WORKING PAPER

Components of uncertainty

NORGES BANK

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.

Components of Uncertainty[*]

Vegard Høghaug Larsen[†]

April 18, 2017

Abstract

Uncertainty is acknowledged to be a source of economic fluctuations. But, does the type of uncertainty matter for the economy's response to an uncertainty shock? This paper offers a novel identification strategy to disentangle different types of uncertainty. It uses machine learning techniques to classify different types of news instead of specifying a set of keywords. It is found that, depending on its source, the effects of uncertainty on macroeconomic variable may differ. I find that both good (expansionary effect) and bad (contractionary effect) types of uncertainty exist.

JEL-codes: D80, E32, E66

Keywords: Newspaper, Topic model, Uncertainty, Business cycles, Machine learning

1 Introduction

A large and growing literature investigates the effect of elevated uncertainty on aggregate macroeconomic fluctuations. Most uncertainty measures tend to be countercyclical, and several studies document that an increase in uncertainty is followed by worsening economic conditions, see e.g., Bloom (2009), Jurado et al. (2015), and Baker et al. (2016). Common to these studies is the construction of uncertainty measures that capture similar types of events related to episodes of financial and economic distress. However, measuring uncertainty that consistently rises in bad times makes it difficult to study potential alternative effects of uncertainty.

This paper offers an identification strategy to disentangle different types of uncertainty. It relies on machine learning techniques to uncover the content of a large set of news articles published in a daily business newspaper. It is shown that depending on the source, uncertainty may have different effects on the same macroeconomic variables. The latter is shown in a structural VAR model.

Specifically, I create measures of uncertainty by first classifying news articles according to theme, and then quantifying uncertainty by the count of uncertainty terms within the different types of news. The method I use belongs to the field of topic modeling, where the objective is to identify hidden patterns in textual data. I estimate the content of news articles using Latent Dirichlet Allocation (LDA), introduced by Blei et al. (2003). The method is an unsupervised learning algorithm, meaning that there is no pre-training of the model or labeling of the news articles before the classification. I identify well-defined uncertainty measures related to categories of high economic relevance such as *Oil price*, *Monetary policy*, *Politics*, and *Stock market*.

Using textual data to extract uncertainty has become popular:[1] For instance Alexopoulos and Cohen (2009) create an uncertainty measure based on the number of New York Times articles about both uncertainty and economic activity. Baker et al. (2016) create Economic Policy Uncertainty (EPU) indices for various countries by counting articles about uncertainty, the economy, and policy. Common to these papers are that they classify articles by a set of pre-determied keywords, and if an article contains words from all categories, it contributes to the index.[2] In contrast, I propose to use a topic model to classify different types of news instead of specifying a set of keywords. An advantage of using the topic model is that the classification does not rely on the article containing a particular set of words. Instead, the mixture of all the words in an article provides information on the theme of that article.[3] I use news articles for more than 28 years from

---

[1]Using text as data has exploded in the recent years, see Gentzkow et al. (2017) for an overview.

[2]Baker et al. (2016) also identify narrower category-specific uncertainty measures by counting articles with words from specific categories such as *National security* and *Health care.*

[3]A related paper using machine learning techniques to extract uncertainty is Manela and Moreira (2016).

Norway's largest business newspaper, *Dagens Næringsliv*. Very few papers in economics use a topic model to extract information from textual data. A related paper is Larsen and Thorsrud (2015) who create a topic-based news index, where the index is used to study the impact of news and noise shocks on the business cycle in Norway. Another example is Hansen et al. (2014), who study how transparency affects monetary policymakers deliberations by using a topic model to classify textual data from the Fed.

I investigate the validity of this topic-based approach by evaluating the uncertainty measures in two ways: First, I do a narrative exercise evaluating whether the uncertainty measures capture known historical events where we expect uncertainty to be high. Second, I compare the uncertainty measures to other proxies for uncertainty such as the US VIX, realized stock market volatility in Norway, and some of the economic policy uncertainty measures created by Baker et al. (2016). Overall, the topic-based measures capture well known historical events and there is a tendency for positive correlations between the topic-based measures and the alternative ones.

Despite the fact that the topic-based measures capture uncertainty in relation to different categories, common periods of high uncertainty, such as during the Global Financial Crisis, are reflected in many of the measures. To examine if distinct types of uncertainty underlie the topic-based measures, I next construct a few orthogonal uncertainty measures using principal component analysis (PCA). Although the PCA makes orthogonal components, it does not assign any label to the components. To be able to give the components an interpretation, I give them a label based on which topics they correlate most with. Doing so I find that the components are related to "economic and financial distress", "the institutional framework of monetary policy", "Norway's relationship with the EU", and "technology and firm expansion".

Uncertainty shocks can have real and substantial negative effects on firm investment and hiring, because firms delay taking action. This is often referred to as "wait and see" behavior, see e.g., Bernanke (1983), McDonald and Siegel (1986) and Bloom (2009). Uncertainty also affects households: Elevated uncertainty can increase precautionary savings and thereby deflate aggregate demand in the economy, see, e.g., Basu and Bundick (2012), Leduc and Liu (2015) and Fernandez-Villaverde et al. (2011). Uncertainty can affect financial markets, where higher firm risk leads to increased cost of capital and more cautionary behavior by investors, see, e.g., Gilchrist et al. (2014) and Arellano et al. (2010). On the other hand, some papers argue for a positive effect of uncertainty, so called "growth options" theories, where willingness to invest can increase due to an improved upside in the economy, see e.g. Segal et al. (2015) and Kraft et al. (2013).

There is a large literature estimating the economic response to an uncertainty shock. There are papers analyzing different, but correlated, types such as macroeconomic uncertainty (Bloom (2009) and Jurado et al. (2015)), economic policy uncertainty (Baker et al. (2016)), and fiscal policy uncertainty (Fernandez-Villaverde et al. (2015)). I analyze the

impact of shocks to different (orthogonal) uncertainty measures, on aggregate economic fluctuations.[4] I find that different types of uncertainty have different implications for the economy. A shock to uncertainty related to "economic and financial distress" foreshadows declines in investment in line with previous studies. The effect is sizable and economically important. I find no effect on the Norwegian economy after an uncertainty shock related to "the institutional framework of monetary policy", while an uncertainty shock related to "Norway's relationship with the EU", gives a large and persistent decline in GDP. A shock to uncertainty related to "technology and firm expansion", leads to a significant increase in GDP. The finding that uncertainty can have both positive and negative effects indicates that both good and bad types of uncertainty exist.

The rest of the paper is organized as follows: Section 2 describes the newspaper data, the topic model and how the uncertainty measures are constructed. Section 3 discusses and evaluates the uncertainty measures. Section 4 creates orthogonal uncertainty components. In Section 5, investigates the effect of uncertainty shocks on aggregate macroeconomic variables. Section 6 concludes.

2   Measuring category-specific uncertainty

This section describes the newspaper data and how the articles are classified according to their underlying content. I describe how the uncertainty of the articles are quantified as well as how these measures of uncertainty is combined with the classification of the news articles to create topic-based measures of uncertainty.

2.1   The newspaper data

The raw data used are articles from *Dagens Næringsliv*, which is Norway's largest business newspaper and also the fourth largest newspaper overall. I use all articles published in the paper version of the newspaper from May 2 1988 to December 31 2016. The data consist of close to 500 000 articles, spread over a period of more than 8000 days. This is a large amount of data that are highly unstructured, and in line with the literature on modeling text, several steps are performed to clean and reduce the data to a more manageable form. First, I remove words that would not convey any important meaning for the underlying theme of a news story, examples of such words are *the*, *is*, and *are*. I also remove common Norwegian surnames and given names. Next, each word is reduced to its word stem.[5] Lastly, I calculate a corpus measure called the *tf–idf* score which

---

[4]For Norway, Gudmundsson and Natvik (2012) create an uncertainty measure in the same way as Alexopoulos and Cohen (2009) and find negative effects of uncertainty shocks on consumption.

[5]The word stem is the part of a word that is common to all the word's inflections, an example is the word *production*, which has the word stem *produc*.

stands for term frequency – inverse document frequency. This is a way of scoring all the words in the corpus based on how important they are in explaining single documents, relative to how frequently the word occurs in the whole text corpus. I select a cutoff for this *tf–idf* score and discard the words with the lowest relative importance in explaining single documents.[6] I keep around 250 000 of the stems with the highest *tf–idf* score, and move on to the classification using the LDA.[7]

2.2   Latent Dirichlet Allocation

The LDA is a model that allow sets of observed documents to be explained by latent structures that explain why documents belong together. It is an unsupervised learning algorithm, meaning that there is no labeling of the articles or training of the model before the articles are classified. It is assumed that all documents are constructed by combining a given set of themes or topics and then drawing words from these topics. Each article is a random mixture of all the topics. The word "topic" is used frequently in this paper and it refers to a distribution over a fixed vocabulary. All the observed words in the newspaper have a positive probability of occurring in all the topics, and all the topics occur with a positive probability in all of the documents. The LDA is a generative model that works as follows:

1. Pick the overall theme of an article by randomly giving it a distribution over topics

2. For each word in the document

    i) From the topic distribution chosen in 1., randomly pick one topic
    ii) Given that topic, randomly choose a word from this topic

Iterating the second step generates a document, while iterating both the first and the second step generates a collection of documents. This is the way we imagine the documents were generated, but in reality, we only observe the outcomes, the published news articles. We use this model of how the articles were generated, together with the realized articles to infer the underlying topic structure. The estimation of the topics is done by starting out with a given set of word distributions where the probabilities of the different words occurring are random. Then we improve these distributions by changing

---

[6]Calculating the *tf-idf* score is not absolutely necessary since the LDA does a similar job when selecting the relevant words for the various topics. The main reason for doing this is to reduce the number of words in the corpus, to ease the computational burden when estimating the LDA.

[7]The corpus reduction and cleaning are standard in the text literature, maybe with the exception of removing the surnames and given names. This choice is made because many persons share the same names, and names often occur in the newspaper; including them will only pollute the underlying meaning of the article since the algorithm gives the same "meaning" to all unique names.

Figure 1. Examples of topic distributions

(a) Macroeconomics

(b) Monetary policy

Note: The 150 words with the highest probabilities are shown, the size of the words corresponds to the probability of that word occurring in the topic distribution. All the word clouds are available at http://www.vegardlarsen.com/Word_clouds/.

the probabilities and evaluating how well they describe the documents. I use a Bayesian approach to estimate the topic model using Gibbs simulations. The estimation procedure follows the algorithm described in Griffiths and Steyvers (2004), and additional details can be found in Appendix A. The topic model is estimated on data up until 2015, and the last two years of data are classified using the previously estimated topics.

Before estimating the topic model, I need to specify the number of topics to be identified, and I set $N^{\text{topics}} = 80$. What makes 80 the right number? I use a model measure called *perplexity* to compare different choices of N^{topics}. The *perplexity* is a predictive likelihood and measures how well the topic model predicts the data. I find that 80 topics are preferable to fewer topics. The goal is not to find the topic model that best describes the documents, but rather a model that delivers topics that give a reasonable description of the newspaper and the Norwegian economy. Increasing the number of topics would likely improve the *perplexity*, but would also give us topics with a narrower meaning. I found that 80 topics gave a good result, where the topics were neither too broad nor too narrow. Chang et al. (2009) show that improving the *perplexity* of a topic model by e.g. increasing the number of topics can lead to semantically less meaningful topics. Increasing the number of topics is also problematic computationally.

The output from the topic model is two sets of distributions: one set of distributions over words, denoted by $\boldsymbol{\theta}_j$, for all topics $j \in \{0, N^{\text{topics}}\}$, and one set of distributions over topics, denoted by $\boldsymbol{\varphi}_i$, for all articles in $i \in \{0, N^{\text{articles}}\}$. In this model both $\boldsymbol{\theta}_j$ and $\boldsymbol{\varphi}_i$ come from a Dirichlet distribution, giving rise to the name Latent Dirichlet Allocation.

I get 80 distributions over words, $\boldsymbol{\theta}_i$, one for each of the topics i. Figure 1 shows two examples where the word distributions are represented as word clouds. The size of the word in the word cloud corresponds to the probability of that word occurring in the given topic. The topics are given by the word distributions, and are not given any label by the topic model. Since referring only to topic numbers gives very little meaning, and

since I want an economic interpretation of the different topics, I label the topics. The labeling is done by visual inspection of the word distributions and then picking a word that gives a reasonable description of the distribution. Most topics convey a clear theme or category. A list of all the 80 topics and their labels, together with a list of the 10 most frequent words occurring in each topic, is given in Table 5 in Appendix B. I get topics related to the aggregate economy such as *Macroeconomics* and *Monetary policy*, topics related to financial markets such as *Banking* and *Funding*, topics related to politics such as *Politics* and *Elections*, international topics such as *USA* and *Asia*. I plot four examples of the topic distributions, φ_j, in Figure 3 in Section 2.4. These distributions tell us how important the different news topics are in describing single news articles.

