Gutenberg School of Management and Economics

& Research Unit “Interdisciplinary Public Policy”

Discussion Paper Series

Anti-poor and anti-rich: Product-downgrading and the distributional effects of UK inflation in the wake of the Brexit vote

Guenter W. Beck, Philipp Harms, Muzammil Hussain, Mark Ruszel

May 03, 2024

Discussion paper number 2408

Johannes Gutenberg University Mainz
Gutenberg School of Management and Economics
Jakob-Welder-Weg 9
55128 Mainz
Germany
https://wiwi.uni-mainz.de/

Contact details

Guenter W. Beck
Chair for European Macroeconomic Studies
University of Siegen
Unteres Schloß 3
57072 Siegen
Germany
guenter.beck@uni-siegen.de

Philipp Harms
Chair of International Economics
Johannes Gutenberg University Mainz
Jakob-Welder-Weg 4
55128 Mainz
Germany
philipp.harms@uni-mainz.de

Muzammil Hussain
Chair for European Macroeconomic Studies
University of Siegen
Unteres Schloß 3
57072 Siegen
Germany
muzammil.hussain@uni-siegen.de

Mark Ruszel
Chair of International Economics
Johannes Gutenberg University Mainz
Jakob-Welder-Weg 4
55128 Mainz
Germany
mark.ruszel@uni-mainz.de

Anti-poor *and* anti-rich: Product-downgrading and the distributional effects of UK inflation in the wake of the Brexit vote

Guenter W. Beck* Philipp Harms† Muzammil Hussain‡ Mark Ruszel§

This version: May 3, 2024.

Abstract

In the second half of 2016, the United Kingdom experienced a strong increase of retail prices which was caused, among other factors, by a massive depreciation of the British pound in the wake of the Brexit vote. In this paper, we analyze the distributional effects of this inflationary episode, examining in particular the role that households' decisions to adjust their consumption behavior at the extensive margin within narrowly defined products have played in this context. Using a very granular scanner data set on purchases of fast-moving consumer goods, we demonstrate that households at an intermediate income level engaged in *product-downgrading*, i.e. they switched from higher-priced varieties of a given product to lower-priced varieties, and thus limited the effect of the overall price increase. By contrast, poor households had no scope for product-downgrading since they already consumed the lowest-priced varieties. Rich households, finally, also did not change the mix of varieties they consumed and thus experienced relatively elevated inflation rates as well – probably because their higher income allowed them to tolerate the price increase.

Keywords

inflation, distributional effects, scanner data, inflation inequality, product substitution

JEL-Codes

E31 · D31 · F15 · F31 · F41 · D12

*University of Siegen. Contact: guenter.beck@uni-siegen.de.
†Johannes Gutenberg University Mainz. Contact: philipp.harms@uni-mainz.de.
‡University of Siegen. Contact: muzammil.hussain@uni-siegen.de.
§Johannes Gutenberg University Mainz. Contact: mark.ruszel@uni-mainz.de.

1 Introduction

In this paper, we explore the distributional effects of inflation, focusing on the surge in retail prices following the massive depreciation of the British pound in the wake of the June-2016 Brexit vote. The "Brexit depreciation" represents a rare example of an exogenous inflation shock since the outcome of the referendum came as a surprise, was not associated with immediate macroeconomic turmoil, and raised the prices of both imported final goods and imported intermediate goods (Breinlich et al., 2022). Hence, it resulted in a price increase across a wide range of goods, and the consequences of this inflationary shock are not confounded by contemporaneous fluctuations of GDP or other macroeconomic variables.

Focusing on the evolution of British inflation between January 2016 and December 2017, we analyze how different income groups were affected by the overall price increase. While it is well-known that effective inflation rates may differ across different parts of the population – e.g. because of group-specific expenditure shares, combined with a heterogeneous evolution of prices, or because of different abilities to alter the composition of overall spending – we zoom in on one particular aspect that can represent an important mechanism leading to different inflation experiences, and that has not obtained much attention in the literature so far: households' ability and willingness to cushion the overall impact of the price increase by engaging in *product-downgrading*, i.e. by replacing more expensive varieties of a given product by less expensive varieties. With the term "product", we refer to a very narrowly and granularly defined product category, e.g. *Semi-Skimmed Milk*. In this case, "product varieties" (*items*) would represent semi-skimmed milk products from different manufacturers. Thus, the varieties we consider are almost perfect substitutes in terms of their basic properties (taste, nutritional value, etc.).

Using a scanner data set on household purchases of fast-moving consumer goods (FMCG), we are able to explore how prices effectively paid at the product level by members of different income groups evolved over time. More specifically, we compute *volume-share-weighted average unit prices* for a wide range of products and analyze whether the evolution of these averages was related to consumers' income. This allows determining to what extent different income groups resorted to lower-priced substitutes when the Brexit inflation shock occurred.

The results of our analysis suggest that, when we focus on the extent of product-downgrading, the distributional consequences of the Brexit depreciation were *anti-poorest* and – to some extent – *anti-rich*: this is because the poorest households in our sample tend to purchase the most affordable varieties within narrowly defined products, limiting their ability to further switch to more affordable varieties during the inflationary period following the Brexit referendum. In contrast, middle income-households have more flexibility to adjust their purchasing habits to evade inflation, resulting in lower inflation rates compared to the poorest households. We demonstrate that this gives rise to quantitatively important differences in effective inflation rates: compared to the poorest households, the ability to adjust the composition of purchases at the product level lowers monthly price increases by up to 0.25 percentage points.[1] Richer households, despite having

[1] A simple example explains the size of these differences and illustrates the importance of product-downgrading: Suppose that

the capacity to substitute away from more expensive varieties, apparently choose not to, leading them to encounter inflation rates higher than those experienced by poor households, but still below the ones experienced by poorest households. While our study focuses on a particular historical episode that lends itself to an event-study design, we believe that the constellation we identify is of general relevance and therefore contributes to a better understanding of the distributional effects of price increases.

The rest of this paper is structured as follows: in Section 2, we review the relevant literature on the distributional effects of inflation. Section 3 shortly summarizes the evolution of exchange rates, prices and macroeconomic aggregates shortly before and after the Brexit vote of June 2016. In Section 4, we present the data set that is underlying our analysis. Section 5 starts by describing our approach to compute volume-share-weighted price averages for different products. It then demonstrates that these averages significantly differ across income groups, with low-income households purchasing a bundle of lower-priced varieties than medium-income and high-income households. In a next step, we compute *price relatives* at the product level, which allow tracing the evolution of price averages over time. We show that price increases for medium-income households are significantly lower than price increases for low-income and high-income households. While these differences could be driven both by changing volume weights and a heterogenous evolution of variety-prices, we show that the difference between a Paasche-type index at the product level (which allows for changing volume weights) and a Laspeyres index (which keeps weights fixed) is lowest for medium-income households. Moreover, we demonstrate that the difference in Paasche-type indices across income groups disappears if we restrict the sample to varieties purchased in adjacent periods. We interpret this as evidence that medium-income households stop purchasing expensive varieties and replace them by cheaper varieties, thus cushioning the overall effect of inflation. Section 6 summarizes our findings and offers some conclusions.

2 Related literature

Our work draws from (and contributes to) various strands of the literature on the distributional effects of inflation.[2]

Ha et al. (2019), offering a comprehensive overview of inflation dynamics over the last 50 years, identify three direct channels through which inflation may vary across the income distribution: the composition of income channel (e.g., labor income vs. profits), the composition of assets channel (e.g., cash vs. stock ownership) and the composition of consumption baskets channel. Our study contributes to a better understanding of this third channel which states that, because households choose different product categories (e.g., meat vs. vegetables), or use differently priced versions within the same product categories (e.g., chicken vs. pork), effective inflation rates may differ.

consumers have access to two varieties of the same product, one costing 1 GBP, the other 2 GBP, and that both prices increase by ten percent, i.e. to 1.1 GBP and 2.2 GBP, respectively. Without an adjustment at the extensive margin, all households would be exposed to the same inflation rate. However, households switching from the second to the first variety actually experience a negative inflation rate of -45 percent.

[2] Inflationary consequences of the Brexit vote, another strand of the literature that we relate to, are discussed in Section 3.

Due to the scarcity of detailed micro data, much of the earlier related literature concentrated on quantifying the extent to which variations in expenditure shares across different product categories result in disparities in inflation rates (Michael, 1979, Hagemann, 1982, Cage & Garner, 2002, Garner & Ruiz-Castillo & Sastre, 2003, Crawford & Oldfield, 2002, Hobijn & Lagakos, 2005, Gürer & Weichenrieder, 2020). However, the omission of the within-category dimension has faced some criticism. Jaravel (2021) emphasizes the importance of considering spending patterns within narrowly defined product categories, as expenditure shares and product choices can vary significantly across income groups. He stresses the necessity of using granular data which capture the *effective* prices paid as well as the expenditure shares and the range of products chosen by different income groups, since such detailed data are crucial to accurately identify disparities in inflation rates. In his contribution, Jaravel (2021) identifies what he labels an 'aggregation bias' – a distortion in inflation inequality that arises when this within-category dimension is overlooked. Adopting a similar perspective, Kaplan & Schulhofer-Wohl (2017) provide empirical evidence supporting the significance of accounting for price and product mix variations within narrowly defined product categories. They demonstrate that, when allowing households to pay different prices for identical products and when acknowledging that the assortment of products within product categories can vary across households, the disparity in inflation rates between the lowest and highest income groups in the US from 2004 to 2013 is five times larger compared to a situation where uniform prices are assumed and the same mix of products is used across all households. Consequently, Kaplan & Schulhofer-Wohl (2017) conclude that standard approaches, which presuppose uniform prices and identical product mixes for all households within product categories, fail to capture the heterogeneity in inflation rates.

As more granular data have become available, recent studies have begun to incorporate the within-category dimension in the estimation of inflation heterogeneity across income groups (Broda & Romalis, 2009, Kaplan & Schulhofer-Wohl, 2017, Jaravel, 2018, Argente & Lee, 2021). A consistent finding emerging from this strand of literature is that poorer households, on average, face significantly higher inflation rates. Various explanations have been proposed to account for this observed inflation inequality. Jaravel (2018) investigates the influence of product innovations on inflation inequality in the United States over the period from 2004 to 2015. He posits that a key driver of this inequality is the accelerated rate of innovation in product categories that are predominantly favored by high-income households. Orhun & Palazzolo (2019) show that liquidity constraints inhibit low-income households from taking advantage of bulk discounts and temporary sales in the US from 2006-2014. Argente & Lee (2021) find that high-income households faced lower effective inflation rates in the US during the Great Recession because they were better able to substitute towards lower-quality products, which experienced lower price increases.[3]

Another area of literature that our paper engages with concerns the heterogeneous price effects of various macroeconomic shocks. Notable studies in this field have examined the impacts of trade liberalization (Porto,

[3]Despite several parallels, our analysis differs from the contribution of Argente & Lee (2021) by considering changes in consumption choices following an inflationary shock (and not during a recession), and by putting a stronger emphasis on product-downgrading as a channel of adjustment. It also presents the non-monotonic relationship between income levels and price increases at the product level as a novel finding.

