Beyond the Headline: How Personal Exposure to Inflation Shapes the Financial Choices of Households

WORKING PAPER SERIES

Eesti Pank
EUROSÜSTEEM

7/2023

The Working Paper is available on the Eesti Pank web site
DOI: 10.23656/25045520/072023/0206
ISBN 978-9916-710-11-1 (pdf)
Eesti Pank Working Paper Series, ISSN 2504-5520; 7/2023 (pdf)

Beyond the Headline: How Personal Exposure to Inflation Shapes the Financial Choices of Households*

Merike Kukk, Jan Toczynski, Christoph Basten

Abstract

Using individual level panel data from a period of volatile inflation in Estonia in 2005-11 and interactive fixed effect estimation, we find individual consumption to respond to personal inflation beyond the headline rate. Households are exposed to different inflation due to different consumption baskets. For each percentage point of higher personal inflation exposure, they increase consumption by more than 1%, and also increase stock market investments. These responses are consistent with backward-looking inflation expectations. They are financed with savings or borrowing, except when the household is liquidity-constrained or over-indebted. Extra demand when inflation is already high can make inflation persistent and dependent on its current distribution.

Keywords: **inflation heterogeneity, personal inflation exposure, consumption, borrowing, interactive fixed effects, intertemporal choices.**
JEL-codes: **D14, D15, E21, E31.**

The views expressed are those of the authors and do not necessarily represent the official views of Eesti Pank or the Eurosystem. Responsibility for the information, data and content in this article lies entirely with the authors.

*Authors affiliation and email: Merike Kukk (corresponding author): Eesti Pank and Tallinn University of Technology, merike.kukk@eestipank.ee; Jan Toczynski, Ecole Polytechnique Federale de Lausanne; Christoph Basten, University of Zurich, Swiss Finance Institute and CESifo.

We would like to thank Michael Funke, Dimitris Georgarakos, Jiyhuan Huang, Ludmila Fadejeva, Matti Keloharju, Felix von Meyerinck, Milda Norkute, Tomasz Olma, Roberto Rigobon, Swapnil Singh, Tairi Rõõm, Melina Vosse, Michael Weber and Yucheng Yang, and participants at the CEBRA 2023 Annual Meeting in NYC, the BEC 2023 Conference in Riga, the ECEE Conference in Tallinn, and the Swiss Society for Financial Market Research (SGF) conference 2022, and participants in seminars at TalTech and at Eesti Pank for useful comments and feedback. We also thank Robin Hazlehurst for the language editing. Merike Kukk acknowledges funding from the EU Horizon 2020 Research and Innovation programme under Grant Agreement No. 952574.

Non-technical summary

Discussions about inflation usually consider the changes in the prices of goods and services that are found from a harmonised index of consumer prices. The index is calculated from the typical consumption basket containing the quantities of goods and services that are consumed by the average household. However, each household has their own consumption basket and so they are exposed to different price changes, meaning different households experience different rates of inflations. In this paper we use data on personally experienced inflation and we analyse how individuals react in their consumption and their financial decisions to the inflation they experience.

Expectations for various economic indicators, including inflation expectations, play an important role in shaping the behaviour of individuals. When people expect prices to rise, they prefer to make their planned purchases before prices go up, meaning that they bring their consumption forward. Several research papers have shown that previous experience is important in forming expectations, because when people notice prices rising, they expect that those prices will continue to rise in the near future. So when people form their expectations from past experience, we can investigate the direct link between the inflation they have experienced and their consumption and financial behaviour.

This paper explores how exposure to inflation affects consumption using two datasets from 2005-2011, which are the Estonian Household Budget Survey and a quarterly, individual-level dataset. We model the consumption baskets of households from the Household Budget Survey and use the model to calculate personal consumption baskets and personal inflation in the individual-level dataset. Quarterly data make it possible to estimate quarterly changes in consumption that are induced by the inflation experienced in the same quarter. We estimate short-term effects that are explained by intertemporal reallocation of consumption.

The baseline model shows that inflation being 1 percentage point higher leads to an increase of 1.4% in quarterly real consumption. The positive coefficient remains when different specifications are used for the fixed-effects models, and it indicates growth in consumption of more than 1%.

The availability of liquid funds determines how much of their consumption people can bring forwards. They can respond more strongly to personal inflation when they have larger amounts in their checking accounts and term deposits, but they cannot increase their consumption when they lack liquid reserves.

We run additional estimations to investigate how people finance this extra consumption. The results show that they do it partly by reducing their savings in term deposits and partly by taking out more in consumer loans and overdrafts. Although borrowing makes it possible to front-load the consumption, indebtedness hampers consumption. The deeper someone is in debt, especially debt from consumer loans or measured as the debt servicing burden, the weaker their consumption response is to the inflation they experience.

The main finding of the paper is the positive reaction of consumption that can be explained by backward-looking inflation expectations. The positive reaction can make inflation more persistent and it is difficult for central banks to affect inflation expectations that are based on past experience. At the same time, liquidity constraints and indebtedness are an important obstacle to the positive consumption response. Policy-makers can evaluate the impact of inflation on consumption more accurately if they take these constraints into account.

Contents

1 Introduction

Inflation has been historically low in developed economies for a long time, but it has now returned to the radars of economists, central bankers, and the general public. Most of the discussions about the causes and consequences of high inflation focus on headline inflation in different geographic areas even though it is widely acknowledged that different parts of the population within an area can actually be facing very different rates of inflation (e.g., Jaravel 2021, Weber et al. 2022). Beyond the obvious implications of this for social welfare (e.g., Yang 2022), the question arises of whether different distributions of inflation exposure may matter for the economic and financial decisions that households make, or more simply for predictions of consumption.

It has long been difficult to identify how different exposures to inflation affect consumption because most countries last experienced episodes of high and variable inflation some time ago when the availability of granular data was more limited. Furthermore current inflation is, or at least can be, endogenous to current consumption in most setups, or both inflation and consumption are affected by common drivers like monetary or fiscal policy. We address these limitations studying a unique episode of high and variable inflation in Estonia in 2005-11. In this small open economy setup inflation is significantly driven by shocks from abroad[1], which affect different groups of households differently depending on their pre-existing consumption basket weights.

But of course different initial consumption baskets may be correlated with other household characteristics which are also associated with different inter-temporal consumption dynamics and which may not be fully absorbed by group and time fixed effects. Therefore we use Interactive Fixed Effects estimators (Bai 2009, Vogt et al. 2022) to absorb any significant time-varying unobservable factors that could be correlated with our regressors of interest. Using a unique set of anonymised quarterly bank account data, we thus obtain three core sets of results.

First, we find that personal inflation exposure impacts real spending by individuals positively. The baseline model shows that an increase of one percentage point in average quarterly inflation exposure translates into a sizeable increase of 1.4% in real expenditure. As we effectively control for aggregate inflation by including time fixed effects, this result should be interpreted as the response to the price changes experienced personally rather than to headline inflation. Those affected by higher inflation increase their real consumption by more than others do. Our finding is consistent with backward looking agents forming expectations on future price growth from their own recent personal experience, as found for house prices

[1] See, e.g., Maćkowiak (2007), Aastveit et al. (2016), Jovičić et al. (2017).

by Kuchler & Zafar (2019), for consumer prices by D'Acunto et al. (2021), or for economic growth by Dovern et al. (2023). It is also consistent with inflation expectations being positively related to subsequent consumption decisions through the intertemporal substitution mechanism (e.g., Coibion et al. 2022, D'Acunto et al. 2022, D'Acunto et al. Forthcoming). However, to the best of our knowledge this study is the first one to link consumption directly to experienced inflation.

Second, people finance this extra consumption partly by reducing their savings in term deposits and partly by borrowing more through consumer loans or overdrafts. We find that individuals with more liquid wealth respond more strongly to exposure to inflation, while there is no front-loading of consumption by households that are observed to have insufficient liquidity in their checking accounts. These results highlight the importance of liquidity constraints acting to block intertemporal substitution since budget-constrained individuals cannot adjust their spending by bringing it forward.

Third, households who are more indebted ex ante respond less to personal inflation exposure. Similarly, the coefficient for the inflation experienced is smaller the higher the debt servicing costs are. This finding suggests that higher debt servicing costs make it harder to front-load consumption in response to inflation. Our findings contribute to the literature considering how debt payments crowd out consumption (e.g., Du Caju et al. 2023, Toczynski 2023) and on the relation between debt and consumption sensitivity (e.g., Baker 2018, Cookson et al. 2022). We further find that exposure to higher inflation leads to an increase in stock market investments that can be explained by portfolios being rebalanced because inflation affects risky assets and safe ones differently (Agarwal et al. 2022).

Our findings have important implications for policy-makers and contribute to the literature on the consequences of inflation inequality (see, e.g., the recent work of Orchard 2022 and Yang 2022). Most importantly, we show that heterogeneous inflation exposure matters for real consumption choices, even after the effect of headline inflation has been controlled for. Our results indicate that inflation leads to a redistribution of consumption toward more exposed households, possibly because it increases their inflation expectations proportionately more. This suggests that there might be a positive feedback loop in which inflation incentivizes spending, which in turn leads to even higher pressure on price growth. Consequently, future inflation might depend on the distribution of current inflation.

Our results further indicate that inflation heterogeneity might also have implications for other components of household balance sheets, as we find increased borrowing by households that have experienced higher inflation. This means that inflation might lead to increased leverage for some household groups, which could have possible consequences for financial stability. Policy-makers should consequently take account of inflation heterogeneity when modelling the effects of monetary policy.

A further contribution of our work is to document that individuals who are neither liquidity-constrained nor over-indebted drive the main impact of experienced inflation. This illustrates that financial constraints might impair a crucial mechanism for inter-temporal substitution. This makes it important when evaluating how inflation expectations affect actual spending to take financial constraints into account. It is possible that some of the studies that find no effect of expectations on consumption might have worked with samples of individuals who are mostly constrained, which prevented them from identifying the effect.

In the remainder, Section 2 discusses the relationship between experienced and expected inflation, Section 3 presents our hypotheses, Section 4 the data and Section 5 our empirical strategy. Then Section 6 discusses our results, Section 7 probes their robustness of results, and Section 8 concludes.

2 Experienced and Expected Inflation

Conventional economic theory, as illustrated by the Euler equation, predicts that individuals will bring consumption forward when they expect prices will be higher, *ceteris paribus*, as they will get less value for the same money in the future (see, e.g., Hall 1988). Rational and forward-looking consumers should thus increase their consumption when they expect prices will rise. For them to do this, two conditions need to be met. First, they need to believe that the increase in prices is not only temporary, as this gives them an incentive to bring their consumption forward. Second, they need to able to bring that consumption forward. This means the effect should be stronger for those who have ample liquid wealth or who are able to borrow, since they will be able to finance the increased consumption. If consumers cannot draw forward a part of their expenditures, such as spending on monthly rent or perishable food, it might be rational for them to save more today so that they will still be able to afford the necessary expenditure in future. Which reasoning dominates will depend on whether more consumption can or cannot be drawn forward and on whether individuals can afford to finance this intertemporal substitution.

Either way, most models that link consumption choices to inflation expectations are agnostic about how these expectations are formed. For simplicity or because of data availability, it has been assumed either that the expectations are formed rationally and anticipate actual future price changes correctly, or that they depend on past headline inflation as reported in the media.

We stipulate in this paper that individuals may be aware that their consumption basket, and the resulting future inflation that is relevant to their inter-temporal utility maximization, may be different from the average household for which underlying headline inflation is calculated. We allow that their recent experience of inflation may be their major simple predictor of the personal inflation they will face.

Our study is motivated in part by recent findings on the experience effect in economics and finance, which show how individuals form their expectations for financial and economic outcomes by looking to their past experience (Malmendier 2021, Malmendier & Wachter 2021). In the literature on experience and inflation expectations, Malmendier & Nagel (2016) show that individuals who have lived through periods of high inflation have persistently higher inflation expectations. Similarly, D'Acunto et al. (2021) and Weber et al. (2022) use granular information on purchases by households that is linked to survey data and find a strong positive correlation between the price changes individuals experience in their daily lives and the inflation expectations they report. This literature further shows that when individuals form inflation expectations, they might think about prices of specific goods (de Bruin et al. 2012), and give excess weight to the prices that they frequently observe while shopping (D'Acunto et al. 2021, Georganas et al. 2014), or to prices that are particularly salient for them such as petrol prices (Coibion & Gorodnichenko 2015).

