

University of Tartu
School of Economics and Business Administration

**INFLATION EXPECTATIONS AND
CONSUMPTION WITH MACHINE LEARNING**

Diana Gabrielyan, Lenno Uusküla

Tartu 2022

ISSN-L 1406-5967
ISSN 1736-8995
ISBN 978-9985-4-1326-5 (pdf)
The University of Tartu FEBA
<https://majandus.ut.ee/en/research/workingpapers>

Inflation expectations and consumption with machine learning

Diana Gabrielyan¹, Lenno Uusküla^{2,3}

Abstract

We extract measures of inflation expectations from online news to build real interest rates that capture true consumer expectations. The new measure is infused to various Euler consumption models. While benchmark models based on traditional risk-free returns rates fail, models built with novel news-driven inflation expectations indices improve upon benchmark models and result in strong instruments. Our positive findings highlight the role played by the media for consumer expectation formation and allow for the use of such novel data sources for other key macroeconomic relationships.

Keywords: Euler equation, expectations, media, machine learning

JEL Classification: C26; C81; E3; E31

¹ Corresponding author: Diana Gabrielyan, gabrielyandiana1992@gmail.com, University of Tartu, School of Economics and Business Administration.

² Lenno Uusküla, lenno.uuskula@ut.ee, University of Tartu, School of Economics and Business Administration and Luminor Bank.

³ The views expressed are those of the authors and do not necessarily represent the official view of Luminor Bank or any other institutions.

We wish to thank the Editor, Amaresh Tiwari, Hakan Eratalay, Jaan Masso, seminar participants at the University of Helsinki, University of Tartu doctoral summer school 2020 participants and conference participants of ICMAIF 2022 for their useful feedback

1. INTRODUCTION

Inflation expectations play a vital role in affecting our everyday decisions. Yet, measuring inflation expectations is one of the complicated tasks in economics, as they are not directly observable. Euler's consumption model is one of the key fundamental equations in modern macroeconomics and is a key ingredient in understanding the relationship between consumer spending and the real interest rate – key to understanding consumption and decisions about saving. The relationship between inflation expectations and consumption is intuitive – an increase in the expected inflation will lead to lower interest rate (given the nominal rate), reducing savings and boosting consumption relative to the future.

In addition to their usefulness in portraying the economic relationship, the Euler equation allows to avoid explicitly solving the optimisation problem, instead focusing on specific first-order conditions of this optimisation problem. One can estimate preference parameters without having to explicitly solving the model. The baseline Euler equation consists of a forward-looking consumption component and an inversely related real interest rate. The general message is that Euler models have a problem of weak identification, poor estimates and that consumption is unresponsive to the real interest rate and is predictable on the basis of the lags of other variables.

Can machine learning and its application in economics change our opinion of Euler's consumption model? Can technological advances help fix one of the key relationships in modern macroeconomics that has been failing over and over? In this paper we address issues related to real interest rates by proposing an alternative measure of inflation expectations that is estimated directly from online news. The real interest rate is determined by the nominal interest rate and inflation expectations. While the nominal interest rate is observable, inflation expectations are not, making the expected real rate unobservable. We suggest that the media and the news they report on have the potential to capture consumer inflation expectations better than traditional survey measures.

Our contribution is therefore twofold. First, we contribute to the literature on Euler models by estimating various extensions of the baseline model and conducting empirical analysis on their performance, while evaluating the strength of instruments and values of structural parameters. Contrary to existing literature that mainly uses survey-based measures of inflation expectations or actual inflation for calculating the real interest rate (see for example Campbell and Mankiw, 1989, Coibon et al., 2020, Dräger and Nghiem, 2021) we use a novel news-based measure of inflation expectations.

In our view, the real interest rate used in the literature can be mis-specified because it does not reflect the agent's true perceptions about the economy. Consumers and households have more or different information on current and future consumption and inflation from news than an econometrician does. Therefore, when estimating the structural parameter of the model, one should not treat current and future consumption as exogenous to avoid correlated residuals and inconsistent estimates.

Compared to the previous literature the paper contributes by adding empirical estimates of the inflation expectations in the Euler equation analysis. The importance of understanding the process of the formation of household inflation expectations for monetary authorities in their attempt to influence household decisions is well documented. When the overall economic prospects are poor, this affects households' perceptions of the economy and leads agents to

defer spending. Consequently, this is reflected in their survey answers. The problem with survey data is that the information is vast and costly to obtain.

Agents receive only partial information while doing everyday shopping and build their expectations through personal experiences and memories, which however can be inaccurate, irrational and diverse. Household surveys often indicate that the perception of current inflation and expectations about the future are different from actual inflation values and differ strongly from surveys of professional forecasters and the inflation rates implied by financial markets (see for example Coibion et al., 2018). A growing body of literature provides evidence in favour of information rigidities rather than full-information rational expectations (Larsen et al., 2021; Armantier et al., 2015; Coibion and Gorodnichenko, 2012 and Doovern et al., 2015).

Differently from other papers this study adds novel empirical data on expectations in estimating the Euler equation. The importance of capturing true consumer expectations for an accurate estimation of the Euler model is also shown in Lamla and Maag (2012), where they find that households and professional forecasters have different ideas about where inflation is heading over the next 12 months. In our proposed solution, we consider the standard theoretical model of the Euler equation proposed by Hall (1988) and use extensions to the baseline model from Ascari et al. (2021) to estimate the equation. The novelty of our approach is possible thanks to technological advances that allow us to build a real-time high frequency indicator that captures true consumer inflation expectations that can be used to estimate the Euler equation.

Our second contribution is to the rapidly growing research on the impact of news on inflation expectations. The news forms a major driver of consumer sentiment and decisions and in recent years a number of studies have used news data for macroeconomic modelling. So far, the focus has largely been on variables like GDP growth, unemployment, business cycles and even cryptocurrency returns (see Thorsrud, 2018; Soric et al., 2019; Corbet et al. 2020 and Saiz et al., 2021) and less on consumption and inflation expectations (Sapiro et al., 2018 and Larsen et al., 2021).

By using novel data from online news, we also take advantage of real-time data and higher frequency. When participating in household surveys, consumers give answers on their planned consumption levels based on their sentiments at the time when the surveys are being conducted. Actual changes in consumption differ from planned levels and this difference is somewhat unpredictable. As major events unfold in the economy (e.g. Brexit vote, Covid19 pandemic, general elections), consumer sentiments and expectations change, but understandably, these changes have not been incorporated into their answers of months ago. The strong co-movement of news-based consumer sentiments and official consumer confidence survey measures for the UK gives us the necessary reassurance for using news-based measures as an alternative to survey-based inflation expectations.

Our approach starts with extracting the textual data from one of the UK's leading online newspapers from January 2000 to June 2021 and performing text selection, pre-processing and cleanup on this data to reduce the dimensionality and 'noisiness'. The resulting transformed textual data is then converted into quantitative indices that capture the intensity of the topics being discussed in the news. To finalise the construction of the novel news-based topic indices, the latter time series are augmented using sentiment indices that reflect the tone expressed by the authors of the news articles. The final indices are used as measures of inflation expectations in a Euler equation. We also incorporate various components of consumption for robustness and analyse which topics have significant impact on household consumption decisions.

The Euler models and specifications in this paper follow closely those suggested by Ascari et al. (2021). They include extensive analysis of the Euler models and the different variants, such as models including consumer habits, hand-to-mouth consumers, recursive preferences etc. They contribute to research addressing the problems of weak identification by identifying structural parameters in both linear and non-linear form. Our aim, however, is different, as we focus on model estimation and performance evaluation using a novel news-based measure of inflation.

The answer to the question of the usefulness of news and machine learning is positive. We show that when building the real interest rate using novel news-based sources of inflation expectations in a Euler model, instead of traditional risk-free returns rates, the Euler models work. While not all models yield precise estimates for the structural parameters, for non-durable goods and services consumption components, most models work and result in strong instruments and good estimates of the elasticity of intertemporal substitution. Our aim is not to find the model with the best fit or find the strongest parameters for it, or the highest and statistically significant value of elasticity of intertemporal substitution. Instead, we show that when a ‘correct’ measure of the real interest rate is infused to the Euler models, the Euler models work, while they fail when traditional real interest rate measures are used.

Key insights from our findings are as follows: not only does the goodness of fit for all three models improve when news-topic driven indices are used instead of official inflations measures, but the elasticity of intertemporal substitution (EIS) values also improve upon benchmark models, albeit with results still quite close to zero. Overall, our results provide evidence in favour of Euler models: Euler models work with news-based topic driven indices.

A topic-based classification of news yields the most important insights on the role of particular news topics on consumers sentiments and inflation perceptions and allows us to evaluate which news affects consumer spending the most. For instance, our results are particularly encouraging for news topics discussing UK, USA and Chinese economies, as well as financial markets. This is not surprising but an important contribution both as a validation of our results from the intuitive point of view, as well as a contribution allowing the use of news-based inflation expectations for macroeconomic modelling and real-time predictions.

The paper proceeds as follows: The next section reviews the literature on Euler models and why often the empirical estimations fail. Section 3 focuses on household inflation expectations and the role of the media in their formation. We also provide a detailed overview of the data. Section 4 focuses on methodologies, from the construction of a news-based inflation expectations measure to its infusion to various Euler model specifications. Section 5 provides the results and robustness analysis. Section 6 concludes.

2. WHY DO EULER MODELS FAIL?

Households maximise a standard intertemporal utility function of the form:

$$U_t = \sum_{i=0}^{\infty} \beta^i E_t[u(c_{t+i})], \quad (1)$$

where c_{t+1} is consumption at period $t + 1$, u is the instantaneous utility function, β is the subjective discount factor and the E_t consumer’s subjective expectations at time t . The utility

function exhibits Constant Relative Risk Aversion (CRRA), that is $u(C_{t+1}) = \frac{C_{t+1}^{1-\frac{1}{\sigma}}}{1-1/\sigma} - 1$, with σ being the inverse of the degree of risk aversion. Emerging from consumers' utility maximization problem are the first order conditions (2), where for the purposes of this paper, $\beta^i \left[\frac{u(C_{t+1})}{u(C_t)} \right]$ is the stochastic discount factor and the r_t is the risk-free real interest rate between t and $t+1$:

$$1 = E_t \left[\beta^i \left[\frac{u(C_{t+1})}{u(C_t)} \right] (1 + r_t) \right]. \quad (2)$$

Once log-linearized, these moment conditions become $\widehat{c}_t = E_t \Delta \widehat{c}_{t+1} - \sigma \widehat{r}_t$, which leads to the simplest case of the Euler model:

$$E_t \Delta \widehat{c}_{t+1} = \sigma \widehat{r}_t, \quad (3)$$

where $\Delta c_{t+1} = c_{t+1} - c_t$ and \widehat{c}_{t+1} and \widehat{r}_t are respectively the log deviations of \widehat{c}_t and \widehat{r}_t . σ is the marginal rate of intertemporal substitution ($\sigma \geq 0$). In this simple model, σ is also the elasticity of intertemporal substitution⁴ between current and future consumptions and tells us how the marginal rate of substitution between these consumptions reacts to the changes in interest rate. In standard theory (see for example Hall, 1988 and Hansen and Singleton, 1996) a lower real interest rate creates an incentive for consumers to spend now and reduce current savings. For riskless assets, the EIS is derived by dividing the elasticity of consumption in periods t and $t+1$ by the relative price of consumption in these two periods, with a minus sign.

Under uncertainty, EIS is computed similarly, except in the nominator the ratio of current-period consumption to the certainty equivalent of future consumption is used (see Epstein and Zin, 1991). For the constant relative risk aversion utility, the expectations of both the consumption growth rate and expected real returns on assets are used for calculating the EIS (see Hansen and Singleton, 1983).

While the standard Euler model of consumption is one of the main building blocks of many macroeconomic models, a sizeable literature, some of which is described in this section, has had difficult times estimating the model, as it often fails to hold at the aggregate level. There are several possible explanations. One is that the real interest rate used in the Euler equation is mis-specified and does not capture the true consumer inflation expectations. Imperfect information may affect the Euler equation leading to inconsistent estimates. In addition, any structural changes arising from policy shifts have weaker impact on real variables, such as consumption. Contrary the impact is stronger on nominal variables, such as future inflation expectations.

Another explanation for failing Euler models is that consumption is difficult to forecast, leading to instruments being weak. When estimating EIS, instruments should be exogenous and relevant; that is, correlated with consumption growth Δc_{t+1} . Intuitively, the estimated value of EIS has important economic implications, but most papers find no evidence of intertemporal substitution. For example, Yogo (2004) estimates EIS for 11 developed countries based on a linearised Euler equation and finds weak identification stemming from the correlation between

⁴ It is one of the most important determinants of the consumers' intertemporal consumption choices, since it measures the elasticity of the marginal substitution between consumption today versus consumption in the next period.

instruments and the dependent variable and resulting in estimates of EIS for all countries ranging from 0 to 0.5; therefore, they are too small to have a significant effect on consumption.

Yogo's paper follows a plethora of research conducted previously that yield similar results. An influential paper by Hall (1988) finds virtually no evidence for intertemporal substitution when estimating the relationship between consumption growth rate and expected real interest rates for the United States. Similarly, Campbell and Mankiw (1989) also find evidence against the permanent income hypothesis for the US when examining both non-durable and durable consumer spending. However, at the same time, the paper also challenges the robustness of Hall (1988) results when introducing the current-income consumers and arguing that the substantial fraction of income goes to rule-of-thumb consumers; therefore, Hall's theory behind the conclusions on the EIS cannot be empirically valid. As Attanasio and Weber (1993) point out, an aggregation bias may lead to such results of a low estimate of the consumption growth response to interest rates, and that as a result, the lagged consumption growth is being invalidated as an instrument. Once this is fixed, their EIS increases.

