



EUROPEAN CENTRAL BANK

EUROSYSTEM

Working Paper Series

Ioana A. Duca, Geoff Kenny, Andreas Reuter

Inflation expectations, consumption
and the lower bound:
micro evidence from a
large euro area survey

No 2196 / November 2018

Abstract

This paper exploits a very large multi-country survey of consumers to investigate empirically the relationship between inflation expectations and consumer spending. We document that for the Euro Area and almost all of its constituent countries this relationship is generally positive: a higher expected change in inflation is associated with an increase in the probability that a given consumer will make major purchases. Moreover, in line with the predictions of macroeconomic theory, the impact is stronger when the lower bound on nominal interest rates is binding. Also, using the estimated spending probabilities from our micro-level analysis, we indirectly estimate the impact of a gradual increase in inflation expectations on aggregate private consumption. We find the effects to be economically relevant, especially when the lower bound is binding.

JEL classification: D12, D84, E21, E31, E52

Keywords: Consumer inflation expectations, Consumption, Lower Bound, Micro data

Non-technical summary

Expectations, and inflation expectations in particular, lie at the centre of modern macroeconomic theory and models of how monetary policy operates. In such models, inflation expectations determine the real interest rate that is relevant to consumption and savings decisions. For example, an expected increase in inflation will lower real interest rates which in turn can boost consumption and reduce incentives to save. With an effective lower bound on nominal interest rates binding in several economies around the world, this relationship between inflation expectations and aggregate demand takes on an increasingly important role. In particular, in such circumstances, the above mechanism may become even more prominent because central banks are deprived of the use of their conventional policy instrument: the short-term interest rate. As a result, economic theory has also emphasized the important stabilization role that higher inflation expectations can play when the lower bound is binding.

Of course, whether actual consumers behave in line with the simple predictions of theory requires empirical evaluation. In this paper, we investigate these theoretical predictions for the Euro Area and 17 of its constituent countries. We benefit from a very rich micro dataset from the EU Consumer Survey which provides direct measures of consumers' expectations in relation to different economic and financial concepts, and which can be broken down according to a set of demographic characteristics. Our evidence is based on over 2,000,000 observations which were carefully collected to reflect Euro Area population over approximately a 14 years period, from May 2003 to December 2016. Most importantly for our analysis, the survey includes quantitative consumer inflation expectations and perceptions and gives information about the consumers' intentions to spend. We focus on an innovative measure of consumers' subjective expected change in inflation, which reflects the difference between consumers' expectations about future inflation and their subjective perceptions about current inflation. Econometrically, this difference helps allow for possible variation in unobserved consumer characteristics that may otherwise distort our results. We find this measure the most economically relevant because it captures the idea that when determining their spending intentions, consumers are likely to consider expected inflation relative to their currently perceived level of inflation. Indeed, the data reveal that the latter can vary widely across households and time and, hence, a failure to take this into account would likely bias our analysis.

Our findings broadly confirm that Euro Area consumers behave in line with the predictions of economic theory. When they expect higher inflation in the future, all other factors held constant, they adjust positively their intention to spend at the current moment. Our pooled results suggest that for a 1 percentage point (pp) expected increase in inflation the consumer's probability to spend increases by between 0.15 pp and 0.39 pp. This is confirmed by country-level results, where for almost all countries, we find a positive relationship between consumer inflation expectations and the propensity to spend, though there is some notable heterogeneity. Another key finding in our study is that the relationship between consumer inflation expectations and consumption becomes stronger when interest rates are bounded from below. This suggests that - once confronted with a lower bound on nominal rates - consumers become more aware of and sensitive to changes in real interest rates and the effect of inflation expectations on real spending behaviour is augmented as a result. We observe this change in the relationship due to the lower bound both in our pooled euro area results and also in the vast majority of our country specific results.

Overall, from a monetary policy perspective, our micro analysis of consumer behaviour provides evidence on the importance of intertemporal substitution linked to the real interest rate channel. Moreover, this channel appears to be stronger when the economy is constrained by the lower bound on nominal interest rates. This provides strong support for central bank concerns about a drop in inflation expectations during lower bound episodes because such developments have the potential to weaken aggregate demand by raising real rates. Conversely, our results point to the important stabilisation role that higher inflation expectations could play when the economy is constrained by the lower bound. In line with this, we link our micro level analysis to macroeconomic observables in a conditional policy simulation. The simulation shows that an increase in consumer inflation expectations by 2.0 pp can raise aggregate private consumption by up to 0.35% over a three year period when the lower bound is binding.

1 Introduction

Expectations, and inflation expectations in particular, lie at the centre of modern macroeconomic theory and models of how monetary policy operates. In such models, optimal intertemporal choices rely on inflation expectations to convert nominal values into the real values that are relevant to decision making and utility (see, for example, Armantier et al. (2015)). Under sticky nominal interest rates, an expected increase in inflation will lower real interest rates due to the well-known Fisher Effect. As a result an increase in expected inflation should boost current consumption or aggregate demand by lowering consumers' incentives to save for the future. With an effective lower bound on nominal interest rates binding in several economies around the world, the predicted relationship between inflation expectations and aggregate demand takes on an increasingly important role. Yet, as emphasised recently in Blanchard et al. (2010) and Blanchard et al. (2013), this relationship may be much less well understood in such an environment. Arguably, with nominal interest rates bounded from below, the above intertemporal mechanism may become even more prominent because central banks are deprived of the use of their conventional policy instrument: the short-term interest rate. In such a context, a rise in expected inflation will transmit one for one to lower ex ante real interest rates. In line with this, as discussed recently in Bachmann et al. (2015), a large theoretical literature has emphasized the important stabilization role of higher inflation expectations at the effective lower bound (see, for example, Krugman et al. (1998), Eggertsson and Woodford (2003), Jung et al. (2005), Eggertsson (2006)). The role of inflation expectations in macroeconomic stabilisation is also central to the debate on the advantages of price level versus inflation targeting monetary frameworks (see, for example, Vestin (2006) and, more recently, Bernanke (2017)). In particular, at the effective lower bound, price level targeting is seen as having more desirable stabilisation properties because it implies a larger rise in short-run inflation expectations which helps lower real interest rates and exerts a stronger stabilising effect on the economy.

Of course, whether actual consumers behave in line with the simple predictions of theory is an empirical question. Existing empirical evidence on the inflation expectations - consumption relationship is still scarce and has often also only been produced at the single-country level. Moreover these studies have brought forward quite conflicting evidence about how inflation

expectations may impact consumption or, indeed, whether there is any such relationship at all. Papers by Ichiue and Nishiguchi (2015) and D'Acunto et al. (2016) find that Japanese and German consumers, respectively, increase consumption in response to higher inflation expectations. However, for US consumers, Bachmann et al. (2015) and Burke and Ozdagli (2013) find that there is no significant positive impact of inflation expectations on durable goods consumption.¹ Indeed, at the lower bound, the findings of Bachmann et al. (2015) suggest that the impact of higher inflation expectations on consumption may be negative. In this paper we exploit a very large multi-country micro dataset for the Euro Area (EA) economies and provide new empirical evidence on this fundamental relationship in macroeconomics.

A necessary requirement for the investigation of the consumer inflation expectations - consumption relationship is to have the right data. Aggregated time series data do not do the trick as both consumption and inflation are determined jointly. Also, through aggregation a lot of information is lost and the heterogeneity of consumer behaviour cannot be taken into account or controlled for. Figure 1 plots consumer inflation expectations, an aggregate indicator of consumers' readiness to spend along with actual real total consumption growth in the Euro Area. With such aggregate time series data, investigating this relationship in the more recent effective lower bound (ELB) period is not practical since there are only few observations available to help estimate an aggregate model's multiple parameters, including parameters linked to other variables for which the econometrician would ideally wish to control (e.g. variables capturing a consumer's real income, wealth or financial constraints). In contrast, because of the much greater sample size and cross-sectional variation, microeconomic data is particularly well-suited to the task of identifying this relationship and how it may differ across individual consumers depending on their economic and social circumstances. This study benefits from a very rich multi-country dataset with individual level consumer data collected on a monthly basis between May 2003 and December 2016 with approximately 26,440 euro area consumers included in each monthly survey round.

[Insert Figure 1 here.]

To help identify the inflation expectations - consumption relationship we estimate an

¹ Armantier et al. (2015) find evidence that most consumers make their investment choices in accordance with the predictions of economic theory and intertemporal choice.

ordered logit model and exploit several novel features in our data set to reduce concerns about endogeneity as a possible driver of the empirical results. Firstly, we focus on the difference between an individual consumer's expectation about future inflation and their perceptions about current inflation as the key driving variable impacting on the consumer's readiness to spend. In economic terms, by focusing on this difference we capture the intuition that when determining their readiness to spend a consumer is more likely to consider their expected future inflation rate not in absolute terms but relative to their perceptions about current inflation. In econometric terms, this transformation is similar to a fixed effects specification which therefore makes less probable the correlation of the covariates of interest with unobserved personality traits such as, for example, pessimism or optimism at the level of individual consumers. Second, we introduce a number of other relevant individual and aggregate controls that should alleviate concerns about omitted variables as a source of endogeneity. For example, we control for variation in individual income and wealth by exploiting survey responses about the expected financial situation, labour market conditions and the outlook for the economy in general. We also control for several important demographic dimensions including gender, educational attainment, employment status and income, to validate that our results are not simply driven by heterogeneity in these consumer characteristics. Lastly, we introduce several interaction terms to explore possible heterogeneity across different consumers in their spending response to a change in inflation expectations.

We find that EA consumers behave in line with conventional macro-economic theories; that is, when they anticipate an increase in future inflation, consumers increase their readiness to spend, holding all other factors constant. There are four main empirical results that support this conclusion. First, pooled EA analysis shows that for a 1.0 percentage point (pp) increase in inflation expectations the likelihood of spending increases by between 0.15 pp and 0.39 pp, depending on the model specification. Second, we find that the relationship between consumers' inflation expectations and the likelihood to spend is stronger when the ELB is binding. This result is robust across all model specifications examined. Third, individual country results confirm the pooled euro area results. With only one exception, all countries in the sample exhibit a positive relationship between consumer inflation expectations and the likelihood of spending today.² Fourth, we use a simple VAR to link our micro analysis of consumer behaviour

²Country heterogeneity does nonetheless show up in the results as the average marginal effect of a change in

to actual aggregate consumption. We then exploit the VAR in a conditional policy simulation which shows how an increase in expected inflation by 2 pp, e.g. from a level of 0% which is commonly seen to be "too low" to 2% - a level which is more in line with the ECB's quantitative definition of price stability - can boost real private consumption growth by 0.35% over a three year horizon when the effective lower bound on nominal interest rates is binding.

Our paper relates to the broader and rapidly growing literature that looks into how consumer inflation expectations are formed and seeks to understand the heterogeneity behind reported inflation expectations. In addition to the papers by Bachmann et al. (2015), Burke and Ozdagli (2013), D'Acunto et al. (2016), Ichiue and Nishiguchi (2015) cited above, our work is also closely linked to a recent study by Crump et al. (2015) who estimate the elasticity of intertemporal substitution using subjective inflation expectations for the US. Several related studies have emphasised the subjectivity of expectations and, in particular, the importance of socio economic and demographic factors in shaping inflation expectations (e.g. Jonung (1981); Bryan and Venkatu (2001); Lombardelli and Saleheen (2003); Souleles (2004); Christensen et al. (2006); Anderson (2008)). Also, Coibion and Gorodnichenko (2015) and Coibion et al. (2018) illustrate how the missing disinflation puzzle associated with the Great Recession period in the US can be explained once one controls for household inflation expectations in a Philips curve setting. Other relevant strands in this literature include Naohito and Yuko (2015) and Armantier et al. (2014) who use survey experiments to investigate the effects on expectations formation of providing information to consumers and Carroll (2003) who fits a model of household inflation expectations in the spirit of the "sticky information" theory of Mankiw and Reis (2001). Malmendier and Nagel (2016) also document the importance of subjective inflation expectations for economic outcomes, finding that households with higher inflation expectations are less likely to invest in long-term bonds and more likely to borrow through fixed-rate mortgages compared to their counterparts with lower inflation expectations. Also Ehrmann et al. (2017) link the observed bias in consumer inflation expectations to household financial difficulties and pessimistic spending attitudes. Finally, Binder (2017) exploits individual consumer data on inflation expectations to propose an inflation uncertainty measure for the US economy.³

inflation expectations, although generally positive, differs in magnitude across countries.