There is some disagreement about the economic mechanism underlying the link between experiences and expectations. One interpretation is that households extrapolating from their personal experiences when they form beliefs about the future might be irrational or constrained by bounded rationality. Kuchler & Zafar (2019) show for example that individuals tend to extrapolate from recently observed changes in local house prices when predicting future aggregate housing prices, but there does not seem to be any relationship between such extrapolations and how informative personal experiences actually are in predicting aggregate trends. In a similar vein, Dovern et al. (2023) study the expectations of German firms and show that they use recent local and industry-specific experience to form their beliefs about aggregate trends.

Dovern et al. (2023) interpret these findings as evidence of rational inattention. It is also possible in our case that individuals care more about their personal price index rather than aggregate ones, in which case it might be rational to use signals received from personal experience, regardless of how informative they are about the aggregate CPI. This branch of literature provides evidence of a strong positive association between personal exposure to inflation and beliefs about future inflation. This highlights that personal exposure to inflation is likely to have an effect on consumption beyond that of aggregate headline inflation.

In this paper, we estimate the consumption responses to recent personal inflation exposure directly,[2] bypassing inflation expectations in the causal chain linking inflation exposure

[2]We are interested in the role of inflation that is "personal" in the sense that it varies at a more granular

and spending decisions. Instead, we analyse whether inflation at the individual level matters for consumption and whether individuals also make spending decisions by considering their personal exposure to inflation rather than relying solely on information about the aggregate Consumer Price Index (CPI) or communication from the central bank that is available to all households. We do this by calculating group-specific exposures to proxy for individual exposure to inflation. If we find a significant relationship between group-specific inflation exposure and consumption decisions, it will mean that the distribution of current inflation in the population matters for the distributional effects of monetary policy.

Unlike Malmendier & Nagel (2016) and D'Acunto et al. (2021), we look at how recent experiences, rather than distant ones, impact spending rather than expectations. The advantages of this approach are firstly that the inflation rates experienced by various demographic groups are more directly observable and straightforward to measure than subjective expectations, and secondly that we highlight the role of inflation inequality and the feedback pathway from current inflation to consumption and then to future inflation. Lastly, our results contribute indirectly to the literature on the relationship between expectations and consumption choices, which we discuss in more detail below.

3 Channels for the relationship between inflation and consumption

Previous studies have provided strong evidence that past experiences are used by individuals when forming inflation expectations; see the overview by Malmendier (2021) and Malmendier & Wachter (2021). This implies that personally experienced inflation is likely to matter for the consumption choices of households, but the existence and sign of such a relationship are ultimately an empirical question. For example, it could be that customers might not be able to compute and track their inflation exposure, or might not be interested in doing so. Instead, they may base their decisions only on the headline inflation that is typically reported in the media or in central bank messages. Thus, our first hypothesis tests whether personal exposure to recent inflation matters even after common aggregate factors, including headline inflation, are controlled for:

level than one single representative consumer per economy. In practice, we observe consumption basket weights for population groups rather than for each individual, and we approximate "personal" by using "group-specific" inflation. Furthermore, we follow statistics agencies in observing average prices for each consumption category and period rather than the prices paid for each purchase, so our regressor strictly captures "inflation exposure". On these grounds, we subsequently use the terms "personal" and "group-specific", and the terms "inflation" and "inflation exposure" interchangeably.

Hypothesis 1. *(Only headline inflation matters) Personal inflation exposure does not affect real consumption beyond the effect of headline CPI inflation.*

Next, it may be that people consciously or subconsciously incorporate their personal exposure to inflation when forming inflation expectations. If that is the case, we want to know the sign of the relationship between personal inflation exposure and spending. Previous studies have shown a positive correlation between actual inflation and expectations. Hence, the sign will depend on the relationship between expectations and consumption decisions.

If the intertemporal substitution channel motivated by the Euler equation dominates, we would expect a positive relationship between inflation exposure and consumption. In this scenario, personally observed changes in price would drive expectations about future inflation, and the spending of individuals seeking to optimise their purchasing power over time.

Hypothesis 1a. *(Intertemporal motives) Personal inflation exposure has a* ***positive*** *effect on real consumption.*

If, however, households deem that a large part of their spending on monthly rent, food or transport for example, cannot be shifted, exposure to higher personal inflation now and the resulting higher inflation expectations may instead lead them to reduce their consumption now. They can increase their precautionary savings, which would allow them still to afford enough of their necessary consumption when prices are higher in the future. In line with this we see that while some studies have found that inflation expectations have a positive effect on consumption (e.g., Coibion et al. 2022, D'Acunto et al. 2022), others have mixed or negative results (e.g., Burke & Ozdagli 2023, Coibion et al. Forthcoming, Weber et al. 2022).

More recently, several authors have argued that the inflation expectations of households are often correlated with them having pessimistic beliefs about their future economic prospects (see, e.g., Kamdar et al. 2018, Weber et al. 2022, Coibion et al. Forthcoming). This supply side view of inflation suggests that the motive of precautionary saving might dominate the intertemporal substitution channel and lead to there being a negative association between inflation expectations and spending. This motivates our last hypothesis:

Hypothesis 1b. *(Precautionary motives) Personal inflation exposure has a* ***negative*** *effect on real consumption.*

Although we organise our reasoning around the link between personal inflation exposure and expectations, there could be other mechanisms that link experienced price changes and consumption decisions. Nominal price changes might influence consumption through behavioural effects such as the money illusion (see e.g., Patinkin 1965, Shafir et al. 1997, Deaton 1977, Branson & Klevorick 1969), while inflation impacts net debtors and borrowers differently by changing the real value of outstanding debt. This reallocation of wealth by inflation could potentially generate additional wealth effects on household consumption. In this study, we estimate the net effect of these channels and do not attempt to measure their relative contribution as our setup is not well suited for that. We focus instead on the role of personal exposure to inflation as opposed to headline inflation. We will, however, use the theoretical setup discussed in this section to guide our empirical analysis and discuss how our results relate to particular theories.

4 Setup and Data

4.1 Inflation Developments in Estonia

It is challenging to study how exposure to inflation affects consumption choices. This is firstly because headline inflation was persistently low and stable in most developed countries for a long time until recently, and so there was little variation to use for the identification. Earlier periods, when there was more variation, lack good quality micro-data though. A second reason is that economic theory suggests that inflation and consumption are inherently interconnected, and so any connection identified may suffer from reverse causality. Our primary motivation for choosing Estonia as the setup for our analysis is that it has a unique environment in which we can aim to overcome these difficulties.

Observed inflation rates in Estonia in 2005-2011 exhibit much larger amplitudes and variability than those in most other advanced economies (see Figure 1). The annual rate of change in the Harmonised Consumer Price Index (HICP) averages 3.46% over the period of our analysis in Estonia ranging from -2.1% to over 11.5% during the financial crisis, while it averaged only 1.72% in the euro area. This variability is crucial for studying the effects of inflation, as agents might react differently in environments of high or low inflation. Similarly, some effects might materialise only when personal exposure to inflation reaches a certain threshold, since individuals might not notice small price increases of say 0.5%, but may notice and react to large ones of say 5%. Moreover, the dynamics of inflation indexes that are consumption category-specific in Estonia show substantial heterogeneity (see Figure 2), providing the variation in personal inflation exposure that we need (we present further category-specific inflation rates in Table A1 in the Appendix).

Estonia is a small and open economy, and so developments in domestic inflation, especially during a crisis, are likely to be driven mainly by external factors, which means that

Figure 1: Inflation rates in Estonia and the EU.

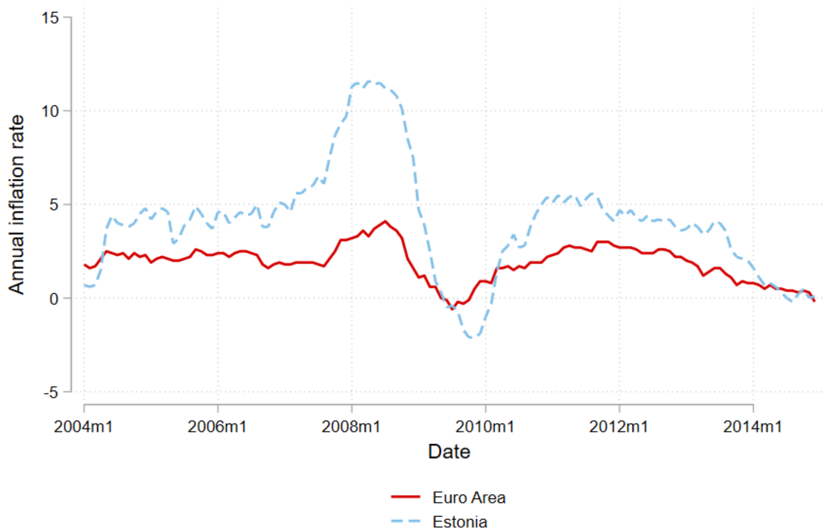

Notes: The figure plots inflation rates (HICP) for Estonia (blue dashed line) and the euro area (red solid line) over time.
Source: Eurostat.

Figure 2: Category-specific inflation rates

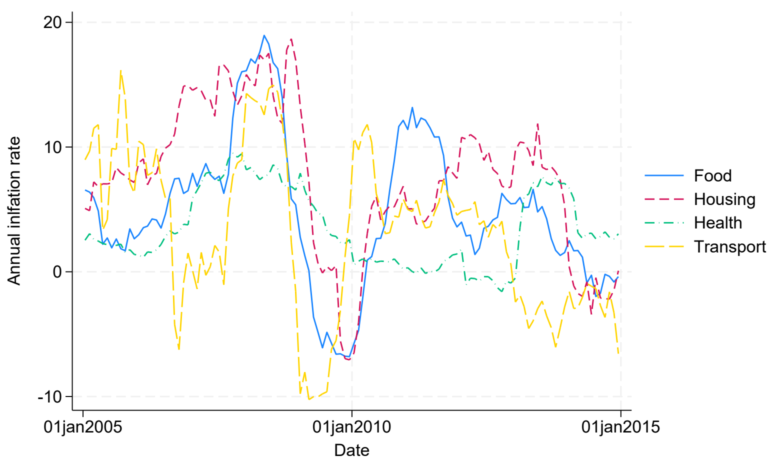

Notes: The figure plots the evolution of the main inflation sub-indexes in Estonia over time. The rates presented here are annual rates calculated for each month.
Source: Statistics Estonia.

our main independent variable is plausibly exogenous to demand from households. This important feature of our empirical design allows us to give our estimates a stronger causal interpretation.

Several papers empirically confirm that domestic inflation in a small open economy is primarily driven externally. Foreign factors might affect domestic inflation through multiple channels such as trade, prices of commodities and especially oil, and spillover effects from foreign monetary policy. Maćkowiak (2007) decomposes the drivers of domestic inflation for a group of emerging economies, and finds that around half of the variability in inflation in these countries is externally driven. Similarly, Aastveit et al. (2016) explore inflation in several small advanced economies and attribute as much as 80% of the variation in it to foreign factors. Jovičić et al. (2017) obtain very similar results from analysing the drivers of inflation in Croatia, another small European country. They find that external factors explain around 50% of the variance in domestic inflation, and they further decompose the domestic factors into supply and demand, concluding that domestic demand contributes only around 20% of the country's inflation. This means that the endogeneity of inflation is less of an issue when data from a small, open economy are used to investigate how inflation affects consumption.

4.2 Account-level Dataset

We use a unique anonymised database from one of Estonia's leading commercial banks.[3] The database contains quarterly data on inflows and outflows from individuals' checking accounts, and the balance of checking and savings accounts, securities, and debt at the end of the quarter, running from Q4 2004 to Q4 2011. The dataset is unique because of its large cross-section of over 100,000 individuals, and its long time dimension at quarterly frequency over seven years.

The dataset is large enough to provide comprehensive coverage of the whole population of the country. The entire dataset covers over 100,000 individuals, accounting for roughly 12% of Estonia's working-age population at the time. Banking concentration is very high in Estonia and the three largest banks had 96% of total banking-sector assets at that time (Cuestas et al. 2020), and customers mainly use just one banking group for all their finances. The sample has been extracted from the regular customer base covering regular income and transactions in the bank, and so we can assume that customers consider the bank to be their main bank. This implies that the dataset captures the complete picture of the financial transactions and financial assets of the individuals concerned. The dataset also contains the

[3]Data processing for scientifc purposes was used in compliance with the General Data Protection Regulation.

age, gender, and region of residence of the individuals.