It is noteworthy that there are a number of papers that also found significant and positive values for the EIS; for example, Attanasio and Weber (1993) and Vissing-Jørgensen (2002), so there is no consensus in the literature as to the value of EIS and how significantly different from zero it should be. There is, however, a relatively wide strand of literature studying the conditions under which the structural preference parameters can be identified in the Euler equation.

Vissing-Jørgensen (2002) uses US Consumer Expenditure Survey micro data to argue that to obtain consistent estimates for EIS in the Euler model, one needs to account for limited asset market participation and that "the Euler equation should hold for a given household only if the household holds a nonzero position in the assets". What this implies, is that, if a household does not hold any assets, then including their consumption in the Euler equation will lead to inconsistent estimates for the EIS, since these agents will not be adjusting their consumption growth rate in response to any expected changes for asset returns. The study eventually finds the EIS values to be between 0.3 and 0.4 for stockholder households and between 0.8 and 1 for bondholder households.

While the households that do not hold any stocks of assets, the EIS are small and close to zero. Similar findings are reported in Attanasio et al. (2002), as well as in Gross and Souleles (2002). The latter finds that EIS is significantly positive for credit card borrowers supported by a significantly negative relationship between the credit card interest rates and the amount of borrowing. While these results are interesting, as Vissing-Jørgensen (2002) notes, one should be cautious about interpreting these results as evidence of heterogeneity in the EIS across households.

The most recent paper published on this topic by Ascari et al. (2021) summarizes the results from various baseline and extension Euler models, using both newly developed robust-to-weak-identification methods and well-established traditional methods. Their results vary depending on the choice of model (e.g. baseline or extension), as well as choice of interest rate parameter. For example, in the case of a risk-free interest rate being used in the estimation of the Euler model, the aggregate EIS is well-identified and low for several loglinear and nonlinear models but is virtually zero for the semi-structural model.

3. INFLATION EXPECTATIONS AND NEWS MEDIA

3.1. The role of the media in the process of forming inflation expectations

The media's role and power in a society is well established and, in most cases, the news provides primary sources and preferred delegates for information. An average consumer does not typically have the resources or time to constantly track the latest statistics and monitor all the events in the economy to get a full understanding of the various economic indicators. In other words, it is primarily through the media (e.g. newspapers, television, online news) that consumers receive and interpret macroeconomic information, form beliefs and opinions, as well as build sentiments about the economy and its future.

Blinder and Krueger (2004) conduct a survey on the determinants of public opinion in the US and find that television is the dominant source of information on economic policy issues, followed by newspapers. Fullone et al. (2007) support these findings through surveys in Italy and Nimark and Pitschner (2019) conclude that agents' beliefs and actions in the economy are affected by the reported information. As agents steadily move away from television and traditional newspapers to online news, more recent research papers particularly focus on examining the relationship between online news and consumer sentiments, see for example Thorsrud (2018) and Bauer (2015). The overall idea from these studies is that news is in one way or another consumed by households through various online channels, whether directly or through other media outlets.

A strand of the above-described literature is specifically focused on how media coverage affects inflation expectations. Intuitively, to some extent, media coverage reflects the current state of the economy. It is possible to understand the importance of this topic for the economy and its future based on the intensity and the extent of how much it is discussed in the news. The frequency of the news and the tone of the text can drive consumer perceptions and allow us to understand consumer inflation expectations. Carroll (2003) contributes to this literature through an analysis of two US newspapers and establishes a link between the amount of news reporting on inflation and the accuracy of consumer expectations. The main findings imply that more news leads to more rational household forecasts. Lamla and Lein (2008) investigate how the media affects inflation expectations through the intensity of the news coverage and the tone of this coverage.

Later Lamla and Maag (2012) adopt a Bayesian learning model to investigate the heterogeneity of inflation expectations and forecast disagreement between German households and professional forecasters, motivated by media reporting on inflation. They challenge Carroll's results and find that media coverage does affect the forecast disagreement and tends to increase with the heterogeneity of media coverage. However, the forecast disagreement declines with the increase of the number of reports pointing to a rise in inflation.

Similar results are reported in Pfajfar and Santoro (2013), where using Michigan Survey data, the authors show that more news coverage may widen the forecast gap between professional forecasters and consumer's mean forecast: more negative news tends to decrease the accuracy of consumer expectations, but favourable news has no statistically significant impact on them. What these results imply is that agents persistently deviate from the mean expectations of professional forecasters and the news is most likely to blame for this and causes distorted expectations.

The relationship between the news and consumer expectations is further analysed by Bauer (2015), where the sensitivity of survey expectations of inflation in response to macroeconomic news is estimated through regression-based models and points to a significant sensitivity. A recent paper by Larsen et al. (2021) analyses about 5 million news articles published over 20 years and finds that many news topics have high predictive power for inflation expectations.

Our Analysis also supports these findings. While more results and discussion will follow in the upcoming sections, Figure 1 illustrates the relationship between the UK's official Consumer Confidence Index (CCI)⁵ and the sentiment index we constructed based on the tone of news coverage (NSI).




Figure 1. UK Consumer Confidence Index (CCI) vs News-based Sentiment Index (NSI) – indices are converted to quarterly series and standardized

A similarity in shape and trend of the curves, as well as a strong visual correlation can be observed from Figure 1. The correlation of about 0.5 is displayed between the NSI and CCI, which is even further improved to 0.7 when the NSI index is shifted 1 month forward, indicating that the sentiments consumers built are reflected in the consumer surveys with slight lags. Both the CCI and sentiment indices built from the news both fell strongly around the financial crisis of 2008, then gained an upwards trend as the economy started recovering.

The confidence started dropping again around 2012 before reaching pre-crisis levels and did not drop until early 2020 when the news about the coronavirus pandemic broke. The CCI also had a slight drop in the periods leading to the Brexit vote (23 June 2016) and for some months afterwards, as expected. However, no significant drop was reflected in the news-based inflation indices. Some divergences between the series are expected and more thorough analysis would be required to identify the causes. Still, the results in Figure 1 support our hypothesis that sentiments from the news are indeed a very strong indicator of consumer expectations about the economy and have a significant impact on them. In section 4.3, we extend and apply this sentiment analysis to the level of topics in the news.

⁵ Source: <https://data.oecd.org/leadind/consumer-confidence-index-cci.htm>

3.2. News as novel data source

While it is clear that the media has a direct impact on consumer sentiments about the economy and the inflation expectation process, the empirical literature on using news-based data for modelling the economy is relatively small, albeit growing consistently. Thorsrud (2018) even adding to his title “words are the new numbers”. Indeed, in the past few years, a number of studies have been conducted and papers published where textual data is used instead of traditional survey-based data.

Applications are numerous, from financial markets to central banking and consumer sentiments. For example, Hendry and Madeley (2010) extract information from Bank of Canada communication statements and using latent semantic analysis analyse which type of information affects returns and volatility in short-term and long-term interest rate markets. El-Shagi and Jung (2015) find that the minutes of the Bank of England’s Monetary Policy Committee have contributed to market expectation formations on the future of monetary policy.

Similarly, Thorsrud (2018) uses articles from the Norwegian daily business newspaper and constructs a perfectly accurate new business cycle index that classifies the phases of the business cycle and provides meaningful insights on which types of news drive or reflect economic fluctuations. Larsen et al. (2021) use large news corpus and machine learning algorithms to investigate the role played by the media in the process of households forming expectations and conclude that certain news topics that the media reports on are good predictors of both inflation and inflation expectations.

We too, use textual data and infuse the ready data to Euler models to analyse their performance. We start by collecting two types of data: those downloaded from traditional published datasets and those manually⁶ extracted from the novel newspaper source. Traditional published datasets were collected from the Bank of England and include data on inflation, consumption and inflation attitude surveys. The consumption data itself includes total household consumption series, as well as its components, such as durable, non-durable and semi-durable goods information.

The novel newspaper data source used comes from a rich textual data environment of online news and is collected from one of the UK’s leading newspapers,⁷ the Guardian, using its open-source API.⁸ The choice of the news outlet is due to its relevance to our research in terms of content and readership. In April 2011, it was the fifth most popular newspaper in the world,⁹ while in May 2013, it was one of the most popular UK newspaper websites with 8.2 million unique visitors per month.

We argue that the news stories relevant for the formation of household expectations are most probably covered by the Guardian (or any other major newspaper for that matter) regardless of the potential skew in the coverage due to political bias or readership. King et al. (2007) performed a real-world randomized experiment to understand the causal effects of news

⁶ Here the word manually means that the data was not readily available for download. Instead, a connection to newspaper’s API is established and some coding is required to extract the data from the newspaper’s website.

⁷ See <https://www.pressgazette.co.uk/uk-newspaper-and-website-readership-2018-pamco/>, as well as <https://pamco.co.uk/pamco-data/latest-results/> for comparison among UK newspapers.

⁸ See <https://open-platform.theguardian.com>

⁹ [Guardian.co.uk most read newspaper site in UK in March.](https://www.theguardian.com/uk-news/2013/may/01/guardian-most-read-newspaper-site-in-uk-in-march)

coverage in various news outlets across the US in nationwide discussions on a range of topics and find that even the news coverage of smaller media outlets can have an impact on increasing public discussion on specific topics and that this increase was uniformly distributed across political affiliations, gender and regions of the US. Similarly, Nimark and Pitschner (2019) highlight in their paper that while different news outlets typically emphasise different topics, major events are covered in all outlets quite homogeneously.

Any news in the Guardian is public and readable by anyone by default. The Guardian API is a public web service for accessing all the content the Guardian creates, categorised by tags and section. Users can query the content database for articles with full content by tags and sections. While different news can drive consumer expectations (e.g., rumours, scandals, entertainment), we consider business section articles to be more suitable for the purposes of the analysis of this paper. Therefore, we take the articles only from the Guardian's business section for the last 21 years. We also filter out articles based on subjectively chosen keywords, such as inflation, deflation, cheaper, cheap, expensive, price, prices, cost, expense, salary, wage, salaries, wages. Arguably, this is only a subset of the news that affects household decisions, yet the main news stories relevant for household sentiment or expectation formation will be undoubtedly covered by articles that include these keywords.

The data extracted from the Guardian comes in unstructured form; that is, it is in a text format and does not have a given structure. Overall, our news corpus consists of around 23,000 English language articles with well above 20 million words in total from January 2000 to June 2021, which is enough data to conduct our analysis. However, this amount of data also makes statistical computations a challenge. We therefore apply data pre-processing steps suggested by Bholat et al. (2015), at the same time adding more steps and more developed methods. We use the text mining bag of word approach when working with textual data, which means all words are analysed as a single token and their structure, grammar or part of lexicon does not matter. Pre-processing results in a so-called document term matrix, which consists of all unique words in the corpus and their respective frequencies. At this step, the dimensionality of the corpus is reduced and we get results that have a clearer meaning. A full description of the steps used to clean up the data is given in Appendix A1. Figure A1 in the appendix visualises the most common words in the Guardian corpus.

4. CAN TEXT ANALYSIS WITH MACHINE LEARNING CHANGE OUR OPINION OF THE EULER CONSUMPTION MODEL

4.1. Modelling news into time series topics

We hypothesise that certain topics written about in the news have different degrees of impact on consumer sentiments and the process of forming expectations. This means that certain events happening in the economy could potentially have economy-wide effects. In turn, this means some topic distribution is needed for the news corpus.

Topic modelling is a branch of unsupervised natural language processing that provides a simple way to analyse large volumes of uncategorised text clustering words that frequently occur together and best explain the underlying information of a particular document. In other words,

it is the process of looking into a large collection of documents and identifying clusters of words based on similarity, patterns and multitude.

Since any document can be assigned to several topics at a time, the probability distribution across topics for each document is therefore needed. For a general introduction to topic modelling see Steyvers and Griffiths (2007) and Blei and Jordan (2003). The latter were the first to suggest Latent Dirichlet Allocation (LDA) for this purpose. LDA is a statistical model that identifies each document as a mixture of topics and attributes each word to one of the document's topics; therefore, clustering words into topics. For more information on how LDA works see Appendix A2.

Generally, in text mining, researchers do not know the topic structure of a set of documents a priori. Different model iterations and parameters result in different document clustering. However, the goal is to find unknown patterns; therefore, there is no perfect value for numbers of topics and the solution will most likely differ for different values. Hence, the choice of the number of topics to be extracted from the corpus is based on the researcher's intuition, domain knowledge and literature.

As such, we classified 80 different topics. To validate this number of topics, we follow the method by Thorsrud (2018) and compare perplexity scores across various LDA models estimated using different numbers of topics, as this allows us to inspect scores across the Markov chain Monte Carlo. The benefit of this approach comes in comparing perplexity across different models with varying topic numbers. The model with the lowest perplexity is generally considered the "best".¹⁰ Once the number of topics is chosen, the LDA procedure derives the topic probability distribution by assigning probabilities to each word and document. Table A2 in Appendix A2 presents the results from topic modelling with LDA for all 80 topics.

Figure 2 presents the frequency distribution of the 10 most probable words of a sample of topics from the LDA procedure, which discuss the future of the economy (classified by the occurrence of future-indicating words, e.g., 'expect', 'future', 'forecast' etc.). Each of the visuals in the figure represents a topic and its top 10 most frequently occurring words in the y axis. These words and corresponding frequency bars are plotted in descending order.

