Inflation risk and heterogeneous trading down

Mar Domenech Palacios *

Preliminary - please do not circulate without permission

See latest version here

December 13, 2023

Abstract

I examine the importance of the quality margin as an insurance mechanism against aggregate shocks, and more specifically the heterogeneous behaviour across household's income distribution. Using household scanner data for Germany, I analyse the extent to which households trade down in the quality goods. First, I document that, on average, lower income households tend to purchase lower quality goods. Furthermore, in the aftermath of an aggregate shock, lower income households exhibit a low propensity to trade down, presumably due to a limited capacity to do so. This is in contrast to the rest of households, who appear to trade further down in the quality of goods. To understand the general equilibrium implications of this shift in aggregate demand towards lower quality goods, I employ a shift-share research design. I find that an aggregate demand shift toward lower quality goods during a recession leads to relatively higher prices of low quality goods compared to the price of higher quality varieties.

JEL Classification :

Keywords: inequality, inflation, uncertainty, quality, substitution, trading down, risk

*University of Cambridge

I am indebted to Elisa Faraglia, Giancarlo Corsetti and Luca Dedola for their patient guidance and support. I thank Vasco Carvalho, Sarur Chaudhary, Meredith Crowley, Davide Debortoli, Fatih Guvenen, Andrew Hannon, Simon Mongey, Galo Nuño, Alba Patozi, Ricardo Reis, Jirka Slacalek, and Macro Workshop seminar participants for helpful discussions and comments.

1 Introduction

Low-income households tend to experience higher levels of inflation compared to their higher-income counterparts. This has been extensively documented in the literature (see, for example, Kaplan & Schulhofer-Wohl (2017)). The reasons behind this phenomenon have been thoroughly explored. One main reason is that a larger portion of the consumption of low-income households is often devoted to essential needs such as energy, which are more prone to price fluctuations. Other researchers (see Jaravel (2019)) shed light on systematic differences in innovation patterns as a contributing factor. Figure 1 presents average inflation rates for each income decile in Germany during the 2005 to 2018 period and shows how lower income households tend to experience, on average, higher levels of inflation.

Figure 1: Average inflation by income decile

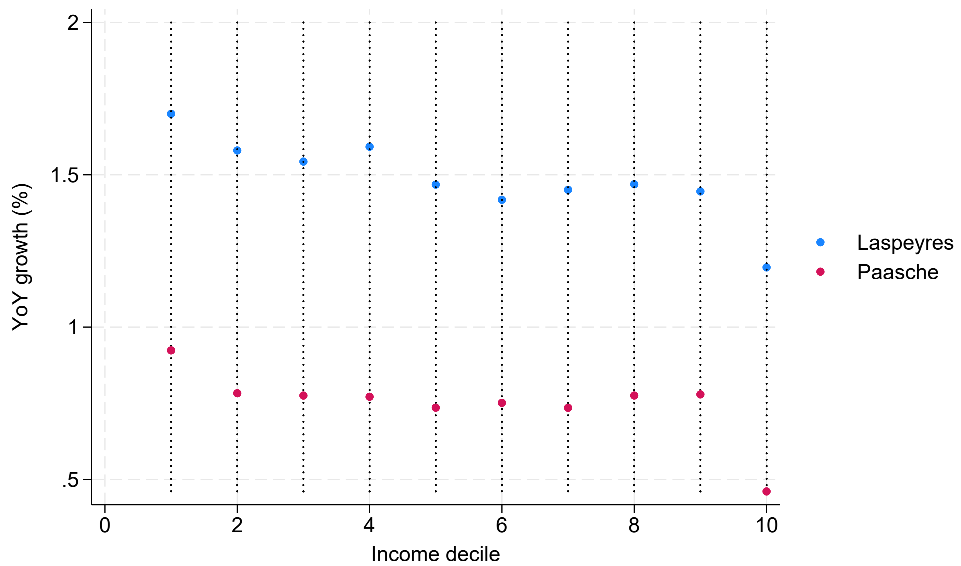

Notes: Inflation is computed at the household level and then averaged by groups across all time periods. The data covers German households, spans from 2005 to 2018, and covers supermarket goods.

This paper aims at focusing on the role of the business cycle in driving these differences. Figure 2 depicts how the median level, percentile 90 and percentile 10 of inflation vary across income deciles and between the great financial crisis period and the rest of the sample. In it, I focus on the first year of the financial crisis: from 2007 to mid 2008. First, the Figure shows a generalised increase in inflation at the onset of the crisis and for all household groups. This is inline with official CPI data on food inflation. However, the increase appears to be larger for lower-income households. More importantly, right-tail risk, that is, the risk of high inflation, appears to change substantially more for the lower income households compared to the rich counterparts. Therefore, inflation risk appears to increase more for lower income households at the onset of a recession. As Figures A.1 and A.2 in the Appendix show, this gap decreases as one incorporates more period of time in the period studied.

Figure 2: Inflation risk by income decile and over the business cycle

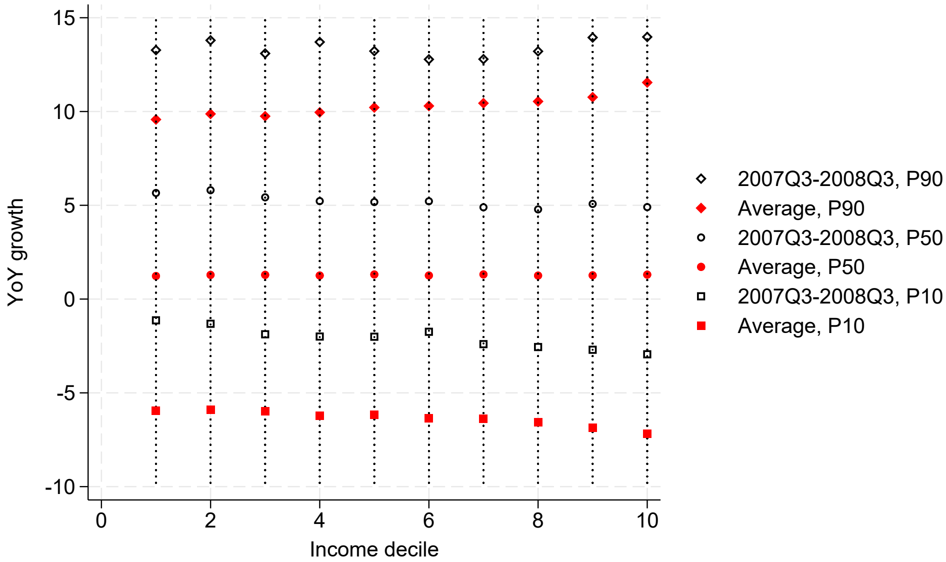

Notes: Laspeyres Inflation is computed at the household level. Income class referes to the income decile of a given household within a given state. The data covers German households and spans from 2005 to 2018. GFC refers to the begining of the financial crisis and includes the second semester of 2007 and first semester of 2008.

Another phenomenon documented in the literature and that is key in understanding household behavior during economic recessions is the tendency of households to lower the quality of goods they purchase, as highlighted by Jaimovich et al. (2019). The paper argues that this adjustment in purchasing behavior amplifies the magnitude of the recession because lower quality goods tend to be less labour intensive and, therefore, the aggregate demand for labour decreases systematically during these periods.

In this paper, I study whether differentials in inflation risk over the cycle and trading down in the quality of goods might be related phenomenons. In particular, I aim at delving deeper into how households tend to trade down in the quality of goods and, more specifically, explore whether this phenomenon occurs heterogeneously across households. Additionally, I analyse whether this effect has implications for price dynamics over the business cycle and across the quality distribution of varieties within products. The underlying hypothesis is that the degree of trading down is heterogeneous across households, and more specifically that low-income households may lack the capacity to engage in this margin of adjustment. Moreover, when the rest of households do trade down, the aggregate demand shift toward lower quality goods leads to an increase in the relative price of low-quality goods compared to higher quality goods. Figure 3 shows the relative price of low versus medium-quality varieties for each product and over time.[1] During the great recession, the average price of lower quality goods relative to that of high quality goods increased. See Figure A.3 the observed graph for selected individual products.

Figure 3: Relative prices of low vs high quality goods

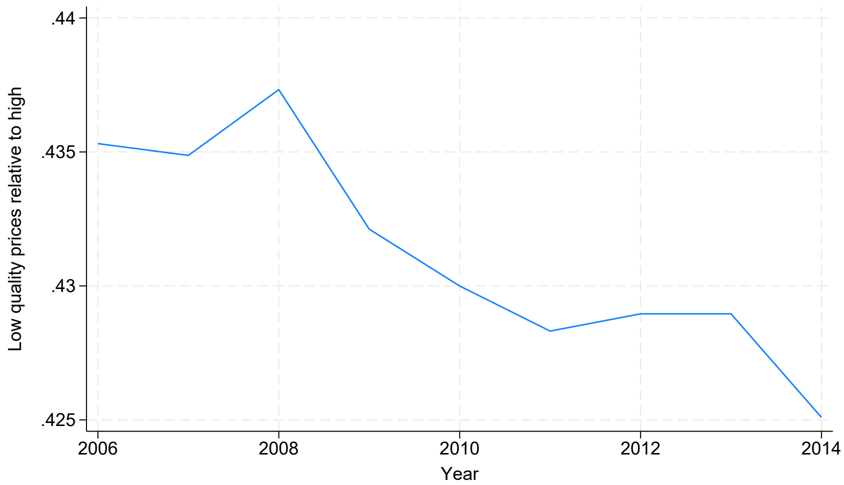

Note: The relative price of low quality goods is defined as the 25th percentile of a product category divided by the median price of that product category within each year. Relative prices are averaged across product categories. Product categories are defined by the COICOP-5 classification.

The aim of this study is therefore to examine the importance of the quality margin as an insurance mechanism against aggregate shocks and, more specifically, to investigate the general equilibrium effect in relative prices of an aggregate demand shift toward lower quality goods in the aftermath of an aggregate shock such as the great recession. To achieve this, I use household scanner data from supermarket expenditure by German households. The data contains a representative sample of German households from 2005 to 2018.

I first document the tendency of lower-income households to purchase lower quality goods on average. Additionally, I analyse how households use the quality margin to decrease their overall expenditures in the aftermath of an aggregate shock, depending on their income group. I find that lower-income households exhibit a limited capacity to engage in trading down given the fact that they are at a lower bound, in contrast to

[1] For each product, a low quality variety is defined as those below the 25th percentile price. Medium quality variety is defined as those whose price is between the 25th and the 75th percentile and those with a price higher than the 75th percentile are defined as high quality varieties.

the rest of households who are likely to trade down further in the quality of goods. To understand the general equilibrium implications of this shift in aggregate demand towards lower quality goods, I employ a shift-share research design based on population growth of narrowly defined groups of households to predict the amount of trading down down when the recession hits, therefore identifying a reasonably exogenous demand shifter. The intuition is that in the regions where the household groups that are more likely to trade down in the quality of goods grow faster, the amount of trading down once the recession happens will be larger. I find that a generalised demand shift toward lower quality goods during recessions leads to an increased price of low quality goods relative to the rest.

By shedding light on the heterogeneous trading down behavior and its effects on household inflation risk, in this paper I aim at contributing to our understanding of the complex dynamics between income distribution, consumption patterns, and inflation risk over the business cycle. This might have important implications for welfare analysis given that, one the one hand, the fraction of households that find themselves in this lower bound could be indicative of the welfare cost of decreased aggregate consumption, given a larger utility cost of decreasing volumes of consumption rather than the quality of the goods, holding volumes constant. Moreover, when households can trade down in the quality of goods, this insures them against shocks by lowering the quality of the purchased goods.