Having comprehensive coverage of an individual's finances lets us construct plausibly accurate measures of our most important variables, which are income, consumption, and savings. An important benefit is that we can construct a measure of spending as the sum of outflows from a checking account while excluding transfers to other individual bank accounts, such as saving accounts or purchases of securities. We can also exclude automatic payments from the checking account made to the bank, which are for loan and interest repayments or banking service fees. Similarly, we can observe the inflows to checking accounts from legal entities such as firms, public institutions, or private accounts. The inflows do not contain transactions between saving accounts, credit amounts issued to customers, or house purchases financed by loans.[4] We use this measure as a proxy for income. Figure A1 in the Appendix shows that mean income and spending in the dataset closely follow the dynamics of earnings and consumption in the aggregate data.

The initial sample covers over 100,000 individuals, but we make several exclusions. We first exclude the self-employed as their transactions may be from their business finances rather than their personal finances. Next we exclude individuals who have zero income or zero spending within a quarter, as missing income or spending indicates that although the bank considers them to be regular customers, they might also run their finances through another bank.

We also drop observations where inflows to checking accounts in four consecutive quarters are smaller than or equal to the automatic transactions with the bank, which are mainly related to debt repayments, as this would indicate that the customer is using the bank account mainly for debt servicing. Finally, we drop extreme values at the top 1% for consumption, income, and the balance of financial assets and debt, and the top 1% for the purchase or sale of securities.

The original dataset is strongly balanced, but these exclusions cause gaps in the panel. We exclude from the final sample all individuals for whom more than five observations were dropped because of outliers. There are 89,507 individuals in the final sample, and the time dimension varies between 24 and 29 quarters.

We convert the nominal values for income, consumption, and financial assets into real values using the individual rate of inflation that we compute in the following section.[5] The distribution of consumption and income is skewed, and that of financial assets and debt is especially so, and so we use a standard approach for addressing the long right tail by

[4] We do not have information about house purchases, but we can deduce the value of housing loans from the subsequent outflow from the checking account, assuming that the loan is used for house purchases.

[5] Alternatively, aggregate HICP can be used to compute real values. The choice of the deflator does not affect the results.

applying log transformation of the variables.[6] The definition of the main variables is given in Table 1 along with their summary statistics.

4.3 Household Budget Survey

The second dataset we use is the Estonian Household Budget Survey (HBS). The account level dataset covers household finances comprehensively, but only gives us a measure of total spending and has no information about disaggregated expenditures. However, we need those expenditures if we are to measure the relative exposure of households to changes in the prices of different goods, and so we use data from the Estonian Household Budget Survey to estimate weights for consumption categories and assign them to observations in the account level dataset using a set of socio-economic characteristics of households that are contained in both datasets. We describe the procedure and list the variables used for this in Section 5.2.

The HBS is conducted annually, and a different group of households is interviewed each month. The information collected by the survey illustrates a rich set of demographic and socio-economic characteristics at the household and individual levels. Households are also asked to take note of all their everyday expenditures for two weeks, or for a month before the survey was redesigned, and the values obtained are then converted into monthly figures. Households are also asked separately to recall other one-off purchases from the previous 12 months, and those figures too are translated into monthly values. Finally, the HBS aggregates the reported expenditures into 12 distinct consumption categories that we can then use to compute consumption baskets and individual inflation in the account-level dataset, as explained in detail in Subsection 5.2 below. Our analysis uses the survey waves from 2005-2007 and 2010-2011. In 2008 and 2009, the survey was suspended as it was redesigned to fit the European Union guidelines more closely.

[6]As we cannot log-transform zero values, we add one unit to all values so the zero values of the original variables remain zero after the log transformation. Alternatively, we can use inverse hyperbolic sine transformation (IHS) to deal with zero values, but the results are very similar, so we use the conventional log-transformation.

Table 1: Summary statistics

	Units	Definition	Obs.	Mean	St.Dev.
Main variables used in regressions					
$IndINFL_{it}$	p.p.	Individual quarterly inflation of individual i in quarter t, calculated from the imputed consumption shares	2458862	1.200	1.110
$\log C_{it}$	log(EUR+1)	Consumption as real quarterly outflows from checking account in EUR with log-transformation $ln(C_{it}+1)$	2479220	7.212	0.856
$\log \mathrm{INC}_{it}$	log(EUR+1)	Income as real quarterly inflows from checking account in EUR with log-transformation $ln(Inc_{it}+1)$	2479220	7.208	0.824
$\log \mathrm{FIN}_{it}$	log(EUR+1)	The real balance of the sum of checking accounts, saving accounts and securities, in EUR with log-transformation $ln(Fin_{it}+1)$.	2479220	6.183	2.463
Other variables					
Total income	EUR	Total inflows to an account.	2479220	2189	1956
Total spending	EUR	Total outflows from an account.	2479220	2260	2235
Checking account	EUR	Real checking account balance	2479220	1130	2588
Term deposit	EUR	Real term deposits balance	316506	4770	6125
Stocks	EUR	Real balance on stocks	79891	2569	4981
Funds	EUR	Real balance on investment funds	50459	2887	4087
Bonds	EUR	Real balance on bonds	915	2675	2082
$\frac{CheckAssets}{Yincome}$	Ratio	The balance on checking account to yearly income	2479220	0.231	0.693
$\frac{TermAssets}{Yincome}$	Ratio	The balance on term deposits to yearly income	2479220	0.084	0.487
$\frac{LiquidAssets}{Yincome}$	Ratio	The balance on checking account and term deposits to yearly income	2479220	0.146	0.426
Total debt	EUR	Real balance of any debt (housing loan, consumer loan, overdraft)	1136363	10729	19561
Add housing loan	1/0	Dummy equal to 1 if the debtor increases balance on mortgage loan in this period.	409844	0.062	0.242
Add consumer loan	1/0	Dummy equal to 1 if the debtor increases balance on consumer loan in this period.	224087	0.125	0.331
Add overdraft	1/0	Dummy equal to 1 if the debtor increases overdraft in this period.	740462	0.479	0.500
$\frac{TotDebt}{Yincome}$	Ratio	Total debt balance to yearly income	1094820	0.945	1.819
$\frac{Housing}{Yincome}$	Ratio	Housing loan balance to yearly income	394112	2.155	2.518
$\frac{ConsLoan}{Yincome}$	Ratio	Consumer loan balance to yearly income	220307	0.259	0.350
DSR	Ratio	Annual debt repayments to yearly income	1094820	0.226	0.218
Age	Years	Age in years.	2479220	46	12
Gender	1/0	Dummy equal to 1 if the account holder is male and 0 otherwise.	2479220	0.363	0.481
Tallinn	1/0	Dummy equal to 1 if the account holder is lives in Tallinn and 0 otherwise.	2479220	0.329	0.470

Notes: This table presents summary statistics for the variables used in the empirical analysis for the period from Q1 2005 to Q4 2011 for which we compute our measure of personal inflation exposure. For a few term deposit and investment accounts, for which no volumes are reported, we convert the missing values to zeros.

5 Empirical Strategy

This section lays out our empirical strategy. We first discuss the measure of personal exposure to inflation that we use, then we explain how we construct that measure by using the available data to assign personalised weights to the consumption categories for each individual account holder. Finally we elaborate on our estimation strategy and discuss the assumptions used for the identification.

5.1 Household Level Inflation

The basis of our estimation strategy is that we exploit the heterogeneity in personal inflation exposure. Headline inflation is usually measured as a single composite index, such as the HICP, that tracks changes in the value of a representative consumption basket. However, there are large differences in how households compose their consumption baskets, and so their specific spending patterns, shopping habits, or geographical location mean that households face different rates of inflation as the prices of various products and services change at different paces.

The concept of inflation inequality has long been recognised in economic research. An early contribution by Michael (1979) for example studies the distribution of individual inflation rates in the United States by combining survey data on consumption with inflation sub-indexes for specific categories of goods. A similar, but more recent, analysis is provided by Hobijn & Lagakos (2005). Even more recent examples come from Weber et al. (2022) who document large inflation heterogeneity during the Covid-19 pandemic in the US, who find that differences in inflation exposure are small during normal times but increase significantly during recessions.

We study inflation heterogeneity in this paper by calculating household-specific inflation rates from different consumption bundles. It typically involves assigning weights to different categories of goods to measure how exposed individuals are to changes in the prices of those categories. A weighted average of category-specific inflation processes is then calculated with these weights (see, e.g., Kaplan & Schulhofer-Wohl 2017, D'Acunto et al. 2021, and Jaravel 2021, for an overview on measuring inflation). Thus, the HouseholdCPI is calculated as:

$$\text{HouseholdCPI}_h = \sum_{i=1}^{W} w_{c,h}\text{CPI}_c, \tag{1}$$

where h indexes households and c indexes different good categories. In the next two sections, we explain how we measure the consumption basket weights and link them to information on price changes. A potential limitation that might cause differences in personal inflation exposure to be underestimated is that this approach focuses solely on the differences in the price increases for wide consumption categories. An interesting novel perspective on the heterogeneity of inflation is provided by Kaplan & Schulhofer-Wohl (2017), who depart from the rest of the literature by focusing on differential price changes *within* consumption categories rather than *between* them. Their analysis uses data at the level of the purchase scanner and is limited to retail supermarket goods. It suggests that variation in prices within categories impacts household-specific inflation rates a great deal. We have to approximate inflation for goods from inflation data at the level of categories of goods.

5.2 Imputation of Consumption Shares

We impute consumption shares for each individual in the accounts dataset from their socio-economic characteristics. We do this using consumption data from the Household Budget Survey (HBS) administered by Statistics Estonia. The HBS provides data for households on their composition, income, and expenditures classified into 12 categories (see Table A2 in the appendix). The data on expenditures allow us to calculate household-level consumption weights $w_{c,h}$ (for household h at time t, we omit the time subscript from now on for simplicity) for the 12 expenditure categories representing the distribution of a household's consumption basket. We further denote a range of observed socio-economic characteristics X_h. Note that since our primary dataset is at the individual level rather than household level, we map account holders to household heads from the HBS.

As we aim to model the distribution of spending across numerous categories, we are effectively modelling the vector of shares $\mathcal{W} = (w_1, ..., w_N)$. Our modelling approach must thus consider the interconnections and dependence between those shares. The weights need to be non-negative ($w_c \leq 1$ for $c \in 1, ..., N$) and sum up to one ($\sum_{c=1}^{N} w_c = 1$). In a regression framework furthermore, the marginal effects of any single independent variable on the weights should add up to zero, as any increase in one weight needs to be balanced by a corresponding decrease in some of the other weights. For these reasons it would not be appropriate to apply the standard regression approach to each share separately, and the weights need to be modelled jointly. To do that we use the Multinomial Fractional Logit framework in which the conditional expectation of weight $w_{c,i}$ (c indicates a good category, while i indicates an individual) is modeled by:[7]

$$E[w_{c,i}|X_i] = \frac{e^{\beta_c X_i}}{\sum_{k=1}^{N} e^{\beta_k X_i}}. \tag{2}$$

This approach jointly estimates the vector of weights by maximum likelihood and meets the criteria for non-negativity and adding-up. It also accommodates the boundary cases of zero weights.

To impute consumption shares for a household from the account level dataset, we estimate equation 2 with the Household Budget Survey using the set of socio-economic variables X_i that is present in both datasets. We then use the estimated coefficients $\hat{\beta}$ to predict the vector of consumption shares $\mathcal{W}$ in the account-level data.

An important consideration when constructing individual-level measures of changes in the cost of living is the period over which basket weights $\mathcal{W}$ should be measured. The typical approach with aggregate data is to use either a Laspeyres or a Paasche-type index, where Laspeyres sets the weights at a fixed base period t and Paasche uses the weights from the period $t + 1$ (see the discussion in Kaplan & Schulhofer-Wohl 2017 and Jaravel 2021). In a micro-setting, the Laspeyres index has the advantage of measuring exposure using only pre-treatment data, though this may be at the cost of possibly overstating inflation by not accounting for substitution between different types of goods.