One characteristic of the LDA procedure is that it does not assign labels to the topics. We do that ourselves based on the two most frequent words for the given topic and based on our subjective understanding of the topics and the economy. By exploring top words within each topic that have the highest probability of belonging to that topic gives a good description of what the topic is about. The exact name, however, plays a minor role in the actual analysis or results.

¹⁰ Additionally, one can choose the number of topics that provide the best statistical decomposition by using the maximum likelihood method to find the model with the best score.

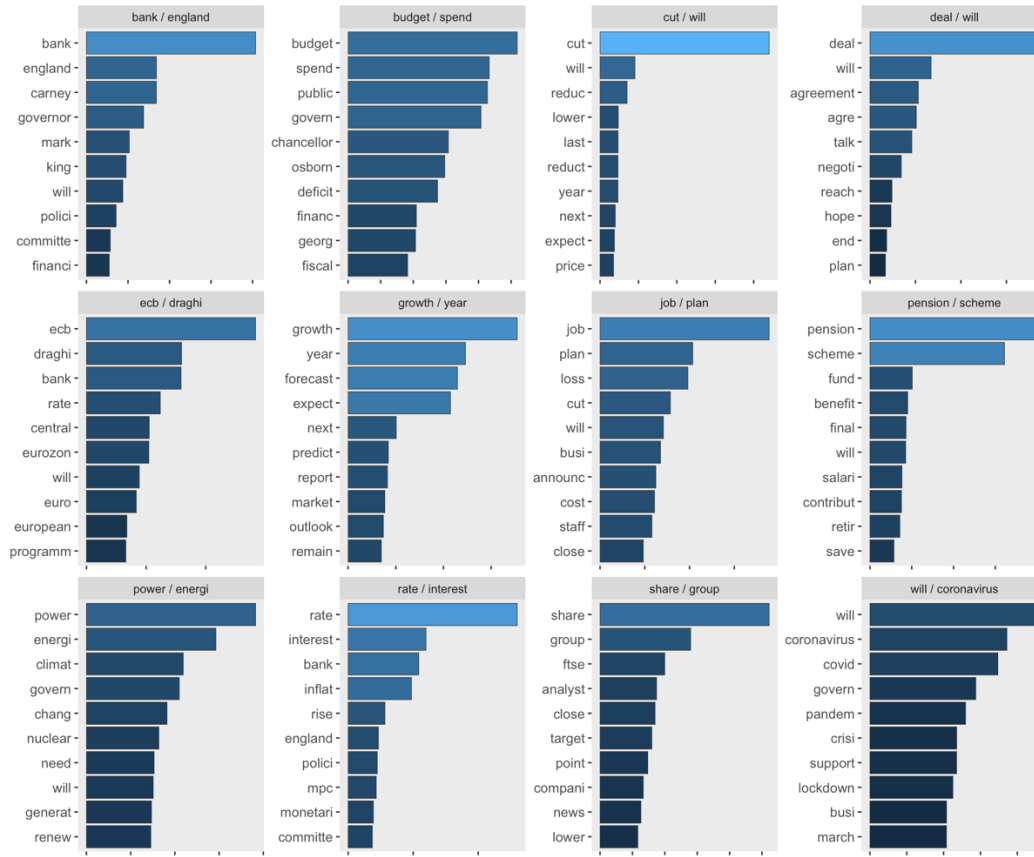


Figure 2. Sample of topics representing the future with the top 10 most frequent words – topic labels are assigned by a concatenation of the two most frequent words within the topic; all words are in stemmed format

4.2. News-topic-driven price index

To proceed with building a high-frequency news-topic-driven inflation index (*NTDI*), we calculate the frequency of each topic, or in other words, the intensity of how much each topic is discussed in the news for a given day or period. Empirically, we first sum together all articles for a given day into one document, grouping them into one plain text. Next, based on the top 20 most frequent words in each topic the article's daily frequency is calculated. The news volume $I_z(t)$ of topic z is given by:

$$I_z(t) = \sum_{d \in I(t)} \sum_w N(d, w, z), \quad (4)$$

where $N(d, w, z)$ is the frequency with which the word w tagged with topic z appears in document d . As such, we build 80 daily series for each topic using topic decompositions and distributions. Figure 3 plots the results of frequency indices built using (4) for topics representing the news on the future of the UK economy.

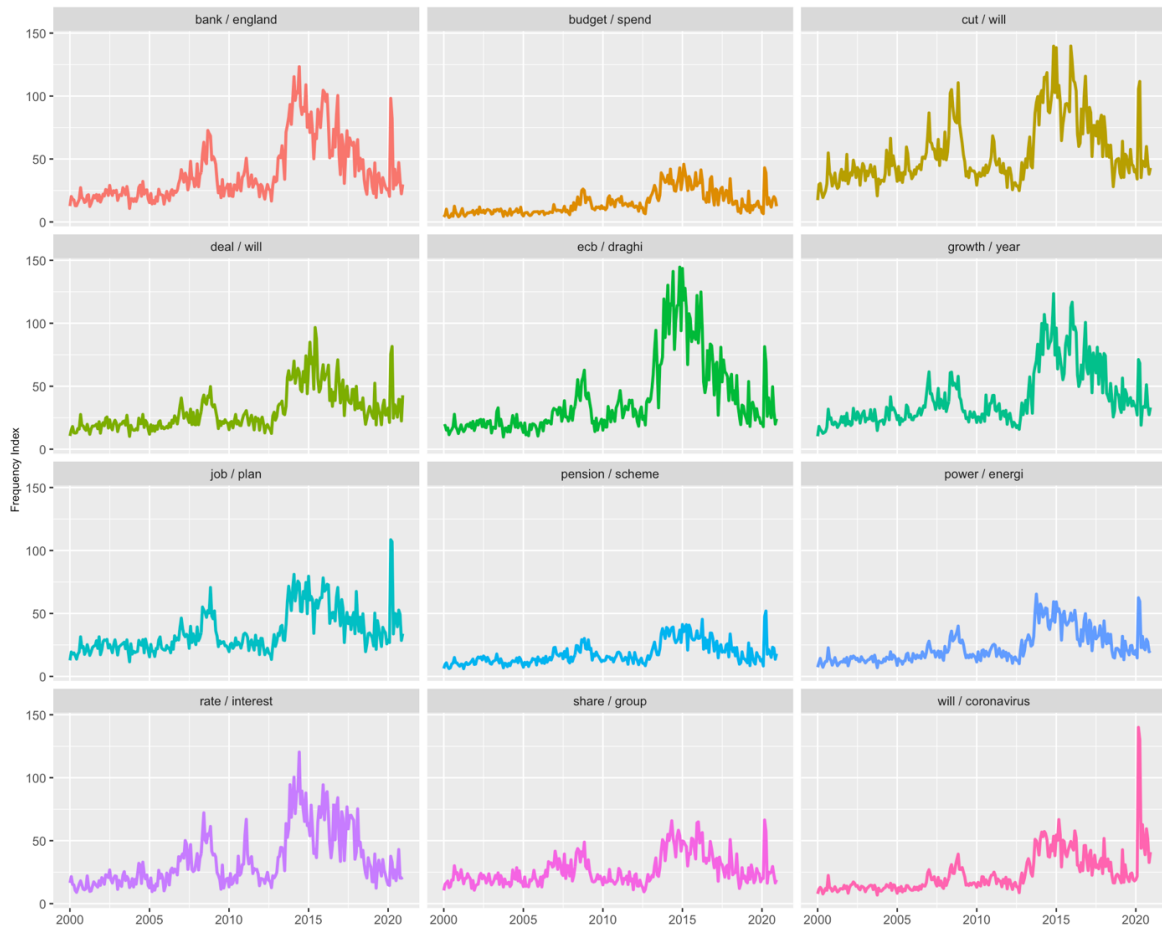


Figure 3. News-based monthly frequency indices for a subset of topics discussing the UK economy

These time series $I_z(t)$, presented in Figure 3, are measures of the volume or intensity of topics discussing the UK economy. Several recent papers find that the combination of the intensity of news topics and corresponding implied sentiment of these news and topics is important for better capturing inflation expectations (see Larsen et al., 2021 and Thorsrud, 2018). Therefore, to get the final measure NTDI, which will capture true consumer inflation expectations, we augment the intensity indices with sentiment indices. The following section describes the method of constructing sentiment indices.

4.3. Adding sentiment

Since our aim is to build the true inflation expectation of consumers, sentiment analysis and its ability to classify articles into positive, negative, or neutral sentiments, is a key step in our methodology. We start by computing thousands of sentiment values which capture the tone expressed by the authors of the news. The problem can be defined as a sentiment prediction problem, where N articles from the Guardian corpus are aggregated as a single text for each day t . What this means is that for each day between January 2000 and June 2021, all D_i s for that day are aggregated into $\{D_1 \dots D_T\}$, where D_t is the collection of documents for day t . Each of these D_i can take a sentiment value s_i , which are calculated as the difference between the frequencies of positive and negative words in the text normalized by the total number of words. As such s_i can take on any value. This approach to calculating sentiment scores is a widespread

approach in the literature as can be seen also in Larsen et al. (2021) and Arslan-Ayaydin et al. (2016).

There are several available methods and ways to conduct sentiment analysis (see Ravi and Ravi, 2015; Ardia et al., 2019; Bai, 2011; Schumacher et al., 2012), each with its own limitations and advantages and some are significant. To make our analysis more robust, we use two different methods. The first method uses a standard dictionary-based sentiment analysis approach to classify words based on their polarity (e.g., positive, negative, or neutral). We chose the Loughran-McDonald (2010) financial dictionary (354 positive words and 2355 negative words) as the most suitable ready dictionary for text analysis in the economic domain. The second method we apply is an extension of the dictionary-based classification, which also considers valence-shifting words. These are words like ‘very’, ‘barely’, ‘mustn’t’, ‘nor’, ‘not’, that may affect the context of nearby words.

Built-in packages in R language provide powerful toolchains facilitating the sentiment analysis of such textual content. We build two sentiment indices using the methods described above, but do not find significant differences in the results¹¹ and chose to proceed with the second approach in further calculations and denote the calculated sentiment index SI. Once SI time series are calculated for each topic, we multiply them with intensity indices from (4) and this concludes the calculation of news-based measures of inflation expectations:

$$\overline{NTDI}_z(t) = I_z(t) * S_z(t). \quad (5)$$

Figure 4 illustrates News-Topic Driven Inflation Indices (NTDI) constructed using (5) for a sample of topics that contain words referring to the future of the economy. The output for all 80 topics is provided in Appendix B. The quantitative values of the inflation expectations series shown in the figure are not of importance for us since the series will be standardised by the mean of consumption series before being used in the Euler model estimations. Instead, the trends and peaks for specific topics and correlations among topics are worth observation. For instance, one can note that the topic titled ‘will/cut’ (labelled based on its two most frequent words) has its lowest peak before and around 2008, corresponding to consumer inflation expectations drop related to the global financial crisis. Similarly, it seems that topics discussing the ECB and its former President Mario Draghi have caused a downward trend in the consumer expectations around 2014.

¹¹ The correlation of sentiment indices built using the second method with the official consumer confidence index is slightly higher than for sentiment indices built using the dictionary-based methods only. The difference, however, is minor.




Figure 4. News-Topic Driven Inflation indices for topics representing news covering the future of the UK economy

It is worth mentioning that more sophisticated sentiment classification techniques (see Pang et al., 2012) exist and our methods for building the inflation indices can be improved. However, as can be seen from figures 1 to 4, our news-based inflation indices provide sufficient confidence on news-based topic indices in their ability to capture the true inflation expectations.

4.4. Infusing a new measure of inflation expectations to the Euler model

The simple Euler model in (3) represents the standard log-linearised Euler equation. The common consensus about this model is that it generally does not work due to the predictability of consumption by other series' lags and its unresponsiveness to the real interest rate. The latter is generally computed as the difference between nominal interest rate i_t and inflation measure π_t . The econometric specification of (3) within the IV-GMM framework can be rewritten as:

$$\Delta c_{t+1} = \alpha + \sigma (i_t - \pi_{t+1}) + \varepsilon_{t+1}. \quad (4)$$

Equation (4) represents the baseline Euler model, where $(i_t - \pi_{t+1})$ is the ex-post real interest rate, α is an unrestricted constant and ε_{t+1} is the error term that can be serially correlated up to order 1. Δc_{t+1} is consumption in the next period. For the benchmark models, π_t is the official UK quarterly CPI inflation, while for all other models, it is replaced with news-topic driven inflation expectations, NTDI (constructed in section 4.3). By our definition, NTDI reflects the true inflation expectations of consumers.

For the results in the main paper, we focus on non-durable goods and services consumption. We present the findings on aggregate and semi-durable goods consumption in additional results. We do not estimate the Euler model for durable consumption: not only are they included in the calculations of aggregate consumption, but we also believe that durables do not respond as significantly to inflation, as does the spending on non-durable or semi-durable goods. It is worth noting, however, that the actual impact of not taking durable goods into our Euler models depends on the elasticity of substitution between durables and non-durables and whether the marginal utility of non-durables consumption is affected by the consumption of durables. We do not check for this in this paper.

