

THE EFFECT OF MONETARY POLICY ON INFLATION HETEROGENEITY ALONG THE INCOME DISTRIBUTION

MIGUEL AMPUDIA, MICHAEL EHLMANN AND GEORG STRASSER

ABSTRACT. This paper studies the effect of monetary policy on inflation along the income distribution in several euro area countries. It shows that monetary policy has differential effects and identifies two channels which point in opposite directions. On the one hand, different consumption shares imply that the inflation experienced by high-income households responds *less* to monetary policy. On the other hand, the paper provides novel evidence that there are substantial differences in shopping behaviour and its reaction to monetary policy, which imply that the inflation experienced by high-income households responds *more* to monetary policy.

JEL codes: E31, E52, D30

Key words: inflation, distributional effects, monetary policy, shopping behaviour, substitution

Date: September 13, 2023.

Ampudia: Bank for International Settlements, miguel.ampudia@bis.org. Ehrmann: European Central Bank and CEPR, michael.ehrmann@ecb.europa.eu. Strasser: European Central Bank, georg.strasser@ecb.europa.eu. This paper presents the authors' personal opinions and does not necessarily reflect the views of the Bank for International Settlements, the European Central Bank or the Eurosystem. The authors would like to thank Romain Aumond and Derry O'Brien for help with the data collection, Lorenz Eichberger and Cesare Dela Pierre for research assistance, and the members of the PRISMA team, Greg Kaplan and seminar participants at the BIS, the BIS Office for the Americas, the ECB and the IARIW-Banca d'Italia conference on Central Banks, Financial Markets and Inequality for comments.

1. INTRODUCTION

Monetary policy has distributional effects. After all, a change in interest rates affects agents differently depending on their interest rate exposure. To give one example, savers and borrowers are exposed in the opposite way, and this is precisely the intention of a monetary policy action. Distributional effects are furthermore triggered because monetary policy affects income and wealth differently across agents. [Coibion, Gorodnichenko, Kueng, and Silvia \(2017\)](#) find that contractionary monetary policy leads to a persistent increase in the inequality in labour earnings and total income, as well as in consumption and total expenditure. This, in turn, affects monetary policy transmission. [Auclert \(2019\)](#), for instance, argues that there is a redistribution channel of monetary policy, whereby the effects of monetary policy are amplified because the winners and losers from a certain policy action have different marginal propensities to consume. [Savignac and Hubert \(2023\)](#) find a U-shaped effect of ECB monetary policy along the labor income distribution. The relevance of heterogeneity for understanding monetary policy transmission is studied extensively in the rapidly growing literature that develops heterogeneous-agent New Keynesian (HANK) models ([Kaplan, Moll, and Violante, 2018](#)).

Beyond the wealth and income effects, monetary policy can trigger distributional effects if inflation itself responds differently across households. This channel has barely been studied, and it is where the current paper aims to contribute. In a first contribution, [Cravino, Lan, and Levchenko \(2020\)](#) show that price stickiness differs along the US income distribution: high-income households tend to consume relatively more goods with stickier and less volatile prices. Accordingly, the inflation rate for their consumption basket responds by around one third less to monetary policy shocks than for middle-income households. As middle-income households experience higher median inflation rates on average, inflation dispersion decreases in response to a monetary contraction ([Lauper and Mangiante, 2021](#)).

Further to a differential response based on the differences in the consumption basket studied by [Cravino et al. \(2020\)](#), inflation might also respond differently to monetary policy because of (changes in) households' purchasing behaviour. [Argente and Lee \(2021\)](#), for instance, show that high-income households had much lower inflation in the aftermath of the Great Recession, because they had more scope for changing their shopping behaviour and for engaging in product quality substitution. Similarly, [Kaplan and Menzio \(2015\)](#) allow for a differential response of shopping behavior to unemployment shocks and document that a transition to unemployment lowers a US household's price index between 1 and 3.5 percent, as food expenditures decline and households spend more time on shopping activities and purchase in less expensive stores. This evidence provides a powerful confirmation of the finding by [Kaplan and Schulhofer-Wohl \(2017\)](#) that the bulk of the cross-sectional variation in US inflation arises due to differences in the prices paid for the same types of goods, not from variation in broadly defined consumption bundles. Such differences in the prices paid arise for various reasons. For instance, high-income households have been shown to exert lower search

effort, which retailers exploit to charge higher markups (Nord, 2022). Also, low-income households tend to rely relatively more on bulk-buying and on purchases of unbranded goods (Griffith, Leibtag, Leicester, and Nevo, 2009). Such differences in the product choice within a product category have also been found to be important in European data (Kiss and Strasser, 2022).

To shed further light on the responsiveness of household-specific inflation to shocks, this paper studies how inflation responds to monetary policy shocks across the income distribution, for the case of the European Central Bank (ECB) and for the six largest euro area countries (i.e., Germany, France, Italy, Spain, the Netherlands and Belgium). Most importantly, it considers both sources of variation identified in the literature, those due to differences in the consumption basket along the income distribution (studied by Cravino et al. (2020)) as well as those arising from differences in purchasing behaviour (highlighted by Kaplan and Menzio (2016), Kaplan and Schulhofer-Wohl (2017) and Argente and Lee (2021), but not studied in relation to monetary policy). As a result, the paper provides a more comprehensive and nuanced picture than the previous literature.

The first key finding is – in line with the earlier literature – that monetary policy affects inflation differently across the income distribution. The second key result is that there are different channels at play. On the one hand, different consumption baskets imply that inflation by high-income households responds *less* to monetary policy. This corroborates the findings of Cravino et al. (2020). On the other hand, the paper provides novel evidence regarding differences in shopping behaviour, which furthermore imply that inflation by high-income households responds *more* to monetary policy, making the overall effect ambiguous.

This paper connects to a large literature on distributional effects of monetary policy. While this is a long-standing question, it has gained prominence following the global financial crisis and the adoption of unconventional monetary policy (UMP) in many advanced economies. Quantitative easing, for instance, pushes up asset prices, thereby disproportionately raising the financial wealth of the small fraction of households that actually hold such assets (Bell, Joyce, Liu, and Young, 2012; Adam and Tzamourani, 2016). While there are many channels at play (for a summary, see Colciago, Samarina, and de Haan (2019)), the income composition and the earnings heterogeneity channel are often found to be the dominating forces. By boosting real activity and lowering unemployment, a monetary easing benefits disproportionately the low- and middle-income households, as these tend to have a higher risk of becoming unemployed in a recession (Mitman, Broer, and Kramer, 2022) and receive a relatively larger share of their income from wages. The importance of these channels has been stressed both in the academic literature (see, e.g., Ampudia, Georgarakos, Slacalek, Tristani, Vermeulen, and Violante, 2018; Lenza and Slacalek, 2018) and by policymakers (e.g. Schnabel, 2020).

Another strand of the literature has studied the distributional effects of inflation in itself. Easterly and Fischer (2001) find that poorer households are more likely to be concerned about inflation than the rich, and Jaravel (2019, 2021) shows that inflation rates do indeed decline with household income

– at least during the recent period 2004-2015. [Albanesi \(2007\)](#) explains the larger vulnerability of low-income households to inflation: as they tend to hold much of their wealth in inflation-sensitive assets such as cash, they see their wealth drop by relatively more when inflation increases unexpectedly. Inflation also tends to benefit borrowers relative to savers, thereby leading to a redistribution from the old to the young ([Doepke and Schneider, 2006](#); [Adam and Zhu, 2016](#); [Cardoso, Ferreira, Leiva, Nuño, Ortiz, Rodrigo, and Vazquez, 2022](#)). Finally, differences in consumption bundles also imply that inflation affects households in different ways. Elderly households, for example, are particularly sensitive to inflation in health care expenditures ([Hobijn and Lagakos, 2005](#)). [Charalampakis, Fagandini, Henkel, Oibat, et al. \(2022\)](#) study the inflation difference between the lowest and highest income quintile households in the euro area more broadly. They find that this difference was hovering around zero as long as aggregate inflation was low, but once aggregate inflation was rising increased to a staggering 1.9 percentage points in September 2022, mainly driven by the different consumption shares related to energy and food. Similarly, [Pallotti, Paz-Pardo, Slacalek, Tristani, and Violante \(2023\)](#) find that households with lower expenditures suffered the largest losses during the recent period of high inflation. In line with this evidence, central banks have consistently argued that stable and low inflation is their best contribution to containing inequality. This is in particular relevant because consumers tend to be unaware of how inflation erodes their nominal asset and debt positions ([Hackethal, Schnorpfeil, and Weber, 2022](#)).

The paper proceeds as follows. Section 2 contains a description of the data that underlies the analysis and develops the econometric framework that we will use. In particular, we provide first evidence regarding the differences in inflation between income groups and validate that our econometric framework properly identifies the effects of monetary policy on inflation. Sections 3 and 4 report and discuss the empirical results, focusing on the response of inflation differentials to monetary policy shocks in Section 3, and providing results for the differences in shopping behavior and its differential response to monetary policy shocks in Section 4. Section 5 concludes.

2. DATA AND ESTIMATION OF THE EFFECTS OF MONETARY POLICY

2.1. Income-specific inflation based on the Household Budget Survey and official inflation statistics. The Harmonised Index of Consumer Prices (HICP) measures the prices of consumer goods and services acquired by households. Prices for a wide array of products are collected on a monthly basis and aggregated to a single index by means of a weighted average. The weight assigned to each product is the ratio of its consumption over total consumption in the economy.

However, individual households experience different inflation rates according to their specific consumption baskets. At a first level, households' consumption differs across broad product categories (e.g. food vs package holidays). But also within a particular broad category, purchases differ – across narrower product categories (e.g. meat or vegetables), across product types within these categories (e.g. pork), across products (e.g. pork sausages) or across items (e.g. a specific flavor from a given brand, identified by a specific barcode). Our first measure of inflation along the income

distribution reflects differences in consumption baskets, but disregards potential differences in the prices paid for the various barcode items.

For this first measure, we construct income-specific inflation rates by combining HICP data with expenditure data from the Household Budget Survey (HBS). Eurostat's HBS provides information on expenditure broken down by consumption purpose for a large sample of households in the European Union. The consumption purposes follow the Classification of Individual Consumption According to Purpose (COICOP) published by the United Nations' Statistics Division. Additionally, a series of socio-demographic characteristics plus the households' income are available from the survey participants. The survey has been conducted every five to six years since 1988. For the gap years between surveys, we linearly interpolate the expenditure data.

The measure of price changes comes from the HICP, broken down by the same COICOP classification as done in the HBS, therefore allowing for an appropriate matching. The disaggregation along COICOP categories is done at the two-digit level, which results in 12 categories.¹ We use seasonally adjusted series, based on [Bańbura and Bobeica \(2020\)](#). Linking the two data sources, we construct monthly inflation rates for different income groups i as

$$(1) \quad \pi_{t-1,t}^i = \sum_{c=1}^{12} \frac{HICP_{c,t}}{HICP_{c,t-1}} \times \frac{C_{c,t}^i}{C_{total,t}^i} - 1,$$

where $HICP_{c,t}$ is the HICP for COICOP category c in month t , $C_{c,t}^i$ is the total consumption of households in the income quintile i for COICOP category c and month t , and $C_{total,t}^i$ is the total consumption of households in the income quintile i and month t (across all COICOP categories).

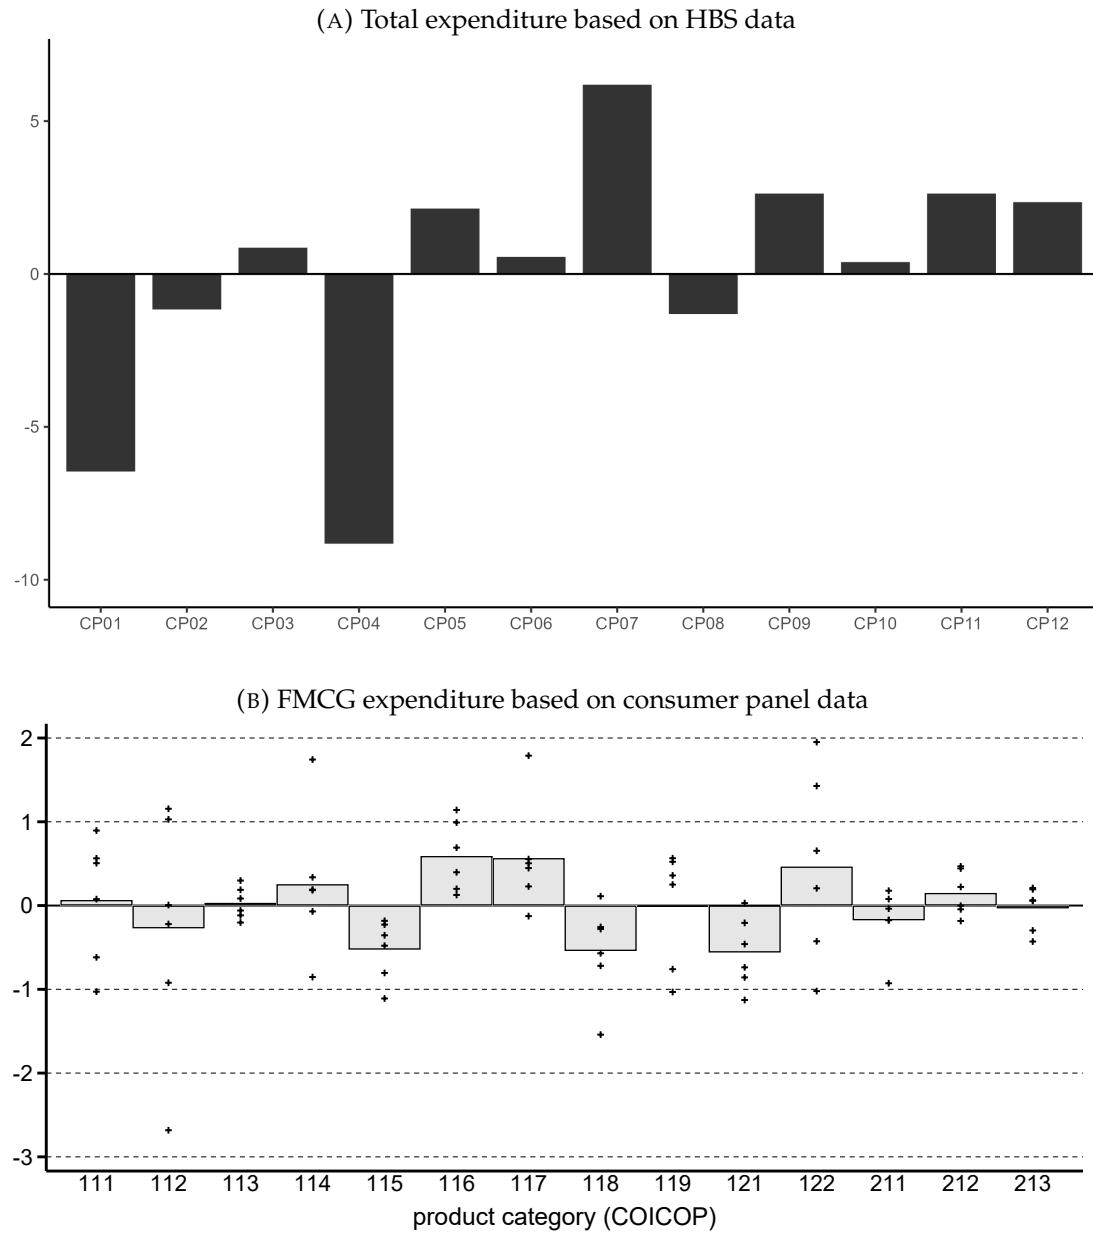
In our subsequent analysis, we will focus on the top and bottom income groups as defined by the highest and lowest income quintiles in the HBS.

The upper panel of Figure 1 provides an overview of the *differences* in the consumption shares between the lowest and the highest income quintile, for the euro area aggregate, and averaged over all HBS surveys. It is apparent that the largest differences arise from consumption patterns related to COICOP categories 01 (food and non-alcoholic beverages) and 04 (housing, water, electricity, gas and other fuels), where low-income households report relatively larger consumption shares, and to COICOP category 07 (transport), which is considerably more important for high-income households.

