DOI: 10.1111/jmcb.12868

SASKIA TER ELLEN
VEGARD H. LARSEN
LEIF ANDERS THORSRUD

Narrative Monetary Policy Surprises and the Media

We propose a simple method to quantify narratives from textual data, and identify "narrative monetary policy surprises" as the difference in narrative focus in central bank communication accompanying interest rate meetings and economic media coverage prior to those meetings. Identifying narrative surprises, using Norwegian data, provides surprise measures that are uncorrelated with conventional monetary policy surprises, and, in contrast to such surprises, have a significant effect on subsequent media coverage. Narrative monetary policy surprises lead to macroeconomic responses similar to what recent monetary policy literature associates with the information component of monetary policy communication, highlighting media's role as information intermediaries.

C01, C55, C82, E43, E52, E58
Keywords: communication, factor identification, monetary policy, textual data

…if researchers are interested in testing market responses to [central bank] communication, it may make sense to focus on statements that actually reach market participants, and on the content as conveyed by the media. (Blinder et al., 2008)

This paper 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. Helene Olsen provided valuable research assistance, while comments from three anonymous referees improved the quality of this paper considerably. We also thank Carola Binder, Jan Qvigstad, seminar participants at Norges Bank, and conference participants of The Changing Economy and the Business Cycle for valuable comments. We are grateful to *Retriever* for sharing data with us for this research project. This work is part of the research activities at the Centre for Applied Macroeconomics and Commodity Prices (CAMP) at the BI Norwegian Business School.

Saskia ter Ellen is a Senior Researcher, *Norges Bank* (*E-mail*: saskia.ter-ellen@norges-bank.no). Vegard H. Larsen is a Senior Researcher, *Norges Bank, and Centre for Applied Macroeconomics and Commodity Prices, BI Norwegian Business School* (*E-mail*: vegard-hoghaug.larsen@norges-bank.no). Leif Anders Thorsrud*is an Associate Professor, Centre for Applied Macroeconomics and Commodity Prices, BI Norwegian Business School* (*E-mail*: leif.a.thorsrud@bi.no).

Journal of Money, Credit and Banking, Vol. 0, No. 0 (October 2021)

The majority of research to date on central bank communication has taken professional observers, such as financial market participants, as the principal receivers of such communication (Blinder et al., 2008). These agents have access to a large set of both hard and soft data when monitoring monetary policy actions. In contrast, and as alluded to by the quote above, the primary source of information for most people, that is, households, is the news media.[1] Thus, although we know a lot about the transmission mechanism between central bank communication and financial markets, less is known about how central bank communication reaches a wider audience through the media.

In this paper, we propose a simple method to extract and identify what we label "narrative monetary policy surprises" from textual, and high-dimensional media coverage and central bank communication data. We then apply this methodology to communication published by Norges Bank, via their Executive Board Assessments (*EBA*), and Norwegian media coverage to study how narrative surprises differ from conventional monetary policy surprises identified using financial market data.

To conceptually identify narrative monetary policy surprises, we informally build on the mechanism proposed in Nimark and Pitschner (2019) where no agent (household) has the resources to monitor all events that are potentially relevant for their decision, and thereby delegate their information choice to specialized news providers. That is, the media works as "information intermediaries" between agents and the state of the world. Following the cognitive psychologist Jerome Bruner, and in particular Bruner (1991), we couple this with a view on narratives as instruments of mind in the construction, or reflection, of reality. Accordingly, we treat media narratives prior to monetary policy announcements as a good proxy for the beliefs the public has about macroeconomic conditions and monetary policy, and assume that such beliefs might be affected by central bank communication that reach the public through the media if it is newsworthy. Thus, narrative monetary policy surprises are identified as the differences in narrative focus between the media and the central bank's communication, that is, a difference in the reflection of reality. An example might be when the media focuses heavily on, for example, (poor) labor market developments, while the central bank focuses almost solely on, for example, (high) inflation developments, when motivating their interest rate decision.

To measure these surprises, we adapt an event study framework and put structure on the problem by focusing on important narrative dimensions that feed into a central bank's decision-making process: inflation, labor market, and exchange rate developments, as well as narratives related to the oil market, financial stability, and uncertainty. As discussed in greater detail later, the method we propose allows us to identify these latent concepts from the different corpora (central bank statements and newspaper articles) using a bag-of-words assumption and a Singular Value Decomposition coupled with a unit rotation identification scheme. The method is fast,

1. See, for example, Blinder and Krueger (2004), Curtin (2007), and Fullone et al. (2007). Interestingly, Hayo and Neuenkirch (2015) find that even professionals such as financial market participants also rely heavily on media reporting when following central bank events.

simple, and requires minimal subjective judgment regarding the size and timing of narrative surprises.

We reach three main conclusions. First, using interest rate market data and the high-frequency identification framework pioneered by Gürkaynak et al. (2005) to identify what we label as conventional monetary policy surprises, we find that there is a weak and insignificant relationship between these surprise measures and the proposed narrative surprises. Hence, the narrative surprises identified here capture a different part of the central bank's communication than conventional monetary policy surprises do, and suggest that the media channel might be complementary to the transmission mechanism traditionally studied.

Second, by nature, the news media reports on newsworthy events, and, by assumption, households' beliefs are shaped by media coverage. Thus, a defining feature of a narrative monetary policy surprise should be that it affects subsequent media coverage, and it does. Our results show that the narrative monetary policy surprises lead to a significant change in media coverage after the interest rate meeting relative to before, whereas conventionally measured surprises do not.

Finally, we show that these discrepancies matter for economic outcomes. Following narrative monetary policy surprises the interest rate, the stock market, consumer confidence, house prices, and industrial production all increase. These response patterns are not in line with conventional monetary policy shock interpretations, but rather in accordance with what the newer monetary policy literature identifies as the information component of monetary policy surprises. In fact, a positive co-movement between the interest rate and the stock market following monetary policy surprises has been the defining identifying feature of "central bank information shocks" in, for example, Jarocinski and Karadi (2020) and Cieslak and Schrimpf (2019). The common interpretation for this information component is that the central bank, through its communication, reveals private information about its views on current and future economic conditions. Thus, the narrative monetary policy surprise is a natural candidate for an information component, both in terms of its estimated impulse responses, and especially in terms of its construction.

These results are important for at least two reasons. First, they suggest that the media, and how they act as information intermediaries, can have a sizable effect on economic outcomes. For central banks trying to manage public expectations, this highlights the role of their media communication strategies. Second, they provide evidence that identifying surprises only through the use of movements in financial market variables might not fully capture how the general public perceives a monetary policy surprise. Moreover, in contrast to existing studies identifying the information component of monetary policy, our methodology explicitly quantifies the different narrative contributions of the information component.

In terms of economics, this paper speaks to a large literature investigating central bank communication and the measurement and content of monetary policy surprises (see, e.g., Gürkaynak et al. 2005, Nakamura and Steinsson 2018, Miranda-Agrippino and Ricco 2020, Jarocinski and Karadi 2020, Andrade and Ferroni 2021). Although influential papers in this literature have focused mainly on hard quantitative information released by the central banks, a growing literature also use written communication like minutes, speeches, and monetary policy reports (see, e.g., Blinder et al. 2008, Hansen and McMahon 2016, Hansen et al. 2018, Ehrmann and Talmi 2020).

The studies most closely related to ours are those focusing on the media channel for central bank communication. Among others, these include Hendry (2012) and Hayo and Neuenkirch (2012) who document that Bank of Canada communication and its media coverage has significant effects on financial markets; Berger et al. (2011) who analyze how favorably the media reports about the European Central Bank's (ECB) monetary policy decisions and find that they are indeed responsive to efforts by the ECB to explain the motivation behind a given decision; Binder (2017) who documents how various U.S. central bank communication tools influence news coverage; Haldane and McMahon (2018) who conduct an experiment showing that there are benefits to clarifying and simplifying central bank messages even to traditional audiences.

We contribute to this literature by proposing a new method for identifying narrative monetary policy surprises, and by explicitly analyzing how media coverage evolves prior to relative to after important central bank communication events. In particular, in line with existing studies showing that households are not well informed about monetary policy per se (Coibion et al., 2019), that they might not use the media to obtain such information (Hayo and Neuenkirch, 2018), but that central bank communication makes it more likely for people to indirectly receive monetary policy related news (Lamla and Vinogradov, 2019), we focus on the narrative dimension of central bank communication, media coverage, and their interaction. This does not require the public to read news with the purpose of acquiring information about the central bank or its policies explicitly, but rather to get general information about, for example, inflation, labor market, and exchange rate developments. As such narratives are key ingredients in any interest rate decision and important for households (Larsen et al., 2020), central bank communication has the potential to affect media coverage, and thus expectations, even though households know little about monetary policy per se.