The specifications of all models used in this paper are provided in Table 1 and follow closely those in Ascari et al. (2021). In addition to the baseline model, we also test two popular variants of the Euler equation that take external and internal consumer habits into consideration. Including consumer habits allows us to account for agents' response to various shocks and for the consumption path persistence in different periods. As commonly described in the literature and highlighted in Ascari et al. (2021), internal habits are those where the consumer is concerned with their current consumption relative to the consumption in the previous period. At the same time, external habits are influenced by external 'factors'¹², and therefore, consumer's current consumption is affected by the aggregate consumption in the previous period instead of its own consumption in the same period, which is the case with internal habits. External habits are characterized by the introduction of the lagged term Δc_t , which affects the forward-looking nature of the Euler equation and the relationship between EIS and aggregate consumption.

Table 1. Models estimated in the paper

Baseline (I)	$\Delta c_{t+1} = \alpha + \sigma (i_t - \pi_{t+1}) + \varepsilon_{t+1}$
Model with External Habits (II)	$\Delta c_{t+1} = \alpha + \gamma \Delta c_t + \sigma(1 - \gamma)(i_t - \pi_{t+1}) + \varepsilon_{t+1}$
Model with Internal Habits (III)	$\Delta c_{t+1} = \alpha + \frac{\gamma}{1 + \gamma^2 \beta} (\Delta c_t + \beta E_t \Delta c_{t+2}) + \sigma \frac{(1 - \gamma)(1 - \gamma \beta)}{1 + \gamma^2 \beta} (i_t - \pi_{t+1}) + \varepsilon_{t+1}$

While there are somewhat contrasting results in the literature on the impact of habits of various consumption components and their importance in explaining the key patterns in household consumption decisions. Still, there seems to be consensus that habits capture consumption durability, and therefore should be accounted for, especially in the presence of time-varying risk premia. Attanasio and Weber (2010) discuss habits and different ways they can be modelled, while more technical details on utility functions describing habit formation can be found in Campbell and Cochrane (1999).

The official data used in the models comes from a variety of sources and undergoes a number of transformations to make all data aligned with each other and in a format comparable to other studies. Detailed descriptions of transformations to the official data is given in Table C1 of Appendix C.

¹² Commonly in the Euler literature, this type of habit formation is described as the 'keeping up with the Joneses' effect, which essentially means that consumers reduce their savings by increasing their consumption to keep up with the level of consumption in their peer group.

News-topic driven indices are in a quarterly format and include a great deal of time-variation due to the nature of the data. To account for high frequency variation and ensure accurate analysis, we standardise the NTDI by means of official inflation series.¹³

To support our empirical analysis, we evaluate the value and sign of EIS, the significance of regressors' coefficients and apply popular robust-to-weak-identification tests to ensure the validity of the instruments. The restriction $E_{t-1}\varepsilon_{t+1} = 0$ is imposed on the models, so that only variables that are determined at time t can be used as instruments. There is an endogeneity problem in our models, which, however, can be addressed using an IV regression approach and taking lagged endogenous variables as instruments. As such our instrument set consists of three lags of both Δc_t and $(i_{t-1} - \pi_t)$, as well as a constant. News driven sentiment indices $S_z(t)$ are also added as additional instruments based on the assumption that they are not only correlated with consumption (correlation coefficient is -0.3), but they would also allow more use of the novel data and potentially have positive impact on the models. This assumption, however, does not find evidential support in our results, as we do not find any significant improvement in model performance when the sentiments are added or not as instruments.

Common tests in the literature for handling weak instruments are used to check the validity, relevance and exogeneity of our instruments. The first of these is the first-stage F-statistic that tests for weakness of instruments and identifies if the instrument has low correlation with the exploratory variable. Next, we use the Hausman test for checking the consistency of OLS estimates under the assumption that IV is consistent. Rejecting the null hypothesis would mean that there is endogeneity present. Lastly, we use Stock and Watson's test of instrument exogeneity using overidentifying restrictions (known also as the J-test or Sargan test), which is only applicable to models where the number of instruments is more than the number of endogenous regressors. Some of our models use only sentiment indices as instruments, hence this test is not applied to them. Overall, if all three tests hold, then we consider the instrument valid and non-weak.¹⁴

5. RESULTS

5.1. Main results

Results presented in Table 2 and 3 show the output of the Euler models with the consumption of non-durable goods and services (NDGS) for benchmark models and models based on NTDI respectively. The results for the consumption of aggregate and semi-durable goods are provided in Appendix D. For benchmark models official CPI inflation data is used for π_t and we compare these to the models infused by news-based topic driven indices as inflation expectation measures. As can be observed from the diagnostic tests in Table 2, Euler models do not seem to work well for NDGS consumption, when using official inflation. The instruments pass the weakness test but fail on the other two. EIS is 0 or negative, indicating that its value is not informative. Therefore, the conclusion at this stage would be that Euler models for UK non-

¹³ To validate our results, we also standardize NTDI using alternative measures of inflation, such as inflation attitudes survey data or 5-year implied inflation rates.

¹⁴ There are other tests available for checking the instruments relevance, such as Kleibergen (2002)'s statistical test for instrument validity and Moreira (2003, 2009)'s coefficient test designed to test coefficients in the structural equation regardless of the strength of identification.

durable goods and services consumption fail. So do the models for the other types of consumption, as is implied from the results in tables D1 and D2 in the appendices.

Table 2. IV regression results from benchmark models: non-durable goods and services

	Benchmark Models		
	(I)	(II)	(III)
$(i_t - \pi_{t+1})$	-0.00005 (0.006)	0.001 (0.006)	0.007 (0.006)
Δc_t		-0.138 (0.274)	-0.290 (0.300)
$E_t \Delta c_{t+2}$			-0.571 (0.618)
Constant	0.004 (0.004)	0.004 (0.004)	0.004 (0.004)
Observations	81	81	80
EIS	0	-0.138	-1.466
R^2	0.0003	0.080	0.215
	Diagnostic tests		
Weak Instruments test	35.424 ***	35.424 *** 0.002 **	18.857 *** 0.002 ** 0.000 ***
Wu-Hausman	0.405 (0.526)	1.214 (0.302)	1.524 (0.215)
Sargan	3.002 (0.809)	2.999 (0.700)	0.581 (0.965)

Notes: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$, all models include 6 instruments: Δc_{t-1} , Δc_{t-2} , Δc_{t-3} , Δr_{t-1} , Δr_{t-2} , Δr_{t-3} , c_t corresponds to non-durable goods and services consumption.

Findings for the Euler models with news-based topic indices are presented in Table 3. Only models where all instruments are valid and an improvement upon benchmark models is observed are shown. An improvement upon the benchmark model is considered as such, when there is an increase in model performance in terms of R^2 (holding the number of variables constant). For the purposes of saving space in the main part of this paper, we will not include individual topic results from the models. Instead, we will cluster topics into larger groups and present the results as a range of minimum and maximum EIS and R^2 for the topics of that cluster.

The logic for clustering is simple: topics discussing any kind of price information are clustered under the group INFLATION. These include topics 6 and 55 (inflation rate), topic 21 (price increase), topics 22 and 25 (house and property prices), topics 24 (energy prices), topic 56 (gold prices) and so on. Similarly, topics containing words about the economy, such as GDP, recession, economic growth, bank rates, unemployment are grouped under either UK or WORLD ECONOMY clusters, depending on the exact keywords. Topics that discuss financial

markets, investments and stocks are grouped under the FINANCIAL MARKETS cluster and so on.

Table 3. IV regressions results from benchmark models: non-durable goods and services

<i>Model</i>	<i>Topic Cluster</i>	R^2	R^2 benchmark	<i>EIS</i>	<i>Improvement</i>
Baseline	INFLATION	0.01 - 0.01	0.0003	0.002 - 0.006	+1 pp
Baseline	OTHER	0.01 - 0.02	0.0003	0.003 - 0.005	+2 pp
Baseline	UK ECONOMY	0.01 - 0.01	0.0003	0.002 - 0.007	+1 pp
Baseline	WORLD ECONOMY	0.01 - 0.01	0.0003	0.006 - 0.007	+1 pp
External Habits	FINANCIAL MARKETS	0.09 - 0.11	0.08	0.001 - 0.006	+3 pp
External Habits	INFLATION	0.09 - 0.11	0.08	0.002 - 0.003	+3 pp
External Habits	OTHER	0.09 - 0.13	0.08	0 - 0.005	+5 pp
External Habits	UK ECONOMY	0.09 - 0.11	0.08	0 - 0.003	+3 pp
External Habits	WORLD ECONOMY	0.12 - 0.12	0.08	0.007 - 0.007	+4 pp
Internal Habits	FINANCIAL MARKETS	0.22 - 0.26	0.21	-0.001 - 0.01	+5 pp
Internal Habits	INFLATION	0.22 - 0.26	0.21	0.001 - 0.013	+5 pp
Internal Habits	OTHER	0.22 - 0.28	0.21	0.003 - 0.01	+7 pp
Internal Habits	UK ECONOMY	0.22 - 0.26	0.21	0.001 - 0.015	+5 pp
Internal Habits	WORLD ECONOMY	0.22 - 0.26	0.21	0.006 - 0.018	+5 pp

For all 3 models, there are around 2–3 topics that we label as outliers, as they do not contain any valuable words and are thus excluded from the results and discussion. All other remaining and potentially insightful topics are clustered into the group OTHER. Detailed results for all three models and individual results for each topic can be found in Appendix D2, tables D3–D5.

Summarizing the findings from Table 3, the Euler models based on the novel data source outperform the benchmark models as highlighted by the value of R^2 . This means, that the data from Table 3 can better explain and fit the Euler consumption models. It is to be noted, however, that while for baseline models the improvement in R^2 is only 1 or 2 percentage points and both R^2 and *EIS* are quite low, in absolute terms the baseline model R^2 using NTDI has improved by more than 30 times, from 0.0003 to 0.01. Similar low R^2 results hold for the other two models with habits but with higher improvement, reaching up to 7 percentage points for the

internal habits model and 5 percentage points for the external habits model. The internal habits model seems to work better with the data than the external habits model. Overall, the goodness of fit and EIS values of all three models improves when NTDI are used instead of official inflation measures. In addition, all instruments are now valid and satisfy our test conditions, as opposed to the benchmark models.

Analysing the results on the topic of non-durable goods and services consumption, interestingly, out of 80 topics, 90% of those selected contain words such as inflation, price increase, economic growth, recession, or losses. The Bank of England, expectedly, comes up often in the UK Economy cluster. Foreign economies, such as China and the USA also have an impact on household consumption decisions as per our results, which is not surprising: with China and the USA being among major world economies, any major news shocks from these countries is expected to have a direct impact on the UK economy.

Topics that include words like retail sales, Australian economy, bonus pays, drug companies, sports, airline flights and other random topics were not selected by the models, indicating that news on these topics are irrelevant or have low impact on consumer spending decisions. A few of the topics that were excluded do contain useful words such as oil prices, job cuts and incomes and this is expected. First, model performance can always be improved by using better models and tools. Second, each topic generally contains around 1,000 words, while in the Euler models we use about the 100 most frequent words, without taking into account that the remaining 900 words might outweigh the top 100 by total quantity.

The topics selected by our algorithm represent news shocks affecting consumer consumption decisions and consequently the relationships between macroeconomic variables in the Euler models. Since we provide evidence that NTDI based Euler models work, this shows that consumers who read the news will build their inflation expectations and adjust their non-durables spending according to the relationship represented by the Euler model framework. Depending on the model, the relationship will be more or less accurate. For instance, from Table 3 and the UK Economy topic, it follows that all three models are able to accurately model consumer non-durable consumption patterns through the Euler models, but EIS is the highest for the external habits model, implying a stronger relationship. This means that when reading the news on the UK economy, households take into the account the spending of others in their own spending and consumption decisions.

For other components of consumption, the results are not as good. The improvement for aggregate consumption is up to 3 points at best, while for the semi-durable goods consumption component, only the baseline model with a full instrument set gives reliable results. The remaining results cannot be considered good, as either R^2 or EIS are negative. The sensitivity of the Euler models' results on the consumption data we observe is common in the literature. For different types of consumption, the same models can yield contrasting results, which is intuitive. Numerous survey-data-based studies on different countries reach the same conclusion, that changes in inflation affect various components of consumption differently both by quantity and the direction of the change (see Dräger and Nghiem, 2021 for the results using German household data and Coibon et al., 2019 for results on Dutch households). It is therefore not surprising, and as expected, that the Euler models would not perform alike for all consumption types.

5.2. Robustness checks and model limitations

Our results, including the value of EIS are almost insensitive to using 3-month interbank rates or end-of-quarter official bank rates.¹⁵ Neither are they particularly sensitive to instrument set combinations. However, they are sensitive to the type of consumption and the number of lags in instruments. For reference, the results from Ascari et al. (2021) on US data are also insensitive to different instruments and specifications. However, in contrast to ours, their results do not change when a different consumption measure is used but instead are highly sensitive to asset returns: the value of EIS changes depending on whether risk-free returns or stock market returns are used. In particular, with stock market returns, EIS is significantly positive, but not precisely estimated. A similar observation about the sensitivity and precision of EIS can also be drawn from our results: it is not always precisely estimated and its value varies a lot when the model and underlying data change.

The value of EIS across all our models is generally low and close to zero. These results are supported in the linear models of Ascari et al. (2021) and Campbell (2003), which find the 95% confidence interval for the estimated EIS to be close to 0 for non-durable consumption. Yogo (2004) robust to weak-identification econometric methods also yield a small EIS that is not significantly different from zero.

There are numerous ways to improve model performance to support our results with even stronger empirical evidence. For example, there are a number of robust-to-weak identification tests for parameter stability or structural change that could be applied to our models. Similarly, there are potentially more efficient methods for further evaluating the sensitivity of our results to different heteroskedasticity and autocorrelation consistent estimates and the variance of moment conditions.