The remaining of the paper is organised as follows. Section 2 reviews the existing related literature, Section 3 presents the dataset, Section 4 compares households habits along the income distribution; Section 5 studies differences in trading down over time and along the income distribution; Section 6 presents a shift share research design to identify exogenous demand shifts; and Section 7 concludes.

2 Related Literature

Previous literature has documented that during economic downturns, households tend to reduce the quality of the goods they purchase. Jaimovich et al. (2019) show that a generalised shift towards lower quality goods has effects, for example, in the labour market, because lower quality goods tend to be less labour intensive. This margin of adjustment potentially insures households against individual or aggregate shocks because it allows to decrease total expenditure without affecting volumes of consumption. On the study of the relevance of the quality margin, Rodnyansky et al. (2022) propose a New Keynesian model with endogenous adjustment in product quality that nests the canonical framework and show that it amplifies the economy's response to productivity shocks, leading to less reactionary monetary policy. On the study of how households react heterogeneously during recessions, Carvalho et al. (2021) consider billions of transactions from card data from BBVA as a source of information for measuring consumption and find strong consumption responses to business closures but a steeper decline in spending in rich neighbourhoods. Nord (2022) formalises a model of frictional product where heterogeneous consumption baskets along the income distribution and higher shopping effort of the poor imply that retailers face different price elasticities and face higher markups for goods consumed by richer households. In the case of durable goods, Gavazza & Lanteri (2021) study a general-equilibrium model of durable consumption and find that, after a tightening of the borrowing limit, debt-constrained households postpone the decision to scrap and upgrade their low-quality cars, which depresses mid-quality car prices. In turn, this effect reduces wealthy households' incentives to replace their mid-quality cars with high-quality ones, thereby decreasing new-car sales. In the same line, Bertolotti et al. (2021) explore from an empirical perspective expenditures on cars during the great recession.

In this paper I examine whether households trade down in the quality of goods in a systematically heterogeneous manner. Specifically, I focus on investigating whether a lower bound exists that prevents certain households from having this margin of adjustment, that is, whether a subset of households, who were already purchasing lower quality goods prior to an aggregate shock, are restricted in the ability to use this margin of adjustment.

A different strand of research increasingly focuses on quantifying and understanding the reasons behind inflation heterogeneity at the household level. More specifically, some authors have attempted to explain why low-income households tend to experience higher levels of inflation. For example, Jaravel (2019) uses scanner data from the retail sector in the US to find that annual inflation for retail products was substantially higher for the bottom income quintile relative to the top income quintile. He investigates the hypothesis that this is due to the fact that firms introduced more products to high-income households due to an increased demand by these (explained by growth and rising inequality), and as a result, the prices of continuing products in these market segments fell due to increased competitive pressure. Kaplan & Schulhofer-Wohl (2017), using scanner data from the United States, compute household level inflation rates and suggest that almost all variability in a household's inflation rate comes from variability in household-level prices relative to average prices. Argente & Lee (2021) construct income-specific price indexes for the period 2004 to 2016 and find that product quality substitution and changes in the shopping behaviour explain around 40% of the gap. A different strand of the literature has focused on the distributional effects of inflation. A few examples include Cardoso et al. (2022), who use bank level data to quantify the three key channels that shape how inflation affects wealth inequality; Orchard (2022), studying how income level inflation rates vary over the course of the business cycle, and documenting that during recessions prices rise more for necessities, and that the aggregate share of spending devoted to necessities is counter-cyclical; Yang (2022) analyses the heterogeneous costs of inflation constructing a HANK model and finds that a utilitarian central bank should adopt an assymmetric monetary policy rule that is accommodative towards inflation and aggressive towards deflation; and Boel et al. (2021) focus on the redistributive effects of expected inflation with heterogeneous discount factors and collateral constraints and find that in their framework inflation is detrimental to capital accumulation and affects borrowing and lending when collateral constraints are present, which they assess to be regressive for the US. Cravino et al. (2020) establish a new mechanism through which monetary policy shocks have distributional consequences: prices of goods consumed by high-income households are more sticky and less volatile than those of the goods consumed by middle-income households (hump-shaped form). They build a New-Keynesian model with Calvo-Style nominal rigidities where sectors are heterogeneous with respect to price stickiness and households are heterogeneous with respect to income levels and consumption baskets. Relatedly, Lauper et al. (2021) studies the effect of monetary policy shocks into household inflation dispersion, and find that contractionary monetary policy significantly and persistently decreases inflation dispersion in the economy, and that middle-income households experience higher inflation rates that are more reactive to a contractionary monetary policy shocks.

On consumption heterogeneity over the business cycle, Michelacci et al. (2022) show that, in response to income shocks, households persistently change their consumption basket by buying varieties never purchased before that were existing. They also find that households search for the varieties they like to consume and when income increases they add more products in their basket. These findings have implications for the macroeconomic effects of fiscal transfers and for the measurement of household-level inflation. Aguiar & Hurst (2005) find that the retired spend less, but obtain the same caloric intake and so although expenditure decreases after retirement, it does not translate into a decreased consumption. Kaplan & Menzio (2015) find that households with fewer employed members pay lower prices and do so by visiting a larger number of stores instead of by shopping more frequently. Michelacci et al. (2022) follow up on this finding and find that about half of the change in U.S. non-durable consumption expenditure is due to changes in the products entering households' consumption basket, that is, the extensive margin.

Finally, literature focuses on depicting the effects of large currency depreciations. Burstein et al. (2005) analyze how a large devaluation leads to large drops in real exchange rates through a slow adjustment in the prices of non-tradable goods and services. With a focus on distributional effects, Colicev et al. (2022) use retailer scanner data to analyse the distributional impact of a large and sudden exchange rate shock on the cost-of-living of consumers. They focus on the 2015 large depreciation of the Kazakh Tenge, after Kazakhstan switched from a fixed to a floating exchange rate regime. they show that marginal costs increase more for foreign varieties than for local varieties. However, retail margins on foreign varieties fall relative to local varieties. Hence, the retailer limits the transmission of the cost shock into consumer prices by adjusting retail margins on foreign and local varieties differently.

In this paper, I aim at studying the effects of heterogeneous trading down onto the relative prices of goods.

3 Data

I use GfK scanner data, a household panel that covers a representative sample of German households for the period from 2005 to 2018. The dataset consists of around 30,000 households per year (between 25,398 and 38,457 depending on the year). A number of household characteristics are observed: age, number of people in the household, income, social class, zipcode and province.

The dataset covers supermarket-purchased goods, that is, mainly nondurables, and it is mainly composed of foods and beverages. It includes 200,000 barcodes of purchased varieties. I classify purchases into products and varieties. A product is characterised by its 2018 5-digit COICOP classification, grouped by narrowly-defined subset classes of a product. A variety is characterised by its barcode and unit of measure, such that a product consists of multiple varieties.

I classify purchased items as product and varieties. A product is characterized by its 5 digit 2018 COICOP (Classification of Individual Consumption According to Purpose)[2]. There are 338 subclasses of products; 186 classes; 63 groups and 15 divisions. A variety in a product is characterised by its barcode.

Households are grouped by their income levels. To do so, I follow three different classifications. First, simply classifying households into their income decile at the country level. Second, by classifying households by their relative income level within their state of residence. The aim is to avoid sorting households geographically. Third, I classify

[2]It is the international reference classification of household expenditure, and is an integral part of the System of National Accounts (SNA). Is is used for household expenditure statistics based on household budget surveys and for consumer price indices. Available here

households according to the social class variable available in the dataset. This variable, however, only has six groups. There are 17 income class groups for Germany. For each household I assume that their income is the centre of the interval provided; for the lowest I assume an income equal to its upper bound and for the highest I assume an income equal to its lower bound.

I build a Laspeyres and Paasche indices of inflation at the year on year level and quarterly frequency at the household level. I construct consumption per household in real terms dividing overall quarterly household level expenditure by household-level price index constructed with Laspeyres year-on-year inflation. To adjust for within-household economies of scale, I follow the modified OECD scale, according to which 1 point is given to the first member, 0.5 points to the rest of adults and 0.3 points for each kid below 14 years old. However, because I do not observe the age of the rest of the members, I assign 1 point to the first one and 0.5 to the rest of members. I then measure consumption growth at the household level as log differences.

4 Decomposition

I first decompose total expenditures for each household with the aim of understanding household habits, how they evolve over time and, more specifically, how they depend on the state of the economy. For it, I modify Nord (2022) decomposition. In this paper, the author decomposes quarterly household expenditures into three different components: the direct effect of shopping behaviour (effort); the differences in substitution among similar goods; and a counterfactual expenditure that measures expenditure if all households purchased same varieties of goods and at the same prices. I further decompose the second term, into temporary differences in the price of products, that is, temporary discounts (temporay substitution) and permanent differences in the price of different varieties, assumed to summarise quality differences between them.

$$e_{i,t} = \sum_k \sum_{j \in J_k} p_{jkit} c_{jkit}$$

$$= \textstyle\sum_k \sum_{j \in J_k} \left[\underbrace{(p_{jkit} - \bar{p}_{jkt}) c_{jkit}}_{Effort} + \underbrace{((\bar{p}_{jkt} - \bar{p}_{jk}) - (\tilde{p}_{kt} - \tilde{p}_k)) c_{jkit}}_{Temp.substitution} + \underbrace{(\bar{p}_{jk} - \tilde{p}_k) c_{jkit}}_{Quality} + \underbrace{\tilde{p}_{kt} c_{jkit}}_{Counterfactual} \right] \quad (1)$$

For household i at time t, k refers to the specific product and j to the barcode; $\bar{p}$ refers to the average price of a barcode at time t in a given state and $\tilde{p}$ is the average price of a product in a given state.

The first term is the difference between what the individual household pays for the same variety relative to other households and therefore can be thought of as a measure of the effort or time invested in shopping, that is, search costs. The second term reflects the extent to which a household takes advantage of temporary discounts of products. The third term can be seen as the substitution between similar varieties within a given product. A larger term indicates that the household is purchasing more expensive varieties of a given product, that is, higher quality goods. How it evolves over time for each household would shed light on potential non-homothetic utility functions. Finally, the last term is the counterfactual expenditure, and indicates how much a given household would spend if they purchase the average-quality variety within each product at the average price of the product.

Income data is available yearly for each household.[3] To test the robustness of the results, I define income groups in three different ways: First, by household income decile at the country level. Second, by income decile at the lander level. The idea with this measure if to avoid sorting households geographically.[4] I smooth household income as the average of the current, the last 4 quarters and the following 4 quarters income and, if not so many observations are available, the maximum amount of observations observed. This avoids sorting certain households (for example, those that become temporarily unemployed) into a given specific group. the third measure uses the variable Social class, available in the dataset, and that depends on the level of education and the profession of the household. There are six social class groups.

4.1 Results

4.1.1 Spending decomposition

I first perform the decomposition at the household level, group households according to their income levels, and investigate the average contribution of each term in either increasing or decreasing the overall spending levels. Figures 4, A.4 and A.5 show the magnitudes of each term, relative to overall spending levels, for each income group. In the first Figure, income decile is classified at the country level; in the second, at the within-state level; and the third one is based on social class.

[3]While one cannot observe exact household income, there are 17 income bins, ranging from below 500 euros to above 5000 euros household income per month and with a range of 250 euros per income bin.

[4]Moreover, this accounts for the fact that there might be strong regional differences in goods available and the supermarket brands, and that households do not typically travel to purchase goods and so the relevant income position within a region would be the relevant measure.