In terms of methodology, this paper relates to the Natural Language Processing (NLP) literature, and in particular the use of computational linguistics to uncover what the themes of documents are (e.g., Deerwester et al. 1990, Blei et al. 2003, Mcauliffe and Blei 2008, Taddy 2013, Le and Mikolov 2014, Kusner et al. 2015). Although this literature is vast, it is mostly applied in either unsupervised settings (without classified training data), or in supervised settings where the researcher has access to large amounts of already classified textual data to train models. In the current setting, as in many cases of interest to economists, such classified data are typically not available. Still, structure is desirable, making purely unsupervised methods unappealing. The method proposed here builds on the factor identification scheme proposed by Bai and Ng (2013), and extends this into the realm of textual analysis, permitting a structural analysis without access to large amounts of already classified data to train models.

The rest of the paper is organized as follows: Section 1 presents the data and methodology. Section 2 presents the results, while Section 3 concludes.

1. DATA AND METHODOLOGY

The methodology we propose is general, while the economic application uses data specific for Norway. In the following, we first describe the raw textual data sources, and then how these data are transformed into quantitative information and identified narrative surprises.

To measure the narrative focus of the central bank, we use communication published by Norges Bank via their *EBAs*. These documents are official statements of roughly two pages, released at the same time as the monetary policy decisions are made public, and serve as a justification for the decision being made (Qvigstad, 2019). Looking at Norway has the advantage that Norges Bank has a long history of being a relatively open and transparent central bank, both in terms of its written communication, but also in terms of releasing, for example, interest rate path predictions, which we explicitly use as hard economic control variables in the analysis later. Between 1999 (October 27) and 2019 (March 21) there have been 152 interest rate decisions, from which the associated *EBAs* are collected from Norges Bank's web pages.

To measure media coverage, we use the entire corpus, that is, all text and articles, published by *Dagens Næringsliv* (DN). This outlet is Norway's largest business newspaper, and is read by roughly 10% of the adult population. Economic news relevant for monetary policy is well covered by DN, and the newspaper publishes explicit stories about the expected outcome of monetary decisions in the days prior to such decisions, as well as commentaries about the actual decision in the days after. Although DN arguably contains only a subset of the news households in Norway consume, it is the fourth largest newspaper in Norway irrespective of subject matter. This makes it likely that its news coverage spills over to other news sources. For example, as documented in King et al. (2017), even news stories published in minor U.S. outlets causes Americans to take public stands on the articles' specific subjects, increasing the discussion in broad policy areas (topics) by roughly 63% relative to a day's normal volume. Moreover, Norwegians generally score high on socioeconomic characteristics known to correlate positively with news consumption (Shehata and Strömbäck, 2011). This is reflected in the fact that approximately 65% of the population reads printed news every day, and, compared to most other (European) countries, Norwegians spend relatively more time reading newspapers than watching television. The news data have been generously provided by the company *Retriever*, and collected manually by us for the latter part of the sample. In total these data consist of roughly 200,000 news articles between 1999 and 2019, and over 80,000 unique words and terms.[2]

2. In unreported results, we have experimented using a larger corpus, including news from other major Norwegian news outlets that do not necessarily specialize in business news reporting. The augmented corpus captures almost 70% of daily newspaper readership in Norway. Our main qualitative conclusions are not affected by these data augmentation. Still, we prefer using the DN corpus as our benchmark data set because the other news data sets only contain a subset of the news published by those sources and also only cover two-thirds of the original sample. We only use printed news, and leave it for future research to explore how changing media consumption habits potentially affects the effect of central bank communication.

1.1 Feature Selection

As is common in the NLP literature, the raw high-dimensional and unstructured textual data are cleaned before further analysis (Gentzkow et al., 2019). The independent feature selection (cleaning) steps taken below are common in most NLP applications, while their combined implementation here is context specific.

First, we define the relevant vocabulary as all the unique words used in the *EBAs*. Because this set of words is much smaller than the vocabulary used in the newspaper, it reduces the dimensionality of the problem and potentially limits the influence of newspaper content completely unrelated to the central bank's function, such as the sports or entertainment sections. The size of the vocabulary is denoted V. Because the newspaper content during weekends differs considerably from that published during business days, that is, featuring more background articles, travel, portrait interviews, etc., all weekends are removed from the news corpus.

Next, we take a bag-of-words view and construct two document term matrices, $\boldsymbol{C}^N$ and $\boldsymbol{C}^{CB}$, for the news media data (*N*) and Norges Bank's *EBAs* (*CB*), respectively. In these matrices, each column represents the unique terms in the vocabulary and each row a unique document. The matrix entries are the number of times term j occurs in document i. The $\boldsymbol{C}^{CB}$ matrix has dimension $T^{CB} \times V$, where $V = 2,716$ and $T^{CB} = 152$. Because there are many more news days than announcement days, the $\boldsymbol{C}^N$ matrix is much larger, and has dimension $T^N \times V$, where $T^N = 6,240$.

To construct a mapping between the information captured in the $\boldsymbol{C}^N$ matrix at event time t and that conveyed by the *EBAs* in $\boldsymbol{C}^{CB}$, the counts in the $\boldsymbol{C}^N$ matrix are summed over a period of w^- days prior to each announcement day t and averaged. Accordingly, smaller values of w^- will potentially capture media's short run focus just prior to the interest rate meeting, while larger values of w^- capture media's more general focus over that period. At the same time, larger values of w^- will incorporate information further away from the event day t, and, as such, challenge the event study identification strategy. For these reasons, and because we do not have any strong prior on what w^- should be, we consider all $w^- = 1, \ldots, 10$, and denote these matrices $\boldsymbol{C}^N_{w^-}$. Similarly, a $\boldsymbol{C}^N_{w^+}$ matrix is constructed, where the only difference between $\boldsymbol{C}^N_{w^+}$ and $\boldsymbol{C}^N_{w^-}$ is that we aggregate w^+ periods forward relative to the announcement day t when constructing $\boldsymbol{C}^N_{w^+}$. However, as the central bank actively engages in various communication strategies following interest rate meetings, we only consider $w^+ = 1, \ldots, 5$.[3] Note here that because of this mapping, all matrices $\boldsymbol{C}^{CB}$, $\boldsymbol{C}^N_{w^-}$, and $\boldsymbol{C}^N_{w^+}$ now have dimensions $T^{CB} \times V$.

3. Following interest rate meetings and the publication of Monetary Policy Reports, central bank officials regularly hold speeches, meet private banks, and give seminars. In the days prior to interest rate meetings, such communication activities are much less prominent. Because weekends are removed from the data set, $w^- = 10$ and $w^+ = 5$ correspond to two and one business week, respectively, and day t news coverage is excluded from the information set used to construct both $\boldsymbol{C}^N_{w^-}$ and $\boldsymbol{C}^N_{w^+}$. We have also experimented with using linearly decaying weights when aggregating the counts. This gives more weight to observations closer to the interest rate meeting date than to those further away, but does not materially affect the main results reported in Section 2.

As a final feature selection step, the different terms in the document term matrices are weighted by the inverse-document-frequency metric implied by the $\boldsymbol{C}^{CB}$ matrix. We do this to put a lower weight on terms the central bank is using frequently in all documents, and thus, a higher weight on terms that might be more representative for particular time periods. In essence, this also considerably downweights stop words. Formally, this is done by first normalizing the $\boldsymbol{C}$ matrices from above such that each matrix entry reflects the relative frequency of that term within each document. Then, the inverse-document-frequency score is computed, denoted $idf = \log(T/d_j^{CB})$ with $d_j^{CB} = \sum_i \mathbf{1}_{C_{ij}^{CB}>0}$, and $C_{ij}^{CB} \times idf_j = \hat{C}_{ij}^{CB}$ and $C_{ij,w}^{N} \times idf_j = \hat{C}_{ij,w}^{N}$.

1.2 Factor Extraction and Identification

Narratives are not captured by the terms in isolation, but rather by how different terms are used in context and together. To capture this, we apply factor modeling techniques to construct numerical approximations to the narratives conveyed in the texts (Larsen and Thorsrud, 2018). In the NLP literature, such factors are commonly referred to as topics, which summarizes the themes of documents in a parsimonious manner.

To estimate the factors, we use the traditional Latent Semantic Indexing (LSI; Deerwester et al. 1990) approach, introduced into the central bank communication literature by Boukus and Rosenberg (2006), coupled with an ex post unit identification step. In particular, while the standard LSI approach is simply a Singular Value Decomposition (SVD) of the document term matrix, we suggest to ex post rotate the estimated factor space such that the factors can be given a narrative interpretation along dimensions of interest.

In the current context, the dimensions of interest are narratives that typically feed into central banks', and in particular Norges Bank's (central bank of a small open economy with oil), decision-making process: inflation, labor market conditions, exchange rate developments, and narratives related to the oil market, financial stability, and uncertainty. Of these, the three former are motivated by a (extended) Taylor rule argument for a small open economy with flexible inflation targeting (Gali and Monacelli 2005, Svensson 2010), whereas the three latter are included to capture the importance of oil for the Norwegian economy (Bjørnland and Thorsrud, 2016), and the increased emphasis on financial stability (Svensson 2014, 2017, Gerdrup et al. 2017) and (political) uncertainty (Bernanke 2007, Bloom 2014, Larsen 2021) in monetary theory and practice.

