movements along the quality ladder

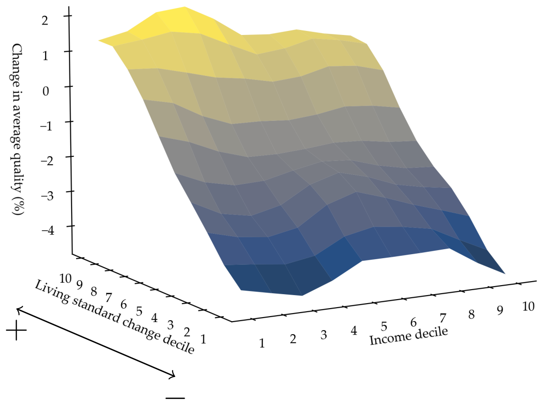

(b) Non-homotheticity cost-of-living bias

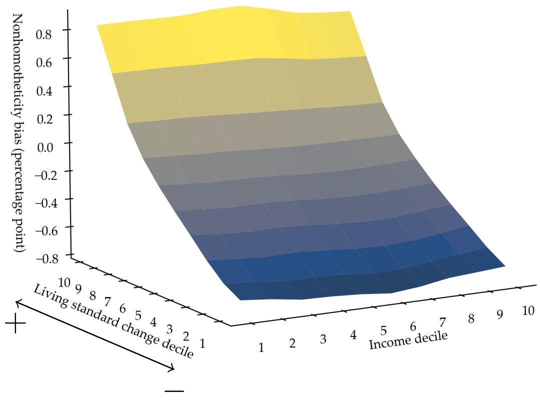

(c) Change in price level of consumption baskets

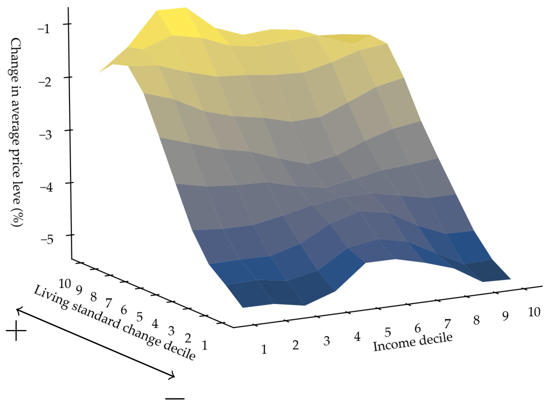

Notes: Authors' calculations using Kantar's Take Home Purchase Panel (2021-2023). For each product category, we compute the percent change in the quality rung of purchases between 2021Q1 to 2023Q3; panel (a) reports the spending-weighted average across categories. Panel (b) shows the difference between our cost-of-living index and a homothetic Törnqvist approximation. For each product category, we compute the percent change in the quantity-weighted average 2023Q3 price, comparing 2021Q1 and 2023Q3 weights; panel (c) reports the spending-weighted average across categories.