2006, Fajgelbaum & Khandelwal, 2016), monetary policy shocks (Ampudia & Ehrmann & Strasser, 2023, Cravino & Lan & Levchenko, 2020) and, most pertinently to our research, the effects of large currency devaluations (Cravino & Levchenko, 2017, Colicev & Hoste & Konings, 2022). A significant focus of the earlier studies has been on understanding how heterogeneity in expenditure shares on product categories generates distributional effects across rich and poor consumers (Porto, 2006, Fajgelbaum & Khandelwal, 2016). Expanding the analytical scope to include the within-category dimension, Cravino & Levchenko (2017) show that heterogeneity in expenditure shares both across and within detailed product categories lead to anti-poor distributional consequences of the 1994 Mexican peso devaluation. They make use of two extremely rich microdata sets. First, on household-level expenditures on detailed product categories, and second, on unique retailer-product level prices. While price data is on a monthly frequency, the expenditure shares are sourced from occasional household surveys, hence the study does not account for the possibility that consumers rapidly adjust their spending in response to price shifts. This aspect, however, is considered to be of particular importance in the ongoing literature concerning the distributional effects of inflation. Cravino & Levchenko (2017)'s work thus sheds light on the anti-poor consequences of the peso devaluation while also highlighting a critical area for further research: the real-time adaptability of consumption behavior.

Colicev & Hoste & Konings (2022), using detailed price and quantity data at the consumer level, study the role of heterogeneity in foreign expenditure shares within a product category and changes in the set of available products before and after the large depreciation of the Kazakh tenge in 2015. They show that heterogeneity in foreign expenditure shares as such does not lead to important distributional inflation effects. By contrast, large changes at the extensive margin as well as heterogeneity in elasticities of substitution lead to a substantially lower increase in inflation for rich consumers relative to poor consumers. Braun & Lein (2020), using Swiss household scanner data, link inflation heterogeneity among households to the heterogeneity in their preferences. They demonstrate that inflation dispersion is time-varying, being lower during periods of near-zero inflation but higher when inflation deviates from zero. Kiss & Strasser (2024), using German and French household scanner data spanning more than a decade, examine the sources of inflation heterogeneity among households. They observe significant and persistent disparities, primarily stemming from two direct sources: variations in the prices paid for identical products and household-specific choices of product varieties within categories. They also identify an indirect source for inflation disparities resulting from income heterogeneity influencing the shopping behavior of households. In a similar vein, Strasser et al. (2023) document large idiosyncratic inflation differences between households using a multi-year household scanner data set for euro area countries. They observe substantial dispersion in inflation across households, which sometimes persists for a year or two but is not permanent. They also note that during recessions, low-income households experience higher inflation compared to high-income households.

Ampudia & Ehrmann & Strasser (2023) study the effect of monetary policy on the effective inflation rates experienced by low- and high-income earners in euro area countries. They utilize a household panel that captures data on prices and quantities purchased, along with socio-demographic details of the purchasing households, and they highlight product substitution as a key aspect of consumer behavior, with its extent varying between high-income and low-income households. Notably, they observe that inflation experienced by high-income households is more sensitive to monetary policy adjustments, which is attributed to changes in shopping behavior; specifically, following a contractionary monetary policy shock, high-income households tend to increase their shopping intensity compared to low-income households and engage in more pronounced product substitution towards those versions that became relatively cheaper. By contrast, our analysis indicates that, confronted with the inflationary shock in the wake of the Brexit vote, medium-income households were most active in product-downgrading.

Our study is also related to the literature on the distributional effects of the British pound depreciation in the wake of the June-2016 Brexit vote. Focusing on this particular episode allows us to examine the distributional effects of foreign exchange shocks on consumer prices in the context of an advanced economy. This is important because existing literature on inflation heterogeneity, driven by exchange rate fluctuations, has predominantly considered developing economies. More generally, the body of literature concerning the distributional effects of inflation – e.g., in times of economic downturns as in Argente & Lee (2021) and Coibion & Gorodnichenko & Hong (2015) – has primarily analyzed the United States. Our study expands this focus by providing empirical evidence from another advanced economy, the United Kingdom.

Our contribution to the existing literature can be summarized as follows: first, by leveraging the granular information on item prices and household purchases in our household scanner data set and by comparing *volume-share-weighted price averages* for different products over time and across households, we can assess how substitution within products impacts inflation rates across income groups. Furthermore, our data set allows us to incorporate the finding that changes in the item availability, such as new introductions and replacements, are important facets for the identification of exchange rate pass-through (Nakamura & Steinsson, 2012, Cavallo & Neiman & Rigobon, 2014, Goetz & Rodnyansky, 2023, Corsetti et al., 2023). This is particularly relevant in the context of scanner data, in which items keep moving in and out of households' consumption baskets. Second, the availability of socioeconomic data, especially household income, enables us to assess how inflation affects different households across the income distribution. In particular, we are not constrained to using transaction-level data to categorize consumers into different income groups (see, e.g., Colicev & Hoste & Konings, 2022). Instead, we can directly divide households into seven distinct income groups. This level of granularity goes beyond the typical binary subdivision into 'rich' and 'poor' consumers commonly seen in many studies. Thus, we can investigate potential non-monotonicities in the heterogeneity of realized inflation across income groups. As we will demonstrate below, this gives rise to some novel and relevant findings.

3 The Brexit vote and the depreciation of the British pound

As already noted by Broadbent (2017), the surprising outcome of the Brexit vote in June 2016 triggered a substantial depreciation of the British pound. Figure 1 demonstrates that, between June and October 2016, the pound (GBP) lost about 15 percent of its value against the US dollar (USD), the Chinese yuan (CNY) and the Euro (EUR). As for the USD and the CNY, some of this depreciation was reversed in subsequent months. By contrast, the GBP remained persistently cheaper vis-a-vis the EUR.

Not surprisingly, the depreciation resulted in a significant increase in consumer prices, as put forward by Gerstein et al. (2019), Breinlich et al. (2022) and Dhingra & Sampson (2022): first, the drop of the pound's value had a direct effect on the prices of imported goods. Second, increasing prices of imported intermediate inputs raised the costs of British producers, and thus also contributed to rising prices of domestically produced goods (Breinlich et al., 2022). While it is well-known that exchange-rate pass-through into consumer prices is much weaker than passthrough into border prices (Burstein & Gopinath, 2014), the evolution of the British CPI in Figure 2 indicates that consumer prices increased by five percent between mid-2016 and the end of 2017. Figure 2 also shows the evolution of sub-indices of the British CPI, focusing on the type of products that are included in our data set on fast-moving consumer goods (FMCG). With the exception of personal care items, the price indices for these goods categories also increased substantially.

Figure 1: British pound (GBP) nominal exchange rates

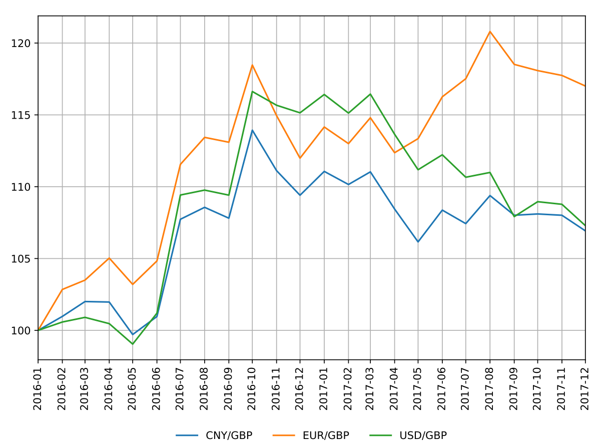

Figure 2: UK prices

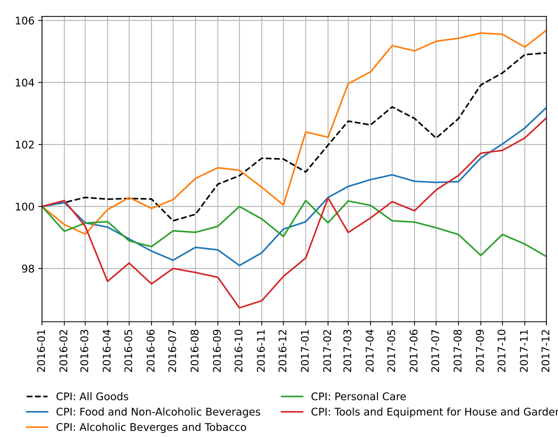

Note: Figure 1 illustrates a selection of bilateral nominal exchange rates involving the British pound. Figure 2 displays price indices in the UK for all goods and specific divisions and subdivisions that closely align with the products encompassed in our scanner data set. Data has been sourced from the UK Office for National Statistics (ONS).

Importantly, the depreciation of the GBP in the wake of the Brexit vote was not associated with major macroeconomic turmoil: while financial markets re-assessed their perspective on the British economy, the UK did not experience an immediate recession or any type of crisis, as evidenced in Figure 3. This distinguishes the Brexit vote-induced depreciation of the British pound from other episodes, which either consider the effects of large depreciations during currency crises (Cravino & Levchenko, 2017) or the evolution of prices during major recessions (Argente & Lee, 2021). This, in turn, implies that both the price changes and the

Figure 3: UK economic indicators shortly before and after the June-2016 Brexit vote

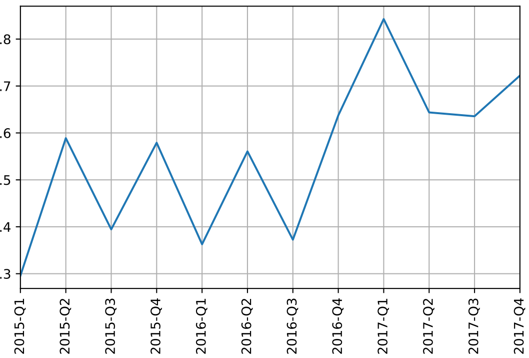

(a) Real GDP, QoQ (in %)

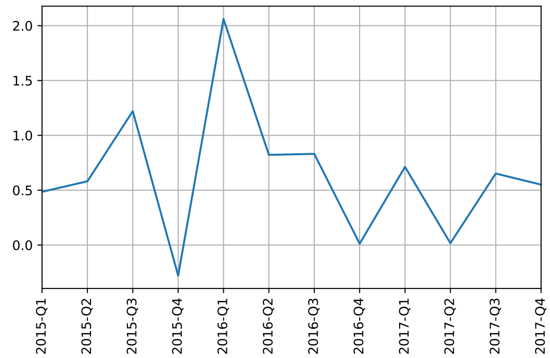

(b) Real hh consumption expenditure, QoQ (in %)

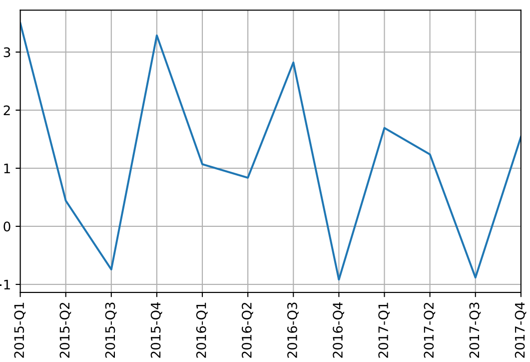

(c) Real investment, QoQ (in %)

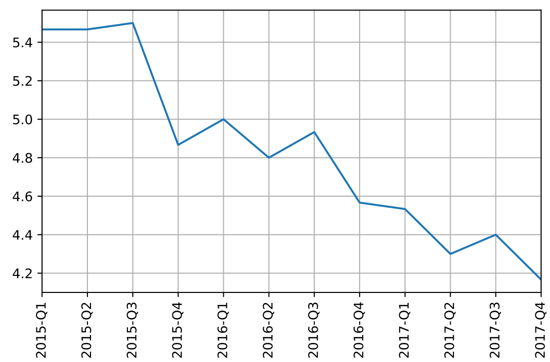

(d) Unemployment rate (in %)

Source: International Monetary Fund

Note: All series are seasonally adjusted except for the unemployment rate.

changes in spending patterns that we observe are unlikely to be driven by forces that go beyond the GBP depreciation.

