

A Service of



Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Saskia ter Ellen; Larsen, Vegard Høghaug; Thorsrud, Leif Anders

Working Paper Narrative monetary policy surprises and the media

Working Paper, No. 19/2019

Provided in Cooperation with: Norges Bank, Oslo

Suggested Citation: Saskia ter Ellen; Larsen, Vegard Høghaug; Thorsrud, Leif Anders (2019) : Narrative monetary policy surprises and the media, Working Paper, No. 19/2019, ISBN 978-82-8379-116-7, Norges Bank, Oslo, https://hdl.handle.net/11250/2652989

This Version is available at: https://hdl.handle.net/10419/246100

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



ND https://creativecommons.org/licenses/by-nc-nd/4.0/deed.no





WORKING PAPER

Narrative monetary policy surprises and the media

NORGES BANK RESEARCH

19 | 2019

SASKIA TER ELLEN, VEGARD H. LARSEN , LEIF ANDERS THORSRUD



Working papers fra Norges Bank, fra 1992/1 til 2009/2 kan bestilles over e-post:

FacilityServices@norges-bank.no

Fra 1999 og senere er publikasjonene tilgjengelige på www.norges-bank.no

Working papers inneholder forskningsarbeider og utredninger som vanligvis ikke har fått sin endelige form. Hensikten er blant annet at forfatteren kan motta kommentarer fra kolleger og andre interesserte. Synspunkter og konklusjoner i arbeidene står for forfatternes regning.

Working papers from Norges Bank, from 1992/1 to 2009/2 can be ordered by e-mail FacilityServices@norges-bank.no

Working papers from 1999 onwards are available on www.norges-bank.no

Norges Bank's working papers present research projects and reports (not usually in their final form) and are intended inter alia to enable the author to benefit from the comments of colleagues and other interested parties. Views and conclusions expressed in working papers are the responsibility of the authors alone.

ISSN 1502-819-0 (online) ISBN 978-82-8379-116-7 (online)

Narrative monetary policy surprises and the media^{*}

Saskia ter Ellen^{\dagger}

Vegard H. Larsen^{\ddagger} Leif Anders Thorsrud[§]

This version October 4, 2019

Abstract

We propose a method to quantify narratives from textual data in a structured manner, and identify what we label "narrative monetary policy surprises" as the change in economic media coverage that can be explained by central bank communication accompanying interest rate meetings. Our proposed method is fast and simple, and relies on a Singular Value Decomposition of the different texts and articles coupled with a unit rotation identification scheme. Identifying narrative surprises in central bank communication using this type of data and identification 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. In turn, narrative monetary policy surprises lead to macroeconomic responses similar to what recent monetary policy literature associates with the information component of monetary policy communication. Our study highlights the importance of written central bank communication and the role of the media as information intermediaries.

JEL-codes: C01, C55, C82, E43, E52, E58

 ${\bf Keywords:} \ {\rm communication, \ monetary \ policy, \ factor \ identification, \ textual \ data$

^{*}This Working 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. We are grateful for research assistance from Helene Olsen. We thank Carola Binder, Jan Qvigstad, an anonymous referee, seminar participants at Norges Bank, and conference participants of The Changing Economy and the Business Cycle for valuable comments. We are grateful to the *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.

[†]Norges Bank. Email: saskia.ter-ellen@norges-bank.no

[‡]Norges Bank, and Centre for Applied Macroeconomics and Commodity Prices, BI Norwegian Business School. Email: vegard-hoghaug.larsen@norges-bank.no

[§]Centre for Applied Macroeconomics and Commodity Prices, BI Norwegian Business School. Email: leif.a.thorsrud@bi.no. BI Norwegian Business School, Nydalsveien 37, 0484 Oslo, Norway.

1 Introduction

"...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)

The quote above emphasizes the role of the media as an important transmission device for central bank communication. For households, unlikely to follow central bank communication directly, this is uncontroversial. Perhaps more surprising is the fact that professionals such as financial market participants also rely heavily on media reporting when following central bank events (Hayo and Neuenkirch (2015)). However, despite the importance of the media as an information source for households and professionals alike, this transmission mechanism is mostly overlooked in research trying to measure the effect of monetary policy communication.

To study this transmission mechanism, we propose a simple method to quantify narratives of economic interest from textual data, without having access to already classified training data. With this new textual model at hand, we adapt an event study framework and investigate; (i) whether the difference in narrative focus in central bank communication accompanying interest rate meetings and economic media coverage prior to those meetings correlates with conventional monetary policy surprises; (ii) whether such narrative differences can explain changes in media coverage after the meeting relative to before; and (iii) whether these explained changes affect important economic aggregates. Our results provide a negative answer to the first question, and a positive answer to the two latter questions.

To reach these conclusions, we take the view that media coverage is a good proxy for public beliefs about macroeconomic conditions and monetary policy, and that such beliefs might be affected by central bank communication that reach the public through the media. Accordingly, we call changes in narrative focus in media coverage explained by central bank communication accompanying interest rate meetings "narrative monetary policy surprises". To measure these surprises we put structure on the problem and focus on important narrative dimensions that feed into a central bank's decision making process: inflation, labor market, and exchange rate developments, as well as issues related to the oil market, financial stability, and uncertainty. As discussed in greater detail later, the method we propose and apply 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, simple, and requires minimal subjective judgment regarding the size and timing of narrative surprises.

Following the high frequency event study framework of Gürkaynak et al. (2005), who identify two-dimensional monetary policy surprises through movements in interest rates around the time of monetary policy announcements, we find that there is a weak and insignificant relationship between these conventional surprises and the difference in narrative focus in central bank communication accompanying interest rate meetings and economic media coverage prior to those meetings. Hence, the narrative differences identified here capture a different part of the central bank's communication than conventional monetary policy surprises do. In turn, we show that the narrative differences lead to a significant change in media coverage after the interest rate meeting relative to before, whereas conventionally measured monetary policy surprises do not.¹ Finally, we show that these discrepancies matter for economic outcomes. Following narrative monetary policy surprises, i.e., changes in narrative focus in media coverage explained by central bank communication accompanying interest rate meetings, 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 labels the information component of monetary policy surprises (e.g. Jarocinski and Karadi (2018), Cieslak and Schrimpf (2018), Nakamura and Steinsson (2018), Andrade and Ferroni (2019)).² 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. 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 monetary policy 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 has the advantage of using information explicitly communicated by the central bank and conveyed by the media. As such, it allows us to quantify the

¹Importantly, both of these results hold when controlling for changing macroeconomic conditions between announcement dates, measured by revisions to forecasts (expectations) published by the central bank alongside the interest rate decision.

²In fact, a positive co-movement between the interest rate and the stock market following monetary policy surprises has been the defining identifying feature of such "central bank information shocks" in Jarocinski and Karadi (2018) and Cieslak and Schrimpf (2018).

different narrative contributions of the information component.

