STOCKHOLM SCHOOL OF ECONOMICS
Department of Economics
5350 Master's Thesis in Economics
Academic year 2020-2021

News-Based Uncertainty and Its Impact on Animal Spirits

How News-based Expectation Uncertainty can give Additional Economic Insights to News-based Economic Policy Uncertainty

Tina Emambakhsh (41561)

Abstract

An ever-unprecedented pandemic has recently skyrocketed uncertainty levels worldwide. This has made the need for measuring uncertainty and understanding its economic role more important than before. Recently, *economic policy uncertainty* expressed in news has been used as a proxy for uncertainty. In this thesis, I propose an alternative approach by using Swedish articles about economic expectations to measure uncertainty and show that they provide additional insights into the economic impact of news-based uncertainty. I construct three uncertainty indices for the Swedish economy. The first is the Economic Forecast Uncertainty (EFU) index and is based on the same method as the Economic Policy Uncertainty (EPU) index. The two other indices are based on latent news *topics* relevant to expectation uncertainty, namely macroeconomic and policy expectations, which are identified with an unsupervised machine learning technique that combines Latent Dirichlet Allocation with Principal Component Analysis. It is demonstrated that news-based expectation uncertainty about the (macro)economy has a long term negative impact on Swedish firm and household confidence. Economic policy and policy expectation uncertainty, in contrast, do not influence Swedish confidence. The empirical findings of this thesis suggest that shocks in news-based economic expectation uncertainty can affect economic fundamentals through falling confidence of economic agents, while news-based economic policy uncertainty passes through a different channel.

JEL: D80, E32, E66, G18, L50
Keywords: Uncertainty, Expectations, Newspaper, Latent Dirichlet Allocation, Economic Policy

Supervisor: Rickard Sandberg
Date Submitted: December 7, 2020
Date Examined: December 17, 2020
Discussant: Valentina Farinelli, Nicolas Powell
Examiner: Matilda Kilström

Acknowledgements

I am thankful for the advice and mentorship received from my supervisor Rickard Sandberg throughout my process. Equal gratitude goes to Sebastian Ankargren for his thought-provoking impulses, our discussions, and for sharing his expertise.

My highest gratitude is also owed to Niklas Thelning and Tatiana Polushina from Retriever who made it possible to realize my idea, and to their colleagues Daniel Svensson and Sigmund Ellingsrud for the engaging discussions. I also thank Joel Smedberg for the technical support and Lina Thomas for her valuable feedback.

Thank you Jean-Philippe Cointet for introducing me into the fascinating world of computational social sciences. Finally, I want to thank my parents, and Albert Planting-Gyllenbåga, Adam Gill, and Mariella Orasch for their dedicated support in my journey.

Contents

1 Introduction **1**

2 Background **3**

2.1 Defining Uncertainty 3

2.2 Evaluating Existing Measures of Uncertainty 5

2.3 Why Uncertainty Shocks Matter 7

3 Foundations of NLP for Macroeconomic Analysis **9**

4 Measurement of News-Based Expectation Uncertainty **11**

4.1 Data Sourcing & Sampling Strategy 11

4.2 Dictionary Method 12

4.3 Semantic Method 14

5 Evaluation of News-Based Expectation Uncertainty **24**

5.1 Lead-versus-Lag Relationship with Swedish Confidence 24

5.2 Shock Response of Swedish Confidence 29

6 Discussion **33**

6.1 Validity and Limitations of Results 35

6.2 Perspectives for Future Research 37

7 Conclusion **39**

A Appendix: Data **48**

B Appendix: Measurement **50**

C Appendix: Additional Results **61**

List of Tables

1 Top Principal Components - Narrative Description. 21

2 Granger Causality from News-Based Uncertainty to Confidence 29

3 Granger Causality from Confidence to News-Based Uncertainty 29

A1 Sources of Macroeconomic Variables Used in the Evaluation. 49

A2 Questions included in the Swedish Confidence Indicators 49

B1 Definition of Key NLP-Terminology. 50

B2 Major uncertainty-related Shocks as labelled in Figure B1. 51

B3 Latent Topics in News about Economic Expectations. 57

C1 Propoerties of the VARs and Breusch-Godfrey Test Results for Serially Correlated Errors. 64