Figure 4: Decomposition by income group

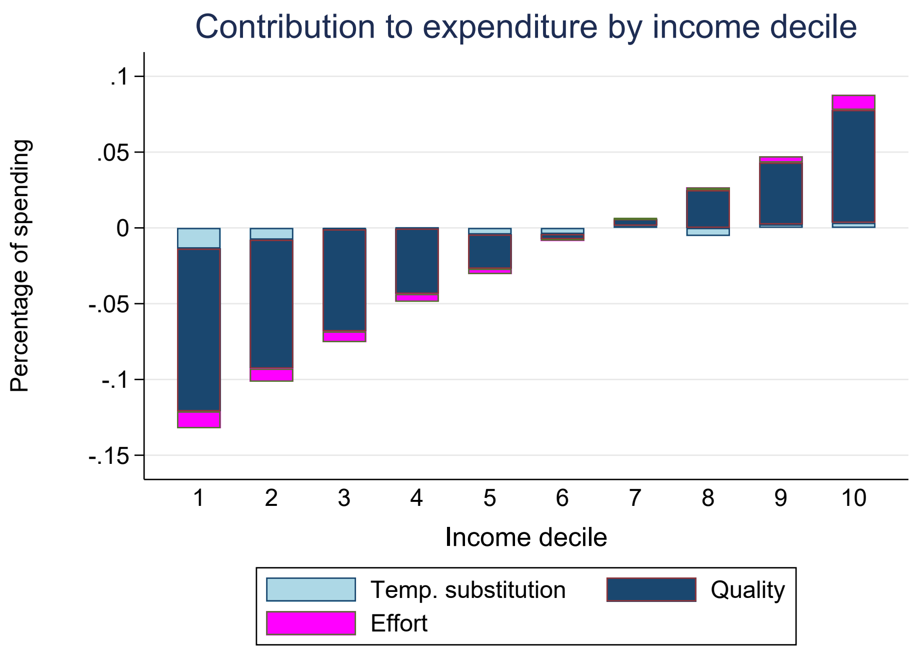

A number of observations are worth mentioning. First, search effort, temporary substitution and quality choice decrease the overall level of spending in lower income households, given a basket of consumption products. This difference monotonically increases with income, such that for income Decile 7 and above, these choices increase the overall spending. This is true for all classification of households into groups. For the lowest income group, the three component together decrease, on average, around 13% of the overall expenditures of the household and, for the highest income decile, they increase by around 8% overall expenditures for a given consumption basket. The second important observation is that this variation is mainly driven by the quality margin, with it being of an order of magnitude larger than the other two components.

The fact that these margins allow households to increase or decrease their overall spending create a wedge between household spending and the counterfactual term in Equation 4.

Figures 5, A.6 and A.7 show the overall level of spending by household group in Euros and the average counterfactual level by group. First, as one could expect, consumption or expenditure levels are larger the higher the income level of households. The same happens with the counterfactual term but, in contrast with the previous case, the curve is steeper for the latter. That means that low income households would be spending more euros if they were purchasing the average product at the average price and, conversely, high income households spend more because of their relatively lower effort, the fact that they take less advantage of temporary discounts and their choice of quality for each product. This implies that consumption inequality is lower than one would observe by directly comparing expenditures.

Figure 5: Expenditures and counterfactual by income group

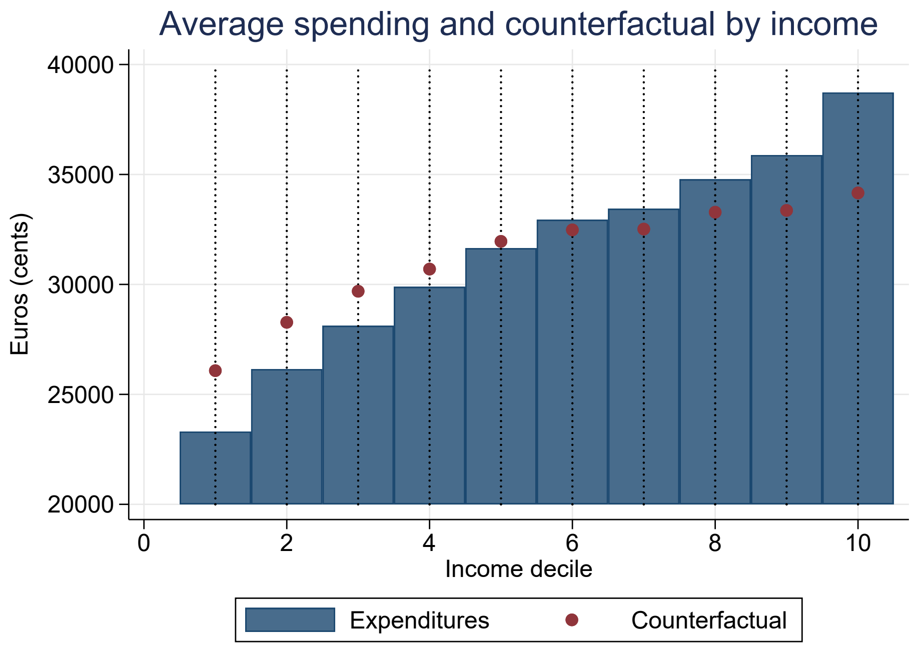

5 Heterogeneous trading down

I next focus on understanding how this household behaviour evolves over the business cycle. In particular, I aim at understanding how and whether households modify the quality of the products they buy over the business cycle as a margin of adjustment and, more specifically, whether this happens heterogeneously across households. For this, I focus on the quality component of the previous decomposition, normalised by household spending, and analyse the evolution over time. I use the following specification:

$$Quality^y_{i,t} = \beta^y \times recession^y_{i,t} + \alpha^y_i + \gamma^y_t + \epsilon^y_{i,t} \tag{2}$$

where y is income group; α_i are household FE and γ_t are time FE, $Quality^y_{i,t}$ is the quality term divided by overall expenditures for household i at time t and $recession_{i,t}$ is a dummy variable equal to 1 if the household resides in a region with negative GDP growth for at least two consecutive quarters.

To focus the analysis on the period of the great financial crisis, I add an interaction term for the great financial crisis and regional crisis, with the aim of exploiting variation between regions in entering and leaving the crisis:

$$Quality^y_{i,t} = \beta^y \times recession^y_{i,t} \times GFC_t + \alpha^y_i + \gamma^y_t + \epsilon^y_{i,t} \tag{3}$$

where GFC_t is a dummy variable equal to 1 for the second half of 2007, all 2008 and 2009.

Table 1 presents the baseline results for the first specification. First, column 1 depicts the result when all households are included and therefore, presents an average at the aggregate level. In particular, during a recession, trading down in the quality of the varieties that households purchase appear to decrease expenditures by an additional 0.2 percentage points on average.

The rest of the columns present the results for a given income group. To improve the granularity of the analysis, I further decompose the lowest income quantile into the first and second deciles. All three tables confirm a differential behaviour of households along the income distribution. While middle and high income households adjust their purchasing behaviour when hit by a recession, this is not the case for low income households. Highest income households also do not exert this additional trading down. Specifically, households in deciles 2 and above of the income distribution appear to exert a higher pressure into lowering expenditures when hit by a recession, whilst those in the lower end of the distribution do not appear to change their behaviour over the business cycle. As exposed in the previous section, these buy relatively lower quality goods on average and, potentially, do not have access to this margin of adjustment when they are hit by a recessionary shock.

Table 2 displays the results when focusing on the Great financial crisis specifically. The results are qualitatively similar to the previous ones, although now all income groups trade further down in the quality of the goods they purchase except for the first decile, compared to the first quantile seen previously. Therefore, a smaller portion of the population seems to be constrained. Moreover, the magnitudes are also larger than those obtained previously: during the recession, trading down in the quality of the varieties that households purchase appear to decrease expenditures by an additional 0.7 percentage points on average.

Tables A.1, A.2, A.3 and A.4 redo the analysis with a different classification of households into groups: first, I classify households by income groups within a state and second, according to their social class. In both cases, the same results prevail and there is evidence of no further trading down for the lower groups. One difference with the previous results is the highest income group when the classification is performed within a given state, where I find that in this case they do trade further down as opposed with the previous result. One explanation might be related to the fact that the classification in 1, at the country level, groups the households with the highest income in the country, whereas in A.1 the highest income group within each state might not group the richest households of the country and therefore might include more variation of income levels within the groups.

Table 1: Heterogeneous trading down, by relative income

VARIABLES	(1) Income All HH	(2) Income decile 1	(3) Income decile 2	(4) Income quantile 1	(5) Income quantile 2	(6) Income quantile 3	(7) Income quantile 4	(8) Income quantile 5
Regional Recession	-0.002***	0.000	-0.001	-0.001	-0.003***	-0.002**	-0.002**	-0.001
	(0.000)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	1,470,702	153,953	147,760	304,093	288,926	293,411	292,569	281,597
R-squared	0.758	0.767	0.786	0.760	0.777	0.778	0.782	0.763
Household FE	YES	YES	YES	YES	YES	YES	YES	YES
Time FE	YES	YES	YES	YES	YES	YES	YES	YES

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: Standard errors are clustered at the household level. Relative income is defined as across country level.

Table 2: Heterogeneous trading down, by relative income GFC

VARIABLES	(1) Income All HH	(2) Income decile 1	(3) Income decile 2	(4) Income quantile 1	(5) Income quantile 2	(6) Income quantile 3	(7) Income quantile 4	(8) Income quantile 5
Regional Recession × GFC	-0.007***	-0.003	-0.009***	-0.006**	-0.008***	-0.005**	-0.007***	-0.006**
	(0.001)	(0.004)	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Observations	1,470,702	153,953	147,760	304,093	288,926	293,411	292,569	281,597
R-squared	0.758	0.767	0.786	0.760	0.777	0.778	0.782	0.763
Household FE	YES	YES	YES	YES	YES	YES	YES	YES
Time FE	YES	YES	YES	YES	YES	YES	YES	YES

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: Standard errors are clustered at the household level. Relative income is defined as across country level.

6 General equilibrium effects: a shift-share research design

6.1 A Shift-Share Research design

In the previous section, I document that households are heterogeneous in their ability to trade down in the quality of the goods they purchase. Specifically, low income households exhibit no change in their quality margin when hit by a recession, presumably due to the fact that they were already purchasing lower quality goods before the recession and, therefore, cannot use this margin of adjustment. This lower bound in the quality margin creates an aggregate demand shift toward lower-quality varieties and, therefore, might have aggregate implications for the price of low compared to high-quality varieties. In this section, I aim at studying this effect in detail focusing on the period of the Great Recession. I develop a shift share research design to assess the causal effects of changes in demand on the price index following the methodology proposed in Jaravel (2019).

6.1.1 Intuition

A regression of the amount of quantities traded down on the price of the varieties toward which goods are traded at the onset of a recession would not identify a causal relationship because of, first, reverse causality (the recession might have implications for how goods are priced in a different way over the quality distribution, that is, causality might run from supply to demand) and, second, omitted variable bias (there might be unobserved heterogeneity in how goods are priced in the quality space, which could happen to coincide with the income patterns).

To address these concerns, a shift-share research design relies on two components. First, the predetermined spending shares across the product space for a large number of sociodemographic groups in each state. I focus on spending on intermediate-quality varieties for each good. The reason is that, on the one hand, the consumers of these goods are likely to trade down in the quality of the goods towards the lower quality varieties at the onset of the recession and, on the other hand, given the lower bound in the quality space, this would translate into a heightened demand for the lowest quality varieties. Second, heterogeneity in the population growth rates for these various groups during the same sample period. The sample focuses on the growth between a period before and a period during the great recession. For the groups whose population growth is largest, the predicted demand of the varieties of goods that these households tend to purchase absent a recession will be largest. Given the recession, the propensity to trade down will be largest and as a consequence so will be the demand for the lowest quality varieties. This identifies an exogenous demand increase for lower quality varieties during the great recession and, therefore, allows for a study on the effect it has on relative prices.