Attanasio and Low (2000) also suggest loglinear approximating Euler models and using a sample long enough to get ‘well-behaved’ estimates. While we believe our sample size is adequate, we are unable to extend the data length due to the limitations of the availability of online news data. It would also be interesting to deep dive into consumer data and better understand the demographics of UK households and the specific characteristics of news readers, such as age, income, wealth, education etc. However, we will leave this for further research.

6. CONCLUSIONS

There is a lot of discussion of Euler models and their potential failures in the literature. We argue that one reason why many authors have reached the conclusion that they fail is because the real interest rate in the model is mis-specified and fails to capture consumers’ true perception of the economy. To tackle this, we propose news as an alternative and more accurate source for capturing inflation expectations.

Our hypothesis is that the news that consumers read has a direct impact on their expectations and recent technological advances allow us to derive these expectations directly from the news using machine learning techniques. Even though we do not solve all the problems related to Euler models, our results are empirically successful. We provide evidence in favour of Euler models when news-based inflation expectations are used to calculate the real interest rate.

¹⁵ Results are not presented in the paper for the purposes of saving space but are available on request.

Using online news data for consumption modelling and predictions is relatively unexplored. Estimating Euler models with news-based inflation expectation measures opens numerous opportunities for macroeconomists to make further progress not only in modelling, but also predicting consumption in real-time. Our positive findings also allow for the use of such novel data sources for other key macroeconomic relationships, for example, the New Keynesian Phillips curve.

REFERENCES

- Ardia, D., Bluteau, K. and Boudt, K. (2019): "Questioning the news about economic growth: Sparse forecasting using thousands of news-based sentiment values," *International Journal of Forecasting* 35(4), pp 1370-1386.
- Armantier, O., Bruine de Bruin, W., Topa, G., Klaauw, W. and Zafar, B. (2015): "Inflation Expectations and Behavior: Do Survey Respondents Act on their Beliefs?", *International Economic Review* 56(2), pp 505–536.
- Arslan-Ayaydin, Ö., Boudt, K. and Thewissen, J. (2016): "Managers set the tone: Equity incentives and the tone of earnings press releases", *Journal of Banking & Finance* 72(S), pp 132-147.
- Ascari, G., Magnusson, L. and Mavroeidis, S. (2021): "Empirical evidence on the Euler equation for consumption in the US", *Journal of Monetary Economics* 117(C), pp 129-152.
- Attanasio, O., Banks, J. and Tanner, S. (2002): "Asset Holding and Consumption Volatility", *Journal of Political Economy* 110 (4), pp 771–92.
- Attanasio, O. and Weber, G. (2010): "Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy", *Journal of Economic Literature* 48 (3), pp 693-751.
- Attanasio, O. and Weber, W. (1993): "Consumption Growth, the Interest Rate and Aggregation", *Review of Economic Studies* 60(3), pp 631–49.
- Attanasio, O. and Low, H. (2004): "Estimating Euler equations", *Review of Economic Dynamics* 7(2), pp 406-435.
- Bauer, M. (2015): "Inflation Expectations and the News", *International Journal of Central Banking* 11(2), pp 1-40
- Blei, D. M., Ng, A. Y. and Jordan, M. I. (2003): "Latent Dirichlet Allocation", *The Journal Of Machine Learning Research* 3, pp 993–1022.
- Blinder, A. S. and Krueger, A. B. (2004): "What Does the Public Know about Economic Policy and How Does It Know It?", *Brookings Papers on Economic Activity, Economic Studies Program, The Brookings Institution* 35(1), pp 327–397.
- Bholat, D., Hansen, S., Santos, P. and Schonhardt-Bailey, Ch. (2015): "Text Mining for Central Banks".
- Breeden, D. T. (1979): "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities", *Journal of Financial Economics* 7, pp 265-296.
- Campbell, John Y. (2003): "Consumption-based asset pricing", *Handbook of the Economics of Finance*, edition 1, volume 1, chapter 13, pp 803-887.
- Campbell, J. Y. and Cochrane, J. H. (1999): "By force of habit: a consumption-based explanation of aggregate stock market behavior", *Journal of Political Economy* 107(2), pp 205-251.
- Campbell, J.Y. and Mankiw, G.N. (1989): "Consumption, income and interest rates: Reinterpreting the time series evidence", *NBER Macroeconomics Annual*, MIT Press, pp. 185–216.
- Carroll, Ch. D. (2003): "Macroeconomic Expectations of Households and Professional Forecasters", *Quarterly Journal of Economics* 118, pp 269–98.
- Coibion, O., Gorodnichenko, Y., Kumar, S. and Pedemonte, M. (2020): "Inflation expectations as a policy tool?", *Journal of International Economics* 124 (C).
- Coibion, O. and Gorodnichenko, Y. (2012): "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy* 120(1), pp 116–59.
- Corbet, Sh., Larkin, Ch., Lucey, B. M., Meegan, A. and Yarovaya, L. (2020): "The impact of macroeconomic news on Bitcoin returns", *The European Journal of Finance* 26(14), pp 1396-1416

- Curtin, R. (2007): “Consumer Sentiment Surveys: Worldwide Review and Assessment,” *Journal of Business Cycle Measurement and Analysis*, OECD Publishing, Centre for International Research on Economic Tendency Surveys (1), pp 7-42.
- Dovern, J., Fritsche, U., Loungani, P. and Tamirisa, N. (2015): “Information rigidities: Comparing average and individual forecasts for a large international panel”, *International Journal of Forecasting* 31(1), pp 144-154.
- Dräger L. and Nghiem G. (2021): “Are Consumers' Spending Decisions in Line with A Euler Equation?”, *The Review of Economics and Statistics* 103 (3), pp 580–596.
- Epstein, L. G. and Stanley, E. Z. (1991): “Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis.” *Journal of Political Economy*, 99(2), pp 263–286.
- Fullone, Flora, Michela Gamba, Enrico Giovannini, and Marco Malgarini. (2007) What Do Citizens Know About Statistics? The Results of an OECD/ISAE Survey on Italian Consumers.
- Fullone, Flora, Michela Gamba, Enrico Giovannini, and Marco Malgarini. (2007) What Do Citizens Know About Statistics? The Results of an OECD/ISAE Survey on Italian Consumers.
- Fullone, Flora, Michela Gamba, Enrico Giovannini, and Marco Malgarini. (2007) What Do Citizens Know About Statistics? The Results of an OECD/ISAE Survey on Italian Consumers.
- Fullone, Flora, Michela Gamba, Enrico Giovannini, and Marco Malgarini. (2007) What Do Citizens Know About Statistics? The Results of an OECD/ISAE Survey on Italian Consumers.
- Fullone, Flora, Michela Gamba, Enrico Giovannini, and Marco Malgarini. (2007) What Do Citizens Know About Statistics? The Results of an OECD/ISAE Survey on Italian Consumers.
- Fullone, Flora, Michela Gamba, Enrico Giovannini, and Marco Malgarini. (2007) What Do Citizens Know About Statistics? The Results of an OECD/ISAE Survey on Italian Consumers.
- Fullone, Flora, Michela Gamba, Enrico Giovannini, and Marco Malgarini. (2007) What Do Citizens Know About Statistics? The Results of an OECD/ISAE Survey on Italian Consumers.
- Fullone, F., Gamba, M., Giovannini, E. and Malgarini, M. (2007): “What Do Citizens Know About Statistics”, The results of OECD/ISAE Survey on Italian Consumers.
- Fuhrer, J.C. and Moore, G.R. (1995): “Monetary policy trade-offs and the correlation between nominal interest rates and real output”, *American Economic Review* 85, pp 219–239.
- Gross, D. B. and Nicholas S. S. (2002): “Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data”, *The Quarterly Journal of Economics* 117(1), pp 149–85.
- Hansen, L. P. and Singleton, K. J. (1982): “Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models”, *Econometrica*, *Econometric Society* 50(5), pp 1269–1286.
- Hansen, L. P. and Singleton, K. J. (1996): “Efficient Estimation of Linear Asset Pricing Models with Moving Average Errors”, *Journal of Business and Economic Statistics* 14(1), pp 53-68.
- Jung, A. and El-Shagi, M. (2015): “Has the publication of minutes helped markets to predict the monetary policy decisions of the Bank of England's MPC?”, Working Paper Series 1808, European Central Bank.
- King, G., Schreer, B. and White, A. (2017): “How the News Media Activate Public Expression and Influence National Agendas”, *Science* 358, 6364, pp 776–780.
- Lamla, M. J. and Lein, S.M (2008): “The Role of Media for Consumers' Inflation Expectation Formation”, KOF Working papers 08-201, KOF Swiss Economic Institute, ETH Zurich.
- Lamla, M. J. and Maag, Th. (2012): “The Role of Media for Inflation Forecast Disagreement of Households and Professional Forecasters”, *Journal of Money, Credit and Banking*, Blackwell Publishing 44(7), pp 1325-1350.
- Larsen, V. H., Thorsrud, L. A. and Zhulanova, J. (2021): “News-driven inflation expectations and information rigidities”, *Journal of Monetary Economics* 117, pp 507-520.
- Loughran, T. and McDonald, B. (2010): “When is a Liability not a Liability? Textual Analysis, Dictionaries and 10-Ks”, *Journal of Finance* 66 (1).

- Lucas, R. E. Jr. (1976): “Econometric Policy Evaluation: A Critique”, *Carnegie-Rochester Conference Series on Public Policy* 1(1), pp 19-46.
- Moreira, M. J. (2009): “Tests with Correct Size When Instruments Can Be Arbitrarily Weak”, *Journal of Econometrics* 152 (2), pp 131-140.
- Moreira, M. J. (2003): “A Conditional Likelihood Ratio Test for Structural Models”, *Econometrica* 71(4), pp 1027–1048.
- Neely, Ch. J., Amlan, R. and Whiteman, Ch. (2001): “Risk Aversion versus Intertemporal Substitution: A Case Study of Identification Failure in the Intertemporal Consumption Capital Asset Pricing Model”, *Journal of Business and Economic Statistics* 19(4),
- Nelson, C. and Startz, R. (1990): “Some further results on the exact small sample properties of the instrumental variables estimator”, *Econometrica* 58, pp 967–976.
- Nimark, K. P. and Pitschner, S. (2019): “News media and delegated information choice”, *Journal of Economic Theory* 181(C), pp 160-196.
- Pang, B., Lee, L. and Vaithyanathan, S. (2002): “Thumbs up?: Sentiment classification using machine learning techniques”, *Proceedings of the ACL-02 Conference on Empirical Methods in Natural Language Processing - Volume 10, EMNLP '02*, Stroudsburg, PA, US
- Pfajfar, D. and Emiliano, S. (2013): “News on Inflation and the Epidemiology of Inflation Expectations.” *Journal of Money, Credit and Banking* 45(6), pp 1045–67.
- Ravi, K. and Ravi, V. (2015): “A survey on opinion mining and sentiment analysis: Tasks, approaches and applications”, *Knowledge-Based Systems* 89, pp 14-46.
- Saiz, L., Ashwin, J. and Kalamara, E. (2021): “Nowcasting euro area GDP with news sentiment: a tale of two crises”, *European Central Bank Working Paper Series* 2616.
- Schumaker, R. and Chen H. (2009): “Textual Analysis of Stock Market Prediction Using Breaking Financial News”, *Association for Computing Machinery Transactions on Information Systems* 27(2).
- Sorić, P., Lolić, I., Claveria, O., Monte, E. and Torra, S. (2019): “Unemployment expectations: A socio-demographic analysis of the effect of news”, *Labour Economics* 60(C), pp 64-74.
- Steyvers, M. and Griffiths, T. (2007): “Probabilistic topic models”, *Handbook of latent semantic analysis*, pp 427–448.
- Stock, J.H., Wright, J. and Yogo, M. (2002): “A survey of weak instruments and weak identification in generalized method of moments”, *Journal of Business and Economic Statistics* 20, pp 518–529.
- Stock, J. H. and Yogo, M. (2003): “Testing for Weak Instruments in Linear IV Regression”, *Harvard University working paper*.
- Shapiro, A. H., Sudhof, M. and Wilson, D. J. (2017): “Measuring News Sentiment”. *Working Paper Series* 2017-1, Federal Reserve Bank of San Francisco.
- Thorsrud, L.A. (2018), “Words are the New Numbers: A Newsy Coincident Index of the Business Cycle”, *Journal of Business & Economic Statistics* 38(2), pp 393-409.
- Vissing-Jorgensen, A. (2002): “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution”, *Journal of Political Economy* 110(4), pp 825-853.
- Yogo, M. (2004): “Estimating the Elasticity of Intertemporal Substitution When Instruments Are Weak”, *The Review of Economics and Statistics* 86(3), pp 797-810.

APPENDIX A. TEXTUAL DATA MINING

A1. Data preparation

All words are analysed as a single token using Natural Language Processing's bag of word (BOW) approach, which means their grammar or structure does not matter. This is a common, if not the most popular, approach applied in the literature (see Thorsrud, 2018; Thorsrud, 2020). Below are the techniques used to clean up the data, which include the most common steps of the BOW approach. However, we extended this approach by also stemming the words. Each of these techniques has its own pros and cons. For example, along with reducing dimensionality, these techniques might obscure meaning for some words or might count words that are written similarly but have different meanings.

