ARTICLE IN PRESS
JID: MONEC [m3Gsc;October 16, 2020;11:58]

Journal of Monetary Economics xxx (xxxx) xxx

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jmoneco

News-driven inflation expectations and information rigidities

Vegard H. Larsen[a], Leif Anders Thorsrud[b,*], Julia Zhulanova[b]

[a] Norges Bank and Centre for Applied Macroeconomics and Commodity Prices (CAMP), BI Norwegian Business School, Norway
[b] Centre for Applied Macroeconomics and Commodity Prices, BI Norwegian Business School, Norway

ARTICLE INFO

Article history:
Received 3 July 2019
Revised 11 March 2020
Accepted 12 March 2020
Available online xxx

Keywords:
Expectations
Media
Machine learning
Inflation

ABSTRACT

Using a large news corpus and machine learning algorithms we investigate the role played by the media in the expectations formation process of households, and conclude that the news topics media report on are good predictors of both inflation and inflation expectations. In turn, in a noisy information model, augmented with a simple media channel, we document that the time series features of relevant topics help explain time-varying information rigidity among households. As such, we provide a novel estimate of state-dependent information rigidities and present new evidence highlighting the role of the media in understanding inflation expectations and information rigidities.

© 2020 The Authors. Published by Elsevier B.V.
This is an open access article under the CC BY license
(http://creativecommons.org/licenses/by/4.0/)

1. Introduction

The fourth estate, i.e., the news media, plays an important role in society and is a primary source of information for most people.[1] In macroeconomics, expectations are center stage. But, expectations are shaped by information, and information does not travel unaffected through the ether. Rather, it is digested, filtered, and colored by the media. Surprisingly, however, the potential independent role of the media in the expectation formation process has received relatively little attention in macroeconomics.

This paper builds on a growing literature providing evidence in favor of information rigidities rather than full-information rational expectations (FIRE; Armantier et al., 2016; Coibion and Gorodnichenko, 2012; Coibion and Gorodnichenko, 2015; Dovern et al., 2015), and investigates the relationship between news and households' inflation expectations in such settings.

In particular, we take the view that agents make endogenous information choices (Mackowiak and Wiederholt, 2009; Sims, 2003; Woodford, 2009), but that no agent has the resources to monitor all the events potentially relevant for her decision, and thus delegate their information choice to specialized news providers, who report only a curated selection of events. As formalized in Nimark and Pitschner (2019), the media act as "information intermediaries" between agents and the state of the world.[2] Two implications of these views are that: i) media coverage should predict households' inflation expectations, and ii) the degree of information rigidity, as defined more precisely below, will be time-varying and a function of media coverage.

* Corresponding author. Department of Economics, BI Norwegian Business School, Nydalsveien 37, 0484 Oslo, Norway.
E-mail addresses: vegard-hoghaug.larsen@norges-bank.no (V.H. Larsen), leif.a.thorsrud@bi.no (L.A. Thorsrud), julia.zhulanova@bi.no (J. Zhulanova).

[1] See, for example, Blinder and Krueger (2004), Curtin (2007), and Fullone et al. (2007).

[2] Rather than agents deciding ex-ante on the expected usefulness of a particular signal, as in, e.g., the costly information literature (Grossman and Stiglitz, 1980), knowledge of events is jointly determined ex-post through a delegated information choice mechanism.

https://doi.org/10.1016/j.jmoneco.2020.03.004
0304-3932/© 2020 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license
(http://creativecommons.org/licenses/by/4.0/)

Please cite this article as: V.H. Larsen, L.A. Thorsrud and J. Zhulanova, News-driven inflation expectations and information rigidities, Journal of Monetary Economics, https://doi.org/10.1016/j.jmoneco.2020.03.004

These implications are tested in two stages. First, the predictive relationship between news and expectations is addressed. To this end, we hypothesize that when the media writes extensively about topics related to, e.g., technology or health, even without explicitly mentioning terms related to inflation, this reflects that something is happening in these areas that could potentially have economy-wide effects and might therefore also affect inflation expectations. In turn, this conjecture is made operational using a Latent Dirichlet Allocation (Blei et al., 2003) model and a large news corpus from the *Dow Jones Newswires Archive* (DJ) to construct 80 time series measures of the news topics the media write about, i.e., the different types of news reporting.

Using penalized linear regressions to handle the high-dimensional predictive problem, and focusing on inflation and households' inflation expectations, measured by the University of Michigan Surveys of Consumers (MSC), we find that many of the news topics written about in the media have high predictive power for both inflation and expectations. There is also a large intersection in the selected topic sets for these two outcome variables, and our results imply that relevant news coverage helps households form more accurate expectations. Furthermore, the narrative realism of the approach is good. Topics about, e.g., (IT) technology and health, significantly affect households' inflation expectations. Additional results strongly indicate that this type of textual data contains information not captured by a large set of roughly 130 conventional economic indicators, suggesting that the media is an important information source for households. In contrast, but following the intuition that the media matters foremost for households and less so for professionals, there is little evidence for a relationship between news topics and inflation expectations from the Survey of Professional Forecasters (SPF).

The MSC micro-data is used to further validate the news-topic-based approach, and shows that the predictive relationship between news topics and expectations align well with conventional stereotypes and what we know about expenditure patterns and media consumption habits. News related to health and politics, for example, tends to be more important for elderly survey respondents than for young people.

Turning to information rigidities, we augment the noisy information framework in Coibion and Gorodnichenko (2015) by allowing for state dependence in the degree of information rigidity, and an explicit, but simple, role for the media. The mechanics of the model are straight forward. When an important event happens, media coverage potentially becomes more persistent, and the signal less noisy, and thereby easier to filter for the agents. Accordingly, information rigidity is reduced as agents put more weight on new information relative to their previous forecasts.

Testing these predictions empirically supports the media channel view. There is high-frequency time variation in information rigidity, and this variation can be explained by the time series properties of relevant news coverage, as the theory predicts. We further show, in a falsification experiment, that this result is unlikely to be obtained by chance, and using properties of inflation itself, or other economic indicators with predictive power for households' expectations, does not deliver theory-consistent results.

The contribution of our analysis is threefold. First, by analyzing media's role in the expectation formation process, our analysis speaks directly to work by Doms and Morin (2004), Pfajfar and Santoro (2013), Lamla and Lein (2014), Dräger and Lamla (2017), and Ehrmann et al. (2017). The epidemiological model of inflation expectations by Carroll (2003) is particularly well known. However, we make an important contribution in how we use text as data in this setting. In contrast to the earlier literature, where analyses have been based on counting inflation terms in the news to measure media (inflation) intensity or survey variables measuring whether people have heard news about prices, we take a topic-based approach. And, indeed, this approach delivers results in accordance with our media mechanism, while the traditional text- and survey-based methods do not.

Second, we are the first to investigate the relationship between information rigidities and news within a well-established theory-based testing framework (Coibion and Gorodnichenko, 2015). This allows us to directly test the null of FIRE versus the alternative news-driven information rigidity view.

Third, we provide direct evidence of high-frequency time-variation in the degree of information rigidity among households in the U.S. As such, our results complement Loungani et al. (2013), Coibion and Gorodnichenko (2015), and Dovern et al. (2015), who document low-frequency changes in information rigidity among professionals and in international panels.

In sum, the analysis conducted here provides positive evidence in favor of the state-dependent information rigidity view, but emphasizes the role of information providers. For this reason, the analysis also speaks to the literature trying to identify the causal effect of the media. This has been relatively unexplored in macroeconomics, but has received much more attention in other branches of the literature and in other sciences (Gentzkow et al., 2011; King et al., 2017; Prat, 2018; Shiller, 2017).

2. Expectations and news

To study the relationship between expectations and news, Section 2.1 describes the news corpus and how the textual data is transformed into quantitative time series. The predictive results are presented in Section 2.2.

Please cite this article as: V.H. Larsen, L.A. Thorsrud and J. Zhulanova, News-driven inflation expectations and information rigidities, Journal of Monetary Economics, https://doi.org/10.1016/j.jmoneco.2020.03.004

2.1. The news

Our news media corpus consists of roughly five million news articles, written in English, from the *Dow Jones Newswires Archive* (DJ), covering the period 1990 to 2016. The database covers a large range of Dow Jones' news services, including content from *The Wall Street Journal*.

Arguable, the DJ includes only a subset of news households consume. Still, news stories relevant for inflation are undoubtedly well covered by this type of business news. The *Dow Jones* company, and its flagship publication, *The Wall Street Journal*, is also one of the largest newspapers in the U.S. in terms of circulation. This means that it has a large footprint in the U.S. media landscape, and it is likely that its news coverage spills over to news sources that households follow more directly, e.g., television, or smaller news outlets (King et al., 2017). While minor news events might not be covered by this data source, major economic or political events are surely covered by both DJ and other media outlets households might follow.

To make our news-topic-based hypothesis operational, we use a Latent Dirichlet Allocation (LDA) model (Blei et al., 2003), where each article is treated as a mixture of topics and each topic is treated as a mixture of words. The LDA is one of the most popular topic models in the Natural Language Processing (NLP) literature because of its simplicity, and because it has proven to classify text in much the same manner as humans would (Chang et al., 2009). Thus, the LDA transforms something large and complex, i.e., the corpus, into something that is relatively small, dense, and interpretable.