6.1.2 IV framework

The goal is to understand how the price index $P_p^{l\in L}$ of all lower quality varieties $l \in L$ within each product p responds to changes in the quantity index $Q_p^{l\in L}$ induced by changes in demand. Conceptually, I aim at finding a demand shifter to vary $Q_p^{l\in L}$ and observe the impact on $P_p^{l\in L}$ across the cells of the product space indexed by p. In other words, I wish to estimate β in the following specification:

$$\Delta log(P_{p,s,t}^{l\in L}) = \beta \Delta log(Q_{p,s,t}^{l\in L}) + \gamma_{t,s} + \delta_{t,p} + \epsilon_{p,s,t} \tag{4}$$

where, $\gamma_{t,s}$ and $\delta_{t,p}$ are time-state and time-product fixed effects, one of which is included in the specification, and only in specifications with more than one period, and $\epsilon_{p,s,t}$ is the unobserved potential outcome that would prevail in p absent changes in demand. Consistent estimation with OLS would require $\mathbb{E}[\Delta log(Q_{p,s,t}^{l\in L}) \times \epsilon_{p,s,t}] = 0$, which is not a plausible assumption because quantities are endogenous to prices. The shift-share design uses variation in $Q_{p,s,t}^{l\in L}$ that comes only from the variation in the size of household groups consuming medium quality goods before the recession.

The shift-share instrument is built to obtain variation in demand from the change in population groups as follows:

$$Z_{p,s,t}^{m\in M} = \sum_{h=1}^{H} s_{h,p,t-1}^{m\in M} \times g^{h,t} \tag{5}$$

Where $g^{h,t} \equiv \Delta log(L^{h,t})$ and H household groups indexed by h are of size L^h , $s_{h,p,t-1}$ denotes the share of sales in p to households of type h *spent in intermediate-quality goods* for a given product in the base period $t-1$.[5] Household groups are defined according to their age, social class, and region. As a consequence, the instrument only uses variation in the demographics, that is, the size of every defined household group, to predict changes in demand. It addresses the concern that changes in demand might be driven by price changes (reverse causality).

Then, I use this instrument in a standard IV framework. The first-stage regression

[5] Note that this term sums 1: among all sales in the medium quality range, a given proportion goes to every household group.

relates the predicted demand for medium-quality goods before the financial crisis with the actual demand for low-quality goods during the financial crisis. The underlying hypothesis is that trading down should lead to a positive relationship between the two variables. The second-stage regression aims at studying the effect of a higher demand for low quality goods, instrumented as detailed, on their relative price.

$$\begin{aligned} \Delta log(P^{l\in L}_{p,s,t}) = \alpha Z^{m\in M}_{p,s,t} + \gamma_{t,s} + \delta_{t,p} + \eta_{p,s,t} \\ \Delta log(Q^{l\in L}_{p,s,t}) = \lambda Z^{m\in M}_{p,s,t} + \gamma_{t,s} + \delta_{t,p} + \epsilon_{p,s,t} \end{aligned} \tag{6}$$

where, $\gamma_{t,s}$ and $\delta_{t,p}$ are time-state and time-product fixed effects, respectively, and are only included in specifications with more than one period, and $\epsilon_{p,s,t}$ is the unobserved potential outcome that would prevail in p absent changes in demand. As before, $l \in L$ denotes the fact that only low quality varieties l within each product p are included and $m \in M$ denotes the fact that only intermediate quality varieties m within each product p are included.

$Q^{l\in L}_{p,s,t}$ denotes the aggregate demand growth for a given product and $P^{l\in L}_{p,s,t}$ the relative price growth of the product. I substract to the price growth of low quality varieties the price growth of the high quality varieties. The reason is twofold: first, the price of high quality varieties serves as a benchmark to compare to and allows to focus on heterogeneous effects in prices across the quality distribution, eliminating all common price shifts. Second, while the middle quality varieties might suffer from more complicated demand shits, given that the households that were likely buying them before the recession might be trading down but other households might begin acquiring these and, therefore, the dynamics of the prices might depend on the relative importance of the two factors. On the other hand, the demand for high quality varieties is likely to not suffer an additional demand shift from households trading down.

In practice, the underlying assumption is that the demographic growth is reasonably exogenous to the crisis and correlated with the amount of trading down in a given region.

Under suitable identification conditions, discussed in the following section, $\frac{\alpha}{\lambda} \to \beta$.

6.1.3 Identification conditions

Instrument relevance requires $\Delta log(Q^{l\in L}_{p})$ and $Z^{m\in M}_{p}$ to be sufficiently correlated and can be directly checked in the first stage.

I refer to the work of Borusyak et al. (2022) to comprehend and verify the exclusion restriction that forms the basis of the instrument validity. Their results show that the exclusion restriction can be expressed as follows:

$$Cov(Z_p^{m\in M}, \epsilon_p) = \mathbb{E}\left[\left(\sum_{h=1}^{H} s_{hp0}^{m\in M} \times g^h\right) \times \epsilon_p\right] = \sum_{h=1}^{H} s_h^{m\in M} \times g^h \mathbb{E}\left[\frac{s_{hp0}^{m\in M}}{s_h^{m\in M}} \times \epsilon_p\right] \to 0 \quad (7)$$

where the covariance and the expectation are taken over the middle-quality varieties $m \in M$ in the product space indexed by p. The key identification condition shown in equation 7 is a weighted covariance (in household space indexed by h, with spending weights s_h between the shocks g^h and the unobservable term $\mathbb{E}\left[\frac{s_{hp0}^{m\in M}}{s_h^{m\in M}} \times \epsilon_p\right]$. This term is a weighted average of product space unobservable potential outcomes ϵ_p.

In Jaravel (2019), a fundamental assumption is that manufacturers possess the foresight to predict shifts in market demand resulting from changes in the population sizes of diverse socio-demographic groups. Under this premise, the instrumental variable (IV) estimates capture the supply reaction to well-anticipated demand changes. In contrast, my focus lies in identifying the short-run supply curve, where manufacturers are not assumed to predict the impact of heterogeneous trading down. If they were to do so, this wouldn't solely involve accounting for population growth trends among households that buy their specific low-quality varieties, but would also encompass accounting for trends among household groups that typically purchase middle-quality varieties. These groups might initially buy such varieties before eventually trading down at the onset of a recession and selecting their own low-quality variety.

In practice, certain household shocks might violate the exclusion restriction. As pointed out in Jaravel (2019), older households tend to grow faster. This would imply a larger g^h for these. Older household groups are more likely to have defined their preferences earlier and, therefore, less likely to adopt new products or vary the quality of the goods they purchase over the cycle. This implies that their $\mathbb{E}\left[\frac{s_{hp0}^{m\in M}}{s_h^{m\in M}} \times \epsilon_p\right]$ might be systematically larger. This would potentially invalidate the exclusion restriction across age groups.

In the following section I discuss the use of fixed effects to address such potential concerns.

6.1.4 Residualised shift-share instrument

To ensure that the aforementioned potential risks to the validity of the instrument are not problematic, I generate more distinctive household population shocks by concentrating on fluctuations within groups, rather than across different household groups. Borusyak et al.

(2022) show that residualising the instrument in the following way is equivalent to running a one-step IV specification with household characteristics onto the product space using initial spending shares.

I consider the following statistical decomposition of the shocks g^h:

$$g_t^h = \mu + g_{age} + g_{socialclass} + g_{region} + \nu_{h,t} \tag{8}$$

This expression suggests that the observed shocks g_t^h can be decomposed into the average shocks along the three dimensions that segment the household space (age, social class, region) as well as a residual component $\nu_{h,t}$.

One can compute a residualised household population shock $\tilde{g}_t^h$ after controlling for age, social class and region either simultaneously or separately. Then one can build the residualised shift-share instrument $\tilde{Z}_p = \sum_{h=1}^{H} s_{hp0} \times \tilde{g}^h$.

Controlling for age fixed effects means that the instrument only relies on variation in household shocks that occur within each age group, addressing the concern about the validity of the exclusion restriction across age groups. I build the residualised shift-share in two steps. First, I regress g^h on household group fixed effects as in Equation 8 to obtain the residualised household population shocks $\tilde{g}_t^h$. Then, I build the shift-share instrument $\tilde{Z}_p$.

Table 3 presents summary statistics on the residualised household population shocks, introducing different controls. To avoid household group changes to be driven by changes in the sampling, I demean household population growth in each year. Therefore, the mean is mechanically zero across all columns. As observed in the standard deviation and interquartile ranges, the amount of variation in household shocks remains very similar across specifications and as controls are added. In particular, the standard deviation drops slightly from 0.085 to 0.081 and the interquartile range from 0.099 to 0.098 as controls are added. This implies that a singular dimension of the data doesn't exclusively drive the variability in household shocks, thus lending support to the notion of employing them within a quasi-experimental framework. With the incorporation of additional fixed effects, the quasi-experimental interpretation gains greater credibility due to the potential reduction in bias. Nonetheless, there is a trade-off, as the instrument's effectiveness might diminish, potentially leading to an increase in variance.

6.2 Implementation

I perform the analysis in different time frames, always including at least one period before the financial crisis and one during the financial crisis. Additionally, to conduct a placebo

Table 3: Changes in population of household groups (2005-2018, yearly averages)

	Annual log change in group population				
	(1)	(2)	(3)	(4)	(5)
Mean	-0.002	0	0	0	0
Standard deviation	0.085	0.082	0.082	0.081	0.081
Interquartile range	0.099	0.092	0.094	0.092	0.089
Residual change after controlling for					
Raw		✓	✓	✓	✓
Linear age			✓	✓	✓
Age fe				✓	✓
Social class fe					✓
Sample sizes					
N total	112				
N age	3				
N social class	3				
N regions	16				

test, I select two periods preceding the recession and run the analysis on these periods as well.

I first adopt two distinct specifications. First, one where I include two periods of growth: second semester of 2006 to second semester of 2007; and first semester of 2007 to first semester of 2008. Under this specification I can include state, time, and product fixed effects. In the second specification I include only the second mentioned period of growth.

The placebo test consists of including data from the first semester of 2006 to the first semester of 2007, two period before the beginning of the crisis. These placebo tests are carried out to establish a comparison with periods where no actual trading down in low-quality goods is expected.

I conduct my analysis at the product-region level, making the fundamental assumption that prices of the same product can vary across regions and respond to different dynamics. This assumption seems reasonable for two main reasons. Firstly, certain product varieties are exclusively offered by regional brands, leading to potential price differences. Secondly, larger firms may strategically adjust prices at the regional level based on varying demand conditions.

I follow the COICOP (Classification of Individual Consumption According to Purpose) classification system. While the dataset provides COICOP-5 product classifications, I employ barcode descriptions to achieve a more refined COICOP-10 classification for products.[6] To achieve this, I implement a matching strategy that involves identifying common words or letters shared between the barcode descriptions and the COICOP-10 classifica-

[6]For it, I follow the mapping from the Federal Statistical Office of Germany (2019) "Consumer price index for Germany. Weighting pattern for base year 2015", available online here.

tions inspired by Beck et al. (2022). Subsequently, I improve the matching accuracy by manually classifying products into specific sub-products. For robustness, I separately also implement the analysis defining products at the COICOP-5 level classification.