Step 1: We remove any metadata such as images, links and any other data in an unknown format contained in the articles and convert any information contained in the article into an appropriate format. Duplication and empty entries should also be accounted for and such documents are removed. This can be done either manually or using methods similar to Echkely (2015). In our analysis, we used R language's powerful commands for duplicate and empty data removal.

Step 2: We then use tokenisation, which is a step which splits longer strings of text into smaller tokens, such as words, numbers, symbols and so on. Tokenisation is usually done by using blank spaces or punctuation marks as delimiters. Tokenisation is sometimes also referred to as lexical analysis. This breakdown process results exclusively in words.

Step 3: Next, all words are normalised; that is, all the words are converted into lower case, punctuation is removed, numbers are converted into their equivalent. This is an important step, otherwise same words, such as Rate and rate, which are written in upper and lowercase respectively will be interpreted as different words. The downside is, however, that when written in uppercase, some words may refer to names of people or places, such as White and white. We assume, however, that the frequency of such words is not significant.

Step 4: A crucial step is removing stop words, otherwise they will appear in the frequently used words and will not give an incorrect picture of the core meaning of the document. Stop words are those words which are filtered out before further processing of text, since these words contribute little to the overall meaning, given that they are generally the most common words in a language. The list of these words is provided in the beginning of the analysis and includes common words in the English language that do not contain any information relating to the article. Examples of such words are the, like, can, I, also, are, in, on, this, that, gmt, pm etc.

Step 5: For further dimensionality reduction and better pre-processing results, we stem words, which involves cutting off affixes and suffixes and reducing all words to their respective word stems. This is a form of linguistic normalization, where the part of speech of each word is identified and each word is converted into its base form; for example, nouns, verbs, pronouns with the same base into base words (e.g. reporting, reported and reporter will be reduced to report).

Step 6: The last step in the pre-processing involves defining the document term matrix (DTM) based on the now clean text and computing the most common words across all the documents. Document Term Matrix (DTM) lists all occurrences of words in the corpus, by document. At this stage, we also remove the sparse terms; that is, terms occurring only in very few documents.

These are the tokens which are missing from more than 90% of the documents in the corpus.¹⁶ The remaining 900,000 stems with the highest TDM score are used in the final analysis.

The visualisation summarising the results described above is given in a word cloud form in Figure A1. Word cloud visualises the most common words in the corpus by differentiating between word colour and size, indicating the frequency intervals by colour and size, with more frequent words having a bigger size.



Figure A1. Word cloud representation of document-term-matrix.

A2. Topic modelling

Latent Dirichlet allocation (LDA) is an approach used in topic modelling based on probabilistic vectors of words, which indicate their relevance to the text corpus. LDA makes it possible to derive the topic probability distribution by assigning probabilities to each word and document. Assigning words and documents to multiple topics also has the advantage of semantic flexibility (e.g. the word ‘rate’ can relate both to inflation and unemployment topics). As Thorstrud (2018) notes, LDA shares many features with Gaussian factor models, with the difference being that factors here are topics and are fed through a multinomial likelihood at the observation.

In LDA, each document is given a probability distribution and for each word in each document, a topic assignment is made. The joint distribution of topic mixture θ , a set of N words w is given by:

$$p(\theta, z, w | \alpha, \beta) = p(\theta | \alpha) * \prod_{n=1}^N p(z_n | \theta) * p(w_n | z_n, \beta), \quad (1)$$

¹⁶ Maximal allowed sparsity is in the range from 0 to 1. For this paper, the sparsity was chosen equal to 0.9, which means the token must appear in at least 10% of the documents to be retained. The sparsity value can be modified to higher or lower value, but that affects the number of terms remained in the corpus.

where parameters α and β are k -vectors with components greater than zero, with k being the dimensionality of the Dirichlet distribution; that is, the directionality of topic variable z . In addition, the topic distribution of each document is distributed as $\theta \sim \text{Dirichlet}(\alpha)$. Term distribution is modelled using $z_n \sim \text{Dirichlet}(\beta)$ and $N \sim \text{Poisson}(\xi)$.

The goal of the LDA model is therefore to estimate θ and φ in order to estimate which words are important for which topic and which topics are important for a given document. For α and β , the higher they are, the more likely each document will contain a mixture of most topics instead of a single topic and the more likely each topic will contain a mixture of most of the words and not just single words. More technical detail and thorough specifications on the LDA model and topic modelling in general are provided in Blei (2003) and Griffiths and Steyvers (2004).

There are different approaches to the LDA algorithm. In this paper, we use the Gibbs sampling method, an algorithm for successively sampling conditional distributions of variables, whose distribution over states converges to the true distribution in the long run. Gibbs sampling makes it possible to improve the topic representations within documents, as well as the word distributions of all the topics. Gibbs method samples from this multinomial posterior distribution on the set of possible subset choices to identify those with higher probability by their more frequent appearance in the Gibbs sample (George and McCulloch, 1993). Each variable from formula (1) is sampled given the full conditional distribution of other variables, which are as follows:

$$p(z_{iv} = k \mid \pi_i, b_k) \propto \exp(\log \pi_{ik} + \log b_{k,y_{iv}}), \quad (2)$$

$$p(\pi_i \mid z_{iv} = k, b_k) = \text{Dirichlet}(\alpha + \sum_k \mathbb{1}(z_{iv} = k)), \quad (3)$$

$$p(b_k \mid z_{iv} = k, \pi_i) = \text{Dirichlet}(\beta + \sum_i \sum_k \mathbb{1}(y_{il} = w, z_{il} = k)), \quad (4)$$

where k is the topic, w is a term, π_i is a vector defining a distribution over T topics and b_k is a vector defining a distribution over N words.

Griffiths and Steyvers (2004) were the first to suggest analytically integrating π_i and b_k and sample z_{iv} to get a better performance by perhaps adding something – better performance of what? Predictive performance? The logic is as follows: for each document d , for each word w , reassign a new topic k to w . The probability of this topic k is equal to the probability of word w given topic k multiplied by the probability of topic k given document d . The mathematical formula is given below:

$$p(z_i = j \mid z_{-i}, w_i, d_i) = \frac{C_{w_{ij}}^{NT} + \beta}{\sum_{w=1}^N C_{w_{ij}}^{NT} + W\beta} \times \frac{C_{d_{ij}}^{DT} + \alpha}{\sum_{t=1}^T C_{d_{it}}^{DT} + T\alpha}, \quad (5)$$

where C^{NT} is a word-topic matrix and C^{DT} is a document-topic matrix, while α and β are parameters that set the topic distribution for the documents and the words respectively.

Different model iterations and different parameters of α and β in (1) result in different document clustering. However, the goal is to find unknown patterns; therefore, there is no perfect value for numbers of topics and the solution will most likely differ for different values. Hence, the choice of the number of topics to be extracted from the corpus is based on the researcher's intuition, domain knowledge and the literature.



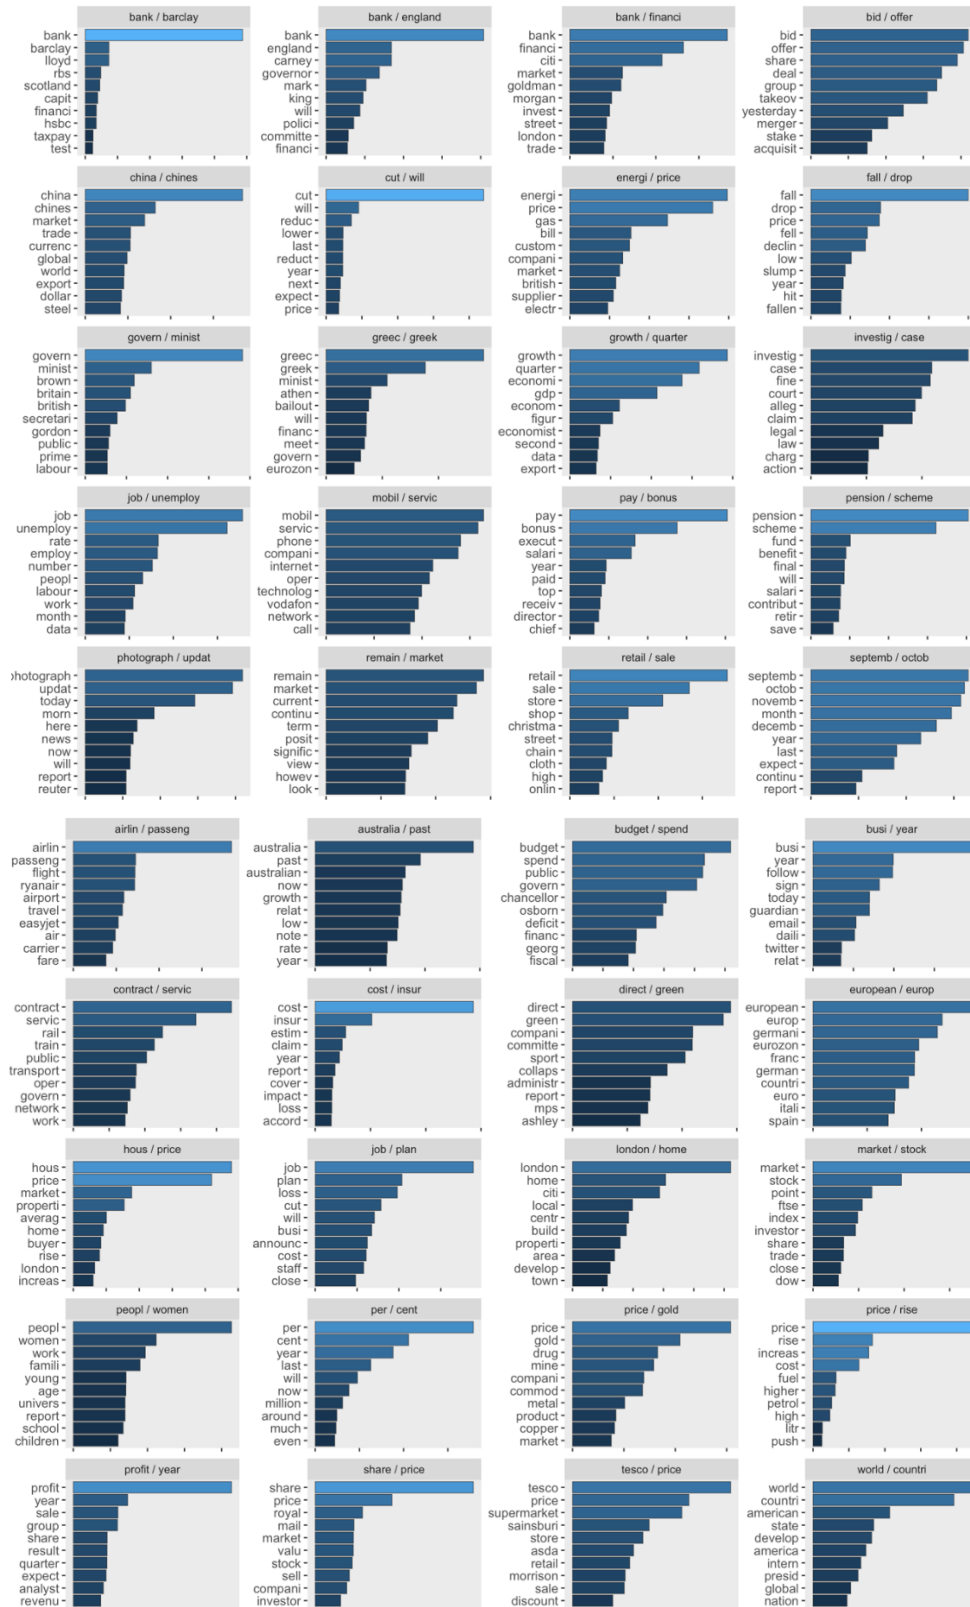
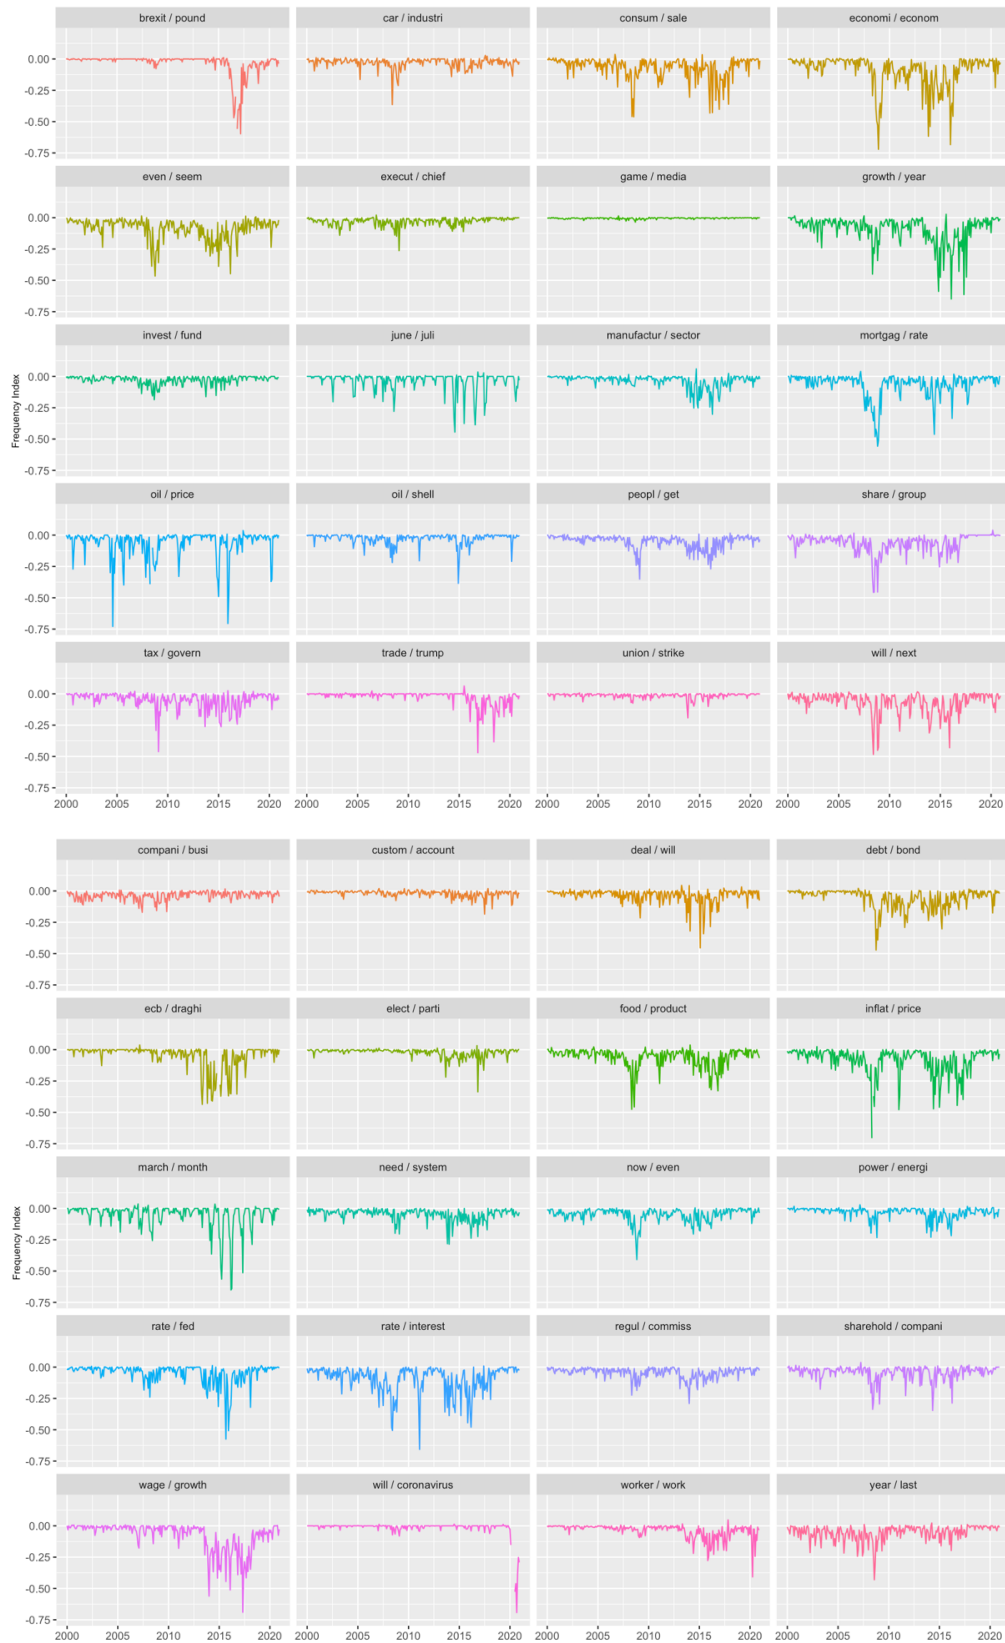