While this pattern is broadly stable across countries, there are some notable differences, as can be seen in Appendix Table B1. For instance, the difference in food consumption (COICOP 01) between low- and high-income households is above average in Spain and Italy, and relatively more muted in the Netherlands.

¹These are: food and non-alcoholic beverages (01), alcoholic beverages, tobacco and narcotics (02), clothing and footwear (03), housing, water, electricity, gas and other fuels (04), furnishings, household equipment and routine household maintenance (05), health (06), transport (07), information and communication (08), recreation and culture (09), education (10), restaurants and hotels (11) and miscellaneous goods and services (12).

FIGURE 1. Differences in expenditure shares between high- and low-income households



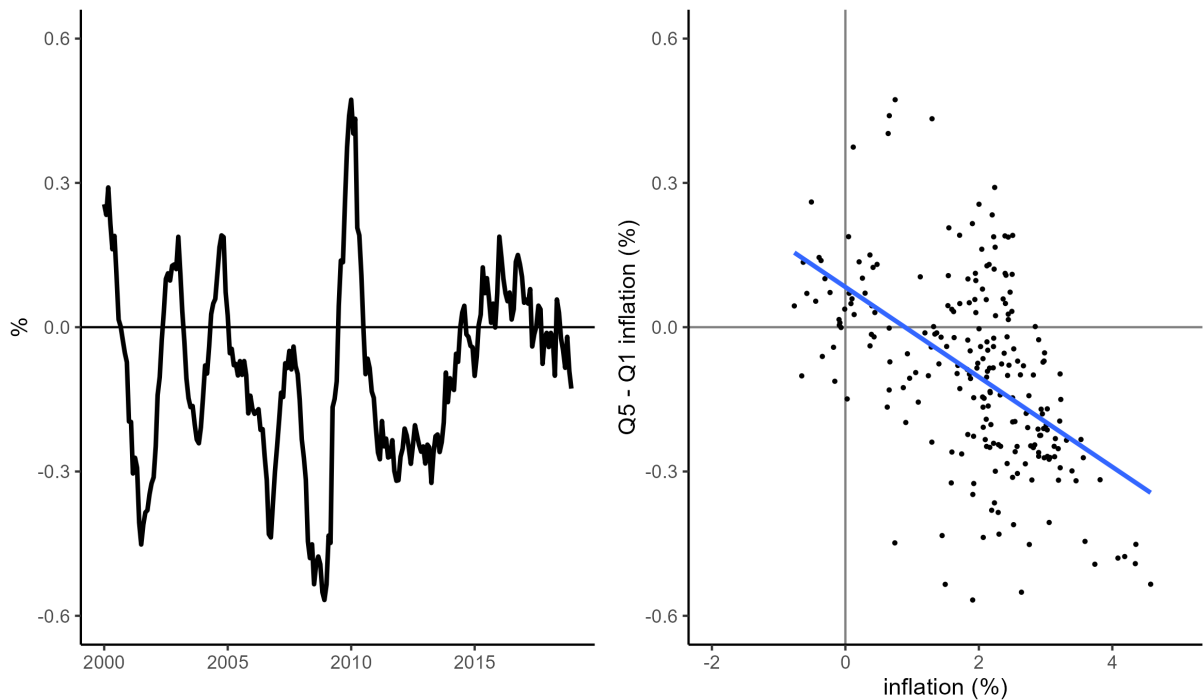
Note: The figure shows the share of expenditure allocated to each COICOP category of high-income households minus the corresponding share of low-income households. Numbers are in percentage points. *Upper panel:* Difference between highest income quintile households and lowest income quintile households. Euro area. HBS expenditure shares for full lineup of COICOPs (01, ..., 12). Shares are relative to the total household expenditure covered by the HBS. Average across all HBS waves (1999, 2004, 2010, 2015). *Lower panel:* Difference between the top income group and bottom income group in the respective consumer panel. Six countries (BE, DE, ES, FR, IT, NL). COICOPs 01.1, 01.2, and 02.1. The bars show the six-country average, whereas the markers show the shares for individual countries. Shares are relative to the household FMCG expenditure covered by the consumer panel. Sample size varies by country (2005/2012–2018), as shown in Appendix Table A1.

To get a sense of the total differences in consumption shares, we can calculate the sum of the positive differences across all COICOP categories (which equals the absolute value of the sum of all negative differences). This sum amounts to 18% in the euro area on average, and varies from around 14% in Spain, France and the Netherlands to above 20% in Germany and Italy.

A look at the evolution of the differences in consumption shares over time (in panels B and C of Appendix Table B1) reveals that the difference in COICOP 04 (housing and utilities) has grown

considerably over time, and has done so in each individual country. In contrast, the difference in the share of food has become smaller in all countries but Italy. Differences in the expenditure share of transport have been stable overall, increasing in some and decreasing in other countries. Overall, consumption patterns have become more unequal: differences have been increasing over time, as judged by the sum of the absolute differences across all COICOP categories. Only in France have they been broadly stable.

FIGURE 2. Differences in euro area HICP inflation between high- and low-income households



Notes: The left panel shows year-on-year HICP inflation of the highest income quintile households in the euro area minus year-on-year HICP inflation of the lowest income quintile households in the euro area. The right panel plots this difference measure against the level of year-on-year HICP inflation. Each dot represents one month.

As a consequence of these different spending shares, also the inflation rates differ between income groups. Figure 2 shows differences in euro area inflation between high- and low-income households, covering the full time sample that we use in the estimations, namely January 2000 until December 2018. The left panel reports the time series of this inflation differential, and the right panel reports a scatter plot of the inflation differential against the level of HICP inflation. Country-specific plots of the inflation differential are provided in Appendix Figure B1. Inflation differences are sizable. They range from approximately -0.6% to around +0.6%, and there are interesting patterns. In particular, as shown in the right panel of Figure 2, inflation for the low-income households is higher than inflation for the high-income households when inflation is high. The opposite is true when inflation is low. This suggests that inflation of the low-income households is relatively more volatile, and that the gap has a cyclical pattern – both in the euro area as well as in the individual countries. On average over our time sample, inflation for low-income households was 0.1 p.p. larger than inflation for

high-income households. This suggests that in this period, the price level has increased by 2.5 p.p. more for low-income households.

The HBS/HICP-based income-specific inflation data are similar in spirit to the US data used in [Cravino et al. \(2020\)](#), who also combine data on expenditure shares for households with category-specific price indices.

2.2. Income-specific inflation based on a household panel. Neither the data underlying [Cravino et al. \(2020\)](#) nor our HBS/HICP-based dataset allow the construction of household-specific inflation. The prices in both datasets are aggregated national averages. As they are not recorded at the barcode item level, they are already an index themselves. By construction, such price indices cannot account for potential differences in the prices paid between households groups. Addressing this shortcoming requires a panel of consumer spending by household type and product category.

For this reason, we turn in the second part of the analysis to a household panel, provided by GfK and Kantar. The panelists record information about their purchases, including the transaction date, the product's barcode, its price and the quantity purchased. In addition, some socio-demographic information on the purchasing household is available, thus allowing us to compare households along the income distribution.²

Using this panel, we track both differences in the consumption basket and differences in the prices that households pay for the same barcode item. The unique barcode item identifier allows for very granular comparisons of consumption baskets. The transaction-nature of the household panel also allows for a faster update of expenditure weights, overcoming a main limitation of the HBS consumption shares.

The main disadvantage relative to HBS/HICP-based income-specific inflation rates is the more limited product scope. Participants in our household panel report only their purchases of fast moving consumer goods (FMCG), i.e. the products typically sold in supermarkets. It is the high frequency ("fast moving") at which these products are purchased, which makes it possible to construct time series in a household panel of limited size. Because the most frequently purchased products are food and beverage items, we restrict our analysis of these data further to food (COICOP 01.1), non-alcoholic beverages (01.2), and alcoholic beverages (02.1). This part of the analysis therefore excludes some other categories that are quantitatively important and show large differences across income groups, such as housing, utilities and fuels (COICOP 04) and transportation (COICOP 07).

Nevertheless, it is fair to say that FMCG are a key area of household heterogeneity, for several reasons. First, they have a relatively large share in the overall consumption basket (they account for around 15% of consumption across euro area households) and according to the HBS, their consumption shares show large differences across income groups (17.3% for the first income quintile vs 12.8% for the fifth income quintile). Second, these are categories where product differentiation is most prominent. Households can choose FMCG from a sheer endless menu of differentiated varieties and

²The income information in the household panel is provided as country-specific income bins. See appendix [A.2](#).

brands. Other components of goods consumption might be equally differentiated but constitute a smaller share of consumption. The consumption of clothing, footwear, furnishings, and household items altogether, for example, adds up to only half of the expenditure on food and beverages in the euro area. Therefore, FMCG is likely to provide general insights about the share of consumption comprising highly differentiated goods, which can potentially extend even to differentiated services, including health, communications, recreation, restaurants and culture. The energy component of consumption, in contrast, is a rather homogeneous good. Households differ considerably in how much they spend on energy as a share of their income, but likely face similar energy prices and inflation. This latter component of inflation heterogeneity is therefore well captured with the HBS/HICP approach described earlier.

Using these household panel data, we can study differences in consumption baskets *within* categories in more granularity than with the HBS data. The lower panel of Figure 1 zooms in on the food and beverage categories. The bars show again the expenditure share difference between the top and bottom income groups. Overall, high-income households in our sample spent a higher share of their total FMCG expenditure on fruits (01.1.6) and vegetables (01.1.7), and in some countries also on mineral waters, soft drinks, and juices (01.2.2). Low-income households instead spent a relatively higher share on oils and fats (01.1.5), sugar, jam, honey, chocolate and confectionery (01.1.8), and on coffee, tea and cocoa (01.2.1). The six-country average masks higher differences within individual countries, as shown by the markers.

The construction of inflation rates for income groups from the GfK/Kantar household panel is explained in detail in Appendix A. In a nutshell, we aggregate the recorded data to generate prices paid by each income group for the specific barcode items. Let us denote the price per unit of a given barcode item b in a given (shopping) transaction s paid by a member of income group i by \tilde{p}_{bs}^i and the quantity purchased by \tilde{x}_{bs}^i . We treat all shopping transactions s during month t by households belonging to income group i as if it was done by one household representative of group i and calculate the quantity-weighted average price paid for barcode item b by the households in this group during month t as

$$(2) \quad p_{bt}^i = \frac{1}{x_{bt}^i} \sum_{s \in S(b,i,t)} \tilde{p}_{bs}^i \tilde{x}_{bs}^i,$$

where $x_{bt}^i = \sum_{s \in S(b,i,t)} \tilde{x}_{bs}^i$ and the set $S(b,i,t)$ consists of all transactions of barcode item b by members of income group i during month t .

Using these prices, we then generate month-on-month inflation indices. Because of very volatile and seasonal consumption patterns at the household and income group level, we use a rolling twelve-month average for quantities. That is, a Laspeyres index at time t (based on the price pair $p_{b,t}^i$

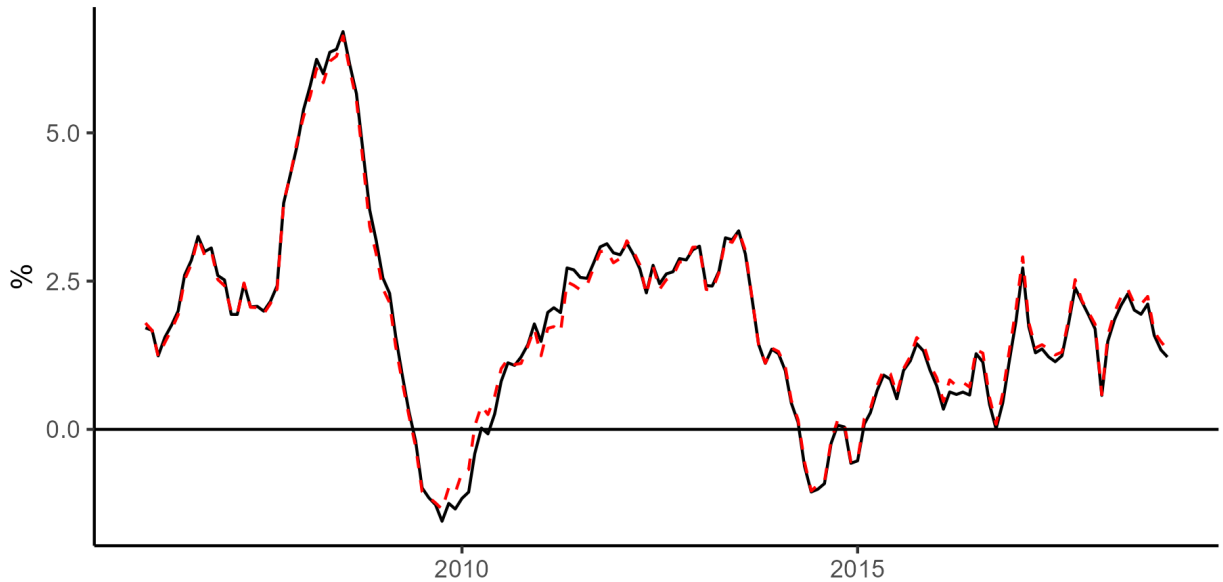
and $p_{b,t-1}^i$) is based on an average $\bar{x}_{b,t-1}^i$ across the 12 quantities $x_{b,t-1}^i, \dots, x_{b,t-12}^i$. The Laspeyres inflation rate of income group i for the set of barcode items $B(i, t-1)$ over a one-month period ending with t is

$$(3) \quad \pi_{t-1,t}^i = \frac{\sum_{b \in B(i,t-1)} p_{b,t}^i \bar{x}_{b,t-1}^i}{\sum_{b \in B(i,t-1)} p_{b,t-1}^i \bar{x}_{b,t-1}^i} - 1.$$

As with the HBS/HICP data, we focus on the top and bottom income groups. In Belgium, France, Germany and the Netherlands, we define income groups directly based on the net income of the household. Because this information is not available for Italy and Spain, we use the reported social class there, to which income is obviously a major contributor. Based on this classification, we can distinguish four or five income groups in each country.³

We have household panel data available until December 2018, which is why we also end our analysis of HBS/HICP data at that point in time. However, the start of the sample period differs across countries. The longest time series are available for Belgium and Germany, where the data start in 2005. In contrast, data for Italy only start in 2012. Full details on the available samples are provided in Appendix A.6.

FIGURE 3. Food and non-alcoholic beverage (COICOP 01) inflation, aggregate of the six largest euro area countries, HICP and GfK/Kantar data



Notes: The figure compares the year-on-year COICOP 01 inflation using HICP (black solid line) with the corresponding inflation as measured by the GfK/Kantar data (red dashed line), each for the six largest euro area countries, aggregated using HICP weights.

³Appendix A.2 describes the information on income available in the household panel in more detail.

As shown in Figure 3 and in line with [Strasser, Messner, Rumler, and Ampudia \(2023\)](#), the inflation rates derived from the transactions reported by the household panel track the aggregate HICP food inflation closely, despite the considerable conceptual differences.

2.3. Monetary policy shocks and the estimation of their effects. In order to identify a causal effect of monetary policy on inflation, it is necessary to control for the response of monetary policy itself to changes in the macroeconomy. This is possible by restricting the analysis to the effect of exogenous monetary policy *shocks*.

[Jarociński and Karadi \(2020\)](#) provide a such a measure of monetary policy that is furthermore purged of possible central bank “information shocks”, i.e. of information that the central bank reveals while it announces its monetary policy decisions. To get at this measure, they study the high-frequency co-movement of interest rates and stock prices in a narrow window around the policy announcement and identify a monetary policy shock through a negative co-movement between interest rates and stock market returns. We use monetary policy shocks ϕ_t identified according to the “poor man’s” method of [Jarociński and Karadi \(2020\)](#). This measure has been used widely in the literature, and therefore allows us to gauge the plausibility of our results.⁴

To capture the effect of monetary policy shocks on inflation heterogeneity, we generate impulse responses based on local projections ([Jordà, 2005](#)). Let P_t denote the Laspeyres inflation index in period t . The cumulative Laspeyres inflation during the period t and $t + h$ is $\pi_{t,t+h} = \log(P_{t+h}/P_t) = \prod_{i=1}^h \pi_{t-1+i,t+i}$.