As is common in this literature, the news corpus is cleaned prior to estimation. We remove stop-words, conduct stemming, and apply term frequency - inverse document frequency calculations. A more detailed description of these steps is given in Appendix A.1. Likewise, in the interest of preserving space, the LDA model is described in Appendix A.2. Note here that, based on results in Larsen and Thorsrud (2019) and Thorsrud (2018), 80 different topics are extracted in total, where the average of the last 10 iterations of the Gibbs simulations, used to estimate the LDA, are used as measures of article weights and topics. Using the output from the LDA, the topic decomposition is transformed into time series, measuring how much each topic is written about at any given point in time. Finally, the tone of the news is computed, using a simple dictionary-based approach and output from the topic model, to sign-adjust the topic frequencies. A more detailed description of this latter step is relegated to Appendix A.3.

To build intuition, Fig. 1 illustrates the output from the above steps for six of the 80 topics. A full list of the estimated topics is given in Table B.1, in Appendix B. First, the LDA produces two outputs; one distribution of topics for each article in the corpus, and one distribution of words for each of the topics. The latter distributions are illustrated using word clouds in Fig. 1. A bigger font illustrates a higher probability for the terms. As the LDA estimation procedure does not give the topics any name, labels are subjectively given to each topic based on the most important terms associated with each topic. How much each topic is written about at any given point in time, and its tone, is illustrated in the graphs below each word cloud. The graphs should be read as follows: Progressively more positive (negative) values means the media writes more about this topic, and that the tone of reporting on this topic is positive (negative).

2.2. News-driven inflation expectations?

The existing literature is largely silent about which types of news households (on average) pay attention to in relation to inflation.[3] Accordingly, we map the high-dimensional news topic dataset to inflation expectations using the Least Absolute Shrinkage and Selection Operator (LASSO; Tibshirani, 1996). The LASSO method shrinks parameter estimates for unimportant variables towards zero, thereby encouraging simple and sparse models.

More formally, we run penalized linear predictive regressions like

$$F_t\pi_{t+12,t+1} = a + \sum_{n=1}^{M} b_n NT_{n,t-1} + \epsilon_t, \tag{1}$$

where $F_t\pi_{t+12,t+1}$ is households' expectations, at time t, of inflation over the next year, and M is the number of news topics $NT_{n,t-1}$. All variables are lagged one period relative to $F_t\pi_{t+12,t+1}$ to avoid simultaneity issues and look-ahead biases. The amount of regularization is optimized by setting the LASSO shrinkage parameter using 5-fold cross-validation, and all variables are standardized prior to estimation. In line with the predictive purpose of the LASSO, we focus on partial R^2 statistics and significance levels, computed using a post-LASSO routine on the selected variable set (Belloni and Chernozhukov, 2013), when reporting the results. Finally, as a benchmark control variable, lagged CPI inflation is always included in the regressions (irrespective of whether or not it is selected by the LASSO).

Column *I* in Table 1 summarizes our first results. Among 80 potential news topics, 20 are selected. Of these, 6 are significant, and the adjusted R^2 statistic is as high as 57%. The selected set includes topics like *Education*, *Trading*, *Health*, *Internet*, *The White House*, and *Transactions*. As also documented in earlier research (see Coibion et al., 2018 for an overview), past inflation is an important variable, while the partial R^2 statistics suggest that the *Health* topic contributes the most to the regression fit among the topics. For future reference, we label the selected topic variable set S^e.

[3] If anything, news coverage about a primary candidate, monetary policy, is potentially ineffective because households in general are found to be poorly informed about central bank policies (see Coibion et al., 2019 and the references therein).

Please cite this article as: V.H. Larsen, L.A. Thorsrud and J. Zhulanova, News-driven inflation expectations and information rigidities, Journal of Monetary Economics, https://doi.org/10.1016/j.jmoneco.2020.03.004

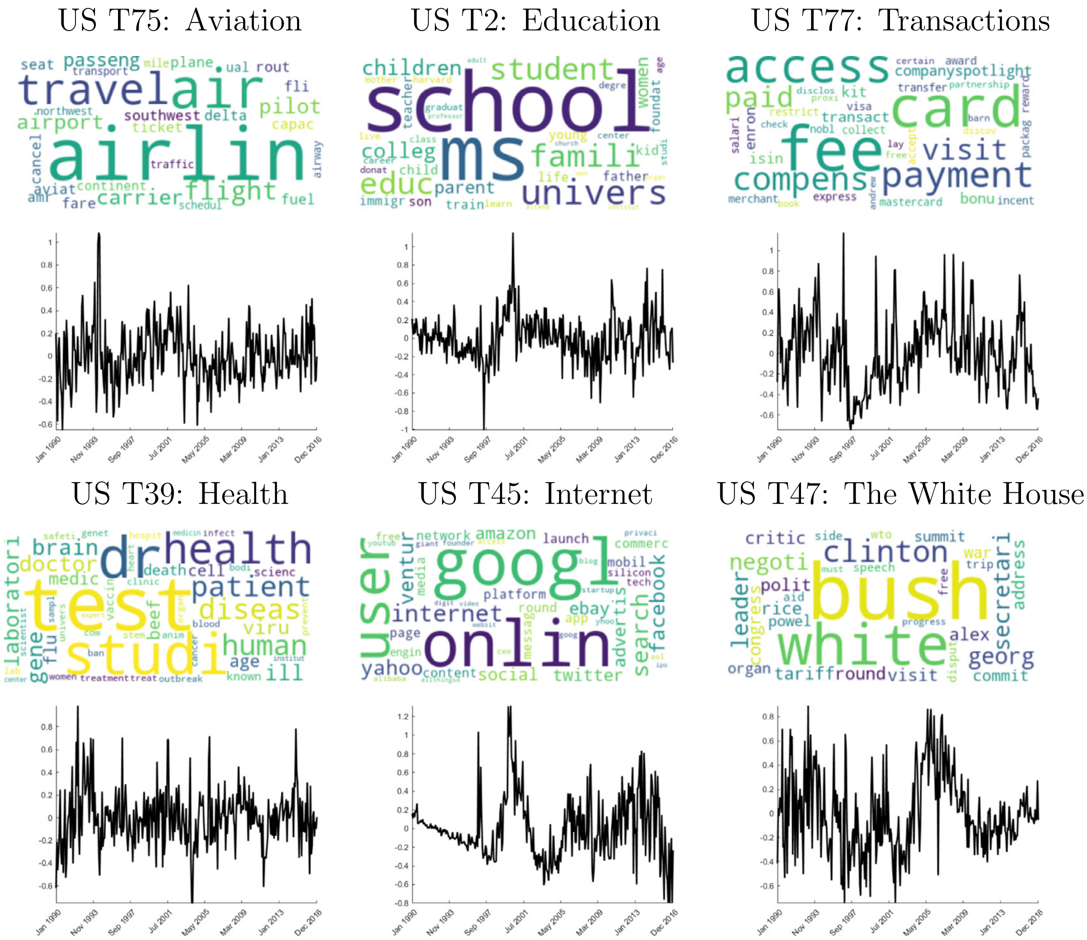

Fig. 1. Topic distributions and time series. For each topic, the size of a word in the word cloud reflects the probability of this word occurring in the topic. Each word cloud only contains a subset of all the most important words in the topic distribution. Topic labels are subjectively given. The topic time series are normalized.

To further test the independent relevance of news topics we augment the regression with roughly 130 hard economic indicators from the *FRED-MD* database and re-estimate the LASSO. The *FRED-MD* is compiled by McCracken and Ng (2016) and is a much used dataset containing (leading) indicators covering the stock market, interest rates and exchange rates, prices, income, consumption, and the labor market (Appendix C). As seen from column *II* in Table 1, the adjusted R^2 statistic increases for this larger model, but not by a very large margin. Fewer topics are also selected, but the significant topics in the news-only regression tend to stay significant.[4]

The topics might not have been given names by us that intuitively link them to inflation expectations. Still, Table 2 shows that the narrative realism of the approach is good. The table contains examples selected by querying the news corpus, of roughly five million articles, for articles where topics important for expectations have a particularly high weight. The *Education* story, for example, talks about expenses, while the *Health* and *Transactions* stories talk about costs and fees. Media coverage related to these types of news might all plausibly affect how households consider inflation developments. As an alternative, to help interpretation, one could interpret each topic as belonging to clusters of higher order abstractions, like, politics, technology, etc. The first columns in Tables 1 and 2 illustrate this, where a clustering algorithm has been used to group the topics into broader categories (Figure B.4 in Appendix B). For example, the *Internet* topic is automatically grouped together with the *Smartphones* and *Software* topics, making it apparent that these news types are (IT) technology related.

In line with our motivation for the topic-based approach, none of the stories listed in Table 2 actually contain explicit *inflation* terms. In contrast, the conventional method used to measure the intensity of media reporting relevant for inflation expectations has been to count the number of terms related to inflation in the corpus' articles. To more formally compare approaches, we construct a traditional media measure by counting terms related to inflation in articles using the wild-card

[4] As illustrated in Table B.2, in Appendix B, among the hard economic indicators selected are many variables already focused on in the earlier literature, such as, production indicators (Ehrmann et al., 2017), volatility measures (Dräger and Lamla, 2017), and consumer sentiment (Doms and Morin, 2004).