Our study is applied to Norwegian data, and communication published by Norges Bank via their Executive Board Assessments, or EBA for short. These documents are official statements released at the same time as the monetary policy decisions are made public, and serve as a justification for the decision being made (Qvigstad and Schei (2018), 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, e.g., interest rate path predictions (Brubakk et al. (2017)). To measure media coverage, we use articles from Norway's most important business newspaper, *Dagens Næringsliv* (DN). This outlet is Norway's fourth largest newspaper irrespective of subject matter, which enables us to capture a representative source of media information for both market participants and households.

In terms of economics, this paper contributes to a large and growing literature investigating the importance of central bank communication and the measurement and content of monetary policy surprises (e.g. Gürkaynak et al. (2005), Miranda-Agrippino and Ricco (2018), Jarocinski and Karadi (2018), Andrade and Ferroni (2019), Nakamura and Steinsson (2018)). While influential papers in this literature have focused mainly on hard quantitative information released by the central banks, an emerging literature has started looking more into written communication like minutes, speeches, and monetary policy reports. Blinder et al. (2008) provides a comprehensive overview of the literature up to 2008, while newer examples include Hansen and McMahon (2016), Ehrmann and Talmi (2017), and Hansen et al. (2018). More specifically, the role of the media for the transmission of central bank communication has received increased attention through work by, e.g., Berger et al. (2011), Hendry (2012), Hayo and Neuenkirch (2012), Binder (2017), and Coibion et al. (2019). Still, to the best of our knowledge, there exists no study looking jointly at monetary policy communication and media coverage to provide a good measure of narrative monetary policy surprises. This paper aims to help fill that void.

In terms of methodology, our paper speaks to the Natural Language Processing 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)). While this literature is vast, it is mostly applied in either unsupervised settings, 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 is 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. As such, it brings together insights from econometrics and Natural Language Processing, providing a useful tool for anyone who wishes to quantify text in a structured manner and analyze its impact.

The rest of the paper is organized as follows: Sections 2 and 3 present our research design and methodology. Section 4 presents the results, while Section 5 concludes.

2 Research design

Our research questions; (i) are differences in narrative focus in central bank communication accompanying interest rate meetings and economic media coverage prior to those meetings correlated with conventional monetary policy surprises?; (ii) do these differences affect media coverage?; and (iii) what are the macroeconomic implications of such narrative surprises?, can be formalized by three simple regressions:

$$s_t^{conv} = b_1 n d_t^{CB,N} + \boldsymbol{b}_2 \boldsymbol{z}_{1t} + e_t \tag{1a}$$

$$nd_t^{N,N} = \delta_1 nd_t^{CB,N} + \boldsymbol{\delta}_2 \boldsymbol{z}_{2t} + u_t \tag{1b}$$

$$y_{t+h} = \phi_1 s_t^{narr} + \phi_2 \boldsymbol{z}_{3t} + \epsilon_{t+h}$$
(1c)

Here, s_t^{conv} is a conventionally measured monetary policy surprise at event day t, and $nd_t^{CB,N}$ (narrative difference) is a measure of the overall difference in narrative focus between news media (N) coverage w^- days prior to the interest rate announcement and the central bank EBA (CB) at day t. $nd_t^{N,N}$ is the overall difference in narrative focus between news media coverage prior to relative to w^+ days after the interest rate announcement, and y_{t+h} is the cumulative change in a macroeconomic outcome variable, measured at monthly frequency h periods forward relative to t. Most importantly, s_t^{narr} is the part of the change in media focus before and after the announcement that can be explained by the surprising content of the central bank communication, i.e., what we define as a narrative monetary policy surprise. It is computed as the monthly aggregation of $\hat{\delta}_1 nd_t^{CB,N}$ from (1b). Finally, the $\mathbf{z}'s$ are vectors of control variables including, e.g., revisions to forecasts published by the central bank alongside the interest rate decision, or lagged values of the dependent variable.

The key variables in (1) are $nd_t^{CB,N}$ and s_t^{narr} . In the next section we describe in greater detail how we compute these latent concepts and identify their narrative dimensions. Before that, it is informative to go through the intuition for the narrative monetary policy surprise and these regressions.

Informally, we take the view that no agent has the resources to monitor all events that are potentially relevant for her 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 (Nimark and Pitschner (2019)).³ Accordingly, we treat media coverage prior to monetary policy announcements as a good proxy for public beliefs about macroeconomic conditions and monetary policy, and think of differences in narrative focus between the media and the central bank EBAs, i.e., $nd_t^{CB,N}$, as the surprising content of these EBAs. An example of such a surprise is when the media focuses heavily on, e.g., labor market developments, while the central bank focuses almost solely on, e.g., inflation developments. However, in line with the assumption that the media works as "information intermediaries" between agents and the state of the world, what we are ultimately after is the part of this surprising content that actually reaches the news readers, i.e., the general public. For this reason we identify the narrative monetary policy surprise as the part of the change in media focus before and after the announcement that can be explained by the surprising content of the central bank communication, i.e., $s_t^{narr} = \hat{\delta}_1 n d_t^{CB,N}$. Alternatively, if one treated the overall difference in narrative focus between news media coverage prior to relative after the interest rate announcement $(nd_t^{N,N})$ as the narrative surprise, one could not forcefully argue that it was the surprising content of the central bank communication that led to potential changes in media coverage.

By focusing our analysis on a window around the monetary policy announcement date, our narrative surprise component shares the event study framework often used to construct conventionally measured monetary policy surprises (s_t^{conv}) . In contrast to such surprises, however, we focus on the narrative dimension while conventional monetary policy surprises are typically derived from movements in specific markets, e.g., the interest rate market, using hard economic statistics and listed prices (see, e.g., Gürkaynak et al. (2005)).

In (1a), (1b), and (1c), the objects of interest are b_1 , δ_1 , and ϕ , respectively. b_1 measures to what extent narrative differences are informative about monetary policy surprises as conventionally measured. If the two objects are highly correlated, the need for additional and perhaps more computationally demanding measures is less pressing. Accordingly, we use (1a) as a means to justify our approach. δ_1 , in (1b), measures whether the surprising content of central bank *EBAs* affects media coverage. This is an important parameter. After all, for the narrative surprise to matter, people need to learn about it, at least in the short run. As we implicitly assume that most people get their information about monetary policy through the media, we expect δ_1 to be positive and

³In a general, but abstract, theoretical model, Nimark and Pitschner (2019) show that this delegation is optimal when the information flow is overwhelming, and that media's news selection functions and distributions of events jointly determine the degree to which knowledge about an event is common among agents.

significant. Finally, equation (1c) is a simple linear projection (Jordà (2005)), measuring how narrative monetary policy surprises affect macroeconomic variables. In essence, ϕ_1 allows us to quantify how, and to what extent, narrative monetary policy surprises matter.