We study the response of the cumulative inflation $\pi_{t,t+h}$ to the monetary policy shock ϕ_t . These shocks reflect only exogenous monetary policy surprises around policy meetings, which are orthogonal to market expectations. Obviously, a large part of monetary policy is anticipated, i.e. it diffuses into the market between meetings, and is thus not reflected in ϕ_t . To control for the effects of this systematic part of monetary policy, we include as control variable lagged values of the one-year overnight interest rate swap (OIS) rate x_t .

For parsimony, we drop lags of the dependent variable (their inclusion does not alter our results) and group lags of shocks and control variables. Defining $\phi_{a,b} = \sum_{j=a}^b \phi_j$ and $x_{a,b} = \sum_{j=a}^b x_j$ our set of

⁴[Jarociński and Karadi \(2020\)](#) also provide a measure of central bank information shocks. In their application to the ECB, these are shown to raise inflation (in contrast to the monetary policy shocks, which lower inflation). This opposite sign also results when we use the central bank information shocks in our analysis: the inflation differential between high- and low-income households moves in the opposite direction than after monetary policy shocks, for both the HBS/HICP and the GfK/Kantar data. As in [Jarociński and Karadi \(2020\)](#), the effects of the central bank information shocks are estimated with somewhat less precision than those for the monetary policy shocks, but they are found to be statistically significantly different from zero. Results are not reported for brevity.

local projections is

$$\begin{aligned}
 (4) \quad \pi_{cty,t,t+h} &= \alpha_h + \theta_h \phi_t \\
 &+ \gamma_h^{1M} \phi_{t-1} + \gamma_h^{2M3M} \phi_{t-2,t-3} + \gamma_h^{4M12M} \phi_{t-4,t-12} + \gamma_h^{2Y} \phi_{t-13,t-24} + \gamma_h^{3Y} \phi_{t-25,t-36} \\
 &+ \kappa_h^{1M} x_{t-1} + \kappa_h^{2M3M} x_{t-2,t-3} + \kappa_h^{4M12M} x_{t-4,t-12} + \kappa_h^{2Y} x_{t-13,t-24} + \kappa_h^{3Y} x_{t-25,t-36} \\
 &+ \delta_{cty} + \epsilon_{cty,t},
 \end{aligned}$$

$\forall h \in \{0, \dots, 48\}$ and with ϵ_t i.i.d.. $\pi_{cty,t,t+h}$ denotes inflation in country cty . δ_{cty} stands for country fixed effects, as the estimation is done for the panel of the six largest euro area countries, given that these are the data that are available for both types of inflation series.⁵ The coefficients are estimated by ordinary least squares. Standard errors are robust to heteroscedasticity and to arbitrary forms of cross-sectional dependence, using Driscoll-Kraay corrections.

Figure 4 shows the response of aggregate HICP inflation to a monetary policy shock in the left panel, and the response of the COICOP 01 category inflation in the right panel. As in all subsequent charts, we report the response over 48 months to a surprise tightening of 10 basis points by the ECB. The solid line represents the estimated coefficients, whereas the dark and light grey areas cover the 1 and 1.645 standard deviation confidence intervals. Results are in line with the conventional findings in the literature. We find that inflation declines in response to a monetary policy tightening. While the decline is estimated to be only marginally significant for HICP inflation, the magnitude is in the same ballpark as the results reported in Jarociński and Karadi (2020). As shown in Appendix Figure B2, the pattern and magnitudes are similar across countries, but there is some heterogeneity. For instance, the response is relatively larger in Italy and Belgium, but given that the confidence bands are overlapping, these differences are likely not statistically significant.

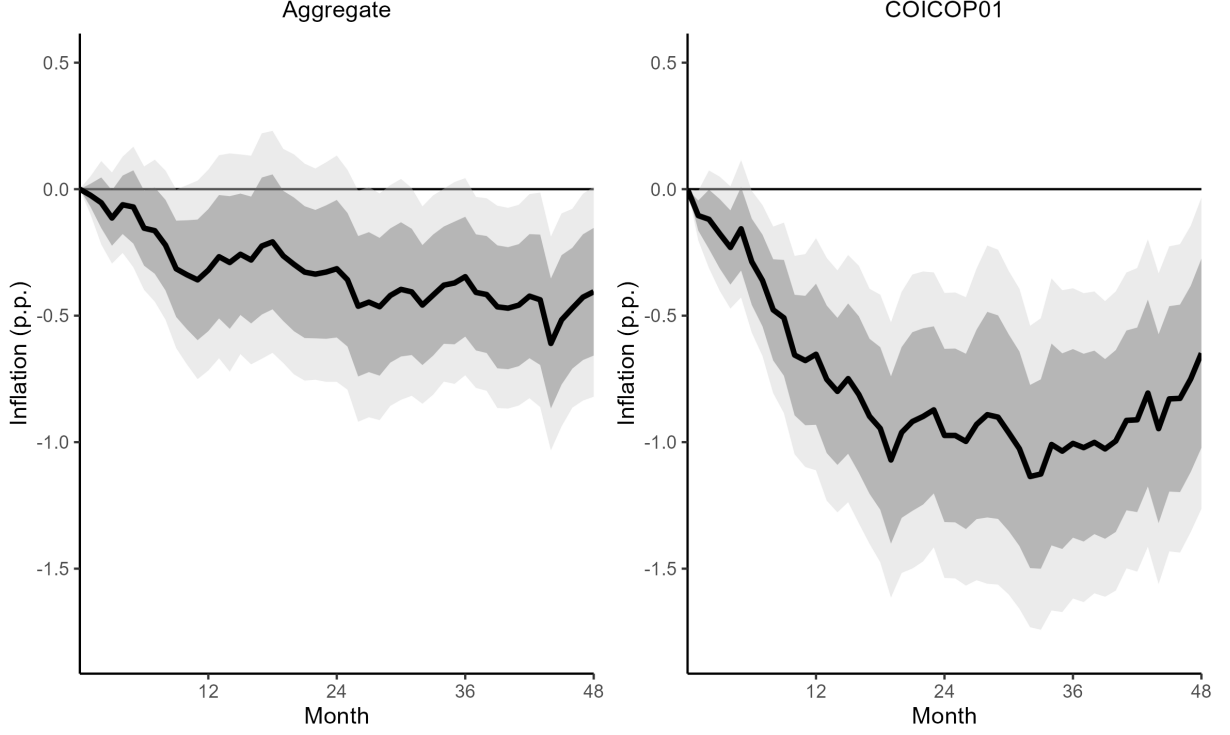
The impulse responses for food and non-alcoholic beverages inflation are estimated at a higher level of statistical significance, both when estimated for COICOP 01 category inflation alone (as in the right panel of Figure 4) or for FMCG inflation based on the household panel data (not shown for brevity). Note also that the magnitude is considerably larger than for overall inflation. Both the stronger response and the tighter estimation are in line with the notion that food inflation is strongly responsive to monetary policy, given its relatively large import content and therefore its more direct exposure to the exchange rate channel.

Appendix Figure B3 repeats the analysis for the COICOP 01 category, for a sample starting in 2005 (to get closer to the categories and the time sample covered in the household panel).⁶ Shortening the sample to 2005, statistical significance is reduced, but the overall pattern and the estimated magnitudes remain similar.

⁵Impulse responses for the euro area aggregate when available (i.e. for the HBS/HICP data) are very similar.

⁶Given the lag structure in specification (4) the estimation using monthly inflation starts in May 2005. Accordingly, the dataset includes monetary policy shocks and controls back until 2002 to account for the lags in equation (4).

FIGURE 4. Response of overall HICP and COICOP category 1 inflation to a monetary policy shock, aggregate of the six largest euro area countries



Notes: The figure shows the impulse response of aggregate HICP inflation (left panel) and HICP COICOP category 1 inflation (right panel) to a 10 basis points tightening ECB monetary policy shock, estimated for a panel of the six largest euro area countries. The dark and light grey areas indicate 1 and 1.645 standard deviation confidence intervals, respectively.

3. THE EFFECT OF MONETARY POLICY ON INFLATION HETEROGENEITY

Following these checks, we are confident that our measure of monetary policy shocks and our estimation setup allow us to trace the effects of monetary policy on inflation, such that we can now start looking into the differential effects along the income distribution.

Let us denote the cumulative Laspeyres inflation of the income group i in country cty over an h months period by $\pi_{cty,t,t+h}^{ICi}$, where "IC1" denotes the lowest, and "IC5" the highest income group in the respective countries. The local projection for inflation differentials is then, in analogy to equation (4),

$$\begin{aligned}
 (5) \quad & \pi_{cty,t,t+h}^{IC5} - \pi_{cty,t,t+h}^{IC1} \\
 &= \alpha_h + \theta_h \phi_t \\
 &+ \gamma_h^{1M} \phi_{t-1} + \gamma_h^{2M3M} \phi_{t-2,t-3} + \gamma_h^{4M12M} \phi_{t-4,t-12} + \gamma_h^{2Y} \phi_{t-13,t-24} + \gamma_h^{3Y} \phi_{t-25,t-36} \\
 &+ \kappa_h^{1M} x_{t-1} + \kappa_h^{2M3M} x_{t-2,t-3} + \kappa_h^{4M12M} x_{t-4,t-12} + \kappa_h^{2Y} x_{t-13,t-24} + \kappa_h^{3Y} x_{t-25,t-36} \\
 &+ \delta_{cty} + \epsilon_{cty,t}.
 \end{aligned}$$

As before, the coefficients are estimated by ordinary least squares. Standard errors are robust to heteroscedasticity and to arbitrary forms of cross-sectional dependence, using Driscoll-Kraay corrections.

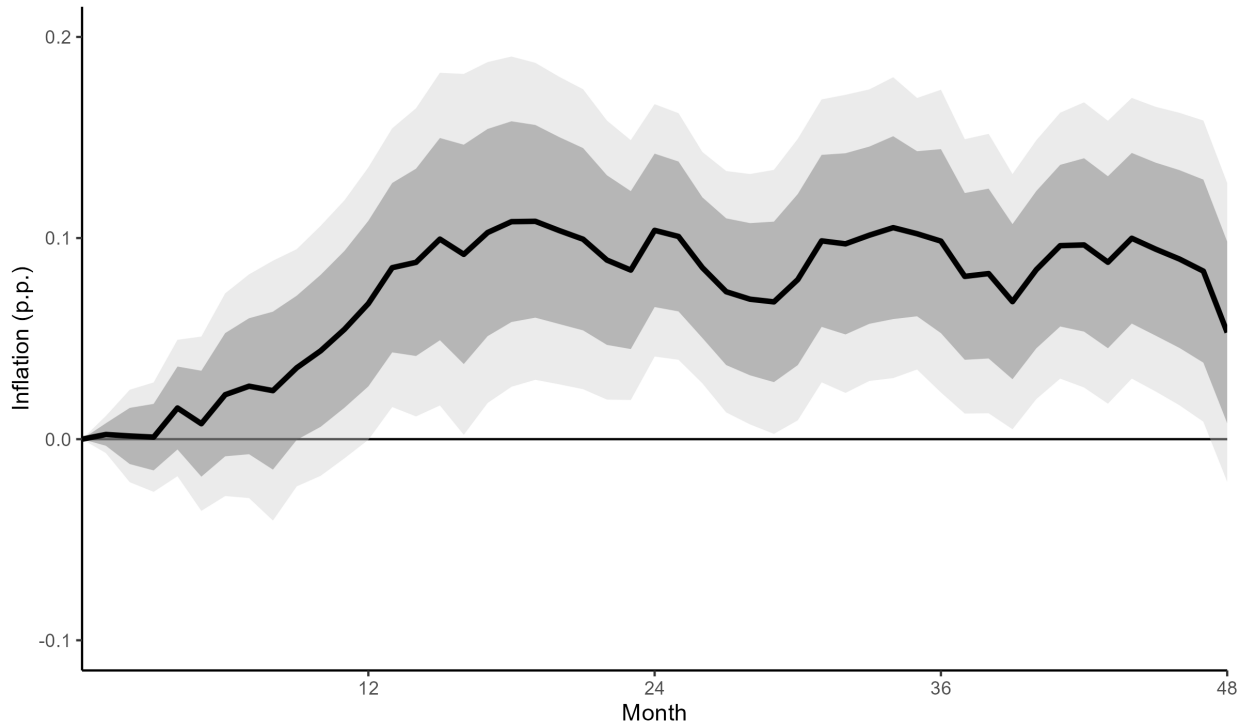
3.1. Broad product coverage, common prices. The first set of results, based on the HBS/HICP inflation series for the six largest euro area countries and equation (5), is reported in Figure 5. It is apparent that – before substitution – inflation for the high-income group responds by less than inflation for the low-income group. The difference increases for slightly more than one year and remains stable thereafter, stabilising at around 0.075 percentage points. In other words, one year after a 10 bp monetary policy tightening both high- and low-income households experience lower inflation (results for inflation of each income group separately confirm this), but the inflation of low-income households declined relatively more. This effect is largely permanent. To put this into perspective, while the shape of the impulse response looks different from the one reported in Cravino et al. (2020), the magnitude is relatively close. In their figure 6, Cravino et al. (2020) compare the high-inflation households (defined as the top 1%) with the middle 20% and with aggregate inflation. They find that the differential keeps increasing over the entire 48 months plotted in their impulse response function, that it becomes statistically significant after around 3 years and that it eventually amounts to around 0.5 percentage points in response to a 100 basis point tightening. The magnitude of the response is therefore broadly in line with our findings (given that the size of the shock in Cravino et al. (2020) is 10 times as large as the one applied in our estimation).

3.2. FMCG only, income-specific prices. Moving on to income-specific inflation based on the GfK-Kantar household panel, we can test to what extent inflation differentials arise once we can control not only for differences in the consumption shares (within the food and beverages category), but also for differences in the prices paid. Here, results are markedly different. As can be seen in Figure 6, we find a *negative* differential, meaning that Laspeyres inflation for high-income households responds by relatively *more* than inflation for low-income households.⁷ This is the exact opposite sign than was found in Cravino et al. (2020) and with our HBS/HICP-based data. These findings are in line with Argente and Lee (2021) who have highlighted that high-income households have more scope to change their shopping behaviour or to engage in product quality substitution than low-income households. The effect is furthermore sizable, with a differential in the same order of magnitude (but the opposite sign) as reported with our HBS/HICP-based data. Statistical significance is slightly weaker than for the previous result. One reason could be that these results are based on a shorter sample period that is dominated by unconventional monetary policy, where our measures of monetary policy shocks are relatively small. As a matter of fact, also the results with the HBS/HICP data become less significant for the shorter sample.⁸

⁷Impulse responses for inflation by income group confirm that these are indeed declining in response to a monetary policy tightening.

⁸This is shown in Appendix Figure B4.

FIGURE 5. Response of the inflation differential to a monetary policy shock, panel of the six largest euro area countries



Notes: The figure shows the impulse response of the inflation differential between high- and low-income households to a 10 basis points tightening ECB monetary policy shock, estimated for a panel of the six largest euro area countries and based on HICP data. The dark and light grey areas indicate 1 and 1.645 standard deviation confidence intervals, respectively.