ARTICLE IN PRESS

Table 1
Expectations, inflation, and news. The table reports the partial R^2 statistics from regressing MSC expectations and inflation on news topics and past inflation ($\pi_{t-1,t-12}$) using the (post-) LASSO algorithm. For readability, the partial R^2 statistics associated with the topics are scaled to sum to one, using the weights reported in the row Σ *partial* R^2. The *FRED-MD* data are included in the regressions when indicated. To conserve space, the selected hard economic indicators are reported in Table B.2, in Appendix B. Topic names are reported in column two, while the group names in the first column are derived using a hierarchical agglomerative clustering algorithm (Figure B.4 in Appendix B). *, **, and ***, indicate that the underlying coefficients are statistically significant at the 10%, 5%, and 1% level, respectively.

Group name	Topic name	Expectations: $F_t\pi_{t+12,t+1}$		Inflation: $\pi_{t+12,t+1}$		
		I	II	III	IV	V
Banking/	Funding				0.18***	
insurance	Insurance				0.04**	
East Asia	East Asia	0.02		0.06**		
Eductation/	Education	0.11***	0.21***	0.20***	0.02	0.09***
Lifstyle	Public safety	0.00	0.05	0.09***	0.02*	0.15***
	Sports				0.03**	0.14***
Europe	Europe				0.06***	0.09***
Food/Retail	Food				0.10***	
Geo. politics	Russia				0.02	
Gov. econ.	Clients	0.01	0.00	0.20***	0.05**	0.12***
policy	Transactions	0.09**	0.14**	0.00		
Health	Health	0.28***	0.19***	0.01	0.00	0.09***
Leadership	Leadership				0.01	0.00
Macro/Market	Agriculture	0.03		0.02		
	Petroleum				0.03**	
	Stock indices	0.00		0.00		
	Volatility				0.01	
	Labor market	0.04		0.05**	0.05**	
	Fear	0.01		0.01		
	Events	0.02		0.01		
	Commodities				0.02	0.06**
News/results	News service	0.00		0.00		
Politics	Strategy	0.03	0.04	0.00		
	Commentary				0.05**	0.03
	The White House	0.04*		0.01		
Regulation	Regulations				0.04**	0.04*
	Documentation				0.00	0.00
Restructuring	M&A	0.01	0.02	0.00		
Stocks	Stocks	0.01		0.04*		
Technology	Smartphones	0.03		0.04*	0.08***	
	Internet	0.11***	0.28***	0.19***	0.11***	0.19***
Trading	Trading	0.12***		0.05**		
Transportation	Automobiles				0.08***	0.00
	Aviation	0.03	0.06	0.01		
	Σ partial R^2 (topics)	0.21	0.12	0.30	0.44	0.32
Controls	Past Inflation	0.26***	0.29***	0.00	0.00[a]	0.11***[a]
	FRED-MD	*False*	*True*	*False*	*False*	*True*
Summary	# topics / # FRED-MD	20.00 /0	9.00 / 13.00	20.00/ 0	21.00 /0	13.00 /7.00
statistics	Adjusted R^2	0.57	0.69	0.28	0.44	0.60
	Estimator	LASSO	LASSO	OLS	LASSO	LASSO
	Name of the (topic) set	S^e	$S^{e\|FMD}$	S^e	S^π	$S^{\pi\|FMD}$

[a] The variable is not chosen by LASSO, but still included in the post-LASSO regression.

search *inflation** (Figure B.1 in Appendix B) and include this variable in the LASSO together with the other variables. Doing so, we observe it is not selected.

2.3. Expectations and inflation

Economic theory, like the New Keynesian Phillips Curve, suggest there should be a strong relationship between inflation expectations and actual inflation, even in the presence of information rigidities (Coibion et al., 2018). This is also the case here. When regressing CPI inflation on the lagged news topics selected in the LASSO regression reported in column *I* in Table 1, half of the significant topics remain significant (column *III*). Allowing all the news topics to enter the variable selection problem results in a similarly sized set of topics, while the adjusted R^2 statistic increases from 0.28 to 0.44 (column *IV*). Thus, using the news topics relevant for households only reduces the model fit by roughly 35%. These results are robust to controlling for all the hard economic indicators in the variable selection problem. As above, topics in the *Macro/Market*

Please cite this article as: V.H. Larsen, L.A. Thorsrud and J. Zhulanova, News-driven inflation expectations and information rigidities, Journal of Monetary Economics, https://doi.org/10.1016/j.jmoneco.2020.03.004

Table 2
Narrative realism and story examples. The examples are found by querying the corpus, of roughly 5 million articles, for articles where the news topics listed in column two have at least a top-ten weight. The date of publication is printed in parenthesis, and only the first sentences of each article are reported.

Group name	Topic name	Story example
Transportation	Aviation	*(2013-01-24) Want a quick 30% discount on your family's trip to Europe or Hawaii? In the crazy airfare world, sometimes buying two tickets is cheaper than one. Pairing two discounted tickets together to create your own connecting itinerary can often be less expensive than flying on one ticket, if you take advantage of airlines' city-specific specials, or create your own route using discount airlines.*
Education/ Lifestyle	Education	*(2014-02-10) It's no secret that one way to reduce the cost of getting a bachelor's degree is to take classes at a less-expensive community college first. What isn't nearly as well-known is how to go about saving that money. For instance, some parents and students may not realize that not all community-college credits can be transferred and applied toward a higher degree at a four-year school. Or they may not know about programs that allow...*
Technology	Internet	*(2011-12-02) Google's plan to partner with major retailers and shippers to help online shoppers get products delivered within a day signals a ratcheting-up its rivalry with e-commerce king Amazon. But the move likely won't come as a surprise to Amazon CEO... Bezos's initial fears about Google were realized when the fast-rising search engine launched its first price-comparison service, Froogle, in 2002...*
Health	Health	*(2006-08-16) An experimental blood test has shown a glimmer of promise of one day addressing a major health-care challenge: detecting lung cancer at an early stage. The test, developed by researchers at the University of Kentucky, is designed to identify... Mr. Cohen of 20/20 GeneSystems estimates the cost of the blood test would be less than $200. CT scans can cost between $300 and $1,000 and usually aren't covered by...*
Gov. econ. policy	Transactions	*(2010-06-25) Retail banking in the U.S. may never be the same again. Proposed legislation limiting debit-card transaction fees paid by merchants will bite into the income of major U.S. card issuers... Main Street will ultimately pay the price, though. Financial institutions have repeatedly said that, to offset the loss of billions of dollars in revenue, they must charge higher fees on basic banking products and water down rewards programs tied to debit-card use...*

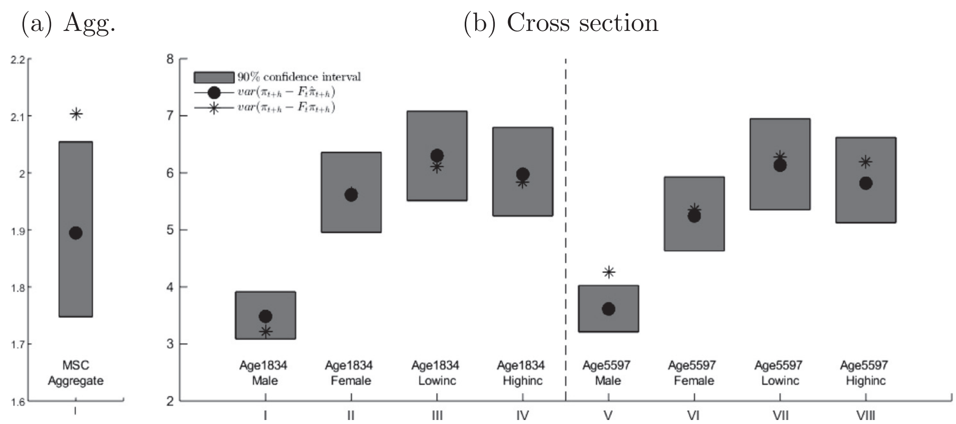

Fig. 2. Forecast error variances. The black stars in the figures report the forecast error variance, i.e., $var(\pi_{t+h} - F_t\pi_{t+h})$, where π_{t+h} denote actual inflation over the next year, and $F_t\pi_{t+h} = F_t\pi_{t+12,t+1}$, i.e., households' expectations of this outcome. Letting $F_t\hat{\pi}_{t+h} \sim N(\alpha + \boldsymbol{\beta}'\boldsymbol{NT}^S_{t-1}, \sigma^2_\epsilon)$, i.e., expectations explained by news topics, the black circles and gray boxes report the median forecast error variance from $var(\pi_{t+h} - F_t\hat{\pi}_{t+h})$, with 90% confidence bands. Fig. 2a and b report these statistics for the aggregated and disaggregated MSC data, respectively. In the former case, the news topics in $\boldsymbol{NT}^S$ are defined by the set listed in column *II* of Table 1. In the latter case, the news topics in $\boldsymbol{NT}^S$ are defined in Table B.3, in Appendix B.

group typically drop out when controlling for the *FRED-MD* data, but almost 60% of the news topics in the selected set $S^{e|FMD}$ for households' expectations (column *II*) are in the selected set $S^{\pi|FMD}$ for inflation (column *V*).