³The measure is built around the idea that round numbers are used by respondents who have high imprecision or uncertainty about inflation. Using this index, Binder (2017) finds that more uncertain consumers are more

The remainder of the paper is structured as follows: In section 2 we provide important details about the micro dataset that we use and shortly illustrate the methodology we adopt to exploit it. Section 3 details all the results of the probabilistic analysis in which we determine a consumer inflation expectation - propensity to spend relationship, including both euro area wide and country-specific results. Section 4 presents the VAR analysis and conditional policy simulation linking the micro-level analysis with macroeconomic data on actual private consumption at the euro area level. Section 5 concludes and discusses the economic relevance of our findings.

2 Data and methodology

Through the design and focus of its survey questionnaire, the EU Consumer Survey provides the ideal micro-information set to study the relationship between inflation expectations and the readiness to spend of the EA consumer. Although in this paper we focus only on EA countries, the survey is carried out at a monthly frequency and covers all European Union economies, as well as four of the five candidate EU countries.⁴ In this section we discuss key features of the dataset that are relevant to the analysis. A Data Appendix provides further detail on data sources and methodology.⁵ Each month we benefit from a sample of 26,440 EA consumers who provide information on their perceptions of current inflation and their expectations about future inflation.⁶ The sample is designed to be representative of the population in each country. Its size and composition vary across countries reflecting socio-economic heterogeneity and differences in country population size. Each month a new sample of consumers is interviewed, implying that we work with a repeated cross-section and not a panel of consumers that can be tracked through time. Although the repeated cross sectional nature of the data has the disadvantage that it impedes the estimation of consumer-specific fixed effects, our analysis focuses on a novel measure of the expected change in inflation which helps overcome this problem.

reluctant to spend on durables, cars and homes, and their spending attitudes are less sensitive to interest rates.

⁴However, our sample does not include Ireland due to data availability.

⁵Arioli et al. (2016) describe the dataset in more detail.

⁶The actual sample size used for estimation is 11,275. This lower sample reflects that we also draw on other questions and the response rate can vary from question to question.

The vast majority of the surveys in the euro area countries are conducted by Computer-Assisted Telephone Interviews (CATI).⁷ Most of the questions in the survey are qualitative and refer to the consumer's own financial situation, their views about the general economic situation, their savings behaviour and intentions with regard to major consumer purchases. Since 2003, the Consumer Survey includes specific quantitative questions about consumers' perceptions of current inflation and their expectations for inflation over the next 12 months.⁸ In addition, the replies can be broken down along several important demographic dimensions (e.g. gender, level of educational attainment, current employment status, income level etc.). The sample period employed in our econometric analysis covers the period between May 2003 and December 2016.

2.1 Survey questions

The main survey questions that we use in our empirical analysis are:

Q51: By how many per cent do you think that consumer prices have gone up/down over the past 12 months? Consumer prices have increased by ____, __% / decreased by ____, __%.

Q61: By how many per cent do you expect consumer prices to go up/down in the next 12 months? Consumer prices will increase by ____, __% / decrease by ____, __%.

Q8: In the view of the general economic situation, do you think that now it is the right moment for people to make major purchases such as furniture, electrical/electronic devices, etc.? Survey respondents can answer: i) yes, it is the right moment now; ii) it is neither the right moment nor the wrong moment; iii) no, it is not the right moment now; iv) don't know.

⁷Only in three countries (Germany, Latvia, Slovakia), interviews take place in a Face to Face (F2F) setting. Two countries apply mixed modes which combine CATI (Austria) or CATI and F2F (Lithuania) with web interviews. The households to be interviewed are determined by random sampling or quota sampling from a frame which, in most cases, is either the country's telephone directory or its population register.

⁸A complete version of the survey questionnaire can be found in the User Guide of The Joint Harmonised EU Programme of Business and Consumer Surveys (pages 36 - 40): http://ec.europa.eu/economy_finance/db_indicators/surveys/method_guides/index_en.htm.

Questions Q51 and Q61 are quantitative and the answers are expressed as percentage changes over a 12 month period and capture consumer inflation perceptions and consumer inflation expectations, respectively. A notable feature of the survey design is the very neutral phrasing of the questions on inflation. Survey respondents are not supplied with information about official price indices nor are they given a range of possible inflation rates from which they could choose. Although this has the advantage that it avoids introducing any framing bias into the responses, it means that the unfiltered responses may be more susceptible to measurement error and extreme and implausible replies. For this reason, in our econometric analysis below, we focus on a filtered dataset which removes very extreme outliers although we also report the results from a completely unfiltered sample.⁹ In contrast to the two questions on inflation expectations, Q8 on consumer spending is qualitative in nature and shows whether, given the prevailing economic context, the consumer considers it to be a good time to spend on "major purchases". Although the question phrasing appears to prompt consumers to respond about planned durable consumption, it is certainly plausible that the replies capture spending plans for non durables. Indeed, as we document later, the responses appear informative for total private consumption in the euro area and not just for durable consumption. Throughout the paper, we refer to the replies to this question as a measure of the so called "readiness to spend" of the consumer. In our empirical analysis, we quantify the answers to this question in the following way: a value of 1 indicates that it is not the right moment to spend, a value of 2 indicates it is neither the right moment nor the wrong moment, while 3 indicates that it is the right moment to spend.

According to the survey results, EA consumers hold very heterogeneous opinions about inflation expectations and perceptions depending on their gender, age, education, income or employment status. Inflation expectations are higher for females, the unemployed, consumers aged below 50, with low income and holding only primary or secondary education (see Table 1). Consumer perceptions about the current rate of inflation follow the same pattern. They are, however, persistently higher than expectations. Also, both consumers' mean expectations and

⁹Manski (2004) reviews recent approaches to the measurement of expectations and highlights the increasing use of histograms where survey respondents indicate probabilities that a particular economic variable will fall within a given interval. One advantage of the histogram approach, used also in the New York Federal Reserve Bank's Survey of Consumer Expectations (SCE) is that it can help centre the responses within a range of plausible values. In contrast, the relatively open questions in the EU Consumer Survey do not offer any such range of response intervals.

mean perceptions of price changes are persistently higher than actual inflation developments, measured by the official Harmonised Index of Consumer Prices (HICP). This positive difference is a common finding across consumer surveys on inflation expectations - see, for example, Kliesen (2015). It might be explained in several ways: the survey questions are open ended with a generic reference to consumer prices and provide no range or other quantitative guidance for the respondent in determining the inflation rate. Also, unusual replies are not probed and respondents are not asked about an objective official price index. Respondents assumedly provide an answer that is based on their own subjective inflation experience which could weight price changes differently compared to an official index or even possibly take account of prices that are excluded from such an index (e.g. house prices). Nevertheless, the size of the difference has narrowed considerably over time.¹⁰ Disregarding this persistent positive difference, both mean expectations and mean perceptions co-move very strongly with actual inflation (see Figure 2). Such a strong co-movement provides solid grounds to use this dataset for investigating the consumer inflation expectations - spending relationship.

[Insert Table 1 here.]

[Insert Figure 2 here.]

2.2 Subjective inflation expectations

When investigating the relationship between consumer inflation expectations and readiness to spend, we focus on a measure of the expected *change* in subjective inflation defined as the difference between an individual's inflation expectations and his/her perceptions:

$$\Delta\pi_{it}^e = \pi_{it}^e - \pi_{it}^p \quad (1)$$

where $\Delta\pi_{it}^e$ is the expected change in inflation of consumer i at time t , π_{it}^e is the subjective inflation expectation over the next 12 months and π_{it}^p the subjective perception of inflation over the past 12 months.

¹⁰The substantially higher inflation perceptions at the beginning of the sample might have been related to the introduction of the euro notes and coins. This was widely seen by the public as being associated with higher prices - as discussed in Ehrmann (2011). Also, the subsequent decline may link to learning on the part of consumers who may have become more informed about the objective and policies of the European Central Bank.

This measure is innovative and contrasts with previous empirical studies which have focused on the level of expected inflation *per se*.¹¹ Our proposed measure has several advantages. First, as it represents a first difference which can be computed for all respondents, it is similar to a fixed effects specification, which may be particularly important given the repeated cross-sectional nature of our data. For example, when changing their individual spending intentions consumers are likely to take into account not just the level of expected inflation but the level of expected inflation relative to their own current subjective perceptions about recent inflation. This transformation also helps to address concerns about possible endogeneity linked to unobserved personality traits that may correlate both with inflation expectations and readiness to spend and thereby give rise to inconsistent parameter estimates. For example, as was shown earlier, both expectations and perceptions about inflation exhibit a positive "bias" compared with official statistics. Hence, by focusing on the difference between the two, we can eliminate possible distortions to our analysis associated with excessively pessimistic or optimistic beliefs about inflation.

A first look at the data (see Figure 3) strongly suggests that conditioning on individual perceptions about current inflation in the way that we propose may be key to better understanding the relationship between expected inflation and consumption behaviour. The Figure shows two scatter plots linking inflation expectations and the consumer's readiness to spend. The scatters are produced by computing a monthly average value for both indicators in each country over the sample period between May 2003 and December 2016. In the first scatter, panel (a), the level of expected inflation is reported on the x-axis while in the second scatter, panel (b), we exploit the expected change in subjective inflation computed according to equation (1). Clearly, when perceptions are controlled for as in equation (1) above, we observe a much stronger positive relationship with the readiness to spend. The relationship is much weaker as demonstrated in panel (a) which does not control for subjective perceptions about past inflation and focuses only on variation in the level of consumers' inflation expectations alone. Such a graphical analysis - while it supports our approach - is, however, only suggestive.

¹¹D'Acunto et al. (2016) do not focus on the expected change in inflation but rather add a qualitative measure of inflation perceptions as a separate regressor. These authors find that the marginal effects of expected inflation would be virtually identical over several specifications both including and excluding inflation perceptions as a separate regressor. Due to data limitations, Bachmann et al. (2015) control for the current official inflation rate which by definition is common across all consumers and therefore does not control for any differences and heterogeneity in perceptions about current inflation.

A more robust economic analysis requires a larger empirical model which can control for the wide array of additional factors impacting on the survey responses. In the next section we detail the empirical model we use for this analysis.

[Insert Figure 3 here.]

2.3 Methodology

Our data dictates our modelling strategy: the discrete nature of the spending attitudes that are retrieved from the survey combined with the fact that we observe a repeated cross-section recommend the use of a discrete choice model. We therefore employ an ordered logit specification to model the relationship between the expected change in inflation and the individual consumer's readiness to spend. This exploits the natural ordering in our dependent variable, the consumer readiness to spend. As it represents the answer to the question whether it is a good moment to spend, it can be ordered into being more or less ready to spend with those consumers responding that it is not the right moment being the least ready to spend. Choosing one response category over another depends on a latent variable (i.e. some continuous measure of the readiness to spend) which - though not observed- can be modeled as:

$$y_{it}^* = X_k \beta + \varepsilon_{it} \quad (2)$$

where i indicates consumer i and t is time, y_{it}^* is the latent variables, X_k is a k -dimensional vector of individual specific, aggregate and cross-sectional controls that will be described in more detail below, β is a k -dimensional vector of coefficients and ε_{it} is the error term.