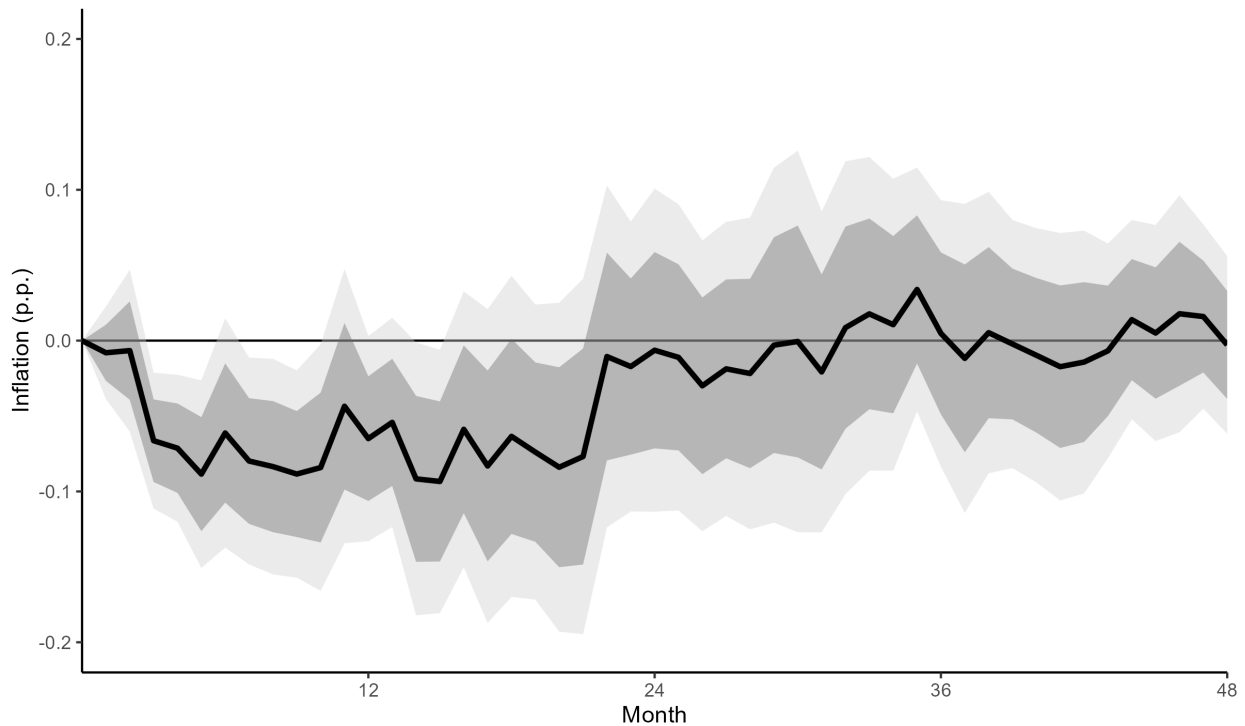
4. THE DIFFERENTIAL RESPONSE OF SHOPPING BEHAVIOR TO MONETARY POLICY SHOCKS

Shopping behavior is known to depend on households' financial situation, implying that low-income households might experience inflation rates substantially different from those experienced by high-income households ([Kaplan and Menzio, 2016](#)). There are different ways in which shopping behavior differs, as outlined by [Griffith et al. \(2009\)](#). For instance, low-income households could purchase the same barcode items but take better advantage of sales, buy in bulk or purchase them from lower-priced outlets. Alternatively, they could buy fewer branded items and more lower-priced unbranded items, i.e. they could purchase different barcode items within the same product category. In this section, we first study to what extent shopping behavior differs across high- and low-income households, before we analyze the differential response to monetary policy shocks.

4.1. Differences in shopping behaviour. One way to understand differences in shopping behavior is to construct inflation based on Paasche indices and to compare these to the Laspeyres indices we have studied so far. Paasche indices use the consumption basket in the *current period* and compare its price in the current period with its price in earlier periods. In contrast, Laspeyres indices use the consumption basket of an earlier *base period* and follow the price for this basket over time.

In other words, inflation based on Paasche indices allows for product substitution, whereas in our preceding analysis of Laspeyres inflation, substitution was ruled out on impact (for both, the

FIGURE 6. Response of the FMCG inflation differential to a monetary policy shock, panel of the six largest euro area countries, Laspeyres index



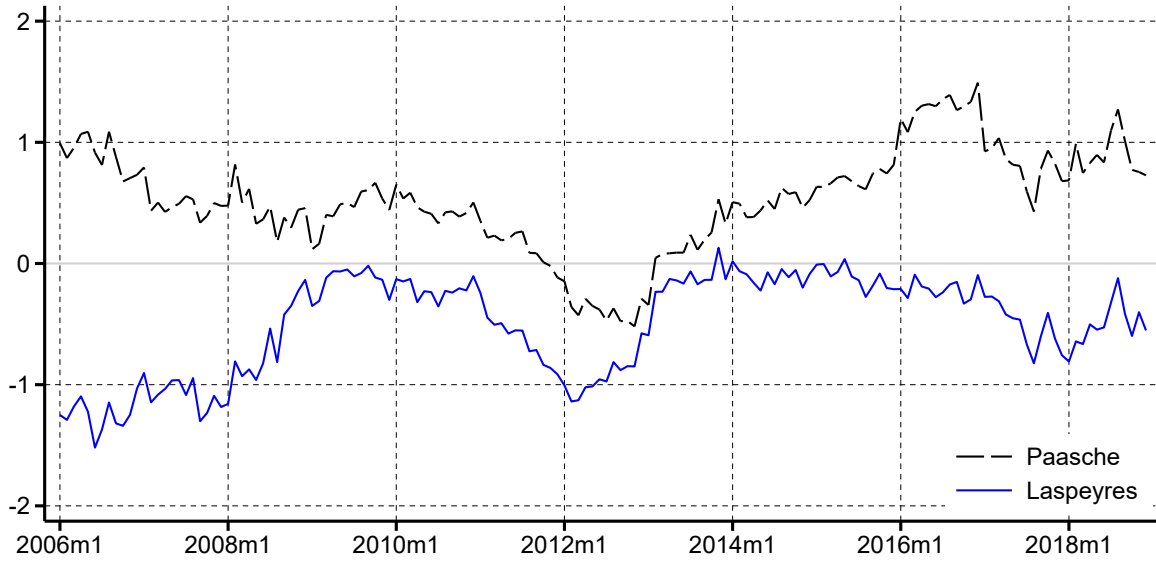
Notes: The figure shows the impulse response of the FMCG (Laspeyres) inflation differential between high- and low-income households to a 10 basis points tightening ECB monetary policy shock, estimated for a panel of the six largest euro area countries and based on GfK/Kantar data. The dark and light grey areas indicate 1 and 1.645 standard deviation confidence intervals, respectively.

HBS/HICP and the GfK/Kantar inflation data).⁹ The construction of the Paasche inflation series is explained in Appendix A.5. Here, it suffices to say that in analogy to the construction of the Laspeyres inflation rates we use a rolling twelve-month average for quantities, which generates some persistence in the consumption baskets. In our construction of Paasche inflation we average quantities over the current and future periods, i.e. only over periods after the monetary policy shock.

Figure 7 shows the difference between inflation in the high- and the low-income group for both, Laspeyres (blue line) and Paasche inflation (black dashed line). Laspeyres inflation of high- and low-income households was approximately the same during half of the sample period (2009-2011, 2013-2016), interrupted by three periods during which inflation for low-income households was up to one percentage point per year higher. After substitution, however, the inflation differential changes sign: The Paasche inflation of high-income households has been higher throughout the sample period, especially after 2014. This implies that low-income households substituted more effectively than high-income households towards products with relatively smaller price increases. Thanks to their effort and willingness to substitute towards other products, low-income households end up with a smaller (Paasche) inflation than high-income households.

⁹The effect of substitution enters the chained Laspeyres indices only gradually as the quantity weights underlying the monthly inflation rates are updated.

FIGURE 7. FMCG inflation differential between high- and low-income households, Laspeyres vs Paasche



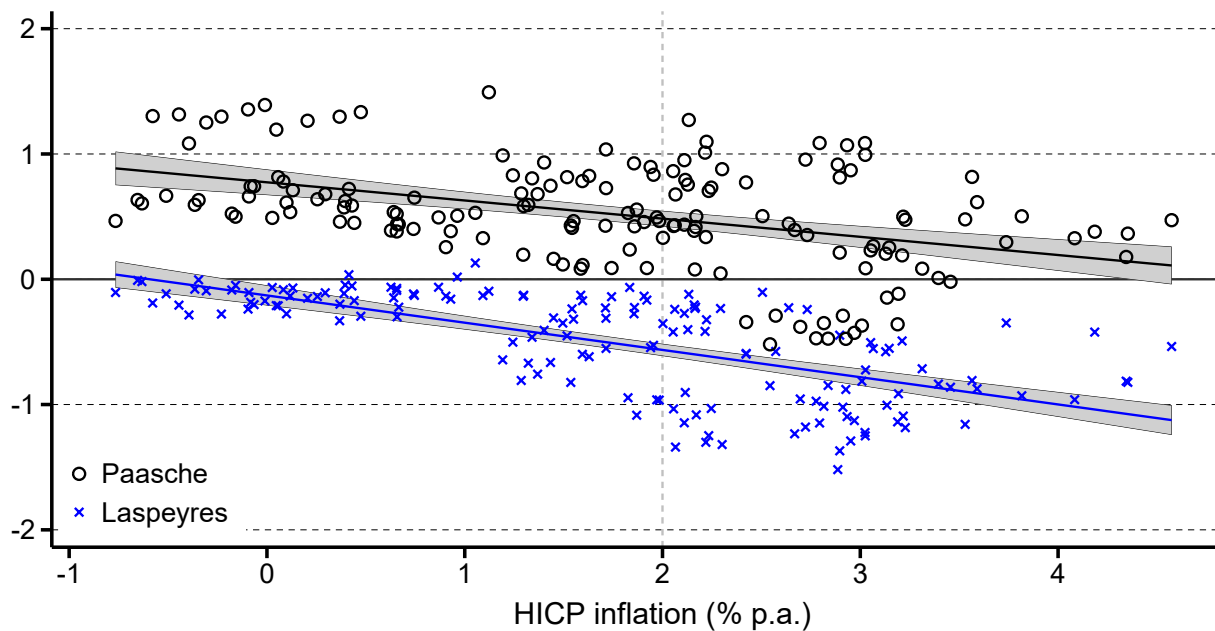
Note: Inflation difference between top and bottom income groups, based on Laspeyres (blue) and Paasche (black) indices in percentage points per year. The underlying inflation indices are the 12-months rolling averages of the six-country weighted average of $\pi_{t-1,t}^i$. COICOPs 01.1, 01.2, and 02.1, based on GfK/Kantar data. The sample grows over time until all six countries are in the sample, as shown in Table A1.

The relative effectiveness of substitution by high- and by low-income households varies over time. In most years (before 2009 and since 2015), the effect of substitution by low-income households considerably exceeds the effect of substitution by high-income households, highlighted by the wide gaps between the two lines in these years. From 2009 until 2013 (i.e. in the aftermath of the global financial crisis and during the European sovereign debt crisis), this gap narrows considerably. Furthermore, during that period the Laspeyres and Paasche inflation differentials move in a synchronized manner, which indicates that both income groups were under similar pressure to substitute towards products getting relatively cheaper in this period.

Another way of understanding the differences in shopping behavior is provided in Figure 8, which contains a scatter plot of the high- minus low-income inflation differential against the contemporaneous HICP inflation, separately for Laspeyres inflation (the blue crosses) and Paasche inflation (the black circles). The figure contains four messages.

First, the blue crosses are below zero. This implies that Laspeyres inflation for high-income households is lower than the one for low-income households. Second, the blue line has a significantly negative slope. At low levels of inflation the difference in Laspeyres inflation between high- and low-income households is effectively zero, while at high levels of inflation Laspeyres inflation of low-income households is substantially higher than the one of high-income households. This is in line with the earlier literature and more recent evidence (Bobasu, Di Nino, and Osbat, 2023), that low-income households are affected more by high inflation than high-income households (at least when measured by Laspeyres inflation, which is the standard measure that the literature has studied).

FIGURE 8. High-income minus low-income FMCG inflation differential versus the level of HICP inflation



Note: Top- minus bottom-income FMCG inflation differential versus the contemporaneous level of euro area HICP inflation, in percentage points per year. Inflation differentials are based on transactions in COICOPs 01.1, 01.2, and 02.1 reported in the GfK/Kantar household panel. Difference of six-country-weighted averages, separately for Laspeyres (blue crosses) and Paasche inflation (black circles). 95% confidence bands of the respective linear regression are shaded in grey.

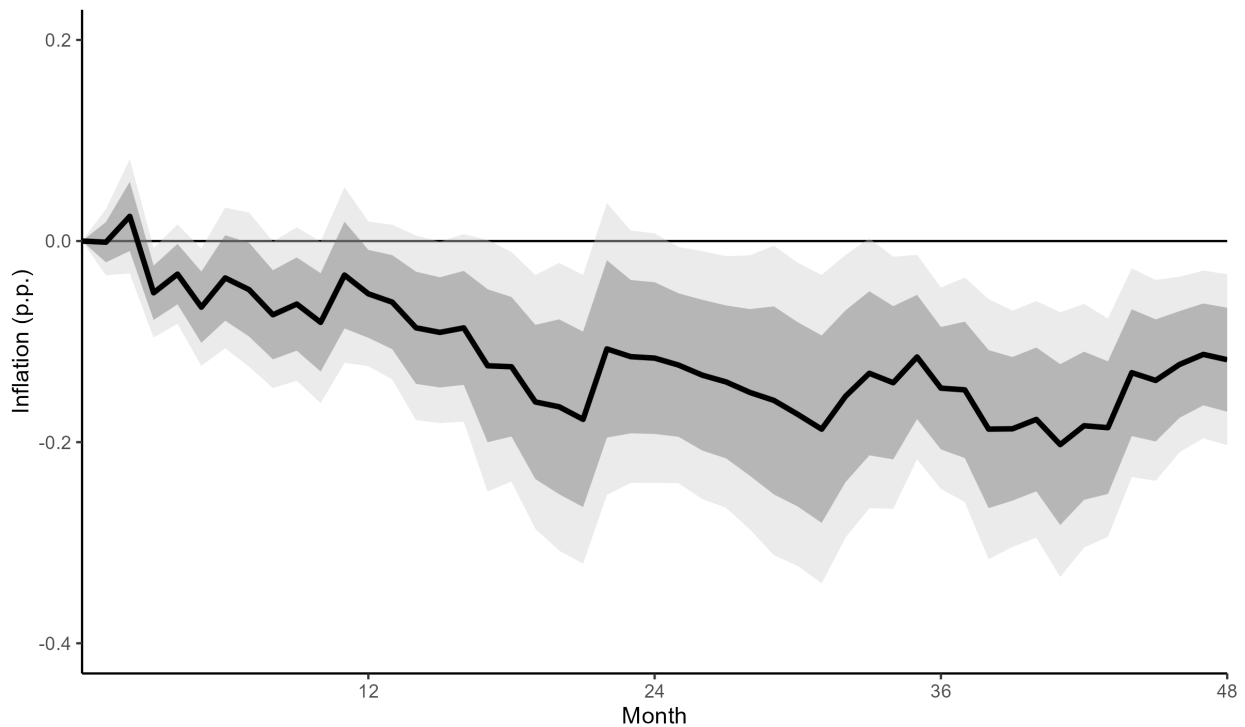
Third, the black circles are above zero, which means that high-income Paasche inflation is *above* low-income Paasche inflation. This shows that low-income households substitute more than their high-income counterparts, and to an extent that they turn a higher Laspeyres inflation into a lower Paasche inflation. While the direction of this effect is well known, its magnitude is surprisingly large. Fourth, also the black line has a negative slope, but that slope is less steep than the slope of the blue line, and significantly so. As inflation increases, low-income households are increasingly unable to offset the higher price increases in their original basket with substitution. Thus, their after-substitution inflation “edge” over high-income households shrinks. When overall HICP inflation reaches 4%, both household groups face about the same Paasche inflation. The differential slopes reflect that while low-income households substitute less and less “successfully” as inflation increases, this is even more the case for high-income households. However, even though low-income households substitute in response to increasing inflation more successfully than high-income households, this is not enough to offset the disproportionately larger price changes in their original (Laspeyres) basket.

4.2. The differential response of shopping behavior to monetary policy shocks. The substantial differences between Laspeyres and Paasche inflation suggest that product substitution is an integral part of shopping behaviour. It is exercised to a different degree by high- and low-income households, and the intensity of product substitution differs over time, and differently so for the two household groups. Applying this to the effect of a monetary policy shock, one could imagine that after a tightening of monetary policy, when interest rates rise and economic prospects worsen, households would

like to cut down on their consumption expenditures. [Argente and Lee \(2021\)](#) propose that in such situations, high-income households have more scope to change their shopping behaviour than low-income households, and do therefore have another margin for adjustment in response to shocks. This could explain why we found in the previous section that the inflation experienced by high-income households responds more to monetary policy.

4.2.1. *Laspeyres vs. Paasche*. To see the extent to which shopping behavior does indeed change in a differential manner between high- and low-income households, the next step in our analysis is to estimate the response of the Paasche inflation differential to monetary policy shocks, and to compare it with the results for the Laspeyres differential.