An alternative approach to classify an article is by identifying specific keywords that are linked to specific categories. By searching through all the articles and looking up these keywords, we can classify the articles according to some pre-specified categories. This is the approach taken by Baker et al. (2016), and I follow this approach in creating an index for Norway to compare against the category-specific uncertainty measures that is the focus in my paper.[8]

2.3 Quantifying uncertainty

I create measures of uncertainty by combining two types of information about the news articles: First, the topic model allows me to classify all the textual content of the newspaper as probability distributions over news categories. Second, I calculate a measure of uncertainty for all the articles in the sample. To quantify the extent to which a news article signals uncertainty, I count the terms related to uncertainty within that article.

I start out by counting the term *uncertain* and its inflections for all the articles.[9] The count of uncertainty terms in article i is given by

$$v_i = \text{number of uncertainty terms in article } i. \tag{1}$$

To control for a varying amount of news coverage over time, I keep track of the total number of words in article i given by:

$$\omega_i = \text{number of total words in article } i. \tag{2}$$

---

[8]The details of this Norwegian version of the Baker et al. (2016) index can be found in Appendix C.

[9]The words that are counted (given in Norwegian): *usikker, usikre, usikkert, usikkerhet, usikkerheter, usikkerheten, usikkerhetene.* I have also experimented with using a broader list of words including terms such as *risk* and *unpredictability*, and this gives indices that lie close to the ones created in this paper. An advantage of using a broader list is that more articles get a non-zero uncertainty term count, which gives us a richer measure. I choose to use only the terms directly affiliated with uncertainty to create a measure that is clean and easy to interpret.

Figure 2. Aggregate newspaper uncertainty
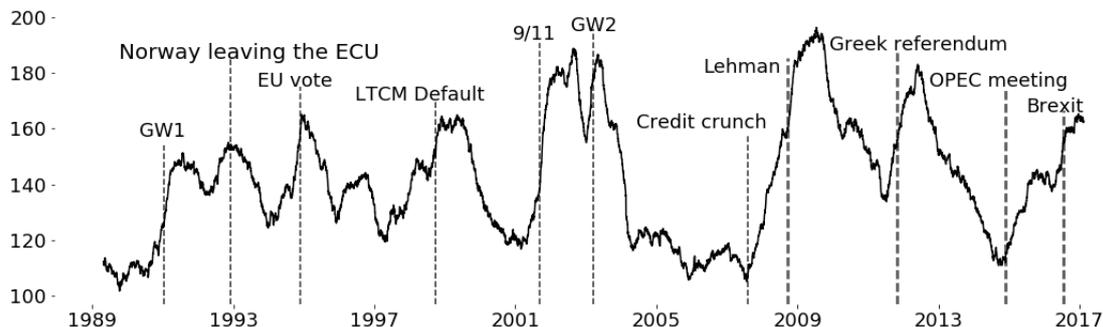
Note: The black line plots the 300 day backward-looking rolling mean. The series gives the share of uncertainty terms per 1 000 000 words in the newspaper.

Then, as a first approach, I calculate an aggregate daily measure, that is the overall uncertainty count in the newspaper each day. Calculating an aggregate count reflects uncertainty about many different underlying concepts, such as sports, the economy, political elections etc. Even if the interpretation of this aggregate uncertainty measure is unclear, it is a point of departure, before looking at the more disaggregated measures. I calculate the aggregate uncertainty measure as follows:

$$\Upsilon_t^{\text{Agg}} = \sum_{i \in \text{ day } t} \left(\frac{\upsilon_i}{\omega_i} \right). \tag{3}$$

On each day, the total count of the uncertainty terms are divided by the total word count that day. Figure 2 plots this aggregate measure as the 300 days backward-looking mean.[10] Over the sample the total daily count of the word *uncertain* and *uncertainty* vary approximately between 100 and 200 out of one million words. From the figure we see that there are large variations in the uncertainty measure and that there are clear episodes where aggregate uncertainty is high. I plot some events that coincided with significant increases in uncertainty. Based on these events it appears that the uncertainty count in *Dagens Næringsliv* is driven mostly by foreign crises such as wars and international financial crises. The episodes that are displayed are the first and the second gulf war (GW1 and GW2), the Long-Term Capital Management (LTCM) default, the 9/11 terrorist attacks, the credit crunch (often considered as the start of the financial crisis), the collapse of Lehman Brothers, the Greek proposed referendum related to a bailout of the Greek government, the OPEC meeting in fall 2014 after the large drop in oil prices, and the UK deciding to leave the EU (Brexit). The only Norwegian events displayed in

---

[10]The reason for plotting the backward-looking mean is because it reduces noise and makes it easy to identify episodes when uncertainty was high. Also, I utilize the data at a daily frequency. This is done for visual clarity, and all empirical results presented are based on the measures at a daily, monthly or quarterly frequency.

Figure 3. Topic distributions for four random articles
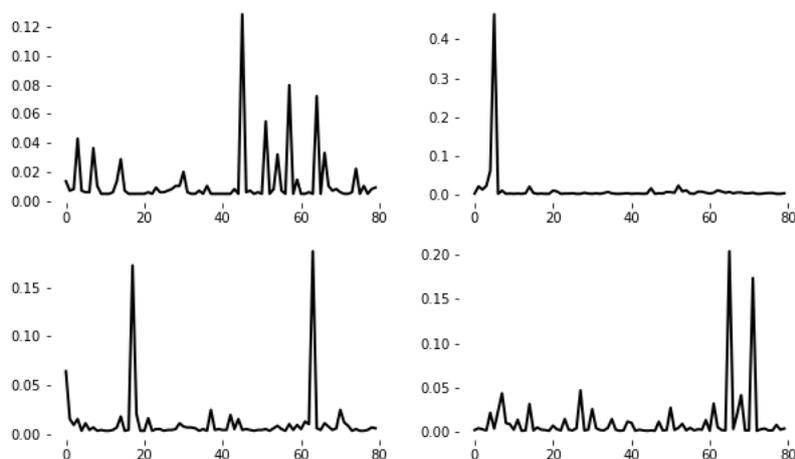
Note: The topic distributions, $\boldsymbol{\varphi}_i$, of four randomly drawn articles, the numbers on the x-axis represent the topics and the corresponding label can be found in Table 5 in Appendix B.

the figure are the referendum on joining the European Union, and Norway depegging its currency from the European Currency Unit (ECU). Of course many of the episodes where uncertainty is high in the figure coincide with other Norwegian events such as the banking crisis in the early 1990s, and a short recession in 2002–2003 and 2008–2009.

2.4   Topic-based measures of uncertainty

The category-specific uncertainty measures are created based on the uncertainty count within the categorized news articles. The topic model delivers the classification of all news articles. This classification is given as a probability distribution over all topics reflecting content in the articles that relates to several topics at once. I calculate an uncertainty measure for all the different news topics. This is done by weighing the uncertainty counts by the relative contribution of all articles to the different topics. That is, article i has an uncertainty count given by v_i, which then contributes by $\boldsymbol{\varphi}_i(\text{topic} = j)$ to topic j. To see what these topic distributions, $\boldsymbol{\varphi}_i$, may look like, Figure 3 plots such topic distributions for four news articles. These distributions tell us how much the uncertainty count from the articles they represent contributes to the uncertainty indexes for the various topics. We see that for some articles there is one or a few topics that explain the content of the article, while others are a broader mix of topics. *Dagens Næringsliv* is a business newspaper, as can be seen by the large share of topics that relate to business and economics.[11] Thus, an article about the economy is likely to be a mix of economy-related topics. On the other

---

[11]The topics are listed in Table 5 in Appendix B.

Table 1. Uncertainty share in different news categories

Top 10	# of words per 1 mil.	Bottom 10	# of words per 1 mil.
Monetary policy	5.7	*Drinks*	0.9
Stock market	4.7	*Movies/Theater*	0.9
Macroeconomics	4.4	*Food*	1.0
Fear	3.7	*Literature*	1.0
Oil price	3.2	*Music*	1.0
Debate	2.9	*Art*	1.0
Negotiation	2.4	*Sports*	1.1
Results	2.4	*Family business*	1.1
Oil production	2.3	*Watercraft*	1.1
Elections	2.3	*Tourism*	1.1

Note: The average number of uncertainty terms used in the different types of news.

hand, there are very few topics related to sports, so a sports-related article is more likely to be described by few topics.

The total amount of content in a newspaper varies over time, as does the coverage of an individual news topic. To control for this, I normalize with respect to the amount of news content on any given day. The more articles and words we observe in one day, the more uncertainty terms we expect to observe in total. For the baseline normalization, I divide the topic-specific uncertainty term count within one day by the total number of words that day.[12] This uncertainty measure is given by:

$$\Upsilon_{j,t} = \frac{\sum_{i \in \text{ day } t} \upsilon_i \boldsymbol{\varphi}_i(\text{topic} = t)}{\sum_{i \in \text{ day } t} \omega_i}. \tag{4}$$

One alternative specification is to divide by the total number of words used within a specific news category.[13] The denominator is important to consider because it in itself causes fluctuations in the uncertainty measure. I choose to use the normalization in Equation 4 as the baseline, because fluctuations in the coverage of a given news topic can vary substantially. Daily fluctuation in topic coverage can have large effects on the alternative uncertainty measure, and this variation is not driven by the uncertainty count.

What types of news categories use the uncertainty terms the most? Table 1 reports the 10 news categories with the largest number of uncertainty terms, and also the 10

---

[12]Dividing by the total daily count is in line with the literature, see e.g. Baker et al. (2016).

[13]This alternative measure is calculated as

$$\tilde{\Upsilon}_{j,t} = \sum_{i \in \text{ day } t} \left(\frac{\upsilon_i}{\omega_i} \right) \boldsymbol{\varphi}_i(\text{topic} = t).$$

I have also computed these measures, and for most topics they give very similar results. The average correlation between the two measures is 0.86.

news categories with the lowest count. The news category where the newspaper writes the most about uncertainty is *Monetary policy*. During the period studied, Norway had five different monetary policy regimes, and this may have led to increased uncertainty. The news category with the second highest uncertainty count is *Stock market*, followed by *Macroeconomics*, *Fear*, and *Oil price*. The *Fear* topic is a news topic where the word *uncertainty* is one of the words with the highest probability and the frequency of the *Fear* topic itself is as a possible proxy for uncertainty. On the other hand, the type of news where the uncertainty terms are the least frequent are *Drinks*, *Movies/Theater* and *Food*.

3 Evaluating the topic-based uncertainty measures

The topics are identified by an unsupervised learning algorithm, and uncertainty is identified as the frequency of uncertainty terms within news related to various topics. There is no subjectivity involved other than the labeling of the topics, and this is only a way of referring to the underlying word distributions. I evaluate whether the uncertainty measures capture what they are supposed to, which is the underlying uncertainty related to various themes or categories. I do a narrative exercise where I plot some of the uncertainty measures together with episodes where it is reasonable to think that uncertainty was high. To conserve space, I discuss eight of the 80 measures.[14] The first four measures are selected based on the type of news that uses the uncertainty terms the most. The remaining four measures are based on news categories that I think are easy to link to well-known historical events. An example is oil price uncertainty, which we expect to be high during episodes with, say, conflicts in regions that produce oil. In the end I evaluate the full set of measures by comparing them to other proxies for uncertainty.

3.1 Narrative evaluation

The first four examples of category-specific uncertainty are chosen by selecting the news topics where the uncertainty terms are used with the highest frequency. These topics are *Monetary policy*, *Stock market*, *Macroeconomics*, and *Fear*. The top four measures are plotted in Figure 4 together with some notable events where it is reasonable to think that uncertainty was high. The exact dates and a short description of the events can be found in Table 6 in Appendix B.

In Panel (a) in Figure 4, I plot the measure for *Monetary policy* uncertainty. It is plotted together with the dates when the monetary policy regime changed, as well as when a new central bank governor assumed office. We see that uncertainty tends to be elevated around these events. Uncertainty was especially high during the second part of the 1990s. This was a period when Norway had a debate on what monetary policy regime

---

[14]Figures for all the 80 topics can be found at http://www.vegardlarsen.com/index.php/ui.