List of Figures

1 The Difference in Impact of Confidence and Uncertainty Shocks on Expectations. 4

2 Number of Articles per Month. 12

3 EFU versus EF and Uncertainty-related Shocks.. 14

4 Identification Process of the Semantic Method. 15

5 News Topics of the EF Article Corpus in Form of Probability Distributions over Words. 17

6 Examples of the Uncertainty-Adjusted Prevalence of Topics over Time. 19

7 Example of Article Selection Procedure for the Computation of EFU-MACRO. 23

8 Final Topic-Based Uncertainty Indices. 23

9 Does News-based Uncertainty lead or lag Confidence? - Visual Inspection. 25

10 Does News-based Uncertainty lead or lag Confidence? - Cross-Correlation Functions. 28

11 VAR Estimation of the Impact of a News-based Uncertainty Shock on Confidence. 31

12 Robust VAR Estimation of the Impact of a News-based Uncertainty Shock on Confidence of each shock response. 32

A1 Bag-of-Words Representation of Articles. 48

A2 Replication of the Swedish EPU-Index as outlined in §4. 48

B1 EFU Index and Uncertainty-related Shocks. 51

B2 Probabilistic Graphical Model of LDA. 53

B3 Prior α and the Symmetric Dirichlet Distribution. 55

B4 Prior α and the Asymmetric Dirichlet Distribution. 55

B5 Optimal K Approximation. 56

B6 Timely Variation Captured by Top Principal Components. 60

C1 VAR Estimation of the Impact of a EFU versus EPU Shock on the Swedish Aggregate Economy. . 62

C2 VAR Estimation of the Impact of a EFU-MACRO and EFU-POL Shock on the Swedish Aggregate Economy. 63

C3 VAR Estimation of the Impact of a News-based Uncertainty Shock on Confidence - Confidence ordered first. 65

C4 VAR Estimation of the Impact of a Confidence Shock on News-based Uncertainty. 65

1 Introduction

Policy makers and macroeconomists alike are pondering on the high levels of uncertainty that have been prevailing to dampen growth in industrial economies since the Great Recession starting in 2008. In the last couple of years, steady employment growth and low interest rates have only been partly successful in eliminating uncertainty, and recently, an ever-unprecedented pandemic has furthermore skyrocketed uncertainty levels worldwide. Empirical literature has persistently found that the economic impact of uncertainty is non-trivial, causing households to pile up precautionary savings and firms to defer investment decisions (Bloom 2009). What is more, uncertainty causes financial markets to become less responsive to short-term interest rates so that policy interventions become less effective (Alessandri & Bottero 2017). These findings combined with the current circumstances of high perceived uncertainty make it especially important to properly measure and understand the economic influence of uncertainty.

One key factor that drives uncertainty are *expectations* about how the economy will develop. Nobel laureate Robert J. Shiller argued in his seminal work about *narrative economics* that uncertainty roots from the shared beliefs about the future economy, which are driven by narratives and social interactions (Shiller 2017). These narratives move personal expectations away from fundamentals and become self-fulfilling, thus causing business cycles to fluctuate. In this regard, news are one of the major channels that can propagate these self-fulfilling narratives. Some critics even claim that the media can *'talk the economy into recessions'* (The Economist 2001).

The purpose of this thesis is to explore in greater detail the aspect of expectations in conveying uncertainty through news. This is done by developing three news-based uncertainty indices on the basis of articles about economic expectations, which are subsequently used for economic analysis. The main motivation behind this approach is to augment the currently most salient method of measuring uncertainty with news that was introduced by Baker et al. (2016). They created an Economic Policy Uncertainty (EPU) index for the US by counting the relative frequency of articles mentioning keywords related to the economy, policy and uncertainty. However, the EPU focuses on uncertainty related to policy aspects only. Moreover, it uses manually handcrafted keywords for the identification of relevant articles which makes it prone to measurement error.

In the first step, I propose a new approach to develop three novel news-based uncertainty indices that modify the EPU index in two aspects. The first modification shifts the underlying subject of interest from policy uncertainty to *expectation uncertainty*. Therefore, I create an Economic Forecast Uncertainty (EFU) index with the methodology proposed by Baker et al. (2016). The second modification undertaken is by identifying *narratives* related to expectation uncertainty to develop meaningful categorical uncertainty indices. Rather than relying on handcrafted keywords or the manual labelling of articles, I make use of Latent Dirichlet Allocation (LDA), an unsupervised machine learning technique from the field of natural language processing (NLP), that can infer latent topics from articles about economic expectations. Subsequently, I identify those topics which are the most relevant for conveying expectation uncertainty over time by making use of Principal Component Analysis. Based on the principal components, I refine the filtering of the EFU articles and obtain two additional indices based on the categories of *Macroeconomic* (EFU-MACRO) and *Policy* (EFU-POL) uncertainty. The focus lies on measuring uncertainty in the Swedish economy, thereby using Swedish newspaper articles for the analysis.