Formally, factor estimation and identification proceeds as follows. First, define K as the total number of factors, and associate each factor with one (subjectively chosen) word representative for each narrative dimension, as illustrated in Table 1. Then, for a given $\hat{\boldsymbol{C}}$ matrix, order these K terms in the K first columns of the matrix and apply the SVD decomposition $\hat{\boldsymbol{C}} = \boldsymbol{U}\boldsymbol{S}\boldsymbol{V}'$ with factors $\boldsymbol{F} = \boldsymbol{U}_{1:K}\boldsymbol{S}_{1:K}$ $(T^{CB} \times K)$ and loadings $\boldsymbol{L} = \boldsymbol{V}_{1:K}$ $(V \times K)$ such that

$$\hat{\boldsymbol{C}} \approx \boldsymbol{F}\boldsymbol{L}'. \tag{1}$$

TABLE 1

Key Word List Used to Identify Factors

Narrative dimension	Inflation developments	Labor market	Exchange rate market	Oil market	Uncertainty	Financial stability
Key word	inflasjonen (inflation)	arbeidsledigheten (unemployment)	kronekursen (exchange rate)	oljeprisen (oil price)	usikkerheten (uncertainty)	kreditten (credit)

Note: In Norwegian, different terms are combined into one word more often than in English. Thus, to be as precise as possible, and to avoid being lost in translation, the words are listed in Norwegian, but with our English translation in parenthesis.

To identify the first factor with an inflation narrative, the second with the labor market, and so on, we rotate the factor space such that we get a so-called unit identification. To do this, partition $\boldsymbol{L}$ from equation (1) as

$$\boldsymbol{L} = \begin{bmatrix} \boldsymbol{L_0} \\ \boldsymbol{L_1} \end{bmatrix} \text{ with } \boldsymbol{L}_0 = \boldsymbol{L}_{1:K} \text{ and } \boldsymbol{L}_1 = \boldsymbol{L}_{K+1:V} \tag{2}$$

and apply the rotation:

$$\tilde{\boldsymbol{F}} = \boldsymbol{F}\boldsymbol{L}_0' \text{ and } \tilde{\boldsymbol{L}} = \boldsymbol{L}\boldsymbol{L}_0^{-1}, \tag{3}$$

where $\tilde{\boldsymbol{F}}$ and $\tilde{\boldsymbol{L}}$ are the identified factor and loading matrices, respectively, and the upper $K \times K$ block of $\tilde{\boldsymbol{L}}$ equals the identity matrix, that is, $\tilde{\boldsymbol{L}}_{1:K} = \boldsymbol{I}_K$. Accordingly, the inflation term loads with one on the first factor, and zero on all other factors, the unemployment term loads with one on the second factor, and zero on all other factors, etc. For this reason, the first factor is associated with an inflation narrative, the second factor with a labor market narrative, etc. Moreover, although the factors in equation (1) are uncorrelated by construction, which in an economic context were narratives interact seems unrealistic, the identified factors in equation (3) are not.

1.3 Constructing Narrative Surprises

To construct measures of the narrative monetary policy surprise and the changes in media coverage around central bank announcements, we proceed in two steps. First, equations (1)–(3) are implemented for each of the three matrices $\hat{\boldsymbol{C}}^{CB}$, $\hat{\boldsymbol{C}}^{N}_{w^-}$, and $\hat{\boldsymbol{C}}^{N}_{w^+}$ separately. Then, difference measures are constructed as

$$\tilde{nd}_t^{CB,N} = \sum_{k=1}^{K} (\tilde{F}_{k,t}^{CB} - \tilde{F}_{k,t:w^-}^{N})^2 \text{ and } \tilde{nd}_t^{N,N} = \sum_{k=1}^{K} (\tilde{F}_{k,w^+:t}^{N} - \tilde{F}_{k,t:w^-}^{N})^2, \tag{4}$$

where $\tilde{\boldsymbol{F}}_t^{CB}$, $\tilde{\boldsymbol{F}}_{t:w^-}^{N}$, and $\tilde{\boldsymbol{F}}_{t:w^+}^{N}$ are derived from the $\hat{\boldsymbol{C}}^{CB}$ $\hat{\boldsymbol{C}}_{w^-}^{N}$, and $\hat{\boldsymbol{C}}_{w^+}^{N}$ matrices, respectively.[4] Thus, $\tilde{nd}_t^{N,N}$ captures changes in media coverage around central bank announcements, while $\tilde{nd}_t^{CB,N}$ measures the narrative differences in central bank communication and news media coverage. Accordingly, large values of $\tilde{nd}_t^{CB,N}$ signal the extent to which the media focuses on different topics than the central bank does in its *EBAs*.

Second, to also capture potential differences in the tonality, that is, sentiment, of reporting, we weight the difference measures in equation (4) using a simple dictionary-based method. This step builds on Larsen and Thorsrud (2019), and is done using an external word list and simple word counts. The word list used here classifies positive/negative words as defined by a Norwegian translation of the *Harvard IV-4 Psychological Dictionary*.[5] For each event day t, the count procedure delivers a statistic containing the normalized difference between positive and negative terms associated with each row of $\hat{\boldsymbol{C}}^{CB}$, $\hat{\boldsymbol{C}}_{w^-}^{N}$, and $\hat{\boldsymbol{C}}_{w^+}^{N}$. For example, $to_t^{CB} =$ (#positive terms − #negative terms) in the tth row of $\hat{\boldsymbol{C}}^{CB}$. These statistics are then normalized across time, denoted $\bar{to}_t^{CB}$, to avoid picking up systematic differences in the use of positive versus negative terms across sources. The tonality differences are computed as

$$to_t^{CB,N} = (\bar{to}_t^{CB} - \bar{to}_{t:w^-}^{N}) \text{ and } to_t^{N,N} = (\bar{to}_{w^+:t}^{N} - \bar{to}_{t:w^-}^{N}) \tag{5}$$

and the to statistics are used to weight (or sign-adjust) the measures computed in equation (4) as

$$n_t = \tilde{nd}_t^{CB,N} to_t^{CB,N} \text{ and } m_t = \tilde{nd}_t^{N,N} to_t^{N,N}, \tag{6}$$

where n_t is what we label the narrative monetary policy surprise and m_t is the change in media coverage around central bank announcements.

1.4 Methodological Discussion

We highlight four points about the methodology. First, using time series data, factor identification like in equations (2) and (3) was first suggested by Bai and Ng (2013). They show that the unit identification scheme yields a unique solution both in terms of the sign and size of the latent factors, and the method is now commonly applied in the

4. By constructing the factors from separate matrices, the exact language and context in which the central bank and the media write about the different terms (used to identify the factors) can differ on average. Instead, the time-variation in the factors is used to identify the surprise component. To simplify the notation, the w subscripts are not included when defining $\tilde{nd}_t^{CB,N}$ and $\tilde{nd}_t^{N,N}$.

5. In the literature, also other English-based word lists have been suggested (see, e.g., Loughran and Mcdonald 2016). For applications using Norwegian language, it is our experience that the exact word list used plays a minor role, and that our Norwegian translation of the *Harvard IV-4* dictionary works well across a range of applications (Larsen and Thorsrud 2017, Thorsrud 2018).

time series literature (Aastveit et al. 2015, Bjørnland and Thorsrud 2016, Stock and Watson 2016). Still, to the best of our knowledge, it has not been applied or suggested in the NLP literature before. In this literature, the usage of supervised Latent Dirichlet Allocation (LDA; Mcauliffe and Blei 2008) models are perhaps more common, but not necessarily better (Kusner et al., 2015), when wanting to extract identified topics. However, a supervised LDA implementation requires a classified data set with identified factors (topics) in the texts prior to training the model. In many macroeconomic applications, including this one, this is not feasible because textual data in the form of *EBAs* are too scarce to appropriately divide the sample into informative training and testing sets.

Second, although the type of factor identification described above could potentially have been achieved simpler by using the counts in the document term matrices associated with the chosen key words (in Table 1) directly, such an approach has several drawbacks. Conceptually, as alluded to above, narratives are not captured by the terms in isolation, but rather by how different terms are used in context and together. Moreover, as described in, for example, Bholat et al. (2015), simple count-based methods cannot handle issues related to synonyms and polysemy, while factor-based methods can. In particular, because a term (not used to identify the factors) potentially loads on all the factors (which represent different contexts), the factor-based approach internalizes that the same word can be used in different contexts (polysemy). Likewise, terms that are similar (synonyms), and used in the same context(s), would likely have very similar factor loadings. In practice, this latter feature also makes the methodology described above relatively robust to changing the exact terms used to identify the factors. We formally show this in Section 2.5.

Third, although related in spirit to narrative identification used in some other macroeconomic applications, the approach taken here differs along several dimensions. For example, in their highly influential work, Romer and Romer (1989, 2004) perform a manual audit of the minutes of the Federal Open Market Committee, made public with a 5-year delay, to single out events they argue represent monetary policy shocks. Similar approaches have since then been applied in the oil market literature (Hamilton 1985, Kilian 2008) and to identify fiscal shocks (Ramey 2011, Mertens and Ravn 2014). In contrast to these approaches, however, the methodology suggested here is more data-driven and automated, and focuses on media's role as information intermediaries. More recently, Zeev (2018) and Antolín-Díaz and Rubio-Ramírez (2018) have suggested to use narrative sign restrictions around key historical events to ensure that the identified shocks agree with the established narrative account of these episodes. Although more data-driven and automated than the manual audit approach, narrative sign restrictions still require the researcher to take a strong stand on both when (timing) and how (sign) historical shocks unfolded.