I classify varieties (barcodes) as being high quality if their price is above the 75th percentile of the prices of a given product (COICOP-10 or COICOP-5) in a given period of time and in a given region. Likewise, I classify varieties as being of low quality if their price is below the 25th percentile of the prices of a given product in a given period of time and in a given region. The rest are medium quality varieties.

To construct the instrument, I first classify households into groups and aim at using the growth of these groups as an indicator for future demand. I classify households according to their social class, their age and their region. For the social class, I rely on the classification available in the dataset that is constructed based on the profession, role within the company and education levels of the head of the household. I divide households into 3 groups based on this variable. Regarding age, I classify households in 3 groups according to the age of the head of the household: less than 45, between 45 and 60, and more than 60. The third dimension is the region (state) of residence of the household. Given that some of these groups are not represented in the dataset, this gives a total of 112 groups to study.[7] The underlying assumption is that the growth of the population in each of these groups is correlated with their growth in the dataset.

I demean population growth at the yearly level to prevent additional noise from sample increases. Figure A.8 compares relative sample population growth and actual population growth in the 5 largest states across the years included in the sample. While the sample growths tend to display larger variation, the differential growths between groups and regions are generally well captured.[8] Figure A.9 shows the average population growth by age group, across regions and across the years around the financial crisis (2007 to 2013). As before, the general trends are well captured although the variability of the data is larger. Finally, A.10 compares the average population growth by age group and year. As before, the dataset follows the general trends decently well, with a generalised decrease in population in the youngest cohort and a decreasing increase in the middle-aged group. It is important to note that the additional noise driven by the differentials in growth between the sample and the actual population growth should not affect the validity of the results other than by debilitating the second stage relationship.

[7]Household groups containing less than 200 observations are dropped.

[8]One exception is the region of Bayern where growth of the oldest population group is larger than growth of second oldest population group, while the statistics show the opposite trend.

6.3 Empirical results

The baseline findings are presented in Table 4. Columns 1, 2, and 3 explore different specifications of the baseline specifications. Notably, the results consistently demonstrate a positive and significant relationship between the demand for low-quality goods (identified through the *predicted* demand growth of middle-quality goods in the absence of a recession and driven by demographic factors) and the *actual* increase in prices of low-quality goods during a recession. This suggests that the aggregate demand shift towards low-quality goods *causes* an increase in their relative prices.

In Columns 1, 2 and 3 the regression includes 2 periods: growth from second semester of 2006 to second semester of 2007 and growth from first semester of 2007 to first semester of 2008. In Column 1 includes Time FE and Product Division FE and portrays the baseline results. On average, a 1% increase in the demand for lower quality varieties is translated into a 0.33% increase in the relative price of low versus higher quality varieties.

Columns 2 and 3 investigate the differential effect of this phenomenon across products in a given region and across regions in a given product. First, in Column 2, Region-Time fixed effects are included, meaning that the variation focuses on the differences across products within a given region and period of time. Specifically, the analysis reveals that a 1% increase in the demand for low-quality goods corresponds to approximately a 0.98% increase in the price of these goods. In Column 3 Product (COICOP-10) $\times$ Time fixed effects are included instead and therefore shows that the effect also is present and significant across regions for the same product. However, the magnitude is considerably smaller: a 1% increase in demand leads to a 0.19% increase in the relative price of goods. Therefore, while most of the variation seems to derive from differential effects across products, there is evidence that the effect also happens across regions for a specific product type.

The specification in Column 4 focuses on one single period of time (growth between the first semester of 2007 and the first semester of 2008). The resulting coefficient is of a similar magnitude to that of the first column. This is in contrast with Column 5, which can be thought of as the placebo experiment where the period studied is from before the great financial crisis. Under this scenario, one would expect to not find a generalised trading down and, therefore, the shift-share would not be able to identify changes in the demand for low quality goods. In other words, the instrument would be irrelevant. As it can be observed, the first stage appears to be significantly weaker than for the other specifications and the coefficient of the second stage regression becomes insignificant in this case.

Table 4: Shift-share instrument results

	(1)	(2)	(3)	(4)	(5)
VARIABLES	Δ log Price low (relative)	Δ log Price low (relative)	Δ log Price low (relative)	Δ log Price low (relative) 2007S1-2008S1	Δ log Price low (relative) 2006S1-2007S1
Δ log demand low q.	0.334**	0.981***	0.185***		
	(0.141)	(0.360)	(0.043)		
Δ log demand low q., 2007S1-2008S1				0.343**	
				(0.149)	
Δ log demand low q., 2006S1-2007S1					0.093
					(0.475)
Observations	9,108	9,108	9,108	4,702	4,624
Time FE	YES	-	-		
Product Division FE	YES	NO	NO		
Time × Region FE	NO	YES	NO		
Time × Product FE	NO	NO	YES		
First Stage F	97.22	16.97	233.6	49.52	5.722
Cragg-Donald F	188.6	27.10	282.1	78.06	7.783

Notes: The table presents the results of the IV estimation specification in Equation **??**. The instrumented variable is demand growth of a product category in a given state between the periods specified in each column, where $S \in 1, 2$ indicates the semester. Columns 1, 2 and 3 include growth data from two periods: 2006S2 to 2007S2 and 2007S1 to 2008S1. Product FE refers to COICOP-10 product classification and Product division refer to COICOP-2 product classification. The instrument is a shift share design as described in Equation 5. Standard errors are clustered at the state-product level.

6.3.1 Robustness checks: Results with residualised shift-share

Tables 5 to 8 present the results once the instrument is controlling for age, social class and region fixed effects as exposed in Section 8. In all, the results remain robust. The magnitudes of the coefficients are also very in line with previous findings. Therefore, the previous results were not driven by the concerns exposed in Section 8.

Table 5: Residualised shift share by age (linear)

	(1)	(2)	(3)	(4)	(5)
VARIABLES	Δ log Price low (relative)	Δ log Price low (relative)	Δ log Price low (relative)	Δ log Price low (relative) 2007S1-2008S1	Δ log Price low (relative) 2006S1-2007S1
Δ log demand low q.	0.308**	0.917***	0.174***		
	(0.139)	(0.346)	(0.042)		
Δ log demand low q., 2007S1-2008S1				0.325**	
				(0.148)	
Δ log demand low q., 2006S1-2007S1					0.008
					(0.463)
Observations	9,108	9,108	9,108	4,702	4,624
Time FE	YES	-	-		
Product Division FE	YES	NO	NO		
Time × Region FE	NO	YES	NO		
Time × Product FE	NO	NO	YES		
First Stage F	99.09	17.17	241.8	49.50	5.735
Cragg-Donald F	191.7	27.22	287.5	78.75	7.836

Table 6: Residualised shift share by age (f.e.)

	(1)	(2)	(3)	(4)	(5)
VARIABLES	Δ log Price low (relative)	Δ log Price low (relative)	Δ log Price low (relative)	Δ log Price low (relative) 2007S1-2008S1	Δ log Price low (relative) 2006S1-2007S1
Δ log demand low q.	0.258*	0.795**	0.153***		
	(0.137)	(0.345)	(0.040)		
Δ log demand low q., 2007S1-2008S1				0.292**	
				(0.146)	
Δ log demand low q., 2006S1-2007S1					-0.171
					(0.456)
Observations	9,108	9,108	9,108	4,702	4,624
Time FE	YES	-	-		
Product Division FE	YES	NO	NO		
Time × Region FE	NO	YES	NO		
Time × Product FE	NO	NO	YES		
First Stage F	101.8	15.46	260.8	48.22	5.517
Cragg-Donald F	196.5	23.98	301.8	78.53	7.710

Notes: The table presents the results of the IV estimation specification in Equation **??**. The instrumented variable is demand growth of a product category in a given state between the periods specified in each column, where $S \in 1, 2$ indicates the semester. Columns 1, 2 and 3 include growth data from two periods: 2006S2 to 2007S2 and 2007S1 to 2008S1. Product FE refers to COICOP-10 product classification and Product division refer to COICOP-2 product classification. The instrument is a shift share design as described in Equation 5. Standard errors are clustered at the state-product level.

Table 7: Residualised shift share by age and social class (f.e.)

VARIABLES	(1) Δ log Price low (relative)	(2) Δ log Price low (relative)	(3) Δ log Price low (relative)	(4) Δ log Price low (relative) 2007S1-2008S1	(5) Δ log Price low (relative) 2006S1-2007S1
Δ log demand low q.	0.266* (0.137)	0.980** (0.424)	0.152*** (0.040)		
Δ log demand low q., 2007S1-2008S1				0.304** (0.146)	
Δ log demand low q., 2006S1-2007S1					-0.196 (0.463)
Observations	9,108	9,108	9,108	4,702	4,624
Time FE	YES	-	-		
Product Division FE	YES	NO	NO		
Time × Region FE	NO	YES	NO		
Time × Product FE	NO	NO	YES		
First Stage F	102.3	12.60	263.2	48.20	5.349
Cragg-Donald F	195.8	19.22	305.7	77.83	7.498

Table 8: Residualised shift share by age, social class and region (f.e.)

VARIABLES	(1) Δ log Price low (relative)	(2) Δ log Price low (relative)	(3) Δ log Price low (relative)	(4) Δ log Price low (relative) 2007S1-2008S1	(5) Δ log Price low (relative) 2006S1-2007S1
Δ log demand low q.	0.344** (0.167)	0.969** (0.419)	0.232*** (0.055)		
Δ log demand low q., 2007S1-2008S1				0.347** (0.173)	
Δ log demand low q., 2006S1-2007S1					-0.262 (0.488)
Observations	9,108	9,108	9,108	4,702	4,624
Time FE	YES	-	-		
Product Division FE	YES	NO	NO		
Time × Region FE	NO	YES	NO		
Time × Product FE	NO	NO	YES		
First Stage F	69.36	12.72	150.1	37.05	4.876
Cragg-Donald F	129.9	19.43	192	56.45	6.792

Notes: The table presents the results of the IV estimation specification in Equation ??. The instrumented variable is demand growth of a product category in a given state between the periods specified in each column, where $S \in 1, 2$ indicates the semester. Columns 1, 2 and 3 include growth data from two periods: 2006S2 to 2007S2 and 2007S1 to 2008S1. Product FE refers to COICOP-10 product classification and Product division refer to COICOP-2 product classification. The instrument is a shift share design as described in Equation 5. Standard errors are clustered at the state-product level.

6.3.2 Sensitivity check: Results at the COICOP-5 level

To confirm that the finer product classification and the matching methodology undertaken in the previous sections are not driving the results, I perform the same analysis at the COICOP-5 product level classification. This translates into a slightly broader definition of products and, therefore, fewer number of observations. In general, the results of these regressions point toward the same conclusion as in the previous section and suggest that the results are robust to the classification used. While the magnitude of the effect on prices is actually larger than in the previous results, this might be driven by a larger variability in prices due to the product classification. While the coefficient in the first column becomes less precisely estimated as the shift share is residualised, the rest of the results largely preserve their significance.

Table 9: Shift-share instrument results at COICOP-5 level

	(1)	(2)	(3)	(4)	(5)
VARIABLES	Δ log Price low (relative)	Δ log Price low (relative)	Δ log Price low (relative)	Δ log Price low (relative) 2007S1-2008S1	Δ log Price low (relative) 2006S1-2007S1
Δ log demand low q.	0.438*	1.104***	0.152**		
	(0.256)	(0.360)	(0.063)		
Δ log demand low q., 2007S1-2008S1				0.500*	
				(0.258)	
Δ log demand low q., 2006S1-2007S1					0.676
					(2.300)
Observations	2,985	2,985	2,985	1,515	1,499
Time FE	YES	-	-		
Product Division FE	YES	NO	NO		
Time × Region FE	NO	YES	NO		
Time × Product FE	NO	NO	YES		
First Stage F	90.20	30.16	202.8	41.68	1.684
Cragg-Donald F	174.3	49.93	267.1	58.46	2.165

Notes: The table presents the results of the IV estimation specification in Equation **??**. The instrumented variable is demand growth of a product category in a given state between the periods specified in each column, where $S \in 1, 2$ indicates the semester. Columns 1, 2 and 3 include growth data from two periods: 2006S2 to 2007S2 and 2007S1 to 2008S1. Product FE refers to COICOP-5 product classification and Product division refer to COICOP-2 product classification. The instrument is a shift share design as described in Equation 5. Standard errors are clustered at the state-COICOP-5 product level.