4 Data set

Our analysis leverages a unique scanner data set in the United Kingdom (UK) spanning the years 2016 and 2017. More specifically, we utilize data from the Kantar Fast-Moving Consumer Goods (FMCG) Purchase Panel, a leading repository for household scanner data. The data set provides detailed transaction information on household purchases of FMCG, which include products typically available in supermarkets, such as food, drinks, alcohol, personal care, household cleaning, cosmetics, etc. Put in the context of the Consumer Price Index (CPI), FMCG predominantly belong to the broad categories *Food and non-alcoholic beverages*, *Alcoholic beverages, Tobacco and narcotics*, *Personal care*, and *Tools and equipment for house and garden*. As per the most recent CPI weights published by the Office for National Statistics (ONS), these four broad categories represent more than 20 percent of total household expenditure in the UK.[4]

The data set provides information on broad product categories (e.g., food and non-alcoholic beverages or personal care), narrower product categories (e.g., dairy products or vegetables), across product types within these categories (e.g., cheese), across products (e.g., gouda cheese) or across items (e.g., a specific gouda cheese from a given brand, identified by a unique code provided by Kantar).[5] For each transaction, we observe details on (i) the item purchased, (ii) the purchasing household, (iii) the retailer where the transaction took place, (iv) the total transaction value, and (v) the total quantity transacted. This transactional information is supplemented with detailed item characteristics, encompassing the associated product, product type, product category, brand and manufacturer as well as volume details. A sample of these items can be seen in Table 1. It is important to note that the products are defined with a high level of granularity. For example, *Pet Food* is subdivided into types of pet like *Dog*, *Cat*, etc., and further segmented into *Daily Nutrition*, *Treats*, and other types. Similarly, *Soft Drinks* are categorized by taste (e.g., *Bitter Lemon*, *Lemon*, *Cola*, etc.) and by calorie content (*Normal*, *Low Calorie*, etc.). *Beer* is also categorized into different types such as *Lager*, *Pilsner*, *Stout*, etc.

Table 1: Some examples of items in our household scanner data set

Item Code*	Product Type	Product	Measurement Unit	Volume
6460	Razor Blades	Double Edge Razor Blades	Piece	5 in a Pack
7439	Cat and Dog Treats	Cat Treats	Gram	55
9892	Milk	Semi-Skimmed Milk	Millilitre	2000
13239	Dry Pasta	Dry Pasta Fusilli	Gram	500
13953	Bitter Lemon	Low Calorie Bitter Lemon	Millilitre	1000
20582	Toilet Tissues	Soft Toilet Rolls	Piece	1
29915	Sun Care	Sun Care Aftersun	Millilitre	400
40997	Spirits	Spirits Rum	Millilitre	1000
41935	Nuts	Nuts Snacks	Gram	50
43879	Beer	Stout Beer	Millilitre	4 X 440
45950	Mineral Water	Mineral Water Flavored	Millilitre	4 X 500
49278	Hair Styling Wax	Hair Styling Wax Creams	Gram	100
59981	Cheese	Gouda Cheese	Gram	200
63315	Sugar Confectionery	Sugar Candy	Piece	50
69938	Popcorn	Popcorn Sweet+Savoury	Gram	30
70285	Sugar	Icing Sugar	Gram	500

Note: The data set also offers information on item name, description, product category, brand and manufacturer. *Item codes are randomized to comply with non-disclosure agreements.

Transactions in the data set occur in over 20 distinct retailer types, ranging from local neighborhood shops to hypermarkets. Retailers are not segmented based on their regional branches or locations; instead, they are aggregated under a unified identifier. For instance, all regional outlets of a specific supermarket

[4] It is interesting to compare expenditure on these broad categories over different income groups. According to data extracted from the *ONS report on family spending in the UK*, in the fiscal year 2015/16, households belonging to the lowest income decile dedicated 17.3 percent of their total expenditure to food and non-alcoholic beverages. In contrast, households belonging to the highest income decile allocated only 7.5 percent of their total expenditure to this broad category. Moreover, it is observed that as one moves down the income distribution, there is a corresponding increase in the proportion of household expenditure allocated to food and non-alcoholic beverages. As a consequence, changes in the prices of these products are likely to disproportionately affect the cost of living of lower-income households.

[5] For products for which a GTIN (global trade identification number, also denoted as EAN in a European context or UPC in a US context) is available, this code corresponds to this GTIN, for other products (such as fresh fruit or meat), Kantar generates a GTIN equivalent code.

chain are collectively treated as a single entity in the data.

To evaluate the accuracy of the scanner data in reflecting price movements, we calculate a price index for the items contained in the food and non-alcoholic beverages (FnB) categories. When doing so, we follow the price of a product at a given retailer type (henceforth denoted as item-retailer pair) as the basic reference unit. This index is then compared to the corresponding official price index provided by the ONS in Figure 4.[6] The close alignment of the two time series highlights that the sample of products and prices covered by our scanner data set does not deviate from the sample underlying the official FnB index.

Figure 4: Food and non-alcoholic beverages price index, ONS and scanner data

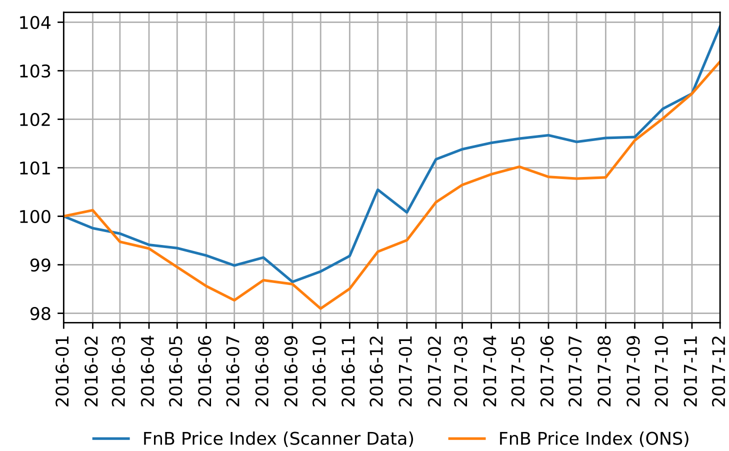

Note: This figure illustrates the food and non-alcoholic beverages (FnB) price index derived from scanner data alongside the official FnB index provided by the ONS. To construct the scanner data price index, we focus exclusively on item-retailer pairs that show transactions throughout 2016 and 2017 on a monthly basis. This selection ensures the inclusion of items that remain on the market for an extended period. Utilizing this subset, we calculate the cumulative product of the period-on-period Fisher price index, which accounts for both previous and current period expenditures.

Households in our data set, which serve as the observational panels, are characterized by demographic information sourced from survey questionnaires. For each household, specific attributes, including the annual income and age of the primary earner, as well as household size and the place of residence, are available. The place of residence is determined using the official UK postcode format, established by the General Post Office (Royal Mail). We utilize information on the outward code, specifically the 124 postcode areas in the United Kingdom as shown in Figure 5. A postcode area corresponds to the first two digits of a zipcode.

With regard to income, Kantar reports that panelists communicate their total labor income before taxes.[7] We utilize Kantar's categorization of households into seven distinct income brackets (in £10,000 increments). Additionally, we apply certain constraints to the households considered in our analysis. First, we keep only those households that report their income either in 2016 or 2017. Second, we include only those households that have consistently reported transactions in every month throughout both 2016 and 2017. This restriction ensures that our results are not distorted by households with specific spending patterns moving into or out of

[6]The price index plotted in Figure 4 corresponds to a Fisher price index. Similar plots are obtained using either a Tornqvist or a Laspeyres index.

[7]It should be noted that the actual household income may exceed the labor income of the primary earner, encompassing various income streams of different household members.

the sample. Third, we restrict our analysis to households that have reported the same postcode in both 2016 and 2017. This criterion is applied to ensure that any variations in expenditure patterns are not influenced by households changing their place of residence. Fourth, our study focuses on households with a size ranging from 1 to 6 members. This range is selected to represent a broad spectrum of household compositions, from single individuals to larger family units, while excluding exceptionally large households that might exhibit atypical consumption behaviors. Fifth, we restrict our analysis to households where the age of the head of the household falls between 18 and 65 years. This age range is chosen to focus on economically active individuals, excluding those typically in full-time education or retirement, as their spending habits may significantly differ from those of the working-age population. Table 2 displays the count of households within each income bracket. The size of these income groups roughly coincides with the actual UK income distribution, as published by the ONS.[8] Moreover, Figure 6 demonstrates that across most regions, households from all seven income groups are present.

Table 2: Distribution of households by income group

Income Group	Income Bracket (in GBP)	Number of hh
1	up to 9,999	716
2	10,000 - 19,999	2,106
3	20,000 - 29,999	2,275
4	30,000 - 39,999	1,997
5	40,000 - 49,999	1,454
6	50,000 - 59,999	928
7	above 60,000	523

Note: This table presents the number of households by income group. We utilize Kantar's categorization of households into seven distinct income brackets (in £10,000 increments).

In our assessment of product-downgrading, we apply two distinct definitions for the unit of analysis. These approaches incorporate specific criteria including the item code and income group, as well as, potentially, the retailer and postcode area dimensions. Accordingly, a *pi*-pair is defined as a specific item (=product variety) p purchased by households belonging to income group $i \in I$. When incorporating the retailer dimension and the postcode area of the purchasing household, we define a *prai*-pairing as a specific item (=product variety) p purchased from retailer $r \in R$ by households residing in postcode area $a \in A$ and belonging to income group $i \in I$.[9] Table 3 provides a summary of the data set, broken down by year. Particularly notable is the high number of products, underscoring the fine granularity with which they are delineated. Furthermore, the table also illustrates that a substantial portion of the *pi*-pairs and *prai*-pairings are accompanied by volume information.

[8]To compare the distribution of incomes in our sample to the income distribution in official statistics, we used the information provided by Office for National Statistics (2023) and deflated nominal income levels by the inflation rate between 2016 and 2022.