Figure 4. Examples of uncertainty measures from uncertainty heavy news

(a) *Monetary policy* uncertainty
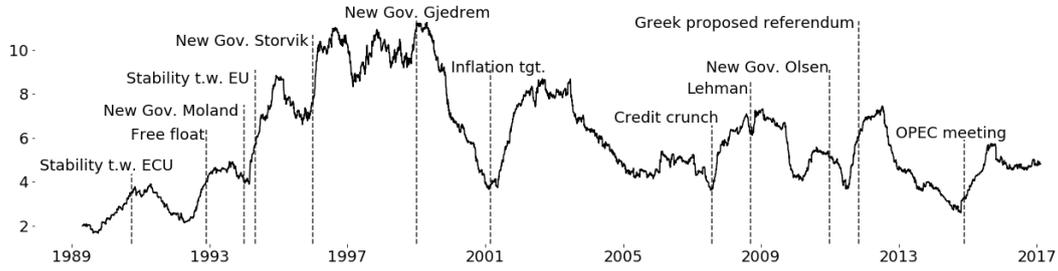
(b) *Stock market* uncertainty
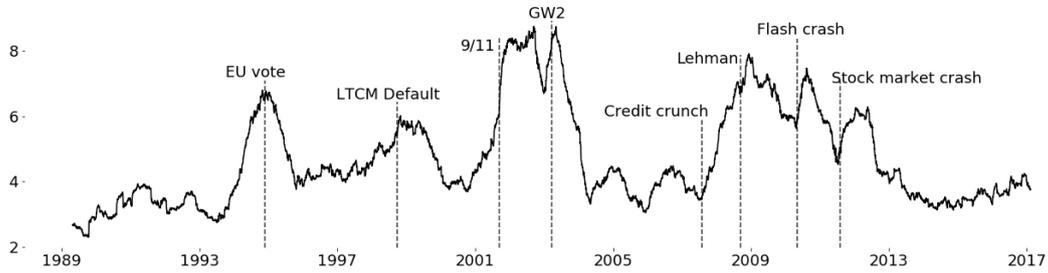

## (c) *Macroeconomics* uncertainty
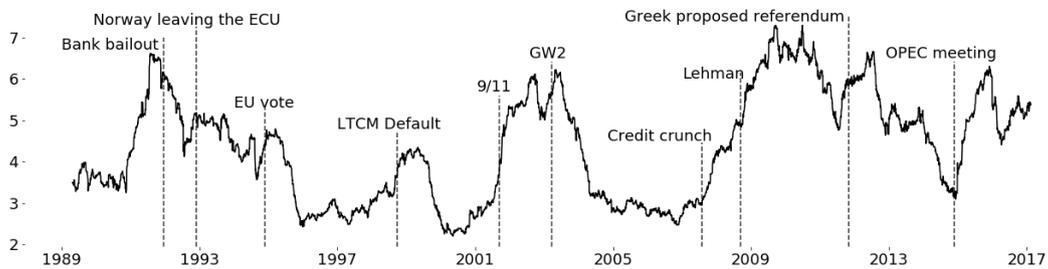
(d) *Fear* uncertainty
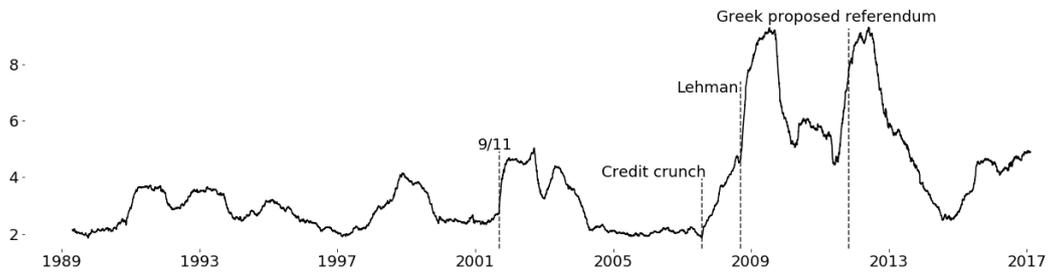
Note: The black line plots the 300 day backward-looking mean. The uncertainty count is the number of uncertainty terms per 1 000 000 words in the full newspaper. Uncertainty heavy news refers to news topics where the uncertainty terms are used with the highest frequency, see Table 1.

Figure 5. Examples of uncertainty measures from distinct types of news

(a) *Oil price* uncertainty
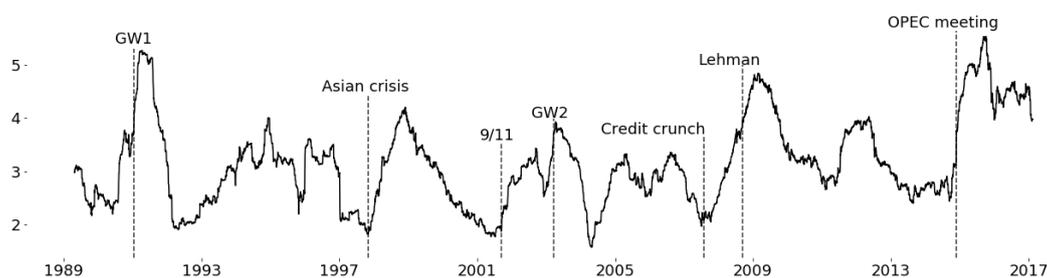
(b) *Telecommunication* uncertainty
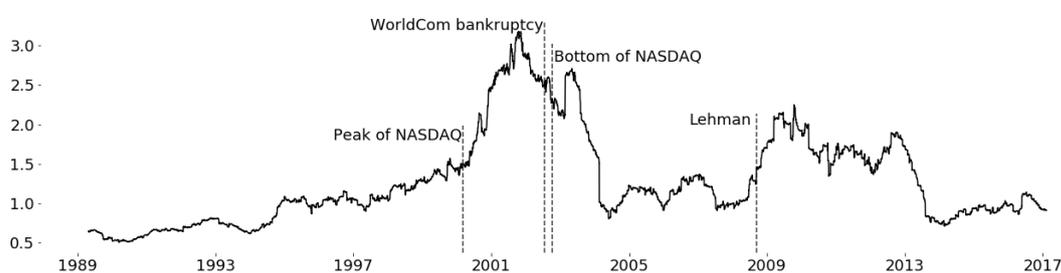
(c) *International conflicts* uncertainty
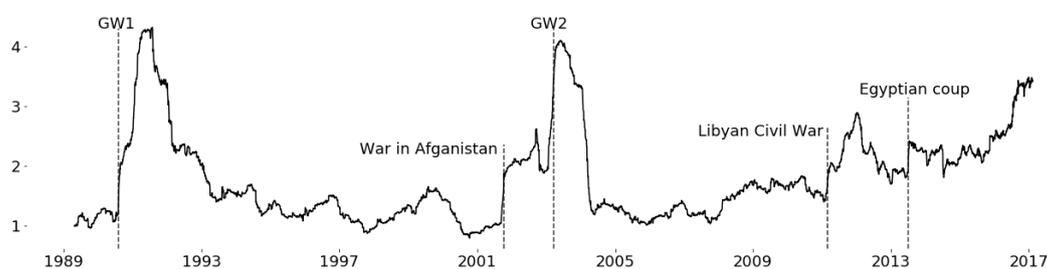
(d) *Politics* uncertainty
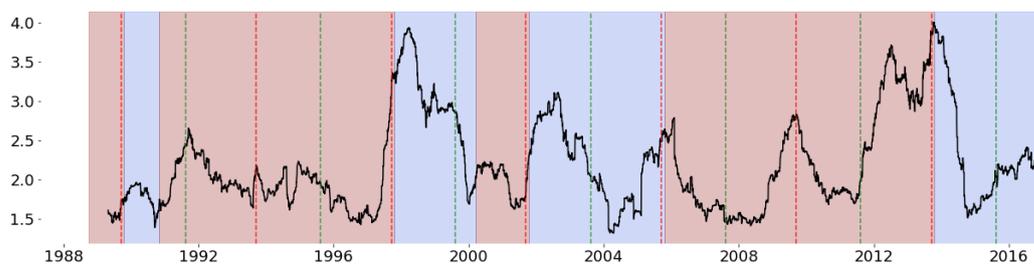
Note: For details see Figure 4. In Panel (d) the vertical red dashed lines represent parliamentary, and the green dashed lines, local elections in Norway. The areas shaded in red represents periods with a left-leaning government, and the blue shaded areas represent right-leaning governments.

that should be implemented. The monetary policy regime in Norway changed four times during the sample studied here.[15] Uncertainty also increase during global events such as the Lehman Brothers bankruptcy and the Greek government-debt crisis. In Panel (b) Figure 4, the series for *Stock market* uncertainty is plotted. This measure captures well-known events of heightened uncertainty, such as the debate in Norway on whether or not to join the EU, the LTCM default, the short Norwegian recession in the early 2000s, and the Global Financial Crisis. *Stock market* uncertainty tends to increase when the stock market is in decline. In, Panel (c) in Figure 4, I plot the *Macroeconomic* uncertainty measure. This series captures many of the same events as *Stock market* uncertainty, but there are a few periods where the two measures diverge: First, the *Macroeconomics* measure captures more uncertainty in the early 1990s during both the Norwegian banking crisis and the episodes of changing monetary policy regimes. Second, we see a large surge in *Macroeconomics* uncertainty after the oil price fall that started in the summer of 2014. The *Macroeconomics* measure is countercyclical and has a negative correlation with the business cycle. Panel (d) in Figure 4 plots the frequency of uncertainty terms within news classified as *Fear*. The *Fear* topic is a type of news that gets considerable coverage during a crisis. The measure is especially high during the Global Financial Crisis and the Greek government-debt crisis.

The first four examples often capture related events. The average correlation between the four measures at a daily (quarterly) frequency is 0.28 (0.43). Turning to measures capturing more distinct types of uncertainty, I plot in Figure 5 uncertainty related to *Oil price*, *Telecommunication*, *International conflicts*, and *Politics* as examples. The average correlation between these four measures is 0.13 (0.20) at a daily (quarterly) frequency.

Panel (a) in Figure 5 displays the series for *Oil price* uncertainty, which is important for Norway being a large oil exporter, cf. Bjørnland and Thorsrud (2016). By inspecting the spikes in *Oil price* uncertainty, it looks like they are driven mostly by foreign events, often related to unrest in the Middle East or global financial crises. Hamilton (2013) identifies historical oil shocks, and note that all his shocks coincide with elevated *Oil price* uncertainty. In Panel (b) of Figure 5, I plot uncertainty related to *Telecommunication*. The 1990s was a period of rapid technological advancements in the IT sector. The value of IT companies on stock markets all over the world rose rapidly. The NASDAQ index, which is a US-based, technology-heavy index, grew from below 1000 in 1995 to over 5000 in 2000. This surge of IT companies was part of the buildup of the dot-com bubble. The *Telecommunication* uncertainty measure grows with the NASDAQ, up until the peak of the index in March 2000, which then is followed by a large buildup of uncertainty. *Telecommunication* uncertainty stayed high for some time after NASDAQ dropped, before the uncertainty measure fell to pre-1995 levels in late 2003. Panel (c) in Figure 5 shows

---

[15]For details on the history of monetary policy in Norway, see the speech by Gjedrem (2008).

uncertainty related to *International conflicts*. The series picks up well-known conflicts such as the first and second Gulf War, and several episodes during the Arab spring. The uncertainty measure is especially high during the first and second Gulf Wars, which likely get much coverage in the business newspaper due to the effect on the oil price.

Lastly, in Panel (d) in Figure 5, I show uncertainty related to *Politics*. The dates for the parliamentary elections and local elections in Norway are indicated by the red and green dashed lines respectively. I also indicate whether there is a left-leaning (red) or right-leaning (blue) central government in office. The uncertainty tends to increase around the parliamentary elections.[16]

3.2 Comparison to alternative uncertainty measures

I compare the topic-based uncertainty measures to some alternative measures of uncertainty. There is limited availability of uncertainty measures for Norway so I generate two alternative measures: First, Norway has no options-based stock market volatility index, and I calculate a realized stock market volatility (RSMV) measure. The RSMV series is calculated as the monthly standard deviation of the Oslo stock exchange benchmark index (OSEBX). The second measure, is a Norwegian version of the EPU created by Baker et al. (2016). The details on how the Norwegian EPU is computed can be found in Appendix C. In addition, I look at seven foreign measures. Those are: the US VIX, the macroeconomic and financial uncertainty measures from Jurado et al. (2015), and the EPU measures for the US, the UK, Europe and China, created by Baker et al. (2016).[17]

Figure 6 displays the correlations between all the 80 topic-based measures and the nine alternative ones. The figure is a heat map where negative correlations are in shades of red, and positive correlations are in shades of blue. The highest correlation, 0.71, is between the *Fear* measure and the US EPU. This observation is placed in the top left corner of the heat map, and I sort the rows and columns in descending order away from this point. The lowest correlation, -0.25, is between the *Europe* and the RSMV measure. Some notable results emerge from Figure 6:

First, almost all the topic-based measures have a positive correlation with the alternative ones. This indicates that most of the measures, both the topic-based and the alternative ones, capture similar events. Two notable exceptions are the *EU* and *Europe*

---

[16]We also see a large surge in uncertainty around the time of the local election in 2011. However, this is likely also capturing an increase in uncertainty in relation to the terrorist attacks on the government headquarters and on the Workers' Youth League summer camp on July 22.

[17]Jurado et al. (2015) create an uncertainty measure based on the unforecastable component of a large set of economic variables. The Jurado et al. (2015) paper focuses on macroeconomic uncertainty. In a related paper, using similar data, Ludvigson et al. (2015) disentangle macro and financial uncertainty. I refer to both the macro and the finance measure as JLN-measures (Jurado, Ludvigson, and Ng), because I downloaded the measures from the supplementary material from Jurado et al. (2015) (`http://www.columbia.edu/~sn2294/pub.html`).