Each response category can then be defined in relation to the latent variable defined in equation 2:

$$y_{it} = \begin{cases} 1 & \text{if } y_{it}^* < \alpha_1 \\ 2 & \text{if } \alpha_1 \leq y_{it}^* < \alpha_2 \\ 3 & \text{if } y_{it}^* \geq \alpha_2 \end{cases} \quad (3)$$

Each alternative response has a probability Pr attached to it:

$$\begin{aligned} Pr(y_{it} = 1) &= Pr(y_{it}^* < \alpha_1) = Pr(X_k\beta + \varepsilon_{it} < \alpha_1) = \\ &= Pr(\varepsilon_{it} < \alpha_1 - X_k\beta) = F(\alpha_1 - X_k\beta) \end{aligned} \quad (4)$$

$$\begin{aligned} Pr(y_{it} = 3) &= Pr(y_{it}^* \geq \alpha_2) = Pr(X_k\beta + \varepsilon_{it} \geq \alpha_2) = \\ &= Pr(\varepsilon_{it} \geq \alpha_2 - X_k\beta) = F(X_k\beta - \alpha_2) \end{aligned} \quad (5)$$

$$\begin{aligned} Pr(y_{it} = 2) &= 1 - Pr(y_{it} = 1) - Pr(y_{it} = 3) = \\ &= 1 - F(\alpha_1 - X_k\beta) - F(X_k\beta - \alpha_2) = \\ &= F(\alpha_2 - X_k\beta) - F(\alpha_1 - X_k\beta) \end{aligned} \quad (6)$$

where F is a function that satisfies $F(-\infty) = 0$, $F(+\infty) = 1$ and $\frac{dF(x)}{dx} > 0$. We model F through a logit function which ensures that the estimates take values between 0 and 1, i.e. the domain of admissible values for a probability. Alternatively, we could have used a probit function. However, in practice, the probit and logit models generally yield very similar results (see, e.g., Davidson and MacKinnon (2004)).

We use maximum likelihood to estimate the parameters of these probability functions, including the thresholds for the latent variables which determine the selected response categories. The parameter vector β is in itself of limited interest; instead we are interested in how the probability of each alternative changes with a change in our controls. Therefore we will focus on the distribution of marginal effects measuring the impact of a change in a given control on our estimated spending probability:

$$\frac{\partial Pr(y_{it} = 3)}{\partial X_k} = f(\alpha_2 - X_k\beta) \beta_k \quad (7)$$

where β_k is the coefficient on regressor X_k and $f = F'$, in our case the probability density function of the logistic distribution.

2.4 Model specification and controls

To model the unobserved spending intentions of consumers we adopt a similar specification to that used in Bachmann et al. (2015). However, in contrast to this study, we replace expected inflation with our proposed forcing variable from equation (1) which controls for consumer-specific perceptions of past inflation. The specification includes a dummy variable (*ELB*) for the period associated with the lower bound on nominal interest rates and allows for both a level shift in spending attitudes associated with the *ELB* period and, via an interaction term, a possible change in the sensitivity of spending intentions to inflation expectations when the *ELB* is binding. As outlined further below, we include additional controls for a series of covariates that are likely to also drive spending intentions independently of inflation expectations as well as additional interaction terms that allow the effect of inflation expectations to vary depending on certain consumers characteristics. Equation (8) below summarizes how we model the latent variable:

$$y_{it}^* = \beta_0 + \beta_1 ELB + \beta_2 \Delta\pi_{it}^e + \beta_3 \Delta\pi_{it}^e ELB + X\gamma + \varepsilon_{it} \quad (8)$$

where $\Delta\pi_{it}^e$ is the expected change in subjective inflation defined according to equation (1), *ELB* is a dummy variable taking a value of 1 from June 2014 to December 2016, X is a vector of additional controls, ε_{it} is the error term and $\beta_1, \beta_2, \beta_3, \gamma$ represent parameters to be estimated. There is some uncertainty concerning the precise date on which the *ELB* became binding in the euro area. On 5 June 2014, the ECB Governing Council reduced the rate on its main refinancing operations by 10 basis points to 0.15% while it also reduced the interest rate on the ECB deposit facility by 10 basis points to -0.10%, with effect from 11 June 2014. Three months later on 4 September 2014, these two policy rates were reduced further to 0.05% and -0.20% respectively. We choose June 2014 as the starting point for the *ELB* period reflecting the fact that, although not all policy rates had been reduced to zero, the issue was very intensively discussed at that time. For example, at the ECB press conference on 5 June 2014, the President

of the ECB was asked if he could "exclude any further interest rate cuts no matter what". He replied ". . . that for all the practical purposes, we have reached the lower bound. However, this doesn't exclude some little technical adjustments and which could lead to some lower interest rates in one or the other or both parts of the corridor. But from all practical purposes, I would consider having reached the lower bound today." ¹²

As mentioned above, when estimating equation (8) we undertake several robustness checks that gradually control for variables that could simultaneously affect both inflation expectations and readiness to spend and we allow several additional interactions among these controls. First, we control for a rich set of consumer characteristics: age, gender, level of educational attainment, employment status and income; which we wrap up together under the heading "Demographics". We have already seen in section 2.1 that there is significant heterogeneity in inflation expectations and perceptions depending on these key consumer characteristics. Souleles (2004) shows that variations in inflation expectations can be explained by such demographic characteristics. The same characteristics may determine different purchasing propensities and we would want to ensure that any impact of inflation expectations on spending - if it is to be interpreted as structural - is not simply driven by these differences.

Second, we consider equally important to control for individual expectations of the general economic and labour market situation, e.g. in a booming economic environment, consumers may increase spending due to a change in expectations about future inflation or, more simply, because of the favourable economic context. Likewise, controlling for the individual consumer's current or expected financial situation may be important as such personal circumstances and perceptions may be an even more important factor compared to the general economic context. The introduction of these controls helps to avoid that any effect of inflation expectations in our regressions is driven by reverse causality linked to a Phillips curve type relationship whereby higher growth or lower unemployment which boost consumption also pushes up consumer's inflation expectations. For example, if a consumer expects that his/her own financial situation may deteriorate, his/her consumption plans will probably decrease even though he/she expects that the general economic situation will get a lot better and inflation will increase. For

¹²See the Introductory statement to the press conference (with Q and A) available here: <https://www.ecb.europa.eu/press/pressconf/2014/html/is140605.en.html>

presentational clarity, we classify these controls into two groups: "Controls I" and "Controls II". "Controls I" includes a consumer's expectations about the general economic situation and their own financial situation. "Controls II" includes a consumer's expectations about the labour market situation and their own financial status, i.e. whether they are a debtor or non-debtor. See also the Data Appendix for further detail in relation to these additional questions that are used as controls.

Third, we also consider further specifications which allow for potential heterogeneity in the transmission of inflation expectations by introducing additional pairwise interactions between the expected change in inflation and the demographic and consumer-specific controls discussed above. Figure 4 shows in a series of additional scatter plots how the relationship between expected inflation and readiness to spend may differ depending on such consumer characteristics. In general, the scatter plots suggest that the relationship is stronger for consumers with a higher level of educational attainment, with a higher income, who are currently employed or who expect an improvement in their personal financial situation. Fourth, we introduce annual time dummies to control for any unobserved aggregate macroeconomic developments that may drive the spending decisions of all consumers. Fifth, to account for the heterogeneity of the economies that constitute the EA, we include country dummies.¹³ Finally, we also include country specific and EA macro aggregates, by drawing on information sources outside the survey. In particular, we control for lending rates to households at the country level and oil prices.

[Insert Figure 4 here.]

3 Empirical results

In this section we report all our empirical findings related to the estimation of the ordered logit model given by equation (8). In subsection 3.1 we show the EA results for the inflation expectations and the propensity to consume relationship, while in subsection 3.2 we analyse how the effect of inflation expectations on consumption may differ across different types of

¹³Although the correlation is mostly positive, a review of the simple scatter plot of the expected change in inflation and consumer readiness to spend at a country level points toward considerable heterogeneity among the EA economies. We investigate this heterogeneity further in section 3.4 below.

consumers. We discuss what our model implies for the role of other factors in the consumption decision in subsection 3.3. Subsection 3.4 shows country-specific results.

3.1 Main results

We find that there is a statistically significant and positive effect of our measure of the expected change in subjective inflation on the probability of being ready to spend. This effect is robust across a variety of specifications in which we allow for different controls and sources of heterogeneity. Table 2 reports average marginal effects of a one unit increase in our measure of the expected change in inflation, across all specifications that we estimate. The average marginal effects are based on the ordered logit estimation and correspond to the alternative that now is the right time to spend, i.e. they show the impact on the probability of being ready to spend.¹⁴ Outside the ELB average marginal effects range between 0.15 to 0.29 pp and thus imply an increase in the probability of being ready to spend in response to a 1.0 pp rise in the expected change in inflation. When the ELB is binding, the effect of inflation expectations is generally larger across all specifications that we consider with the marginal effects ranging between 0.19 to 0.39 pp (see Table 2). The last column of Table 2 reports the p-value from the likelihood ratio test that the average marginal effects away from the ELB are identical to the average marginal effects when the ELB is binding. According to the results, such equality is rejected for all model specifications at the 10% level of significance, while for the two specifications using a winsorized data sample this equality is even more strongly rejected at the 1% level of significance. Also in line with this result, the coefficient interacting the ELB dummy with the expected change in inflation is positive and statistically significant at the 1% level. This is very much in line with what one would expect: in an environment where the ELB is binding, nominal interest rates are bounded from below. As a result, the real interest rate may respond one-for-one with an expected change in inflation. This may make the latter even more relevant for spending decision than would otherwise be the case. Our results therefore offer support to the potential stabilisation role of inflation expectations at the ELB, e.g. because a higher expected change in inflation imparts a stronger stimulus to consumption when policy is constrained by the ELB.¹⁵

¹⁴With the ordered logit model one can separately estimate the probabilities of each alternative, i.e. being a good moment to spend, not a good moment to spend, neither good nor bad moment to spend, and the corresponding marginal effects for all controls.

¹⁵We have run a separate set of ordered logit regressions where, instead of the expected change in inflation as our main regressor, we use inflation expectations on their own and then add inflation perceptions as a separate

Overall, therefore, across all specifications, EA consumers appear to behave in line with the predictions of mainstream economic theory: a higher expected rate of inflation is associated with an increase in individual spending intentions. Interestingly, the largest and most economically plausible marginal effects are obtained when we remove outliers by winsorizing the data on inflation expectations at the 5th and 95th percentiles. As mentioned earlier, given the very open nature of the survey questions on inflation expectations, such a procedure may help control for measurement error. Once extreme outliers are removed, the estimated effect is also remarkably stable across the very parsimonious model with only few controls (second row of Table 2) and a model where all controls and interactions are included (last row of Table 2).¹⁶

In terms of the sign of the impact, our results are qualitatively consistent with results reported by D'Acunto et al. (2016). These authors demonstrate that for German consumers¹⁷ an increase in inflation expectations leads to a 6 to 9 pp increase in the probability that consumers are ready to spend. However, this result refers to qualitative data about inflation expectations and perceptions and, hence, the magnitude of the effect cannot be compared with the marginal effect of inflation expectations in our model. Our findings are also in accordance with Ichiue and Nishiguchi (2015) who use micro data for Japan and find a significant positive impact of inflation expectations on consumption. The Japanese case is of particular interest because, in contrast with many other advanced economies, it has experienced a prolonged period of near zero interest rates. As a result, the authors argue that, even if they expect higher inflation, consumers are less likely to expect a similar simultaneous movement in nominal rates.

Our results contrast with the main findings of Bachmann et al. (2015) which relate to the US economy. Using the Michigan Survey, these authors find that the effect of higher inflation expectations for US consumers is very close to zero and statistically not significant during

control. The results of this estimation in the case in which we use the full set of controls, i.e. corresponding to the last specification reported in Table 2, indicate a positive and statistically significant effect of inflation expectations on the readiness to spend outside the ELB and a negative and non statistically significant effect at the ELB. Moreover, these marginal effects are much smaller than the ones reported in the case of the expected change in inflation measure.