3 Constructing narrative differences

To measure the narrative focus of market participants and the central bank, we use Norges Bank's Executive Board Assessment, or EBA for short, and the entire corpus, i.e., text and articles, published by *Dagens Næringsliv* (DN). Each EBA is a roughly two-page document published at the same time as the interest rate decision is made public. Between 1999 (Oct. 27) and 2019 (Mar. 21) there have been 152 interest rate decisions, for which we collect the associated EBAs from Norges Bank's web pages. The news data has been generously provided to us by the company *Retriever* through their "Atekst" database, and collected manually by us for the latter part of the sample. In total this data consists of roughly 200 000 news articles between 1999 and 2019, and over 80 000 unique words and terms.⁴

Importantly, both data sources are high-dimensional and unstructured, i.e., containing many words and documents, and none of the textual data sources have been classified as being about particular economic narratives. In the following, we first describe how we transform the raw data into quantitative information, and then how we extract identified narratives and differences from the texts.

3.1 Feature selection

As is common in the Natural Language Processing (NLP) literature, the raw textual data is cleaned before further analysis (Gentzkow et al. (2017)). 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. This set of words is much smaller than the vocabulary used in the newspaper, but reduces the dimensionality of the problem considerably. Note that this also potentially limits the newspaper content that is completely unrelated to the central bank's function, such as the sports or entertainment sections. We denote the size of this vocabulary as V. Next, because the newspaper content during weekends differs considerably from that published

⁴Although more and more news media consumption nowadays happens online, we only use printed news, and leave it for future research to explore how the changing media landscape might affect relationships such as those investigated in the current analysis.

during business days, i.e., featuring more background articles, travel, portrait interviews, etc., we remove all weekends from the news corpus.

Based on these steps we take a bag-of-words view and construct two document term matrices, \mathbf{C}^N and \mathbf{C}^{CB} , for the news media data (N) and the *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 \mathbf{C}^{CB} matrix has dimension $T^{CB} \times V$, where V = 2716 and $T^{CB} = 152$. Because there are many more news days than announcement days, the \mathbf{C}^N matrix is much larger, and has dimension $T^N \times V$, where $T^N = 6240$.

To construct a mapping between the information captured in the \mathbb{C}^N matrix at event time t and that conveyed by the *EBAs* in \mathbb{C}^{CB} , we sum the counts in the \mathbb{C}^N matrix over a period of w^- days prior to each announcement day t and take the mean of these counts. 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 into the matrices, and, as such, challenge the event study identification strategy. For these reasons, and because we do not have any strong exact prior of what w^- should be, we consider all $w^- = 1, \ldots, 10$, and denote these matrices $\mathbb{C}_{w^+}^N$. Similarly, we construct a $\mathbb{C}_{w^+}^N$ matrix, where the only difference between $\mathbb{C}_{w^+}^N$ and $\mathbb{C}_{w^-}^N$ is that we aggregate w^+ periods forward relative to the announcement day t when constructing $\mathbb{C}_{w^+}^N$. However, since the central bank actively engages in various communication strategies following interest rate meetings, we only consider $w^+ = 1, \ldots, 5.^5$

The final feature selection step we take is to weigh the different terms in the document term matrices by the inverse-document-frequency metric implied by the 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 downweighs stop words. Formally, we do this by first normalizing the C matrices from above such that each matrix entry reflects the relative frequency of that term within each document. Then, we compute the inverse-document-frequency score, denoted $idf = log(T/d_j^{CB})$, where $d_j^{CB} = \sum_i \mathbf{1}_{C_{ij}^{CB}>0}$, and construct $C_{ij}^{CB} \times idf_j = \hat{C}_{ij}^{CB}$ and $C_{ij,w}^N \times idf_j = \hat{C}_{ij,w}^N$.

⁵For example, 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. We also note that, because weekends are removed from the dataset, $w^- = 10$ 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 $C_{w^-}^N$ and $C_{w^+}^N$.

Notice here that because of the mapping constructed above, all matrices \hat{C}^{CB} , $\hat{C}^{N}_{w^{-}}$, and $\hat{C}^{N}_{w^{+}}$ now have dimensions $T^{CB} \times V$.

3.2 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. In the NLP literature, such factors are commonly referred to as topics, allowing us to identify what the different documents thematically are about in a parsimonious manner.⁶

Two commonly used factor modeling approaches used in the NLP literature are Latent Dirichlet Allocation (LDA) and Latent Semantic Indexing (LSI), introduced by Blei et al. (2003) and Deerwester et al. (1990), respectively. In both approaches, a document is thought of as consisting of potentially many topics/factors, but with different weights. Generically, the LSI-based approach is unsupervised, while the LDA method exists both in unsupervised and supervised versions (Mcauliffe and Blei (2008)). In our case, where we are interested in a specific set of narratives directly related to the central bank's decision making problem, the supervised version would be the appropriate choice. However, a supervised LDA implementation requires the researcher to construct a classified data set with identified factors (or topics) in the texts prior to training the model. In many macroeconomic applications, including this one, this is not feasible because we do not have enough textual data, i.e., *EBAs*, in the sample to appropriately divide it into informative training and testing sets.⁷ For this reason, we build on the LSI approach, noting that although the LDA method is widely looked upon as a better description of a text generating model than the LSI approach (Gentzkow et al. (2017)), the latter can still be highly competitive in terms of classifying text and documents (Kusner et al. (2015)).

The methodological contribution we make in this paper is to apply the LSI method in a supervised manner using textual data. In particular, while the standard LSI approach is simply a Singular Value Decomposition (SVD) of the document term matrix, introduced into the central bank communication literature by Boukus and Rosenberg (2006), we suggest to ex-post rotate the estimated factor space such that the factors can be given a narrative interpretation along dimensions of interest.

⁶An alternative approach to this problem could be to apply regularization techniques, e.g., LASSO (Tibshirani (1996)), to impose sparsity and implicit dimension reduction. Because we are interested in capturing differences in narrative focus, we do not follow this route. See Larsen and Thorsrud (2018) for a richer discussion about how factors/topics derived from textual data can be interpreted as narratives.

⁷Likewise, newer popular methods involving neural network architectures and word embeddings, like, e.g., Le and Mikolov (2014) and Kusner et al. (2015), are mostly unsupervised algorithms which in addition require large corpora for training.

Table 1. Key word list used to identify factors. 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.

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

While there are potentially a plethora of narratives one could consider, we put structure on the problem by focusing on narrative dimensions 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, issues 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 (see for example Gali and Monacelli (2005), Svensson (2010)), while 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 (2017)) in monetary theory and practice.