Figure C.6: *By three year average annual expenditure*

(a) Movements along the quality ladder

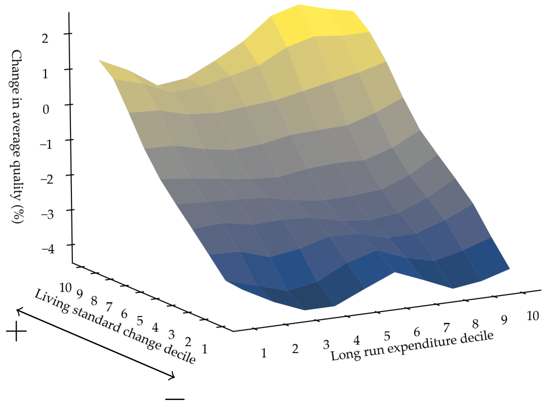

(b) Non-homotheticity cost-of-living bias

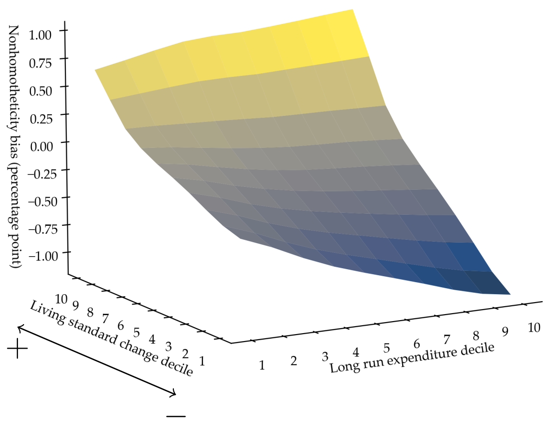

(c) Change in price level of consumption baskets

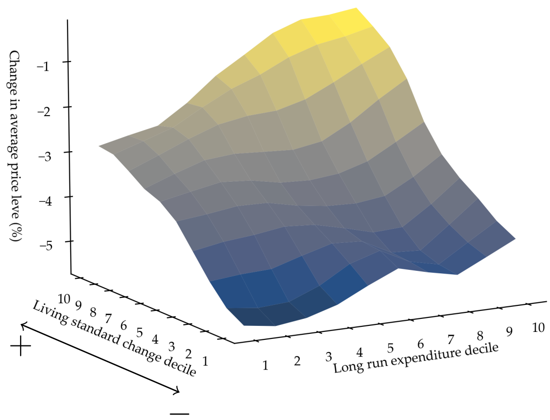

Notes: Authors' calculations using Kantar's Take Home Purchase Panel (2021-2023). For each product category, we compute the percent change in the quality rung of purchases between 2021Q1 to 2023Q3; panel (a) reports the spending-weighted average across categories. Panel (b) shows the difference between our cost-of-living index and a homothetic Törnqvist approximation. For each product category, we compute the percent change in the quantity-weighted average 2023Q3 price, comparing 2021Q1 and 2023Q3 weights; panel (c) reports the spending-weighted average across categories.

Figure C.7: *By changes in annual expenditure*

(a) Movements along the quality ladder

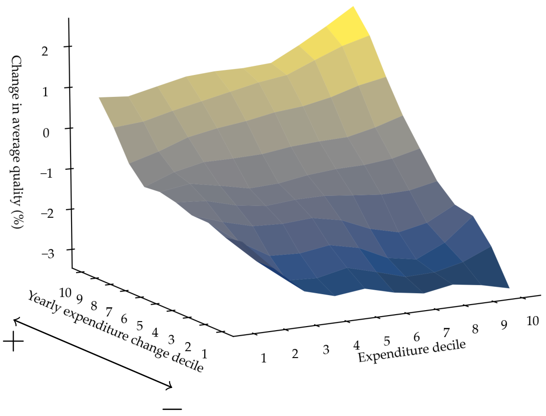

(b) Non-homotheticity cost-of-living bias

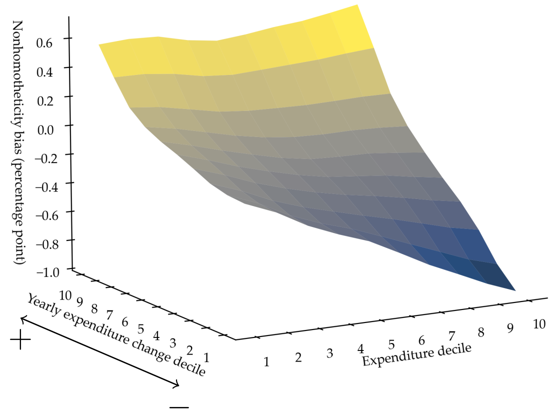

(c) Change in price level of consumption baskets

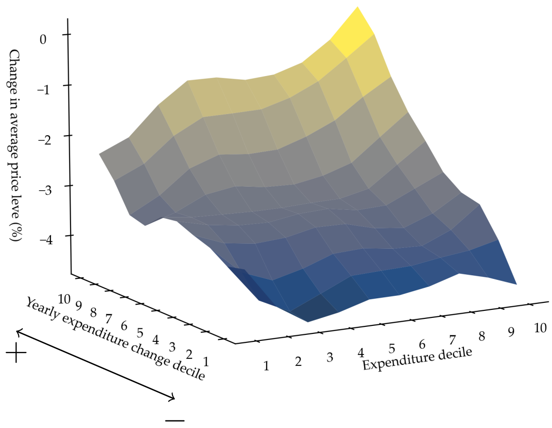

Notes: Authors' calculations using Kantar's Take Home Purchase Panel (2021-2023). For each product category, we compute the percent change in the quality rung of purchases between 2021Q1 to 2023Q3; panel (a) reports the spending-weighted average across categories. Panel (b) shows the difference between our cost-of-living index and a homothetic Törnqvist approximation. For each product category, we compute the percent change in the quantity-weighted average 2023Q3 price, comparing 2021Q1 and 2023Q3 weights; panel (c) reports the spending-weighted average across categories.