¹⁶This result does not depend on the approach used to adjust for outliers. When we use an alternative approach which eliminates observations that are more than 3 standard deviations away from the mean, we obtain very similar results to the winsorized dataset. In particular, the results of this robustness check indicate that the marginal effects remain positive and are larger at the ELB.

¹⁷They use data provided by the market research firm GfK, which conduct the consumer survey for Germany.

normal times. Moreover, in periods when the ELB is binding, the relationship is generally shown to be negative, i.e. higher expected inflation which reduces real interest rates is associated with a drop in consumption.¹⁸ Also with respect to US consumers, Burke and Ozdagli (2013) find that consumers do not increase their spending on large home appliances and electronics in response to an increase in inflation expectations. However, Burke and Ozdagli (2013) do find that consumers are more likely to purchase a car and increase spending on non-durable goods.¹⁹

What might explain these contrasting results for US and Euro Area consumers? A first possible explanation relates to our focus on the subjective expected change in inflation, which, as argued previously, may better help control for unobserved sources of consumer heterogeneity. A second possible explanation is that the different results may relate to the considerably larger sample size used in our study. For example, the study of Bachmann et al. (2015) is based on a total of 67,855 observations covering a time span of 24 years, although the amount of cross-sectional units that can be obtained in any given month is relatively modest.²⁰ In working with the EU Consumer Survey, we benefit from a total sample size of over 2,000,000, albeit over a shorter period of 13¹/₂ years. This amounts to an average of 11,275 matched cross-sectional observations per month.

[Insert Table 2 here.]

3.2 Cross-sectional heterogeneity

An enormous virtue of our consumer dataset is that it permits an analysis of how the effect of inflation expectations on consumption may differ across different types of consumers. In this section we examine how the distribution of individual marginal effects differs depending on particular consumer characteristics. This gives us a comprehensive picture of the cross-sectional

¹⁸These authors do find an interesting exception to these general results. In particular they obtain a positive effect of inflation expectations at the lower bound for highly accurate inflation forecasters. Given the small share of total households with such forecasts, this results is less important in macroeconomic terms.

¹⁹These conclusions are based on data from the New York Fed/RAND-American Life Panel household expectations survey. This dataset contains detailed information on actual consumer spending, including the ability to distinguish durable and non-durable consumption.

²⁰The modest number of cross-sectional units available is a consequence of the fact that only a subsample of first interviews (the Michigan Survey has a rotating panel structure, i.e. about 40% of the respondents are interviewed also in the next round) is used and observations that are larger than 20 percent in absolute value are excluded. As a consequence, the authors are left on average with a sample of 195 consumers out of the 500 who are interviewed.

heterogeneity of our results and builds on the graphical evidence reported previously in Figure 4. Figure 5 plots the probability density functions (PDFs) of the individual marginal effects distinguishing between several different groups of consumers. The different groups of consumers are formed based on characteristics such as gender, age, level of educational attainment, income quartile, employment status and expected financial situation. In estimating these PDFs, and in further analysis reported below, we use the most general model specification which includes several groups of controls: demographics, expectations and financial status, time and country dummies, macro economic aggregates and adjusting for outliers, i.e. from the last row of Table 2. To construct the distribution of marginal effects for each consumer, we use this specification to fit the probability of being ready to spend at the ELB. Then we augment the individual expected change in inflation by 1 pp and re-compute a second set of individual probabilities from the model. The individual marginal effect of a 1 pp increase in the expected change in inflation is then estimated as the difference between the second and first set of individual probabilities. We then sort each of these marginal effects according to demographic and other economic characteristics and fit a Gaussian kernel density to the distribution of marginal effects for each of the groups. For ease of interpretation the figure shows on the X-axis the individual marginal effects represented as probabilities (i.e. multiplied by 100).

[Insert Figure 5 here.]

According to the results, age and gender do not seem to significantly influence the distribution of marginal effects across individuals. However differentiating according to other characteristics such as educational attainment, employment status, income level and financial situation has a more noticeable impact on the distribution of marginal effects across individuals. For example, the distribution is centred much more toward a smaller marginal impact for consumers who have attained only a primary education, who are located in the lowest quartile of the income distribution or who are currently unemployed. As discussed previously in relation to Figure 4, each of these consumer characteristics may correlate highly with the existence of constraints on consumers which impede their access to financial markets and, hence, their ability to engage in intertemporal substitution. In line with this, we observe a very sharp concentration of the distribution of marginal effects at very low or near zero levels for consumers indicating that they expect their financial situation to get a lot worse. Indeed, for these latter four groups, the

results in Figure 5 suggest that a change in expected inflation tends to have a much lower impact on their readiness to consume and the distribution of the marginal effects for these categories attaches a much higher probability to lower, or close to zero, marginal effects. Each of these distributions are shifted to the left, show signs of bi-modality with one of the peaks near zero or they become more skewed with the peak of the PDFs close to zero. In contrast, the distribution of marginal effects for higher educated, employed or higher income consumers and consumers who are more optimistic about their current financial situation are much more concentrated at positive values. Overall, therefore, these results confirm the graphical impression obtained previously from Figure 4 and the existence of substantial heterogeneity amongst consumers in the response to a change in inflation expectations.

3.3 Detailed model results

To get a broader sense of how successful we have been in modelling consumer spending intentions, Table 3 shows in detail the average marginal effects associated with all controls that we have used in the estimation. As can be seen from the table, many of the additional controls included in the model exhibit a significant and economically plausible impact on the consumer's readiness to spend. We discuss each of these in turn below.

Regarding the two common macro controls, both higher oil prices and higher nominal lending rates tend to reduce consumers' readiness to spend with the effects being broadly similar across ELB and non-ELB regimes. According to the model estimates, the effects of a change in oil prices are quantitatively much smaller than the effects of a change in lending rates, with the latter variable having a slightly larger overall effect compared with a change in expected inflation (also reported in the table). Amongst the demographic controls included in our model, there is evidence that being older tends to decrease readiness to spend with the 30-49 age cohort being significantly less ready to spend compared with the youngest 16-29 age cohort. However, for the two older 50+ cohorts, there is no significant effect of age on readiness to spend. Also, the model identifies a gender-specific negative effect, with females on average being less ready to spend compared with their male counterparts. In contrast our model estimates show how having a higher level of education, higher income, being employed as opposed to unemployed all increase the consumers' readiness to spend. Importantly, the impact of these discrete demographic controls on the consumer probability to spend cannot be

compared directly in quantitative terms with the marginal effects associated with a change in the continuous variables such as oil prices, lending rates or the expected change in inflation.

Table 3 also reports the marginal effects associated with the other respondent-specific controls measuring expectations about a respondent's own financial situation, their debt burden or more general macroeconomic and labour market developments. According to the results, readiness to spend increases with expectations of an improved own financial situation, an expected improvement in the general economic situation, but it declines as the consumer takes on debt or if he or she expects a deterioration in labour market conditions (an expected increase in unemployment rates). Such results suggest that the model captures well the effects of the key economic factors that can influence directly income and wealth expectations and therefore drive spending behaviour independently of the intertemporal substitution motive. Interestingly, the average marginal effects of expectations of a better *individual* financial situation are generally quite close to the average marginal effects related to expectations for a better *general* economic situation. For example, according to the estimated effects, expecting that the general economic situation gets a lot better relative to getting a lot worse increases the propensity to spend by approximately 12.7 pp, while expecting that one's own financial situation gets a lot better relative to getting a lot worse increases the propensity to spend by a comparable 10.7pp. Finally, the model results also highlight the important negative impact that deteriorating labour market expectations and increasing individual debt burdens can have on readiness to spend. For example, the expectation that unemployment levels will increase sharply, as opposed to falling, is estimated to decrease the propensity to spend by 7.7 pp. Also, when a consumer is accumulating a higher debt burden, or running into debt, the overall readiness to spend declines by 4.0pp. These latter effects are present and broadly the same both when the lower bound is binding and when it is not.

[Insert Table 3 here.]

3.4 Country heterogeneity

All of the above empirical evidence relates to pooled estimates for the Euro Area sample as a whole. However, it is of clear interest to consider in detail how the estimated impact

of inflation expectations may differ across individual countries in the Euro Area. We have therefore re-estimated the ordered logit model individually for each euro area country included in our sample, again using the general specifications with the broadest set of controls and outlier correction. In general, we find that country results²¹ confirm aggregate EA results. All countries except one, show a positive relationship between the expected change in subjective inflation and the readiness to spend (see Figure 6). The only exception is Malta for which we find average marginal effects of -0.22 outside the ELB and an even stronger negative one of -0.36 at the ELB. Of course, among the countries which exhibit a positive marginal effect, we do find considerable heterogeneity as the range for the effects stands between 0.06 and 0.94 pp. Many countries show effects around the EA estimates both outside the ELB and at the ELB. Portugal and Greece seem to have a weak, close to zero impact, meaning that for consumers in these countries inflation expectations do not matter much in the consumption decision. Finland exhibits the highest marginal effects of 0.7 outside the ELB period and 0.94 when the ELB is binding, making Finish consumers the most sensitive to inflation expectations. However, in the case of several large euro area countries, e.g. Germany, France, Italy, Spain and the Netherlands, we observe marginal effects that are consistently higher when the ELB is binding. Indeed, for 14 of the 17 country cases that we consider, the impact when positive is stronger when the ELB is binding. The results are therefore very consistent with the findings in Ichiue and Nishiguchi (2015) for Japanese consumers, suggesting that once confronted with a prolonged period of low and close to zero deposit interest rates on their savings, consumers in many euro area countries have become more sensitive to the importance of inflation expectations when deciding between current consumption relative to saving and consuming in the future.

[Insert Figure 6 here.]

4 Linking micro and macro: A VAR analysis

In the previous section we have estimated the impact of consumer inflation expectations on the propensity to spend. Nevertheless, the question of whether or not this relationship translates into an impact on actual consumption remains unanswered. For this purpose, we

²¹At a country level we do not report results for Estonia, as information for the consumer inflation perceptions was available only at the beginning of the sample.

use a bi-variate VAR to model the relationship between aggregate real consumption and the average estimated consumer propensity to spend derived from our ordered logit model. Such a VAR provides a bridge which can link the micro evidence in our discrete choice regressions with macroeconomic aggregates and, as we illustrate below, can also enable macro level simulations and scenario analysis.

We first estimate the spending readiness of the "representative consumer" in the euro area as a weighted average of the individual fitted probabilities obtained from our logit regressions. In constructing the average probability we use individual consumer weights based on the representativeness of a consumer in total population and therefore control for variation in survey samples across countries. This representative probability summarizes all the micro and macro level information that we have included in our ordered logit specification. Hence, a bivariate VAR capturing the interaction between the log of real total private consumption and this spending propensity measure seems most appropriate and a multivariate analysis which adds other relevant macroeconomic controls in the VAR is not necessary.

[Insert Figure 7 here.]

Our derived measure of aggregated readiness to spend co-moves quite strongly with real consumption (see Figure 7). Moreover, it also seems to anticipate the drop in consumption during the Great Recession (2008 - 2009) and the sovereign debt crisis (2011 - 2012) as well as the subsequent recoveries. Figure 8 shows the impulse response functions based on a Cholesky decomposition, with the average propensity to spend variable (Prob) ordered first in the bivariate VAR. Thus, we assume that in the first period the consumer's readiness to spend does not react to a shock in log consumption ($\ln C$). This identifying restriction is justified by the fact that we average spending probabilities using only data from the first month of each quarter. Hence we should not expect survey responses to be contemporaneously impacted by real consumption shocks which cannot be anticipated at the start of the quarter. We employ the Schwarz and Hannan-Quinn information criteria to determine the appropriate lag length to use in the VAR. Both criteria suggest a relatively parsimonious specification with only one lag of each variable needed. The estimated impulse response functions behave in an economically intuitive manner: an initial positive shock in the spending probability leads to a persistent

increase which slowly decays toward zero. Importantly, consumption slowly increases after a shock in the spending probability and the positive effect lasts for eight to ten quarters after the initial shock. Interestingly, the propensity to spend does not react at all to a shock in real consumption which suggests that it can largely be treated as exogenous within the VAR. After a consumption shock, consumption increases with high persistence with the effect lasting close to ten quarters.