From a forecasting perspective, Fig. 2a shows that households have a forecast error variance above 2. Using the part of expectations explained by the news topics when computing this statistic improves forecasting performance by roughly 10%, i.e., media coverage helps households form more accurate expectations. Additional results presented in Appendix D document that the significant predictive relationship between inflation, expectations, and news topics withstands out-of-sample evaluation.

2.4. The survey of professional forecasters and cross sectional differences

It is interesting to contrast our results with those obtained if households' expectations in (1) are replaced by expectations from SPF. A priori we conjecture that professional forecasters surely know and follow actual CPI inflation and have much less need to delegate their information choice to the media. And, indeed, when predicting quarterly SPF CPI inflation expectations using news topics aggregated to quarterly frequency, none of the news topics are selected.

Please cite this article as: V.H. Larsen, L.A. Thorsrud and J. Zhulanova, News-driven inflation expectations and information rigidities, Journal of Monetary Economics, https://doi.org/10.1016/j.jmoneco.2020.03.004

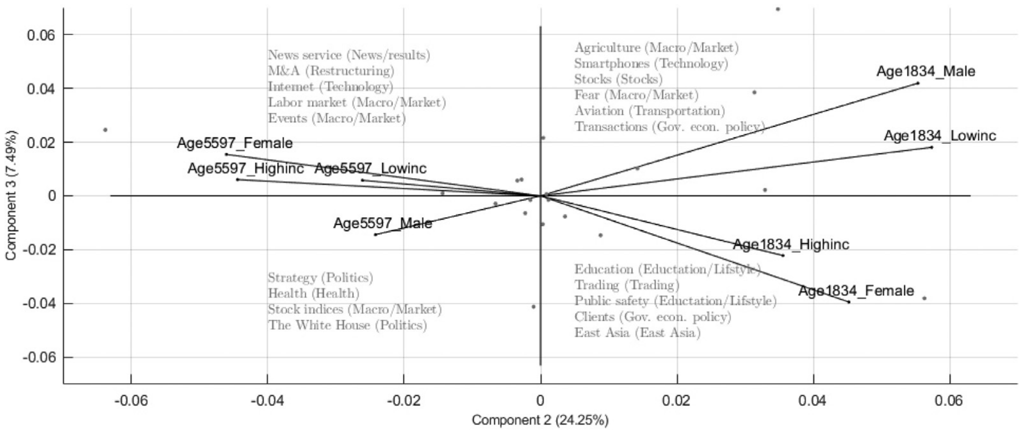

Fig. 3. Biplot of partial R^2. The figure reports a principal-component-based biplot of the normalized partial R^2 weights in Table 3. Three common components explain roughly 90% of the overall variation in the data, and all loadings point in the same (positive) direction for the first common component. For visual clarity, only components two and three are graphed. The gray dotted markers are the factors. Their labels have been stacked together in the lists displayed in each quadrant of the figure depending on their position in the space covered by the decomposition. The text in parenthesis is the associated group name, derived using a hierarchical agglomerative clustering algorithm (Figure B.4 in Appendix B).

As more carefully described in Appendix E, earlier literature and data on expenditure patterns and media consumption habits suggest there is substantial heterogeneity among survey respondents, especially along the age, but also gender and income, dimensions. We use this knowledge and the MSC micro-data to further validate the topic-based approach.

The disaggregated data is naturally more volatile than aggregated expectations. To discipline the analysis, Table 3 shows the results from simply regressing inflation expectations for eight different survey cohorts on the lagged news topics in the set S^e from above. There is a large common component in terms of which news topics predict expectations. Topics related to the *Macro/Market* and *Health* groups often receive particularly high scores and significance. Fig. 3, a biplot of the normalized partial R^2 weights in Table 3, illustrates the more subtle nuances in these results. In line with data on expenditure patterns and media consumption habits, there is a clear distinction along the age dimension. Expectations among elderly are more associated with news about health and politics, while expectations among the young are more associated with news related to education and lifestyle. Whereas the elderly are a relatively homogeneous group, young males relate more to news about, e.g., transportation and technology, while young females relate more to education and lifestyle news. Along the income dimension, however, we do not find any clear mapping to expenditure patterns and media consumption habits in these data.

Expanding the analysis by allowing the LASSO to select the relevant news topics results in more dispersed variable sets. Still, the decomposition of the normalized partial R^2 statistics from the LASSO regressions continue to indicate that the elderly are more associated with health-related news than the young (Figure B.2 and Table B.3 in Appendix B). Likewise, the earlier finding suggesting news increases forecasting accuracy in aggregated expectations tend to hold for elderly survey respondents, but is only statistically significant for elderly male respondents (Fig. 2b).

3. Information rigidities in theory

We now turn to address whether the degree of information rigidity among households is state-dependent and a function of media coverage. To structure the analysis, the easy to implement noisy information model suggested by Coibion and Gorodnichenko (2015) is augmented with a simple reduced form media channel.

We start by making the assumption that households do not follow inflation as measured by the statistical agency per se, but get information about future prices primarily through the media, which operate as information intermediaries between agents and the state of the world (Nimark and Pitschner, 2019). While this information object is high-dimensional, letting π_t^N denote an aggregated measure of relevant media coverage, the signal agent i receives about inflation at time period t is

$$s_{it} = \pi_t^N + \omega_{it} \quad \omega_{it} \sim N(0, \sigma_{\omega t}^2), \tag{2}$$

where ω_{it} is idiosyncratic noise. The noise term captures differences in how agents weigh and interpret different news sources and items, while the precision of the signal is state-dependent.

News coverage has persistence, and the time series properties of media coverage, as perceived by the agents, are modeled as an autoregressive process

$$\pi_t^N = \rho_t^N \pi_{t-1}^N + \nu_t^N \quad \nu_t^N \sim N(0, \sigma_{\nu t}^2), \tag{3}$$

Please cite this article as: V.H. Larsen, L.A. Thorsrud and J. Zhulanova, News-driven inflation expectations and information rigidities, Journal of Monetary Economics, https://doi.org/10.1016/j.jmoneco.2020.03.004

Table 3
Cross-sectional expectations and news. The table reports the partial R^2 statistics from regressing MSC expectations on news topics and past inflation ($\pi_{t-1,t-12}$) using OLS. The set of news topics is chosen based on the results in column *I* in Table 1. For readability, the partial R^2 statistics associated with the topics are scaled to sum to one, using the weights reported in the row Σ *partial* R^2. See Appendix E for a detailed description of how the survey cohorts are constructed from the MSC micro-data. Topic names are reported in column two, while the group names in the first column are derived using a hierarchical agglomerative clustering algorithm (Figure B.4 in Appendix B). *, **, and ***, indicate that the underlying coefficients are statistically significant at the 10%, 5%, and 1% level, respectively.

Group name	Topic name	Age1834				Age5597			
		Male I	Female II	Lowinc III	Highinc IV	Male V	Female VI	Lowinc VII	Highinc VIII
East Asia	East Asia	0.00	0.03	0.00	0.04	0.02	0.00	0.01	0.05
Eductation/	Education	0.00	0.04	0.00	0.01	0.01	0.00	0.00	0.00
Lifstyle	Public safety	0.02	0.03	0.01	0.03	0.00	0.03	0.05	0.00
Gov. econ.	Clients	0.01	0.04	0.01	0.09*	0.00	0.00	0.01	0.00
policy	Transactions	0.13**	0.10**	0.09	0.14**	0.01	0.01	0.01	0.00
Health	Health	0.09*	0.23***	0.03	0.19***	0.25***	0.15***	0.22***	0.13**
Macro/Market	Agriculture	0.00	0.01	0.03	0.00	0.02	0.01	0.02	0.00
	Stock indices	0.00	0.02	0.01	0.00	0.02	0.04	0.00	0.00
	Labor market	0.01	0.03	0.01	0.01	0.22***	0.36***	0.21***	0.37***
	Fear	0.24***	0.00	0.19**	0.00	0.02	0.01	0.05	0.00
	Events	0.02	0.01	0.03	0.00	0.03	0.01	0.01	0.09*
News/results	News service	0.01	0.00	0.00	0.00	0.00	0.04	0.00	0.02
Politics	Strategy	0.00	0.02	0.00	0.10**	0.01	0.07*	0.07*	0.02
	The White House	0.00	0.02	0.05	0.01	0.11*	0.03	0.00	0.11**
Restructuring	M&A	0.00	0.00	0.02	0.03	0.04	0.05	0.09**	0.02
Stocks	Stocks	0.23***	0.08*	0.10	0.10**	0.00	0.06*	0.04	0.03
Technology	Smartphones	0.01	0.00	0.16**	0.05	0.02	0.00	0.00	0.04
	Internet	0.01	0.01	0.02	0.00	0.14**	0.05	0.11**	0.06
Trading	Trading	0.12**	0.32***	0.24**	0.16***	0.09*	0.00	0.00	0.05
Transportation	Aviation	0.09*	0.00	0.00	0.05	0.00	0.09**	0.09**	0.01
	Σ partial R^2	0.13	0.14	0.08	0.14	0.11	0.15	0.15	0.13
	Past inflation	0.09***	0.08***	0.05***	0.07***	0.01	0.01*	0.00	0.02**
	Adjusted R^2	0.15	0.24	0.12	0.15	0.15	0.16	0.11	0.18
	Estimator	OLS	OLS	OLS	OLS	OLS	OLS	OLS	OLS

where ρ_t^N and $\sigma_{\nu t}^2$ depend on the time index t. Variation in ρ_t^N can be due to major economic or political events that become extensively covered by the media, while a higher $\sigma_{\nu t}^2$ implies that news reporting becomes less predictable, e.g., in times of abrupt economic or political changes.