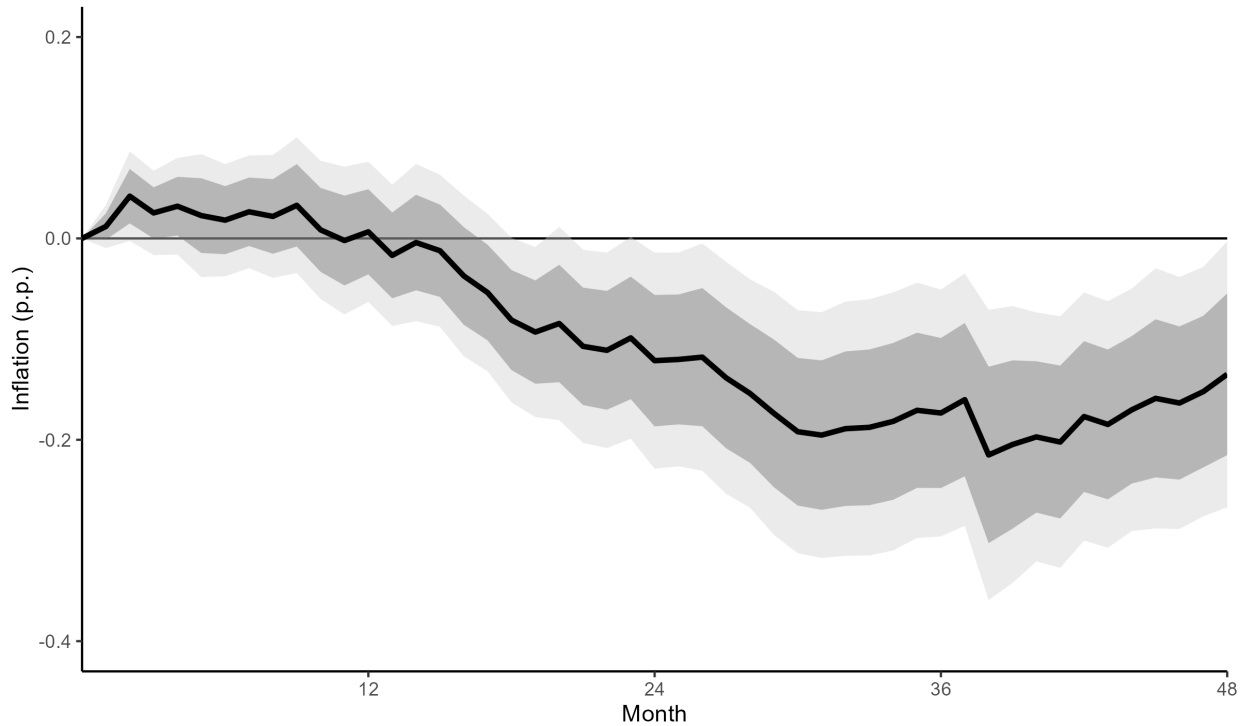
FIGURE 9. Response of the FMCG inflation differential to a monetary policy shock, panel of the six largest euro area countries, Paasche index



Notes: The figure shows the impulse response of FMCG (Paasche) inflation differential between high- and low-income households to a 10 basis points tightening ECB monetary policy shock, estimated for a panel of the six largest euro area countries and based on GfK/Kantar data. The dark and light grey areas indicate 1 and 1.645 standard deviation confidence intervals, respectively.

When we estimate impulse responses for high-income Paasche inflation and low-income Paasche inflation separately, we confirm that these decline after a monetary policy tightening. Figure 9 shows the impulse response of the high-minus-low income Paasche inflation differential, and reveals a pronounced pattern which is qualitatively in line with the results for Laspeyres inflation in the sense that high-income inflation responds by relatively *more* than low-income inflation. There are noteworthy differences, though. First, the results are estimated at higher levels of statistical significance. Second, the timing of the response resembles more closely the one found for HICP inflation, as there is a gradual decline that then levels out after around 18 months. The Laspeyres results, instead, show a

FIGURE 10. Response of the FMCG inflation differential to a monetary policy shock, panel of the six largest euro area countries, difference between Laspeyres and Paasche index



Notes: The figure shows the impulse response of FMCG inflation differential between high- and low-income households and between Laspeyres and Paasche indices, to a 10 basis points tightening ECB monetary policy shock, estimated for a panel of the six largest euro area countries and based on GfK/Kantar data. The dark and light grey areas indicate 1 and 1.645 standard deviation confidence intervals, respectively.

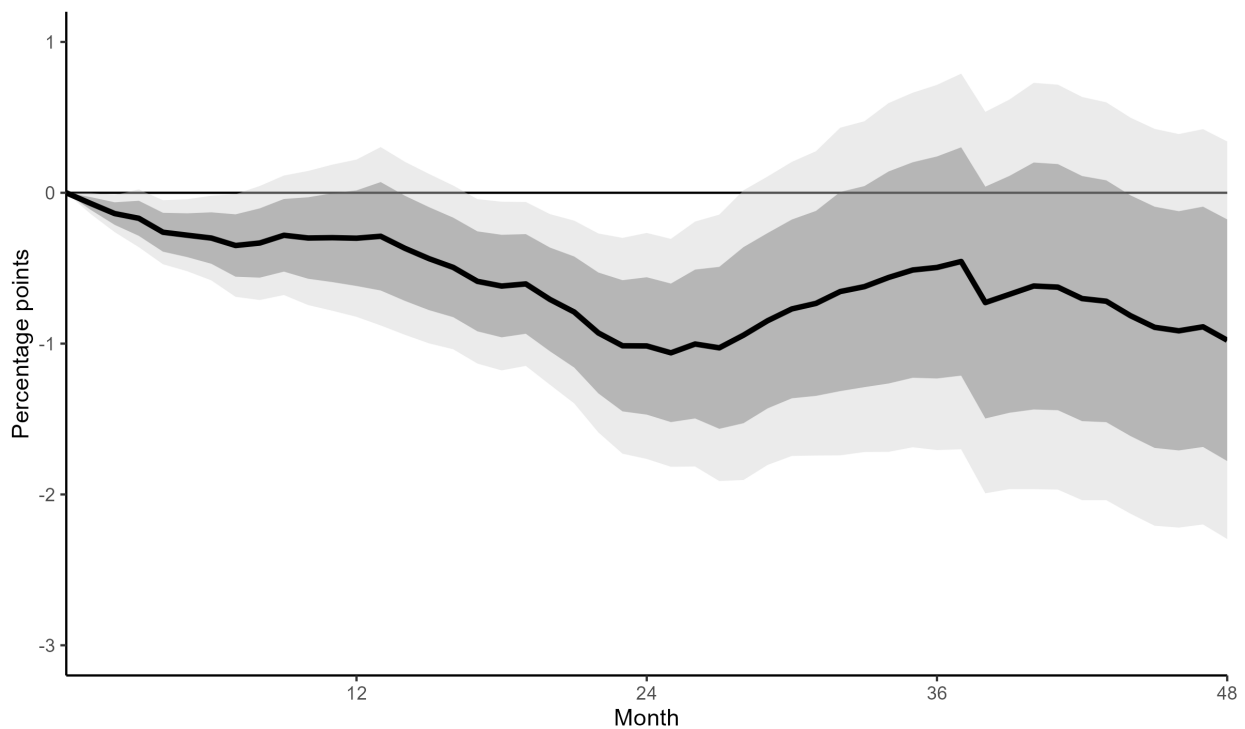
relatively quick response which then turns insignificant after around 20 months. Note that the impulse responses based on Laspeyres and Paasche indices are significantly different after 18 months, as can be seen in Figure 10.

These results are in line with the notion that there is an initial response whereby high-income households change their shopping behaviour relative to low-income households, but still purchase the same barcode items. Only over time, they engage in product substitution and therefore lower their individual inflation rates relative to low-income households also in a more persistent manner.

4.2.2. Quantities purchased. Our granular data also allows for other ways to better understand the changes in shopping behavior. For instance, we can single out to what extent the differences in the impulse responses arise because of a change in quantities purchased. We base our analysis on a decomposition of inflation into an expenditure change component and a quantity change component. The quantity change component is the ratio of the quantity of barcode item b purchased per household in income group i during month t relative to the respective quantity in the previous month, where the quantities in both periods are weighted by the price vector of month t . This quantity change, in turn, contains three parts: first, barcode items purchased by income group i in month $t - 1$ that have not been repurchased by that income group in month t ; second, barcode items that

have been purchased in month t , but had not been purchased in the previous month; and third, repurchased barcode items. Focusing on the third part, Figure 11 reports how the quantities of repurchased barcode items respond, as usual looking at the difference between high and low-income households (after confirming that the quantities decline for both, high and low-income households separately).¹⁰ The relative decline in quantities for high-income households corroborates the findings obtained with the Paasche index: high-income households reduce their consumption of those barcode items that they had purchased previously by relatively more. Also the timing of the impulse response function is in line with the result that substitution proceeds with some delay.

FIGURE 11. Response of the differential quantities purchased to a monetary policy shock, panel of the six largest euro area countries



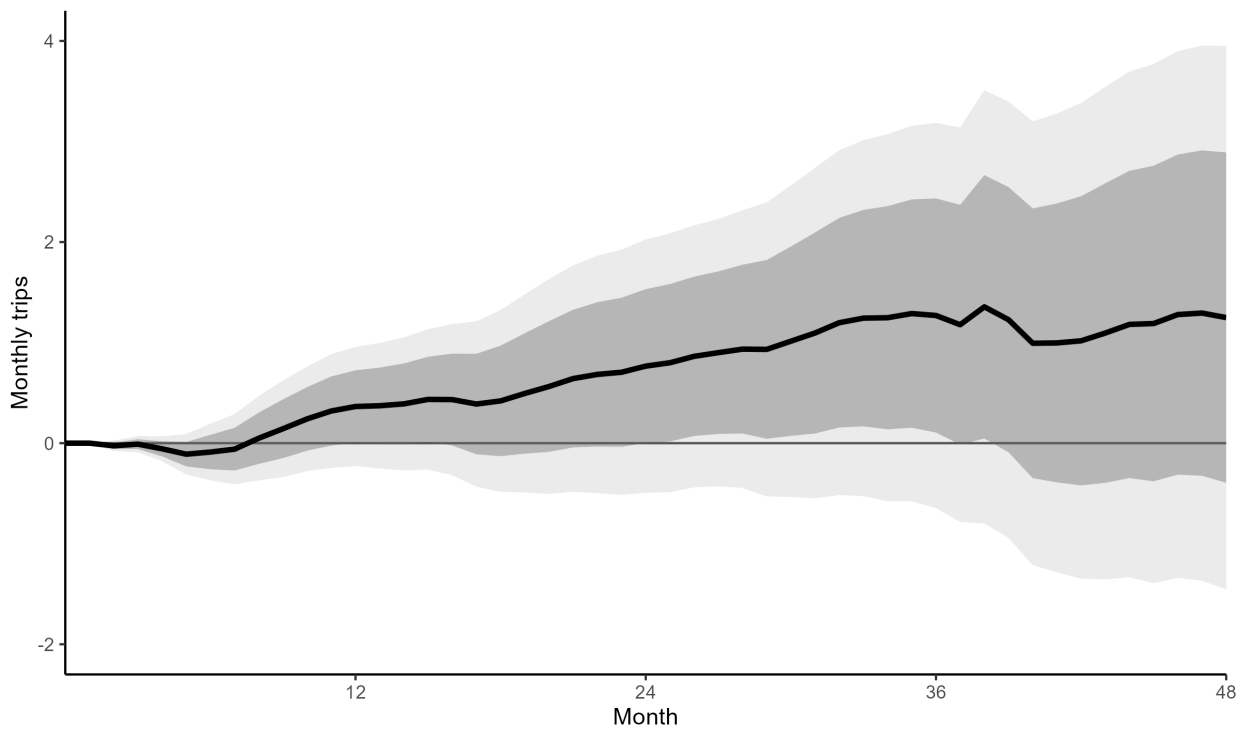
Notes: The figure shows the impulse response of the difference in the quantities purchased between high- and low-income households to a 10 basis points tightening monetary policy shock, estimated for a panel of the six largest euro area countries and based on household panel data. Quantities purchased refer to repurchased products only, weighted by period t prices. See Appendix A.7 for details. The dark and light grey areas indicate 1 and 1.645 standard deviation confidence intervals, respectively.

4.2.3. *Shopping intensity.* Our next test studies households' shopping intensity, where we analyse how the number of shopping trips by the two household groups responds to a monetary policy tightening. (Recall that [Kaplan and Menzio \(2016\)](#) find that shopping time is responsive to unemployment shocks, a finding which is in line with the one of [Krueger and Mueller \(2010\)](#) that unemployed households spend considerably more time on shopping than employed households, not only in the United States but also in Europe.)

¹⁰Results for *all* quantity changes, i.e. not only the repurchased barcode items, are similar, in magnitude and in the overall shape of the impulse response. They are estimated at slightly higher levels of statistical significance, and the impulse response is less smooth.

We conjecture that in response to a monetary policy tightening, when the economic outlook worsens, households increase their shopping intensity. This is in line with a large literature on shopping time, which has shown that during the Great Recession the frequency of sales doubled (Kryvtsov and Vincent, 2021), households spent more time on shopping (Aguiar, Hurst, and Karabarbounis, 2013), purchased more on sale, larger sizes, and generic products, and increased coupon usage and shopping at discount stores (Nevo and Wong, 2019). To test whether this response differs across income groups, we count the number of unique day-store visits per month for each household.¹¹ Based on this, we calculate the average number of shopping trips per month per household in a given income group. We confirm that this number increases after a monetary policy tightening for both, high and low-income households. Looking at the differential response, the results in Figure 12 suggest that high-income households increase the number of shopping trips relative to low-income households in several countries. This relationship is only marginally significant and relatively small in magnitude, but in line with our hypothesis.

FIGURE 12. Response of the differential number of shopping trips to a monetary policy shock, panel of the six largest euro area countries



Notes: The figure shows the impulse response of the cumulated difference in the average number of monthly shopping trips between high- and low-income households to a 10 basis points tightening monetary policy shock, estimated for a panel of the six largest euro area countries and based on household panel data. The dark and light grey areas indicate 1 and 1.645 standard deviation confidence intervals, respectively.

To summarise the various results reported in this subsection, they are in line with the notion that household-level inflation of high-income households responds relatively *more* to monetary policy

¹¹We only consider shopping for food and beverages (COICOPs 01.1, 01.2, 02.1). Several purchases in stores of the same name (and chain) on the same day count as a single trip. If the purchases are on different days, each day counts as one trip. If on a given day a household reports transactions in multiple stores (of different name), we count the number of unique store names on that day.

because of an adjustment in shopping behavior: In response to a contractionary shock, high-income households increase their shopping intensity relatively to low-income households and in particular engage in more intense product substitution.

5. CONCLUSION

Monetary policy has distributional effects — because a change in interest rates affects agents differently depending on their interest rate exposure, because income and wealth are affected differently, and possibly also because inflation responds differently across agents. While there is a large literature on the wealth and income effects of monetary policy, the effects on inflation differentials have barely been studied.

This paper studies the effect of monetary policy on inflation along the income distribution. Building on the paper by [Cravino et al. \(2020\)](#), it examines this question for the six largest euro area countries, and does so based on two different datasets. The first one has broad coverage across different consumption categories, spanning the entire spectrum that is underlying the measurement of the HICP, but covers variation along the income distribution only based on slow-moving differences in consumption shares. The second accounts for high-frequency changes in consumption patterns as well as for differences in the prices paid for the same goods, but is restricted to food and beverages.

The results reported in the paper suggest a more nuanced picture than in the previous literature. While we confirm that monetary policy affects inflation differently across the income distribution, there are different channels at play, which are important to understand. On the one hand, different consumption shares imply that inflation by high-income households responds *less* to monetary policy. On the other hand, we provide evidence for substantial differences in shopping behaviour, which furthermore imply that inflation by high-income households responds *more* to monetary policy.

This suggests that the effects identified in the earlier literature likely are an upper bound. It will therefore be important that future research quantifies the *overall* effect, also to understand to what extent differential inflation responses are an important additional channel through which monetary policy has distributional implications beyond those on wealth and income.