[Insert Figure 8 here.]

Based on this VAR, we implement two scenarios capturing the effects of a 2.0 pp change in inflation expectations. These scenarios can be interpreted as illustrating the impact of moving from a situation of undesirably low inflation (e.g. at 0%) to a rate of inflation more in line with price stability (e.g. of 2%) whilst holding perceptions about past inflation constant. A first scenario depicts the effects of such a change when the ELB is binding while a second is intended to capture the impacts when the economy is away from the ELB. The 2.0 pp expected increase in inflation is implemented gradually as a 0.5pp increase that takes place over four consecutive quarters. Once again, we compute the change in the spending probability associated with an 0.5 pp increase in inflation expectations based on the specification which (i) includes the broadest set of possible controls and (ii) adjusts for outliers. According to this specification, a 0.5 pp increase in inflation expectations is associated with approximately 0.19 pp increase in the spending probability when the ELB is binding and a smaller 0.15 pp increase when it is not.

We then use the estimated results of the above VAR to trace out the impacts of the corresponding "shocks" to the spending probability on real consumption under each scenario.²² Figure 9 shows the difference in the annual real consumption growth rates relative to a baseline scenario where inflation expectations and consumers' spending probabilities are kept constant throughout the 12 quarter simulation period. In cumulative terms, real consumption rises by 0.36% over the three year horizon when the ELB is binding and by just over 0.26% when it is not. Overall, we would therefore conclude that the impact of inflation expectations on actual

²²In the ELB scenario, the path of the aggregate consumer propensity to spend increases by 0.19 pp in each of the four consecutive quarters and then remains constant for the next eight quarters. Outside the ELB scenario, in line with the above estimates, the shock to the consumer propensity to spend is lower and is fixed at 0.15 in each of the first four quarters.

private consumption is not just significant in a statistical sense, but it is also quantitatively relevant in economic terms.

[Insert Figure 9 here.]

The above approach, focusing on the estimated probabilities from our micro level model differs somewhat from previous attempts in the literature to convert survey findings into quantitative effects on actual consumption spending. Bachmann et al. (2015) also use a bivariate VAR in which they include an aggregate index for buying conditions, which is measured by the fraction of people saying that now is a good moment to buy durable goods minus those reporting that now is a bad moment to buy, and the HP-filtered natural logarithm of real durable consumption expenditures. These authors then report impulse response functions for which they calibrate the size of the innovation corresponding to the aggregate index such that it corresponds to the marginal effect of a 1pp increase in inflation expectations as computed based on their micro-data analysis. In line with their negligible and insignificant estimated marginal effects, they find that the impact is almost zero outside the ELB and about -0.1% at the ELB. In order to estimate the impact on real consumption, D'Acunto et al. (2016) perform a "back-of-the-envelope" calculation and regress the natural logarithm of real durable consumption expenditure on the end of quarter value of the average durable purchasing propensity and quarterly dummies. They find 4.8% higher real durable consumption if all Germans would expect higher inflation as opposed to prices not changing. This impact is in line with the direction of the effect that we estimate but appears at first pass to be considerably higher in magnitude. However the quantitative differences with our results should not be overstated given differences in the underlying scenario. In particular, the higher impact of D'Acunto et al. (2016) relates to a general and qualitative increase in inflation expectations and not to the 2.0 pp quantitative increase that we have implemented over four quarters. Moreover, their results relate only to durable consumption.

5 Conclusion

In this paper we analyse the relationship between consumer inflation expectations, consumer's readiness to spend and actual private consumption. Although this relationship stands at the very heart of mainstream macroeconomic models, there are very few papers that have provided robust empirical evidence on its nature and magnitude, or in a way that allows for a comparative analysis across countries. We investigate this relationship for the Euro Area and 17 of its constituent countries. To the best of our knowledge, we are the first to provide such a comprehensive multi-country view on this central issue for understanding Euro Area consumer behaviour. We benefit from a very rich micro dataset provided by the EU Consumer Survey which provides information about consumer's expectations for economic and financial conditions, and which can be broken down according to a set of demographic characteristics. Most importantly for our analysis, the survey includes quantitative consumer expectations and perceptions and gives information about the consumer's intentions to spend.

Our evidence is based on just over 2,000,000 observations,²³ which were carefully collected to reflect Euro Area population over a 13¹/₂ year period, from May 2003 to December 2016. As the survey does not include information about actual consumption, we perform our analysis in two steps. First, we estimate the relationship between consumers' inflation expectations and their readiness to spend based on the survey data. In doing so, we use an innovative measure of the expected change in subjective inflation, which reflects the difference between a consumer's expectation about future inflation and his/her subjective perception about current inflation. Econometrically, this difference helps take account of possible variation in unobserved consumer characteristics that may otherwise distort our results. We find this measure the most economically relevant because it captures the idea that when determining their spending intentions, consumers are likely to consider expected inflation relative to their currently perceived level of inflation. Indeed, the data reveal that the latter can vary widely across households and time and, hence, a failure to take this into account would likely bias our analysis. In a second step, using macro level information, we employ a simple VAR framework to translate the estimated impact of inflation expectations on current spending readiness into an effect on actual private consumption expenditure in the Euro Area.

²³This is the number of observations that we are left with after eliminating those belonging to consumers that did not reply to the full set of questions that we use in this analysis.

Our results suggest that Euro Area consumers behave in line with the predictions of economic theory. When they expect higher inflation in the future, all other factors held constant, they adjust positively their intention to spend at the current moment. The result is robust across several specifications, in which we gradually control for demographics, other consumer expectations, the consumers' financial situation, interactions of inflation expectations with various controls, time dummies, country dummies and other common factors such as oil prices and bank lending rates. Our pooled results suggest that for a 1 pp expected increase in inflation the consumers' probability to spend increases by between 0.16 pp and 0.39 pp. This is confirmed by country-level results, where for almost all countries, we find a positive relationship between consumer inflation expectations and propensity to spend, though there is some notable heterogeneity. This result complements and extends the existing empirical literature on consumer behaviour using survey data: Bachmann et al. (2015), Ichiue and Nishiguchi (2015), D'Acunto et al. (2016), Burke and Ozdagli (2013), Armantier et al. (2015). While our results for the EA and most of its constituent countries confirm previous findings for Germany and Japan, they differ from recent findings for the US, as in Bachmann et al. (2015) and Burke and Ozdagli (2013).

Another key finding in our study is that the relationship between consumer inflation expectations and consumption becomes stronger at the ELB when interest rates are bounded from below. At the ELB, fluctuations in the real interest rate respond one for one to changes in expected inflation because nominal rates can no longer adjust to mitigate the real effects of a change in inflation expectations. Our results thus suggest that once confronted with a lower bound on nominal rates, consumers become more aware of and sensitive to changes in real interest rates and the effect of inflation expectations on spending behaviour is augmented. We observe this change in the relationship due to the lower bound both in our pooled results and also in a large majority of our country-specific results. These differences are also shown to be economically relevant in our VAR simulations. When comparing a scenario in which inflation expectations increase by 2.0pp relative to a scenario where inflation expectations remain unchanged, we find that real consumption rises by 0.36% over a three year period at the ELB and by 0.26% outside the ELB.

Overall, from a monetary policy perspective, our micro analysis of consumer behaviour provides evidence on the importance of intertemporal substitution linked to the real interest rate channel. Moreover, this channel appears to be stronger when the economy is constrained by the lower bound on nominal interest rates. This provides strong support for central bank concerns about a drop in inflation expectations during lower bound episodes because such developments have the potential to weaken aggregate demand further by raising real rates. Equally, our results point to the important stabilisation role that higher inflation expectations could play when the economy is constrained by the lower bound.

References

- Anderson, R. (2008). US consumer inflation expectations: evidence regarding learning, accuracy and demographics, *Centre for Growth and Business Cycle Research, The University of Manchester Discussion Paper Series No 99*.
- Arioli, R., Bates, C., Dieden, H., Duca, I., Friz, R., Gayer, C., Kenny, G., Meyler, A. and Pavlova, I. (2016). EU consumers' quantitative inflation perceptions and expectations: An evaluation, *European Economy Discussion Paper No 38, DG ECFIN, European Commission*.
- Armantier, O., Bruine de Bruin, W., Topa, G., van der Klaauw, W. and Zafar, B. (2015). Inflation expectations and behavior: Do survey respondents act on their beliefs?, *International Economic Review* **56**(2): 505–536.
- Armantier, O., Nelson, S., Topa, G., van der Klaauw, W. and Zafar, B. (2014). The price is right: Updating inflation expectations in a randomized price information experiment, *Review of Economics and Statistics* **98**(3): 503–523.
- Bachmann, R., Berg, T. O. and Sims, E. R. (2015). Inflation expectations and readiness to spend: Cross-sectional evidence, *American Economic Journal: Economic Policy* **7**(1): 1–35.
- Bernanke, B. (2017). Temporary price-level targeting: an alternative framework for monetary policy, *Brookings blog post, October 12, 2017*.

- Binder, C. C. (2017). Measuring uncertainty based on rounding: New method and application to inflation expectations, *Journal of Monetary Economics* **90**: 1–12.
- Blanchard, O., Dell’Ariccia, G. and Mauro, P. (2010). Rethinking macroeconomic policy, *Journal of Money, Credit and Banking* **42**(Supplement): 199–215.
- Blanchard, O., Dell’Ariccia, G. and Mauro, P. (2013). Rethinking macro policy II: Getting granular, *IMF Discussion Notes No 13/03, International Monetary Fund* .
- Bryan, M. F. and Venkatu, G. (2001). The demographics of inflation opinion surveys, *Federal Reserve Bank of Cleveland, Economic Commentary, 10.15.2001* .
- Burke, M. A. and Ozdagli, A. (2013). Household inflation expectations and consumer spending: Evidence from panel data, *Federal Reserve Bank of Boston Working Paper Series No 13-25, December* .
- Carroll, C. D. (2003). Macroeconomic expectations of households and professional forecasters, *The Quarterly Journal of Economics* **118**(1): 269–298.
- Christensen, C., van Els, P. and van Rooij, M. (2006). Dutch households’ perceptions of economic growth and inflation, *De Economist* **154**(2): 277–294.
- Coibion, O. and 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. and Kamdar, R. (2018). The formation of expectations, inflation and the Phillips curve, *Journal of Economic Literature* **Forthcoming**.
- Crump, R. K., Eusepi, S., Tambalotti, A. and Topa, G. (2015). Subjective intertemporal substitution, *Federal Reserve Bank Of New York Staff Reports No 734, July* .
- D’Acunto, F., Hoang, D. and Weber, M. (2016). The effect of unconventional fiscal policy on consumption expenditure, *NBER Working Paper No 22563, August* .
- Davidson, R. and MacKinnon, J. G. (2004). *Econometric theory and methods*, Oxford University Press New York.