We take an intermediate approach and we pool observations from all the available waves of the HBS covering the period 2005-2011 that we are considering. To ensure that the variation that consumption models pick up comes from changes in inflation rather than from changes in weights, we use constant consumption weights calculated over the entire sample period.[8] Weights $\mathcal{W}$ consequently measure the average exposure to a given category of goods over the sample period, and in addition, any seasonal patterns in the composition of consumption are averaged out. To make sure that this choice does not introduce any systematic biases by incorporating future information into the current measure of exposure, we repeat the baseline estimations using only data from 2005 as the first year to calculate the weights as a robustness check, effectively using a conservative approach with a Laspeyres-type index with the weights set at the beginning of the sample. The results are very similar, indicating that this design choice does not excessively influence our results (see Table A5 in

[7]We do not directly map group-level weights from HBS to the account dataset since our explanatory variables define around 400 groups and some of them would only include few observations in HBS.

[8]The use of constant weights implies that we do not use all information on inflation heterogeneity across time, hence allowing for higher noise by not updating households' consumption weights. Initial estimations with time-varying weights indicate stronger results of consumption to inflation. However, we hold to the conservative approach with constant weights to avoid reverse causality from consumption to the change in the weights.

the Appendix showing robustness).

The explanatory variables we use to estimate the consumption weights are: 1) *IncDecile*, for the income decile of the household in a given survey year; 2) *Age*, with ten dummies for the age group of the household head in five-year intervals; 3) *Tallinn*, which is a dummy indicating residence in the capital city Tallinn; and 4) *Gender*, for the gender of the household head.

We impose four restrictions on the available HBS waves. The first is that we require the household head to be between 20 and 70 years of age to match the age groups in our account-level data. The second is that we exclude households headed by a self-employed person, as we did in the account database, because we might wrongly classify self-employed business income as personal income.[9] The third is that we drop observations in which households report zero values for at least eight of the twelve spending categories to alleviate any biases from misreporting. We also drop any household that reports zero spending on food as those values are virtually certain to be misreported. Our fourth restriction is that we trim all the consumption categories at the 99th percentile to eliminate extreme values that are likely to be incorrect or to result from unusual expenditures. Our final sample contains 11,541 households.

We use the coefficients obtained from estimating Equation 2 on the HBS data to impute consumption weights $\hat{\mathcal{W}} = (\hat{w_1}, ..., \hat{w_N})$ for the account dataset. Table A2 in the Appendix compares the consumption shares obtained originally from the HBS data and those imputed. We observe that our procedure accurately recreates the means of the consumption shares, especially for the largest categories. The leading consumption weights in the HBS are "Food and non-alcoholic beverages" at 0.306 and "Housing" at 0.189, which together make up half of the average consumption basket. Figure A2 plots the mean of the group-specific inflation rates from the transaction dataset against headline CPI. The two variables co-move closely, as should be expected. We run additional robustness checks of the imputation strategy in Section 7.

Our imputation strategy necessarily leads to us losing a substantial part of the variation. As we are using two categorical variables and two dummy ones to impute consumption weights, we are essentially estimating mean inflation for 400 groups in the sample. While the procedure reconstructs the means quite accurately, the imputation compresses the standard deviations in the imputed shares. However, the size of our dataset allows us to overcome the problem of a relatively high noise-to-signal ratio. Imputing the estimated consumption shares

[9]There is another strand of literature that estimates misreporting of income and consumption by the self-employed in the household budget surveys (Kukk et al. 2020), so excluding the group would make the results more robust.

could also introduce considerable measurement error into our main independent variable. Conventional econometric theory suggests that measurement error would lead to attenuation bias, which would work against us finding any significant results by biasing the coefficients towards zero. An alternative approach would be to extract the mean consumption weights of the 400 groups directly from the HBS to the transactional dataset, but this would add more noise into the means of small groups in the HBS, which we can avoid with the imputation.[10] In our analysis, we are interested in the existence and signs of the relationships studied rather than in finding their accurate magnitudes. Ultimately, our main estimates are statistically and economically significant despite the possible bias towards a zero effect.

5.3 Imputing Personal Inflation Exposure

Equipped with the estimated shares $\hat{\mathcal{W}} = (\hat{w_1}, ..., \hat{w_N})$, we construct our main independent variable, which is the person-level inflation rate for individual i in period t, as the consumption-weighted average of the category-specific inflation rates:

$$IndCPI_{i,t} = \sum_{c=1}^{W} \hat{w}_{c,i} \text{CPI}_{c,t}. \tag{3}$$

As we impute weights from a set of characteristics of the household head, $IndCPI_{i,t}$, should be interpreted as a group-specific price index, as we essentially instrument a person-level price index with a group-specific index. Since our financial data come at quarterly frequency, we calculate personal CPI as a quarterly variable by using category-specific monthly HICPs averaged over a quarter.[11] From the individual CPI we calculate individual quarterly inflation: $IndINFL_{i,t} = \frac{(IndCPI_{i,t} - IndCPI_{i,t-1})}{IndCPI_{i,t-1}}$.

We are aware that using the household-level HBS data to impute person-level consumption weights in the transactional data adds noise to the inflation variable if the individual consumption weights are different from the weights at the household level. Although we cannot link household members in the transactional data, we can assume that there is income and consumption sharing within couples, so that even when individual household members are responsible for purchases in different product categories, they take spending decisions jointly while also sharing information about price changes in the categories concerned. We use the deciles for equalised income in the HBS in the imputation model, as these are adjusted to the number of household members and match better with the deciles of individual income in the account-level data.

[10] In five groups there is only one observation while in 54 groups there are up to five observations.

[11] We opt for averaging HICP in a quarter instead of using the end-of-quarter values to account for the fact that consumption is a flow variable.

Figure A2 in the Appendix shows that our measure of personal quarterly inflation follows the dynamics of headline inflation closely. It averages 1.2% over the entire period with a within standard deviation of 1.11, and a between standard deviation of 0.07. There is substantial heterogeneity in personal inflation across individuals, with the largest difference being two percentage points in 2008Q4 while the average difference over the period is 0.72 percentage point. Personal inflation hits its maximum level at 4% in 2008Q1 at the start of the global financial crisis, while the lowest personal rate of inflation is -2.71% in 2009Q2. The actual heterogeneity is expected to be even larger as our calculated rate of inflation does not capture differences in price changes caused by different products being consumed *within* consumption categories.

Note that $IndINFL_{i,t}$ is a quarterly average of monthly inflation rates, so that annualised values can be obtained by accumulating quarterly inflation. Figure 3 depicts the evolution of yearly $IndINFL_{i,t}$ over time, showing there to be larger variation across individuals and over time, with the highest yearly inflation at 13.8% and the lowest at 3.7%. Figure 3 clearly shows that the heterogeneity in inflation is greater during periods when there are large price changes, so it matters more during more turbulent periods.

Figure 3: Personal inflation rates.

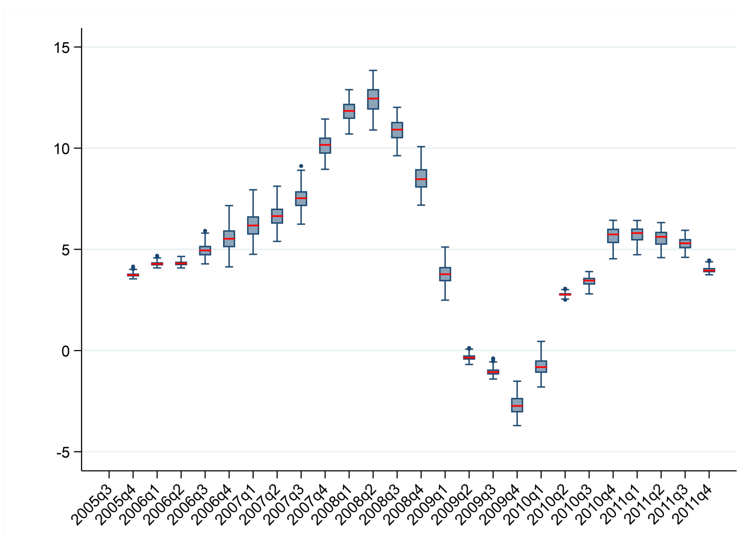

Notes: The figure plots the dynamics of the mean, maximum and minimum values of the yearly personal inflation. Personal inflation is defined as a weighted average of consumption categories inflation sub-indexes with individuals' consumption shares used as weights.

Figure 4: Personal inflation rates by age and income groups

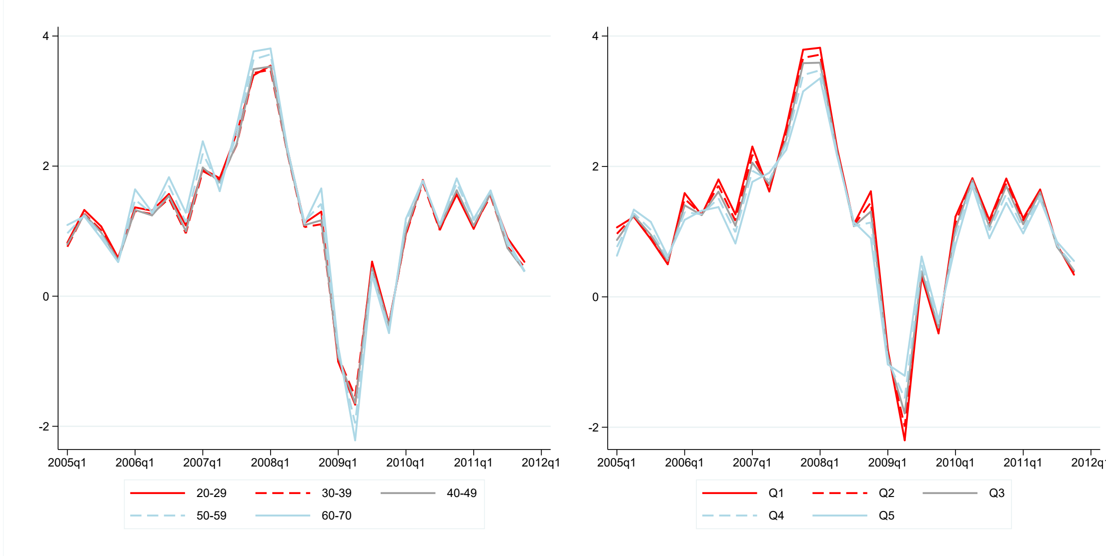

Notes: The figure plots the dynamics of the personal inflation rates broken down by age groups (right panel) and income quintiles (left panel).

Figure 4 shows the dynamics of average personal inflation in income quintiles for five age groups. Low-income groups and elderly age groups experienced higher inflation when prices increased, but were exposed to lower inflation when it was negative in 2009Q2. Differences in personal quarterly inflation reached 0.5 percentage point both for age groups and for income quintiles, while the differences in yearly inflation reached 1.2 points for age groups and 1.4 points for income deciles. Similar patterns in inflation inequality have been found by other studies (Jaravel 2021).

We use the quarterly inflation rate $IndINFL_{i,t}$ as the main independent variable in our regressions to study how personal inflation affects the financial decisions of individuals.

5.4 Two-Way Fixed Effects (TWFE) Estimation

In our baseline model, we use a parsimonious empirical consumption function. The dependent variable in our two-way fixed effects (TWFE) regressions is the quarterly growth rate of real total spending expressed as log differences, $\Delta \log C_{i,t}$, and the main independent variable of interest is group-specific inflation, $IndINFL_{i,t}$. As we impute consumption weights from a limited set of explanatory variables, our independent variable is at a higher level of aggregation than the dependent variable. We consequently run the risk of overstating the statistical significance of our results by erroneously assuming that we are observing more independent individuals than we actually can. To correct this, we cluster standard errors at the level of cells that are defined as interactions of all the time-invariant variables used in the imputation procedure. An alternative would be to run our models at the group level, which would yield similar results but stop us using the rich information from the account-level dataset on the finances of individuals.

We control for changes in real income and financial wealth, which are the main triggers for changes in consumption found in the literature.[12] The final model that we estimate is of the following form:

$$\Delta \log C_{i,t} = \beta_1 \Delta \log \mathrm{INC}_{i,t} + \beta_2 \Delta \log FIN_{i,t-1} + \beta_3 IndINFL_{i,t} + \alpha_i + \gamma_t + \epsilon_{it} \tag{4}$$

With $\Delta \log \mathrm{INC}_{i,t}$ we control for a response by consumption to the change in real income,[13] while the wealth effect is captured by the change in real financial assets $\Delta \log FIN_{i,t-1}$, but we do not have any data on real assets. Inflation in period t denotes a change in the average price level from quarter $t-1$ to quarter t. The lagged change in financial assets depicts a change in wealth by the beginning of quarter t as the balances for financial assets are from the end of each period.