15

Figure 6. Correlations with alternative measures
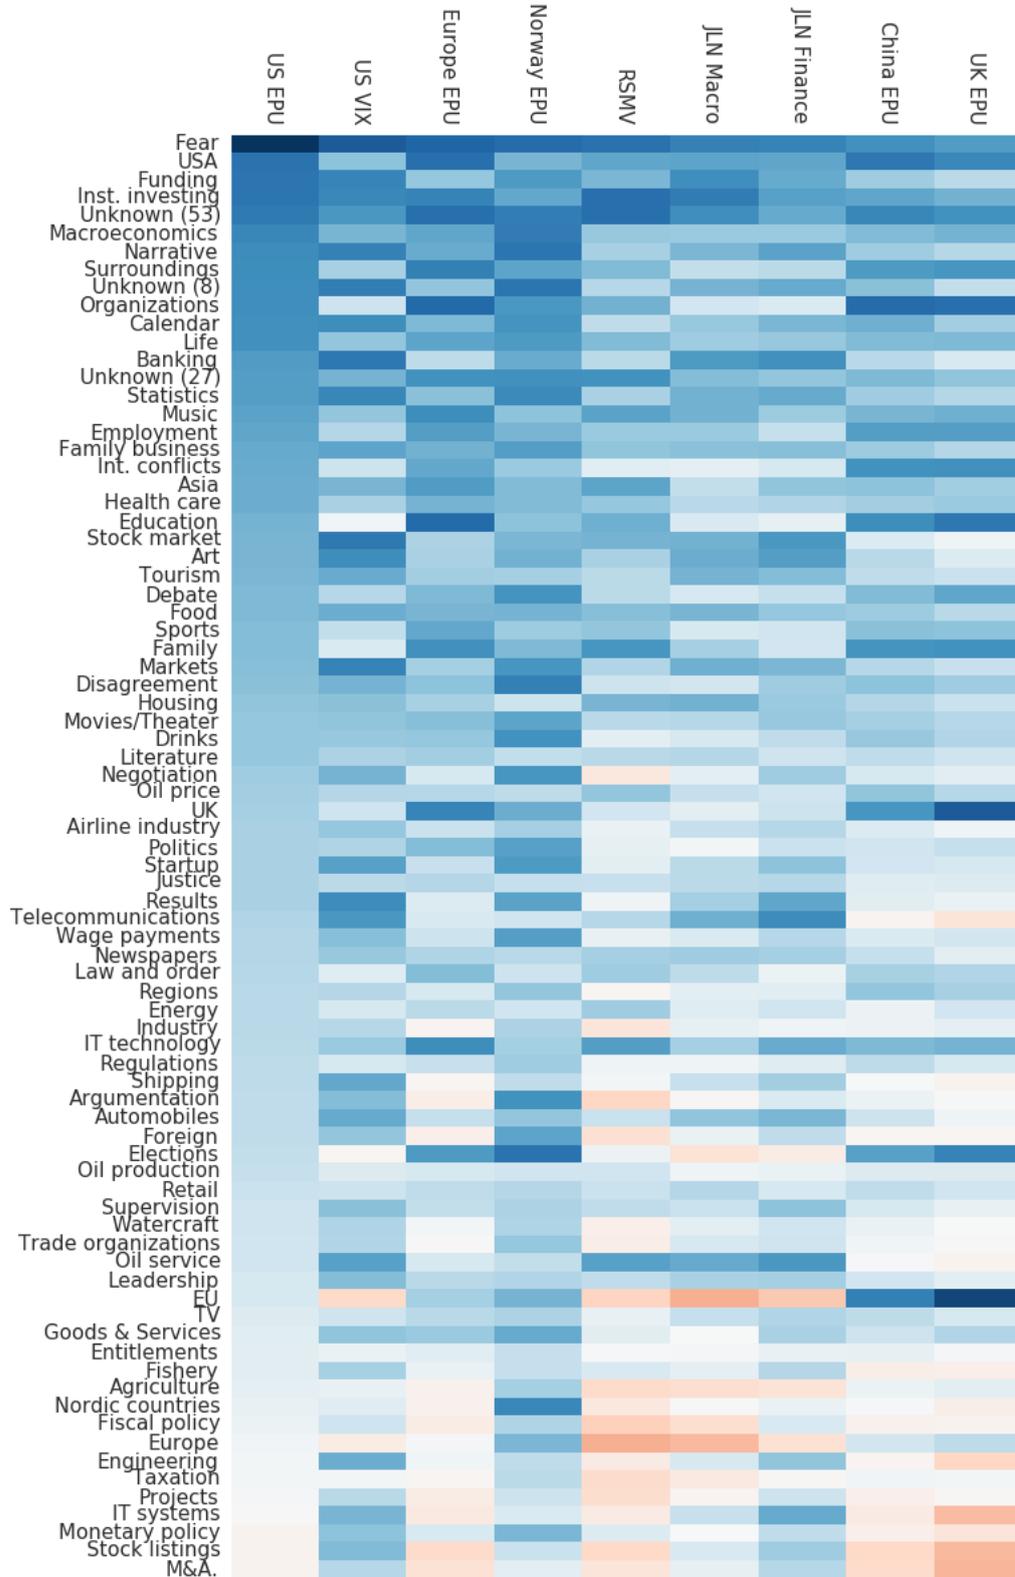
Note: The correlations are computed at a monthly frequency. Blue represents a positive correlation while red represents a negative one. The topics are sorted by the correlation with the US EPU, where the correlations range from 0.71 to -0.03.

measures, which have a negative correlation with several of the alternative measures. In part of the sample, these measures capture a Norway-specific type of uncertainty related to the referendum on membership of the European Union.

Second, the *Fear* measure captures a type of uncertainty that is common to all the alternative measures. The topic-based measures do not seem to capture much heterogeneity between the alternative ones, but there are some exceptions: Finance related measures such as *Funding*, *Banking*, and *Stock market* have a relative high correlation with the US VIX of 0.47, 0.51, and 0.51 respectively. Political measures such as *Politics* and *Elections*, on the other hand, capture more Norway-specific events and have a relatively high correlation with the Norwegian EPU of 0.38 and 0.52 respectively. The *Elections* measure has a high correlation with all the EPU measures. The *USA* uncertainty measure has a high correlation with the US EPU of 0.52, while the *UK* measure has a high correlation with both the UK EPU and the EU EPU of 0.59 and 0.65 respectively.

Third, given that the topic-based measures capture relevant types of uncertainty, the RSMV measure does not look like a good measure for uncertainty: the average correlation between the topic measures and the RSMV is 0.15. Given that no options-based volatility measure exists for Norway, a measure such as *Stock market* uncertainty can be a good alternative as a proxy for a Norwegian VIX.

The topic-based uncertainty measures do capture the type of events we expected, and different measures capture category-specific events. Most measures are positively correlated with the alternative measures of uncertainty, which suggests that there are some common components captured across uncertainty measures. This motivates an analysis of the underlying components of uncertainty.

4 The underlying components of uncertainty

In times of economic distress, the uncertainty count in most types of news tends to increase. We saw in the previous sections that during the Global Financial Crisis, uncertainty increased in many of the topic-based measures. The uncertainty measures are not orthogonal, and they have a between-topic correlation varying from -0.32 to 0.87.[18] I use Principal Component Analysis (PCA) to extract uncorrelated components from the topic-based measures. PCA is a method for reducing a set of potentially correlated variables down to a set of linearly uncorrelated variables.

I run the PCA on data at a quarterly frequency. I later include the components in a VAR using quarterly data, and extracting quarterly components ensures that the

---

[18]A heat map of all the between-topic correlations is given in Figure 10 in Appendix B. The correlations are calculated at a quarterly frequency. At a daily frequency they vary between 0.03 and 0.56.

Table 2. The component measures – descriptive statistics

	Component 1	Component 2	Component 3	Component 4
Explained Variation	34	15	9	6
Cumulative E.V.	34	49	58	64
AR(1)	0.50	0.81	0.77	0.50
Skewness	1.15	0.82	1.44	0.37
Kurtosis	1.98	0.32	2.28	1.34

Note: Fisher's definition of kurtosis is used where the kurtosis of a normal is zero. The components are normalized (mean zero and a standard deviation of one).

Figure 7. Components of uncertainty
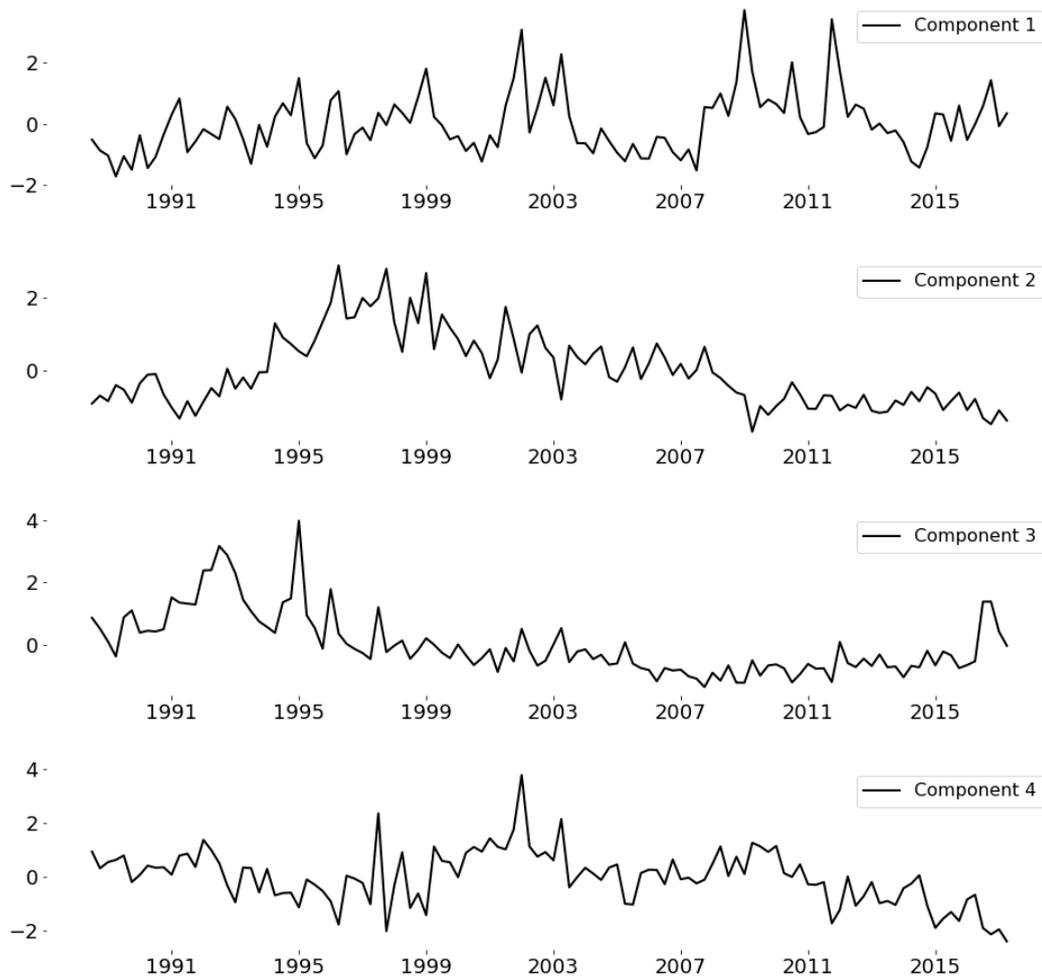
Note: A plot of the four principal components extracted from 80 topic-based uncertainty measures.

Figure 8. Correlations with alternative measures

	Component 1	Component 2	Component 3	Component 4
Norway EPU	0.69	-0.16	0.25	-0.2
US VIX	0.66	-0.033	-0.24	0.33
US EPU	0.55	-0.59	-0.094	0.0053
JLN Finance	0.52	-0.06	-0.29	0.37
EU EPU	0.48	-0.41	-0.11	-0.4
JLN Macro	0.42	-0.28	-0.41	0.34
China EPU	0.41	-0.52	0.24	-0.32
RSMV	0.38	-0.33	-0.56	-0.19
UK EPU	0.27	-0.62	0.34	-0.49

Note: The correlations are computed at a quarterly frequency. Blue is for positive correlations and red is for negative correlations. The correlation coefficients are reported in the rectangles.

components are orthogonal when included in the VAR.[19] I focus on the first four principal components, motivated by keeping the components that explain five percent or more of the total variation in the topic-based measures. The four components explain a total of 64 percent of the underlying measures. The first component is, by definition, the most important one, and it explains 34 percent of the total variation. The explained variation for all the components is given in Table 2 along with some descriptive statistics for the measures. The principal components are not identified with a sign, so whether an increase in the component measures corresponds to more or less uncertainty is not defined. To deal with this I normalize the sign of the four components so they have a positive correlation with the topic-based uncertainty measure where they have the highest correlation, see the next subsection for details. Figure 7 plots the measures using the final normalizations.