REFERENCES

- Adam, K. and P. Tzamourani (2016). Distributional consequences of asset price inflation in the euro area. *European Economic Review* 89, 172–192.
- Adam, K. and J. Zhu (2016). Price-level changes and the redistribution of nominal wealth across the euro area. *Journal of the European Economic Association* 14(4), 871–906.
- Aguiar, M., E. Hurst, and L. Karabarbounis (2013). Time use during the Great Recession. *American Economic Review* 103(5), 1664–1696.
- Albanesi, S. (2007). Inflation and inequality. *Journal of Monetary Economics* 54(4), 1088–1114.
- Ampudia, M., D. Georgarakos, J. Slacalek, O. Tristani, P. Vermeulen, and G. Violante (2018). Monetary policy and household inequality. Working Paper 2170, European Central Bank.
- Argente, D. and M. Lee (2021). Cost of living inequality during the Great Recession. *Journal of the European Economic Association* 19(2), 913–952.
- Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review* 109(6), 2333–2367.
- Bañbura, M. and E. Bobeica (2020, October). PCCI: A data-rich measure of underlying inflation in the euro area. Statistics Paper 38, European Central Bank.
- Bell, V., M. Joyce, Z. Liu, and C. Young (2012). The distributional effects of asset purchases. Quarterly Bulletin Q3, Bank of England.
- Bobasu, A., V. Di Nino, and C. Osbat (2023). The impact of the recent inflation surge across households. *Economic Bulletin* 3, 99–117. European Central Bank.
- Cardoso, M., C. Ferreira, J. M. Leiva, G. Nuño, Á. Ortiz, T. Rodrigo, and S. Vazquez (2022). The heterogeneous impact of inflation on households' balance sheets. Working Paper 176, Red Nacional de Investigadores en Economía.
- Charalampakis, E., B. Fagandini, L. Henkel, C. Osbat, et al. (2022). The impact of the recent rise in inflation on low-income households. Economic Bulletin Boxes 7, European Central Bank.
- Coibion, O., Y. Gorodnichenko, L. Kueng, and J. Silvia (2017). Innocent bystanders? Monetary policy and inequality. *Journal of Monetary Economics* 88, 70–89.
- Colciago, A., A. Samarina, and J. de Haan (2019). Central bank policies and income and wealth inequality: A survey. *Journal of Economic Surveys* 33(4), 1199–1231.
- Cravino, J., T. Lan, and A. A. Levchenko (2020). Price stickiness along the income distribution and the effects of monetary policy. *Journal of Monetary Economics* 110, 19–32.
- Doepke, M. and M. Schneider (2006). Inflation and the redistribution of nominal wealth. *Journal of Political Economy* 114(6), 1069–1097.
- Easterly, W. and S. Fischer (2001). Inflation and the poor. *Journal of Money, Credit and Banking* 33(2), 160–178.
- Griffith, R., E. Leibtag, A. Leicester, and A. Nevo (2009). Consumer shopping behavior: how much do consumers save? *Journal of Economic Perspectives* 23(2), 99–120.

- Hackethal, A., P. Schnorpfel, and M. Weber (2022). Inflation, net nominal positions, and consumption. Working paper, Goethe University Frankfurt. Available at <https://ssrn.com/abstract=4277575>.
- Hobijn, B. and D. Lagakos (2005). Inflation inequality in the United States. *Review of Income and Wealth* 51(4), 581–606.
- Jaravel, X. (2019). The unequal gains from product innovations: Evidence from the U.S. retail sector. *Quarterly Journal of Economics* 134(2), 715–783.
- Jaravel, X. (2021). Inflation inequality: Measurement, causes, and policy implications. *Annual Review of Economics* 13, 599–629.
- Jarociński, M. and P. Karadi (2020). Deconstructing monetary policy surprises—the role of information shocks. *American Economic Journal: Macroeconomics* 12(2), 1–43.
- Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review* 95(1), 161–182.
- Kaplan, G. and G. Menzio (2015, November). The morphology of price dispersion. *International Economic Review* 56(4), 1165–1206.
- Kaplan, G. and G. Menzio (2016). Shopping externalities and self-fulfilling unemployment fluctuations. *Journal of Political Economy* 124(3), 771–825.
- Kaplan, G., B. Moll, and G. L. Violante (2018). Monetary policy according to HANK. *American Economic Review* 108(3), 697–743.
- Kaplan, G. and S. Schulhofer-Wohl (2017). Inflation at the household level. *Journal of Monetary Economics* 91, 19–38.
- Kiss, R. and G. Strasser (2022). Inflation heterogeneity across households. Mimeo. European Central Bank.
- Krueger, A. B. and A. Mueller (2010). Job search and unemployment insurance: New evidence from time use data. *Journal of Public Economics* 94(3-4), 298–307.
- Kryvtsov, O. and N. Vincent (2021). The cyclicalities of sales and aggregate price flexibility. *Review of Economic Studies* 88(1), 334–377.
- Lauper, C. and G. Mangiante (2021). Monetary policy shocks and inflation inequality. Working Paper 21.02a, Lausanne University.
- Lenza, M. and J. Slacalek (2018). How does monetary policy affect income and wealth inequality? Evidence from quantitative easing in the euro area. Working Paper 2190, European Central Bank.
- Mitman, K., T. Broer, and J. Kramer (2022). The curious incidence of monetary policy across the income distribution. Discussion Paper 17589, CEPR.
- Nevo, A. and A. Wong (2019). The elasticity of substitution between time and market goods: Evidence from the Great Recession. *International Economic Review* 60(1), 25–51.
- Nord, L. (2022). Shopping, demand composition, and equilibrium prices. Available at SSRN 4178271. European University Institute.

- Pallotti, F., G. Paz-Pardo, J. Slacalek, O. Tristani, and G. Violante (2023). Distributional effects of inflation in the euro area. Mimeo. European Central Bank.
- Savignac, F. and P. Hubert (2023, May). Monetary policy and labor income inequality: the role of extensive and intensive margins. Discussion Paper 18130, CEPR.
- Schnabel, I. (2020). Narratives about the ECB's monetary policy – reality or fiction? Speech at the Juristische Studiengesellschaft, available at www.ecb.europa.eu/press/key/date/2020/html/ecb.sp200211_1_b439a2f4a0.en.html.
- Strasser, G., T. Messner, F. Rumler, and M. Ampudia (2023, July). Inflation heterogeneity at the household level. Occasional Paper 325, European Central Bank.

APPENDIX A. CALCULATION OF INFLATION BY INCOME GROUP FROM MICRO DATA

All of the following is done separately for each country in our sample, namely Belgium, France, Germany, Italy, the Netherlands, and Spain. For this reason, the country index is suppressed.

A.1. Household panel. Participants in the GfK/Kantar household panel report their purchases of fast moving consumer goods (FMCG), i.e. the items typically sold in supermarkets. The GfK/Kantar panel contains – among other information – the transaction date, the item’s barcode, its price, the quantity purchased, and some socio-demographic information on the purchasing household. Our sample spans six countries with differing starting years. The second column of Table A1 reports the first month for which month-on-month inflation rates can be calculated. All time series end in December 2018.

TABLE A1. Summary statistics

	sample period first month	income groups #	Laspeyres inflation			
			food & beverage avg. % p.a.	s.dev. % p.a.	all FMCG avg. % p.a.	s.dev. % p.a.
Belgium	Feb 2005	4	1.94	0.29	1.77	0.26
France	Feb 2008	4	0.23	0.29	0.15	0.29
Germany	Feb 2005	5	0.79	0.34	0.59	0.30
Italy	Feb 2012	4	-0.35	0.48	-0.59	0.50
Netherlands	Feb 2008	5	2.18	0.81	2.03	0.83
Spain	Feb 2008	4	0.69	0.23	0.58	0.21

Note: “First month” of the sample period is the first month for which month-on-month Laspeyres inflation can be calculated, i.e. for which price and quantity data are available already one month earlier. The reported number of income groups are those used in the analysis, after suitable combining thinly populated income groups in the raw data. The inflation rates are the annualized geometric averages of monthly inflation rates for the full sample period available for the respective country. In the case of “food & beverage” this inflation is based on sampled products in COICOP 01.1, 01.2 and 02.1, in the case of “all FMCG” it is based on sampled products in food, beverages, household maintenance, hobby, and personal care (COICOPs 01.1, 01.2, 02.1, 05.6, 09.3, 12.1).

We clean the raw transaction data as described in Appendix A.3. In particular, we restrict the sample to mainland locations, remove outliers, and remove items of varying weight. Because the most frequently purchased items are food and beverage items, we restrict our benchmark analysis to food (COICOP 01.1), non-alcoholic beverages (01.2), and alcoholic beverages (02.1). The right columns of Table A1 show that in all countries the inflation of food and beverage products (column 4) was between 0.08 and 0.24 percentage points higher than the overall FMCG inflation (column 6). In each country of our sample the standard deviation of the two series is very similar.

A.2. Information on income. Within each country we group the households by income. The availability of income information varies by country, but the data allows us nevertheless to distinguish four or five income groups in each country. Because of the large income differences between countries, the income thresholds are country-specific. For most countries (Belgium, France, Germany and the Netherlands), information on the net income of the household per month from all sources

is available. No income information is available for Italy and Spain, and we therefore group households in these countries instead by their reported social class. In Spain, social class captures the situation of the head of the household, whereas in Italy it reflects the situation of the entire household. The definition of social class changes in Italy at the beginning of 2016, which implies a break in the assignment of Italian households to income groups. Transactions by households which are not classified by income (or social class) in the dataset are excluded from the analysis.¹²

The third column of Table A1 reports the number of income groups we have available for each country. The high minus low inflation differentials used in this paper are based on the inflation difference between the highest and the lowest income group in the respective country.

A.3. Data cleaning. The raw data undergo a five-step cleaning and filtering procedure. We subject all countries to the same set of filters to make the datasets as comparable as possible.

- (1) Basic filters: As we want to study inflation experienced by typical (“representative”) households, we exclude overseas or non-mainland territories, as well as islands which are distant to the respective mainland (France: Corsica, Italy: Sardinia & surrounding islands, Spain: Balearic Islands, Canary Islands). The retail markets and with it the FMCG pricing on these islands appears to be quite different from the mainland. Furthermore, we drop items for which we have no information on the brand or the store it was bought at. We drop observations which entail seemingly unreliable price information. In many cases this is due to erroneous recording of multi-packs. For example, the price for a six-pack of beers is reported by some households as a price per bottle, whereas by others as a price for the entire six-pack. A filter which excludes observations of prices outside a factor of five of a product’s median price over the entire sample captures these cases reasonably well.¹³ This step reduces the number of distinct items (barcodes) in the sample by 33% (column “base” of Table A2).
- (2) COICOP: The product coverage in the raw data differs between countries. To improve the cross-country comparability and facilitate interpretation of inflation rates, we limit the product scope in our benchmark to the COICOP codes 01.1, 01.2 and 02.1. We drop observations which have no information on value or unit sales. This step reduces the number of distinct items (barcodes) in the sample compared to the previous stage by 20 percentage points (column “coicop” of Table A2).
- (3) Price pairs: We fill gaps of one month in the income group price series by linear interpolation. If an item has not been repurchased by a given income group – in the Laspeyres case, for example, if somebody in income group i purchased the item b in period $t - 1$, but nobody from that income group purchased it in period t and $t + 1$ – then we use the average price paid by households (in any income group) in that period. If neither is possible, we drop the

¹²In 2018, this affects 8% of households in Belgium and 3% in the Netherlands.

¹³At this stage of the cleaning process, one additional filter is applied for France: We found that some items entail a wrong mapping to their product category (specifically, hygiene and beauty products which are wrongly assigned to the categories popcorn or sugar products). We drop these items from our sample.

item from the index calculation. This step reduces the number of distinct items in the sample compared to the previous stage by 11 percentage points (column “pairs” of Table A2).

- (4) Volume per unit: In a subset of countries (Belgium, Germany and the Netherlands), the volume-per-unit of some items changes over time. This happens not only when the items are indeed individually packaged, but also when volume-per-unit is erroneously reported by some panelists. We therefore drop the items whose volume per unit in month t is different from its volume in month $t - 12$. The complementary set of countries (France, Italy and Spain) is not affected by this filter. For these countries, volume per unit is part of the time-invariant information set on the items. Step 4 reduces the number of distinct items (barcodes) in the sample only marginally (column “volume” of Table A2)
- (5) Outliers: First, we exclude observations with a price changing by more than a factor of four between t and $t - 12$. Second, we eliminate observations which show a price change beyond a factor of two while the observed quantity changes in the same (i.e. counter-intuitive) direction (i.e. up for a price increase or down for a price decrease) by more than one unit. Step 5 reduces the number of distinct items (barcodes) in the sample only marginally (column “final” of Table A2).

TABLE A2. Number of distinct items by stage of the data filtering process (to be completed)

c	raw	base	coicop	pairs	volume	final
Belgium	646262	318680	229088	174015	174013	173995
France	587490	416054	261978	226427	–	226274
Germany	901918	716773	602019	420700	420485	420163
Italy	387955	326325	237005	167037	–	166918
Netherlands	924489	416892	216038	180306	180305	180222
Spain	486790	358050	236577	175209	–	175178
%	100	67.06	46.64	35.11	35.11	35.09

Note: The table reports tallies of distinct items identified at the barcode level after each stage of the filtering algorithm as described in section A.3. Counts over the entire sample period, which differs by country. Column *raw* reports the count of distinct barcodes of the raw dataset. *base* reports the tally after a basic set of filters have been applied: items with unknown brand or store are dropped, a multi-pack price filter is applied and island/non-mainland territories are dropped from the sample. In the column *coicop* we choose only a subset of product categories, namely the 3-digit-COICOP codes 01.1, 01.2 & 02.1, and we drop observations with missing values in value or unit sales. Column *pairs* requires that a pair of prices is available for that item in both month t and month $t - 1$, either directly or by imputation. *volume* reports the number of remaining items after dropping all with a changing volume per unit. This affects only a subset of countries where this information is reported by GfK/Kantar panelists and is not part of the time-invariant information set on the items. *final* reports the set of barcodes on which the estimation of our inflation indices is based on. In this stage, two final filters are applied: an extreme-value filter and a filter that drops observations for which we observe a severe counter-intuitive change in price and quantity of the item.

A.4. Aggregation of prices. We treat all shopping during month t by households belonging to income group i as if it was done by one household representative of group i .

Let us denote the price per unit of a given barcode item b in a given (shopping) transaction s paid by a member of income group i by \tilde{p}_{bs}^i and the quantity purchased by \tilde{x}_{bs}^i . The set $S(bit)$ consists of all transactions of barcode item b by members of income group i during month t . The quantity of

item b purchased during month t by any household in group i is $x_{b,t}^i = \sum_{s \in S(b,t)} \tilde{x}_{bs}^i$. We calculate the quantity-weighted average price paid for item b by the households in this group during month t as

$$(A.1) \quad p_{bt}^i = \frac{1}{x_{bt}^i} \sum_{s \in S(b,t)} \tilde{p}_{bs}^i \tilde{x}_{bs}^i.$$

A.5. FMCG inflation indices. We focus on food and beverage products (COICOP 01.1, 01.2, and 02.1) because the GfK/Kantar household panel covers these categories most comprehensively. Our benchmark are month-on-month inflation indices.

Consumption at the household and income group level is very volatile and seasonal. In order to avoid spurious patterns in the inflation index, we use a rolling twelve-month average for quantities. We correct for the time variation in the panel size by using per-household quantities. Let h_t^i be the number of households in income group i with nonzero expenditure (in COICOP 01.1, 01.2, or 02.1) during month t . The average expenditure on item b of a household in income group i during month t is therefore given by $x_{b,t}^i/h_t^i$.

We start the analysis in the paper with the Laspeyres index, because of its analogy to the HICP. Our FMCG Laspeyres index in month t (based on the price pair p_{bt}^i and $p_{b,t-1}^i$) is based on the quantity during the past 12 months, i.e. $\bar{x}_{b,t-1}^i = \sum_{\tau=1}^{12} x_{b,t-\tau}^i/h_{t-\tau}^i, \forall t > 12$. As during the first year of the sample earlier quantities are unavailable, we use during that year the average across all monthly quantities, i.e. $\bar{x}_{b,1}^i = \dots = \bar{x}_{b,12}^i = \sum_{\tau=1}^{12} x_{b,\tau}^i/h_{\tau}^i$.