[9]This implies that we incorporate the income group dimension into the definition of the unit of analysis. As a consequence, the prices for items sold within the same month, and, potentially, at the same retailer within the same postcode area, can vary across income groups.

Figure 5: Map of UK postcode areas

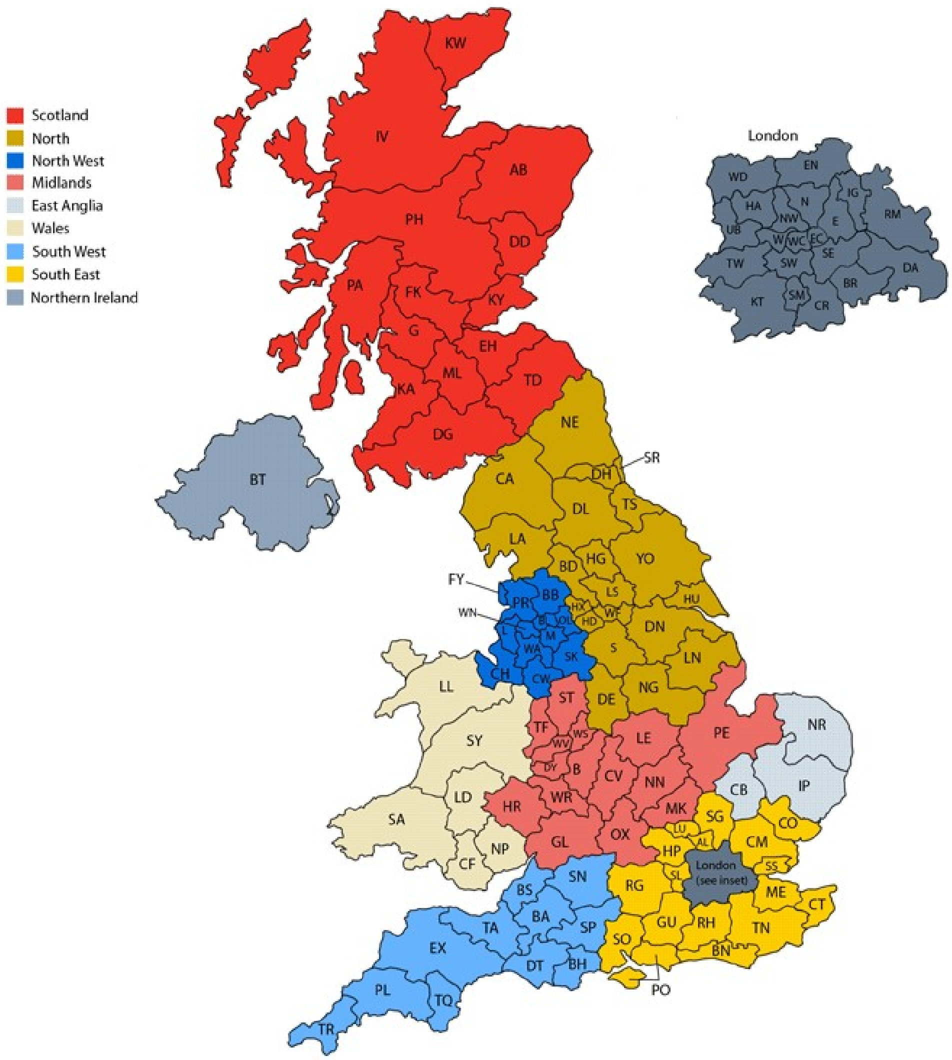

Source: www.electricmarketing.co.uk/map-uk-postcodes

Note: This figure shows the 124 postcode areas in the United Kingdom. A postcode area corresponds to the first two digits of a zipcode.

5 Analysis

5.1 Computing volume-weighted price averages

Our goal is to explore whether – and to what extent – British households cushioned the overall price increase in the wake of the Brexit vote by engaging in *product-downgrading*, i.e. by shifting the composition of their consumption at the product level towards less expensive varieties. Note that this may imply assigning changing weights to varieties purchased in two adjacent periods, but also completely dropping certain varieties from the shopping list or including others that had not been purchased before.

Figure 6: Distribution of income groups across postcode areas

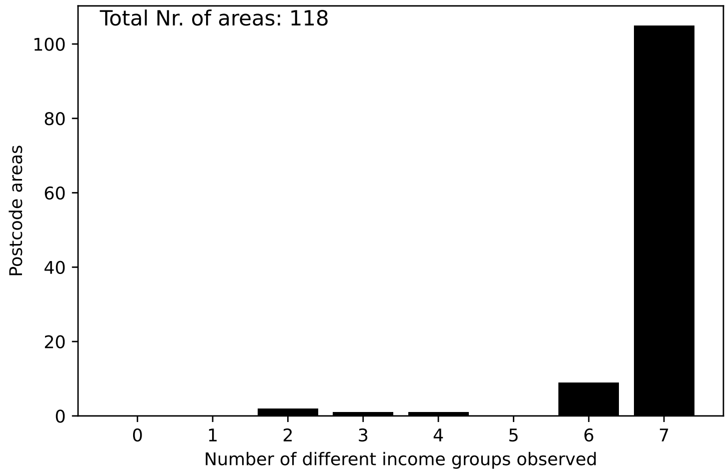

Note: This figure gives the number of different income groups observed in our data set across postcode areas in the United Kingdom. In the majority of postcode areas, households from all seven income groups are observed.

Table 3: Summary of the data set by year

Year	**2016**	**2017**
Number of Items	122,689	125,221
Number of Retailers	24	24
Number of Areas	118	118
Number of Income Groups	7	7
Number of Product Types	491	487
Number of Products	1701	1678
Number of *prai* -pairings	6,301,362	6,178,300
(of which with Volume Info)	(4,942,730)	(4,464,695)
Number of *pi* -pairs	627,918	632,555
(of which with Volume Info)	(506,197)	(438,058)

Note: We adopt a multi-dimensional approach to define the unit of analysis, incorporating the item and income group, as well as – potentially – the retailer and postcode area dimension. Accordingly, a *pi*-pair is defined as an item (=product variety) p purchased by households belonging to income group $i \in I$. When incorporating the retailer dimension and the postcode area of the purchasing household, we define a *prai*-pairing as an item p purchased from retailer $r \in R$ by households residing in postcode area $a \in A$ and belonging to income group $i \in I$.

To assess whether average prices paid for a given *variety-mix* within a product differed across income groups and how these differences evolved over time, we start by computing *volume-weighted price averages*. Our data set offers information on unit prices (price per, e.g., gram or milliliter) and on purchased volumes (in, e.g., grams or milliliters) of a *prai*-pairing from which we can derive volume-weighted price averages at the *gai-product-group* level. This product group *gai* is defined to include all items p within a specific product $g \in G$, purchased by households living in a specific postcode area $a \in A$ and belonging to a specific income group i.[10]

The *volume share* μ of a specific item-retailer-area-income group combination (i.e. of a *prai*-pairing) at

[10]For example, consider a particular *prai*-pairing, which could be a specific wheat bread from a given brand (e.g., '$WheatBread_1$') purchased from a specific retailer ('$Retailer_1$') by households living in a particular postcode area ('$Area_1$') and belonging to a specific income group (e.g., the lowest income group). Expanding this concept to the *gai*-product-group, we encompass a range of wheat bread items purchased by members of the lowest income group living in Area 1 at any retailer.

a given point in time t within the product g is given by the following expression:

$$\mu^{g}_{prai,t} = \frac{V_{prai,t}}{\sum_{p\in g}\sum_{r\in R} V_{prai,t}} \tag{1}$$

In (1), $V_{prai,t}$ represents the volume purchased (for instance, in grams or milliliters) of the *prai*-pairing at time t. The denominator, $\sum_{p\in g}\sum_{r\in R} V_{prai,t}$, sums the volumes purchased for all items p that belong to product g at time t across all retailers r. To illustrate, our approach allows us to calculate the time-dependent volume share of an item p, say a specific wheat bread from a given brand, purchased at a specific retailer r, within its product g (wheat bread) across all retailers r, taking into account both the area a that purchasing households live in and the income group i they belong to.

Following the calculation of volume shares, we proceed to compute the *volume-share-weighted average unit price* $\bar{P}^U$ for each *gai*-product-group at time t:

$$\bar{P}^{U}_{gai,t} = \sum_{p\in g}\sum_{r\in R}(\mu^{g}_{prai,t} \cdot P^{U}_{prai,t}) \tag{2}$$

where $P^{U}_{prai,t}$ is the *unit price* of a *prai*-pairing at time t. $\sum_{p\in g}$ and $\sum_{r\in R}$ indicate the summation over all varieties that are part of product g and across all retailers r, respectively. As the average unit price is not simply the arithmetic mean, but instead weighted by the volume share of each *prai*-pairing, it is representative of the actual unit price households of a certain income group living in a certain area pay when purchasing product g (for instance, what lowest income households living in South East London pay for one volume unit of wheat bread). Therefore, each *gai*-product-group in our analysis represents a collection of all items, differentiated by retailer, that belong to the same product, transacted by households from a particular area and income group. Note that we deliberately add over all retailers when computing product-group averages, since we allow households to substitute items not only *within* a certain retailer, but potentially also *across* retailers.

The analysis of the following subsections will focus on the question whether, shortly before and after the Brexit vote, the level and evolution of average prices at the product level differed across income groups. Note that the use of *volume shares* allows for a more direct measure of the average price paid for a (composite) unit of a product than, e.g., *expenditure weights*, which combine information on both quantities *and* prices and assign greater importance to varieties with higher prices.[11] Of course, volume-share weighting would

[11] A simple example with two varieties illustrates the point: the weighted average of the varieties' prices can be written as $\bar{P}_t = P_{2,t}\left(1 - \lambda_{1,t}\frac{P_{2,t}-P_{1,t}}{P_{2,t}}\right)$, with $\lambda_{1,t}$ denoting the weight assigned to the price of variety 1. While $\lambda^{VS}_{1,t} = \frac{V_{1,t}}{V_{1,t}+V_{2,t}}$ in the case of *volume-share weighting* , $\lambda^{ES}_{1,t} = \frac{V_{1,t}}{V_{1,t}+\frac{P_{2,t}}{P_{1,t}}V_{2,t}}$ in the case of *expenditure-share* weighting. Defining $v_t = \frac{V_{2,t}}{V_{1,t}}$ and $p_t = \frac{P_{2,t}}{P_{1,t}}$ and inserting the alternative weights into the definition of the average price yields $\bar{P}^{VS}_t = P_{2,t}\left[1-(1+v_t)^{-1}(1-\frac{1}{p_t})\right]$ for volume-share weighting and $\bar{P}^{ES}_t = P_{2,t}\left[1-(1+p_tv_t)^{-1}(1-\frac{1}{p_t})\right]$ for expenditure-share weighting. It is easy to show that $\bar{P}^{ES}_t > \bar{P}^{VS}_t$ unless $p_t = 1$, irrespective of the volumes actually purchased. The intuitive explanation is that, with expenditure-share weighting, the higher price tends to dominate the mean. If an overall increase of the price level is associated with a change in relative prices, this result is likely to carry over into the computation of inflation rates.