Formally, we achieve this as follows. First, define K as the total number of factors, and associate each factor with one particular (subjectively chosen) word, as illustrated in Table 1. Then, for a given \hat{C} matrix, order these K terms in the K first columns of the matrix and apply the SVD decomposition $\hat{C} = USV'$ with factors $F = U_{1:K}S_{1:K}$ $(T^{CB} \times K)$ and loadings $L = V_{1:K}$ $(V \times K)$ such that:

$$\hat{C} \approx FL'$$
 (2)

Now, the decomposition in (2) does not permit giving the factors and loadings an economic interpretation along the narrative dimensions discussed above. To identify the first factor with an inflation narrative, the second with the labor market, and so on, we therefore propose to rotate the factor space such that we get a so-called unit identification. To do this, we partition L from (2) as:

$$\boldsymbol{L} = \begin{bmatrix} \boldsymbol{L}_0 \\ \boldsymbol{L}_1 \end{bmatrix} \text{ with } \boldsymbol{L}_0 = \boldsymbol{L}_{1:K} \text{ and } \boldsymbol{L}_1 = \boldsymbol{L}_{K+1:V}$$
(3)

and apply the rotation:

$$\tilde{\boldsymbol{F}} = \boldsymbol{F} \boldsymbol{L}_0' \text{ and } \tilde{\boldsymbol{L}} = \boldsymbol{L} \boldsymbol{L}_0^{-1}$$
 (4)

where \tilde{F} and \tilde{L} are the identified factor and loading matrices, respectively. The upper $K \times K$ block of \tilde{L} equals the identity matrix, i.e., $\tilde{L}_{1:K} = I_K$. Accordingly, focusing on narrative dimensions that typically feed into central banks' decision making process, with

K = 6 as illustrated in Table 1, 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, we associate the first factor with an inflation narrative, the second factor with a labor market narrative, etc.

To construct measures of the narrative differences $nd_t^{CB,N}$ and $nd_t^{N,N}$ in (1), we proceed in two steps. First, we implement the SVD decomposition and (2)-(4) for each of the three matrices \hat{C}^{CB} , $\hat{C}_{w^-}^N$, and $\hat{C}_{w^+}^N$ separately. Then, difference measures are constructed as:

$$\tilde{nd}_{t}^{CB,N} = \sum_{k=1}^{K} (\tilde{F}_{k,t}^{CB} - \tilde{F}_{k,t:w^{-}}^{N})^{2} \text{ and } \tilde{nd}_{t}^{N,N} = \sum_{k=1}^{K} (\tilde{F}_{k,w^{+}:t}^{N} - \tilde{F}_{k,t:w^{-}}^{N})^{2}$$
(5)

i.e., the sum of the squared differences between each of the identified factors. Accordingly, large values of, e.g., $\tilde{nd}_t^{CB,N}$, signal the extent to which the media focuses on different topics than the central bank does in its *EBAs*. Note here that by constructing the factors from separate matrices, we allow the exact language and context in which the central bank and the media write about the different terms (used to identify the factors) to differ on average, and instead use the time-variation in the factors to identify the surprise component.

Second, to also capture potential differences in the tonality, i.e., sentiment, of reporting, we sign-adjust the $n\tilde{d}'s$ in (5) 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*.⁸ 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{C}^{CB} , $\hat{C}^N_{w^-}$, and $\hat{C}^N_{w^+}$. For example, $to_t^{CB} = (\#\text{positive terms} - \#\text{negative terms})$ in the t^{th} row of \hat{C}^{CB} , and these statistics are normalized across time, denoted to_t^{CB} , to ensure that we do not pick up systematic differences in the use of positive versus negative terms across sources. Then, the tonality difference across sources is computed as:

$$to_t^{CB,N} = (\bar{to}_t^{CB} - \bar{to}_{t:w^-}^N) \text{ and } to_t^{N,N} = (\bar{to}_{w^+:t}^N - \bar{to}_{t:w^-}^N)$$
 (6)

and the to statistics are used to sign-adjust the topic frequencies computed in (5) as:

$$nd_t^{CB,N} = \tilde{nd}_t^{CB,N} to_t^{CB,N} \text{ and } nd_t^{N,N} = \tilde{nd}_t^{N,N} to_t^{N,N}$$
(7)

⁸In recent economic research, and particularly in finance, 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 (international) word list used plays a minor role, and that our Norwegian translation of the *Harvard IV-4 Psychological Dictionary* works well across a wide range of applications (Larsen (2017), Larsen and Thorsrud (2017), Thorsrud (2018)).

3.3 Methodological discussion

We highlight four points about this methodology. First, using time series data, factor identification like in (3) and (4) 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 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.

Second, although the type of factor identification described above could potentially have been achieved much more simply 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, e.g., Bholat et al. (2015), simple count-based methods can not 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, whereas the simple count-based method is not. We formally show this in Section 4.3.

Third, while related in spirit to narrative identification used and proposed 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 (FOMC), made public with a five-year delay, to single out events that they argue represent monetary policy shocks. Similar approaches have since then been applied in both 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 we focus on media's role as "information intermediaries" by letting the discrepancy between media coverage and the EBA define narrative monetary policy shocks. Relatedly, and more recently, 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 in a time series context. While more data-driven and automated than the pure 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 the extent that the researcher is interested in identifying the difference in tonality for specific narratives, e.g., with respect to *inflation*, and not only the overall contribution, as in (6), our approach falls short. We leave it for future research to design approaches that can also identify the tonality of the individual components of, e.g., narrative monetary policy shocks.⁹

4 Results

In the following we first present the estimated factors and our measure of the narrative differences. We then turn to the regression analysis and our estimates of the equations in (1).

4.1 Factors and narrative differences

Figure 1 reports the identified factors $\tilde{F}_{k,t}^{CB}$ and $\tilde{F}_{k,t:w^-}^N$, as well as $to_t^{CB,N}$ and $nd_t^{CB,N}$. Starting with the latent factors estimated from the news media dataset ($\tilde{F}_{k,t:w^-}^N$), colored in gray in Figures 1a–1f, the overall picture is that they seem to capture well the conventional narrative held about economic developments and monetary focus the last two decades. For example, the estimates suggest that the media focused more on unemployment related issues around 2003, 2009 and 2015. All of these periods were 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 went from a fixed exchange rate regime to inflation targeting

⁹A related concern can be raised with respect to equation (7), where the potential case $\tilde{nd}_t^{CB,N} = 0$ (or $\tilde{nd}_t^{N,N} = 0$), i.e., perfectly equal narrative focus, yields the unrealistic result $nd_t^{CB,N} = 0$ (or $nd_t^{N,N} = 0$), irrespective of any differences in tonality. As a response to this, we show in Section 4.3 that our main results are robust to working with the unsigned narrative differences (from equation (5)) and hence our results are not driven by the peculiarities of the tone measure. Still, we prefer the tone-adjusted difference measures as our benchmark specification because it allows is to compute meaningful impulse response functions using equation (1c).