Figure A2. 80 topics resulting from LDA with top 10 frequent words in them. Topic labels are assigned by a concatenation of two most frequent words within the topic. All words are in stemmed format.

APPENDIX B. NEWS-BASED TOPIC DRIVEN INFLATION EXPECTATIONS



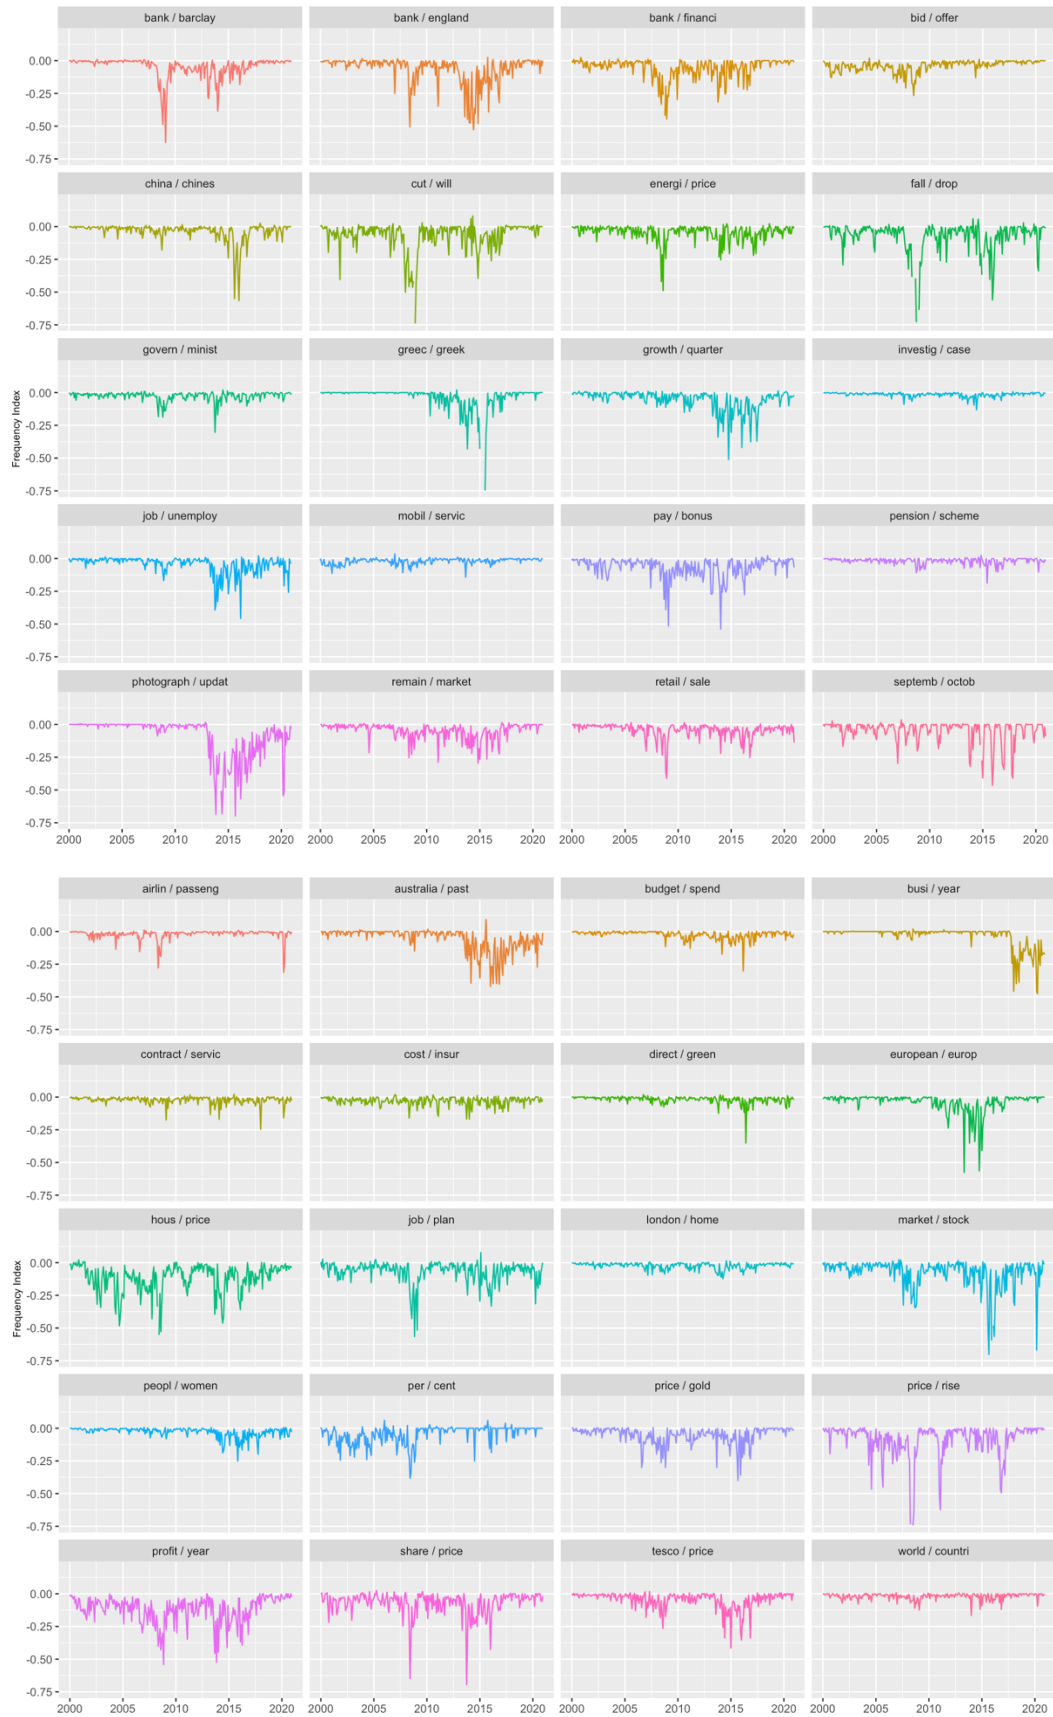


Figure B1. News-based topic driver indices

APPENDIX C. DATA DESCRIPTION

Table C1 lists the series used in the Euler consumption model. We use quarterly data covering the sample period between January 2000 and June 2021 and apply seasonal adjustment and growth transformation where needed. Like (Ascari, Magnussen and Mavroeidis, 2021), we use per head measures of consumption. Household final consumption, HFCE_PH is available both as a total and per capita measure and is directly exported from the source, while for the consumption components we manually transform the series to per head measures using population time series POP from the UK Labour Force Survey. We also transform nominal measures of consumption, HFCE, NDE, NDG, SERV to real measures using corresponding implied deflators. The formulae are as follows:

For final consumption expenditure:

$$RHFCE = \frac{HFCE_PH}{HFCE_DEFL}. \quad (1)$$

For non-durable goods and services, we combine the two components into one real measure as follows:

$$RNDGSERV = \frac{NDG_PH + SERV_PH}{P_{NDGSERV}}. \quad (2)$$

Implicit deflator in (2) for nondurable goods and services $P_{NDGSERV}$ is computed as follows:

$$P_{NDGSERV} = \frac{NDG+SERV}{\frac{NDG}{NDG_DEFL} + \frac{SERV}{SERV_DEFL}}. \quad (3)$$

And the per capita measures for consumption components are computed using:

$$NDG_{PH} = \frac{NDG}{POP}, \quad (4)$$

$$SERV_{PH} = \frac{SERV}{POP}. \quad (5)$$

As a last step for consumption related data, the per head measures RHFCE and RNDGSERV are log transformed to be used as consumption proxies in the Euler model.

For the interest rate we use 3-Month Interbank Rates for the United Kingdom for the main results and End of Quarter Official Bank Rates for robustness checks in additional results (see Appendix D). Both time series are monthly and are consequently converted to quarterly series by averaging.

To compare our results with the benchmark model we use official inflation data from the Bank of England, as well as 5-year Inflation Implied Forward and Inflation attitude surveys data in additional results for comparison.

Overall, at the final stage all data described in Table C1 is in quarterly values at annual rates, seasonally adjusted. Some are also log transformed.

Table C1. Data used in the paper

Mnemonic in the dataset	Description	Transformation	Source
<i>INFL_ATT</i>	Median value from survey indicating public attitudes to inflation and general expectation on inflation change over the next 12 months	Seasonal adjustment	Office for National Statistics
<i>CPI_INFL</i>	Official CPI inflation	Seasonal adjustment	Bank of England
<i>INFL_EXP</i>	Quarter average of yield from British Government Securities, 5-year Inflation Implied Forward	Seasonal adjustment	Bank of England
<i>HFCE_DEFL</i>	Real Household final consumption expenditure: Implied deflator	Seasonal adjustment	Office for National Statistics
<i>HFCE</i>	Household final consumption expenditure at current prices	-	Office for National Statistics
<i>HFCE_PH</i>	Household final consumption expenditure per head at current prices	-	Office for National Statistics
<i>IB_3M</i>	3-Month Interbank Rates for the United Kingdom, growth rates	Seasonal adjustment	Federal Reserve Economic Data
<i>IB_EQR</i>	End of Quarter Official Bank Rate	Seasonal Adjustment	Bank of England
<i>NDG</i>	Nominal non-durable goods expenditure at current prices		Office for National Statistics
<i>NDG_DEFL</i>	Implied Deflators for Nondurable goods, 2008 Index	Seasonal adjustment	Office for National Statistics
<i>POP</i>	LFS: Population aged 16+: UK: All: 4 quarter average	Seasonal adjustment	Labor Force Survey, ONS
<i>SERV</i>	Nominal services expenditure at current prices	-	Office for National Statistics
<i>SERV_DEFL</i>	Implied Deflators for Services, 2008 Index	Seasonal adjustment	Office for National Statistics
Calculated variables			
<i>SERV_PH</i>	Nominal services expenditure per head		
<i>NDE_PH</i>	Nominal non-durable goods expenditure per head		
<i>RHFCE</i>	Real Household final consumption expenditure per head	Log	
<i>RNGSERV</i>	Real Household non-durables and services consumption expenditure per head	Log	

APPENDIX D. RESULTS

D1. Results from benchmark Euler models

Table D1. IV regression results from benchmark models / total consumption

	Benchmark Models		
	(I)	(II)	(III)
$(i_t - \pi_{t+1})$	0.002 (0.007)	0.005 (0.007)	0.012* (0.007)
Δc_t		-0.206 (0.265)	-0.366 (0.278)
$E_t \Delta c_{t+2}$			-0.512 (0.339)
Constant	0.002 (0.004)	0.002 (0.004)	0.001 (0.004)
Observations	81	81	80
EIS	0.002	-0.207	-1.12
R ²	-0.003	0.114	0.280
	Diagnostic tests		
Weak Instruments	22.56 ***	322.56 / 6.46 *** / ***	20.37 / 2.92 / 27.96 *** / * / ***
Wu-Hausman	0.159 (0.691)	0.582 (0.561)	0.982 (0.406)
Sargan	6.726 (0.455)	5.857 (0.32)	0.585 (0.964)

Notes: *p<0.1; **p<0.05; ***p<0.01, All models include 6 instruments: Δc_{t-1} , Δc_{t-2} , Δc_{t-3} , Δr_{t-1} , Δr_{t-2} , Δr_{t-3} , c_t corresponds to total consumption.