not be possible if we did not have such granular data, since adding quantities that refer to different products would not make sense.[12]

5.2 Differences in *price levels* across income groups

In this subsection, we demonstrate that the average prices paid by households for the mix of varieties within a given product significantly differ across income groups. Table 4 presents the results of regressing the log of the volume-share-weighted average unit price $\bar{P}^U_{gai,t}$ from equation (2) on income group dummies, as well as product, area and time fixed effects, i.e.

$$\ln \bar{P}^U_{gai,t} = \sum_{i=2}^{I} \beta_i \cdot \text{Income\_Group}_i + \delta_g + \sigma_a + \gamma_t + \epsilon_{gai,t} \tag{3}$$

Note that taking the logarithm of price averages transforms absolute differences into percentage deviations, thus addressing the problem that FMCG prices may differ substantially with respect to scale. The product and area dummies make sure that the effect of belonging to a certain income group is not confounded with the effect of consuming specific products or living in a specific area. Finally, the time dummy controls for the overall price increase in the wake of the Brexit depreciation. The estimated coefficients β_i thus indicate by how much average prices paid by a representative member of income group i differed from the average prices paid by a representative member of income group 1, which is the omitted category. Figure 7 plots the coefficients displayed in Table 4 with 95-percent confidence intervals. These results indicate that average prices paid by members of the lowest income group are significantly lower than the prices paid by the other income groups. Moreover, the average prices paid increase considerably as incomes rise. This suggests that middle-income households and high-income households purchase a relatively expensive mix of varieties to start with. As a consequence, their consumption bundle exhibits a large scope for within-product adjustment.

5.3 Differences in *price changes* across income groups

In a next step, we compute the gross growth rate of the volume-share-weighted average unit price for each gai-product-group at time t:

$$R^P_{gai,t} = \frac{\bar{P}^U_{gai,t}}{\bar{P}^U_{gai,t-1}} = \frac{\sum_{p \in g} \sum_{r \in R} (\mu_{prai,t} \cdot P^U_{prai,t})}{\sum_{p \in g} \sum_{r \in R} (\mu_{prai,t-1} \cdot P^U_{prai,t-1})} \tag{4}$$

In this equation, $R^P_{gai,t}$ represents a Paasche-type *price relative* for a gai-product-group at time t. It is derived by dividing the volume-share-weighted average unit price of the current period, $\bar{P}^U_{gai,t}$, by the volume-share-weighted average unit price of the preceding period, $\bar{P}^U_{gai,t-1}$.[13] This approach allows tracing the transition

[12] Ampudia & Ehrmann & Strasser (2023) also integrate volume-share weighted averages of product prices into their analysis of income group-specific inflation rates.

[13] Note that we are careful to call the expression in (4) a Paasche-*type* price relative, since it does not assign the weights of period t to the prices of periods $t-1$ and t (as a Paasche index *strictu sensu* would), but uses the current weights for both periods. We are choosing this strategy since it allows accounting for adjustments at both the intensive margin – i.e. changing

Table 4: Income groups and price levels at the product level

	Dependent variable
	Log volume-share-weighted average unit price
Income Group 2	0.0227***
	(0.0018)
Income Group 3	0.0290***
	(0.0021)
Income Group 4	0.0455***
	(0.0024)
Income Group 5	0.0629***
	(0.0027)
Income Group 6	0.0825***
	(0.0030)
Income Group 7	0.0879***
	(0.0033)
Observations	6,486,678
Product Fixed Effects	1,724
Area Fixed Effects	118
Time Fixed Effects	24
Adj. R-Squared	0.835
Standard Errors: Clustered (Product) in Parentheses	
***p<0.01, **p<0.05, *p<0.1	

Note: We assess whether households with lower incomes tend to purchase more affordable items compared to those with higher incomes. To do so, we regress the log of the volume-share-weighted average unit price $\bar{P}^U_{gai,t}$ from equation (2) on income group dummies, product, area and time fixed effects: $\ln \bar{P}^U_{gai,t} = \sum_{i=2}^{I} \beta_i \cdot \text{Income\_Group}_i + \delta_g + \sigma_a + \gamma_t + \epsilon_{gai,t}$. The β_i coefficients represent the estimated percentage difference in the volume-share-weighted average unit price associated with being in the respective income group i compared to the lowest income group (the reference income group).

towards other varieties within a *gai*-product-group. For example, if households of a certain income group that live in a certain area start buying more of a cheaper variety of a certain product, this substitution is reflected in the current period's volume shares, and thus, in the calculated $R^P_{gai,t}$. The transition potentially includes substitution among varieties transacted in two adjacent periods $t-1$ and t, but also towards a newly transacted variety in period t, and away from a variety transacted in period $t-1$ but not in period t. Therefore, this approach offers a simple way to deal with *newly transacted* and *discontinued* varieties. The Paasche-type price relative $R^P_{gai,t}$ is representative of the actual gross growth rate of the unit price households from a specific income group and area pay when purchasing one volume unit of product g (for instance, the gross growth rate of what households of a specific income group and area pay for one volume unit of wheat bread).

Table 5 presents the coefficients of regressing the price relative defined in (4) for the *gai*-product-group

weights of items purchased in both periods – and at the extensive margin – i.e. some items disappearing from and some items appearing on the shopping list. We invoke the term *Paasche* to signal that the evolution of average prices includes an adjustment of spending behavior.

Figure 7: Income groups and price levels at the product level: Estimated coefficients and 95% CI

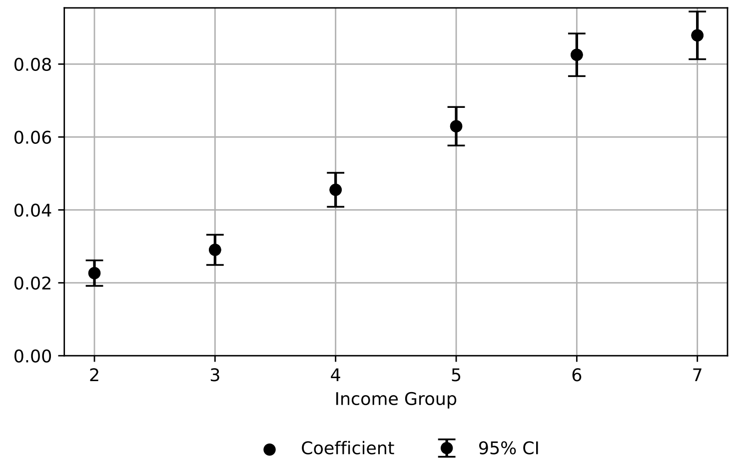

Note: This figure depicts the β_i coefficients and the corresponding 95% confidence intervals when regressing the log of the volume-share-weighted average unit price $\bar{P}^U_{gai,t}$ from equation (2) on income group dummies, product, area and time fixed effects: $\ln \bar{P}^U_{gai,t} = \sum_{i=2}^{I} \beta_i \cdot \text{Income\_Group}_i + \delta_g + \sigma_a + \gamma_t + \epsilon_{gai,t}$. The β_i coefficients represent the estimated percentage difference in the volume-share-weighted average unit price associated with being in the respective income group i compared to the lowest income group (the reference income group).

at time t on income group dummies, product, area and time fixed effects, i.e.

$$R^P_{gai,t} = \sum_{i=2}^{I} \beta_i \cdot \text{Income\_Group}_i + \delta_g + \sigma_a + \gamma_t + \epsilon_{gai,t} \tag{5}$$

Again, the product group and area dummies account for the potential correlation between income group-specific and product-specific or area-specific price increases. Moreover, the time dummies capture the overall price increases in the wake of the Brexit depreciation.

Figure 8 plots the coefficients displayed in Table 5 with 95-percent confidence intervals. The findings displayed in Table 5 and Figure 8 suggest that, controlling for product g, area a, and time t, the price increases experienced by income groups 2 to 4 were significantly lower than the price increases experienced by the lowest income group 1. More specifically, the price increase experienced for a given product g by a member of income group 3 was 0.52 percentage points lower, on average, than the price increase experienced by a member of income group 1 (which is the omitted category). Our point estimates indicate that average price increases were also lower for income groups 5 to 7 than for group 1, but not significantly so.

Note, however, that the above results may be driven by heterogeneous price increases at the variety level – medium-income households purchasing varieties whose prices changed by less – but also by differences in how various income groups adjusted their spending patterns – in particular, differences in the extent of product-downgrading. In the following subsections, we will explore to what extent changes in the variety-mix at the product level contributed to the results displayed in Table 5 and Figure 8.

Table 5: Income group-specific inflation differences at the product level

	Dependent variable
	$R^P_{gai,t}$
Income Group 2	-0.0041***
	(0.0010)
Income Group 3	-0.0052***
	(0.0011)
Income Group 4	-0.0049***
	(0.0011)
Income Group 5	-0.0018*
	(0.0010)
Income Group 6	-0.0003
	(0.0010)
Income Group 7	-0.0005
	(0.0009)
Observations	4,087,583
Product Fixed Effects	1,617
Area Fixed Effects	118
Time Fixed Effects	23
Adj. R-Squared	0.008
Standard Errors: Clustered (Product) in Parentheses	
***p<0.01, **p<0.05, *p<0.1	

Note: We assess whether households with higher incomes tend to experience lower inflation rates compared to those with lower incomes. To do so, we regress the Paasche-type price relative $R^P_{gai,t}$ from equation (4) on income group dummies, product, area and time fixed effects: $R^P_{gai,t} = \sum_{i=2}^{I} \beta_i \cdot \text{Income\_Group}_i + \delta_g + \sigma_a + \gamma_t + \epsilon_{gai,t}$. The β_i coefficients – when multiplied by 100 – represent the percentage-point difference in inflation associated with being in income group i compared to the lowest income group (the reference income group). Note that after having created the price relatives, we have excluded the 1st and 99th percentiles to mitigate the effect of outliers.

5.4 Identifying the role of product-downgrading: Comparing Laspeyres and Paasche-type indices

In order to assess how much households' adjusting of the variety-mix at the product level contributed to differences in income group-specific inflation rates, we compare Laspeyres price relatives to Paasche-type price relatives. As we will show below, this allows quantifying the price increase that households would have experienced without an adjustment of the variety-mix.[14]

To achieve this goal, we aggregate unit price levels at the *gi*-product-group level and then compute their period-on-period changes over time using two different weighting schemes. The primary reason for excluding the retailer and area dimensions from the definition of the unit of observation is that aggregating over retailers and areas significantly reduces fluctuations in the monthly purchasing sample over time: while the *prai*-pairing could be a specific wheat bread from a given brand purchased from a specific retailer by households living in a particular postcode area and belonging to a specific income group, the *pi*-pair would

[14] Ampudia & Ehrmann & Strasser (2023) follow a similar approach to assess the relevance of different margins of adjustment in generating heterogenous inflation rates.