Figure 1. Identified factors and the narrative monetary policy surprise. In Figures 1a–1f the solid black lines illustrate the evolution of the narrative focus in the *EBAs* (\tilde{F}_t^{CB}), while the broken gray lines illustrate the evolution of the narrative focus in the media ($\tilde{F}_{t:w^-}^N$) for different values of w^- . All topics are normalized (mean of zero, and standard deviation of one). Announcement dates are reported on the horizontal axis. See Figure 5a in Appendix A for an illustration of $nd_t^{N,N}$.

in 2001. The particular peak in the exchange rate factor around 2003 is also natural, and 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 the oil market got a lot of attention in the mid 2000s when this sector was a key engine for growth (Bjørnland and Thorsrud (2016)), as well as since 2014, when the decline in oil price led to concerns about the Norwegian economy, and that focus on credit conditions and uncertainty peaked during the financial crisis. 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, i.e. the news aggregation window.

To get an alternative impression of the contexts the different terms used to identify



Figure 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. The font size of each term represents the degree of similarity. 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 we use an underscore to illustrate such cases.

the factors represent, we report in Figure 2 word clouds constructed based on the cosine similarity between the word vector for key word k = 1, ..., K and term j = 1, ..., 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, e.g., 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, the results align well with the results reported for the factors themselves, and suggest that the method presented in Section 3.2 is able to extract meaningful information from the textual data.

The narrative focus in the central bank EBAs ($\tilde{F}_{k,t}^{CB}$), as we estimate it, is reported in black in Figures 1a–1f. For the inflation and exchange rate factors 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. This was a period when oil prices fell sharply, triggering discussions in Norway about recession risks and future economic prospects. For the other factors the differences between the two sources seem quite large. It is, for example, striking that the EBAs during the financial crisis did not contain more narrative information about uncertainty and credit.

Finally, Figures 1g and 1h summarize the overall difference in tonality and narrative focus between what is written about in the media and in the *EBAs*. In Figure 1g we have plotted the tonality contribution $(to_t^{CB,N})$ separately 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. While our approach for identifying the difference in sentiment between the different sources undoubtedly is simple, we conclude from this that it at least seems able to capture important features of the evolution of the actual monetary policy instrument.

Our measure of the overall narrative difference $(nd_t^{CB,N})$ is reported in Figure 1h. 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. Such unprecedented events are likely to cause some disagreement between the central bank and the public, and potentially surprising central bank communication.

4.2 Regression results

To gauge whether the proposed narrative surprises capture something different than what conventional monetary policy surprises do, we start by estimating (1a), repeated here for convenience: $s_t^{conv} = b_1 n d_t^{CB,N} + b_2 z_{1t} + e_t$. $n d_t^{CB,N}$, i.e., the narrative difference in central bank communication relative to media coverage, is the main explanatory variable, and we identify the dependent variable s_t^{conv} following the method pioneered by Gürkaynak et al. (2005). In particular, to construct a measure of s_t^{conv} , we use a high-frequency event study identification strategy, and extract movements in interest rates around the monetary policy announcement time on day t. The way we do this for the Norwegian data is described in detail in Brubakk et al. (2017).¹⁰ Note that this methodology allows us to decompose the surprise into two components, namely a "target" (T) and "path" (P) component. 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^{conv} , surprises are graphed in Figure 5b in Appendix A.

To control for changing macroeconomic conditions between announcement dates and other quantitative information that potentially explains monetary policy surprises we include in the vector \mathbf{z}_{1t} revisions in forecasts published by the central bank at the interest

¹⁰Our event window is 90 minutes: it captures the change in interest rates between 15 minutes before the announcement and 75 minutes after the meeting. This captures 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 minutes and a day.

rate announcement time. As Norges Bank has published its own interest rate path since 2005, 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. Thus, the control vector z_{1t} contains nine elements, which are all collected from Norges Bank's *Monetary Policy Report*.¹¹

To favor a small model size, and reduce noise and potential biases, we follow Belloni et al. (2014) and implement a double selection procedure for selecting the relevant control variables in z_t . In short, the double selection algorithm is implemented as follows: First, we regress the treatment $(nd_t^{CB,N})$ and the dependent (s_t^{conv}) variables separately on all the variables in the z_t vector using the LASSO estimator.¹² Next, after these two penalized regressions, we run an OLS regression on the dependent variable, including the treatment variable and the union of the control variables selected in step one.

Figure 3a reports the t-value associated with b_1 in (1a), for all values of w^- . Numbers reported above and below the lines are the adjusted R^2 statistics and the number of chosen control variables in each regression, respectively. One feature stands out: Irrespective of whether we measure conventional monetary policy shocks using s_t^T or s_t^P , their correlation with $nd_t^{CB,N}$ is weak and insignificant. Thus, in terms of question (i), we conclude that the narrative differences capture something different than conventional monetary policy shocks do.

Turning to question (ii), namely whether narrative differences in central bank communication affect media coverage, we estimate equation (1b), which was: $nd_t^{N,N} = \delta_1 nd_t^{CB,N} + \delta_2 \mathbf{z}_{2t} + u_t$. The upper plane in Figure 3b 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 \mathbf{z}_{2t} includes \mathbf{z}_{1t} , as well as s_t^T and s_t^P . We observe that the narrative differences have a positive and highly significant effect on the change in media coverage when we construct $nd_t^{CB,N}$ and $nd_t^{N,N}$ using small values of w. For $w^- = 1$ and $w^+ = 1$, the adjusted R^2 statistic is roughly 16 percent. 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 $nd_t^{CB,N}$ with s_t^T in equation (1b), and re-do the double selection estimation routine. As seen from the lower plane in Figure 3b, the conventional monetary policy shock (s_t^T) has

¹¹We 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 z_t vector.

¹²We estimate the LASSO for 100 different penalization parameters λ , and use the BIC to chose the one with minimum loss. See Tibshirani (1996) for details about the LASSO algorithm itself.



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



(c) Media spillovers and narrative contributions

Figure 3. Figure 3a reports the t-values of \hat{b}_1 in equation (1a). 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 3b reports the t-values of $\hat{\delta}_1$ in equation (1b) when either $nd_t^{CB,N}$ or s_t^T is used as the treatment variable. Figure 3c shows s_t^{narr} decomposed into narrative contributions.

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" and contained within the interest rate market.¹³

Figure 3c reports a bar plot of the narrative monetary policy surprise $s_t^{narr} = \hat{\delta}_1 n d_t^{CB,N}$ for each event day, i.e., monetary policy announcement day, in the sample. The figure also highlights an additional advantage with our narrative methodology, relative to con-

¹³To the extent that a "silent" shock propagates through the economy it can of course eventually have important economic consequences. For example, a "silent" contractionary monetary policy shock might eventually lead private banks to increase their interest rates, because their borrowing costs have in effect gone up, and thereby increase, e.g., the public's incentives to save.

ventional identification strategies, namely that we can decompose the surprise component into what it is about. In particular, since $nd_t^{CB,N}$ is a linear combination of the different narrative contributions, we can use the $\hat{\delta}_1$ estimate from (1b) to decompose the regression results into the contributions from each narrative. Using the $\hat{\delta}_1$ estimate associated with $w^- = w^+ = 1$, three periods stand out; the late 1990s and early 2000s, 2008/2009, and 2011/2012. The peculiarity of each of these episodes have been commented on earlier in this analysis. The important point here is that the methodology gives a plausible decomposition of the narrative monetary policy surprise contribution given what we now know about historical developments.