Finally, although the NLP literature has come a long way in terms of classifying the sentiment, or tonality, of written text (Pang et al. 2002, Taboada et al. 2011, Howard and Ruder 2018, Merity et al. 2018), doing so is still very much a supervised machine learning problem. Accordingly, for the same reasons as discussed earlier, with limited amount of training data available, alternative approaches are needed. The dictionary-based approach adapted here is simple (and naive), but well suited in that respect. However, to identify the difference in tonality for specific narratives, for example, with respect to *inflation*, and not only the overall contribution, as in equation (5), our approach falls short. We leave it for future research to design approaches that can also identify the tonality of the individual narrative components.[6]

2. RESULTS

In the following, we first present the estimated factors and the narrative monetary policy surprises. We then study how these surprises differ from conventional monetary policy surprises and how they affect media coverage and macroeconomic developments.

2.1 Factors and Narrative Differences

Figure 1(a)–(f) reports, in gray and dotted black, the factors estimated from the news media data set ($\tilde{\boldsymbol{F}}^{N}_{t:w^-}$) as well as key macroeconomic aggregates for the Norwegian economy. The frequency on the x-axis is event dates and the quarterly observable variables are mapped to this frequency using the first event date in the quarter. Overall, the factors seem to capture reasonably well the conventional narrative held about economic developments the last two decades. For example, the estimates suggest that the media focused more on unemployment related issues around 2003, 2009, and 2015, which are periods associated with downturns, or recessions, in the Norwegian economy. Likewise, the enhanced focus on exchange rates and inflation during the earlier parts of the sample, relative to the latter part, is natural given that Norges Bank formally went from a fixed exchange rate regime to inflation targeting in 2001. The particular peak in the exchange rate factor around 2003 is likely due to the broad discussion of the changing market for global trade and its impact on Norway at the time (Bjørnland et al. 2004). We further observe that focus on credit conditions and uncertainty was high during the financial crisis, and that the oil market got a lot of attention in the mid 2000s when this sector was an engine for growth (Bjørnland and Thorsrud 2016), as well as between 2014 and 2016, when low oil prices led to concerns about future growth prospects. Although there is some high-frequency variation, it is also noteworthy that these broad patterns seem to be relatively robust to the choice of w, that is, the news aggregation window.

6. A related concern can be raised with respect to equation (6), where the potential case $\tilde{nd}_t^{CB,N} = 0$ (or $\tilde{nd}_t^{N,N} = 0$), that is, perfectly equal narrative focus, yields the unrealistic result $n_t = 0$ (or $m_t = 0$), irrespective of any differences in tonality. As a response to this, we show in Section 2.5 that our main results are robust to working with the unsigned narrative differences and hence our results are not driven by the tone adjustment (weighting). Still, we prefer the tone-adjusted difference measures as a benchmark specification because it allows computing meaningful impulse response functions in Section 2.4.

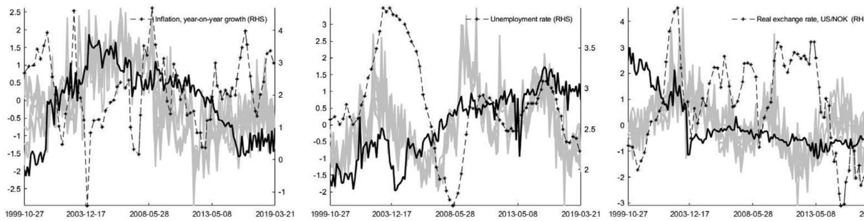

(a) Inflation **(b)** Unemployment **(c)** Exchange rate

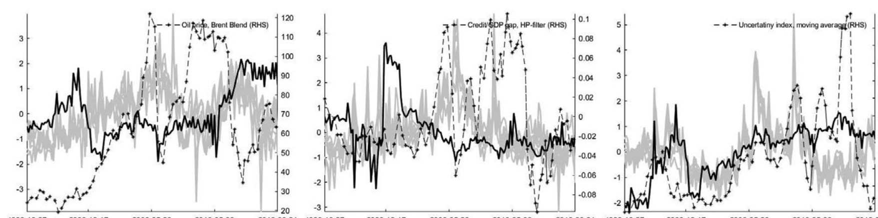

(d) Oil **(e)** Credit **(f)** Uncertainty

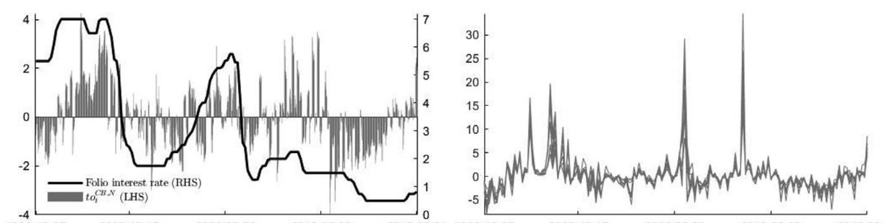

(g) Folio interest rate and sentiment ($to_t^{CB,N}$) **(h)** Narrative monetary policy surprises (n_t)

Fig 1. Identified Factors and the Narrative Monetary Policy Surprise. (a) Inflation. (b) Unemployment. (c) Exchange rate. (d) Oil. (e) Credit. (f) Uncertainty.

NOTES: In Figure 1(a)–1(f), the solid black lines illustrate the evolution of the narrative focus in the *EBAs* ($\tilde{\boldsymbol{F}}_t^{CB}$), while the broken gray lines illustrate the evolution of the narrative focus in the media ($\tilde{\boldsymbol{F}}_{t:w^-}^{N}$) for different values of w^-. The dotted black lines report, on the right-hand side y-axis (RHS), the evolution of key macroeconomic aggregates for Norway. All topics are normalized (mean of zero, and standard deviation of one). Announcement dates are reported on the horizontal axis. See Figure A1a, in the online Appendix, for an illustration of m_t.

To get an alternative impression of the contexts, the different terms used to identify the factors represent, Figure 2 reports word clouds constructed based on the cosine similarity between the word vector for key word $k = 1, \ldots, K$ and term $j = 1, \ldots, V$ in $\tilde{L}^N$. In the figure, a larger font represents a higher degree of similarity. Naturally, each key word vector has the biggest similarity with itself. However, as seen in the figure, inflation is typically written about in the media in the same context as, for example, energy prices, the inflation report, and Asia. Unemployment, on the other hand, is typically talked about in the context of recessions, the outlook, and the labor market. Similar information can be extracted from the other word clouds. In short,

Fig 2. For each Narrative, the Word Cloud is Constructed Based on the Cosine Similarity between the Identifying Word Vector in the $\tilde{L}^N$ Loadings Matrix, and All Other Word Vectors in that Matrix.

NOTES: A larger font indicates a larger similarity. For visual clarity, only the 50 most similar terms are reported. In Norwegian, different terms are combined into one word more often than in English. In the translation used for the graph, an underscore is used to illustrate such cases.

the results align well with the results reported for the factors themselves, and suggest that the method presented in Section 1.2 is able to extract meaningful narrative information from the textual data.

The narrative central bank factors ($\tilde{\boldsymbol{F}}_t^{CB}$) are reported in black in Figure 1(a)–(f). For the inflation factor, the low-frequency patterns seem to be relatively similar to those estimates for the news media. Moreover, for the oil-related narrative, the two sources seem to be sharing an upwardly drifting trend starting around 2014 when oil prices fell sharply. For the other factors, and for other time periods, the differences between the two sources are sometimes large. For example, during the financial crisis period, the differences in narrative focus related to oil, credit, and uncertainty is substantial.

Figure 1(g) reports the tonality contribution ($to_t^{CB,N}$) together with the actual key policy rate set by the central bank. As seen from the figure, there is a clear correlation between the two: when the interest rate increases or is high, the tonality of the central bank *EBAs* tend to be more positive than the media, and vice versa. Although our approach for identifying the difference in sentiment between the different sources is simple, we conclude that it at least captures important features of the actual monetary policy rate.

Finally, the narrative monetary policy surprise (n_t) is reported in Figure 1(h). Three time periods stand out as particularly striking, namely the late 1990s and early 2000s, 2008/2009, and 2011/2012. As discussed above, the former period was associated with large terms of trade effects and the early years of inflation targeting in Norway, and 2008/2009 and 2011/2012 capture the financial crisis and the European debt crisis, respectively. Large and potentially sudden changes in the economic or political landscape might be associated with either divergence or convergence in beliefs. The results reported in Figure 1(h) suggests that the narrative monetary policy surprise identified here peaks around at least some noticeable events.