Let us denote the set of all food and beverage items b (i.e. items within the COICOPs 01.1, 01.2, and 02.1) that were purchased by – potentially different – households in group i in month t by $\tilde{B}(it)$. In order to calculate a price change in period t , prices or price proxies must be available for both month $t-1$ and month t . The set of items within income group i for which this is possible based on our procedure in step (3) of section A.3 is $B(i, t-1) = \tilde{B}(i, t-1) \cap [\tilde{B}(i, t) \cup \tilde{B}(i, t+1) \cup \tilde{B}(i^C, t)]$, where i^C denotes the complement of income group i .

Definition 1 (Laspeyres inflation, month-on-month). *The FMCG Laspeyres inflation rate of income group i over a one-month period ending with month t is*

$$(A.2) \quad \pi_{t-1,t}^i = \frac{\sum_{b \in B(i,t-1)} p_{b,t}^i \bar{x}_{b,t-1}^i}{\sum_{b \in B(i,t-1)} p_{b,t-1}^i \bar{x}_{b,t-1}^i} - 1.$$

In order to capture the effect of substitution, we analyse the corresponding Paasche indices as well. The underlying quantities are defined analogously, but now based on a leading moving average to capture substitution. More precisely, our Paasche index in month t (based on the price pair p_{bt}^i and $p_{b,t-1}^i$) is based on the average quantity over the next 12 months, i.e. $\bar{x}_{b,t}^i = \sum_{\tau=0}^{11} x_{b,t+\tau}^i/h_{t+\tau}^i, \forall t \leq T-11$. As during the last year of the sample later quantities are unavailable, we use for the Paasche index during that year the average across all monthly quantities, i.e. $\bar{x}_{b,T-11}^i = \dots = \bar{x}_{b,T}^i = \sum_{\tau=T-11}^T x_{b,\tau}^i/h_{\tau}^i$.

Definition 2 (Paasche inflation, month-on-month). *The FMCG Paasche inflation rate of income group i over a one-month period ending with month t is*

$$(A.3) \quad \pi_{t-1,t}^i = \frac{\sum_{b \in B(i,t)} p_{b,t}^i \bar{x}_{b,t}^i}{\sum_{b \in B(i,t)} p_{b,t-1}^i \bar{x}_{b,t}^i} - 1.$$

The imputation of prices proceeds analogously to the Laspeyres index. The set of items within income group i for which this is possible is therefore $B(i, t) = \tilde{B}(i, t) \cap [\tilde{B}(i, t-1) \cup \tilde{B}(i, t-2) \cup \tilde{B}(i^C, t-1)]$.

A.6. Summary Statistics. Table A3 provides an overview of the inflation experienced by high- and low-income households in the household panel. With the exception of Italy and Belgium, the (Laspeyres) inflation experienced by low-income households during the sample period has been higher than the inflation experienced by high income households. The upper two panels on the left side of Table A3 show a one percentage point higher inflation for food and beverage items in the Netherlands, and only a marginally smaller difference for the “all FMCG” basket.

TABLE A3. Inflation (% p.a.)

	Laspeyres inflation				Paasche inflation			
	food & beverage avg.	s.dev.	all FMCG avg.	s.dev.	food & beverage avg.	s.dev.	all FMCG avg.	s.dev.
(a) high income								
Belgium	1.56	0.32	1.47	0.30	0.03	0.32	-0.30	0.31
France	-0.05	0.33	-0.16	0.35	-1.00	0.35	-1.24	0.36
Germany	0.29	0.44	0.10	0.38	-0.55	0.43	-0.87	0.37
Italy	0.12	0.53	-0.09	0.52	-2.08	0.52	-2.22	0.51
Netherlands	2.35	1.03	2.25	1.07	-0.75	1.02	-1.01	1.06
Spain	0.76	0.26	0.62	0.23	0.00	0.26	-0.12	0.23
(b) low income								
Belgium	1.45	0.35	1.25	0.32	-0.09	0.34	-0.25	0.31
France	0.51	0.36	0.37	0.34	-1.31	0.36	-1.56	0.36
Germany	0.92	0.46	0.71	0.42	-1.29	0.43	-1.51	0.39
Italy	0.04	0.63	-0.17	0.61	-3.30	0.71	-3.61	0.68
Netherlands	3.37	1.01	3.18	1.02	-1.56	0.99	-1.75	1.01
Spain	0.85	0.26	0.73	0.23	0.07	0.26	-0.09	0.23
(c) difference high-low income (p-value)								
Belgium	0.66		0.35		0.56		0.83	
France	0.00		0.00		0.09		0.08	
Germany	0.00		0.00		0.00		0.00	
Italy	0.88		0.87		0.05		0.02	
Netherlands	0.00		0.00		0.00		0.00	
Spain	0.53		0.39		0.63		0.81	

Note: The reported inflation rates are the annualized geometric average of monthly inflation rates during the common sample period February 2012 – December 2018. In the case of “food & beverage” this inflation is based on sampled products in COICOPs 01.1, 01.2 and 02.1, in the case of “all FMCG” it is based on sampled products in food, beverages, household maintenance, hobby, and personal care (COICOPs 01.1, 01.2, 02.1, 05.6, 09.3, 12.1).

This difference in Laspeyres inflation is highly significant as shown in panel (c) of Table A3, which shows the p -values of a test of the null hypothesis of equal inflation between the highest and the lowest income group for the joint sample period. It rejects equal (Laspeyres) inflation during the common sample period 2012–2018 at the 1% significance level for France, Germany and the Netherlands. The same test in our full sample returns the same results.

At the same time, we do not find a significant difference in the respective inflation between high- and low-income households during 2012–2018 in Belgium, Italy and Spain.¹⁴

Households respond to inflation by substitution. The intensity of substitution differs by income group, which is vividly uncovered by the Paasche inflation reported on the right side of Table A3. After substitution, in all countries but Spain it is the low-income households which end up with a lower inflation rate. According to panel (c), this difference is significant at the 5% level in Germany, Italy and the Netherlands, and at the 10% level in France. This suggests that – unconditionally – low-income households adjust their basket more effectively than high income households in response to inflation.

Inflation rates differ considerably between countries. Within each country, however, all income groups broadly share main inflation fluctuations for the set of goods covered in the panel. But whereas there is little evidence for a *persistent* deviation of the inflation of an income group from the national average, there are for many countries extended periods during which specific income groups experience notably higher or lower inflation.

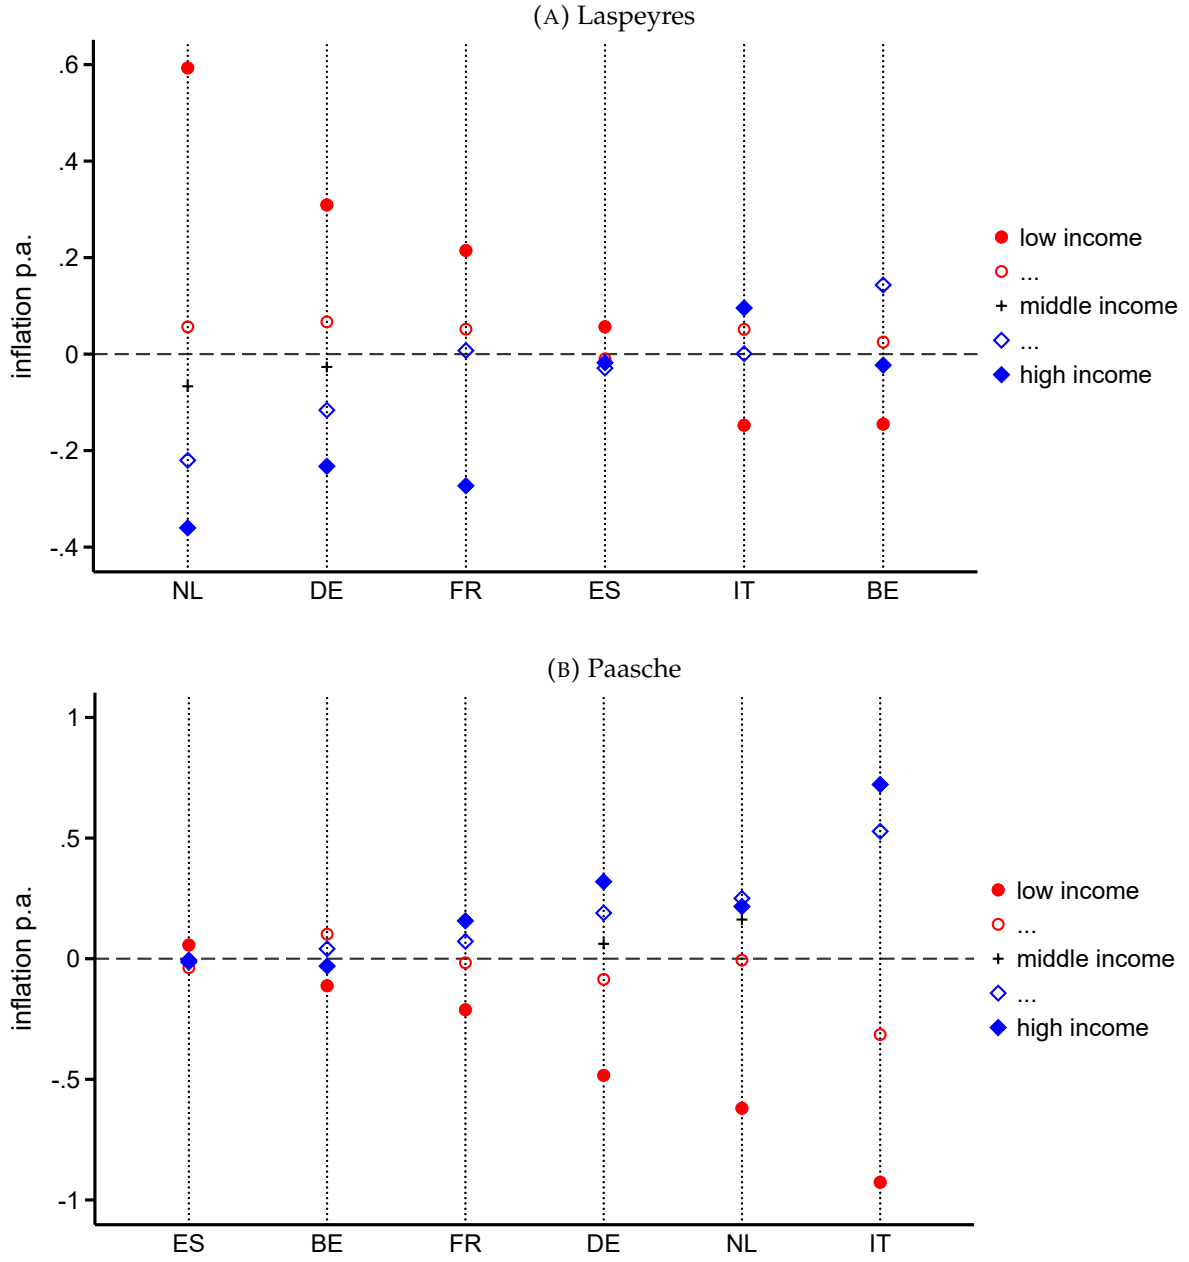
The relation between income and inflation is not limited to tail income groups. Figure A1 compares the inflation differences between the top and the bottom income group for the common sample period. Among our six countries, especially the Netherlands, France, and Germany display higher Laspeyres inflation for low income groups (upper panel). In these countries, inflation during 2012–2018 decreases monotonically with income. The difference between income groups varies between countries, however. In Spain, the differences are very small, whereas in Italy and Belgium there is no persistent relationship between income and average Laspeyres inflation during the sample period. The latter, of course, does not imply that such a relation never exists. In fact, the time variation in figure 7 gives reason to suspect that there might not be other periods during which such a relationship emerges.

A comparison with the lower panel of figure A1 unveils in France, Germany, the Netherlands and Italy a relatively more effective substitution by low-income than by high-income households. As a result high-income households are left with higher Paasche inflation than low-income households in these countries.

A.7. Decomposition. Laspeyres inflation $\pi_{t-1,t}^i$ as defined in equation (3) can be decomposed into an expenditure change component $\varphi_{t-1,t}^i$ and a quantity change component $\Xi_{t-1,t}^i$ as $1 + \pi_{t-1,t}^i = \frac{1 + \varphi_{t-1,t}^i}{1 + \Xi_{t-1,t}^i}$.

¹⁴One potential reason is that in Italy and Spain households are classified by social class instead of income.

FIGURE A1. Inflation by income group (p.p. p.a., demeaned, 2012-2018)



Note: Average y-o-y inflation by income group based on the 12-months rolling average of $\pi_{t-1,t}^i$ over the common sample period February 2012 – December 2018 in percentage points per year. COICOPs 01.1, 01.2, and 02.1. The series are demeaned country-by-country and sorted by the difference between inflation experienced by higher vs. the one experienced by lower income groups.

The *expenditure change* component $\varphi_{t-1,t}^i$ is the ratio of expenditure per household in income group i during month t relative to the respective expenditure in the previous month. The expenditure change component can be written as

$$(A.4) \quad \varphi_{t-1,t}^i = \frac{\sum_{b \in B(i,t)} p_{bt}^i \bar{x}_{bt}^i}{\sum_{b \in B(i,t-1)} p_{b,t-1}^i \bar{x}_{b,t-1}^i} - 1,$$

which can be further decomposed according to

$$1 + \varphi_{t-1,t}^i = \frac{\sum_{b \in B(i,t-1) \cap B(i,t)} p_{b,t-1}^i \bar{x}_{b,t-1}^i}{\sum_{b \in B(i,t-1)} p_{b,t-1}^i \bar{x}_{b,t-1}^i} \times \frac{\sum_{b \in B(i,t-1) \cap B(i,t)} p_{b,t}^i \bar{x}_{b,t}^i}{\sum_{b \in B(i,t-1) \cap B(i,t)} p_{b,t-1}^i \bar{x}_{b,t-1}^i} \times \frac{\sum_{b \in B(i,t)} p_{b,t}^i \bar{x}_{b,t}^i}{\sum_{b \in B(i,t-1) \cap B(i,t)} p_{b,t}^i \bar{x}_{b,t}^i}.$$

The first term captures the month-to-month expenditure drop due to products that have not been repurchased by the income group in month t , and the third term, conversely, captures the month-to-month expenditure increase due to products that have been purchased in month t , but had not been purchased in month $t - 1$. The second term captures the expenditure change in repurchased products (within a 12-month rolling average).