Individual fixed effects α_i are applied to account for any unobserved time-invariant heterogeneity and ensure that our results are for within-person variation over time. Similarly, time fixed effects γ_t account for any developments within a quarter that affect all individuals. Importantly, headline inflation is absorbed by time fixed effects by design, so the estimated coefficient β_3 captures the response of consumption to personal inflation conditional on some level of headline inflation. The TWFE model estimates short-term effects, as given by the Euler equation.

By estimating Model 4 in this way we make two identifying assumptions. The first is that inflation in Estonia is for the most part not driven by short-term domestic demand, which seems reasonable given our discussion in Subsection 4.1. The second is that there are no significant unobserved variables that are correlated with personal exposure to inflation and consumption and that vary over time so that they are not differenced away by the unit fixed effects. This problem could arise if there was, for example, a common unobserved factor such as a demand shock that affects consumption by various social groups differently and so is correlated with both individual consumption and consumption weights. This identification

[12] Higher inflation lowers real income and real wealth, hence affecting also real spending. These channels are controlled by the inclusion of income and wealth variables.

[13] Studies on income expectations add also future income into consumption model to capture advanced information, see Pedroni et al. (2022). As income variable is not the main focus of the paper while the inclusion of future income does not alter the results for inflation, we keep the model parsimonious without lead variables.

challenge essentially lies in the endogeneity of the consumption weights and is equivalent to the parallel trends assumption not being satisfied.

This problem is greatly mitigated by our strategy of using fixed weights so that any changes in the independent variable are due only to price changes, while the weights play the role of a continuous exposure variable. It is nevertheless possible that the TWFE fail to capture the confounding effects of some time-varying unobservable factors. To alleviate this concern, we alternatively estimate a series of recently proposed interactive fixed effects (IFE) models (see Bai 2009 and Vogt et al. 2022).

5.5 Interactive Fixed Effects (IFE) Estimation

Interactive fixed effects (IFE) models generalise the standard TWFE framework and are designed to control for unobserved time-varying heterogeneity. In IFE, the error term is assumed to have a factor structure $\epsilon_{it} = \lambda_i \times f_t + e_{it}$ with f_t representing unobserved time-varying common regressors or factors, and λ_i representing factor loadings, or the differential exposure of cross-sectional units to the factors. In our case, the underlying model is assumed to have an IFE structure:

$$\Delta \log C_{i,t} = \beta_0 + \beta_1 IndINFL_{it} + \beta_2 X_{it} + \lambda_i \times f_t + e_{it}, \tag{5}$$

where $\lambda_i \times f_t$ and e_{it} are unobserved. Taken together, $\lambda_i \times f_t$ represents a confounding variable that varies along both dimensions. The standard panel TWFE model is a special case of IFE for $f = 1$, and so interactive fixed effects allow unobserved heterogeneity to be controlled for more flexibly than in standard TWFE models.

That factors f_t (and their loadings) are unobserved and time-varying poses a more profound estimation challenge than that of time-invariant unobserved variables in TWFE settings. Several approaches have been proposed for obtaining consistent estimates of parameters in such frameworks (e.g., Bai 2009, Callaway & Karami 2023). The IFE methods effectively use data-driven approaches to back out, and control for, the main unobserved factors. The approach of Bai (2009) treats f_t as parameters to be estimated, which is the *controlling by estimating* approach. Bai (2009) proposes an iterative procedure whereby an optimal $\widehat{\beta}$ in a given iteration can be obtained given a pair $(\widehat{f_t}, \widehat{\lambda_i})$ from the previous step and *vice versa* and shows that this procedure allows for consistently estimating the model β even though f_t and λ_i are unobserved.

However, this method requires the number of time periods to be large. Since our sample period covers quite a short period of 27 quarters, we complement this exercise with the method recently proposed by Vogt et al. (2022), which allows for small T settings.

Instead of estimating the vector of coefficients and the unobserved factors simultaneously, the method of Vogt et al. (2022) takes the idea of eliminating the factors from the model by transforming it with an appropriate projection matrix (see also Pesaran 2006, who introduces this idea in a low dimensional case).

We present the results of our baseline specification using both the standard TWFE and the IFE. We show that both approaches give very similar results, indicating that unobserved factors are not a major concern in our case. For that reason, we default in our subsequent analysis to using the TWFE because of its simplicity and because it is easier to interpret.

6 Results

In this section, we present and explain the results of our analysis. We first discuss the results of the baseline specification and explore how the results vary in the cross-section. We then investigate in Section 6.2 how other components of the household balance sheet respond to the inflation experienced. Lastly we look at how indebtedness impacts the responses of individuals.

6.1 Baseline model: Effects on Consumption

Table 2 presents the results of the estimations of Equation 4 (TWFE) and Equation 5 (IFE). Personal inflation has a strong positive relationship with real consumption growth. The first column includes the results for a linear model, and the coefficient for IndINFL is 0.014, which means that an increase of one percentage point in mean quarterly personal inflation increases real consumption by 1.4%. The second column in Table 2 shows the estimated coefficients in the model, allowing for a non-linear relationship between inflation and spending, including a quadratic term. We find that the coefficient for the quadratic term is 0.005, which is positive and significant, indicating that individuals respond to price changes stronger at higher levels of inflation.

The positive quadratic term reflects the idea that individuals might take little notice of inflation when it is at a sufficiently low level, and pay more attention to it when it is at a high level. This is in line with the findings of Cavallo et al. (2017), that individuals have weaker priors about inflation in low inflation environments. Such behaviour is consistent with rational inattention, which argues that individuals only pay attention to inflation when it gets sufficiently high.

Table 2: The effect of inflation exposure on consumption growth

	TWFE		IFE Vogt et al. (2022)	IFE Bai (2009)		
	(1)	(2)	(3)	(4)	(5)	(6)
IndINFL_{it}	0.014***	0.007**	0.015***	0.014***	0.014***	0.011***
	(0.003)	(0.003)	(0.004)	(0.003)	(0.003)	(0.003)
IndINFL^2_{it}		0.005***				
		(0.001)				
$\Delta \log \text{INC}_{it}$	0.481***	0.481***	0.490***	0.515***	0.517***	0.525***
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
$\Delta \log \text{FIN}_{it-1}$	0.108***	0.108***	0.107***	0.106***	0.105***	0.104***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Individual FE	Yes	Yes	X	X	X	X
Time FE	Yes	Yes	X	X	X	X
Number of factors	X	X	1	1	2	3
Observations	2,352,678	2,352,678	1,979,235	1,781,325	1,781,325	1,781,325
R^2	0.200	0.200				

Notes: The table presents the results of the baseline model regressing total individual spending growth on personal inflation exposure. The dependent variable is the log difference of real total spending between two consecutive quarters. All the variables except inflation are in log differences. Columns (1) and (2) present the baseline results using two-way fixed effects (TWFE) estimation. In column (3), we estimate the interactive fixed effects (IFE) model using the approach of Vogt et al. (2022), where the method chooses the optimal number of factors. In columns (4)-(6), we estimate the IFE model with the approach of Bai (2009) with different manually chosen numbers of factors. The sample in column (3) is restricted to be a balanced panel. Due to computational complexity, columns (4)-(6) use a random 90% of that sample. Standard errors clustered at the group level are presented in parentheses and in column (3) are calculated with a block-bootstrap procedure. Groups are the same as used in the imputation model. *, **, *** denote statistical significance at the 10%, 5%, and 1% level.

When we replace the TWFE model with the IFE model by Vogt et al. (2022), which chooses the optimal number of time-varying and cross-sectionally varying factors, we find that personal inflation has a marginal effect of 1.5% on consumption growth (Column 3), where the TWFE found 1.4%. Using the IFE model proposed by Bai (2009) instead and manually choosing 1, 2 or 3 factors (Columns (4)-(6)), gives estimates of 1.4%, 1.4%, and 1.1% respectively. All 4 investigations suggest that using the simpler TWFE methodology does not lead to significantly biased estimates, and so we can proceed with TWFE estimators alone in the rest of the analysis.

Overall, those estimates in Table 2 are consistent with Hypothesis 1a, which is that individuals respond to greater exposure to personal inflation by increasing their current consumption. They are, however, not consistent with either Hypothesis 1 that people respond only to headline inflation, or with Hypothesis 1b that people respond but on average reduce their consumption in response to experiencing higher personal inflation.

The most plausible explanation for this finding seems to be the expectations channel, although our core estimates do not provide directly any channel. We note the empirical evidence that inflation expectations have the strongest correlation with current inflation rather than with previous or future inflation (Binder & Kamdar 2022).[14] We conjecture that individuals who are experiencing higher inflation increase their expectations for future inflation proportionately more and so they increase their current consumption. To support this explanation, we carry out an additional exercise using survey data from the series of business and consumer surveys of the European Central Bank (ECB), which include qualitative measures of the 12-month ahead inflation expectations of Estonian households, and their 12-month backwards perceptions of inflation divided by several demographic characteristics. We show in Table 3 that there is a strong positive correlation between perceptions and expectations, even after controlling for cohort and time fixed effects, and this supports our reasoning.

[14] The expectation channel can consequently not be tested just by including future inflation in the model, as also noted by Rudd & Whelan (2005). However, we have estimated the model after adding future inflation, and the estimated coefficient for the inflation in the next period is either slightly negative or insignificant, depending on the model specification, suggesting that actual inflation exposure in period $t + 1$ indeed does not reflect inflation expectations in period t (the results are not reported but are available upon request).

Table 3: Inflation perceptions and expectations

	Dependent variable: 12 month ahead expectations			
	Age	Income	Gender	Education
12 month perception	0.437***	0.604***	0.385***	0.156
	(0.115)	(0.075)	(0.075)	(0.106)
Time FE	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes
Observations	104	104	104	104

Notes: This table presents the results from regressing the qualitative 12-month ahead inflation expectations score on 12 month back qualitative inflation perceptions score. Both measures come from European Central Bank's business and consumer surveys data series for Estonia. As the data are reported along several cross-sectional splits, we present the result for each split in a separate column. Robust standard errors are presented in parentheses. *, **, *** denote statistical significance at the 10%, 5%, and 1% level.

If the intertemporal substitution channel is indeed the driver of our findings, as suggested by the exercises above, the effect should be particularly strong for people who have the financial capacity to bring their consumption forward. To verify this, we augment our model by including an interaction between personal inflation and a ratio of liquid wealth to annual income (Table 4). We find that individuals respond more strongly to personal inflation when they have larger liquid reserves, and the amount of the most liquid assets in the form of cash balances on checking accounts matter most. At the same time, the interaction term is smaller but still significant for balances on term deposits. The magnitude of the coefficients suggests that the relationship between individual inflation and consumption is negative when individuals do not have any liquid buffers on their checking account. When they do not have any resources on term deposits, they can still increase their consumption, because households usually open term deposits when they already have sufficient buffers on their checking accounts. This suggests that we are observing liquidity constraints when individuals cannot finance to bring consumption forward. Overall, this exercise further supports our preferred interpretation that our results are driven by the expectations channel.

Table 4: Liquid assets and consumption responses

	Checking account	Term deposit	Total
IndINFL_{it}	-0.011^{**}	0.012^{***}	-0.004
	(0.004)	(0.003)	(0.003)
$\text{IndINFL}_{it} \times (\frac{LiqAssets}{Yincome})_{t-1}$	0.083^{***}	0.007^{***}	0.033^{***}
	(0.010)	(0.001)	(0.003)
$\Delta \log \text{INC}_{it}$	0.480^{***}	0.479^{***}	0.479^{***}
	(0.006)	(0.006)	(0.006)
$\Delta \log \text{Fin}_{i,t-1}$	0.107^{***}	0.109^{***}	0.108^{***}
	(0.001)	(0.001)	(0.001)
Individual FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
Observations	2,261,849	2,261,849	2,261,849

Notes: The table presents the results of fixed effect models estimation. In all columns, the dependent variable is the (log) real consumption growth. Relative to the baseline model, we also include an interaction term between individual inflation and a ratio of liquid assets to yearly income (lagged by one period). In each column, we limit our attention to balances on a specific type of account: checking accounts, term deposits, and the total of the two. FE stands for additive fixed effects. Standard errors are clustered at the group level and they are presented in parentheses. *, **, *** denote statistical significance at the 10%, 5%, and 1% level. Clustering groups are the same as those used in the imputation model.