2.2 Correlations

To gauge whether the proposed narrative surprises capture something different than what conventional monetary policy surprises do, we start by estimating:

$$s_t = b_1 n_t + \boldsymbol{b}_2 \boldsymbol{z}_{1t} + e_t, \tag{7}$$

where s_t is a conventionally measured monetary policy surprise at event day t, and n_t is the narrative monetary policy surprise. For completeness, equation (7) is estimated for all w^- versions of n_t.

We identify s_t using the high-frequency identification framework pioneered by Gürkaynak et al. (2005), and extract movements in interest rates around the monetary policy announcement time on day t. The way this is done using Norwegian data is described in Brubakk et al. (2017).[7] The methodology decomposes the surprise into two components, namely a "target" (T) and "path" (P) component, where the former is seen as a response to the actions of a central bank, while the latter is thought of as capturing unexpected central bank communication and unconventional policy. Going forward, we label these s_t, and s_t^T and s_t^P when the difference is relevant. In the interest of preserving space, the s_t surprises are graphed in Figure A1b in the Online Appendix.

First, the unconditional correlation between s_t and n_t is low. In fact, depending on w^-, the correlation between s_t^T and n_t is in the range -0.01 to 0.10, while the correlation between s_t^P and n_t is the range -0.12 to -0.03. Still, comparing the narrative surprise in Figure 1(h) to the conventional ones in Figure A1b, in the Online Appendix, we observe that the series share increased volatility patterns around the three time periods discussed above.

To control for changing macroeconomic conditions between announcement dates and other quantitative information that potentially explains monetary policy surprises, we include in the vector $\boldsymbol{z}_{1t}$ revisions in forecasts published by the central bank at the interest rate announcement time. The vector includes revisions to both GDP and inflation projections as well as revisions to the interest rate path for the current quarter and up until two quarters ahead. All variables are collected from Norges Bank's *Monetary Policy Report*.[8] In addition, to the extent that important international developments are not already incorporated into these projections (Bjørnland et al. 2020), we also control for the high-frequency euro area monetary policy surprises constructed by Altavilla et al. (2019) (using the 1 week and 1 year surprises from the press release window). Thus, the vector $\boldsymbol{z}_{1t}$ contains 11 elements.

7. In short, the event window is 90 min and captures the change in interest rates between 15 min before the announcement and 75 min after the meeting. This includes both the actual announcement time, as well as the press conference. Brubakk et al. (2017) show that the target factor is robust to event window size, and that the path factor is robust for event windows between 90 min and 1 day.

8. Norges Bank was one of the first central banks to publish its own interest rate path, starting in 2005. We look at revision to the projections, and not their level, to capture the new information in the projections. Only roughly every other interest rate meeting is accompanied by a publication of updated projections. For meeting dates where there are no updated projections, we fill in with zeros in the $\boldsymbol{z}_t$ vector.

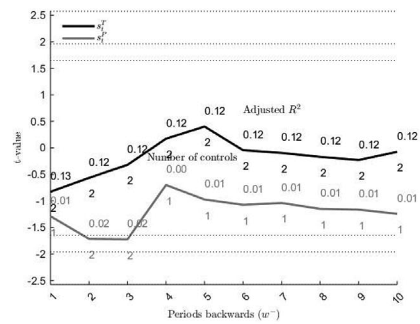

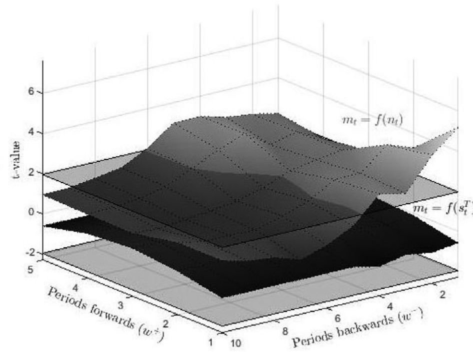

(a) Conventional and narrative surprises **(b)** Media spillovers

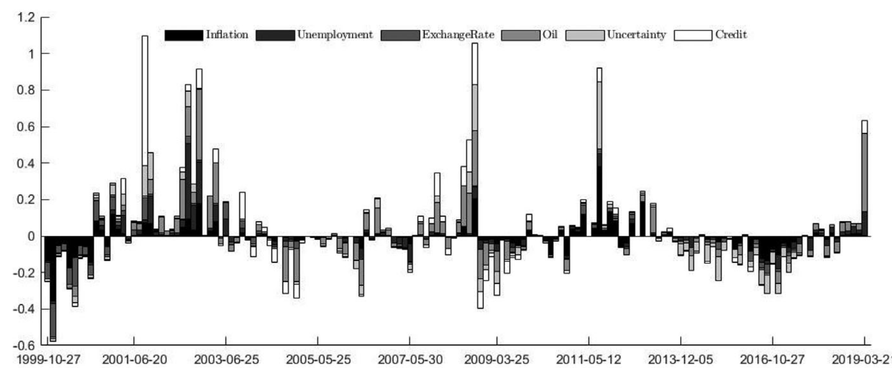

(c) Media spillovers and narrative contributions

Fig 3. Figure 3(a) Reports the t-values of $\hat{b}_1$ in equation (7).

NOTES: Numbers reported above and below the lines are the adjusted R^2 statistics and the number of chosen control variables in each regression, respectively. The x-axis reports the aggregation window w. Figure 3(b) reports the t-values of $\hat{\delta}_1$ in equation (8) when either n_t or s_t^T is used as the treatment variable. Figure 3(c) shows the scaled narrative monetary policy surprise, that is, $\hat{\delta}_1 n_t$, decomposed into narrative contributions.

To favor a small model size, and reduce noise and potential biases, a double selection procedure for selecting the relevant control variables in z_t is implemented (Belloni et al. 2014). First, the n_t and s_t variables are regressed separately on all the variables in the z_t vector using the LASSO estimator (Tibshirani 1996). Next, after these two penalized regressions, we estimate equation (7) using n_t and the union of the control variables selected in step one.[9]

Figure 3(a) reports the t-value associated with b_1 in equation (7), for all values of w^-. Numbers reported above and below the lines are the adjusted R^2 statistics and

9. In the first step, LASSO 100 different penalization parameters together with the BIC are used to tune the amount of regularization. As reported in Figure A2a, in the Online Appendix, our results are robust to estimating equation (7) using simple OLS.

the number of chosen control variables in each regression, respectively. One feature stands out; Irrespective of whether conventional monetary policy shocks are measured using s_t^T or s_t^P, their correlation with n_t is weak and insignificant. The fact that few control variables are selected also suggest that the variables in the z_t vector are relatively uninformative about both n_t and s_t.

In sum, we conclude that the narrative surprises clearly capture something different than conventional monetary policy shocks do. Thus, the media channel might be complementary to the transmission mechanism traditionally studied.

2.3 Impact on Media Coverage

Under the maintained assumption that the media works as information intermediaries between agents and the state of the world, a defining feature of a narrative monetary policy surprise should be that it affects subsequent media coverage. After all, for the narrative surprise to matter, people need to learn about it. We address this hypothesis by estimating:

$$m_t = \delta_1 n_t + \boldsymbol{\delta}_2 \boldsymbol{z}_{2t} + u_t, \tag{8}$$

where m_t is the overall difference in narrative focus between news media coverage prior to, relative to w^+ days after, the interest rate announcement, and δ_1 measures whether the narrative monetary policy surprise affects media coverage.

The upper plane in Figure 3(b) reports our estimate of $\hat{\delta}_1$ when this equation is estimated with the double selection procedure described above, and for all the indicated combinations of w^- and w^+. Here, the control vector $\boldsymbol{z}_{2t}$ includes $\boldsymbol{z}_{1t}$, as well as s_t^T and s_t^P. We observe that the narrative surprises have a positive and highly significant effect on the change in media coverage when we construct n_t and m_t using small values of w. For $w^- = 1$ and $w^+ = 1$, the adjusted R^2 statistic is roughly 16%. For larger window sizes, the R^2 statistic rapidly falls towards the range 4 to 5 percent.

Still, these results stand in sharp contrast to what we obtain if we instead replace n_t with s_t^T in equation (8), and re-do the double selection estimation routine. As seen from the lower plane in Figure 3(b), the conventional monetary policy shock (s_t^T) has an insignificant effect on media coverage. We have also done this analysis using s_t^P instead of s_t^T, finding similar insignificant results. As such, to the extent that households follow the news, the narrative differences contain information they will receive. Conventionally measured monetary policy surprises, on the other hand, seem to be more "silent."

Figure 3(c) reports a bar plot of the scaled narrative monetary policy surprise, that is, $\hat{\delta}_1 n_t$, for each event day in the sample and using $w^- = w^+ = 1$. The figure highlights an additional advantage with our narrative methodology, relative to conventional identification strategies, namely that one can decompose the surprise component into what it is about. In particular, as n_t is a linear combination of the different narrative contributions, one can decompose the regression results into the contributions from each narrative. In that respect, the last oil-driven observation in the sample is illuminating. Although Norges Bank increased their policy rate by 25 basis point in March 2019, motivated in parts by what they described in their *EBA* as positive petroleum-related activity, international business cycles were cooling down and foreign central banks were cutting rates. In a DN summary published the day after the meeting, and written by one of DN's main commentators in response to the central bank communication, the title was "There are good reasons to not be worried about too high interest rates", following up by stating "The Norwegian economy performs well due to better times in the oil and gas industry."