Together, the results presented in Figure 3 suggest that the narrative surprises presented here reflect information that is not already present in existing surprise measures, and that this type of information has an effect on media coverage. Do these differences also matter for macroeconomic outcomes? Figure 4 answers this question (iii), and reports our estimates of ϕ_1 from (1c), repeated here for convenience: $y_{t+h} = \phi_1 s_t^{narr} + \phi_2 z_{3t} + \epsilon_{t+h}$. We consider $h = 0, \ldots, 24$, and six important monthly financial and 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, for comparison, we also include response functions from estimating equation (1c) with s_t^T instead of s_t^{narr} . In both cases, the shocks are aggregated to monthly frequency and normalized to a one standard deviation innovation, and we report 95 percent confidence bands as well as the mean response.¹⁴ 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 percent of the variation in the stock market can be explained by the narrative monetary policy surprise on the 5 months horizon, while the conventional monetary policy shocks explains only roughly 6 percent at the same horizon.

The differences in macroeconomic outcomes between a conventional monetary pol-

¹⁴All dependent variables are (log) differenced prior to estimation. The control vector z_{3t} in (1c) includes up to 12 lags of the dependent variable as well as a linear trend. The lag length is selected by the BIC. We have also experimented with including additional macroeconomic control variables in the z_{3t} vector, observing that this only adds noise to the estimation and does not affect our qualitative conclusions.



(a) 3-month interest rate

(b) Stock market



Figure 4. The figures report the estimates of $\hat{\phi}^h$ from equation (1c) for $h = 0, \ldots, 24$ months. The mean estimates and 95 percent 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 s_t^{narr} or s_t^T .

icy surprise and the narrative surprises are striking, but speak directly to newer monetary studies emphasizing the information component of monetary policy surprises (e.g. Jarocinski and Karadi (2018), Cieslak and Schrimpf (2018), Nakamura and Steinsson (2018), Andrade and Ferroni (2019)). In fact, our results are also qualitatively in line with the macroeconomic responses obtained in Brubakk et al. (2019), who analyze the effects of information shocks, in addition to conventional and forward guidance shocks, on Norwegian data using a modified version of the methodology developed by Jarocinski and Karadi (2018).¹⁵

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, 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. (2019)).

4.3 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 differences $(nd_t^{CB,N})$ and those identified through movements in the interest rate market (s_t^{conv}) , might seem surprising. However, as we show in Figure 6a, in Appendix A, if we instead focus on the absolute size of the surprises, and disregard their sign, we obtain a more significant link. In particular, using (1a) and regressing $n\tilde{d}_t^{CB,N}$ (from equation (5)) on the absolute value of the conventional surprise measures $(|s_t^{conv}|)$, 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 $n\tilde{d}_t^{N,N}$ as the dependent variable, and $|s_t^{conv}|$ as the treatment variable in equation (1b), we obtain more or less the same insignificant result as before, see Figure 6b in Appendix A. In contrast, $n\tilde{d}_t^{CB,N}$ has a positive and significant effect on $n\tilde{d}_t^{N,N}$, confirming that also this (unsigned) measure of a narrative surprise in central bank communication affects media coverage.¹⁶

One might argue that it is the "path" factor, rather than the "target" factor, that captures central bank communication and hence should be more similar to our narrative surprise component in terms of macroeconomic responses. We have also computed the

¹⁵See also Bjørnland et al. (2019) for additional evidence pointing towards the information component of (Norwegian) monetary policy surprises.

¹⁶Figure 6c in Appendix A replicates the decomposition graph in Figure 3c using the unsigned \tilde{nd} measures, and confirms the same narrative impression discussed earlier.

macroeconomic responses following a conventional monetary policy shock using s_t^P as the shock of interest, i.e., the "path" factor identified through movements in mediumterm interest rates, orthogonal to movements in the near-term interest rate. Figure 7, in Appendix A, shows that this is not the case.

As a final experiment, we show in Table 2, in Appendix A, how the factor-based identification scheme proposed here is relatively robust to changing the exact key terms used to identify the factors. In particular, we cycle through 30 unique alternative combinations of key words, listed in Table 2, and re-do the calculations of $n\tilde{d}_t^{CB,N}$. As seen from the table, depending somewhat on the window w^- used in the calculations, the correlation between these alternative shock estimates, and our benchmark estimate, is seldom lower than 0.40, very often above 0.70, and sometimes as high as 0.90. In Figure 8, in Appendix A, we also confirm that the macroeconomic effects of narrative monetary policy surprises remain robust to these changes. In contrast, if we instead compute the factors as simple counts, as discussed in Section 3.3, we observe from the second row in Table 2 that the resulting narrative differences would have been much more sensitive to the exact key words used to identify the narrative dimensions.

5 Conclusion

In this paper we propose a fast, simple, and automated method for identifying what we label as "narrative monetary policy surprises" using textual data. Taking the view that central bank communication that actually reaches the general public might have a different effect on the economy than conventionally measured monetary policy surprises, we identify the narrative surprises as the change in media coverage that can be explained by the surprise in narrative focus in central bank communication accompanying interest rate meetings. We put structure on the problem by focusing on narrative dimensions that typically feed into a central bank's decision making process and propose to identify these from the different corpora (central bank statements and newspaper articles) by applying a Singular Value Decomposition and an ex-post unit rotation identification scheme.

Our results suggest that the narrative monetary surprises have a weak and insignificant correlation with conventionally measured monetary policy surprises, indicating that they capture a different part of the central bank's communication than conventional monetary policy surprises do. We further show that the narrative surprises in central bank communication lead to a significant change in media coverage after the interest rate meeting relative to before, while monetary policy surprises measured using conventional methods do not. In turn, these differences are shown to matter for the evolution of macroeconomic aggregates following monetary policy surprises, where narrative surprises cause response patterns in line with what newer monetary policy studies label the information component of monetary policy. As such, our study highlights the importance of written central bank communication and the role of the media as information intermediaries.

The method we have proposed and applied is simple, fast, 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. Accordingly, similar analysis can easily be undertaken in other applications where narrative focus is relevant.