To link inflation to news about inflation, we build on the results presented in the previous section and assume the media fulfills its purpose in informing the public about important developments in society, and work with a tractable and simple editorial function

$$\pi_t^N = \pi_t + \alpha_t, \tag{4}$$

where α_t is a time-fixed effect, capturing for example potential media biases. Importantly, under the maintained assumption that agents do not follow inflation per se, they are not in the position to bias-adjust the news signal towards actual inflation. Likewise, combining (3) and (4) implies that agents' perceived time series properties of news are a composite of actual inflation developments and time-fixed media effects. However, since agents' have delegated their information choice to the media, they are not able to discriminate between these two factors.[5]

As agents do not observe relevant news coverage directly (π_t^N), but only a noisy measure of it, the fundamental model feature is a signal extraction problem. The agents use the Kalman filter for this purpose. Given (2) and (3), the Kalman Gain is

$$K_t = \rho_t^N \Psi_t (\Psi_t + \sigma_{\omega t}^2)^{-1}, \tag{5}$$

and captures the weight assigned to new information about π_t^N in the prediction error (with variance Ψ_t). Averaging across agents, iterating h periods forward, and using (4), gives

$$\pi_{t+h} - F_t\pi_{t+h} = c_t + \beta_t (F_t\pi_{t+h} - F_{t-1}\pi_{t+h}) + e_t, \tag{6}$$

where $F_t\pi_{t+h}$ is households' expectation of future inflation, $\beta_t = \frac{1-K_t}{K_t}$, $c_t = -\alpha_{t+h}$, and $e_t = \sum_{j=1}^{h} (\rho_t^N)^{h-j} \nu_{t+j}^N$.

[5] In general, these assumptions are consistent with Nimark and Pitschner (2019), who establish optimality conditions for the delegated information choice mechanism, and they are consistent with a substantial literature showing that people are not fully informed about important expenses (Carter and Milon, 2005; Chetty and Saez, 2013; Jensen, 2010). Moreover, most of the variation in household level inflation is disconnected from movements in aggregate inflation (Kaplan and Schulhofer-Wohl, 2017), making it perfectly rational for households to not follow aggregate inflation directly, but rather use the news media for this purpose.

As in Coibion and Gorodnichenko (2015), Eq. (6) describes the relationship between ex-post forecast errors and ex-ante mean forecast revisions. Although individuals form their forecasts rationally conditional on their information set, the ex-post mean forecast error across agents is systematically predictable using ex-ante mean forecast revisions due to gradual adjustment of beliefs to new information. A higher value of β_t implies a higher degree of information rigidity. Conversely, if $\beta_t = c_t = 0$, we have FIRE. The media effect comes through (5). β_t decreases if media persistence (ρ_t^N) is high and increases if the amount of noise in the signal ($\sigma^2_{\omega t}$) is high (relative to $\sigma^2_{\upsilon t}$). In contrast, in the conventional model, where agents are assumed to follow inflation directly, the degree of information rigidity is determined by properties of inflation itself.[6,7]

4. Information rigidities in the data

The theoretical predictions from the model in the previous section are tested using a two-step estimation approach. In Section 4.1 (6) is used to estimate β_t, while we test if the underlying time series features of media coverage help explain the evolution of β_t, as predicted by (5), in Section 4.2.

4.1. Time-varying information rigidities?

The MSC survey only contains households' forecast of inflation over the course of the next year, resulting in non-overlapping time periods in observed forecast revisions. For this reason, we follow Coibion and Gorodnichenko (2015), and instrument the forecast revisions using the (log) change in the monthly price of oil. In a time-varying parameter setting, however, a traditional instrumental variable (IV) estimator will still be biased due to the induced correlation between the time-varying parameters and the error term. This issue has often been ignored in the literature (Chang-Jin et al., 2010), but can be solved using a control function approach. As described in Appendix H.1, this implies the following system

$$y_t = c_t + \beta_t x_t + \gamma v_t^* + w_t \qquad w_t \sim i.i.d.N(0, \sigma_w^2) \tag{7}$$

$$x_t = \delta_t z_t + \sigma_\upsilon v_t^* \qquad v_t^* \sim i.i.d.N(0, 1), \tag{8}$$

where $y_t = \pi_{t+12,t+1} - F_t\pi_{t+12,t+1}$ and $x_t = F_t\pi_{t+12,t+1} - F_{t-1}\pi_{t+11,t}$ denote households' forecast errors and revisions, respectively, of U.S. headline CPI inflation over the next year, and z_t is the instrument. $v_t^* = \sigma_\upsilon^{-1}(x_t - z_t\delta_t)$ is the control function, and the disturbance term w_t is uncorrelated with x_t and β_t conditional on v_t^*.

To be faithful to the null hypothesis of full information, i.e., $\beta_t = 0$, we use the Latent Threshold Model (LTM) idea by Nakajima and West (2013) to enforce dynamic sparsity on the system through the time-varying parameters. For β_t the LTM structure can be written as

$$\beta_t = \beta_t^* \varsigma_{\beta,t} \quad \varsigma_{\beta,t} = I(|\beta_t^*| \geq d_\beta) \quad \beta_t^* = \beta_{t-1}^* + \upsilon_{\beta^*,t}, \tag{9}$$

where β_t^* follows a random walk process, with $\upsilon_{\beta^*,t} \sim i.i.d.N(0, \sigma^2_{\beta^*\upsilon})$, and $\varsigma_{\beta,t}$ is a zero one variable, whose value depends on the indicator function $I(|\beta_t^*| \geq d_\beta)$. If $|\beta_t^*|$ is above the threshold value d_β, then $\varsigma_{\beta,t} = 1$, otherwise $\varsigma_{\beta,t} = 0$, and β_t shrinks to zero. For the c_t parameter, a similar, but independent, structure is assumed. For the δ_t parameter in (8), sparsity is not enforced. Doing so would go against the standard IV relevance criterion. Instead, we let δ_t follow a regular random walk process with error term $\upsilon_{\delta,t} \sim i.i.d.N(0, \sigma^2_{\delta\upsilon})$. Finally, $\upsilon_{\beta^*,t}$, $\upsilon_{c^*,t}$, and $\upsilon_{\delta,t}$ are assumed to be independent of each other and w_t and v_t^*.

Eqs. (7) and (8), together with the law-of-motion for c_t, β_t, and δ_t, are used to estimate all the parameters of the model jointly in a state space system using MCMC simulations. This avoids concerns about generated regressors in two-stage approaches, and allows us to sample the model's latent states jointly with the hyper-parameters. In the interest of conserving space, details about priors, initialization, and the estimation algorithm are relegated to Appendix H.1. We note here, however, two points about the prior specification which are particularly relevant in this setting. First, we set the prior variance for $\sigma^2_{\beta^*\upsilon}$ equal to $(0.2)^2$. This results in a roughly 95% prior probability of a sevenfold cumulative change in β_t^* over the sample length considered here, which is well inside the range of low frequency change in information rigidity documented in Coibion and Gorodnichenko (2015) for professional forecasters.[8] Second, to obey the IV relevance criterion, we a priori allow for much less variation in δ_t and set the prior variance of $\sigma^2_{\delta\upsilon}$ equal to $(0.01)^2$ (and initialize δ_t at the OLS solution).

The MSC forecast errors and revisions are reported in Fig. 4a, while the time-varying posterior estimates of β_t are reported in Fig. 4b. Given that we work in a high-dimensional time-varying parameter setting, the posterior uncertainty in

[6] Appendix F shows that this is also the case here if agents form an expectation about α_t in (4). However, as documented in Section 4.3, using properties of inflation gives results at odds with theory, suggesting that this assumption is questionable.

[7] As noted by one referee, another plausible mechanism for assessing how the media affects information rigidities is to assume news coverage provides noisy signals about inflation developments directly. Although this line of reasoning does not map fully into the framework presented above, it still captures the underlying idea, where information rigidity should be lower in times of higher precision. Appendix G expands on this reasoning, and shows that the main conclusion presented below also holds under this alternative view.

[8] This prior assumption only affects the cumulative change we might observe, not the time evolution of the parameter itself. Our main results are fairly robust to other reasonable prior choices (Appendix H.1.3).