We have four distinct types of variation in the uncertainty measures, and Figure 8 reports the correlations between the component measures and the alternative ones. The first component has a positive correlation with all the alternative measures. Looking at Component 1 in Figure 7, we see that it captures well-known events of heightened uncertainty, such as the Asian crisis, the 9/11 attacks, and the collapse of Lehman Brothers. This type of uncertainty is common in all the alternative measures. The second component has a negative correlation with all the alternative measures. The third component has a positive correlation with the EPU measures for Norway, the UK and China, and a negative correlation with the rest. The fourth component has a negative correlation with all the EPU measures and a positive correlation with US VIX as well as the Jurado et al. (2015) measures.

---

[19]If I perform the PCA on daily or monthly values, and then do the aggregation, I get very similar results, but the aggregated components will not be strictly orthogonal.

Table 3. The "content" of the components

	Component 1		Component 2	
		Correlation		Correlation
1st	*Narrative*	0.87	*Monetary policy*	0.75
2nd	*Fear*	0.83	*Employment*	-0.56
3rd	*Stock Market*	0.81	*Organizations*	-0.56
4th	*Statistics*	0.81	*Macroeconomics*	-0.46
5th	*Unknown*	0.81	*Weekdays*	0.46
	Component 3		**Component 4**	
		Correlation		Correlation
1st	*EU*	0.85	*Mergers & Acquisitions*	0.55
2nd	*Europe*	0.72	*Stock listings*	0.55
3rd	*Agriculture*	0.68	*IT systems*	0.50
4th	*Argumentation*	0.59	*Engineering*	0.47
5th	*Fiscal policy*	0.55	*Telecommunication*	0.43

Note: The five topic-based measures with the highest absolute correlation with the four components.

4.1   Labeling the components

The topic-based measures, discussed in Section 3, have a direct link to the news categories, giving the measures a straight-forward interpretation. However, the measures are not orthogonal, and they often capture the same type of uncertainty. The component measures on the other hand, are uncorrelated measures, capturing different types of variation in the uncertainty count, but they have no direct link to different types of news. To deal with this, I give the components a label by relating them back to the topic-based uncertainty measures. The relationship is based on the correlation between the components and the topic-based measures. Table 3 reports the five topic-based measures that have the highest absolute correlation with the four components. The topic labels reported in Table 3 are only a way of referring to the underlying topic distributions. To get a deeper understanding of the content of the components, I show the three topic distributions with the highest absolute correlation for all the components in Figure 11 in Appendix B.

The first principal component is strongly correlated with many of the topic-based measures. The topic-based measures that have the highest correlation with Component 1 are the *Narrative* and *Fear* measures, with a correlation coefficient of 0.87 and 0.83. We see in Figure 8 that Component 1 uncertainty captures many of the well-known events the literature has focused on, see the high correlation with e.g. the US VIX in Figure 8. I label the first component as uncertainty related to "economic and financial distress".

The second principal component increases with uncertainty related to *Monetary policy*, and decreases with uncertainty related to *Employment*, *Organizations*, and *Macroeconomics*. Having topic-based measures with a high correlation of opposite signs complicates

the labeling of the component. Component 2 uncertainty was particularly high during the second part of the 1990s, see Figure 7. This was a period where there was a debate about what type of monetary policy regime that should be implemented in Norway. The uncertainty terms used in news about monetary policy can emerge for at least three reasons: First, the usage of uncertainty terms in news about monetary policy increases during economic and financial distress because there is uncertainty related to how the monetary authorities will react. This type of uncertainty is captured by Component 1, which is orthogonal to Component 2. Second, when the newspaper writes about monetary policy, the uncertainty terms are common vocabulary, without relating to an uncertain environment. Third, in the sample studied here, there have been several changes in the monetary policy regime, so part of the monetary policy uncertainty can be related to what type of monetary policy regime that will be implemented. Given that Component 2 is highly correlated with the *Monetary policy* measure, it can capture increases in uncertainty that are caused by the two last reasons. However, this does not tell us why Component 2 increase with lower uncertainty related to *Employment*, *Organizations*, and *Macroeconomics*. After 2001, there has been no changes in monetary policy regime, and if the type of news related to monetary policy has changed after 2001, this may have polluted the measure. Although the labeling of Component 2 is challenging, I select a label based on the topic-based measure with the highest correlation. I label Component 2 as uncertainty related to "the institutional framework of monetary policy".

In Figure 7 we see that Component 3 uncertainty was high during the first part of the 1990s and that the elevated uncertainty coincides with two events during this period: First, the spike in 1992 coincides with Norway depegging its currency from the ECU. Second, this was a period when Norway considered joining the EU, and the second spike in 1994 coincides with the Norwegian referendum on joining the European Union, which ended with Norway not joining. The topic-based measures with the highest correlation with Component 3 are *EU*, *Europe* and *Agriculture*.[20] Also this type of uncertainty has increased again after the UK voted to leave the EU. I label Component 3 as uncertainty related to "Norway's relationship with the EU".

The fourth component has a high correlation with *Mergers & Acquisitions*, *Stock listings*, and *IT systems*. Within the sample studied, the IT sector has grown exponentially. An example is the Norwegian telecommunications company Telenor, which is Norway's second largest company and also one of the largest telecommunications companies in the world, with 211 million mobile subscriptions.[21] I label the fourth component as uncertainty related to "technology and firm expansion".

---

[20]The agriculture sector was strongly opposed to Norway joining the EU.

[21]See "3Q 2016 Telenor Group Report", https://www.telenor.com/wp-content/uploads/2016/10/Telenor-Group-Q3-2016-report-4a2186a1d4d0e6d568e2088b637e2baa.pdf.

5 Uncertainty and the economy

A robust finding in the uncertainty literature is that uncertainty proxies are countercyclical. Bad economic times are also times of high uncertainty. I investigate the response of investment and GDP after shocks to the different components identified in the previous section.

The standard modeling framework in the literature is to estimate the effects of uncertainty shocks in a structural VAR model, using a recursive identification scheme. The main finding from these studies is that, using various proxies, uncertainty shocks are followed by a decline in real activity, see e.g. Bloom (2009), Jurado et al. (2015), and Baker et al. (2016). Using a VAR to investigate uncertainty shocks is challenging due to the endogeneity issues between the uncertainty measures and the macroeconomic variables. The impulse responses presented in this section cannot necessarily be interpreted as causal, but the goal is to learn something about how different types of uncertainty affect aggregate quantities in the economy.

I follow Baker et al. (2016) and specify a structural VAR model where the identification is achieved using a Cholesky decomposition. The uncertainty measure is ordered on top in the VAR, implying that uncertainty cannot react to the other variables within the same period. The VAR model is specified as follows

$$\mathbf{A}_0\mathbf{y}_t = \sum_j \mathbf{A}_j\mathbf{y}_{t-j} + \mathbf{B}\boldsymbol{\varepsilon}_t, \tag{5}$$

$$\text{where } \mathbf{y}_t \equiv \begin{bmatrix} \text{Uncertainty} \\ \log(\text{OSEBX}) \\ \text{Interest rate} \\ \log(\text{Investment}) \\ \log(\text{GDP}) \end{bmatrix}_t \text{ and } \boldsymbol{\varepsilon}_t \sim i.i.d.N(0,1).$$

The $\mathbf{A}_0$ matrix is lower triangular and the $\mathbf{B}$ matrix is diagonal. The topic-based uncertainty measures are available daily, but for most macroeconomic time series, data are available only at a lower frequency. I estimate a model using quarterly data so I can include investment and GDP. The variable OSEBX, is the Oslo stock exchange benchmark index, downloaded from Yahoo Finance. I include a nominal interest rate where I use the 3-month Norwegian interbank offered rate (NIBOR), downloaded from Norges Bank. I also include gross investment in mainland Norway and GDP in mainland Norway both variables downloaded from Statistics Norway. In the baseline specification, the model includes three lags and the data sample used is 1988Q2–2016Q4. I report impulse responses from a one standard deviation shock to the uncertainty measures.

As a point of reference, I start out by including the aggregate uncertainty measure,

Υ_t^{Agg}, in the VAR. Since the newspaper covers business and economics news, this measure is a proxy for the overall uncertainty related to business and economics. Figure 12 in Appendix B plots the impulse responses to investment and GDP after a shock to the aggregate uncertainty measure. There is no effect on investment nor GDP after an uncertainty shock to the aggregate measure. This motivates the next step, where I include the topic-based measures, discussed in Section 3, directly in the VAR. This approach yields a range of different responses, but given the high correlation among many of them, the responses often look similar. Figure 13 in Appendix B plots the eight most and the eight least negative responses of GDP and investment using the topic-based measures directly in the VAR. The goal in this section is to capture the effect of distinct types of uncertainty and I estimate the baseline VAR by including the component measures in the model one at a time.

First, a Component 1 shock, labeled as uncertainty related to "economic and financial distress", gives a significant fall in investment by more than one percent. The response reaches a minimum after 3–4 quarters. We also observe a fall in GDP but this effect is not significant at the 90 percent level. Part of Component 1 uncertainty is likely an endogenous response to economic and financial distress, an example is the large increase in uncertainty after the Lehman Brothers bankruptcy. But, it might also capture uncertainty that is causing economic downturns, uncertainty during the Global Financial crisis might have depressed economic growth, see e.g. Baker et al. (2016). Component 1 likely mixes together these types of uncertainty. For a study that disentangles the exogenous and endogenous part of this type of uncertainty see Ludvigson et al. (2015). Component 1 uncertainty does capture the same type of uncertainty as the literature has focused on and the negative impulse responses after a Component 1 shock do resemble those in the literature, see e.g. Bloom (2009), Jurado et al. (2015), and Baker et al. (2016).

Second, a Component 2 shock, labeled as uncertainty related to "the institutional framework of monetary policy", gives no significant responses in either investment or GDP. This type of uncertainty might be especially noisy, since the uncertainty terms are used frequently in this type of news for various reasons (see the discussion on labeling Component 2 in the previous section). The uncertainty about the institutional framework of monetary policy was prevalent in the 1990s. In 2001, Norway implemented inflation targeting, which is in place as of 2017. There has not been much discussion on moving away from this regime, and uncertainty regarding the institutional framework of monetary policy has been low in this period.[22]

---

[22]We see in Figure 7 that Component 2 uncertainty declines in late 2007 and during 2008, coinciding with the start of the Global Financial Crisis. Lower Component 2 uncertainty might be driven by increased uncertainty related to *Macroeconomics* and *Employment* in this crisis period, see Table 3. Component 2 has both a strong negative and a strong positive correlation with some of the topic-based measures, making the interpretation of the uncertainty shock difficult.

**Figure 9.** Impulse responses using the four components

(a) Component 1 − "economic and financial distress"
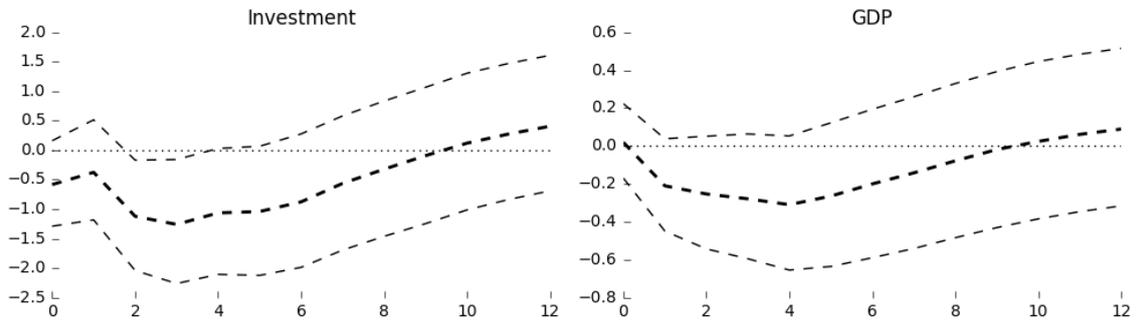
(b) Component 2 − "the institutional framework of monetary policy"
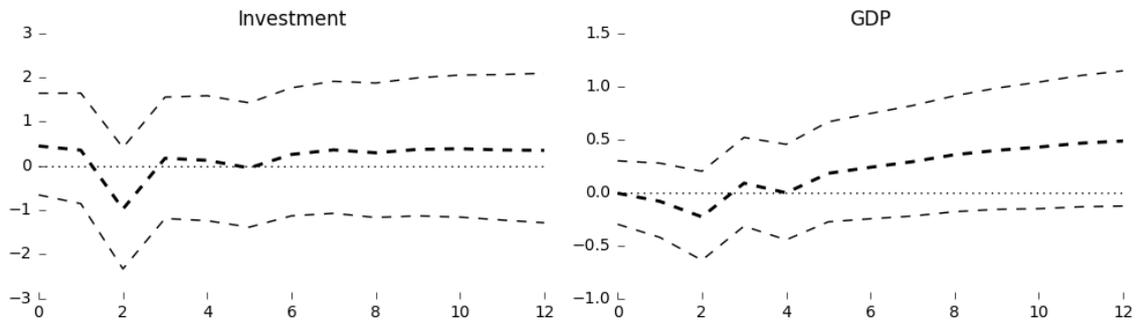
(c) Component 3 − "Norway's relationship with the EU"
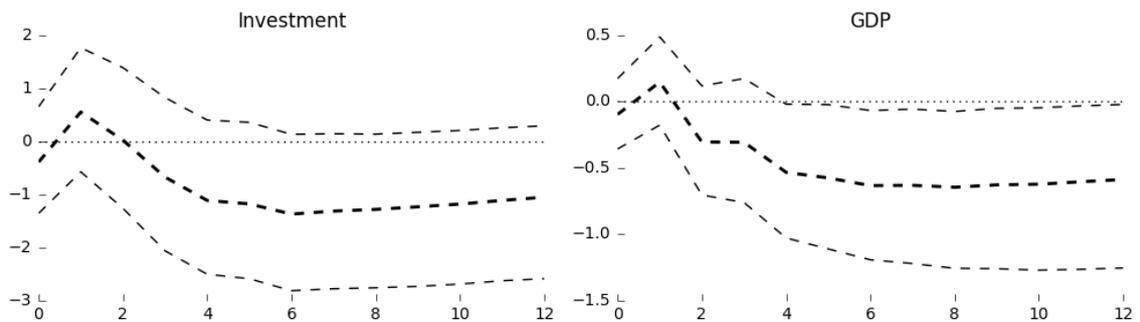
(d) Component 4 − "technology and firm expansion"
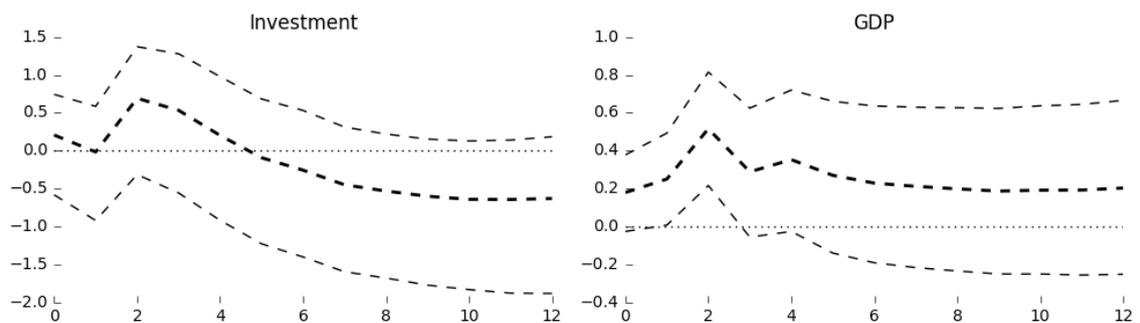
Note: The uncertainty components are introduced in the model one at a time. The 90 percent confidence bands are plotted.