2.4 Macroeconomic Implications

Together, the results presented in Figure 3 suggest that the narrative surprises contain information not already present in existing surprise measures, and that this type of information has an effect on media coverage. Do these differences matter for macroeconomic outcomes? Figure 4 answers this question, and reports estimates of ϕ_1 from simple linear projections (Jordá 2005):

$$y_{t+h} = \phi_1 n_t + \boldsymbol{\phi}_2 \mathbf{z}_{3t} + \epsilon_{t+h}, \tag{9}$$

where y_{t+h} is the cumulative change in a monthly macroeconomic outcome variable, measured h periods forward relative to month t, n_t is the monthly aggregation of the event date specific narrative monetary policy surprise, and the control vector $\mathbf{z}_{3t}$ includes up to 12 lags of the dependent variable as well as a linear trend.[10]

We consider $h = 0, \ldots, 24$, and six important monthly macroeconomic aggregates (y_{t+h}): the 3-month interest rate, the stock market, house prices, consumer confidence, industrial production, and consumer prices. In the figure, we also include results from estimating equation (9) with monthly aggregations of s_t^T instead of n_t. In all cases, the shocks are normalized to a one standard deviation innovation, and 95% confidence bands as well as the mean responses are reported. Two main findings stand out.

First, following a narrative monetary policy surprise, close to all macroeconomic aggregates increase. The response paths of the interest rate, the stock market, consumer confidence, and industrial production are also significantly different from zero (at least on some horizons). In contrast, a conventional monetary policy surprise leads to an increase in the interest rate, but a decrease in returns, house prices, consumer confidence, and industrial production, as one would expect.

Second, with the exception of house prices, the narrative monetary policy surprise explains a much larger degree of the forecast error variance decomposition in the variables than the conventional monetary policy shock does. For example, up to 37% of the variation in the stock market can be explained by the narrative monetary policy

10. The lags are included to control for potential auto-correlation in the estimated residuals and the lag length is selected using the BIC. All dependent variables are (log) differenced prior to estimation, while the linear trend is included in the specification to control for potential slow moving drifts in the series. In unreported results, we confirm that excluding the trend from $\mathbf{z}_{3t}$, as well as including additional macroeconomic control variables in the $\mathbf{z}_{3t}$ vector, does not affect our main conclusions.

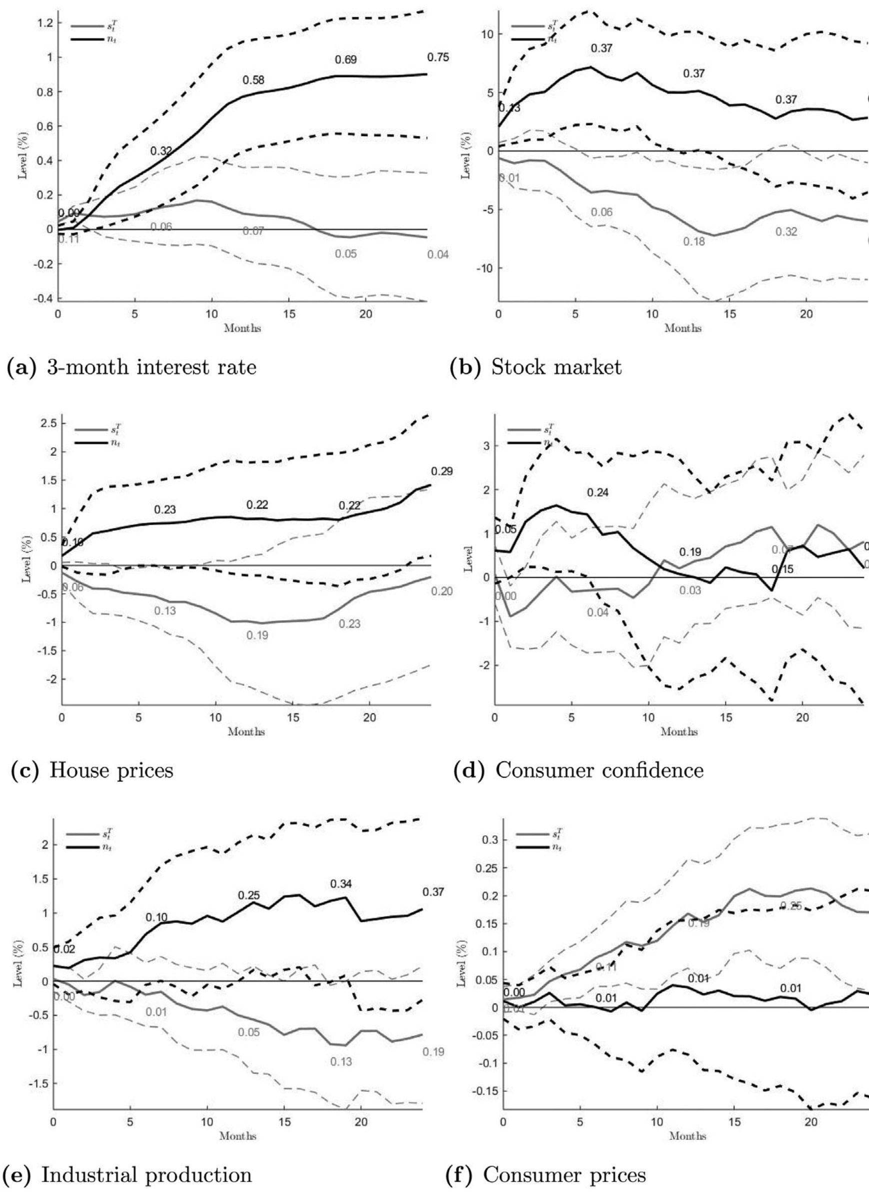

(a) 3-month interest rate

(b) Stock market

(c) House prices

(d) Consumer confidence

(e) Industrial production

(f) Consumer prices

Fig 4. The Figures Report the Estimates of $\hat{\phi}^h$ from Equation (9) for $h = 0, \ldots, 24$ months.

Notes: The mean estimates and 95% confidence bands are reported using Newey and West (1987) corrected standard errors. The responses are normalized to one standard deviation of the original shock, and to increase the 3-month interest rate on impact. Numbers reported along the curves are variance decompositions, computed as $v^h = (\sum_{i=0}^{h} \phi_1^i)\sigma_{a_t}^2 / ((\sum_{i=0}^{h} \phi_1^i)\sigma_{a_t}^2 + \sigma_{\epsilon_{t+h}}^2)$, where a_t equals either n_t or s_t^T.

surprise on the 5 months horizon, while the conventional monetary policy shocks explains only roughly 6% at the same horizon.

The differences in macroeconomic outcomes between a conventional monetary policy surprise and the narrative surprises speak directly to newer monetary studies emphasizing the information component of monetary policy surprises (see, e.g., Nakamura and Steinsson 2018, Cieslak and Schrimpf 2019, Jarocinski and Karadi 2020, Andrade and Ferroni 2021). In fact, our results are qualitatively in line with the macroeconomic responses obtained in Brubakk et al. (2019), who analyze the effects of information shocks using Norwegian data and a modified version of the methodology developed by Jarocinski and Karadi (2020).[11]

The common interpretation for this information component is simple: Through its communication, the central bank reveals private information about its views on current and future economic conditions. Under the assumption that central bank communication affects the market, a release of positive (negative) information should then, all else equal, increase (decrease) returns, interest rates, and the general economic outlook. As such, the narrative monetary policy surprise is a natural candidate for an information component, both in terms of its estimated impulse responses, and especially in terms of its construction. In contrast to other ways of identifying this monetary policy information component, however, the methodology suggested here allows the researcher to decipher what the information is mostly about, and highlights the role of the media as information intermediaries (Nimark and Pitschner 2019, Larsen et al. 2020).

2.5 Additional Results and Robustness

To the extent that financial market participants and journalists follow the same central bank communication, the lack of correlation between the narrative monetary policy surprises (n_t) and those identified through movements in the interest rate market (s_t), might seem surprising. However, focusing on the absolute size of the surprises, and disregarding their sign, we obtain a more significant link (Figure A3a in the Online Appendix). In particular, using equation (7) and regressing $\tilde{nd}_t^{CB,N}$ from equation (4), that is, the unsigned n_t measure, on the absolute value of the conventional surprise measures ($|s_t|$), we obtain a positive and mostly significant relationship. Accordingly, in terms of timing, but not in terms of sign, agents in the interest rate market and the media share surprise patterns. Still, using $\tilde{nd}_t^{N,N}$ from equation (4), that is, the unsigned m_t measure, as the dependent variable, and $|s_t|$ as the treatment variable in equation (8), we obtain more or less the same insignificant result

11. See also Bjørnland et al. (2020) for evidence regarding the information component of (Norwegian) monetary policy surprises. On a similar note, we have also estimated a Vector Autoregressive model (VAR), including the six macroeconomic variables discussed above, and identified monetary policy shocks using the narrative monetary policy shock as an external instrument following the methodology introduced by Stock (2008) and used in, for example, Gertler and Karadi (2015). Controlling for the dynamic correlation between the macroeconomic aggregates changes the size and evolution of the impulse response functions somewhat, but does not change our qualitative conclusions.

as before (Figure A3b in the Online Appendix). In contrast, $\tilde{nd}_t^{CB,N}$ has a positive and significant effect on $\tilde{nd}_t^{N,N}$, confirming that also this unweighted narrative monetary policy surprise measure affects media coverage, whereas conventional monetary policy surprises do not.