As noted in Subsection 5.3, individuals in lower income groups and older age groups are exposed to higher inflation, so it is important to understand the extent to which the differences feed into consumption responses. To this end, we define five time-invariant income quintiles based on the average income of individuals over the sample period, and we rerun the baseline model after adding interactions of income groups with all the regressors. The right panel in Figure A3 in the Appendix shows the results of the estimations for the interaction with personal inflation. Similarly, we interact five age groups for age at the beginning of the sample with the regressors, and the interaction with inflation is shown in the left panel in Figure A3. The coefficients are similar across income and age groups, suggesting that all the groups respond similarly to the change in personal inflation, so that inflation heterogeneity across the groups that we observed in Figure 4 is not amplified or alleviated when it is fed into the economy. The results suggest that the aggregate dynamics of consumption are not driven by any specific groups.

6.2 Effects on Savings, Loans and Investments

The positive response from consumption raises the question of how the increased consumption is financed. Individuals might increase their borrowing to get additional funds, and they might aslo reduce their savings to free up additional resources for increased consumption. Our analysis shows that a combination of these two factors is at work. Large exposure to personal inflation seems to encourage new consumer borrowing and makes individuals reduce the amounts they hold in bank accounts, and more specifically in term deposits.

We can show that personal inflation is positively related to increased use of consumer loans and overdrafts. We estimate a panel logit model in which the dependent variable is a dummy variable indicating an increase in the volume of consumer debt (Consumer Loan), overdrafts (Overdraft) or housing loans (Housing Loan), covering both new loans shown by an increase from 0 in the previous quarter, and additional borrowing shown as an increase in the volume of loans.[15]

$$Logit(NewLoan_{it}) = \beta_1 IndINFL_{it} + \beta_3 X_{it} + \psi \overline{IMR_{it}} + \alpha_i + \gamma_t + \epsilon_{it} \tag{6}$$

We control for changes in income and financial assets as we did in the baseline model, but we also control for the selection to debt ownership by including the inverse Mills ratio IMR_{it}. Only a subset of individuals hold the various loans and the selection depends on individual preferences, and perceived and actual credit constraints, and possibly also on other unobserved factors. We follow a two-step procedure and first estimate the probability of owning each type of loan in each quarter, allowing the probability to vary across time. We calculate the inverse Mills ratio for each individual as $IMR_{it} = \phi(Z)/\Phi(Z)$, where $\phi(Z)$ is the probability density function and $\Phi(Z)$ is the cumulative density function. The vector of regressors Z in the probability function contains the second-step regressors together with categorical variables for income deciles and age groups, and dummy variables for gender, nationality, and the capital region. The two-step procedure is estimated by bootstrap with 1000 replications.

The model is estimated only for individuals who experienced a change in the dependent variable over the sample period, meaning that those who have not taken a new or additional loan are excluded from the estimations.

[15] As we do not have data on the balance of credit cards, we cannot investigate financing by this loan type.

Table 5: Taking out new or additional loan - odds ratios

	Add Housing Loan	Add Consumer Loan	Add Overdraft
IndINFL$_{it}$	0.949	1.142***	1.055***
	(0.057)	(0.053)	(0.017)
Δ log INC$_{it}$	1.199***	0.987	0.636***
	(0.020)	(0.015)	(0.004)
Δ log FIN$_{it-1}$	1.005***	0.989***	0.978***
	(0.006)	(0.004)	(0.002)
IMR$_{it}$ (Mortgage)	12.978***		
	(2.822)		
IMR$_{it}$ (Consumer Loan)		7.017***	
		(1.161)	
IMR$_{it}$ (Overdraft)			1.769***
			(0.084)
Individual FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
Obs. in 2nd step	231,025	207,067	693,707

Notes:The table presents the results of fixed effect conditional logit models estimation where the probability model is estimated in the first step to compute inverse Mills Ratios for debt ownership. The dependent variables are dummies for new loans in the form of a housing or a consumer loan. Income and wealth changes are in log differences. FE stands for additive fixed effects. The standard errors are obtained from bootstrapping with 1000 replications. *, **, *** denote statistical significance at the 10%, 5%, and 1% level.

Table 5 indicates a strong positive relationship between the propensity to take out new consumer loans or overdrafts and personal exposure to inflation, suggesting that higher inflation triggers borrowing for consumption. The estimated inverse Mills ratio is statistically significant and positive in the models for all loan types, indicating that the tendency to take new or additional debt is stronger among debtors than it would be in the total sample.

We also investigate the use of savings products, and we observe consumers' balances in their checking, term, and investment accounts. To analyse what savings decisions people make to finance the increased consumption driven by their personal inflation exposure, we separately estimate four TWFE models with changes in the log volumes kept in checking accounts, term deposits, investment funds, stocks, and bonds as our dependent variables and personal inflation as the main independent variable:

$$\Delta log(BalAccount)_{it} = \beta_1 IndINFL_{it} + \beta_2 X_{it} + \psi \overline{IMR_{it}} + \alpha_i + \gamma_t + \epsilon_{it}, \tag{7}$$

where $BalAccount$ denotes the balance on the checking account or term account, or the volume of investment funds, stocks or bonds. We also control for changes in income and for balances in deposits, and for selection into the group that holds savings on the accounts or investments by including the inverse Mills ratio IMR_{it}. Not all individuals own term deposits, investment funds, stocks, and bonds and the ownership of those products is not random, so we use a two-step procedure similar to the one we used for ownership of different types of loans. In each case, the subsample used in the second-step regression consists of observations where the balance on the account or investment concerned is positive, indicating selection into the group.

The results presented in Table 6 indicate that personal inflation is not related to checking account balances. All transactions are executed from checking accounts, but inflation-induced consumption is not explicitly related to holdings in these accounts. However, personal inflation appears to have a strong negative effect on the balances kept in term deposits, which seem to be the main source of funding besides consumer loans and overdrafts for increased consumption spending. An increase in quarterly personal inflation of one percentage point translates into a fall of roughly 12% in the balances in term deposits. Individuals seem to prefer to keep a certain buffer on their checking account while using resources from term deposits to finance their additional consumption. This is in line with the finding that individuals respond more strongly to consumption when they have more liquidity on term deposits.

Lastly, we find that personal inflation has a positive and statistically significant effect on stock investments, indicating that those who are more exposed to inflation increase their share of riskier investment. This is consistent with the results of Agarwal et al. (2022), who postulate that this effect arises because the returns on risky assets are more responsive to inflation than deposit rates are. By this reasoning, individuals who expect prices to rise even further in future decide to rebalance their portfolios in order to hedge against inflation. The result confirms that households view stock investments as a kind of inflation hedging strategy. At the same time, the result for investments is at odds with the recent research by Braggion et al. (2023), who show in a historical setting that local inflation negatively affects net investments, and interpret this as a symptom of the money illusion (Deaton 1977).

Table 6: Experienced inflation and the balance of savings products

	Checking acc.	Term deposit	Inv. funds	Bonds	Stocks
IndINFL_{it}	0.008	-0.119^{***}	0.050	-0.079	0.093^{***}
	(0.010)	(0.020)	(0.035)	(0.491)	(0.018)
$\Delta \log \text{INC}_{it}$	0.952^{***}	0.307^{***}	0.116^{***}	-0.114	0.068^{***}
	(0.012)	(0.009)	(0.016)	(0.091)	(0.009)
$\Delta \log \text{Term}_{i,t-1}$	0.047^{***}		0.009^{**}	0.012	0.005
	(0.001)		(0.005)	(0.021)	(0.003)
$\Delta \log \text{Check}_{i,t-1}$		0.135^{***}	0.062^{***}	-0.012	0.026^{***}
		(0.002)	(0.005)	(0.035)	(0.003)
$\Delta \log \text{Sec}_{i,t-1}$	0.023^{***}	0.008			
	(0.003)	(0.006)			
IMR_{it} (Term dep.)		0.376^{***}			
		(0.080)			
IMR_{it} (Funds)			0.656^{***}		
			(0.159)		
IMR_{it} (Bonds)				-0.138	
				(0.293)	
IMR_{it} (Stocks)					1.047^{***}
					(0.124)
Individual FE	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes
Observations in 2nd step	2,351,269	301,180	46,339	846	72,507

Notes: The table presents the results of fixed effect models estimation where the probability model is estimated in the first step to compute inverse Mills Ratios for savings in columns (2)-(4). The dependent variables are log changes in the real balance in checking accounts, term deposits, holdings of stocks and bonds. FE stands for additive fixed effects. The standard errors are obtained from bootstrapping with 1000 replications. *, **, *** denote statistical significance at the 10%, 5%, and 1% level.

We do not see a significant relationship between personal inflation and investments in bonds or investment funds. There were next to no government bonds in Estonia during the sample period, so the ownership rate of bonds was very low, and we consequently do not get any clear results for the financial choices about bonds. Investment funds typically contain both bonds and stocks and we cannot distinguish different types of funds. As it was for the regressions on debt, the inverse Mills ratio is statistically significant in all the regressions where it is included, confirming the need to control for selection.

6.3 Effects by Indebtedness Levels

Inflation affects savers and debtors differently, as higher inflation reduces the real value of assets but also reduces the real value of liabilities. Indebted agents are net beneficiaries from higher inflation because of debt depreciation, and it could be hypothesised that the mechanism may positively affect their consumption response to inflation (Lieb & Schuffels 2022). Equally though, recent research shows that individuals are rarely aware of the debt-erosion channel of inflation (Schnorpfeil et al. 2023), and the cost of debt servicing might crowd out consumption when the budget constraint is binding (see e.g. Dynan et al. 2012, Toczynski 2023, Kukk 2016), and so the net effect of debt in our case is not initially obvious.

To test the mechanism linking the strength of the inflation sensitivity of consumption to the level of indebtedness, we re-estimate the baseline model after adding a series of interactions of the personal inflation with debt ratios for total debt, housing debt and consumer debt, and the debt servicing ratio (DSR), while also controlling for the direct effect of debt on consumption. As before, we control for selection into holding a particular type of debt by including the inverse Mills ratio in the model. The two-step procedure is again estimated with 1000 bootstrap replications.

Table 7 presents the results, where each column includes an interaction term with three debt ratios. As expected, debt in itself is associated with lower consumption, which is consistent with the crowding-out mechanism. The interaction term between personal inflation and each type of debt is negative and statistically significant for total indebtedness and more clearly for indebtedness from consumer loans, indicating that the positive response of consumption is hampered by higher levels of indebtedness. This suggests that a large debt acts as a binding constraint for the ability to bring consumption forward and the mechanism prevails over the positive debt depreciation effect. The negative and significant interaction with the debt servicing ratio is similarly consistent with this interpretation, implying there is a smaller consumption response when regular debt repayments are larger. The suppressed response of consumption can be explained by the lack of liquidity and the presence of constraints on further borrowing, particularly on borrowing for consumption purposes.

Table 7: Indebtedness and consumption responses

	Total	Housing	Consumer Loan	DSR
IndINFL_{it}	0.023^{***}	0.023^{***}	0.089^{***}	0.030^{***}
	(0.003)	(0.007)	(0.008)	(0.003)
$\Delta \log \text{INC}_{i,t}$	0.535^{***}	0.514^{***}	0.616^{***}	0.535^{***}
	(0.002)	(0.004)	(0.006)	(0.002)
$\Delta \log \text{FIN}_{it-1}$	0.080^{***}	0.102^{***}	0.076^{***}	0.081^{***}
	(0.001)	(0.002)	(0.001)	(0.001)
ΔDebt_{it-1}	-0.065^{***}	-0.038^{***}	-0.124^{***}	-0.065^{***}
	(0.001)	(0.002)	(0.002)	(0.001)
$\text{IndINFL}_{it} \times (\frac{TotDebt}{Yincome})_{it-1}$	-0.003^{*}			
	(0.0015)			
IMR_{it} (Total debt)	0.036^{***}			0.031^{***}
	(0.011)			(0.011)
$\text{IndINFL}_{it} \times (\frac{Housing}{Yincome})_{it-1}$		-0.002		
		(0.001)		
IMR_{it} (Housing)		0.046^{**}		
		(0.016)		
$\text{IndINFL}_{it} \times (\frac{ConsLoan}{Yincome})_{it-1}$			-0.184^{***}	
			(0.010)	
IMR_{it} (Consumer Loan)			0.131^{***}	
			(0.023)	
$\text{IndINFL}_{it} \times (\frac{DSR}{Yincome})_{it-1}$				-0.039^{***}
				(0.001)
Individual FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Observations in 2nd step	1,029,230	359,715	213,326	1,074,958

Notes: The table presents the results of estimations featuring interactions of experienced inflation and indebtedness. The dependent variable is (log) consumption growth. In columns (1)-(3) we include interactions with debt ratios of total debt, housing debt, and consumer debt respectively. Income and wealth changes are in log differences. DSR stands for debt servicing ratio, Yincome for yearly income, and IMR for the inverse Mills ratio for debt holdings. FE stands for additive fixed effects. The standard errors are obtained from bootstrapping with 1000 replications. *, **, *** denote statistical significance at the 10%, 5%, and 1% level.