Figure 8: Income group-specific inflation differences at the product level: Estimated coefficients and 95% CI

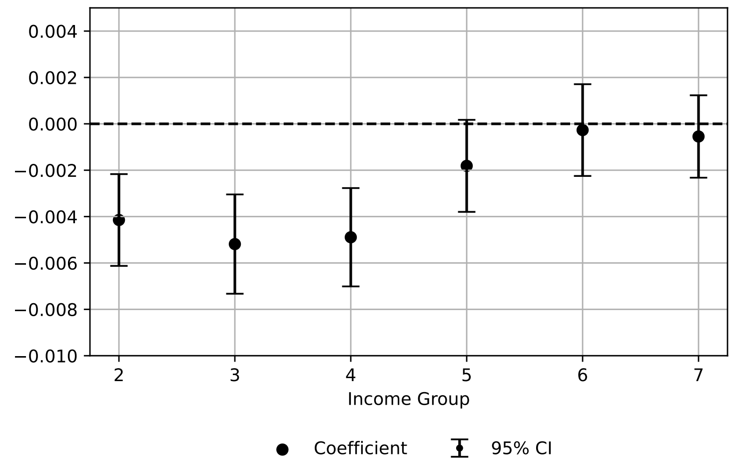

Note: This figure depicts the estimated β_i coefficients and the corresponding 95% confidence intervals when regressing the Paasche-type price relative $R^P_{gai,t}$ from equation (4) on income group dummies, product, area and time fixed effects: $R^P_{gai,t} = \sum_{i=2}^{I} \beta_i \cdot \text{Income\_Group}_i + \delta_g + \sigma_a + \gamma_t + \epsilon_{gai,t}$. The β_i coefficients – when multiplied by 100 – represent the percentage-point difference in inflation associated with being in income group i compared to the lowest income group (the reference income group). Note that after having created the price relatives, we have excluded the 1st and 99th percentiles to mitigate the effect of outliers.

just depict a specific wheat bread from a given brand purchased by households belonging to a specific income group. Consequently, the majority of *pi*-pairs are purchased in two adjacent periods. More specifically, approximately 85 percent of total expenditure is allocated to those *continued pi*-pairs, thereby diminishing the importance of the extensive margin – i.e., newly transacted and discontinued varieties. This allows restricting the data set to a sub-sample of continued *pi*-pairs and facilitates the comparison of period-on-period inflation rates based on Paasche-type and Laspeyres indices without sacrificing a significant number of observations.

The average unit *price relative* for each *gi*-product-group at time t, utilizing a weighting scheme based on the volume shares from the *previous* period, is given by:

$$R^L_{gi,t} = \frac{\sum_{p \in g}(\mu_{pi,t-1} \cdot P^U_{pi,t})}{\sum_{p \in g}(\mu_{pi,t-1} \cdot P^U_{pi,t-1})} \tag{6}$$

In this equation, $R^L_{gi,t}$ represents the *Laspeyres* price relative for the *gi*-product-group at time t. The numerator gives the volume-share-weighted average unit price of all *pi*-pairs within the *gi*-product-group at time t, each weighted by its respective *previous*-period volume share, defined as $\mu_{pi,t-1} = \frac{V_{pi,t-1}}{\sum_{p \in g} V_{pi,t-1}}$. The denominator depicts the volume-share-weighted average unit price of all *pi*-pairs within the *gi*-product-group at time $t-1$, each weighted by its volume share in perod $t-1$ This Laspeyres calculation gives us an insight into how the prices would have changed over time if the consumption pattern (in terms of volume distribution among varieties) had remained stable. As it relies on the consumption pattern of the previous period, it does not incorporate substitution within the *gi*-product-group that may have occurred in the current period. Note that if a variety was purchased in period $t-1$ but not in period t by members of a specific income group, it does not have a unit price $P^U_{pi,t}$ in period t. These varieties are referred to as *discontinued* varieties because they were purchased in the previous period but are no longer available or relevant in the current period. If a variety was purchased in period t but not in period $t-1$, its volume share $\mu_{pi,t-1}$ would be zero in period $t-1$. These varieties are referred to as *newly transacted* varieties because they were not available or relevant in the previous period, but are purchased in the current period. As both discontinued and new varieties impede the computation of $R^L_{gi,t}$, they are removed from the sample and we restrict our attention to those varieties that are transacted in two adjacent periods, i.e. to the sub-sample of *continued* varieties. As mentioned above, the aggregation over areas and retailers guarantess that a large share of varieties is continued at each point in time.

In order to assess the extent of substitution within products, we also compute Paasche-type price relatives, which are given by the average unit price relative for each gi-product-group at time t and $t-1$, weighted by the respective period's volume shares:

$$R^P_{gi,t} = \frac{\bar{P}^U_{gi,t}}{\bar{P}^U_{gi,t-1}} = \frac{\sum_{p \in g}(\mu_{pi,t} \cdot P^U_{pi,t})}{\sum_{p \in g}(\mu_{pi,t-1} \cdot P^U_{pi,t-1})} \tag{7}$$

Note that this approach allows tracing the shift of volume shares within a gi-product-group. For example, if consumers start buying more of a cheaper variety within a certain gi-product-group, this substitution is reflected in the current period's volume shares, and thus in the calculated $R^P_{gi,t}$. To retain the same sub-sample of varieties as in the Laspeyres price relatives, we keep our attention limited to the sub-sample of continued varieties.

The difference between the Laspeyres and Paasche-type price relatives, which quantifies the extent to which the inflation rate for a gi-product-group is influenced by substitution within that gi-product-group in the specified time period, is computed as:

$$\Delta \hat{P}_{gi,t} = R^L_{gi,t} - R^P_{gi,t} \tag{8}$$

Note that we can rewrite (8) as

$$\Delta \hat{P}_{gi,t} = \frac{\sum_{p \in g}(\mu^g_{pi,t-1} - \mu^g_{pi,t}) \cdot P^U_{pi,t}}{\bar{P}^U_{gi,t-1}} \tag{9}$$

The value yielded by (9) is high if, on average, households *reduce* the volume share of varieties within a gi-product-group that are *expensive* in period t. The volume-weighted average in the denominator ($\bar{P}^U_{gi,t-1}$) serves as a scaling variable, which makes sure that $\Delta \hat{P}_{gi,t}$ is comparable across products, and that the expression is not dominated by a few high-price products. Note that (9) stands in contrast to approaches that focus on households substituting away from varieties whose prices *grow* at a higher rate.[15] A simple

[15] When considering the role of quality substitution, Argente & Lee (2021) focus on the substitution from products with higher

example illustrates the difference: suppose that a household from a certain income group has access to two varieties of a certain product, with the first one initially costing one monetary unit and the second one two monetary units. In period $t-1$, the household primarily purchases the expensive variety. Between period $t-1$ and period t, *both* prices increase by ten percent, and the household reacts by switching in parts from the expensive to the inexpensive variety. This would not be detected in an analysis concentrating on growth rates, since both varieties' prices increase by the same percentage. By contrast, equation (9) would yield a positive value, indicating some product-downgrading.

To explore whether the extent of product-downgrading differs across income groups, we estimate the parameters of the following regression equation:

$$\Delta \hat{P}_{gi,t} = \sum_{i=2}^{I} \beta_i \cdot \text{Income\_Group}_i + \delta_g + \gamma_t + \epsilon_{gi,t} \tag{10}$$

As before, Income_Group_i is a dummy variable with the subscript starting from 2 because income group 1 is the omitted reference group. δ_g and γ_t are product and time fixed effects, respectively. The results of estimating this regression are presented in Table 6. Furthermore, Figure 9 illustrates the estimated coefficients along with their respective 95-percent confidence intervals. The results indicate that the middle-income groups engage significantly more in product-downgrading than both the lowest and highest income groups. More specifically, relative to income group 1, the ability or willingness to adjust the mix of varieties lowers the average product-specific inflation rate by 0.25 percentage points for members of income group 2. Given that we consider month-on-month price changes, this is substantial. To rationalize the size of this difference, we point out that, in the presence of product-downgrading, the effective prices paid for some products may actually *decrease* during an inflationary phase, since the overall price increase is dominated by consumers switching to cheaper varieties (see footnote 1 in the introduction for a simple example illustrating this effect).

5.5 Identifying the role of product-downgrading: The role of the extensive margin

While comparing Laspeyres and Paasche-type inflation at the *gi*-product-group level yields important insights on the extent of product-downgrading across income groups, this approach comes with at least three drawbacks. First, the exclusion of regional and retailer-specific price effects may influence the results. In other words, factors attributed to income group-specific dynamics may actually be influenced by the omitted retailer or area dimensions. For instance, if lower-income households predominantly reside in rural areas, and if price dynamics in rural areas significantly differ from those in urban areas, inflation differentials may be misinterpreted as being driven by the income group rather than the location. Similarly, if lower-income households mainly shop at discount retailers, and if price dynamics in these stores significantly differ from

price growth rates towards products with lower price growth rates.

Table 6: The difference between Laspeyres and Paasche-type inflation at the product level – continued *pi*-pairs

	Dependent variable
	$\Delta\hat{P}_{gi,t}$
Income Group 2	0.0025***
	(0.0004)
Income Group 3	0.0020***
	(0.0004)
Income Group 4	0.0026***
	(0.0005)
Income Group 5	0.0024***
	(0.0005)
Income Group 6	0.0011**
	(0.0005)
Income Group 7	0.0004
	(0.0005)
Observations	195,625
Product Fixed Effects	1,656
Time Fixed Effects	23
Adj. R-Squared	0.003
Standard Errors: Clustered (Product) in Parentheses	
***p<0.01, **p<0.05, *p<0.1	

Note: This table shows the results of regressing the difference between the Laspeyres and Paasche-type price relatives $\Delta\hat{P}_{gi,t}$ on income group dummies, product and time fixed effects. Note that after having computed $\Delta\hat{P}_{gi,t}$, we have excluded the 1st and 99th percentiles to mitigate the effect of outliers.