References

- Aastveit, K. A., H. C. Bjørnland, and L. A. Thorsrud (2015). What drives oil prices? Emerging versus developed economies. *Journal of Applied Econometrics* 30(7), 1013– 28.
- Andrade, P. and F. Ferroni (2019). Delphic and Odyssean monetary policy shocks: Evidence from the euro area. *Working Paper*.
- Antolín-Díaz, J. and J. F. Rubio-Ramírez (2018). Narrative sign restrictions for svars. American Economic Review 108(10), 2802–29.
- Bai, J. and S. Ng (2013). Principal components estimation and identification of static factors. *Journal of Econometrics* 176(1), 18 29.
- Belloni, A., V. Chernozhukov, and C. Hansen (2014). High-dimensional methods and inference on structural and treatment effects. *Journal of Economic Perspectives* 28(2), 29–50.
- Berger, H., M. Ehrmann, and M. Fratzscher (2011). Monetary policy in the media. *Journal* of Money, Credit and Banking 43(4), 689–709.
- Bernanke, B. S. (2007). Inflation expectations and inflation forecasting. Speech 306, Board of Governors of the Federal Reserve System (U.S.).
- Bholat, D., S. Hansen, P. Santos, and C. Schonhardt-Bailey (2015). Text mining for central banks: Handbook. *Centre for Central Banking Studies 33*, pp. 1–19.
- Binder, C. (2017). Federal reserve communication and the media. Journal of Media Economics 30(4), 191–214.
- Bjørnland, H. C. and L. A. Thorsrud (2016). Boom or gloom? Examining the Dutch disease in two-speed economies. *The Economic Journal* 126 (598), 2219–2256.

- Bjørnland, H. C., T. Ekeli, P. M. Geraats, and K. Leitemo (2004). Norges Bank Watch 2004. Technical Report 5, Centre for Monetary Economics BI Norwegian Business School.
- Bjørnland, H. C., L. A. Thorsrud, and S. K. Zahiri (2019). Do Central Banks Respond Timely to Develoments in the Global Economy? Oxford Bulletin of Economics and Statistics (Forthcoming).
- Blei, D. M., A. Y. Ng, and M. I. Jordan (2003). Latent Dirichlet Allocation. J. Mach. Learn. Res. 3, 993–1022.
- Blinder, A. S., M. Ehrmann, M. Fratzscher, J. De Haan, and D.-J. Jansen (2008). Central bank communication and monetary policy: A survey of theory and evidence. *Journal* of Economic Literature 46(4), 910–45.
- Bloom, N. (2014). Fluctuations in Uncertainty. *Journal of Economic Perspectives* 28(2), 153–76.
- Boukus, E. and J. V. Rosenberg (2006). The information content of fomc minutes. Available at SSRN 922312.
- Brubakk, L., S. ter Ellen, Ø. Robstad, and H. Xu (2019). The macroeconomic effects of forward communication. *Working Paper*.
- Brubakk, L., S. ter Ellen, and H. Xu (2017). Forward guidance through interest rate projections: does it work? *Norges Bank Working Paper*.
- Cieslak, A. and A. Schrimpf (2018). Non-Monetary News in Central Bank Communication. Journal of International Economics forthcoming.
- Coibion, O., Y. Gorodnichenko, and M. Weber (2019). Monetary policy communications and their effects on household inflation expectations. *Working Paper*.
- Deerwester, S., S. T. Dumais, G. W. Furnas, T. K. Landauer, and R. Harshman (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science* 41(6), 391–407.
- Ehrmann, M. and J. Talmi (2017). Starting from a blank page? Semantic similarity in central bank communication and market volatility. Working Paper Series 2023, European Central Bank.
- Gali, J. and T. Monacelli (2005). Monetary Policy and Exchange Rate Volatility in a Small Open Economy. *Review of Economic Studies* 72, 707–734.

- Gentzkow, M., B. T. Kelly, and M. Taddy (2017). Text as data. Working Paper 23276, National Bureau of Economic Research.
- Gerdrup, K., F. Hansen, T. Krogh, and J. Maih (2017). Leaning Against the Wind When Credit Bites Back. *International Journal of Central Banking September*, 385–408.
- Gürkaynak, R. S., B. Sack, and E. 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.
- Hamilton, J. D. (1985). Historical causes of postwar oil shocks and recessions. The Energy Journal 6(1), pp. 97–116.
- Hansen, S. and M. McMahon (2016). Shocking language: Understanding the macroeconomic effects of central bank communication. *Journal of International Economics* 99(S1), 114–133.
- Hansen, S., M. McMahon, and A. Prat (2018). Transparency and deliberation within the fomc: A computational linguistics approach. The Quarterly Journal of Economics 133(2), 801–870.
- Hayo, B. and M. Neuenkirch (2012). Bank of Canada communication, media coverage, and financial market reactions. *Economics Letters* 115, 369–372.
- Hayo, B. and M. Neuenkirch (2015). Self-monitoring or reliance on media reporting: How do financial market participants process central bank news? Journal of Banking & Finance 59(C), 27–37.
- Hendry, S. (2012). Central Bank Communication or the Media's Interpretation: What Moves Markets? Staff Working Papers 12-9, Bank of Canada.
- Howard, J. and S. Ruder (2018). Universal Language Model Fine-tuning for Text Classification. In *ACL*.
- Jarocinski, M. and P. Karadi (2018). Deconstructing Monetary Policy Surprises The role of Information Shocks. *CEPR Discussion Paper Series DP12765*.
- Jordà, s. (2005). Estimation and Inference of Impulse Responses by Local Projections. American Economic Review 95(1), 161–182.
- Kilian, L. (2008, May). Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy? The Review of Economics and Statistics 90(2), 216–240.

- Kusner, M. J., Y. Sun, N. I. Kolkin, and K. Q. Weinberger (2015). From Word Embeddings to Document Distances. In Proceedings of the 32Nd International Conference on International Conference on Machine Learning - Volume 37, ICML'15, pp. 957–966. JMLR.org.
- Larsen, V. H. (2017). Components of uncertainty. Working Paper 2017/5, Norges Bank.
- Larsen, V. H. and L. A. Thorsrud (2017). Asset returns, news topics, and media effects. Working Paper 2017/17, Norges Bank.
- Larsen, V. H. and L. A. Thorsrud (2018). Business Cycle Narratives. Working Paper 2018/03, Norges Bank.
- Larsen, V. H. and L. A. Thorsrud (2019). The value of news for economic developments. Journal of Econometrics 210(1), 203–218.
- Larsen, V. H., L. A. Thorsrud, and J. Zhulanova (2019, February). News-driven inflation expectations and information rigidities. Working Paper 2019/5, Norges Bank.
- Le, Q. and T. Mikolov (2014). Distributed representations of sentences and documents. In *International conference on machine learning*, pp. 1188–1196.
- Loughran, T. and B. Mcdonald (2016). Textual analysis in accounting and finance: A survey. *Journal of Accounting Research* 54(4), 1187–1230.
- Mcauliffe, J. D. and D. M. Blei (2008). Supervised topic models. In J. Platt, D. Koller, Y. Singer, and S. Roweis (Eds.), Advances in Neural Information Processing Systems 20, pp. 121–128. Curran Associates, Inc.
- Merity, S., N. S. Keskar, and R. Socher (2018). Regularizing and optimizing lstm language models. CoRR abs/1708.02182.
- Mertens, K. and M. Ravn (2014). A reconciliation of SVAR and narrative estimates of tax multipliers. *Journal of Monetary Economics* 68, S1–S19.
- Miranda-Agrippino, S. and G. Ricco (2018). The Transmission of Monetary Policy Shocks. Working paper.
- Nakamura, E. and J. Steinsson (2018). High Frequency Identification of Monetary Non-Neutrality: The Information Effect. Quarterly Journal of Economics 133(3), 1283– 1330.
- Newey, W. K. and K. D. West (1987, May). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55(3), 703–08.