Table 10: Residualised shift share by age (linear) at the COICOP-5 level

	(1)	(2)	(3)	(4)	(5)
VARIABLES	Δ log Price low (relative)	Δ log Price low (relative)	Δ log Price low (relative)	Δ log Price low (relative) 2007S1-2008S1	Δ log Price low (relative) 2006S1-2007S1
Δ log demand low q.	0.399	1.066***	0.142**		
	(0.258)	(0.351)	(0.062)		
Δ log demand low q., 2007S1-2008S1				0.478*	
				(0.260)	
Δ log demand low q., 2006S1-2007S1					0.402
					(0.775)
Observations	2,985	2,985	2,985	1,515	1,499
Time FE	YES	-	-		
Product Division FE	YES	NO	NO		
Time × Region FE	NO	YES	NO		
Time × Product FE	NO	NO	YES		
First Stage F	89.12	30.39	203.7	40.60	3.709
Cragg-Donald F	171.9	50.14	265.3	57.19	2.107

Table 11: Residualised shift share by age (f.e.) at the COICOP-5 level

	(1)	(2)	(3)	(4)	(5)
VARIABLES	Δ log Price low (relative)	Δ log Price low (relative)	Δ log Price low (relative)	Δ log Price low (relative) 2007S1-2008S1	Δ log Price low (relative) 2006S1-2007S1
Δ log demand low q.	0.315	0.970***	0.119*		
	(0.262)	(0.339)	(0.061)		
Δ log demand low q., 2007S1-2008S1				0.432	
				(0.267)	
Δ log demand low q., 2006S1-2007S1					-0.213
					(1.309)
Observations	2,985	2,985	2,985	1,515	1,499
Time FE	YES	-	-		
Product Division FE	YES	NO	NO		
Time × Region FE	NO	YES	NO		
Time × Product FE	NO	NO	YES		
First Stage F	85.83	29.35	206.9	37.25	2.683
Cragg-Donald F	166	47.55	264.8	53.52	1.960

Notes: The table presents the results of the IV estimation specification in Equation **??**. The instrumented variable is demand growth of a product category in a given state between the periods specified in each column, where $S \in 1, 2$ indicates the semester. Columns 1, 2 and 3 include growth data from two periods: 2006S2 to 2007S2 and 2007S1 to 2008S1. Product FE refers to COICOP-5 product classification and Product division refer to COICOP-2 product classification. The instrument is a shift share design as described in Equation 5. Standard errors are clustered at the state-COICOP-5 product level.

Table 12: Residualised shift share by age and social class (f.e.) at the COICOP-5 level

	(1)	(2)	(3)	(4)	(5)
VARIABLES	Δ log Price low (relative)	Δ log Price low (relative)	Δ log Price low (relative)	Δ log Price low (relative) 2007S1-2008S1	Δ log Price low (relative) 2006S1-2007S1
Δ log demand low q.	0.318	1.041***	0.117*		
	(0.260)	(0.366)	(0.061)		
Δ log demand low q., 2007S1-2008S1				0.439*	
				(0.266)	
Δ log demand low q., 2006S1-2007S1					-0.230
					(1.316)
Observations	2,985	2,985	2,985	1,515	1,499
Time FE	YES	-	-		
Product Division FE	YES	NO	NO		
Time × Region FE	NO	YES	NO		
Time × Product FE	NO	NO	YES		
First Stage F	86.69	27.23	207.4	37.43	2.764
Cragg-Donald F	166.3	42.62	268.3	53.23	1.951

Table 13: Residualised shift share by age, social class and region (f.e.) at the COICOP-5 level

	(1)	(2)	(3)	(4)	(5)
VARIABLES	Δ log Price low (relative)	Δ log Price low (relative)	Δ log Price low (relative)	Δ log Price low (relative) 2007S1-2008S1	Δ log Price low (relative) 2006S1-2007S1
Δ log demand low q.	0.456	1.033***	0.252***		
	(0.309)	(0.365)	(0.084)		
Δ log demand low q., 2007S1-2008S1				0.518*	
				(0.298)	
Δ log demand low q., 2006S1-2007S1					-0.267
					(1.372)
Observations	2,985	2,985	2,985	1,515	1,499
Time FE	YES	-	-		
Product Division FE	YES	NO	NO		
Time × Region FE	NO	YES	NO		
Time × Product FE	NO	NO	YES		
First Stage F	60.11	27.32	116.5	29.95	2.600
Cragg-Donald F	114	42.79	172.7	40.27	1.821

Notes: The table presents the results of the IV estimation specification in Equation **??**. The instrumented variable is demand growth of a product category in a given state between the periods specified in each column, where $S \in 1, 2$ indicates the semester. Columns 1, 2 and 3 include growth data from two periods: 2006S2 to 2007S2 and 2007S1 to 2008S1. Product FE refers to COICOP-5 product classification and Product division refer to COICOP-2 product classification. The instrument is a shift share design as described in Equation 5. Standard errors are clustered at the state-COICOP-5 product level.

7 Conclusions

In this study, I explore the significance of the quality margin as a protective mechanism against aggregate shocks, focusing specifically on the heterogeneous behavior observed among households across the income distribution. Using household scanner data from Germany, I analyze the degree to which households engage in trading down. I provide evidence that, on average, lower income households tend to opt for lower quality goods. Moreover, in the aftermath of an aggregate shock, lower income households demonstrate a limited inclination to engage in trading down, presumably due to their constrained capacity to do so. This stands in contrast to other households, who exhibit a greater propensity to trade down by selecting lower quality goods. To comprehensively understand the broader implications of this shift in aggregate demand towards lower quality goods, I employ a shift-share research design. I find that this aggregate demand shift toward lower quality varieties in the aftermath of a recession increases the relative price of low versus high quality varieties. This might have implications for inflation risk over the business cycle for lower income households.

8 Bibliography

References

Aguiar, M. & Hurst, E. (2005), 'Consumption versus expenditure', *Journal of political Economy* **113**(5), 919–948.

Argente, D. & Lee, M. (2021), 'Cost of living inequality during the great recession', *Journal of the European Economic Association* **19**(2), 913–952.

Beck, G., Carstensen, K., Menz, J.-O., Schnorrenberger, R. & Wieland, E. (2022), Using high-frequency scanner data to evaluate german food prices in real time, Technical report, Working Paper.

Bertolotti, F., Gavazza, A. & Lanteri, A. (2021), 'Dynamics of expenditures on durable goods: the role of new-product quality'.

Boel, P., Diaz, J. & Finocchiaro, D. (2021), 'Liquidity, capital pledgeability and inflation redistribution', *Available at SSRN 3961028* .

Borusyak, K., Hull, P. & Jaravel, X. (2022), 'Quasi-experimental shift-share research designs', *The Review of Economic Studies* **89**(1), 181–213.

Burstein, A., Eichenbaum, M. & Rebelo, S. (2005), 'Large devaluations and the real exchange rate', *Journal of political Economy* **113**(4), 742–784.

Cardoso, M., Ferreira, C., Leiva, J. M., Nuño, G., Ortiz, Á., Rodrigo, T. & Vazquez, S. (2022), 'The heterogeneous impact of inflation on households' balance sheets', *Red Nacional de Investigadores en Economıa Working Paper* **176**.

Carvalho, V. M., Garcia, J. R., Hansen, S., Ortiz, Á., Rodrigo, T., Rodríguez Mora, J. V. & Ruiz, P. (2021), 'Tracking the covid-19 crisis with high-resolution transaction data', *Royal Society Open Science* **8**(8), 210218.

Colicev, A., Hoste, J. & Konings, J. (2022), 'The impact of a large depreciation on the cost of living of rich and poor consumers'.

Cravino, J., Lan, T. & Levchenko, A. A. (2020), 'Price stickiness along the income distribution and the effects of monetary policy', *Journal of Monetary Economics* **110**, 19–32.

Gavazza, A. & Lanteri, A. (2021), 'Credit shocks and equilibrium dynamics in consumer durable goods markets', *The Review of Economic Studies* **88**(6), 2935–2969.

Jaimovich, N., Rebelo, S. & Wong, A. (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 us retail sector', *The Quarterly Journal of Economics* **134**(2), 715–783.

Kaplan, G. & Menzio, G. (2015), 'The morphology of price dispersion', *International Economic Review* **56**(4), 1165–1206.

Kaplan, G. & Schulhofer-Wohl, S. (2017), 'Inflation at the household level', *Journal of Monetary Economics* **91**, 19–38.

Lauper, C., Mangiante, G. et al. (2021), *Monetary policy shocks and inflation inequality*, Université de Lausanne, Faculté des hautes études commerciales (HEC ....

Michelacci, C., Paciello, L. & Pozzi, A. (2022), 'The extensive margin of aggregate consumption demand', *The Review of Economic Studies* **89**(2), 909–947.

Nord, L. (2022), 'Shopping, demand composition, and equilibrium prices', *Available at SSRN 4178271* .

Orchard, J. (2022), 'Cyclical demand shifts and cost of living inequality', *Available at SSRN 4033572* .

Rodnyansky, A., Van der Ghote, A. & Wales, D. (2022), 'Product quality, measured inflation and monetary policy'.

Yang, Y. (2022), 'Redistributive inflation and optimal monetary policy', *Available at SSRN 4275770* .

A Appendix

Table A.1: Heterogeneous trading down, by relative income (within state)

VARIABLES	(1) Income All HH	(2) Income decile 1	(3) Income decile 2	(4) Income quantile 1	(5) Income quantile 2	(6) Income quantile 3	(7) Income quantile 4	(8) Income quantile 5
Regional Recession	-0.002***	0.000	0.000	0.000	-0.002**	-0.003***	-0.002**	-0.002**
	(0.000)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	1,470,702	155,838	145,809	304,044	290,158	291,251	291,919	283,298
R-squared	0.758	0.767	0.783	0.759	0.778	0.778	0.780	0.763
Household FE	YES	YES	YES	YES	YES	YES	YES	YES
Time FE	YES	YES	YES	YES	YES	YES	YES	YES

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: Standard errors are clustered at the household level. Relative income is defined as within-state level.

Table A.2: Heterogeneous trading down, by social class

VARIABLES	(1) Social class all	(2) Social class 1	(3) Social class 2	(4) Social class 3	(5) Social class 4	(6) Social class 5	(7) Social class 6
Regional Recession	-0.002***	0.001	-0.002	-0.002***	-0.002***	-0.002	0.001
	(0.000)	(0.005)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
Observations	1,470,702	15,575	242,561	474,601	554,436	119,253	62,288
R-squared	0.758	0.773	0.759	0.767	0.771	0.786	0.792
Household FE	YES	YES	YES	YES	YES	YES	YES
Time FE	YES	YES	YES	YES	YES	YES	YES

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: Standard errors are clustered at the household level. Households are grouped by their social class.