Figure 9: The difference between Laspeyres and Paasche-type inflation at the product level – continued *pi*-pairs: Estimated coefficients and 95% CI

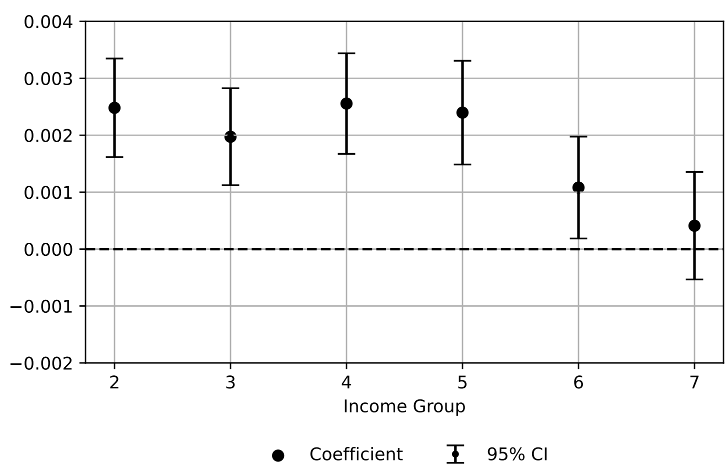

Note: This figure depicts the β_i coefficients and the corresponding 95% confidence intervals when regressing the difference between the Laspeyres and Paasche-type price relatives $\Delta\hat{P}_{gi,t}$ on income group dummies, product and time fixed effects: $\Delta\hat{P}_{gi,t} = \sum_{i=2}^{I} \beta_i \cdot \text{Income\_Group}_i + \delta_g + \gamma_t + \epsilon_{gi,t}$ Note that after having computed $\Delta\hat{P}_{gi,t}$, we have excluded the 1st and 99th percentiles to mitigate the effect of outliers.

other types of retailers, inflation differentials may be incorrectly attributed to the income group rather than retailer type. Second, the inability to identify the role of the extensive margin in driving inflation differentials arises from the reliance on a sub-sample of *continued* varieties – which does not account for all available varieties, although it is large if we aggregate across areas and retailers. Lastly, what may initially appear as a continued variety under the *pi*-definition of the unit of analysis could actually be a newly transacted variety under the *prai*-definition. In other words, when a specific item is purchased in two adjacent periods by households within a particular income group across various retailers and postcode areas, we cannot be sure that it was part of the same income group's consumption basket in the previous period within a specific retailer and postcode area. Therefore, the inclusion of the retailer and postcode area dimensions, despite increasing fluctuations in the purchasing sample over time, enables us to account for location and retailer effects and identify extensive margin effects more clearly. Since it is not possible to compare period-on-period inflation indices based on previous period volume weighting and current period volume weighting for each *prai*-pairing, we propose an alternative strategy to assess inflation differentials across income groups.

We begin by analyzing the *continued prai*-pairings sub-sample, whereby we replicate the computation of volume shares $\mu^{g,cont}_{prai,t}$ and volume-share-weighted average unit prices $\bar{P}^{U,cont}_{gai,t}$, as described in equations (1) and (2), respectively. Following this, we compute the average unit price relative for each *gai*-product-group at time t:

$$\dot{R}^{P}_{gai,t} = \frac{\bar{P}^{U,cont}_{gai,t}}{\bar{P}^{U,cont}_{gai,t-1}} \tag{11}$$

Subsequently, we repeat the regression analysis in which we regress this price relative on income group dummies, product, area, and time fixed effects, akin to the procedure outlined in equation (5):

$$\dot{R}^{P}_{gai,t} = \sum_{i=2}^{I} \beta_i \cdot \text{Income\_Group}_i + \delta_g + \sigma_a + \gamma_t + \epsilon_{gai,t} \tag{12}$$

The results of this regression are presented in Table 7, while Figure 10 illustrates the estimated coefficients along with their respective 95-percent confidence intervals.

The result that there are no inflation differences between income groups when only *continued* varieties are considered (see Figure 10), but that such differences exist when we include all varieties (see Figure 8) suggests that the differences are driven by the discontinued and newly transacted varieties. This indicates that a shift in the *variety-mix* at the *extensive margin* – which includes changes in the item p and/or retailer r – is key to understanding the observed inflation disparities.

To further corroborate this finding, we compute the previous-period volume shares $\mu^{g,discont}_{prai,t-1}$ and volume-share-weighted average unit prices $\bar{P}^{U,discont}_{gai,t-1}$ based on the sub-sample of *discontinued prai*-pairings. Next, we compute the ratio $\tilde{R}^{P}_{gai,t}$, which compares the average unit price of *all* varieties in period t to the average

Table 7: Inflation disparities on the product level – continued *prai*-pairings

	Dependent variable
	$R^{P,cont}_{gai,t}$
Income Group 2	0.0005*
	(0.0003)
Income Group 3	0.0004
	(0.0003)
Income Group 4	0.0004
	(0.0003)
Income Group 5	0.0003
	(0.0003)
Income Group 6	0.0002
	(0.0003)
Income Group 7	0.0001
	(0.0003)
Observations	2,391,455
Product Fixed Effects	1,606
Area Fixed Effects	118
Time Fixed Effects	23
Adj. R-Squared	0.002
Standard Errors: Clustered (Product) in Parentheses	
***p<0.01, **p<0.05, *p<0.1	

Note: We assess whether households with higher incomes tend to experience lower inflation rates compared to those with lower incomes within the continued *prai*-pairings sub-sample. To do so, we regress the unit price relative $R^{P,cont}_{gai,t}$ for continued varieties on income group dummies, product, area and time fixed effects: $R^{P,cont}_{gai,t} = \sum_{i=2}^{I} \beta_i \cdot \text{Income\_Group}_i + \delta_g + \sigma_a + \gamma_t + \epsilon_{gai,t}$. The β_i coefficients – when multiplied by 100 – represent the percentage-point difference in inflation associated with being in the respective income group i compared to the lowest income group (the reference income group). Note that after having created the price relatives, we have excluded the 1st and 99th percentiles to mitigate the effect of outliers.

unit price of *discontinued* varieties in the previous period $t-1$:

$$\tilde{R}^{P}_{gai,t} = \frac{\bar{P}^{U}_{gai,t}}{\bar{P}^{U,discont}_{gai,t-1}} \tag{13}$$

The *more expensive* the discontinued varieties are in period $t-1$, the *lower* is the ratio $\tilde{R}^{P}_{gai,t}$. If we find that $\tilde{R}^{P}_{gai,t}$ significantly differs across income groups (controlling for products, areas, and time), this indicates that income groups differ in the extent to which they drop expensive varieties from their consumption basket.

Similarly, using the *newly transacted prai*-pairings sub-sample, we compute the current period volume shares $\mu^{g,new}_{prai,t}$ and volume-share-weighted average unit prices $\bar{P}^{U,new}_{gai,t}$. Subsequently, we compute the average unit price of new *prai*-pairings in period t relative to all *prai*-pairings in period $t-1$ for each *gai*-product-group:

$$\hat{R}^{P}_{gai,t} = \frac{\bar{P}^{U,new}_{gai,t}}{\bar{P}^{U}_{gai,t-1}} \tag{14}$$

The ratio $\hat{R}^{P}_{gai,t}$ serves as a comparison between the average unit price of newly transacted varieties in period t

Figure 10: Inflation disparities on the product level – continued *prai*-pairings: Estimated coefficients and 95% CI

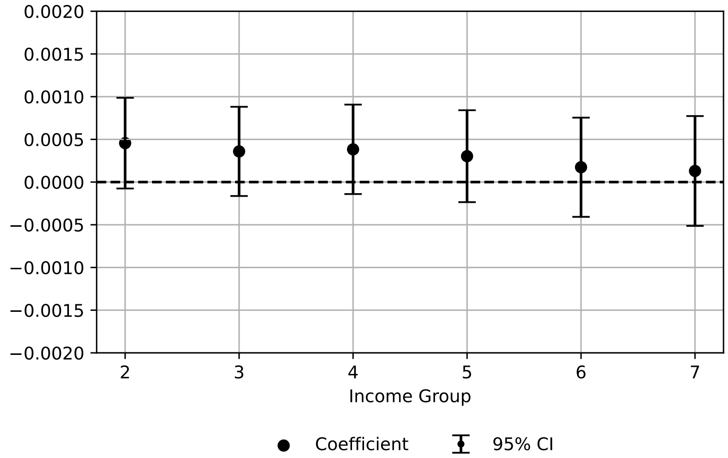

Note: This figure depicts the β_i coefficients and the corresponding 95% confidence intervals when regressing the Paasche-type price relative $R^{P,cont}_{gai,t}$ on income group dummies, product, area and time fixed effects: $R^{P,cont}_{gai,t} = \sum_{i=2}^{I} \beta_i \cdot \text{Income\_Group}_i + \delta_g + \sigma_a + \gamma_t + \epsilon_{gai,t}$. Note that after having created the price relatives, we have excluded the 1st and 99th percentiles to mitigate the effect of outliers.

and the average unit price of all varieties in the previous period $t-1$. The *less expensive* the newly transacted varieties are in period t relative to the average unit price based on all varieties in period $t-1$, the *lower* is the ratio $\hat{R}^{P}_{gai,t}$. Significant differences of this ratio across income groups indicate differences in the extent to which households newly include less expensive varieties in their consumption basket.

To identify possible differences across income groups, we regress the above price relatives on income group dummies, product, area, and time fixed effects:

$$\tilde{R}^{P}_{gai,t} = \sum_{i=2}^{I} \beta_i \cdot \text{Income\_Group}_i + \delta_g + \sigma_a + \gamma_t + \epsilon_{gai,t} \tag{15}$$

and

$$\hat{R}^{P}_{gai,t} = \sum_{i=2}^{I} \beta_i \cdot \text{Income\_Group}_i + \delta_g + \sigma_a + \gamma_t + \epsilon_{gai,t} \tag{16}$$

The regression results are presented in Table 8, while Figure 11 shows the resulting coefficients along with their respective 95-percent confidence intervals. The table and figure demonstrate that adjustments at the extensive margin play a significant role in explaining the inflation disparities across income groups, as depicted in Figure 8. Additionally, the significantly negative coefficients in column 1 of Table 8 suggest that (complete) substitution away from expensive varieties constitutes the primary channel of adjustment that gives rise to differences across income groups. Again, it is members of the *medium-income* groups for whom this effect is most pronounced, i.e. for whom average prices for a product relative to the average price of discontinued varieties is lowest, which indicates that these households are most effectively engaged in product-downgrading. The results in column 2 of Table 8 demonstrate that the inclusion of new (inexpensive) varieties in the product mix also plays a role in generating inflation differences across income groups. However, this channel does not some seem as quantitatively important as the dropping of expensive varieties.

Table 8: Inflation disparities at the product level – discontinued and new *prai*-pairings

	Dependent variable	
	$\tilde{R}^P_{gai,t}$	$\hat{R}^P_{gai,t}$
Income Group 2	-0.0166***	-0.0054**
	(0.0015)	(0.0024)
Income Group 3	-0.0196***	-0.0047*
	(0.0017)	(0.0027)
Income Group 4	-0.0186***	-0.0062**
	(0.0017)	(0.0027)
Income Group 5	-0.0135***	-0.0030
	(0.0015)	(0.0025)
Income Group 6	-0.0054***	-0.0054**
	(0.0014)	(0.0023)
Income Group 7	-0.0016	-0.0071***
	(0.0014)	(0.0022)
Observations	3,329,896	3,308,937
Product Fixed Effects	1,470	1,468
Area Fixed Effects	118	118
Time Fixed Effects	23	23
Adj. R-Squared	0.007	0.019
Standard Errors: Clustered (Product) in Parentheses		
***p<0.01, **p<0.05, *p<0.1		

Note: This table shows the results of regressing the unit price relatives $\tilde{R}^P_{gai,t}$ and $\hat{R}^P_{gai,t}$ on income group dummies, product, area and time fixed effects. Note that after having created the price relatives, we have excluded the 1st and 99th percentiles to mitigate the effect of outliers.