Table D2. IV regression results from benchmark models / semi durable goods consumption

	Benchmark Models		
	(I)	(II)	(III)
$(i_t - \pi_{t+1})$	0.005 (0.007)	0.005 (0.008)	0.002 (0.009)
Δc_t		-0.022 (0.344)	0.441 (0.422)
$E_t \Delta c_{t+2}$			-0.097 (0.310)
Constant	0.002 (0.004)	0.002 (0.004)	0.001 (0.004)
Observations	81	81	80
EIS	0.005	-0.022	-3.289
R ²	0.044	0.056	-0.345
Diagnostic tests			
Weak Instruments	3.698 (0.003) **	3.698 / 0.742 (0.003) ** / (0.618)	4.399 / 0.581 / 7.788 *** / / ***
Wu-Hausman	1.180 (0.280)	0.148 (0.862)	0.894 (0.448)
Sargan	16.121 (0.013) *	16.325 (0.006) **	1.320 (0.858)

Notes: *p<0.1; **p<0.05; ***p<0.01, All models include 6 instruments: Δc_{t-1} , Δc_{t-2} , Δc_{t-3} , Δr_{t-1} , Δr_{t-2} , Δr_{t-3} , c_t corresponds to total consumption.

D2. Results from NTDI based Euler models

Table D3. IV regression results from NTDI models, total consumption

Topic	Model	Instruments	EIS	R ²	R ² benchmark	Improvement
6	Internal Habits	All	-0.00	0.28	0.28	+0 pp
8	Internal Habits	All	-0.01	0.3	0.28	+2 pp
21	Internal Habits	All	-0.00	0.28	0.28	+0 pp
43	External Habits	Sentiment index only	-0.33	0.13	0.11	+2 pp
60	External Habits	Sentiment index only	-0.31	0.14	0.11	+3 pp

Table D4. IV regression results from NTDI models, non-durable goods and services consumption

<i>Topic</i>	<i>Model</i>	<i>Instruments</i>	<i>EIS</i>	<i>R²</i>	<i>R² benchmark</i>	<i>Improvement</i>
1	External Habits	All	0.00	0.12	0.08	+4 pp
3	External Habits	All	0.00	0.13	0.08	+5 pp
3	Internal Habits	All	0.00	0.24	0.21	+3 pp
6	Internal Habits	All	0.00	0.23	0.21	+2 pp
7	External Habits	All	0.00	0.1	0.08	+2 pp
9	Internal Habits	All	0.01	0.25	0.21	+4 pp
11	Baseline	All	0.00	0.01	0.00	+1 pp
11	Baseline	Sentiment index only	0.00	0.01	0.00	+1 pp
11	Baseline	All excl. sentiment index	0.00	0.01	0.00	+1 pp
11	Internal Habits	All	0.01	0.26	0.21	+5 pp
11	Internal Habits	All excl. sentiment	0.01	0.26	0.21	+5 pp
13	Baseline	All	0.01	0.01	0.00	+1 pp
13	Baseline	All excl. sentiment index	0.01	0.01	0.00	+1 pp
13	Internal Habits	All	0.01	0.26	0.21	+5 pp
13	Internal Habits	All excl. sentiment	0.01	0.26	0.21	+5 pp
16	Internal Habits	All excl. sentiment	0.00	0.24	0.21	+3 pp
17	Internal Habits	All	0.01	0.24	0.21	+3 pp
17	Internal Habits	All excl. sentiment	0.01	0.24	0.21	+3 pp
18	External Habits	All	0.00	0.09	0.08	+1 pp
20	Internal Habits	All	0.02	0.24	0.21	+3 pp
20	Internal Habits	All excl. sentiment	0.02	0.23	0.21	+2 pp
21	Internal Habits	All	0.00	0.25	0.21	+4 pp
21	Internal Habits	All excl. sentiment	0.00	0.23	0.21	+2 pp
22	Baseline	All	0.00	0.01	0.00	+1 pp
22	Baseline	Sentiment index only	0.00	0.01	0.00	+1 pp
22	Baseline	All excl. sentiment index	0.00	0.01	0.00	+1 pp
22	Internal Habits	All	0.01	0.24	0.21	+3 pp
23	Internal Habits	All	0.01	0.24	0.21	+3 pp
23	Internal Habits	All excl. sentiment	0.01	0.24	0.21	+3 pp
24	Internal Habits	All	0.00	0.24	0.21	+3 pp
25	Baseline	All excl. sentiment index	0.01	0.01	0.00	+1 pp
25	Internal Habits	All	0.01	0.22	0.21	+1 pp
27	Internal Habits	All	0.00	0.26	0.21	+5 pp
27	Internal Habits	All excl. sentiment	0.00	0.26	0.21	+5 pp
28	External Habits	All	0.00	0.09	0.08	+1 pp
31	Internal Habits	All	0.00	0.25	0.21	+4 pp
31	Internal Habits	All excl. sentiment	0.00	0.22	0.21	+1 pp
32	Baseline	All	0.00	0.01	0.00	+1 pp
32	Baseline	All excl. sentiment index	0.00	0.01	0.00	+1 pp
32	External Habits	All	0.00	0.09	0.08	+1 pp
33	Internal Habits	All	0.01	0.24	0.21	+3 pp
33	Internal Habits	All excl. sentiment	0.01	0.24	0.21	+3 pp
34	External Habits	All	0.00	0.09	0.08	+1 pp
37	External Habits	All	0.01	0.12	0.08	+4 pp
37	Internal Habits	All	0.01	0.25	0.21	+4 pp
39	Internal Habits	All	0.01	0.24	0.21	+3 pp
39	Internal Habits	All excl. sentiment	0.01	0.23	0.21	+2 pp
40	Internal Habits	All	0.01	0.24	0.21	+3 pp
41	External Habits	All	0.00	0.1	0.08	+2 pp

42	External Habits	All	-0.00	0.11	0.08	+3 pp
42	Internal Habits	All	0.00	0.25	0.21	+4 pp
43	External Habits	All	0.00	0.13	0.08	+5 pp
43	Internal Habits	All	0.01	0.25	0.21	+4 pp
44	Internal Habits	All excl. sentiment	0.00	0.25	0.21	+4 pp
46	External Habits	All	0.00	0.1	0.08	+2 pp
48	Baseline	All	0.00	0.01	0.00	+1 pp
48	Baseline	All excl. sentiment index	0.00	0.01	0.00	+1 pp
48	External Habits	All	0.00	0.1	0.08	+2 pp
48	Internal Habits	All	0.01	0.26	0.21	+5 pp
50	Internal Habits	All	-0.00	0.24	0.21	+3 pp
50	Internal Habits	All excl. sentiment	0.00	0.22	0.21	+1 pp
51	Baseline	All	0.01	0.01	0.00	+1 pp
51	Baseline	All excl. sentiment index	0.01	0.01	0.00	+1 pp
51	Internal Habits	All	0.01	0.24	0.21	+3 pp
51	Internal Habits	All excl. sentiment	0.01	0.22	0.21	+1 pp
55	Baseline	All	0.00	0.01	0.00	+1 pp
55	Baseline	Sentiment index only	0.00	0.01	0.00	+1 pp
55	Baseline	All excl. sentiment index	0.00	0.01	0.00	+1 pp
55	Internal Habits	All	0.01	0.24	0.21	+3 pp
55	Internal Habits	All excl. sentiment	0.00	0.23	0.21	+2 pp
56	Internal Habits	All	0.00	0.24	0.21	+3 pp
57	Internal Habits	All	0.00	0.25	0.21	+4 pp
57	Internal Habits	All excl. sentiment	0.00	0.25	0.21	+4 pp
59	Baseline	All	0.00	0.01	0.00	+1 pp
59	Baseline	Sentiment index only	0.01	0.01	0.00	+1 pp
59	Baseline	All excl. sentiment index	0.00	0.01	0.00	+1 pp
59	Internal Habits	All	0.01	0.26	0.21	+5 pp
60	External Habits	All	0.01	0.09	0.08	+1 pp
61	External Habits	All	0.00	0.11	0.08	+3 pp
62	Baseline	All	0.01	0.01	0.00	+1 pp
62	Baseline	All excl. sentiment index	0.01	0.01	0.00	+1 pp
62	Internal Habits	All	0.01	0.26	0.21	+5 pp
62	Internal Habits	All excl. sentiment	0.01	0.26	0.21	+5 pp
64	Baseline	All	0.01	0.01	0.00	+1 pp
64	Baseline	All excl. sentiment index	0.01	0.01	0.00	+1 pp
64	Internal Habits	All excl. sentiment	0.01	0.22	0.21	+1 pp
66	Baseline	All	0.00	0.01	0.00	+1 pp
66	Baseline	All excl. sentiment index	0.00	0.01	0.00	+1 pp
66	External Habits	All	0.00	0.09	0.08	+1 pp
68	External Habits	All	0.01	0.09	0.08	+1 pp
69	Baseline	All	0.01	0.01	0.00	+1 pp
69	Baseline	Sentiment index only	0.01	0.01	0.00	+1 pp
69	Baseline	All excl. sentiment index	0.01	0.01	0.00	+1 pp
69	Internal Habits	All	0.01	0.22	0.21	+1 pp
70	Baseline	All	0.00	0.01	0.00	+1 pp
70	Baseline	All excl. sentiment index	0.01	0.01	0.00	+1 pp
70	Internal Habits	All	0.01	0.28	0.21	+7 pp
70	Internal Habits	All excl. sentiment	0.01	0.26	0.21	+5 pp
72	External Habits	All	0.00	0.11	0.08	+3 pp
72	Internal Habits	All	0.00	0.26	0.21	+5 pp
72	Internal Habits	All excl. sentiment	0.00	0.26	0.21	+5 pp
74	Baseline	All	0.01	0.01	0.00	+1 pp

74	Baseline	All excl. sentiment index	0.01	0.01	0.00	+1 pp
74	Internal Habits	All	0.01	0.25	0.21	+4 pp
74	Internal Habits	All excl. sentiment	0.01	0.25	0.21	+4 pp
76	Internal Habits	All	0.00	0.27	0.21	+6 pp
77	Baseline	All	0.00	0.01	0.00	+1 pp
77	Baseline	All excl. sentiment index	0.00	0.01	0.00	+1 pp
77	External Habits	All	0.00	0.09	0.08	+1 pp
78	Baseline	All	0.00	0.02	0.00	+2 pp
78	Baseline	All excl. sentiment index	0.00	0.02	0.00	+2 pp
78	Internal Habits	All	0.01	0.26	0.21	+5 pp
78	Internal Habits	All excl. sentiment	0.01	0.26	0.21	+5 pp
79	Internal Habits	All excl. sentiment	0.02	0.22	0.21	+1 pp
80	External Habits	All	0.00	0.1	0.08	+2 pp

Table D5. *IV regressions results from NTDI models / semi-durable goods consumption*

<i>Topic</i>	<i>Model</i>	<i>Instruments</i>	<i>EIS</i>	R^2	R^2 benchmark	<i>Improvement</i>
22	Internal Habits	All	-0.02	0.02	-0.34	+0.02 pp
22	Internal Habits	All excl. sentiment	-0.01	-0.02	-0.34	-
32	Internal Habits	All	-0.03	-0.12	-0.34	-
32	Internal Habits	All excl. sentiment	-0.01	-0.21	-0.34	-
55	Internal Habits	All	-0.04	-0.08	-0.34	-
55	Internal Habits	All excl. sentiment	-0.03	-0.06	-0.34	-
65	External Habits	Sentiment index only	-0.17	0.06	0.06	0 pp

KOKKUVÕTE

Inflatsiooniootused ja tarbimine masinõppega

Käesolev artikkel analüüsib inflatsiooniootuseid kasutades uudiseid ning kasutab seda tarbijate oodatava reaalse intressimäära lähendina. Reaalse intressimäära ootus on võti mõistmaks tarbimist ja seepärast kasutatakse seda erinevate Euleri võrrandite hindamisel. Uued inflatsiooniootustel põhinevad reaalsed intressimäärad parandavad märgatavalt Euleri võrrandi hinnanguid ja parandavad hinnangutel kasutatud instrumente. Tulemused näitavad Euleri võrrandi hindamisel inflatsiooniootuste mõõtmise olulisust ning meedia rolli majapidamiste tarbimisotsuste tegemisel.