- Nimark, K. P. and S. Pitschner (2019). News media and delegated information choice. Journal of Economic Theory 181, 160 – 196.
- Pang, B., L. Lee, and S. Vaithyanathan (2002). Thumbs up?: Sentiment classification using machine learning techniques. In *Proceedings of the ACL-02 Conference on Empirical Methods in Natural Language Processing - Volume 10*, EMNLP '02, Stroudsburg, PA, USA, pp. 79–86. Association for Computational Linguistics.
- Qvigstad, J. F. (2019). How to draft "good" minutes of monetary policy decisions. Economic Affairs 39(1), 44–64.
- Qvigstad, J. F. and T. Schei (2018). Criteria for "good" justifications. Working Paper 2018/6, Norges Bank.
- Ramey, V. A. (2011). Identifying government spending shocks: It's all in the timing. *The Quarterly Journal of Economics* 126(1), 1–50.
- Romer, C. D. and D. H. Romer (1989). Does monetary policy matter? a new test in the spirit of friedman and schwartz. *NBER macroeconomics annual* 4, 121–170.
- Romer, C. D. and D. H. Romer (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review* 94(4), 1055–1084.
- Stock, J. H. and M. W. Watson (2016). Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics. In J. B. Taylor and H. Uhlig (Eds.), *Handbook of Macroeconomics*, Volume 2, pp. 415–525. Elsevier.
- Svensson, L. E. O. (2010). Open-economy inflation targeting. Journal of International Economics 50(1), 155–183.
- Svensson, L. E. O. (2014). Inflation Targeting and 'Leaning against the Wind'. International Journal of Central Banking June, 103–114.
- Svensson, L. 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 September*, 385–408.
- Taboada, M., J. Brooke, M. Tofiloski, K. Voll, and M. Stede (2011). Lexicon-based methods for sentiment analysis. *Comput. Linguist.* 37(2), 267–307.
- Taddy, M. (2013). Multinomial inverse regression for text analysis. Journal of the American Statistical Association 108(503), 755–770.

- Thorsrud, L. A. (2018). Words are the new numbers: A newsy coincident index of the business cycle. *Journal of Business & Economic Statistics*, 1–17.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society. Series B (Methodological) 58(1), 267–288.

Appendices

Appendix A Additional results

Table 2. Sensitivity of the narrative surprise measure to changing the terms used to identify the factors. 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. From all unique combinations of the keywords listed in *Panel A* in the table, we randomly draw 30 sets, where each set contains one key word for each narrative dimension. We then re-do the calculations described in Section 3.2, and construct a new narrative difference series $n\tilde{d}_{a,t}^{CB,N}$, for $a = 1, \ldots, 30$. *Panel B* of the table reports the correlation (distribution) between the benchmark estimate of $n\tilde{d}_{t}^{CB,N}$ and the alternatives $n\tilde{d}_{a,t}^{CB,N}$. The last rows of the table, associated with the *Count-based* label, reports the same type of correlations. Here, however, the narrative dimensions have instead been identified by a simple count-based method, as discussed in Section 3.3.

Panel A: N	arrative dime	ension and	key word	s								
Narrative	Inflation		Labor		Exchange rate		Oil		Uncertainty		Financial	
dimension	developments		market		market		market				sta	bility
Key word(s)	vord(s) inflasjon		arbeidsledig		kronekurs		olje		usikker		kreditt	
	inflasjonsutsiktene		arbeidsledighet		valutamarkedet		oljeindustrien		usikkerhet		kredittmarkedet	
	$inflas_{jonsforventningene}$		arbeidsmarkedet		kronekursutviklingen		oljemarkedet		usikkert		kredittveksten	
Panel B: Correlations with benchmark surprise												
						w^{-}						
Method	Percentile	1	2	3	4	5	6	7	8	9	10	Average
	5%	0.46	0.48	0.44	0.42	0.42	0.40	0.44	0.53	0.53	0.55	0.46
Factor-based	50%	0.56	0.66	0.67	0.66	0.70	0.76	0.81	0.80	0.80	0.79	0.72
	95%	0.70	0.75	0.78	0.84	0.85	0.88	0.91	0.90	0.91	0.91	0.84
	5%	-0.02	-0.02	0.01	0.03	-0.00	-0.03	-0.04	-0.04	-0.03	-0.06	-0.02
Count-based	50%	0.17	0.21	0.26	0.24	0.20	0.22	0.20	0.20	0.24	0.25	0.22
	95%	0.23	0.27	0.33	0.34	0.26	0.29	0.27	0.27	0.31	0.33	0.29



Figure 5. Announcement dates are reported on the horizontal axis, where the difference between Figure 5a and 5b is due to differences in data availability. See Figure 1 for further details.



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



(c) Media spillovers and narrative contributions

Figure 6. Figure 6a reports the t-values of \hat{b}_1 in equation (1a) when $|s_t^T|$ or $|s_t^P|$ is used as the dependent variable. 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 6b reports the t-values of $\hat{\delta}_1$ in equation (1b) when either $n\tilde{d}_t^{CB,N}$ or $|s_t^T|$ is used as the treatment variable. Figure 6c shows $s_t^{\tilde{T}}$ from (1b) decomposed into narrative contributions.



(a) 3-month interest rate

(b) Stock market



Figure 7. The figures report the estimates of $\hat{\phi}$ from equation (1c), for $h = 0, \ldots, 24$ months. The mean estimates and 95 percent 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 s_t^{narr} or s_t^P .



Figure 8. The figures report the estimates of $\hat{\phi}$ from equation (1c), for $h = 0, \ldots, 24$ months. For the s_t^{narr} impulse, the mean estimates and 95 percent confidence bands are reported using Newey and West (1987) corrected standard errors. All responses are normalized to one standard deviation of the original shock, and to increase the 3-month interest rate on impact. $s_t^{narr,a}$, for $a = 1, \ldots, 30$, denotes the alternative narrative monetary policy shocks, identified using 30 different key word combinations. Their macroeconomic responses are graphed in dotted gray lines in the figures.