6 Summary and conclusions

In this paper, we have focused on the role of *product-downgrading* as one particular mechanism, due to which different income groups may experience different effective inflation rates, and which may thus contribute to the distributional effects of inflation: households who are able and willing to replace more expensive varieties of a given product by less expensive alternatives cushion the impact of the overall price increase. Conversely, households who are unable or unwilling to do so – either because their initial consumption basket already consists of the cheapest varieties, or because their income allows them to stick to the previous spending pattern – experience higher effective inflation rates. The granular data set of household-level purchases of fast-moving consumer goods allows tracing the level and evolution of volume-share-weighted price averages for a large number of products, and thus to precisely identify the extent of product-downgrading for different income groups.

Figure 11: Inflation disparities at the product level – discontinued and new *prai*-pairings: Estimated coefficients and 95% CI

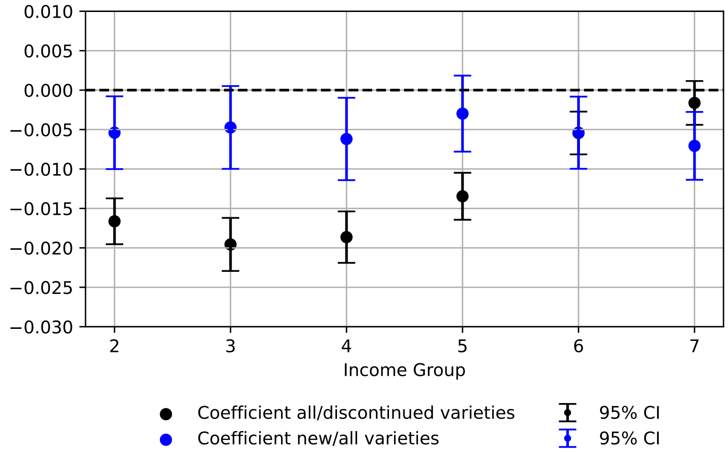

Note: This figure depicts the β_i coefficients and the corresponding 95% confidence intervals when regressing the unit price relatives $\tilde{R}^P_{gai,t}$ and $\hat{R}^P_{gai,t}$ on income group dummies, product, area and time fixed effects. Note that after having created the price relatives, we have excluded the 1st and 99th percentiles to mitigate the effect of outliers.

Focusing on the evolution of prices in the United Kingdom shortly before and after the depreciation of the British pound in the wake of the Brexit referendum, we have come up with the following results: first and not surprisingly, the average unit price paid for a given product increases in households' income. Second, and more importantly, the relationship between product-specific inflation rates and income levels exhibits a U-shaped pattern: while medium-income households experience significantly lower price increases for the product-mix they actually purchase than low-income households, there is no significant difference between high-income and low-income households.

To demonstrate the relevance of product-downgrading for different effective inflation rates across income groups, we computed the differences between Laspeyres and Paasche-type price relatives, thus juxtaposing the price evolution without an adjustment of the variety-mix with an evolution that allows for substitution. The results of this analysis suggest that the scope for product downgrading allowed middle-income households to reduce the average price increase of products by up to 0.25 percentage points, relative to members of the lowest income group. We also demonstrated that income groups differ in the extent to which the average price of discontinued varieties relates to the average price of all varieties consumed, and to which the average price of newly transacted varieties differs from the average price of all varieties consumed in the previous period. Our findings suggest that it is predominantly medium-income households' ability and willingness to completely remove expensive varieties from their product mix that gives rise to the observed pattern of product-downgrading.

Hence, from this perspective, inflation is both anti-poor *and* anti-rich – but, of course, for different reasons: while the scope for product-downgrading is limited for poor households, who already consume the lowest-priced varieties to start with, rich households' reluctance to switch to less expensive varieties may reflect their ability to afford the price increase. In contrast to households at the lower and upper ends of the British income distribution, those in the middle seem to take advantage of the possibility to engage in product-downgrading and to thus cushion the effect of the overall price increase. We are, of course, aware, that the observation of lower effective inflation rates for these groups does not exhaust all the welfare effects of the "Brexit inflation": any substitution – be it across or within products – potentially reduces the utility of consumers who replace components of their favored consumption bundle by possibly *inferior* items. However, we point out that the varieties we consider are almost perfect substitutes in terms of nutritional value, cleaning performance etc. – i.e. the welfare effects of substituting varieties within those products are likely to be rather small, and essentially reflect the *brand value* of more expensive varieties.[16]

Finally, we emphasize that our paper focused on *one* aspect that may potentially contribute to heterogeneous effects of inflation – namely, substitution within narrowly defined products. We are aware that these effects are augmented by other sources of heterogeneity – most importantly, different *expenditure shares* and different elasticities of substitution *across* products. However, we believe that the results we have presented add an important insight on the distributional effects of inflation, which – for lack of appropriate data – had to be neglected so far.

[16] This is also the reason why we preferred to use the term *product-downgrading* instead of *quality adjustment*, referred to by, e.g., Argente & Lee (2021).

Bibliography

Ampudia, Miguel & Michael Ehrmann & Georg Strasser (2023). *The Effect of Monetary Policy on Inflation Heterogeneity Along the Income Distribution*. Working Paper 1124. Bank for International Settlements.

Argente, David & Munseob Lee (2021). “Cost of Living Inequality During the Great Recession”. In: *Journal of the European Economic Association* 19.2, 913–952.

Braun, Rahel & Sarah M Lein (2020). *Heterogeneity in Inflation and Preferences across Households*. Working Paper. Mimeo.

Breinlich, Holger & Elsa Leromain & Dennis Novy & Thomas Sampson (2022). “The Brexit Vote, Inflation and U.K. Living Standards”. In: *International Economic Review* 63.1, 63–93.

Broadbent, Ben (2017). *Brexit and the Pound*. Speech Given at Imperial College, London. Bank of England.

Broda, Christian & John Romalis (2009). *The Welfare Implications of Rising Price Dispersion*. Working Paper. University of Chicago.

Burstein, Ariel & Gita Gopinath (2014). “International Prices and Exchange Rates”. In: *Handbook of International Economics*. Elsevier, 391–451.

Cage, Robert & Thesia Garner (2002). *Constructing Household Specific Consumer Price Indexes: An Analysis of Different Techniques and Methods*. Working Paper 354. Bureau of Labor Statistics.

Cavallo, Alberto & Brent Neiman & Roberto Rigobon (2014). “Currency Unions, Product Introductions, and the Real Exchange Rate”. In: *Quarterly Journal of Economics* 129.2, 529–596.

Coibion, Olivier & Yuriy Gorodnichenko & Gee Hee Hong (2015). “The Cyclicality of Sales, Regular and Effective Prices: Business Cycle and Policy Implications”. In: *American Economic Review* 105.3, 993–1029.

Colicev, Anatoli & Joris Hoste & Jozef Konings (2022). *The Impact of a Large Depreciation on the Cost of Living of Rich and Poor Consumers*. Discussion Paper 16869. Centre for Economic Policy Research.

Corsetti, Giancarlo & Meredith A. Crowley & Lu Han & Huasheng Song (2023). *Markets and Markups: A New Empirical Framework and Evidence on Exporters from China*. Working Paper 2023/27. Refugee Studies Centre.

Cravino, Javier & Ting Lan & Andrei Levchenko (2020). “Price Stickiness along the Income Distribution and the Effects of Monetary Policy”. In: *Journal of Monetary Economics* 110, 19–32.

Cravino, Javier & Andrei A. Levchenko (2017). “The Distributional Consequences of Large Devaluations”. In: *American Economic Review* 107.11, 3477–3509.

Crawford, Ian & Zoe Oldfield (2002). *Distributional Aspects of Inflation*. Working Paper. Institute for Fiscal Studies.

Dhingra, Swati & Thomas Sampson (2022). “Expecting Brexit”. In: *Annual Review of Economics* 14.1, 495–519.

Fajgelbaum, Pablo D. & Amit K. Khandelwal (2016). “Measuring the Unequal Gains from Trade”. In: *Quarterly Journal of Economics* 131.3, 1113–1180.

Garner, Thesia I. & Javier Ruiz-Castillo & Mercedes Sastre (2003). “The Influence of Demographics and Household-Specific Price Indices on Consumption-Based Inequality and Welfare: A Comparison of Spain and the United States”. In: *Southern Economic Journal* 70.1, 22–48.

Gerstein, Neil & Bart Hobijn & Fernanda Nechio & Adam Shapiro, et al. (2019). *The Brexit Price Spike*. Economic Letter. Federal Reserve Bank of San Francisco.

Goetz, Daniel & Alexander Rodnyansky (2023). “Exchange Rate Shocks and Quality Adjustments”. In: *Review of Economics and Statistics* 105.1, 86–100.

Gürer, Eren & Alfons Weichenrieder (2020). “Pro-rich Inflation in Europe: Implications for the Measurement of Inequality”. In: *German Economic Review* 21.

Ha, Jongrim & Anna Ivanova & Franziska Lieselotte Ohnsorge & Derya Filiz Unsal Portillo Ocando (2019). *Inflation: Concepts, Evolution, and Correlates*. The World Bank.

Hagemann, Robert P. (1982). “The Variability of Inflation Rates Across Household Types”. In: *Journal of Money, Credit and Banking* 14.4, 494–510.

Hobijn, Bart & David Lagakos (2005). “Inflation Inequality in the United States”. In: *Review of Income and Wealth* 51.4, 581–606.

Jaravel, Xavier (2018). “The Unequal Gains from Product Innovations: Evidence from the U.S. Retail Sector”. In: *Quarterly Journal of Economics* 134.2, 715–783.

— (2021). “Inflation Inequality: Measurement, Causes, and Policy Implications”. In: *Annual Review of Economics* 13.1, 599–629.

Kaplan, Greg & Sam Schulhofer-Wohl (2017). “Inflation at the Household Level”. In: *Journal of Monetary Economics* 91.C, 19–38.

Kiss, Regina & Georg Strasser (2024). *Inflation Heterogeneity across Households*. Working Paper 2898. European Central Bank.

Michael, Robert T. (1979). “Variation Across Households in the Rate of Inflation”. In: *Journal of Money, Credit and Banking* 11.1, 32–46.

Nakamura, Emi & Jón Steinsson (2012). “Lost in Transit: Product Replacement Bias and Pricing to Market”. In: *American Economic Review* 102.7, 3277–3316.

Office for National Statistics (2023). *Distribution of UK Household Disposable Income, Financial Year Ending 2022*. Statistical bulletin. URL: `https://www.ons.gov.uk/peoplepopulationandcommunity/personalandhouseholdfinances/incomeandwealth`.

Orhun, A. Yesim & Mike Palazzolo (2019). “Frugality Is Hard to Afford”. In: *Journal of Marketing Research* 56.1, 1–17.

Porto, Guido G. (2006). “Using Survey Data to Assess the Distributional Effects of Trade Policy”. In: *Journal of International Economics* 70.1, 140–160.

Strasser, Georg & Teresa Messner & Fabio Rumler & Miguel Ampudia (2023). *Inflation Heterogeneity at the Household Level*. Occasional Paper Series Working Paper 325. European Central Bank.