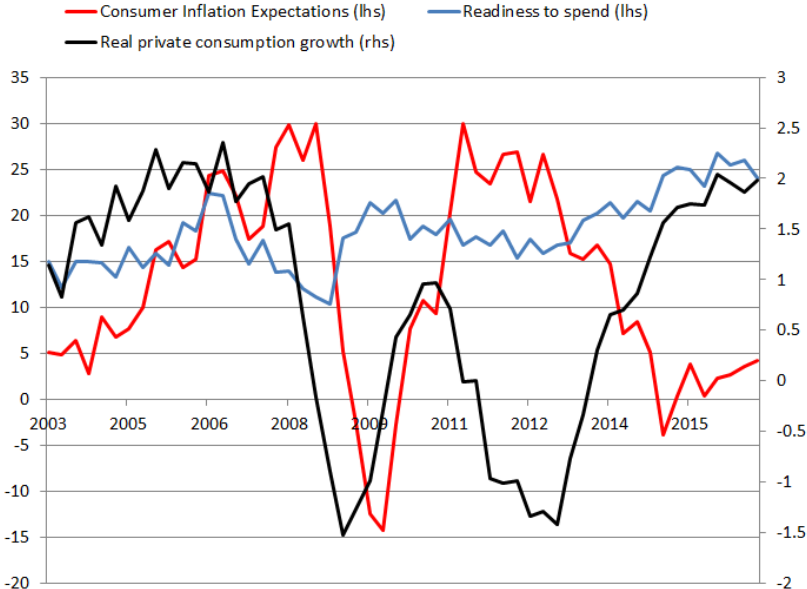
- Eggertsson, G. B. (2006). The deflation bias and committing to being irresponsible, *Journal of Money, Credit and Banking* **38**(2): 283–321.
- Eggertsson, G. B. and Woodford, M. (2003). Zero bound on interest rates and optimal monetary policy, *Brookings Papers on Economic Activity* (2003(1)): 139–233.
- Ehrmann, M. (2011). Inflation developments and perceptions after the euro cash changeover, *German Economic Review* **12**(1): 33–58.
- Ehrmann, M., Pfajfar, D. and Santoro, E. (2017). Consumers' attitudes and their inflation expectations, *International Journal of Central Banking* **13**(1): 225–259.
- Ichiue, H. and Nishiguchi, S. (2015). Inflation expectations and consumer spending at the zero bound: Micro evidence, *Economic Inquiry* **53**(2): 1086–1107.
- Jonung, L. (1981). Perceived and expected rates of inflation in Sweden, *The American Economic Review* **71**(5): 961–968.
- Jung, T., Teranishi, Y. and Watanabe, T. (2005). Optimal monetary policy at the zero-interest-rate bound, *Journal of Money, Credit and Banking* **37**(5): 813–835.
- Kliesen, K. L. (2015). How accurate are measures of long-term inflation expectations?, *Federal Reserve Bank of Saint Louis Economic Synopses No 9, 2015* .
- Krugman, P. R., Dominquez, K. M. and Rogoff, K. (1998). It's baaack: Japan's slump and the return of the liquidity trap, *Brookings Papers on Economic Activity* (1998(2)): 137–205.
- Lombardelli, C. and Saleheen, J. (2003). Public expectations of UK inflation, *Bank of England Quarterly Bulletin, Autumn 2003* .
- Malmendier, U. and Nagel, S. (2016). Learning from inflation experiences, *The Quarterly Journal of Economics* **131**(1): 53–87.
- Mankiw, N. G. and Reis, R. (2001). Sticky information: A model of monetary nonneutrality and structural slumps, *NBER Working Paper No 8614, December* .
- Manski, C. F. (2004). Measuring expectations, *Econometrica* **72**(5): 1329–1376.
- Naohito, A. and Yuko, U. (2015). Measuring inflation expectations: Consumers' heterogeneity and nonlinearity, *RCESR Discussion Paper Series No 15-5, June* .

Souleles, N. S. (2004). Expectations, heterogeneous forecast errors, and consumption: Micro evidence from the Michigan consumer sentiment surveys, *Journal of Money, Credit and Banking* **36**(1): 39–72.

Vestin, D. (2006). Price-level versus inflation targeting, *Journal of Monetary Economics* **53**(7): 1361–1376.

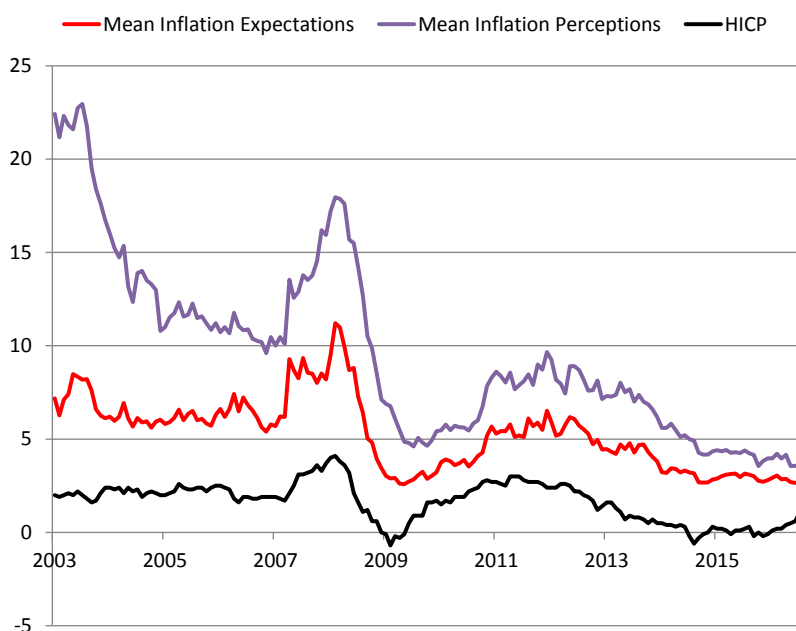
Charts and tables

Figure 1: Consumer inflation expectations, real private consumption and readiness to spend



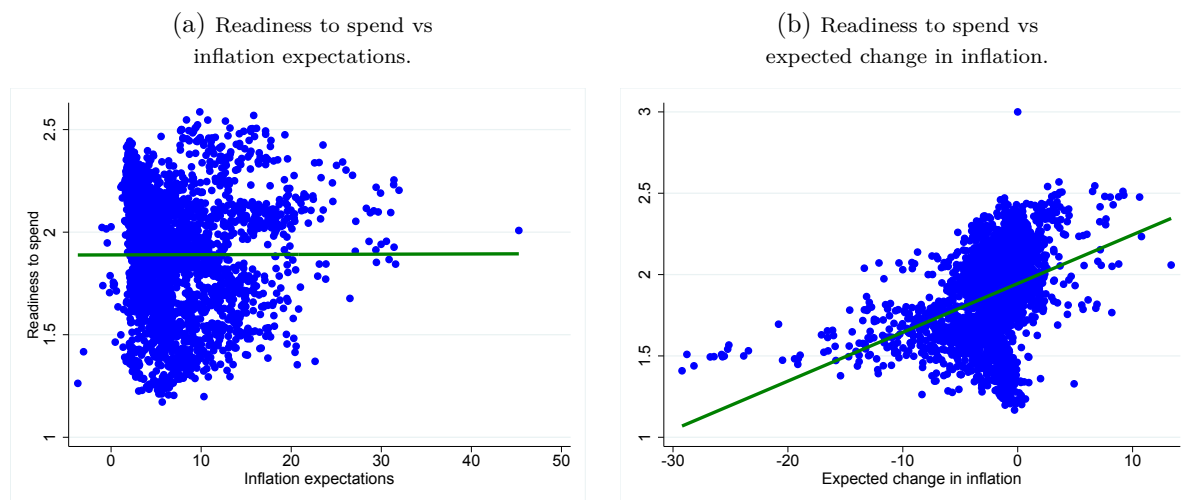
Note: Quarterly data, 2003Q4 - 2016Q4. Consumer Inflation Expectations - price trends over next 12 months (balance statistics); Readiness to spend - major purchases at present, percent positive replies, source: DG-ECFIN Consumer Survey; Real private consumption - household consumption expenditure, Eurostat.

Figure 2: Mean inflation expectations and perceptions vs HICP



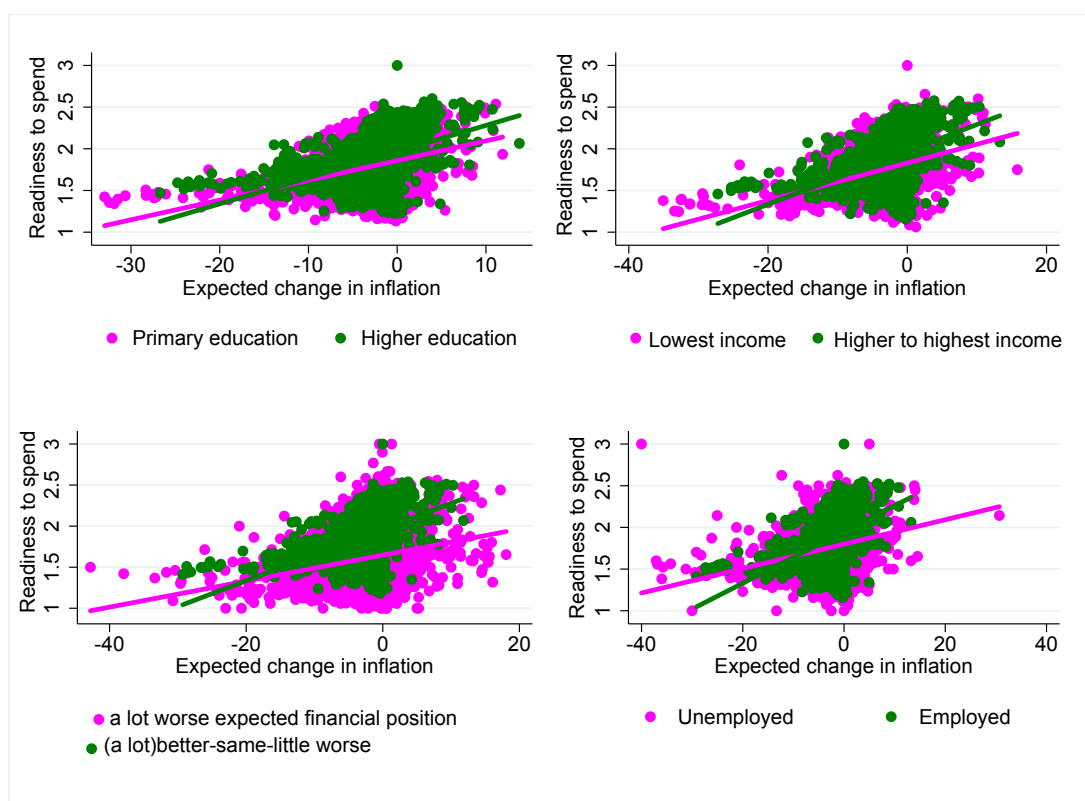
Note: Individual and country weights used for aggregation. Time period covered: May 2003 - December 2016.

Figure 3: Scatterplots of readiness to spend vs inflation expectations and expected change in inflation