7 Robustness Checks

To test how informative our imputation strategy is for measuring a person's exposure to inflation we randomly divide the HBS sample used for estimating the shares model into two equally sized groups, which are an estimation sample and a testing sample. We estimate the coefficients of the shares using the estimation sample and then impute them onto the testing sample. We subsequently calculate the actual personal inflation and the imputed inflation for the testing sample and compare the two by regressing the actual level on the imputed one.

Table A3 presents the results and shows the coefficient for imputed inflation is 1.005 and highly statistically significant, which means that the imputed and actual measures of experienced inflation co-move very closely on average. One problem with testing the quality of fit in our case is that the time dimension of the data is likely to drive most of the relationship. We therefore control for time fixed effects in the second column. The coefficient for imputed inflation is still close to one and highly statistically significant, which means that the relationship between the two measures is also strong in the cross-sectional dimension.

To get a more detailed measure of how much of the variation in personal exposure to inflation in the cross-section is explained by imputed inflation, we need to get rid of the time dimension. We repeat the exercise for each of the time points separately and look at the R^2 of the regressions. As expected, the R^2s fluctuate across time with values around 0.2 for periods when there is a lot of heterogeneity in the inflation rates across different consumption categories, and low values when the inflation rates were similar for all categories. The per-period regressions in the out-of-sample exercises have an average R^2 of 0.084, which is the same as the in-sample result. However, we mainly care about the R^2 at times of wide inflation heterogeneity, as it is supposed to be low by construction in the other periods.

Importantly for our analysis, we do not find any patterns in the imputation error that is defined as the difference between the actual and the imputed rates of personal inflation. Regressing the difference on the imputed inflation rate (with time fixed effects, not reported) yields insignificant results, showing that the imputation errors are on average randomly distributed.

Overall, the tests suggest that the imputation strategy is informative for the true personal inflation and on average captures the individual exposure to inflation quite well. While our strategy allows us to explain only a fraction of the overall variation, the imputation errors seem to be reasonably random, allowing us to use the size of our main dataset to overcome the high noise-to-signal ratio. In either case, measurement errors should, if anything, depress our coefficients towards zero.

We run another sensitivity analysis for the imputation by comparing the results found from Equation 4 when different approaches are used for imputing the consumption shares. The results in Table A5 in the Appendix are very similar to those from the baseline estimations.

Lastly, although the TWFE model assumes stationarity for the variables, there is usually no discussion about the stationarity of the variables because the time period of the microdata is generally short. Our dataset has a relatively long time dimension though of up to 28 quarters, and so ignoring its time-series properties could lead to spurious results even after first differencing. To avoid this, we explicitly test for non-stationarity and can reject it for all variables.[16]

8 Conclusion

We have used a rich anonymised account-level dataset from a major commercial bank in Estonia that covers an episode of high and variable inflation. We investigate the responses to personal exposure to inflation, which is computed from personal consumption categories and category inflation rates. Our study provides robust evidence that personal inflation exposure affects personal consumption beyond the effect of headline inflation.

More specifically, we find that individuals respond to greater exposure to inflation by increasing their contemporaneous real spending, which is consistent with the motive of intertemporal substitution when individuals form inflation expectations from their personal experience of inflation. Moreover, the effect is non-linear and increases with the level of inflation, which is in line with the idea that individuals notice significant price movements more.

Looking beyond consumption, we also analyse how this extra consumption is financed. We find that part of this extra consumption is financed through a reduction in savings on term deposits and another part by increased use of consumer loans and overdrafts. Individuals with more liquid wealth, particularly in term deposits, respond strongly to personal inflation, while liquidity-constrained individuals cannot bring their consumption forward.

Going even further, we find that the response of consumption to personal inflation does not vary significantly across income groups or age groups. Although net debtors should

[16] We implement two panel data unit root tests that are suitable for a panel with N$\rightarrow \infty$; these are the Harris-Tzavalis (HT) test and Im-Pesaran-Shin (IPS) test. The HT test assumes a common autoregressive parameter and requires a balanced dataset, while the IPS test allows heterogeneous AR parameters and permits gaps in the beginning or at the end of the period but no gaps within a panel. We can therefore only run the test for data series with no missing observations. We also take cross-sectional dependence into account. The results of the tests are provided in the Appendix in Table A4, which shows that the hypothesis of non-stationarity is rejected for all the variables.

benefit more from inflation, we do not find evidence of a positive debt depreciation effect. Rather the opposite, as the front-loading of consumption is hampered by higher levels of indebtedness, suggesting that a lack of liquidity and credit constraints on further borrowing play a role.

Our results support the idea that inflation heterogeneity matters for welfare and for predicting consumption responses, which in turn are likely to feed back into future inflation. Moreover, the liquidity and indebtedness determine how consumers respond to the inflation they experience. On these grounds, policy-makers in central banks and beyond may want to give a more prominent role to the distribution of inflation across consumption categories and to the actual inflation experienced across population sub-groups.

References

Aastveit, K. A., Bjørnland, H. C. & Thorsrud, L. A. (2016), 'The world is not enough! small open economies and regional dependence', *The Scandinavian Journal of Economics* **118**(1), 168–195.

Agarwal, S., Chua, Y. H., Ghosh, P. & Song, C. (2022), 'Inflation expectations and portfolio rebalancing of households: Evidence from inflation targeting in india', *Available at SSRN 4069564* .

Bai, J. (2009), 'Panel data models with interactive fixed effects', *Econometrica* **77**(4), 1229–1279.

Baker, S. R. (2018), 'Debt and the response to household income shocks: Validation and application of linked financial account data', *Journal of Political Economy* **126**(4), 1504–1557.

Binder, C. & Kamdar, R. (2022), 'Expected and realized inflation in historical perspective', *Journal of Economic Perspectives* **36**(3), 131–155.

Braggion, F., von Meyerinck, F. & Schaub, N. (2023), 'Inflation and Individual Investors' Behavior: Evidence from the German Hyperinflation', *The Review of Financial Studies* .

Branson, W. H. & Klevorick, A. K. (1969), 'Money illusion and the aggregate consumption function', *The American Economic Review* **59**(5), 832–849.

Burke, M. A. & Ozdagli, A. (2023), 'Household inflation expectations and consumer spending: evidence from panel data', *Review of Economics and Statistics* **105**(4), 948–961.

Callaway, B. & Karami, S. (2023), 'Treatment effects in interactive fixed effects models with a small number of time periods', *Journal of Econometrics* **233**(1), 184–208.

Cavallo, A., Cruces, G. & Perez-Truglia, R. (2017), 'Inflation expectations, learning, and supermarket prices: Evidence from survey experiments', *American Economic Journal: Macroeconomics* **9**(3), 1–35.

Coibion, O., Georgarakos, D., Gorodnichenko, Y. & van Rooij, M. (Forthcoming), 'How does consumption respond to news about inflation? field evidence from a randomized control trial', *American Economic Journal: Macroeconomics* .

Coibion, O. & Gorodnichenko, Y. (2015), ‘Is the phillips curve alive and well after all? inflation expectations and the missing disinflation’, *American Economic Journal: Macroeconomics* **7**(1), 197–232.

Coibion, O., Gorodnichenko, Y. & Weber, M. (2022), ‘Monetary policy communications and their effects on household inflation expectations’, *Journal of Political Economy* **130**(6), 1537–1584.

Cookson, J. A., Gilje, E. P. & Heimer, R. Z. (2022), ‘Shale shocked: Cash windfalls and household debt repayment’, *Journal of Financial Economics* **146**(3), 905–931.

Cuestas, J. C., Lucotte, Y. & Reigl, N. (2020), ‘Banking sector concentration, competition and financial stability: the case of the baltic countries’, *Post-Communist Economies* **32**(2), 215–249.

de Bruin, W. B., Van der Klaauw, W., Topa, G., Downs, J. S., Fischhoff, B. & Armantier, O. (2012), ‘The effect of question wording on consumers’ reported inflation expectations’, *Journal of Economic Psychology* **33**(4), 749–757.

Deaton, A. (1977), ‘Involuntary saving through unanticipated inflation’, *The American Economic Review* **67**(5), 899–910.

Dovern, J., Müller, L. S. & Wohlrabe, K. (2023), ‘Local information and firm expectations about aggregates.’, *Journal of Monetary Economics* .

Du Caju, P., Périlleux, G., Rycx, F. & Tojerow, I. (2023), ‘A bigger house at the cost of an empty stomach? the effect of households’ indebtedness on their consumption: micro-evidence using belgian hfcs data’, *Review of Economics of the Household* **21**(1), 291–333.

Dynan, K., Mian, A. & Pence, K. M. (2012), ‘Is a household debt overhang holding back consumption?’, *Brookings Papers on Economic Activity* pp. 299–362.

D’Acunto, F., Hoang, D., Paloviita, M. & Weber, M. (Forthcoming), ‘Iq, expectations, and choice’, *Review of Economic Studies* .

D’Acunto, F., Hoang, D. & Weber, M. (2022), ‘Managing households’ expectations with unconventional policies’, *The Review of Financial Studies* **35**(4), 1597–1642.

D’Acunto, F., Malmendier, U., Ospina, J. & Weber, M. (2021), ‘Exposure to grocery prices and inflation expectations’, *Journal of Political Economy* **129**(5), 1615–1639.

Georganas, S., Healy, P. J. & Li, N. (2014), 'Frequency bias in consumers perceptions of inflation: An experimental study', *European Economic Review* **67**, 144–158.

Hall, R. E. (1988), 'Intertemporal substitution in consumption', *Journal of Political Economy* **96**(2), 339–357.

Hobijn, B. & Lagakos, D. (2005), 'Inflation inequality in the united states', *Review of Income and Wealth* **51**(4), 581–606.

Jaravel, X. (2021), 'Inflation inequality: Measurement, causes, and policy implications', *Annual Review of Economics* **13**, 599–629.

Jovičić, G., Kunovac, D. et al. (2017), What is driving inflation and gdp in a small european economy: the case of croatia, Technical report.

Kamdar, R. et al. (2018), 'The inattentive consumer: Sentiment and expectations', *Manuscript. https://rupalkamdar. github. io/pdfs/Inattentive Consumer. pdf* .

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

Kuchler, T. & Zafar, B. (2019), 'Personal experiences and expectations about aggregate outcomes', *The Journal of Finance* **74**(5), 2491–2542.

Kukk, M. (2016), 'How did household indebtedness hamper consumption during the recession? evidence from micro data', *Journal of Comparative Economics* **44**(3), 764–786.

Kukk, M., Paulus, A. & Staehr, K. (2020), 'Cheating in europe: Underreporting of self-employment income in comparative perspective', *International Tax and Public Finance* **27**(2), 363–390.

Lieb, L. & Schuffels, J. (2022), 'Inflation expectations and consumer spending: the role of household balance sheets', *Empirical Economics* **63**(5), 2479–2512.

Maćkowiak, B. (2007), 'External shocks, us monetary policy and macroeconomic fluctuations in emerging markets', *Journal of Monetary Economics* **54**(8), 2512–2520.

Malmendier, U. (2021), 'Fbbva lecture 2020 exposure, experience, and expertise: Why personal histories matter in economics', *Journal of the European Economic Association* **19**(6), 2857–2894.

Malmendier, U. & Nagel, S. (2016), 'Learning from inflation experiences', *The Quarterly Journal of Economics* **131**(1), 53–87.

Malmendier, U. & Wachter, J. A. (2021), 'Memory of past experiences and economic decisions', *Available at SSRN 4013583* .

Michael, R. T. (1979), 'Variation across households in the rate of inflation', *Journal of Money, Credit and Banking* **11**(1), 32–46.

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

Patinkin, D. (1965), Money, interest, and prices; an integration of monetary and value theory, Technical report.