In terms of the feature selection steps described in Section 1.1, and as commented by one referee, the *idf* scaling might lead to a misinterpretation of the narrative shocks as communication mistakes rather than surprises in the traditional sense and the removal of weekends from the corpus might create a bias if it leads to an interruption of the continuous flow of information. However, our main results in Sections 2.2 and 2.3 remain virtually unchanged if we skip the *idf* scaling step (Figure A4 in the Online Appendix). We still use and describe it here because of the favorable properties it might have in other applications. Likewise, our identified narrative monetary policy shock remains close to unchanged if we do not remove weekends from the news corpus (Figure A5 in the Online Appendix).

Table A1, in the Online Appendix, shows how the factor-based identification scheme proposed here is relatively robust to changing the key words used to identify the factors. Cycling through 30 unique alternative combinations of key words, listed in Table A1, and re-doing the calculations of $\tilde{nd}_t^{CB,N}$ from equation (4), reveals that the correlation between these alternative estimates, and our benchmark estimate, is seldom lower than 0.40, very often above 0.70, and sometimes as high as 0.90. We also confirm that the macroeconomic effects of narrative monetary policy surprises remain robust to these changes (Figure A6 in the Online Appendix). In contrast, if we instead compute the factors as simple counts, as discussed in Section 1.4, we observe that the resulting narrative surprise would have been much more sensitive to the exact key words used to identify the narrative dimensions (Table A1).

In terms of conventional monetary surprises, one might argue that it is the "path" factor (s_t^P), rather than the "target" factor (s_t^T), that captures central bank communication and hence should be more similar to the narrative surprises in terms of macroeconomic responses. Re-estimating equation (9) using s_t^P as the shock of interest shows that this is not the case. The macroeconomic responses following a s_t^P shock resembles those following a s_t^T shock (Figure A7 in the Online Appendix).

3. CONCLUSION

Although a large literature has provided knowledge about the transmission mechanism between central bank communication and financial markets, less is known about how such communication is transmitted to a wider audience, that is, households and nonprofessional market observers.

In this paper, we take the view that the news media works as information intermediaries between agents (households) and the state of the world, and propose a method for identifying what we label as "narrative monetary policy surprises." Using textual data from news media and central bank communication, we put structure on the problem by focusing on narrative dimensions that typically feed into a central bank's decision-making process and identify these from the different corpora by applying a Singular Value Decomposition and an ex post unit rotation identification scheme.

The empirical evaluation, using data from Norwegian news media and Norges Bank, shows that the narrative monetary policy surprises have an insignificant correlation with conventionally measured monetary policy surprises, suggesting that the media channel might be complementary to the transmission mechanism traditionally studied. We further document that the narrative surprises in central bank communication lead to a significant change in media coverage after the interest rate meeting relative to before, whereas monetary policy surprises identified using conventional methods do not. In turn, these differences are shown to matter for the evolution of macroeconomic aggregates, where narrative surprises lead to response patterns in line with what newer monetary policy studies label the information component of monetary policy.

Our study highlights the importance of written central bank communication and the role of the media as information intermediaries, at least for the case of Norway. The proposed method is fast, simple, automated, and language-agnostic. It is particularly useful in the current context, where access to large amounts of classified training data makes more sophisticated supervised algorithms less suited.

LITERATURE CITED

Aastveit, Knut Are, Hilde C. Bjørnland, and Leif Anders Thorsrud. (2015) "What Drives Oil Prices? Emerging Versus Developed Economies." *Journal of Applied Econometrics*, 30, 1013–28.

Altavilla, Carlo, Luca Brugnolini, Refet S. Gürkaynak, Roberto Motto and Giuseppe Ragusa. (2019) "Measuring Euro Area Monetary Policy." *Journal of Monetary Economics*, 108, 162–79.

Andrade, Philippe, and Filippo Ferroni. (2021) "Delphic and Odyssean Monetary Policy Shocks: Evidence from the Euro Area." *Journal of Monetary Economics*, 117, 816–32.

Antolín-Díaz, Juan, and Juan F. Rubio-Ramírez. (2018) "Narrative Sign Restrictions for SVARs." *American Economic Review*, 108, 2802–29.

Bai, Jushan, and Serena Ng. (2013) "Principal Components Estimation and Identification of Static Factors." *Journal of Econometrics*, 176, 18–29.

Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen. (2014) "High-Dimensional Methods and Inference on Structural and Treatment Effects." *Journal of Economic Perspectives*, 28, 29–50.

Berger, Helge, Michael Ehrmann, and Marcel Fratzscher. (2011) "Monetary Policy in the Media." *Journal of Money, Credit and Banking*, 43, 689–709.

Bernanke, Ben S. (2007) "Inflation Expectations and Inflation Forecasting." Speech 306, Board of Governors of the Federal Reserve System (U.S.).

Bholat, David, Stephen Hansen, Pedro Santos, and Cheryl Schonhardt-Bailey. (2015) "Text Mining for Central Banks: Handbook." *Centre for Central Banking Studies*, 33, 1–19.

Binder, Carola. (2017) "Federal Reserve Communication and the Media." *Journal of Media Economics*, 30, 191–214.

Bjørnland, Hilde C., and Leif Anders Thorsrud. (2016) "Boom or Gloom? Examining the Dutch Disease in Two-Speed Economies." *Economic Journal*, 126, 2219–56.

Bjørnland, Hilde C., Thomas Ekeli, Petra M. Geraats, and Kai Leitemo. (2004) "Norges Bank Watch 2004." Technical Report 5, Centre for Monetary Economics BI Norwegian Business School.

Bjørnland, Hilde C., Leif Anders Thorsrud, and Sepideh K. Zahiri. (2020) "Do Central Banks Respond Timely to Developments in the Global Economy?." *Oxford Bulletin of Economics and Statistics*, 82, 285–310.

Blei, David M., Andrew Y. Ng, and Michael I. Jordan. (2003) "Latent Dirichlet Allocation." *Journal of Machine Learning Research*, 3, 993–1022.

Blinder, Alan S., and Alan B. Krueger. (2004) "What Does the Public Know about Economic Policy, and How Does It Know It?" NBER Working Paper No. 10787.

Blinder, Alan S., Michael Ehrmann, Marcel Fratzscher, Jakob De Haan, and David-Jan Jansen. (2008) "Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence." *Journal of Economic Literature*, 46, 910–45.

Bloom, Nicholas. (2014) "Fluctuations in Uncertainty." *Journal of Economic Perspectives*, 28, 153–76.

Boukus, Ellyn, and Joshua V. Rosenberg. (2006) "The Information Content of FOMC Minutes."

Brubakk, Leif, Saskia ter Ellen, and Hong Xu. (2017) "Forward Guidance through Interest Rate Projections: Does it Work?" Norges Bank Working Paper.

Brubakk, Leif, Saskia ter Ellen, Ørjan Robstad and Hong Xu. (2019) "The Macroeconomic Effects of Forward Communication." Working Paper.

Bruner, Jerome. (1991) "The Narrative Construction of Reality." *Critical inquiry*, 18, 1–21.

Cieslak, Anna, and Andreas Schrimpf. (2019) "Non-Monetary News in Central Bank Communication." *Journal of International Economics*, 118, 293–315.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber. (2019) "Monetary Policy Communications and their Effects on Household Inflation Expectations." NBER Technical Report.

Curtin, Richard. (2007) *What US Consumers Know About Economic Conditions*. Paris, France: OECD.

Deerwester, Scott, Susan T. Dumais, George W. Furnas, Thomas K. Landauer, and Richard Harshman. (1990) "Indexing by Latent Semantic Analysis." *Journal of the American Society for Information Science*, 41, 391–407.

Ehrmann, Michael, and Jonathan Talmi. (2020) "Starting from a Blank Page? Semantic Similarity in Central Bank Communication and Market Volatility." *Journal of Monetary Economics*, 111, 48–62.

Fullone, Flora, Gamba, Michela, Giovannini, Enrico, and Malgarini, Marco. (2007) *What Do Citizens Know about Statistics*. Paris, France: OECD.

Gali, Jordi, and Tommaso Monacelli. (2005) "Monetary Policy and Exchange Rate Volatility in a Small Open Economy." *Review of Economic Studies*, 72, 707–34.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy. (2019) "Text as Data." *Journal of Economic Literature*, 57, 535–74.

Gerdrup, Karsten, Frank Hansen, Tord Krogh and Junior Maih. (2017) "Leaning Against the Wind When Credit Bites Back." *International Journal of Central Banking*, 13(3): 287–320.