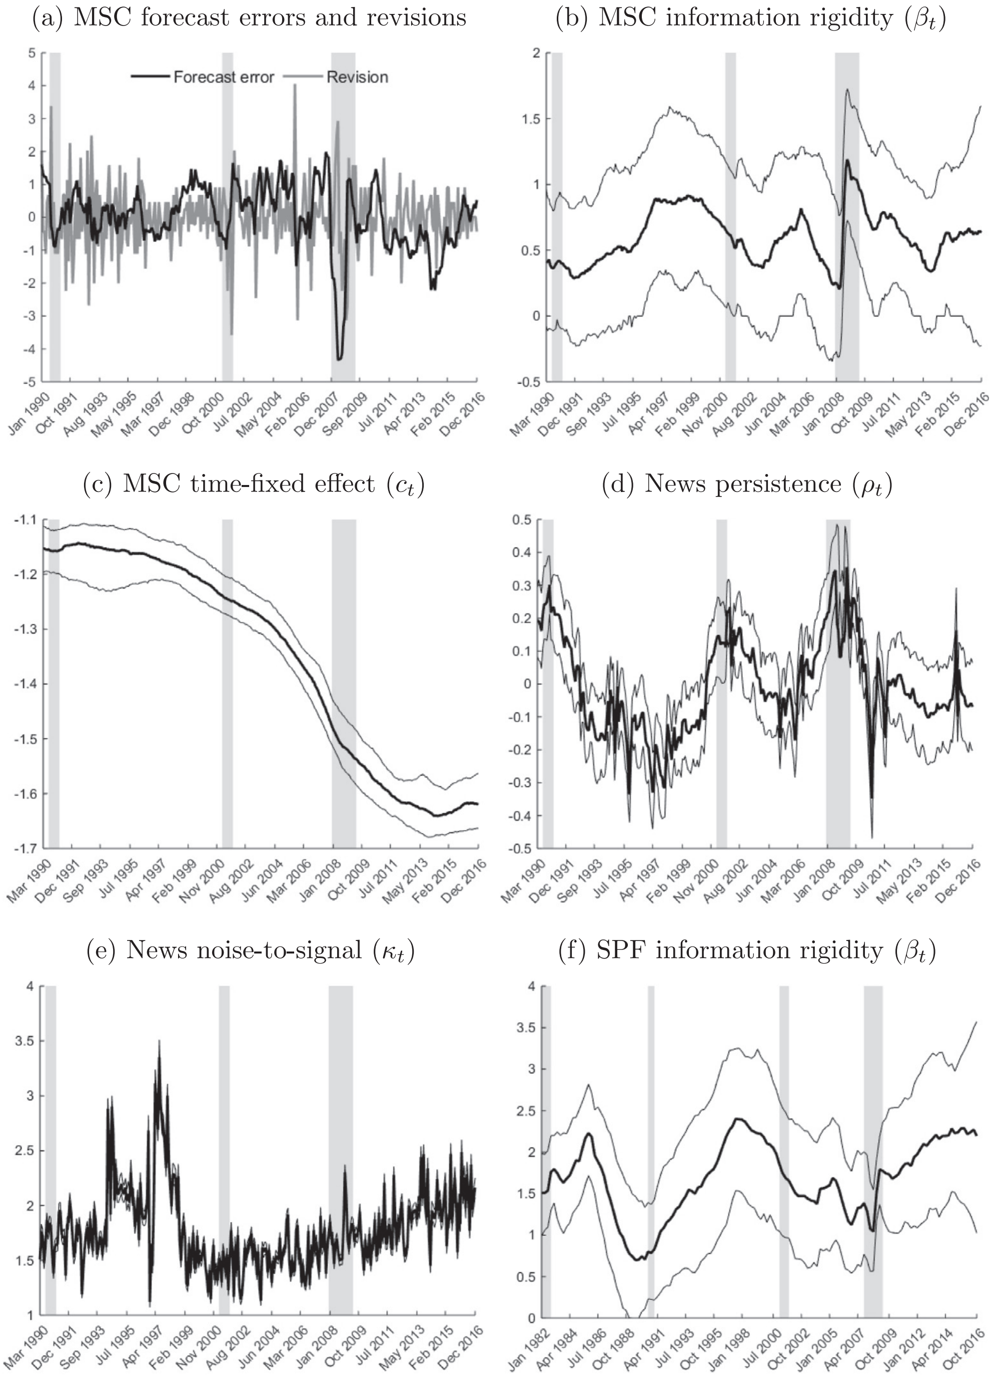

Fig. 4. Data and estimation output. In Fig. 4b–f, the black solid lines are the median estimates, while the gray solid lines are 68% probability bands. The gray shaded areas are recession periods defined by NBER (U.S.). In Fig. 4f, the x-axis shows quarterly dates. For the other graphs, the x-axis shows monthly dates.

the β_t estimate is naturally large. Still, three periods stand out as having a particular high degree of information rigidity, namely the late 1990s, mid 2000s, and the financial crisis years.[9]

[9] Although not our primary focus, the c_t parameter is negative and downward trending (Fig. 4c). This indicates that media biases are not constant across time, as also suggested by findings in, e.g., Souleles (2004), and a full departure from FIRE.

Interpreted through the lens of the model in Section 3, our estimates suggest that the weight agents put on new information, i.e., $\hat{K}_t = 1/(1+\hat{\beta}_t)$, varies from roughly 0.7, during the early 1990s and prior to the Great Recession period, to less than 0.5 during the financial crisis, with an average across the sample of roughly 0.6. As discussed in Coibion and Gorodnichenko (2015), information rigidity of this magnitude has profound macroeconomic effects in theoretical models incorporating information frictions.

To bridge our analysis with the earlier literature, Fig. 4f reports estimates of (7) using quarterly SPF expectations, with details about estimation relegated to Appendix H.1. Starting in the early 1990s, information rigidity among professionals shares some of the same time-varying features as those obtained for households, and, as in Coibion and Gorodnichenko (2015), it also contains an upward-drifting trend during this time period. Moreover, the time series average of the series is well in line with constant parameter estimates obtained in earlier studies, indicating that professionals only put roughly 40% weight on new relative to old information. This number is lower than for households, which might be surprising. Still, an extensive evaluation of a constant parameter version the model, for both MSC and SPF data, yields the same conclusion (Tables B.4 and B.5 in Appendix B).

4.2. News-driven information rigidities?

To test whether the degree of information rigidity among households is a function of media coverage, as predicted by (5), we estimate

$$\beta_t = c + \gamma_1 \rho_t + \gamma_2 \kappa_t + u_t, \tag{10}$$

where β_t is the median time-varying information rigidity, reported in Fig. 4b, while ρ_t and κ_t are the persistence and noise-to-signal ratio in the underlying information set S. S, ρ_t and κ_t are defined as follows:

First, we use the LASSO results from Section 2 to define S. At an abstract level, the idea here is to construct an approximation to the high dimensional object π_t^N in (2). Under the assumption that only news topics with predictive power for expectations are relevant for describing the information households care about, we use the set $S^{e|FMD}$ from Table 1 as our *Benchmark* selection.

Second, for each of the news topics i in S, we estimate time-varying autoregressive models, of order one, to obtain posterior draws of $\hat{\rho}_{i,t}$ and $\hat{\sigma}_{i,t}$, i.e., the time-varying persistence and volatility for news topic i. The model structure, together with the Gibbs simulations used for estimation, is standard in the time series literature and is described in greater detail in Appendix H.2. Next, a measure of the noise in the signal, denoted $\hat{\omega}_{i,t}$, is constructed using the sum of the standard deviation in the posterior article weight distributions, and (therefore also) in the articles selected to tone-adjust the news topic time series. Both estimates are easily available from the topic model output. Intuitively, this noise measure can be interpreted as follows: If the topic proportions assigned to each article by the topic model algorithm have high posterior variance, it would likely be difficult for humans to assign accurate topic weights as well. Thus, uncertainty regarding what the news is about increases. Likewise, uncertainty increases if articles differ in terms of their tone.

Finally, aggregating across all the news topics i in the set S, and combining the output from step two above, one obtains

$$\hat{\rho}_t = \sum_{i=1}^{n} \varpi_i \hat{\rho}_{i,t} \quad \hat{\kappa}_t = \sum_{i=1}^{n} \varpi_i \hat{\kappa}_{i,t} \text{ with } \hat{\kappa}_{i,t} = \frac{\hat{\omega}_{i,t}}{\hat{\sigma}_{i,t}} \text{ and } \varpi_i = \frac{R_i^2}{\sum_{i=1}^{n} R_i^2}, \tag{11}$$

where ϖ_i is the normalized partial R^2 statistic from the relevant regression in Table 1. Thus, variables are weighted according to their relative importance when constructing $\hat{\rho}_t$ and $\hat{\kappa}_t$, while uncertainty in these estimates comes from the posterior distribution of $\hat{\rho}_{i,t}$ and $\hat{\sigma}_{i,t}$.

To control for the generated regressor issue, we use non-informative natural conjugate priors and draw from (10) using the posterior estimates of $\hat{\rho}_t$ and $\hat{\kappa}_t$. Accordingly, the parameter estimates γ_1 and γ_2 are drawn from the OLS solution, but taking into account the generated regressors issue by sampling from the full distribution of $\hat{\rho}_t$ and $\hat{\kappa}_t$ (Bianchi et al., 2017).