Pedroni, M., Singh, S. & Stoltenberg, C. A. (2022), *in* 'Advance Information and Consumption Insurance: Evidence and Structural Estimation'.

Pesaran, M. H. (2006), 'Estimation and inference in large heterogeneous panels with a multifactor error structure', *Econometrica* **74**(4), 967–1012.

Rudd, J. & Whelan, K. (2005), 'New tests of the new-keynesian phillips curve', *Journal of Monetary Economics* **52**(6), 1167–1181.

Schnorpfeil, P., Weber, M. & Hackethal, A. (2023), Households' response to the wealth effects of inflation, Technical report, National Bureau of Economic Research (WP 31672).

Shafir, E., Diamond, P. & Tversky, A. (1997), 'Money illusion', *The Quarterly Journal of Economics* **112**(2), 341–374.

Toczynski, J. (2023), Spend or invest? analyzing mpc heterogeneity across three stimulus programs, Technical report, Swiss Finance Institute.

Vogt, M., Walsh, C. & Linton, O. (2022), 'Cce estimation of high-dimensional panel data models with interactive fixed effects', *arXiv preprint arXiv:2206.12152* .

Weber, M., Gorodnichenko, Y. & Coibion, O. (2022), 'The expected, perceived, and realized inflation of us households before and during the covid19 pandemic', *IMF Economic Review* pp. 1–43.

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

Appendix

Table A1: Annual inflation rates by category

	2005	2006	2007	2008	2009	2010	2011
Food and non-alcoholic beverages	3.5	5.1	9.3	14.2	−4.0	3.0	9.7
Alcohol and tobacco	4.2	3.4	4.1	15.9	10.7	4.5	6.3
Clothing and footwear	1.9	2.6	3.6	3.8	0.8	2.5	3.5
Housing	7.0	10.4	14.6	15.8	1.0	2.9	5.9
Household goods	0.3	2.1	3.5	4.4	2.5	−1.1	0.7
Health	2.2	2.5	7.9	7.6	4.2	0.8	0.4
Transport	9.3	4.4	2.7	11.1	−6.5	7.1	5.17
Communication	−4.1	−5.3	−1.2	−0.9	−0.2	3.7	−4.2
Recreation, entertainment	0.7	4.2	3.3	2.4	−0.2	−0.6	0.7
Education	3.4	3.5	5.4	8.5	4.7	1.7	2.2
Hotels, cafes, restaurants	4.0	4.4	10.1	13.3	0.5	−1.0	6.0
Miscellaneous goods and services	2.9	4.1	4.7	8.9	7.6	1.5	2.84
Total	4.1	4.4	6.6	10.4	−0.1	3.0	5.09
Max	9.3	10.4	14.6	15.9	10.7	7.1	9.70
Min	−4.1	−5.3	−1.2	−0.9	−6.5	−1.1	−4.2

Notes: The table compares annual inflation rates for different categories of goods or services over time. Source: Statistics Estonia.

Table A2: Imputed consumption shares

	HBS				Imputed			
	mean	sd	min	max	mean	sd	min	max
Food and non-alcoholic beverages	0.306	0.126	0.018	0.767	0.297	0.043	0.186	0.372
Alcohol and tobacco	0.040	0.052	0.000	0.327	0.036	0.008	0.020	0.064
Clothing and footwear	0.053	0.071	0.000	0.410	0.057	0.019	0.021	0.100
Housing	0.189	0.128	0.000	0.699	0.197	0.036	0.118	0.304
Household goods	0.052	0.070	0.000	0.565	0.048	0.010	0.026	0.088
Health	0.037	0.056	0.000	0.373	0.034	0.012	0.017	0.070
Transport	0.092	0.107	0.000	0.557	0.088	0.025	0.038	0.169
Communication	0.072	0.052	0.000	0.310	0.074	0.007	0.050	0.088
Recreation, entertainment	0.075	0.073	0.000	0.525	0.076	0.016	0.045	0.118
Education	0.004	0.023	0.000	0.332	0.006	0.004	0.001	0.021
Hotels, cafes, restaurants	0.021	0.041	0.000	0.316	0.024	0.014	0.004	0.086
Miscellaneous goods and services	0.059	0.061	0.000	0.453	0.062	0.015	0.033	0.106
Observations	11,470				2,567,205			

Notes: The table compares the consumption shares obtained from the Household Budget Survey with those imputed to the account level dataset.

Table A3: Actual vs. imputed experienced inflation

	Dependent variable: Actual experienced inflation	
	(1)	(2)
Imputed $IndINFL$	1.005***	0.996***
	(0.009)	(0.015)
Constant	0.000	0.001***
	(0.000)	(0.000)
Time FE	No	Yes
Observations	5740	5740
R^2	0.683	0.686

Notes: The table presents the results of regressing the actual experienced inflation observed for a subsample of households in the Household Budget Survey on the imputed value of experienced inflation obtained in an out-of-sample exercise. Standard errors are presented in parentheses. *, **, *** denote statistical significance at the 10%, 5%, and 1% level.

Table A4: Panel unit-root tests

	Harris-Tzavalis	p-value	Im-Pesaran-Shin	p-value
$IndINFL$	−1846.955	0.000	−641.626	0.000
$\Delta \log \text{INC}$	−3402.852	0.000	−997.219	0.000
$\Delta \log \text{FIN}$	−3237.137	0.000	−927.328	0.000
$\Delta \log \text{CONS}$	−3314.862	0.000	−976.527	0.000

The table presents the results of panel unit root tests: the Harris-Tzavalis test and Im-Pesaran-Shin test. The results make us reject the null hypothesis of stationarity in the panel.

Table A5: Different imputation models

	Dependent variable: $\Delta \log C_{it}$		
	Model 1	Model 2	Model 3
IndINFL_{it}	0.010^{***}	0.019^{***}	0.007^{***}
	(0.002)	(0.003)	(0.002)
$\Delta \log \text{INC}_{it}$	0.481^{***}	0.481^{***}	0.481^{***}
	(0.006)	(0.006)	(0.006)
$\Delta \log \text{FIN}_{i,t-1}$	0.108^{***}	0.108^{***}	0.108^{***}
	(0.001)	(0.001)	(0.001)
Individual FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
Observations	2,352,678	2,328,785	2,318,096

The table presents the results of re-estimating the effect of experienced inflation on consumption using different imputation models. Model 1 presents the estimates when consumption deciles are used instead of income deciles for the imputation. In Model 2 only data from the beginning of the sample (from 2005) are used to impute consumption weights. Model 3 presents the results when group-specific consumption weights are not imputed but taken directly from the HBS. All the variables except for inflation are in log differences. Standard errors clustered at group level are presented in parentheses while groups are the same as those used in the imputation model. *, **, *** denote statistical significance at the 10%, 5%, and 1% level.

Figure A1: Dynamics of consumption and income in micro data and aggregate data

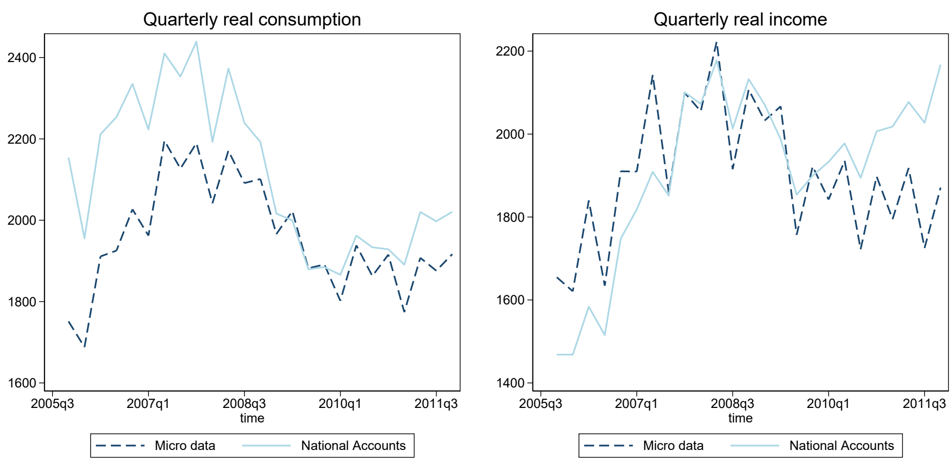

Notes: Aggregate consumption and income are converted to real values in 2010 prices. For micro data, mean real consumption and mean real income are provided, while for aggregate data total consumption of the household sector and compensation to employees are scaled by 15-75 population.

Figure A2: Individual inflation rates.

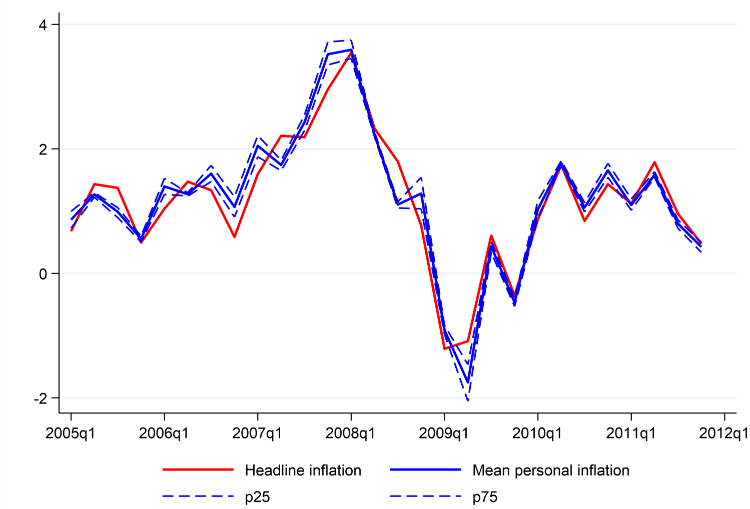

Notes: The figure plots the dynamics of the mean personal inflation rate (with corresponding 25th and 75th percentiles) against the headline inflation rate over time.

Figure A3: Cross-sectional heterogeneity.

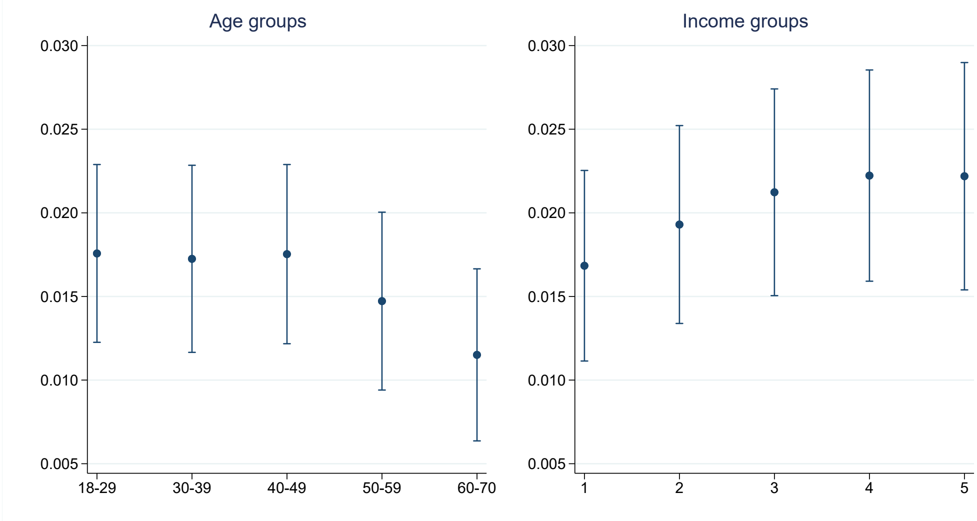

Notes: The figure plots the estimated coefficients for the consumption response across income and age groups.

Working Papers of Eesti Pank 2023

No 1
Karsten Staehr. Economic Growth, Current Account Dynamics and Growth Regimes in the Baltic States

No 2
Łukasz Wiktor Olejnik. Economic growth and military expenditure in the countries on NATO's Eastern flank in 1999–2021

No 3
Karsten Staehr, Oļegs Tkacevs, Katri Urke. Fiscal Performance under Inflation and Inflation Surprises: Evidence from Fiscal Reaction Functions for the Euro Area

No 4
Merike Kukk, Natalia Levenko and Nicolas Reigl. Measuring the effects of borrower-based policies on new housing loans in Estonia

No 5
Eva Branten. Income expectations, risk attitudes and household borrowing decisions

No 6
Dmitry Kulikov, Nicolas Reigl. The natural rate of unemployment in Estonia: empirical determinants and a new semi-structural model