The *quantity change* component $\Xi_{t-1,t}^i$ is the ratio of the quantity purchased per household in income group i during month t relative to the respective quantity in the previous month, where the quantities in both periods are weighted by the price vector of month t . The quantity change component can be written as

$$(A.5) \quad \Xi_{t-1,t}^i = \frac{\sum_{b \in B(i,t)} p_{b,t}^i \bar{x}_{b,t}^i}{\sum_{b \in B(i,t-1)} p_{b,t}^i \bar{x}_{b,t-1}^i} - 1,$$

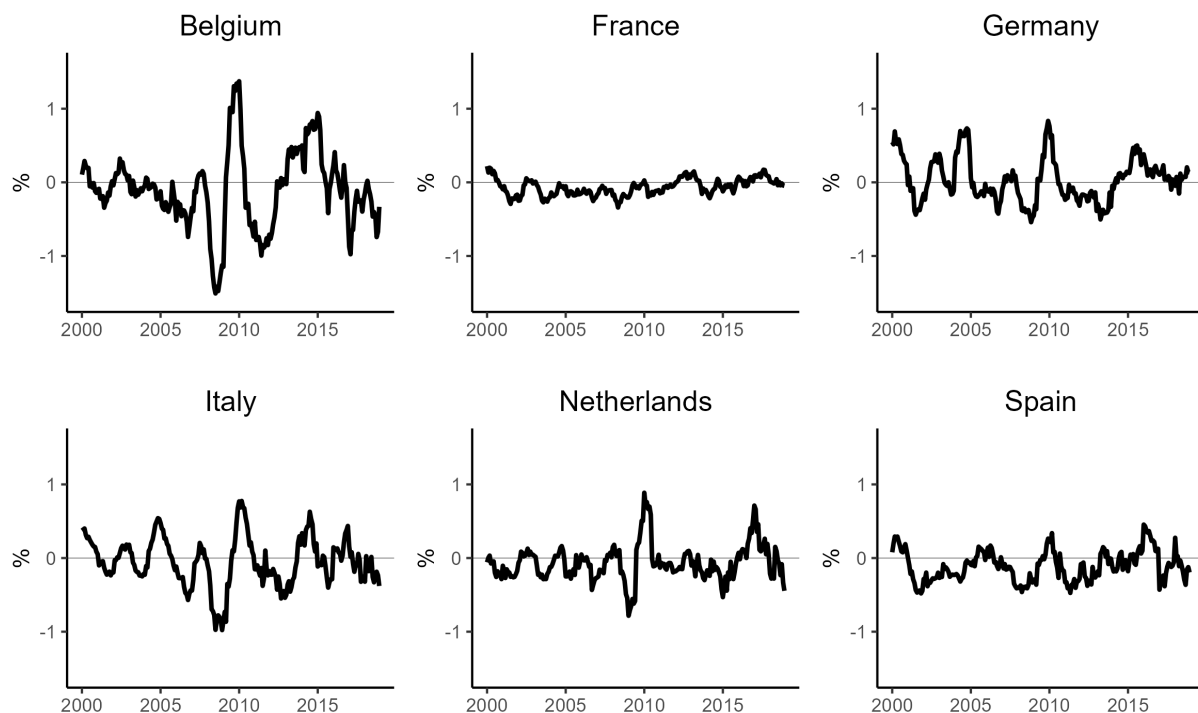
which can be further decomposed according to

$$1 + \Xi_{t-1,t}^i = \frac{\sum_{b \in B(i,t-1) \cap B(i,t)} p_{b,t}^i \bar{x}_{b,t-1}^i}{\sum_{b \in B(i,t-1)} p_{b,t}^i \bar{x}_{b,t-1}^i} \times \frac{\sum_{b \in B(i,t-1) \cap B(i,t)} p_{b,t}^i \bar{x}_{b,t}^i}{\sum_{b \in B(i,t-1) \cap B(i,t)} p_{b,t}^i \bar{x}_{b,t-1}^i} \times \frac{\sum_{b \in B(i,t)} p_{b,t}^i \bar{x}_{b,t}^i}{\sum_{b \in B(i,t-1) \cap B(i,t)} p_{b,t}^i \bar{x}_{b,t}^i}.$$

The first term captures the quantity drop due to products that have not been repurchased by the income group i in month t , and the third term, conversely, captures the quantity increase due to products that have been purchased in month t , but had not been purchased in month $t - 1$. The second term captures the quantity change in repurchased products.

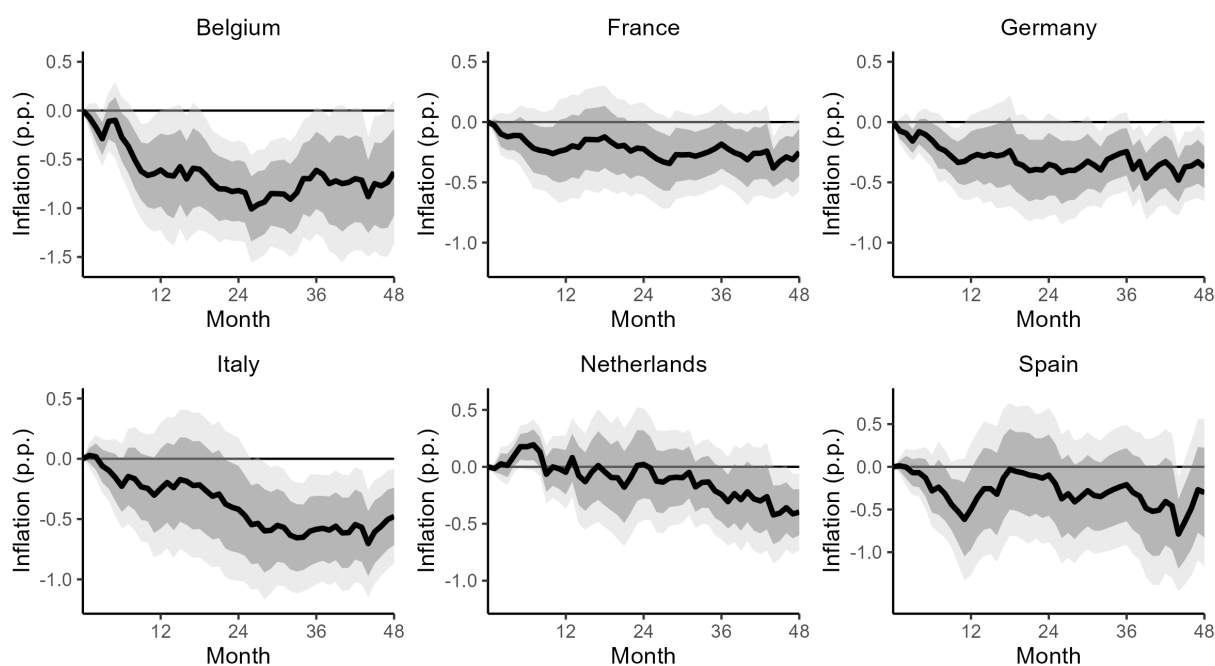
APPENDIX B. ADDITIONAL RESULTS AND ROBUSTNESS CHECKS

FIGURE B1. Differences in HICP inflation between high- and low-income households



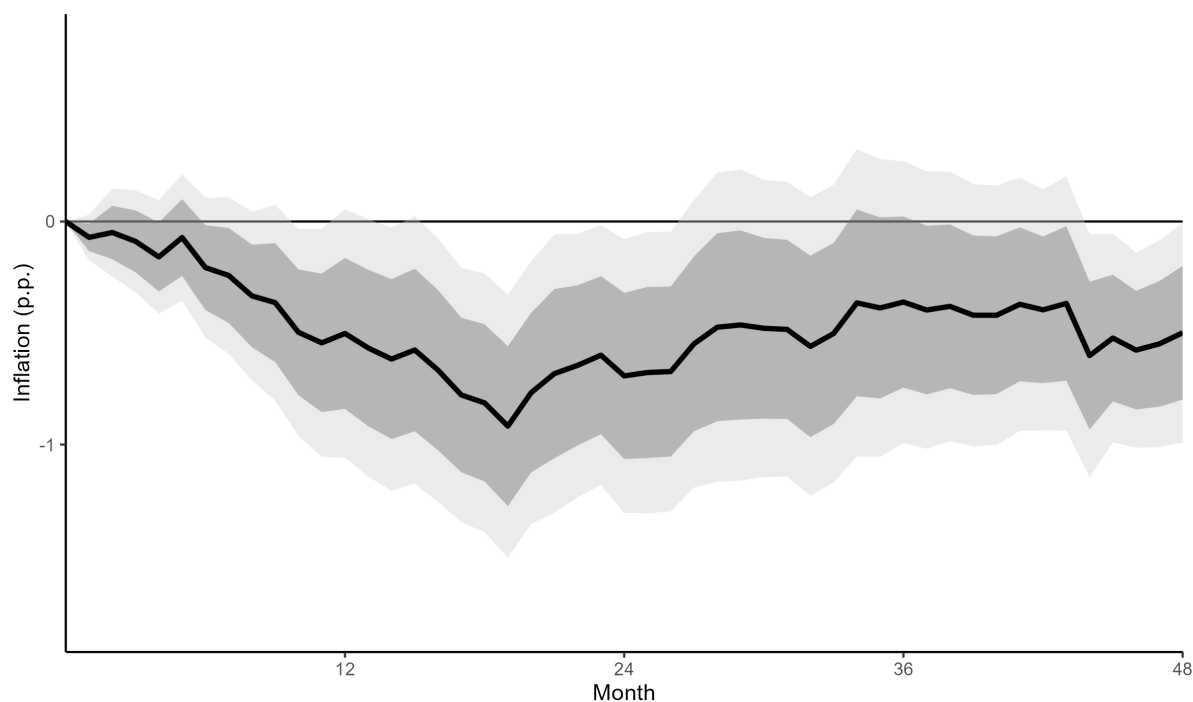
Notes: The figure shows year-on-year HICP inflation of the highest income quintile households in the various countries minus year-on-year HICP inflation of the lowest income quintile households in those countries.

FIGURE B2. Response of national HICP inflation to a monetary policy shock



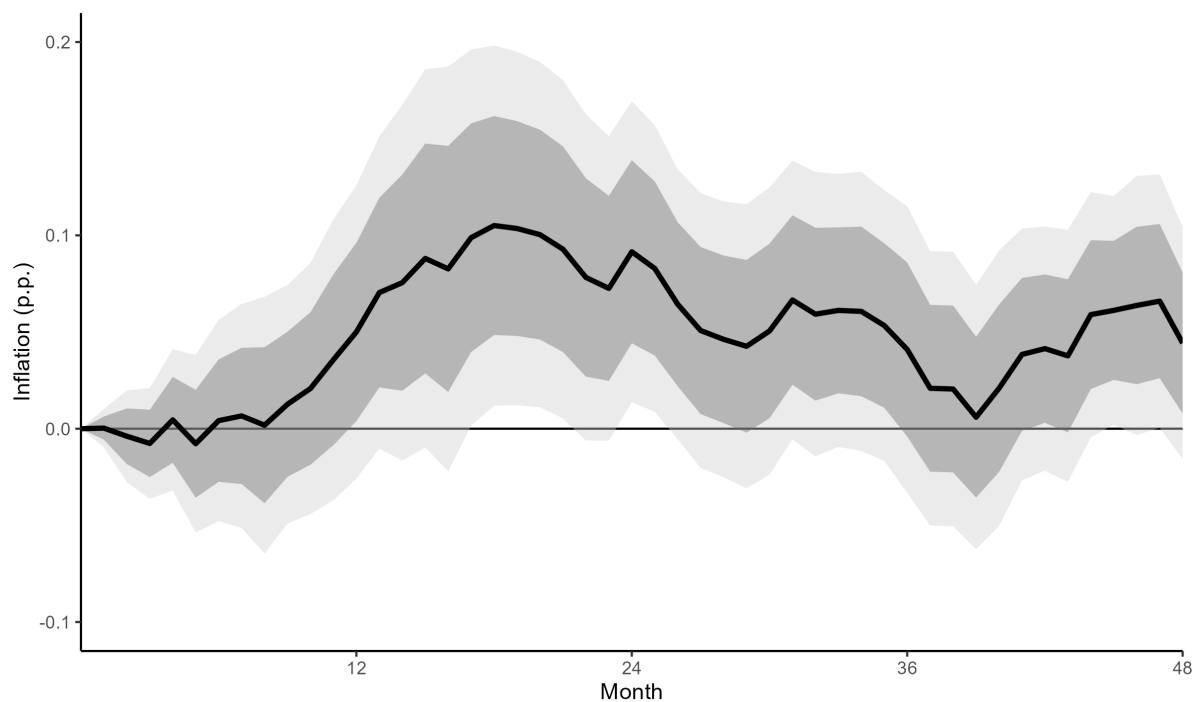
Notes: The figure shows the impulse response of country-specific aggregate HICP inflation to a 10 basis points tightening monetary policy shock. The dark and light grey areas indicate 1 and 1.645 standard deviation confidence intervals, respectively.

FIGURE B3. Response of HICP COICOP1 inflation to a monetary policy shock, panel of the six largest euro area countries, sample starting in 2005



Notes: The figure shows the impulse response of HICP COICOP category 1 inflation to a 10 basis points tightening ECB monetary policy shock, for a sample starting in 2005, estimated for the largest 6 euro area countries. The dark and light grey areas indicate 1 and 1.645 standard deviation confidence intervals, respectively.

FIGURE B4. Response of the inflation differential to a monetary policy shock, panel of the six largest euro area countries, sample starting in 2005



Notes: The figure shows the impulse response of inflation differential between high- and low-income households to a 10 basis points tightening ECB monetary policy shock, for a sample starting in 2005, estimated for a panel of the six largest euro area countries. The dark and light grey areas indicate 1 and 1.645 standard deviation confidence intervals, respectively.

TABLE B1. Expenditure shares differences between high- and low-income households

Panel A: Average across all Household Budget Survey waves													
Country	01	02	03	04	05	06	07	08	09	10	11	12	Overall
BE	-3.54	-0.95	1.90	-12.58	2.94	-1.45	6.47	-1.05	3.19	0.32	3.03	1.72	19.56
DE	-6.22	-1.18	0.67	-12.46	2.03	2.57	8.58	-1.89	2.01	0.17	1.93	3.81	21.72
ES	-9.69	-1.36	0.50	-2.07	1.71	0.11	2.44	-0.50	2.92	1.32	3.31	1.31	13.62
FR	-3.04	-1.25	0.21	-8.61	2.48	0.30	5.04	-1.51	2.36	-0.04	2.72	1.34	14.46
IT	-13.93	-0.93	2.25	-8.34	4.61	1.07	7.18	-1.33	3.66	0.50	4.58	0.68	24.54
NL	-1.96	-0.99	1.41	-9.65	2.19	0.35	5.06	-1.11	2.38	-0.60	2.91	0.01	14.34
EA	-6.46	-1.16	0.86	-8.82	2.14	0.56	6.19	-1.31	2.63	0.39	2.63	2.35	17.76
Panel B: Household Budget Survey 1999													
Country	01	02	03	04	05	06	07	08	09	10	11	12	Overall
BE	-6.30	-0.81	0.87	-5.42	2.86	-2.40	4.97	-0.50	3.43	-0.01	2.19	1.14	15.42
DE	-7.29	-1.83	0.25	-8.15	2.10	2.95	9.35	-1.45	1.42	-0.04	1.79	0.89	18.78
ES	-11.68	-1.62	0.01	3.54	1.87	-0.26	1.50	0.11	2.60	1.27	0.97	1.68	13.56
FR	-5.81	-0.90	0.55	-7.18	2.41	0.30	5.94	-0.95	2.23	-0.06	2.95	0.53	14.88
IT	-11.09	-1.01	0.28	3.12	1.02	-0.83	3.19	-0.81	2.32	-0.11	3.32	0.61	13.86
NL	-4.56	-0.75	0.75	-6.28	2.05	0.59	2.70	-0.96	1.90	-0.99	4.82	0.72	13.56
EA	-8.17	-1.33	0.42	-4.06	1.93	0.73	6.00	-0.95	2.10	0.07	2.38	0.87	14.52
Panel C: Household Budget Survey 2015													
Country	01	02	03	04	05	06	07	08	09	10	11	12	Overall
BE	-0.60	-1.20	2.40	-17.58	3.40	-0.70	7.09	-1.80	3.30	0.40	3.60	1.70	21.90
DE	-4.70	-0.90	1.40	-15.32	2.30	2.50	9.11	-1.80	2.90	0.40	2.40	1.70	22.74
ES	-6.74	-0.90	1.09	-9.58	1.49	0.19	3.58	-1.01	2.79	1.70	5.99	1.39	18.24
FR	-1.91	-1.50	-0.20	-9.23	2.39	-0.00	4.69	-1.70	2.29	0.10	3.09	1.99	14.52
IT	-14.87	-0.90	2.90	-12.16	5.80	1.70	8.51	-1.50	4.10	0.70	5.00	0.71	29.46
NL	-0.26	-2.09	1.82	-14.77	2.42	0.20	9.45	-1.29	3.03	-0.89	2.42	-0.05	19.32
EA	-4.49	-1.10	1.11	-11.88	1.91	0.20	6.62	-1.50	2.91	0.60	2.81	2.81	18.96

Notes: Columns 01 to 12 show the the share of total expenditure allocated to each COICOP category by households in the higher income quintile minus the share allocated by households in the lower income quintile in percentage points. The column "overall" is one-half of the sum of the absolute value of the preceding 12 columns. Panel A shows average numbers across all Household Budget Survey waves (1999, 2004, 2010, 2015), panel B shows numbers for the 1999 Household Budget Survey and panel C shows numbers for the 2015 Household Budget Survey.