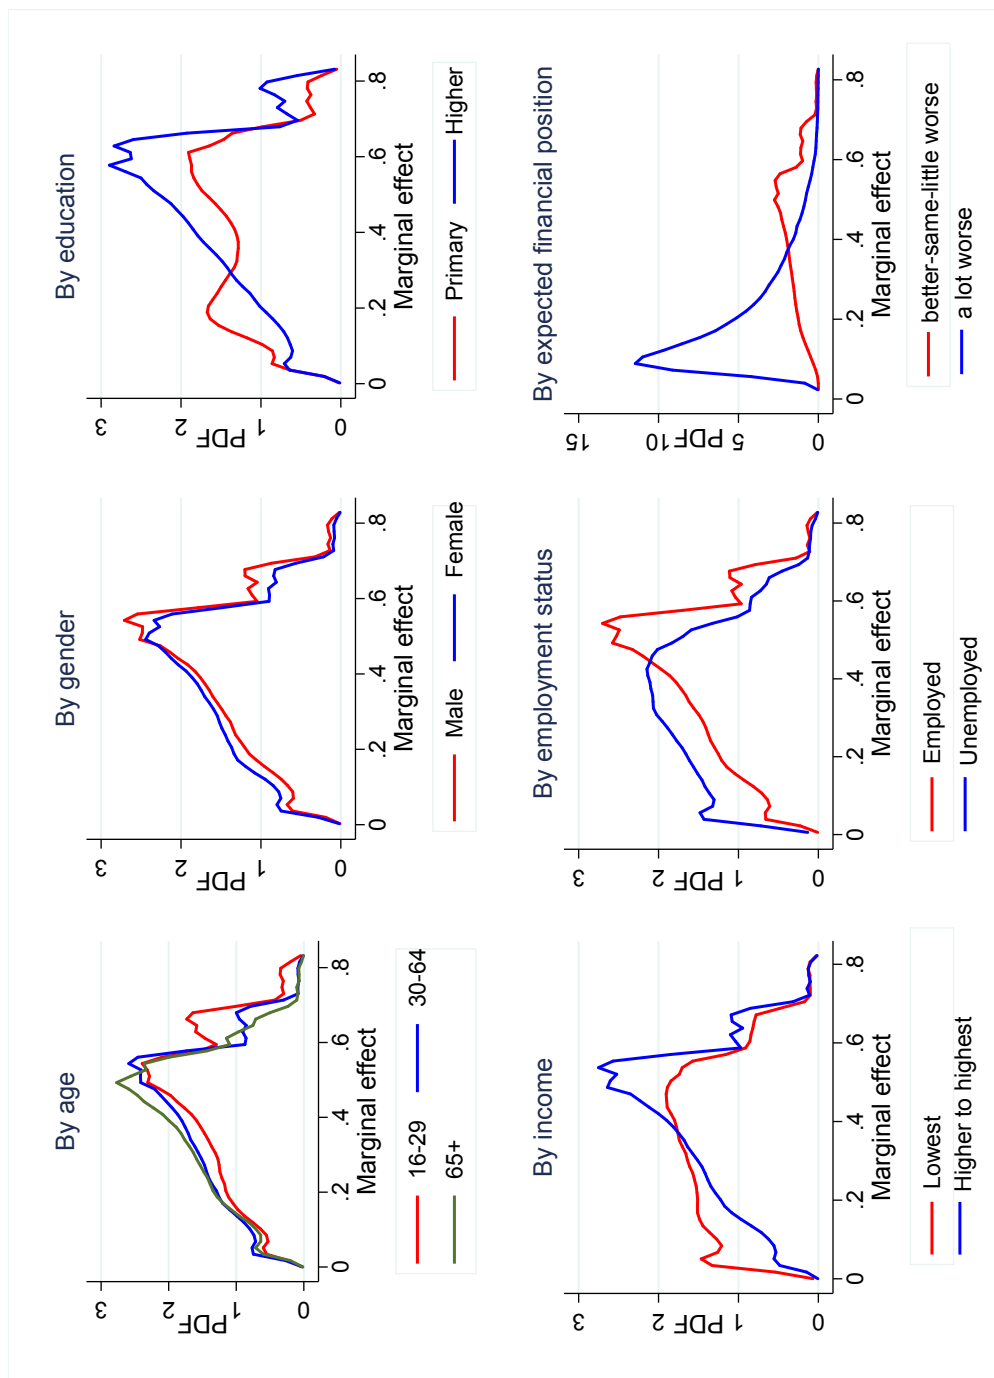
Note: One dot represents a country aggregate (weighted by individual weights) at one moment in time (identified by month and year) of (a) inflation expectations and (b) expected changes in inflation. Readiness to spend is coded 1 for not being the right moment to spend, 2 for being neither the right moment nor the wrong moment and 3 for being the right moment to spend.

Figure 4: Scatterplots of expected change in inflation vs readiness to spend differentiating by consumer characteristics



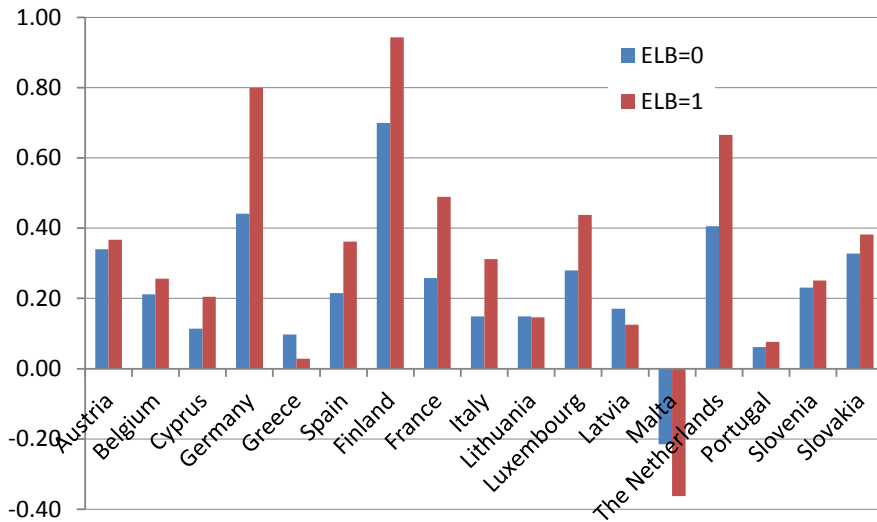
Note: One dot represents a simple average of a particular category of consumer (as defined by education, income, expected financial situation and employment status) at one moment in time (identified by month and year).

Figure 5: Probability density functions of marginal effects distinguishing different consumer characteristics



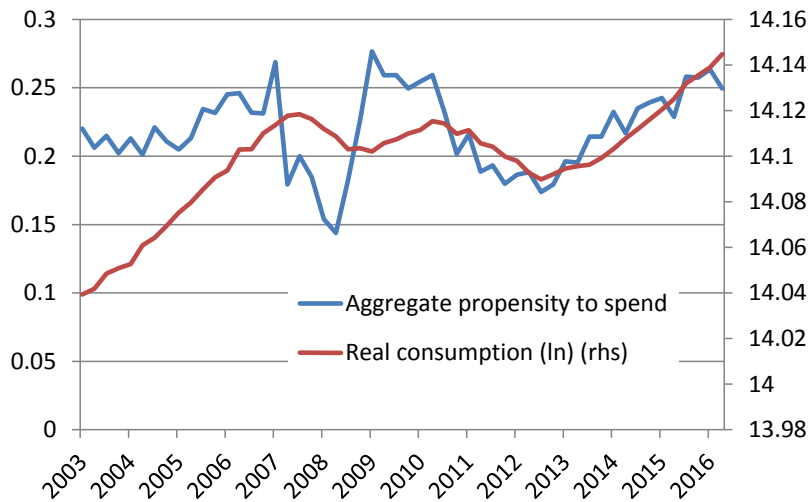
Note: PDF stands for Probability Density Function, Marginal effects refer to the effect of a 1 pp increase in the expected change in inflation on a consumer's readiness to spend.

Figure 6: Country average marginal effects



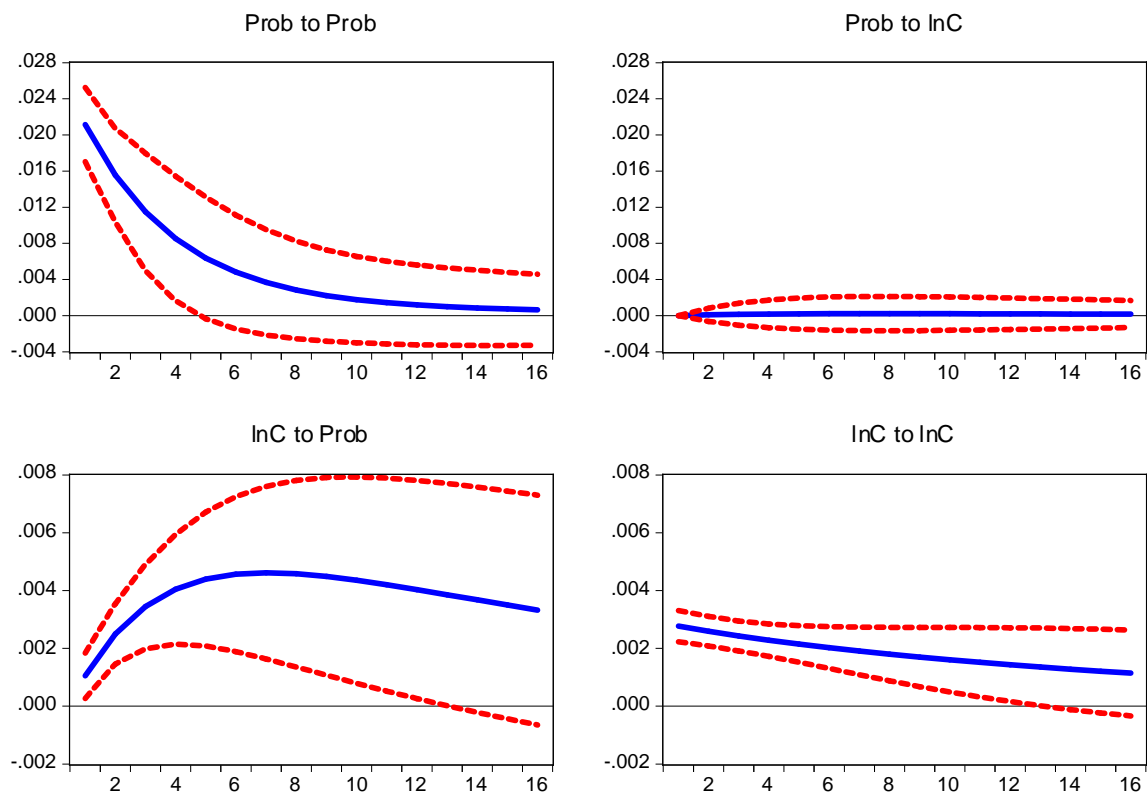
Note: This figure presents average marginal effects based on the ordered logit specification which includes the following groups of controls: "Demographics", "Controls I", "Time dummies", "Controls II", "Interactions" and "Macro aggregates". The 5% and 95% cuts are used for winsorizing. For Malta we report results based on an ordered logit specification which excludes "Macro-aggregates" due to lack of data on the lending rates. All average marginal effects are statistically significant at 1% level, except the average marginal effect at the ELB for Greece.

Figure 7: Aggregated propensity to spend and log real consumption



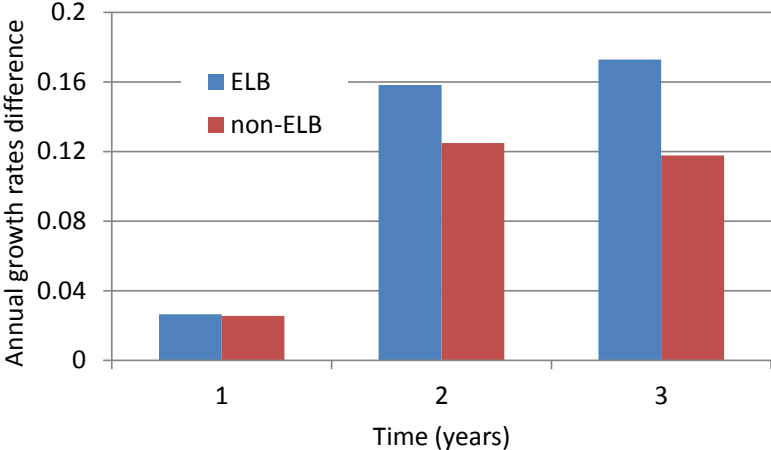
Quarterly frequency, 2003Q1 - 2016Q4. The aggregate propensity to spend is computed as the weighted average of the individual fitted probabilities obtained from the logit regression in Table 3 (full specification, winsorized data sample) and reduced to quarterly frequency by taking the aggregate corresponding to the first month of the quarter.

Figure 8: Bivariate VAR: impulse response functions



Note: Impulse response functions based on a Cholesky factorisation using a bivariate VAR(1) including consumer aggregate propensity to spend and log real total consumption (quarterly frequency), in this order. Prob refers to the fitted aggregated probability, i.e. the aggregated propensity to spend and lnC refers to log real total consumption. Confidence bands represent ± 2 standard errors.

Figure 9: Impact on real consumption of a gradual increase in consumer inflation expectations



Note: Based on the conditional forecasts from a VAR(1) model which includes a quarterly aggregate measure of the consumer propensity to spend and the natural logarithm of real total consumption. Each forecast assumes a 2.0 percentage point increase in consumer inflation expectations spread over four consecutive quarters. Each forecast conditions on a shock to the consumer propensity to spend which corresponds to the assumed rise in inflation expectations. The size of the shock is calibrated using the estimated marginal effects from a change in inflation expectations derived from the ordered logit model both when the ELB is binding and when it is not.