Table 4. Contribution to the variance

	Component 1				Component 2			
<u>Horizon</u>	0-year	1-year	2-year	5-year	0-year	1-year	2-year	5-year
Investment	1.6%	8.5%	9.0%	7.4%	0.4%	0.9%	0.7%	0.5%
GDP	0.0%	5.0%	4.8%	2.9%	0.0%	0.5%	1.3%	7.6%
	Component 3				Component 4			
<u>Horizon</u>	0 year	1-year	2-year	5-year	0-year	1-year	2-year	5-year
Investment	0.4%	1.9%	6.5%	10.1%	0.0%	0.7%	1.5%	5.2%
GDP	0.4%	4.5%	11.6%	17.8%	1.9%	8.9%	8.5%	9.1%

Note: The contribution to the variance of investment and GDP from the four components. The contribution to variance is computed as the forecast error variance decomposition.

Moving to a a Component 3 shock, labeled as uncertainty related to "Norway's relationship with the EU", this gives a significant drop in GDP, reaching a bottom of more than -0.5 percent. It takes some time before the full effect materializes, and we see the peak response after 1.5 years. The responses after a Component 3 shock are very persistent and seems to have a considerable negative effect on the economy.

Fourth, a Component 4 shock, labeled as uncertainty related to "technology and firm expansion", gives a significant increase in GDP of around 0.5 percent, peaking after two quarters and the effect lasts for several years. This result may be positive evidence for "growth options" theories: If increased uncertainty represents a higher potential upside in the economy, we have a good type of uncertainty that leads to increased activity. There is no significant effect on investment after a Component 4 shock.

The contribution to the variance of investment and GDP from the four components is reported in Table 4. Uncertainty related to "Norway's relationship with the EU" (Component 3) explains the most of investment and GDP evaluated at the 5-years horizon. Also Component 1 and 4 has a considerable contribution.[23]

To sum up a component 1 uncertainty, related to "economic and financial distress", gives a significant fall in investment. Given the negative response, Component 1 uncertainty can be categorized as a bad type of uncertainty. This type of uncertainty has a similar profile as other uncertainty measures used in the literature, and it gives similar responses as the alternative measures when included in a structural VAR. However, higher uncertainty can also yield a positive economic response. An uncertainty shock related to "technology and firm expansion" gives a boom in GDP. Given the positive response, Component 4 uncertainty can be categorized as a good type of uncertainty.

Figure 14 in Appendix B shows that the baseline VAR presented in this section is robust to a variety of alternative specifications, with a few exceptions: First, as alluded

---

[23]Note that the uncertainty measures are included in separate models. Including all four measures in the same system yields similar results.

to earlier, the monetary policy regime changed several times during the sample studied here, and we observe excessive uncertainty related to monetary policy in the second part of the 1990s. In 2001, Norway adopted inflation targeting and this regime has prevailed for the rest of the sample. Estimating the model on the pre-2001 sample yields impulse responses that mostly strengthen the results relative to the baseline specification with the exception that we now see an increase in GDP after a Component 2 shock. The high negative correlation between Component 2 and *Macroeconomics* and *Employment* uncertainty (see Table 3) may be an explanation for the positive effect on GDP in the post-2001 sample. Second, I estimate the model using an alternative ordering, where the uncertainty measure is placed below asset prices in the VAR. This implies that uncertainty cannot react to asset prices within the same quarter. This alternative ordering does not alter the responses after shocks to the last three components. However, this specification gives no contraction after a Component 1 shock. I take this finding as an indication of a potential endogeneity problem between this first component and the rest of the system.

6   Conclusion

This paper introduces a text-based approach to create category-specific measures of uncertainty, taking advantage of text classification tools from the machine learning literature. I classify more than 28 years of newspaper articles from Norway's largest business newspaper. The articles are classified according to their underlying meaning using a topic model. I measure the degree of uncertainty conveyed by the different articles by counting the uncertainty terms within the articles. I obtain uncertainty measures related to a wide range of categories, often related to the economy, such as *Oil price*, *Monetary policy*, *Politics*, and *Stock market*.

The uncertainty measures capture well-known episodes of heightened uncertainty both at an aggregate level and at a more category-specific level. To be able to capture distinct types of uncertainty, I do an orthogonalization of the topic-based uncertainty measures. I identify four distinct types of uncertainty related to "economic and financial distress", "the institutional framework of monetary policy", "Norway's relationship with the EU", and "technology and firm expansion".

Using a structural VAR model to investigate the effect of the four different uncertainty components on investment and GDP for the Norwegian economy, I find that shocks to the different components have different effects. A shock to uncertainty related to "economic and financial distress" is followed by a contraction in the Norwegian economy; this type of uncertainty resembles the bad type of uncertainty the empirical literature has focused on. In contrast to this bad type of uncertainty, a shock to uncertainty related to "technology and firm expansion" leads to increased GDP, indicating that we also can have a good type of uncertainty.

References

Alexopoulos, M. and J. Cohen (2009). Uncertain times, uncertain measures. *University of Toronto Department of Economics Working Paper 352*.

Arellano, C., Y. Bai, and P. Kehoe (2010). Financial markets and fluctuations in uncertainty. *Federal Reserve Bank of Minneapolis Working Paper*.

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics 131*(4), 1593.

Basu, S. and B. Bundick (2012). Uncertainty shocks in a model of effective demand. Technical report, National Bureau of Economic Research.

Bernanke, B. S. (1983, February). Irreversibility, Uncertainty, and Cyclical Investment. *The Quarterly Journal of Economics 98*(1), 85–106.

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.

Blei, D. M., A. Y. Ng, and M. I. Jordan (2003). Latent dirichlet allocation. *Journal of machine Learning research 3*(Jan), 993–1022.

Bloom, N. (2009, May). The Impact of Uncertainty Shocks. *Econometrica 77*(3), 623–685.

Chang, J., S. Gerrish, C. Wang, J. L. Boyd-graber, and D. M. Blei (2009). Reading tea leaves: How humans interpret topic models. In Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, and A. Culotta (Eds.), *Advances in Neural Information Processing Systems 22*, pp. 288–296. Curran Associates, Inc.

Fernandez-Villaverde, J., P. Guerron-Quintana, K. Kuester, and J. Rubio-Ramirez (2015, November). Fiscal volatility shocks and economic activity. *American Economic Review 105*(11), 3352–84.

Fernandez-Villaverde, J., P. Guerron-Quintana, J. F. Rubio-Ramirez, and M. Uribe (2011, October). Risk matters: The real effects of volatility shocks. *American Economic Review 101*(6), 2530–61.

Gentzkow, M., B. T. Kelly, and M. Taddy (2017, March). Text as data. Working Paper 23276, National Bureau of Economic Research.

Gilchrist, S., J. W. Sim, and E. Zakrajšek (2014). Uncertainty, financial frictions, and investment dynamics. Technical report, National Bureau of Economic Research.

Gjedrem, S. (2008). Monetary policy from a historical perspective. Address at the conference to mark the 100th anniversary of the Association of Norwegian Economists in Oslo, Norway, 16 September 2008.

Griffiths, T. L. and M. Steyvers (2004). Finding scientific topics. *Proceedings of the National academy of Sciences of the United States of America 101* (Suppl 1), 5228–5235.

Gudmundsson, J. and G. J. Natvik (2012). That uncertain feeling - how consumption responds to economic uncertainty in norway. *Staff Memo, Norges Bank*.

Hamilton, J. D. (2013). Historical oil shocks. In R. E. Parker and R. M. Whaples (Eds.), *Routledge Handbook of Major Events in Economic History*, pp. 239–265. New York: Routledge Taylor and Francis Group.

Hansen, S., M. McMahon, and A. Prat (2014, June). Transparency and Deliberation within the FOMC: A Computational Linguistics Approach. CEP Discussion Papers dp1276, Centre for Economic Performance, LSE.

Jurado, K., S. C. Ludvigson, and S. Ng (2015). Measuring uncertainty. *American Economic Review 105* (3), 1177–1216.

Kraft, H., E. S. Schwartz, and F. Weiss (2013, February). Growth options and firm valuation. Working Paper 18836, National Bureau of Economic Research.

Larsen, V. H. and L. A. Thorsrud (2015, June). The Value of News. Working Papers 0034, Centre for Applied Macro- and Petroleum economics (CAMP), BI Norwegian Business School.

Leduc, S. and Z. Liu (2015). Uncertainty shocks are aggregate demand shocks. Federal Reserve Bank of San Francisco Working Paper 2012-10.

Ludvigson, S. C., S. Ma, and S. Ng (2015). Uncertainty and business cycles: Exogenous impulse or endogenous response? Technical report, National Bureau of Economic Research.

Manela, A. and A. Moreira (2016). News implied volatility and disaster concerns. *Journal of Financial Economics*.

McDonald, R. and D. Siegel (1986). The value of waiting to invest. *The Quarterly Journal of Economics 101* (4), 707–727.

Segal, G., I. Shaliastovich, and A. Yaron (2015). Good and bad uncertainty: Macroeconomic and financial market implications. *Journal of Financial Economics 117* (2), 369–397.

Appendices

Appendix A   Latent Dirichlet Allocation Model

The LDA model was developed in Blei et al. (2003). I follow the model setup presented in Griffiths and Steyvers (2004). Let T be the number of topics. The probability of word i occurring in a given document is written as

$$P(w_i) = \sum_{j=1}^{T} P(w_i|z_i = j)P(z_i = j), \tag{6}$$

where w_i is word i, z_i is a latent variable denoting which topic word i was drawn from. The term $P(w_i|z_i = j)$ denotes the probability that word i is drawn from topic j. The last term $P(z_i = j)$ gives the probability that we draw a word from topic j in the current document. Different documents will have different probabilities for drawing words from the various topics.

Let D be the number of documents in our corpus and W is the number of unique words. Then we can represent the importance of the words for the different topics as

$$P(w_i|z = j) = \phi_w^{(j)}, \text{ for all } j \in [1, T] \text{ and } w_i \in \{w_1, w_2, \ldots, w_W\} \tag{7}$$

where ϕ is a set of T multinomial distributions over the W words. The importance of a topic within a given document can be represented as

$$P(z = j) = \theta_j^{(d)}, \text{ for all } j \in [1, T] \text{ and } d_i \in \{d_1, d_2, \ldots, d_D\} \tag{8}$$

where θ is a set of D multinomial distributions over the T topics.

With the new notation we can imagine that we have the generating algorithm for the documents, this is a two step algorithm

1. Pick a distribution over topics by randomly choosing θ from a Dirichlet distribution, which then determines $P(z)$

2. For each word in the document

   (a) Pick a topic j from θ

   (b) Given that you have topic j, pick a word from $\phi^{(j)}$, which is assumed to be fixed.