Fig. 4d and e graphs the posterior distribution of the estimates in (11). News persistence varies significantly across time and tends to be especially high around recession periods. The noise-to-signal ratio displays a more surprising pattern. It associates the mid 1990s as a particularly "noisy period" and contains an upward-drifting trend starting around year 2000.[10]

Column *I* in Fig. 5a reports the results from estimating (10) using the *Benchmark* selection. The time series properties of relevant topics help explain the time-varying information rigidity among households, and the coefficient estimates have the correct sign. A higher persistence and lower noise-to-signal ratio lead to a reduction in information rigidity.[11]

To control for potential omitted variable biases the model specification in (10) is augmented with 10 factors extracted from the economic indicators in the *FRED-MD* database and a double selection algorithm (Belloni et al., 2014), controlling for all the roughly 130 variables in this database, is implemented. Computational details are described in Appendix H.4. As

[10] The method used to construct this variable is intuitive, but also somewhat sensitive to the raw corpus data. If some time periods contain news extracts from fewer, or different types, of articles, this might contaminate our noise-to-signal measure.

[11] Because the most important topics are the same in both sets, this result also holds when defining S based on S^e instead of $S^{e|FMD}$. Likewise, this result is not driven by the Great Recession. Augmenting the model with a dummy variable for this time period only increases the explanatory power of the model.

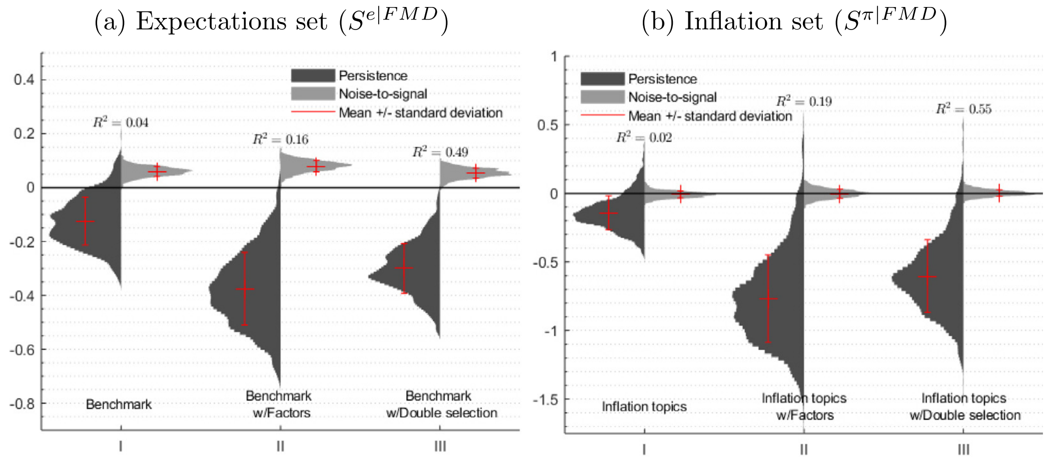

Fig. 5. Violin plots of posterior distributions. Fig. 5a, column *I*, reports the estimates from the *Benchmark* selection ($S^{e|FMD}$). Dark and light gray areas correspond to the distributions for γ_1 and γ_2, from Eq. (10), respectively. The red crosses mark the mean estimate +/− one standard deviation. The mean adjusted R^2 statistic is reported above the distributions. Columns *II* and *III* report estimates from an augmented version of the *Benchmark*, using either: 10 factors from the *FRED-MD* database, or a double selection algorithm controlling for all the variables in this database. Fig. 5b reports the same type of estimates as Fig. 5a, but now with $\hat{\rho}_t$ and $\hat{\kappa}_t$ constructed based on the set $S^{\pi|FMD}$.

Table 4
Cumulative histogram. Eq. (10) is estimated for 100 randomly selected sets of news topics (not in the set of topics used to generate the *Benchmark* parameter distributions). The table reports the fraction of draws that have posterior probabilities $Pr(\gamma_1 < 0) \geq x$ and $Pr(\gamma_2 > 0) \geq x$, where x refers to a bin in the histogram. The bin associated with the posterior probability for the *Benchmark* is marked in gray.

x	10	20	30	40	50	60	70	80	90
$Pr(\gamma_1 < 0) \geq x$	0.54	0.48	0.34	0.24	0.20	0.15	0.11	0.07	0.02
$Pr(\gamma_2 > 0) \geq x$	0.96	0.90	0.89	0.86	0.82	0.81	0.79	0.77	0.71

seen from columns *II* and *III* in Fig. 5a, taking aboard potentially omitted variables strengthens the result that media matters further.

These results are very unlikely to be obtained by chance, at least for the persistence parameter. We show this by running a falsification experiment. First, 100 different sets of news topics, not including topics in the union of S^e and $S^{e|FMD}$, are constructed. Then, for each of these alternative sets, we calculate $\hat{\rho}_t$ and $\hat{\kappa}_t$, and redo the estimation of (10). As illustrated in Table 4, only roughly 2% of the alternative regressions have a posterior probability of $\hat{\gamma}_1 < 0$ that is equal to, or greater, than 90%. For the noise-to-signal ratio, the result is less strong, and 71% of the alternative regressions have roughly the same high posterior probability of $\hat{\gamma}_2 > 0$ as the *Benchmark* selection. This suggests that there is a large common component in the noise-to-signal variable, which might be attributed to the measurement issue mentioned earlier.

The following section describes how our results are robust along a number of additional dimensions.

4.3. Inflation, alternative news measures, and robustness

First, since it is difficult to argue that news relevant for expectations, but not actual inflation, should help lower information rigidities, we also consider the set of topics in $S^{\pi|FMD}$, from Table 1, when constructing $\hat{\rho}_t$ and $\hat{\kappa}_t$, and re-estimate (10). The noise-to-signal component becomes insignificant for this specification, while the earlier result for the persistence parameter is largely confirmed (Fig. 5b). Given the results from Section 2.3, these findings are as expected. The intersection of $S^{e|FMD}$ and $S^{\pi|FMD}$ is large.

This finding does not imply that there is something inherent with our method forcing there to be a significant relationship between information rigidity and the persistence in the variables selected by the LASSO. Using the same methodology as above to compute the persistence in the hard economic variables relevant for households' inflation expectations, confer column *II* in Table B.2 in Appendix B, and including this persistence measure together with the news-based one in the regression, we find that it plays an insignificant role (column *I* in Fig. 6a).[12]

Likewise, the results presented thus far do not necessarily imply that one can use the persistence of inflation itself to explain the degree of information rigidity among households. The no-media channel version of the noisy information model

[12] Coibion and Gorodnichenko (2015) use real-time revisions in the series of interest as a proxy for noise. As most of the hard economic series selected by the LASSO are not revised, the noise-to-signal components are not included in this and the following regressions.

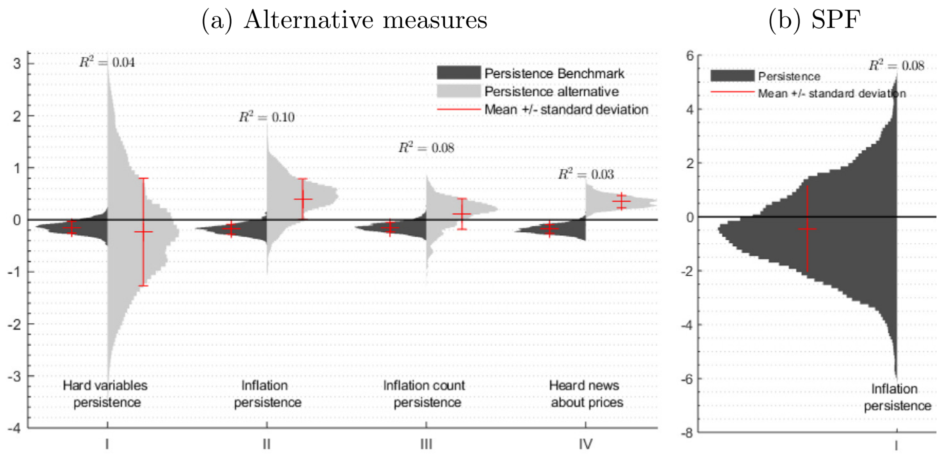

Fig. 6. Violin plots of posterior distributions for alternative model specifications. In Fig. 6a Eq. (10) is estimated without the noise-to-signal term, but with; The persistence in hard economic variables relevant for inflation expectations (*I*); Using the persistence in inflation (*II*); Using the persistence in the alternative inflation count series (*III*); Using the *Heard news about prices* series from the Michigan Survey of Consumers (*IV*). In Fig. 6b Eq. (10) is estimated without the noise-to-signal term, but using β_t for the SPF data as dependent variable and the persistence in inflation as the explanatory variable. See the text and the notes to Fig. 5 for additional details.

in Section 3 suggests that one can, while the media augmented version of the model suggest that one cannot. The result presented in column *II* in Fig. 6a supports the latter view. When the benchmark regression is augmented with the time-varying persistence in inflation, the inflation-based persistence parameter has high posterior variance, and if anything, the wrong sign. In contrast, if we instead regress β_t for the SPF data on inflation persistence, the parameter estimate has the correct sign, although uncertainty is high (Fig. 6b). Accordingly, in line with our findings in Section 2.3, the media channel matters for households, but less so for professionals.

The results presented in columns *III* and *IV* in Fig. 6a document that the topic-based approach provides a more theory-consistent description of information rigidities among households than more traditional news- and survey-based methods do; Regressing households' information rigidity on the topic-based persistence measure and the persistence in the alternative count-based media measure (Section 2.2) results in an uncertain point estimate with the wrong sign for the latter variable; Including the variable measuring if people have heard news about prices from the MSC itself (Figure B.3 in Appendix B) in the regression, we observe, as in Ehrmann et al. (2017), that it actually increases information rigidities.