Table A.3: Heterogeneous trading down, by relative income (within state) GFC

VARIABLES	(1) Income All HH	(2) Income decile 1	(3) Income decile 2	(4) Income quantile 1	(5) Income quantile 2	(6) Income quantile 3	(7) Income quantile 4	(8) Income quantile 5
Regional Recession × GFC	-0.007***	-0.003	-0.006*	-0.004*	-0.005*	-0.008***	-0.007***	-0.005**
	(0.001)	(0.004)	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Observations	1,470,702	155,838	145,809	304,044	290,158	291,251	291,919	283,298
R-squared	0.758	0.767	0.783	0.759	0.778	0.778	0.780	0.763
Household FE	YES	YES	YES	YES	YES	YES	YES	YES
Time FE	YES	YES	YES	YES	YES	YES	YES	YES

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: Standard errors are clustered at the household level. Relative income is defined as within-state level.

Figure A.1: Inflation risk during the recession

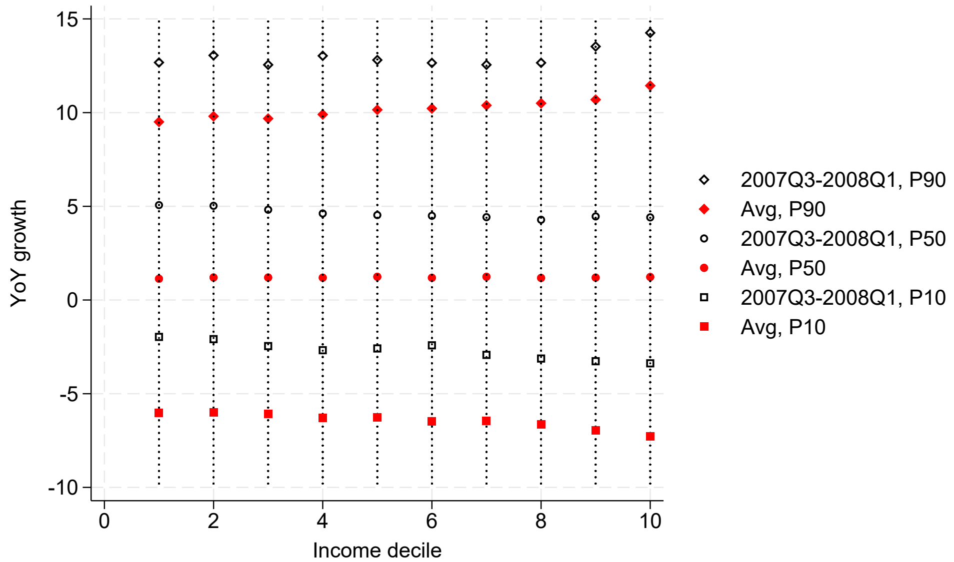

Notes: Laspeyres Inflation is computed at the household level. The data covers German households and spans from 2005 to 2018. GFC refers to the beginning of the financial crisis and includes the second semester of 2007 and all 2008.

Figure A.2: Inflation risk during the recession: all GFC

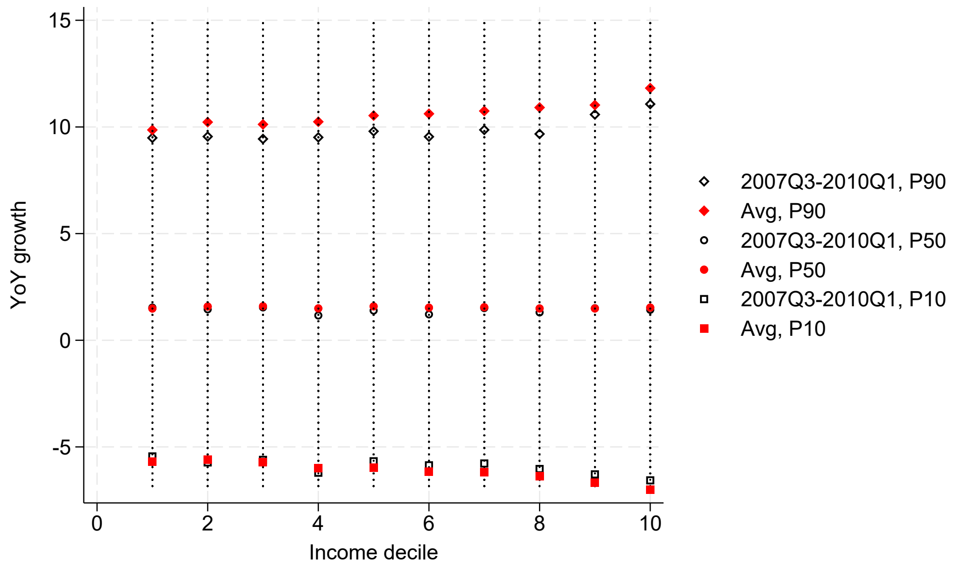

Notes: Laspeyres Inflation is computed at the household level. The data covers German households and spans from 2005 to 2018. GFC refers to the beginning of the financial crisis and includes the second semester of 2007, all 2008 and 2009.

Figure A.3: Relative prices of low vs high quality goods

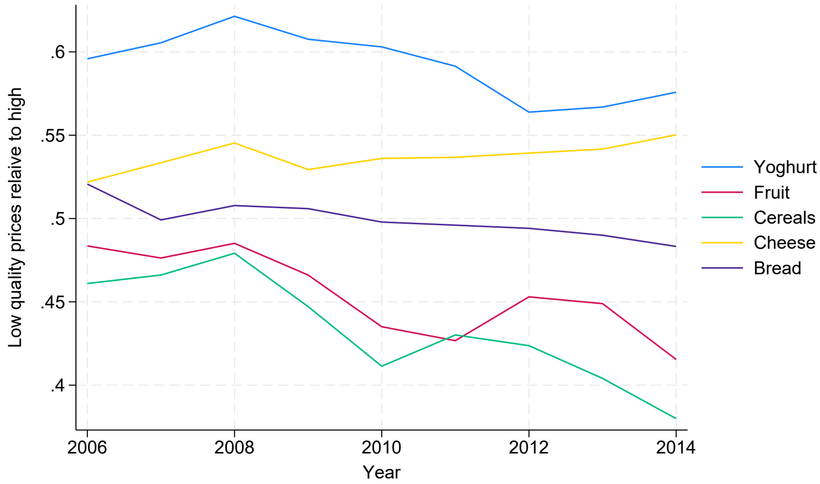

Note: The relative price of low quality goods is defined as the 25th percentile of a product category divided by the median price of that product category within each year. Product categories are defined by the COICOP-5 classification.