Then if we know the algorithm the documents was generated from, and we have the final documents, it is possible to estimate the distribution ϕ and θ. We use Gibbs sampling to estimate the distributions.

Estimating the LDA Model

Griffiths and Steyvers (2004) gives a complete specification LDA-model with the additional Dirichlet prior on ϕ as

$$w_i | z_i, \phi^{(z_i)} \sim \text{Discrete}(\phi^{(z_i)}) \tag{9a}$$

$$\phi \sim \text{Dirichlet}(\beta) \tag{9b}$$

$$z_i | \theta^{(d_i)} \sim \text{Discrete}(\theta^{(d_i)}) \tag{9c}$$

$$\theta \sim \text{Dirichlet}(\alpha) \tag{9d}$$

where α and β are hyperparameters specifying the prior distribution for ϕ and θ. When estimating a topic model we need to set three parameters; That is the number of topics, T, and the two hyperparameters of the Dirichlet priors, α and β. We follow Griffiths and Steyvers (2004) when selecting priors

$$\alpha = \frac{50}{T}, \quad \text{and} \quad \beta = \frac{200}{W}.$$

We select 80 topics and we can start sampling topics, we use a burn in period before we start keeping samples. When keeping samples we use a thinning interval of 50. Then we keep as many samples as is computational feasible. We have a very large text corpus and sampling topics from this corpus demands allot of memory. The LDA model is estimated on a server with 50gb of memory.

Model selection

Given values for the hyperparameters α and β, the only variable we can change is the number of topics, T. Choosing the right value for T is a model selection problem. A measure we use to evaluate the performance of our topic model is perplexity (equivalent to predictive likelihood) defined as follows

$$\text{Perplexity}(\boldsymbol{w}) = \exp\left\{ -\frac{\mathcal{L}(\boldsymbol{w})}{W} \right\}, \tag{10}$$

where

$$\mathcal{L}(\boldsymbol{w}) = \log P(\boldsymbol{w}|\boldsymbol{z}). \tag{11}$$

Appendix B   Additional results

Table 5. Estimated topics and labeling

Topic	Label	First words
Topic 0	Calendar	january, march, october, september, november, february
Topic 1	Family business	family, foundation, name, dad, son, fortune, brothers
Topic 2	Institutional investing	fund, investments, investor, return, risk, capital
Topic 3	Justice	lawyer, judge, appeal, damages, claim, supreme court
Topic 4	Surroundings	city, water, meter, man, mountain, old, outside, nature
Topic 5	Housing	housing, property, properties, apartment, square meter
Topic 6	Movies/Theater	movie, cinema, series, game, producer, prize, audience
Topic 7	Argumentation	word, besides, interesting, i.e., in fact, sure, otherwise
Topic 8	Unknown	road, top, easy, hard, lift, faith, outside, struggle,fast
Topic 9	Agriculture	industry, support, farmers, export, production, agriculture
Topic 10	Automobiles	car, model, engine, drive, volvo, ford, møller, toyota
Topic 11	USA	new york, dollar, wall street, president, usa, obama, bush
Topic 12	Banking	dnb nor, savings bank, loss, brokerage firm, kreditkassen
Topic 13	Leadership	position, chairman, ceo, president, elected, board member
Topic 14	Negotiation	solution, negotiation, agreement, alternative, part, process
Topic 15	Newspapers	newspaper, media, schibsted, dagbladet, journalist, vg
Topic 16	Health care	hospital, doctor, health, patient, treatment, medication
Topic 17	IT systems	it, system, data, defense, siem, contract, tandberg, deliver
Topic 18	Stock market	stock exchange, fell, increased, quote, stock market
Topic 19	Macroeconomics	economy, budget, low, unemployment, high, increase
Topic 20	Oil production	statoil, oil, field, gas, oil company, hydro, shelf, stavanger
Topic 21	Wage payments	income, circa, cost, earn, yearly, cover, paid, salary
Topic 22	Regions	trondheim, llc, north, stavanger, tromsø, local, municipality
Topic 23	Family	woman, child, people, young, man, parents, home, family
Topic 24	Taxation	tax, charge, revenue, proposal, remove, wealth tax, scheme
Topic 25	EU	eu, eea, commission, european, brussel, membership, no
Topic 26	Industry	hydro, forest, factory, production, elkem, industry, produce
Topic 27	Unknown	man, he, friend, smile, clock, evening, head, never, office
Topic 28	Mergers and acquisitions	orkla, storebrand, merger, bid, shareholder, acquisitions
Topic 29	UK	british, london, great britain, the, of, pound, england
Topic 30	Narrative	took, did, later, never, gave, stand, happened, him, began
Topic 31	Shipping	ship, shipping, dollar, shipowner, wilhelmsen, fleet, proud
Topic 32	Projects	project, nsb, development, fornebu, entrepreneurship
Topic 33	Oil price	dollar, oil price, barrel, oil, demand, level, opec, high
Topic 34	Sports	olympics, club, football, match, play, lillehammer, sponsor
Topic 35	Organizations	leader, create, organization, challenge, contribute, expertise
Topic 36	Drinks	wine, italy, taste, drinks, italian, fresh, fruit, beer, bottle
Topic 37	Nordic countries	swedish, sweden, danish, denmark, nordic, stockholm
Topic 38	Airline industry	sas, fly, airline, norwegian, braathens, airport, travel

Table 5 – continued from previous page

| Topic	Label	First words
Topic 39	Entitlements	municipality, public, private, sector, pension, scheme
Topic 40	Employment	cut, workplace, measures, salary, labor, working, employ
Topic 41	Politics	conservatives, party, ap, labor party, stoltenberg, frp
Topic 42	Funding	loan, competition, creditor, loss, bankruptcy, leverage
Topic 43	Literature	book, books, read, publisher, read, author, novel, wrote
Topic 44	Statistics	count, increase, investigate, share, average, decrease
Topic 45	Watercraft	ship, boat, harbor, strait, shipowner, on board, color
Topic 46	Results	quarter, surplus, deficit, tax, group, operating profit, third
Topic 47	TV	tv, nrk, channel, radio, digital, program, media
Topic 48	International conflicts	war, africa, irak, south, un, army, conflict, troops, attack
Topic 49	Elections	election, party, power, politics, vote, politician, support
Topic 50	Music	the, music, record, of, in, artist, and, play, cd, band, song
Topic 51	Oil service	rig, dollar, contract, option, offshore, drilling, seadrill
Topic 52	Tourism	hotel, room, travel, visit, stordalen, tourist, guest Â
Topic 53	Unknown	no, thing, think, good, always, pretty, actually, never
Topic 54	Engineering	aker, kværner, røkke, contract, shipyard, maritime
Topic 55	Fishery	fish, salmon, seafood, norway, tons, nourishment, marine
Topic 56	Europe	german, russia, germany, russian, west, east, french, france
Topic 57	Law and order	police, finance guards, aiming, illegal, investigation
Topic 58	Weekdays	week, financial, previous, friday, wednesday, tdn, monday
Topic 59	Supervision	report, information, financial supervision, enlightenment
Topic 60	Retail	shop, brand, steen, rema, reitan, as, group, ica, coop
Topic 61	Startups	bet, cooperation, establish, product, party, group
Topic 62	Food	food, restaurant, salt, nok, pepper, eat, table, waiter
Topic 63	Stock listings	shareholder, issue, investor, holding, stock exchange listing
Topic 64	Asia	china, asia, chinese, india, hong kong, south, authorities
Topic 65	Art	picture, art, exhibition, gallery, artist, museum, munch
Topic 66	Disagreement	criticism, express, asserting, fault, react, should, alleging
Topic 67	Debate	degree, debate, context, unequal, actually, analysis
Topic 68	Life	man, history, dead, him, one, live, church, words, strokes
Topic 69	Goods and services	customer, post, product, offers, service, industry, firm
Topic 70	Telecommunication	telenor, mobile, netcom, hermansen, telia, nokia, ericsson
Topic 71	IT technology	internet, net, pc, microsoft, technology, services, apple
Topic 72	Monetary policy	interest rate, central bank, euro, german, inflation, point
Topic 73	Education	school, university, student, research, professor, education
Topic 74	Regulations	rules, authorities, competition, regulations, bans
Topic 75	Trade organizations	lo, nho, members, forbund, strike, organization, payroll
Topic 76	Fear	fear, emergency, hit, severe, financial crisis, scared
Topic 77	Fiscal policy	suggestions, parliamentary, ministry, selection, minister
Topic 78	Energy	energy, emissions, statkraft, industry, environment
Topic 79	Foreign	foreign, abroad, japan, japanese, immigration, games

Note: The topics are labeled based on the meaning of the most important words, see the text for details. The "# of articles" column reports the number of articles, in the full sample which, according to the model, belong to that specific topic. The words are translated from Norwegian to English using Google Translate.

Table 6. Historical events of heightened uncertainty

Event	Date	Description
1990s		
GW1	1990-08-02	Gulf War 1 (2 Aug. 1990 – 28 Feb 1991)
Fixed rate (ECU)	1990-10-01	MPR: Fixed exchange rate pegged to ECU
Bank bailout	1991-12-20	The nationalization of several Norwegian banks
Free float	1992-12-01	MPR: Free float exchange rate (leaving the ECU)
Moland	1994-01-01	New central bank governor: Moland
Stability t.w. EU	1994-05-01	MPR: Stability t.w. EU currencies
EU vote	1994-11-28	Norwegian referendum for EU membership
Storvik	1996-01-01	New central bank governor: Storvik
Asian crisis	1997-11-01	The Asian financial crisis
LTCM Default	1998-09-23	Collapse of Long-Term Capital Management
Gjedrem	1999-01-01	New central bank governor: Gjedrem
2000s		
Peak of NASDAQ	2000-03-10	NASDAQ peak before the burst of the dot-com bubble
Inflation tgt.	2001-03-01	MPR: Inflation target and floating FX
9/11	2001-09-11	The Al-Qaeda attack on 9/11
War in Afghanistan	2001-10-07	War in Afghanistan (7 Oct 2001 – 28 Dec 2014)
WorldCom bankruptcy	2002-07-21	WorldCom goes bankrupt
Bottom of NASDAQ	2002-10-01	NASDAQ bottom after the burst of the dot-com bubble
GW2	2003-03-20	Gulf War 2 (20 Mar 2003 – 01 May 2003)
Credit crunch	2007-08-01	Start of the Global Financial Crisis
Lehman	2008-09-15	The collapse of Lehman Brothers
2010s		
Libyan civil war	2011-02-15	15 Feb 2011 – 23 Oct 2011
Olsen	2011-01-01	New central bank governor: Olsen
Stock market crash	2011-08-01	Stock market crash
Greek prop. refer.	2011-10-31	Greek proposed economy referendum
Egyptian coup	2013-07-03	Egyptian coup d'état
OPEC meeting	2014-11-28	OPEC chose not to reduce production
Brexit	2016-06-26	UK voted to leave the EU

Note: Details on the historical events that are indicated in the plots of the uncertainty indexes. MPR is an abbreviation for Monetary Policy Regime.