Finally, in the interest of conserving space, a detailed analysis using the MSC micro-data is relegated to Appendix E. In short, information rigidity is time-varying also among survey cohorts, at least when the level of aggregation is not too low, and the relationship between information rigidities and relevant media coverage described above largely holds.

5. Conclusion

Media's role in the expectation formation process has received relatively little attention in macroeconomics. Using a novel news topic-based approach, this paper contributes by investigating the role of the media for inflation expectations and information rigidities among U.S. households. Taking the view that the degree of information rigidity is state-dependent, and that the media act as information intermediaries between agents and the state of the world, we find empirical support for the following: First, the news types the media choose to report on are good predictors of both inflation and inflation expectations, and news coverage helps households form more accurate expectations. Second, in a standard noisy information model, augmented with a simple media channel, we document that the degree of information rigidity among households varies across time, and that relevant media coverage helps explain this variation. These results are robust to numerous alternative experiments and also largely hold when analyzing the cross-sectional dimension of households' expectations. Thus, our analysis should be useful for future theoretical and empirical work investigating media's important role in the expectation formation process.

Declaration of Competing Interest

None.

CRediT authorship contribution statement

Vegard H. Larsen: Data curation, Visualization, Methodology, Writing - review & editing, Software. **Leif Anders Thorsrud:** Methodology, Investigation, Validation, Visualization, Writing - review & editing, Software, Project administration. **Julia Zhulanova:** Methodology, Investigation, Writing - review & editing, Software.

ARTICLE IN PRESS

Acknowledgments

This article should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank. Comments by the editor, one anonymous referee, and Kristoffer Nimark improved the quality of this paper considerably. We also thank Hilde C. Bjørnland, Paul Beaudry, Drago Bergholt, Gunnar Bårdsen, Olivier Coibion, Òscar Jordà, Michael McMahon, Francesco Ravazzolo, and Mathias Trabandt for valuable comments. This work is part of the research activities at the Centre for Applied Macroeconomics and Commodity Prices (CAMP) at the BI Norwegian Business School. We are grateful to the *Dow Jones Newswires Archive* for sharing their data with us for this research project.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.jmoneco.2020.03.004.

References

Armantier, O., Nelson, S., Topa, G., van der Klaauw, W., Zafar, B., 2016. The price is right: updating inflation expectations in a randomized price information experiment. Rev. Econ. Stat. 98 (3), 503–523.
Belloni, A., Chernozhukov, V., 2013. Least squares after model selection in high-dimensional sparse models. Bernoulli 19 (2), 521–547.
Belloni, A., Chernozhukov, V., Hansen, C., 2014. High-dimensional methods and inference on structural and treatment effects. J. Econ. Perspect. 28 (2), 29–50.
Bianchi, D., Guidolin, M., Ravazzolo, F., 2017. Macroeconomic factors strike back: a Bayesian change-point model of time-varying risk exposures and premia in the u.s. cross-section. J. Bus. Econ. Stat. 35 (1), 110–129.
Blei, D.M., Ng, A.Y., Jordan, M.I., 2003. Latent Dirichlet allocation. J. Mach. Learn. Res. 3, 993–1022.
Blinder, A.S., Krueger, A.B., 2004. What Does the Public Know about Economic Policy, and How Does It Know It? Working Paper. National Bureau of Economic Research. 10787
Carroll, C.D., 2003. Macroeconomic expectations of households and professional forecasters. Q. J. Econ. 118 (1), 269–298.
Carter, D.W., Milon, J.W., 2005. Price knowledge in household demand for utility services. Land Econ. 81 (2), 265–283.
Chang, J., Gerrish, S., Wang, C., Boyd-graber, J.L., Blei, D.M., 2009. Reading tea leaves: how humans interpret topic models. In: Bengio, Y., Schuurmans, D., Lafferty, J., Williams, C., Culotta, A. (Eds.), Advances in Neural Information Processing Systems 22. The MIT Press, Cambridge, MA, pp. 288–296.
Chang-Jin, K., et al., 2010. Dealing with endogeneity in regression models with dynamic coefficients. Found. Trends® Econom. 3 (3), 165–266.
Chetty, R., Saez, E., 2013. Teaching the tax code: earnings responses to an experiment with EITC recipients. Am. Econ. J. 5 (1), 1–31.
Coibion, O., Gorodnichenko, Y., 2012. What can survey forecasts tell us about information rigidities? J. Polit. Econ. 120 (1), 116–159.
Coibion, O., Gorodnichenko, Y., 2015. Information rigidity and the expectations formation process: a simple framework and new facts. Am. Econ. Rev. 105 (8), 2644–2678.
Coibion, O., Gorodnichenko, Y., Kamdar, R., 2018. The formation of expectations, inflation and the phillips curve. J. Econ. Lit. 56 (4), 1447–1491.
Coibion, O., Gorodnichenko, Y., Weber, M., 2019. Monetary Policy Communications and their Effects on Household Inflation Expectations. Working Paper. National Bureau of Economic Research. 25482
Curtin, R., 2007. What US Consumers Know About Economic Conditions. OECD.
Doms, M., Morin, N.J., 2004. Consumer Sentiment, the Economy, and the News Media. Finance and Economics Discussion Series 2004-51. Board of Governors of the Federal Reserve System (US).
Dovern, J., Fritsche, U., Loungani, P., Tamirisa, N., 2015. Information rigidities: comparing average and individual forecasts for a large international panel. Int. J. Forecast. 31 (1), 144–154.
Dräger, L., Lamla, M.J., 2017. Imperfect information and consumer inflation expectations: evidence from microdata. Oxf. Bull. Econ. Stat. 79 (6), 933–968.
Ehrmann, M., Pfajfar, D., Santoro, E., 2017. Consumers' attitudes and their inflation expectations. Int. J. Central Bank. 13 (1), 225–259.
Fullone, F., Gamba, M., Giovannini, E., Malgarini, M., 2007. What Do Citizens Know about Statistics. OECD.
Gentzkow, M., Shapiro, J.M., Sinkinson, M., 2011. The effect of newspaper entry and exit on electoral politics. Am. Econ. Rev. 101 (7), 2980–3018.
Grossman, S.J., Stiglitz, J.E., 1980. On the impossibility of informationally efficient markets. Am. Econ. Rev. 70 (3), 393–408.
Jensen, R., 2010. The (perceived) returns to education and the demand for schooling. Q. J. Econ. 125 (2), 515–548.
Kaplan, G., Schulhofer-Wohl, S., 2017. Inflation at the household level. J. Monet. Econ. 91 (C), 19–38.
King, G., Schneer, B., White, A., 2017. How the news media activate public expression and influence national agendas. Science 358 (6364), 776–780.
Lamla, M.J., Lein, S.M., 2014. The role of media for consumers' inflation expectation formation. J. Econ. Behav. Organ. 106, 62–77.
Larsen, V.H., Thorsrud, L.A., 2019. The value of news for economic developments. J. Econom. 210 (1), 203–218.
Loungani, P., Stekler, H., Tamirisa, N., 2013. Information rigidity in growth forecasts: some cross-country evidence. Int. J. Forecast. 29 (4), 605–621.
Mackowiak, B., Wiederholt, M., 2009. Optimal sticky prices under rational inattention. Am. Econ. Rev. 99 (3), 769–803.
McCracken, M.W., Ng, S., 2016. FRED-MD: a monthly database for macroeconomic research. J. Bus. Econ. Stat. 34 (4), 574–589.
Nakajima, J., West, M., 2013. Bayesian analysis of latent threshold dynamic models. J. Bus. Econ. Stat. 31 (2), 151–164.
Nimark, K.P., Pitschner, S., 2019. News media and delegated information choice. J. Econ. Theory 181, 160–196.
Pfajfar, D., Santoro, E., 2013. News on inflation and the epidemiology of inflation expectations. J. Money Credit Bank. 45 (6), 1045–1067.
Prat, A., 2018. Media power. J. Polit. Econ. 126 (4), 1747–1783.
Shiller, R.J., 2017. Narrative economics. Am. Econ. Rev. 107 (4), 967–1004.
Sims, C.A., 2003. Implications of rational inattention. J. Monet. Econ. 50 (3), 665–690.
Souleles, N.S., 2004. Expectations, heterogeneous forecast errors, and consumption: micro evidence from the Michigan consumer sentiment surveys. J. Money Credit Bank. 36 (1), 39–72.
Thorsrud, L.A., 2018. Words are the new numbers: a newsy coincident index of the business cycle. J. Bus. Econ. Stat. 1–17.
Tibshirani, R., 1996. Regression shrinkage and selection via the lasso. J. R. Stat. Soc. Series B 58 (1), 267–288.
Woodford, M., 2009. Information-constrained state-dependent pricing. J. Monet. Econ. 56 (S), 100–124.

Please cite this article as: V.H. Larsen, L.A. Thorsrud and J. Zhulanova, News-driven inflation expectations and information rigidities, Journal of Monetary Economics, https://doi.org/10.1016/j.jmoneco.2020.03.004