Figure A.4: Decomposition by income group within state

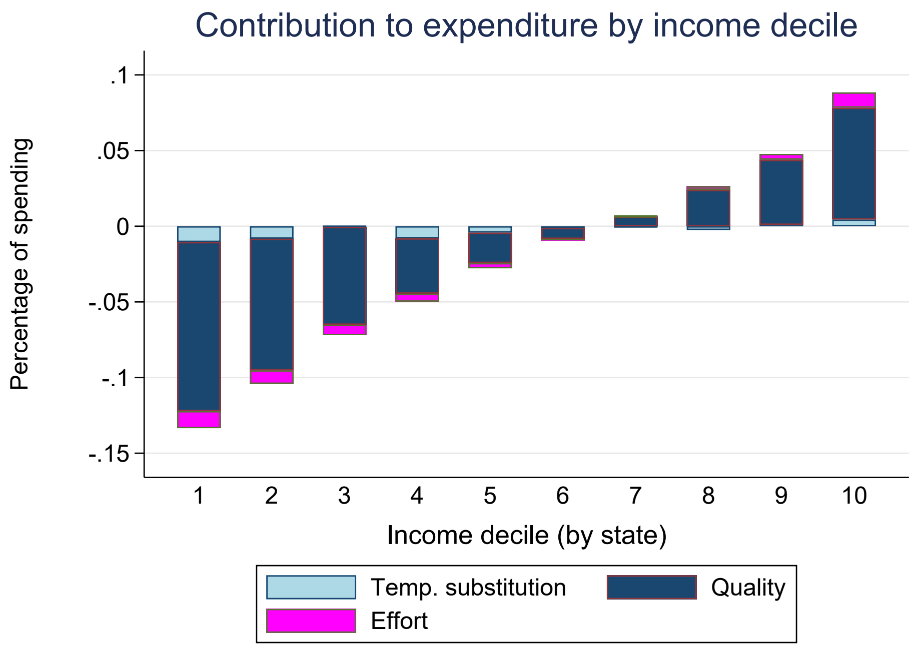

Figure A.5: Decomposition by social class

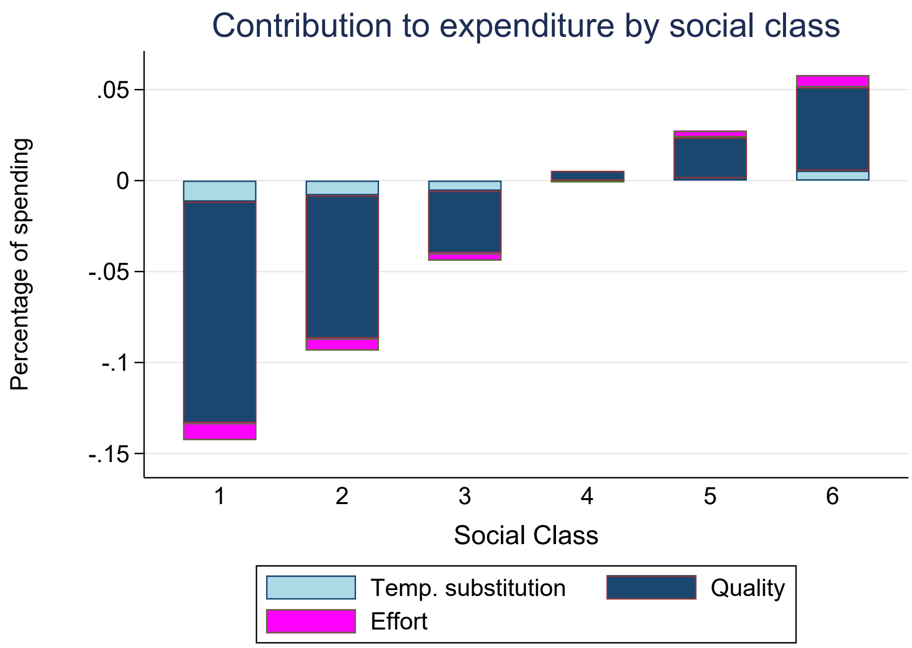

Figure A.6: Expenditures and counterfactual by income group within state

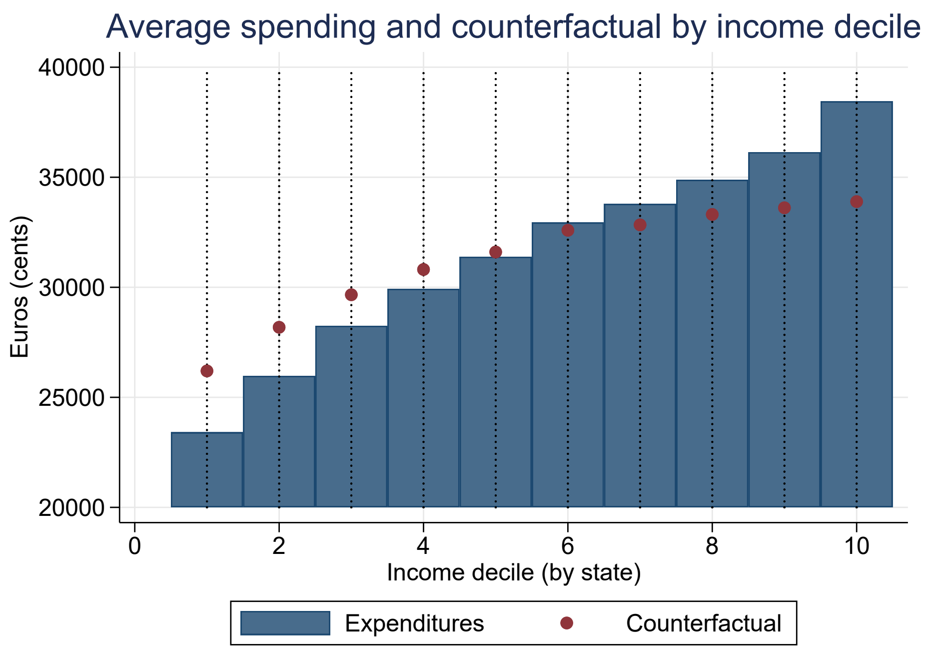

Figure A.7: Expenditures and counterfactual by social class

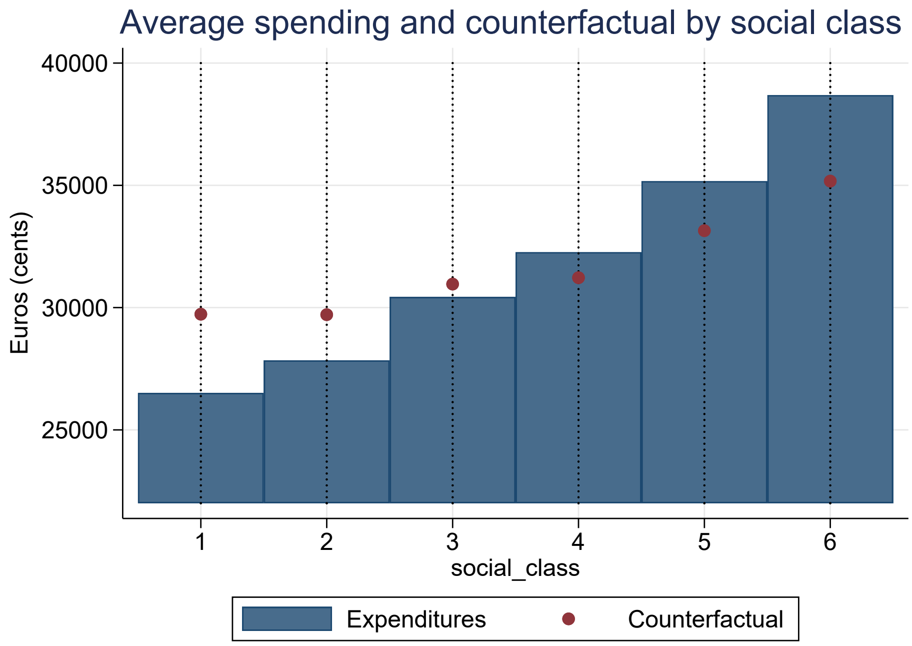

Table A.4: Heterogeneous trading down, by social class GFC

VARIABLES	(1) Income All HH	(2) Social class 1	(3) Social class 2	(4) Social class 3	(5) Social class 4	(6) Social class 5	(7) Social class 6
Regional Recession × GFC	-0.007***	0.000	-0.008***	-0.008***	-0.005***	-0.006	-0.004
	(0.001)	(0.012)	(0.003)	(0.002)	(0.002)	(0.004)	(0.004)
Observations	1,470,702	15,575	242,561	474,601	554,436	119,253	62,288
R-squared	0.758	0.773	0.759	0.767	0.771	0.786	0.792
Household FE	YES	YES	YES	YES	YES	YES	YES
Time FE	YES	YES	YES	YES	YES	YES	YES

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: Standard errors are clustered at the household level. Households are grouped by their social class,

Figure A.8: Population and sample population growth by region and age group

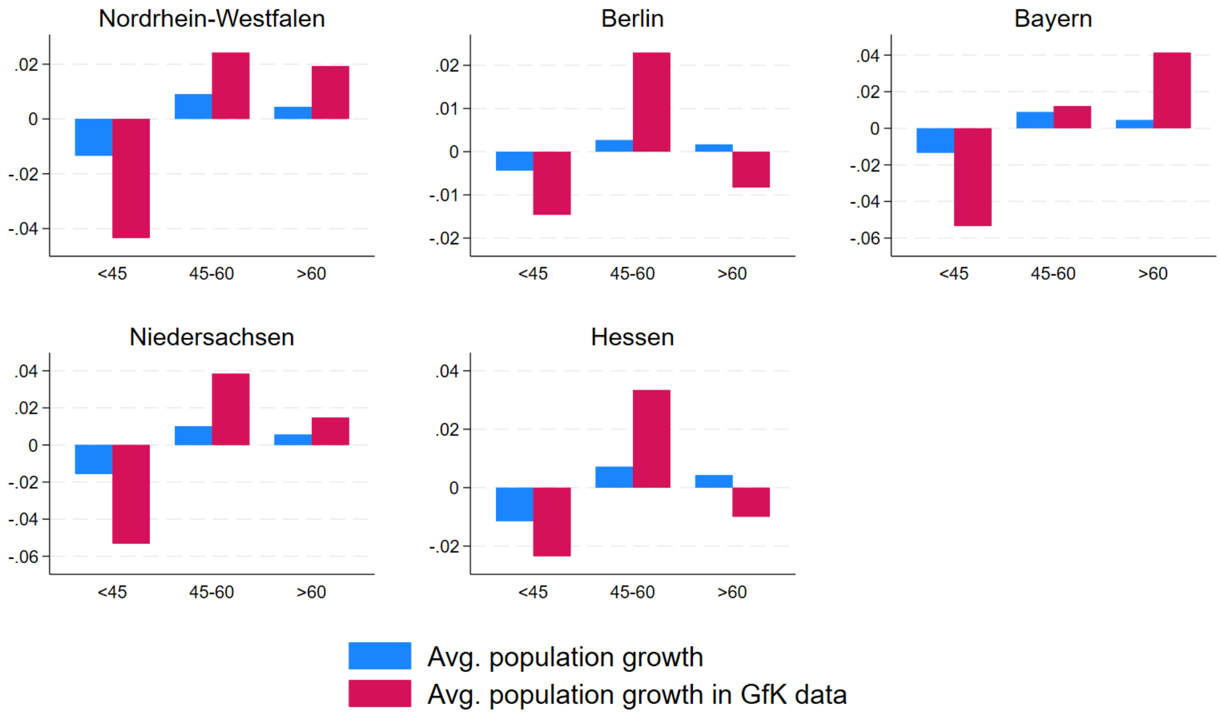

Note: 2006 to 2018 average annual population and sample growth in the 5 largest German regions by age group.

Figure A.9: Population and sample population growth by age group

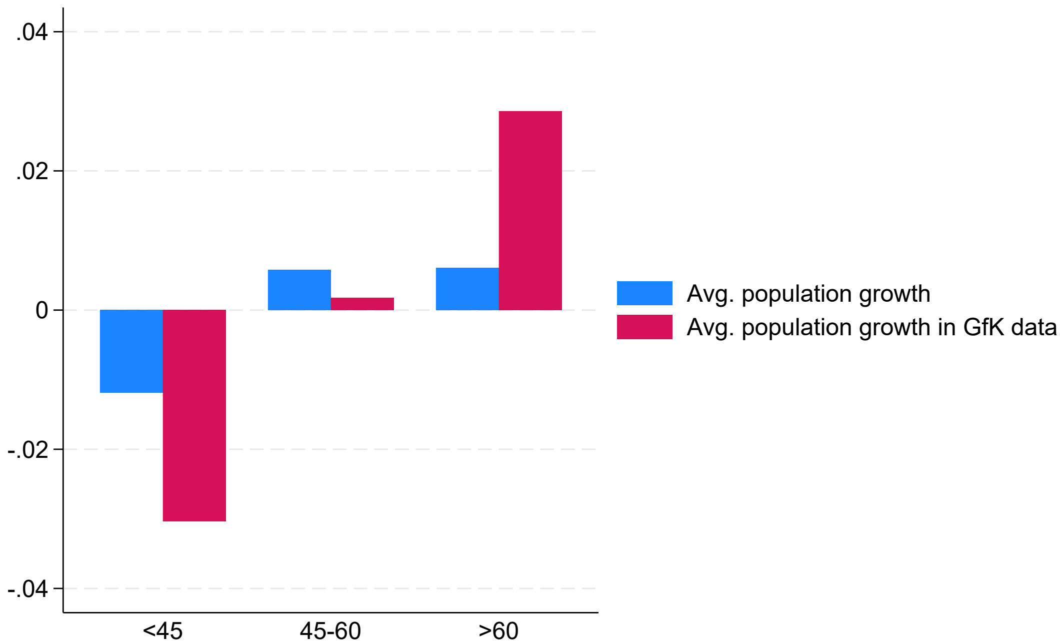

Note: 2007 to 2013 average annual population and sample growth by age group.

Figure A.10: Population and sample population growth by age group and year

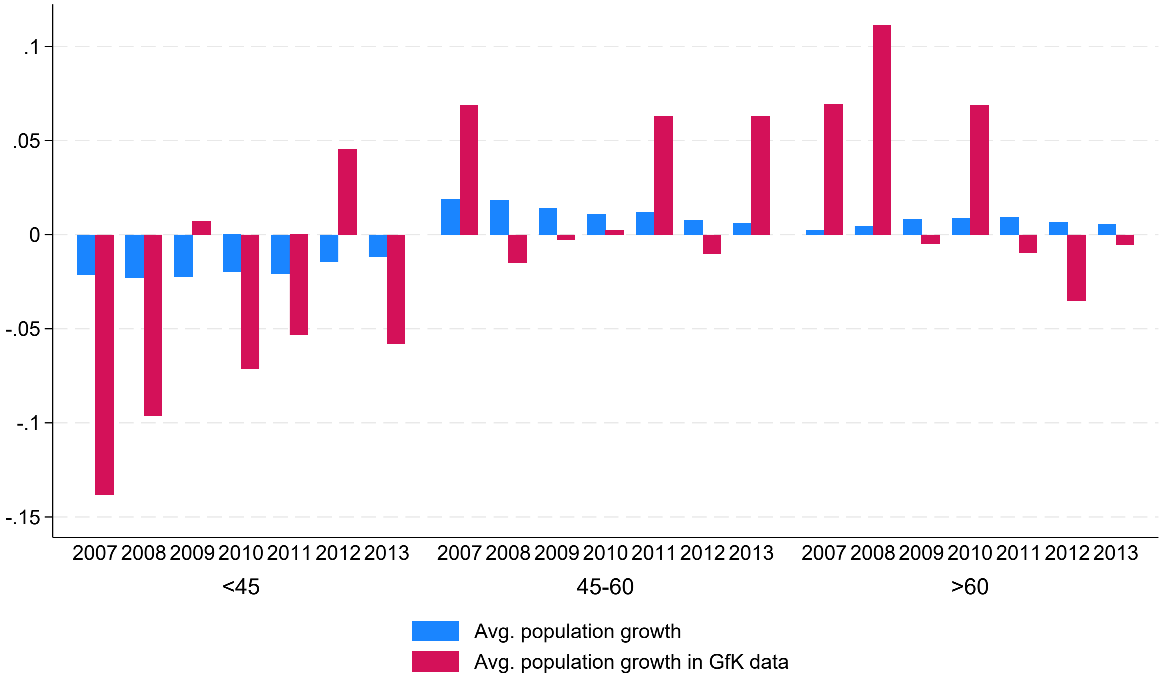

Note: 2006 to 2018 average annual population and sample growth by age group and year.