Figure 10. Correlation between the topic-based measures
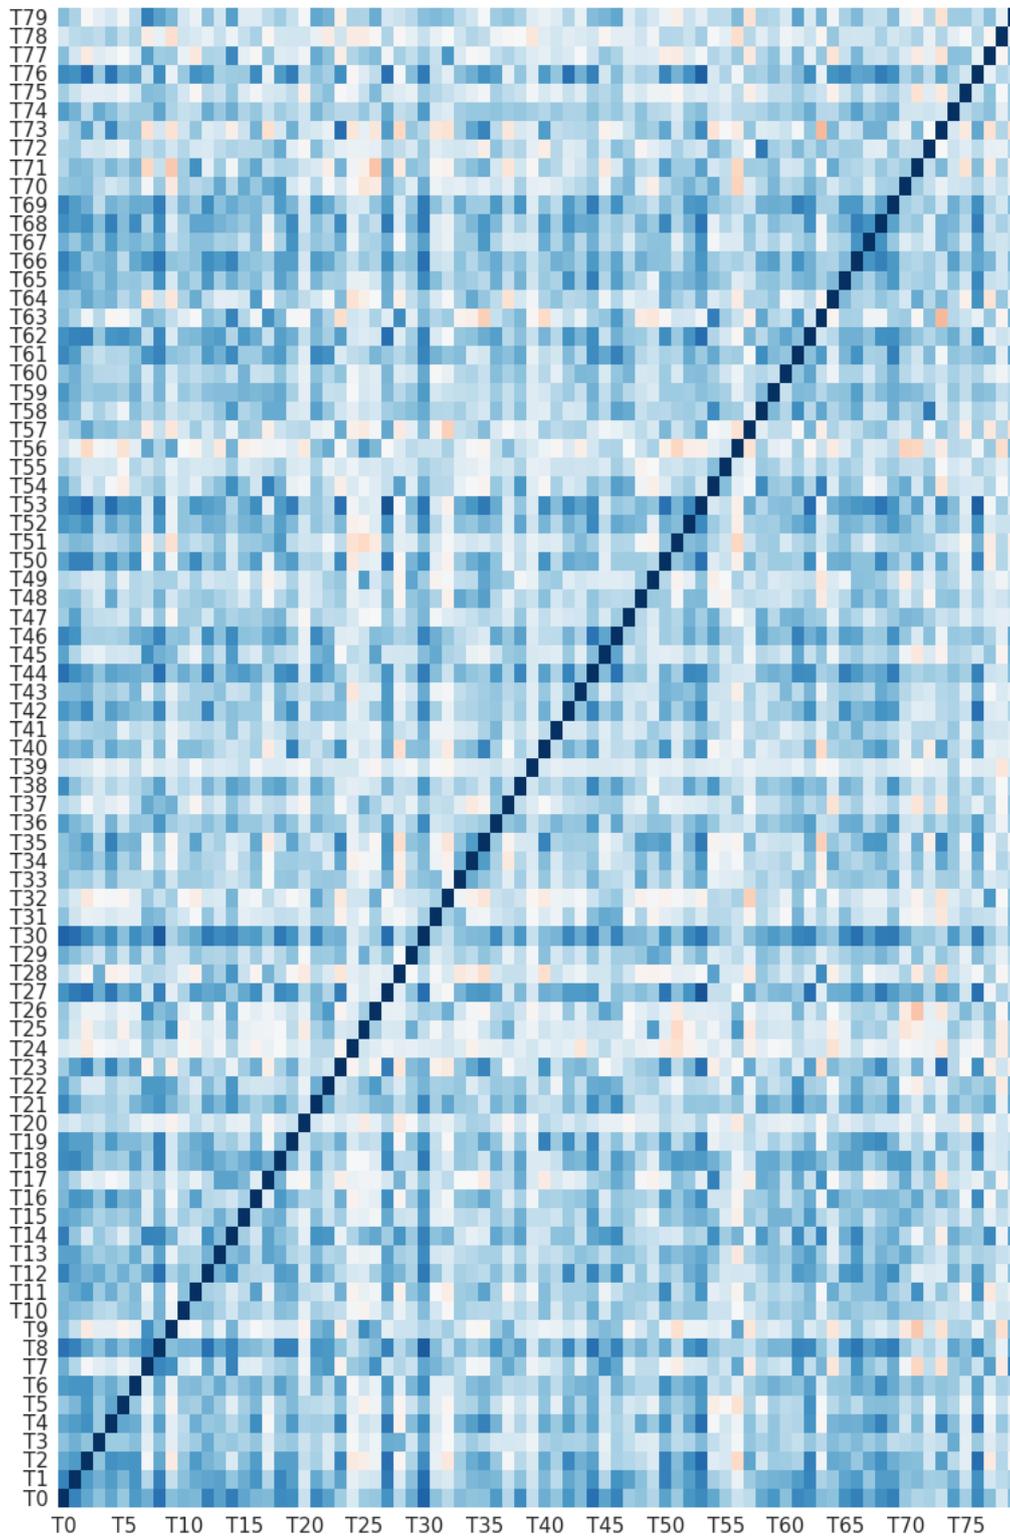
Note: The correlations are computed at a quarterly frequency. See Table 5 for the corresponding topic labels. Blue represents a positive correlation while red represents a negative one.

Figure 11. Giving content to the components
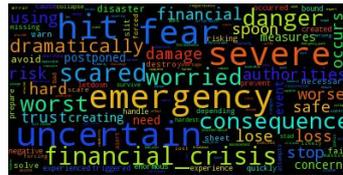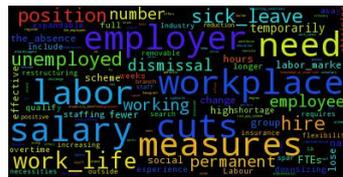
*Note:* The 150 words with the highest probabilities are shown, the size of the words corresponds to the probability of that word occurring in the topic distribution. The correlation with the components is given in the parentheses.

Figure 12. IRFs using the aggregate uncertainty measure
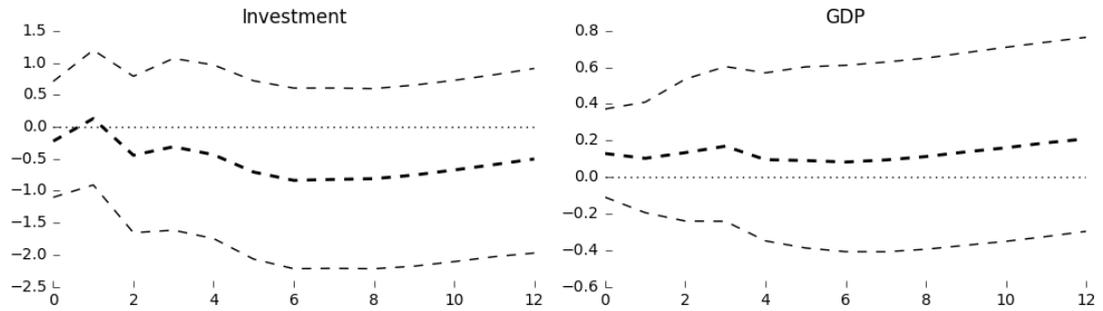
Note: The 90 percent confidence bands are plotted.

**Figure 13.** The most and least negative responses

Note: The eight most, and the eight least, negative responses in investment and GDP using the topic-based measures.

**Figure 14.** Impulse responses using alternative specifications

(a) Component 1 − "economic and financial distress"
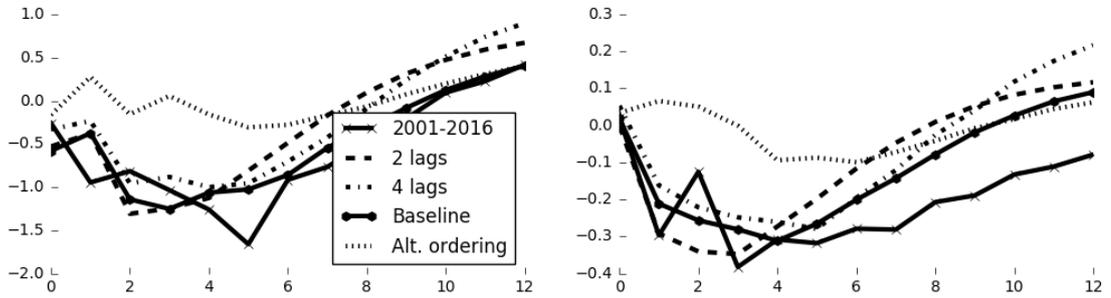
(b) Component 2 − "the institutional framework of monetary policy"
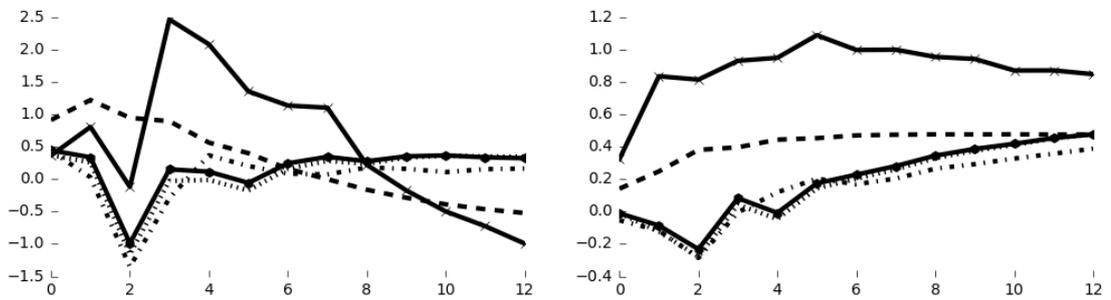
(c) Component 3 − "Norway's relationship with the EU"
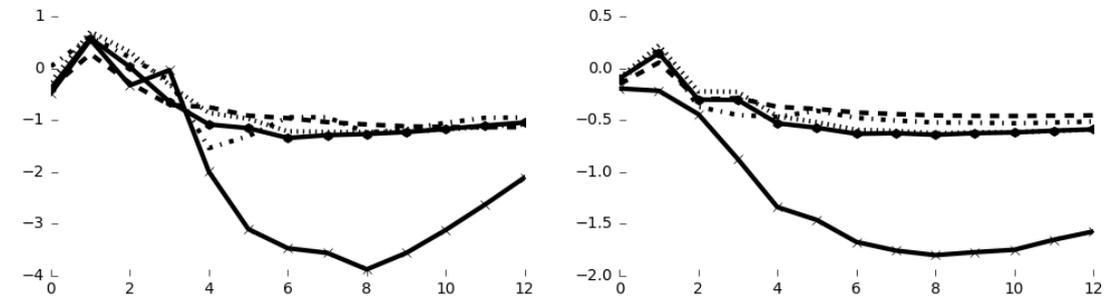
### (d) Component 4 − "technology and firm expansion"
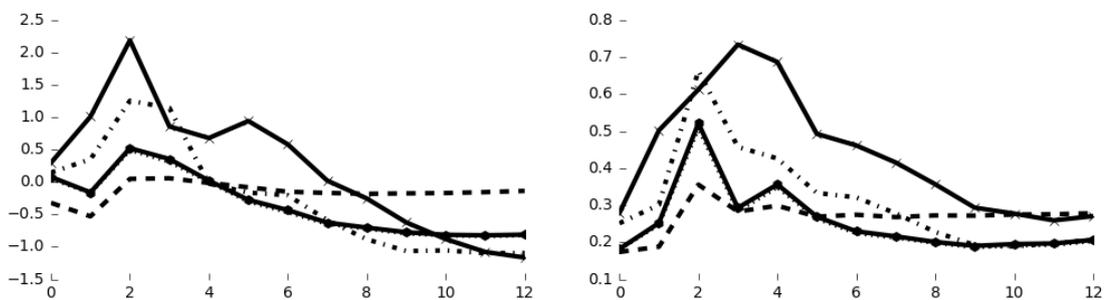
Note: Impulse responses using four alternative specifications together with the baseline specification outlined in Section 5.

Appendix C Norwegian index of economic policy uncertainty

Baker et al. (2016) create an economic policy uncertainty (EPU) index for 11 countries. These indices have proven to be popular and are available through commercial data resources such as Bloomberg, FRED, and Reuters. This index is not available for Norway, and this section will follow Baker et al. (2016) and create this index based on the Norwegian newspaper DN. The series is plotted in Figure 15.

The index is created by counting the articles that contain fords from the following three categories: *uncertainty* or *uncertain*; *economic* or *economy*; and Baker et al. (2016) use the following policy terms: *congress*, *deficit*, *Federal reserve*, *legislation*, *regulation* or *white house*. These three categories are named: uncertainty, economy, and policy. The counted articles contain words from all of them. Each day the final count is divided by the total number of articles that day to control for changes in total news coverage over time. The words need to be translated into their Norwegian counterparts to suit a Norwegian setting, the translations used are given in Table 7.

Figure 15. Economic Policy Uncertainty Index for Norway
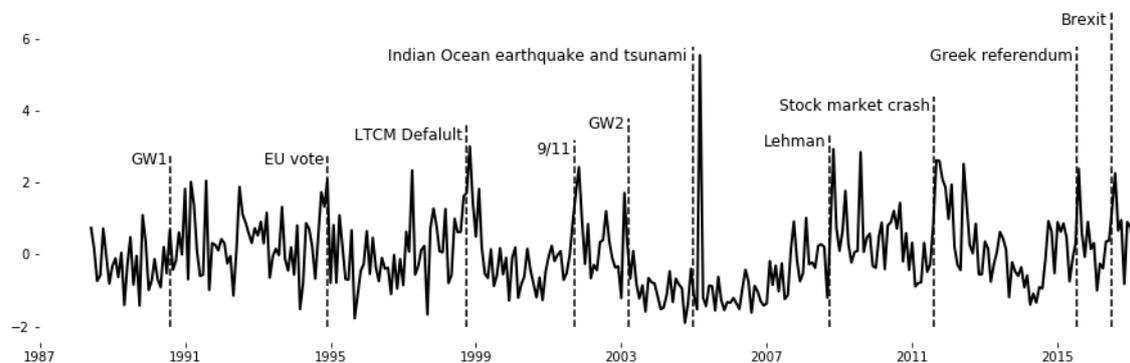
Note: The index is calculated by counting articles that contain words from the *uncertainty* terms, the *economy* terms, and the *policy* terms.

Table 7. Term Sets for the Norwegian EPU index

Category	English	Norwegian
Uncertainty	*uncertainty* or *uncertain*	*usikker* or *usikkerhet*
Economy	*economic* or *economy*	*økonomisk* or *økonomi*
Policy	*government*	*regjering*
	parliament	*storting*
	authorities	*myndigheter*
	tax	*skatt*
	regulation	*regulering*
	budget	*budsjett*
	deficit	*underskudd*
	ministry of finance	*Finansdepartementet*
	central bank	*sentralbank*

Note: I also include variations of the words given in this table such as *taxation* and *regulations*.