Table 1: Mean and median inflation expectations and perceptions over May 2003 - December 2016

	Inflation expectations		Inflation perceptions	
	Mean	Median	Mean	Median
Gender				
Male	6.4	3.0	8.8	5.0
Female	8.1	5.0	11.4	6.0
Age				
16-29	7.8	4.0	10.4	5.0
30-49	7.5	3.5	10.4	5.0
50-64	7.2	3.0	9.9	5.0
65+	6.6	2.5	9.7	5.0
Education				
Primary	7.5	3.0	11.4	5.0
Secondary	7.7	4.0	10.3	5.0
Further	6.4	3.0	8.4	5.0
Income				
1st quartile	8.5	4.0	11.5	6.0
2nd quartile	7.3	3.0	10.2	5.0
3rd quartile	7.0	3.0	9.4	5.0
4th quartile	6.6	3.0	8.9	5.0
Employment status				
Unemployed	8.0	3.0	10.7	5.0
Employed	7.3	3.0	10.0	5.0
Euro Area	5.2	2.0	9.4	3.5

Table 2: Propensity to spend: average marginal effects of the expected change in inflation, Euro Area

Average marginal effects											
ELB=0	ELB=1	Demographics	Controls I	Time and country dummies	Controls II	Interactions	Macro aggregates	Winsorize	Pseudo R ²	Observations	P-value
0.00152*** (0.000242)	0.00199*** (0.000352)	x	x	x					0.0759	2,055,368	0.031
0.00280*** (0.000465)	0.00398*** (0.000661)	x	x	x				x	0.0762	2,055,368	0.002
0.00155*** (0.000228)	0.00195*** (0.000351)	x	x	x	x				0.0796	2,013,956	0.062
0.00170*** (0.000242)	0.00210*** (0.000336)	x	x	x	x	x			0.0797	2,013,956	0.045
0.00160*** (0.000231)	0.00199*** (0.000336)	x	x	x	x	x	x		0.0809	1,849,151	0.061
0.00293*** (0.000406)	0.00396*** (0.000644)	x	x	x	x	x	x	x	0.0812	1,849,151	0.002

Note: Clustered standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. The table shows the average marginal effect of a unit increase in the expected change in inflation on the probability that consumers are ready to spend given current conditions, estimated by the ordered logit model. "Demographics" includes age, gender, education, employment status, income; "Controls I" includes expectations of an individual's financial situation and of the general economic situation; "Time and country dummies" includes quarterly time dummies from 2003 to 2016 and country dummies; "Controls II" includes expectations of the general unemployment situation and consumer current financial status, i.e. debtor or non-debtor; "Interactions" includes pairwise interactions as follows: expected change in inflation _ the expected financial situation, expected change in inflation _ debt status, expected change in inflation _ employment status, expected change in inflation _ income, expected change in inflation _ education; "Macro aggregates" includes oil prices and lending rates; ZLB dummy takes value 1 from June 2014 to December 2016. When winsorizing, the 5% and 95% cuts are used. The last column "P-value" indicates the p-value from a likelihood ratio test of the hypothesis that the marginal effects at the ELB and outside the ELB are identical.

Table 3: Full specification: average marginal effects

Variables	ELB=0	ELB=1
Expected change in inflation	0.00293*** (0.000406)	0.00396*** (0.000644)
ELB	0.0123** (0.00619)	0.0123** (0.00619)
Oil price	-0.000322*** (0.000116)	-0.000331*** (0.000118)
Lending rates	-0.00487*** (0.00169)	-0.00501*** (0.00175)
Age (30 - 49)	-0.0173*** (0.00357)	-0.0178*** (0.00365)
Age (50 - 64)	-0.00425 (0.00905)	-0.00437 (0.00931)
Age (65+)	0.00562 (0.0132)	0.00577 (0.0136)
Gender (Female)	-0.0165*** (0.00468)	-0.0170*** (0.00467)
Education (Secondary)	0.0243*** (0.00296)	0.0249*** (0.00309)
Education (Further)	0.0345*** (0.00442)	0.0354*** (0.00472)
Income (2nd Quartile)	0.0227*** (0.00480)	0.0234*** (0.00480)
Income (3rd quartile)	0.0371*** (0.00827)	0.0382*** (0.00839)
Income (4th quartile)	0.0655*** (0.0114)	0.0673*** (0.0116)

Table 3 Continued: Full specification: average marginal effects

Variables	ELB=0	ELB=1
Employment status (Employed)	0.0241*** (0.00557)	0.0248*** (0.00573)
Debt status (debtor)	-0.0400*** (0.0106)	-0.0412*** (0.0110)
Expected financial situation (a little worse)	0.0375*** (0.00662)	0.0390*** (0.00694)
Expected financial situation (same)	0.0712*** (0.0105)	0.0737*** (0.0108)
Expected financial situation (a little better)	0.0844*** (0.0123)	0.0874*** (0.0128)
Expected financial situation (a lot better)	0.105*** (0.0130)	0.109*** (0.0135)
Expected general economic situation (a little worse)	0.0355*** (0.00499)	0.0367*** (0.00509)
Expected general economic situation (the same)	0.0505*** (0.00649)	0.0522*** (0.00647)
Expected general economic situation (a little better)	0.0884*** (0.0109)	0.0910*** (0.0111)
Expected general economic situation (a lot better)	0.127*** (0.0164)	0.130*** (0.0164)
Expected general unemployment situation (fall slightly)	0.0162 (0.0104)	0.0165 (0.0105)
Expected general unemployment situation (the same)	-0.0141 (0.0127)	-0.0145 (0.0132)

Table 3 Continued: Full specification: average marginal effects

Variables	ELB=0	ELB=1
Expected general unemployment situation (increase slightly)	-0.0342** (0.0134)	-0.0352** (0.0141)
Expected general unemployment situation (increase sharply)	-0.0767*** (0.0153)	-0.0791*** (0.0164)
Observations	1,849,151	1,849,151

Note: Clustered standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The table shows the average marginal effect estimated with an ordered logit model. The results show the average marginal effect of one unit increase in inflation expectations on the probability that consumers are ready to spend. For the discrete variables, the reported marginal effect shows the discrete change from a base alternative, e.g. Income (1st Quartile) to another alternative, e.g. Income (4th Quartile). The ordered logit regression includes all groups of controls: "Demographics", "Controls I", "Time and country dummies", "Controls II", "Interactions" and "Macro aggregates". The 5% and 95% cuts are used for winsorizing.

Data Appendix

In this appendix we describe our main survey and other data sources.

EU Consumer Survey: The EU Consumer survey data used in the study are collected under the framework of the Joint Harmonised EU Programme of Business and Consumer Surveys (BCS). The programme is administered by the European Commission (EC). Its consumer survey is the largest of its kind, covering the 28 European Union (EU) member states,²⁴ as well as all five candidate countries, with up to 41,060 respondents included in each monthly round (26,440 if only euro area EU countries are included). For comparison, the University of Michigan Consumer Survey includes 500 respondents, while the Federal Reserve Bank of New York Survey of Consumer Expectations (SCE) has 1300 households in its sample.

The BCS is carried out at a monthly frequency. Our sample starts in May 2003 as this represents the first month in which quantitative questions on inflation perceptions and inflation expectations were introduced into the Consumer Survey and it runs until December 2016. A complete description of the survey is available on the European Commission's website at [EC Consumer Survey](#). As described in further detail on this website, the surveys are carried out at a national level by partner institutes such as ministries, statistical offices, central banks, research institutes and private companies. A list of the partner institutes is available on the Commission's website here: [Partner Institutes](#). These partner institutes retain ownership of the data for each of their countries. However, for a large majority of the countries studied in this paper, the underlying country data can be accessed directly by contacting these partner institutes. The anonymised micro data set on quantitative inflation perceptions and expectations is described in Arioli et al. (2016). Based on the agreement with all national partner institutes, it was also provided to the ECB by the European Commission's DG ECFIN for joint research purposes. Weights used for the construction of euro area averages from country data were obtained from Eurostat.

Additional survey questions: As described in detail in Section 2.1, our econometric analysis makes extensive use of survey questions 51, 61 and 8 on, respectively, perceived

²⁴In the case of Ireland, the consumer survey currently does not cover the full list of harmonised questions.

inflation, expected inflation and the readiness to make major purchases. However, the ordered logit specification also makes use of additional questions as controls and interaction variables to check for heterogeneity in the impact of inflation expectations on readiness to spend. In particular, we exploit the following additional questions on the expected household financial position, the consumer's expectations about the general economic situation, their expectations about the level of unemployment and whether they are currently running into debt.

Q2: How do you expect the financial position of your household to change over the next 12 months? It will i) get a lot better; ii) get a little better; iii) stay the same; iv) get a little worse; v) get a lot worse and vi) don't know.

Q4: How do you expect the general economic situation in this country to develop over the next 12 months? It will i) get a lot better; ii) get a little better; iii) stay the same; iv) get a little worse; v) get a lot worse and vi) don't know.

Q7: How do you expect the number of people unemployed in this country to change over the next 12 months? The number will i) increase sharply; ii) increase slightly; iii) remain the same; iv) fall slightly; v) fall sharply and vi) don't know.

Q12: Which of these statements best describes the current financial situation of your household? i) we are saving a lot, ii) we are saving a little, iii) we are just managing to make ends meet on our income, iv) we are having to draw on our savings, v) we are running into debt.

As with Q8 on consumer readiness to spend, these questions are all qualitative in nature. We therefore exploit the responses relative to a base response category whilst excluding all consumers who reply "don't know". For each of the four questions category v) is chosen as this base category.

HICP Inflation: Annual change in the euro area Harmonised Index of Consumer Prices, neither seasonally, nor working day adjusted. Source: Eurostat.

Euro area private consumption: World concept, households and non-profit institutions serving households, Euro, chain linked volumes, calendar and seasonally adjusted. Source: Eurostat.

Household lending rates: Lending rates for loans for consumption excluding revolving loans and overdrafts, convenience and extended credit card debt. ECB MFI Statistics - Statistical Data Warehouse from national sources.

Oil prices: Brent oil commodity price index (2010=100). Source: International Monetary Fund.

Acknowledgements

We thank Olivier Armantier, Rudi Bachmann, Heinz Dieden, Maarten Dossche, Lena Dräger, Michael Ehrmann, Dimitris Georganakos, Christian Gayer, Yuriy Gorodnichenko, Luc Laeven, Michele Lenza, Bartosz Mackowiak, Michael Mc Mahon, Aidan Meyler, Francesco Mongelli, Jirka Slacalek and Georg Strasser for helpful comments as well as seminar participants at the ECB, the University of Hamburg, the European Commission, the 2017 RCEA, the 2017 Money-Macro-Finance Annual Conference and the 8th Ifo Conference on Macroeconomics and Survey Data. The views expressed in the paper belong to the authors and are not necessarily the views of either the European Central Bank or the European Commission. Any errors are the sole responsibility of the authors.

Ioana A. Duca

European Central Bank, Frankfurt am Main, Germany; email: ioana.duca@ecb.europa.eu

Geoff Kenny

European Central Bank, Frankfurt am Main, Germany; email: geoff.kenny@ecb.europa.eu

Andreas Reuter

European Commission, Brussels, Belgium; email: andreas.reuter@ec.europa.eu

© European Central Bank, 2018

Postal address 60640 Frankfurt am Main, Germany

Telephone +49 69 1344 0

Website www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the [Social Science Research Network electronic library](#) or from [RePEc: Research Papers in Economics](#). Information on all of the papers published in the ECB Working Paper Series can be found on the [ECB's website](#).

PDF

ISBN 978-92-899-3301-8

ISSN 1725-2806

doi:10.2866/67230

QB-AR-18-076-EN-N