Gertler, Mark, and Peter Karadi. (2015) "Monetary Policy Surprises, Credit Costs, and Economic Activity." *American Economic Journal: Macroeconomics*, 7, 44–76.

Gürkaynak, Refet S., Brian Sack, and Eric T. Swanson. (2005) "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements." *International Journal of Central Banking*, 1, 55–93.

Haldane, Andy, and Michael McMahon. (2018) "Central Bank Communications and the General Public." *American Economic Review: Papers and Proceedings*, 108, 578–83.

Hamilton, James D. (1985) "Historical Causes of Postwar Oil Shocks and Recessions." *Energy Journal*, 6, 97–116.

Hansen, Stephen, and Michael McMahon. (2016) "Shocking Language: Understanding the Macroeconomic Effects of Central Bank Communication." *Journal of International Economics*, 99, 114–33.

Hansen, Stephen, Michael McMahon, and Andrea Prat. (2018) "Transparency and Deliberation Within the FOMC: A Computational Linguistics Approach." *Quarterly Journal of Economics*, 133, 801–70.

Hayo, Bernd, and Edith Neuenkirch. (2018) "The Influence of Media use on Laymen's Monetary Policy Knowledge in Germany." *Scottish Journal of Political Economy*, 65, 1–26.

Hayo, Bernd, and Matthias Neuenkirch. (2012) "Bank of Canada Communication, Media Coverage, and Financial Market Reactions." *Economics Letters*, 115, 369–72.

Hayo, Bernd, and Matthias Neuenkirch. (2015) "Self-Monitoring or Reliance on Media Reporting: How do Financial Market Participants Process Central Bank News?." *Journal of Banking & Finance*, 59, 27–37.

Hendry, Scott. (2012) "Central Bank Communication or the Media's Interpretation: What Moves Markets?." Staff Working Papers 12-9, Bank of Canada.

Howard, Jeremy, and Sebastian Ruder. (2018) "Universal Language Model Fine-tuning for Text Classification." In "ACL," .

Jarocinski, Marek, and Peter Karadi. (2020) "Deconstructing Monetary Policy Surprises - The role of Information Shocks." *American Economic Journal: Macroeconomics*, 12, 1–43.

Óscar, Jordá. (2005) "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review*, 95, 161–82.

Kilian, Lutz. (2008) "Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?." *Review of Economics and Statistics*, 90, 216–40.

King, Gary, Benjamin Schneer and Ariel White. (2017) "How the News Media Activate Public Expression and Influence National Agendas." *Science*, 358, 776–80.

Kusner, Matt J., Yu Sun, Nicholas Kolkin, and Kilian Weinberger. (2015) "From Word Embeddings to Document Distances." *Proceedings of Machine Learning Research*, 37, 957–66.

Lamla, Michael, and Dmitri Vinogradov. (2019) "Central Bank Announcements: Big News for Little People?." *Journal of Monetary Economics*, 108, 21–38.

Larsen, Vegard H. (2021) "Components of Uncertainty." *International Economic Review*, 62, 769–788.

Larsen, Vegard H., and Leif A. Thorsrud. (2019) "The Value of News for Economic Developments." *Journal of Econometrics*, 210, 203–18.

Larsen, Vegard H., and Leif Anders Thorsrud. (2017) "Asset Returns, News Topics, and Media Effects." Working Paper 2017/17, Norges Bank.

Larsen, Vegard H., and Leif Anders Thorsrud. (2018) "Business Cycle Narratives." Working Paper 2018/03, Norges Bank.

Larsen Vegard H., Thorsrud Leif Anders, Zhulanova Julia. 2021. News-driven inflation expectations and information rigidities. *Journal of Monetary Economics* 117, 507–520. http://doi.org/10.1016/j.jmoneco.2020.03.004

Le, Quoc, and Tomas Mikolov. (2014) "Distributed Representations of Sentences and Documents." *Proceedings of Machine Learning Research*, 32, 1188–96.

Loughran, Tim, and Bill Mcdonald. (2016) "Textual Analysis in Accounting and Finance: A Survey." *Journal of Accounting Research*, 54, 1187–230.

Mcauliffe, Jon D., and David M. Blei. (2008) "Supervised Topic Models." In *Advances in Neural Information Processing Systems, edited by D. Koller, D. Schuurmans, Y. Bengio, L. Bottou.*, pp. 121–128. Vancouver, British Columbia, Canada: Neural Information Processing Systems Foundation, Inc. (NIPS). https://papers.nips.cc/paper/2008

Merity, Stephen, Nitish Shirish Keskar, and Richard Socher (2018) "Regularizing and Optimizing LSTM Language Models." Conference paper at ICLR 2018.

Mertens, Karel, and Morten Ravn. (2014) "A Reconciliation of SVAR and Narrative Estimates of Tax Multipliers." *Journal of Monetary Economics*, 68, S1–19.

Miranda-Agrippino, Silvia, and Giovanni Ricco. (2020) "The Transmission of Monetary Policy Shocks." *American Economic Journal: Macroeconomics*, Forthcoming.

Nakamura, Emi, and Jon Steinsson. (2018) "High Frequency Identification of Monetary Non-Neutrality: The Information Effect." *Quarterly Journal of Economics*, 133, 1283–330.

Newey, Whitney K., and Kenneth D. West. (1987) "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55, 703–8.

Nimark, Kristoffer P., and Stefan Pitschner. (2019) "News Media and Delegated Information Choice." *Journal of Economic Theory*, 181, 160–96.

Pang, Bo, Lillian Lee, and Shivakumar Vaithyanathan. (2002) "Thumbs Up?: Sentiment Classification Using Machine Learning Techniques." In "*Proceedings of the ACL-02 Conference on Empirical Methods in Natural Language Processing*", pp. 79–86, Vol. 10, Stroudsburg, PA, USA: Association for Computational Linguistics. https://dl.acm.org/doi/proceedings/10.5555/1118693; https://doi.org/10.3115/1118693.1118704

Qvigstad, Jan Fredrik. (2019) "How to Draft "Good" Minutes of Monetary Policy Decisions." *Economic Affairs*, 39, 44–64.

Ramey, Valerie A. (2011) "Identifying Government Spending Shocks: It's all in the Timing." *Quarterly Journal of Economics*, 126, 1–50.

Romer, Christina D., and David H. Romer. (1989) "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz." *NBER Macroeconomics Annual*, 4, 121–70.

Romer, Christina D., and David H. Romer. (2004) "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review*, 94, 1055–84.

Shehata, Adam, and Jesper Strömbäck. (2011) "A Matter of Context: A Comparative Study of Media Environments and News Consumption Gaps in Europe." *Political Communication*, 28, 110–34.

Stock, James H. (2008) "What's New in Econometrics: Time Series, Lecture 7." Short course lectures, NBER Summer Institute.

Stock, James H., and Mark W. Watson. (2016) "Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics. " In *Handbook of Macroeconomics*, edited by John B. Taylor and Harald Uhlig, Vol. 2, pp. 415–525. Amsterdam: Elsevier.

Svensson, Lars E O. (2010) "Open-economy Inflation Targeting." *Journal of International Economics*, 50, 155–183.

Svensson, Lars E O. (2014) "Inflation Targeting and 'Leaning Against the Wind'." *International Journal of Central Banking*, June, 10(2): 103–114.

Svensson, Lars E O. (2017) "Leaning Against the Wind: Costs and Benefits, Effects on Debt, Leaning in DSGE Models, and a Framework for Comparison of Results." *International Journal of Central Banking*, 13(3): 385–408.

Taboada, Maite, Julian Brooke, Milan Tofiloski, Kimberly Voll, and Manfred Stede. (2011) "Lexicon-Based Methods for Sentiment Analysis." *Computational Linguistics*, 37, 267–307.

Taddy, Matt. (2013) "Multinomial Inverse Regression for Text Analysis." *Journal of the American Statistical Association*, 108, 755–70.

Thorsrud, Leif Anders. (2018) "Words are the New Numbers: A Newsy Coincident Index of the Business Cycle." *Journal of Business & Economic Statistics*, 38, 1–17.

Tibshirani, Robert. (1996) "Regression Shrinkage and Selection via the Lasso." *Journal of the Royal Statistical Society. Series B (Methodological)*, 58, 267–88.

Zeev, Nadav Ben. (2018) "What can we Learn About News Shocks from the Late 1990s and Early 2000s Boom-Bust Period?." *Journal of Economic Dynamics and Control*, 87, 94–105.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Table A1. Sensitivity of the narrative surprise measure to changing the terms used to identify the factors.

Figure A1. Announcement dates are reported on the horizontal axis, where the difference between Figure 1a and 1b is due to differences in data availability

Figure A2. Figure A2a reports the t values of $\hat{b}1$ in equation(7)

Figure A3. Figure A3a reports the t values of $\hat{b}1$ in equation(7)

Figure A4. Figure A4a reports the t values of $\hat{b}1$ in equation(7)

Figure A5. Figure A5a reports the t values of $\hat{b}1$ in equation(7)

Figure A6. The figures reports the estimates of $\hat{\phi}$ from equation(9)

Figure A7. The figures reports the estimates of $\hat{\phi}$ from equation(9)