The second step comprises the evaluation of the newly obtained uncertainty indices. I validate their time-series against notable uncertainty-related events in the Swedish economy. Furthermore, I evaluate how much they improved the Swedish EPU. *Improvement* is measured through two aspects. First, I analyze whether the novel indices impact Swedish economic fundamentals differently than the EPU. Second, I investigate whether news-based uncertainty can indeed influence the beliefs of economic agents. Therefore, I conduct a systematic analysis of the relation between news-based uncertainty and beliefs, the latter being proxied by firm and household confidence indicators. I compare whether news-based expectation uncertainty impacts Swedish confidence more strongly than the EPU and support my analysis by conducting Granger causality tests and assessing the dynamic relationship between news-based uncertainty and confidence in a vector autoregression (VAR) model.

My findings show that both policy and expectation uncertainty show a similar impact on economic fundamentals. However, the difference in impact on Swedish firm and household confidence is highly substantial. Rising economic and macroeconomic expectation uncertainty in the news, measured by the EFU and EFU-MACRO, are a leading indicator for falling Swedish confidence, while policy-related news, measured by the EPU and EFU-POL, mostly rise only after confidence falls. Moreover, only shocks in news-based economic and macroeconomic expectation uncertainty cause Swedish confidence to fall, from which they only recover in the long-term. Overall, these empirical findings show that news coverage about expectation uncertainty related to the (macro)economy can affect economic fundamentals through changing the beliefs of consumers and firms. Uncertainty related to policy or policy expectations passes through a different channel. An added value of the new indices is therefore their predictive potential for Swedish confidence, which itself is a predictor for fluctuations in aggregate economic activity.

With my thesis, I make several contributions to the economic literature. First, I am, to the best of my knowledge, the first to focus on economic expectations for measuring uncertainty in news. Second, my thesis systematically analyses the role of news-based uncertainty on firm and household confidence for the first time, thereby contributing empirical evidence to Shiller (2017)'s theory that narratives (in the news) change the beliefs and ultimately the behaviour of economic agents. Finally, I provide categorical uncertainty indices based on Swedish newspaper articles, which can be incorporated in future Swedish macroeconomic models.

The rest of the thesis is organized as follows. §2 introduces the topic by defining uncertainty, evaluating existing uncertainty measures, and presenting stylized facts about the impact of uncertainty shocks. §3 provides a conceptual understanding of the use of natural language processing for macroeconomic analysis. This shall facilitate the understanding of §4, where the methodology for measuring uncertainty in news is explained. Since the measurement is an iterative process, the section simultaneously incorporates the results of the measurement by visually presenting the obtained indices. In §5, the new uncertainty indices are evaluated by comparing their impact on Swedish firm and household confidence with the impact of the existing Swedish EPU index. Finally, §6 summarizes and discusses the validity and limitations of the findings, and §7 concludes.

2 Background

Before delving into the measurement of news-based uncertainty, this section gives relevant background information about measuring and analyzing uncertainty from an economic perspective. The first subsection explains how uncertainty is defined in economic terms and determines its definition for the application of this thesis. The second subsection evaluates existing time-series measures of uncertainty that have been used in the economic literature, and outlines their drawbacks which motivated the subject of and approach taken in this thesis. Finally, the last subsection explains why measuring uncertainty is important for economic research. It presents empirical evidence about the economic impact of uncertainty, which guides the analysis undertaken in §5 of the thesis, where the new uncertainty measures will be evaluated.

2.1 Defining Uncertainty

Uncertainty can have different meanings depending on the domain in which it is used. This section presents the most salient definitions of the term uncertainty in the economic literature and clarifies which concept lies in the focus of this thesis.

The first differentiation can be done between uncertainty *in measures* and uncertainty *in beliefs*. The former is used to express the quality of a measurement or the doubt that exists about the preciseness of what is measured. This concept is of rather technical nature, although one can still argue that it provides information about, for instance, an econometrician's belief about the economic object she measures, as will be seen in §2.2. In contrast, the concept of uncertainty in beliefs is concerned with people's perception about current and future economic developments, being what is steering *animal spirits* in their investment decisions as Keynes (1936) theorizes. Hence, by this definition, uncertainty can be understood as a cognitive state (Eckley 2015).

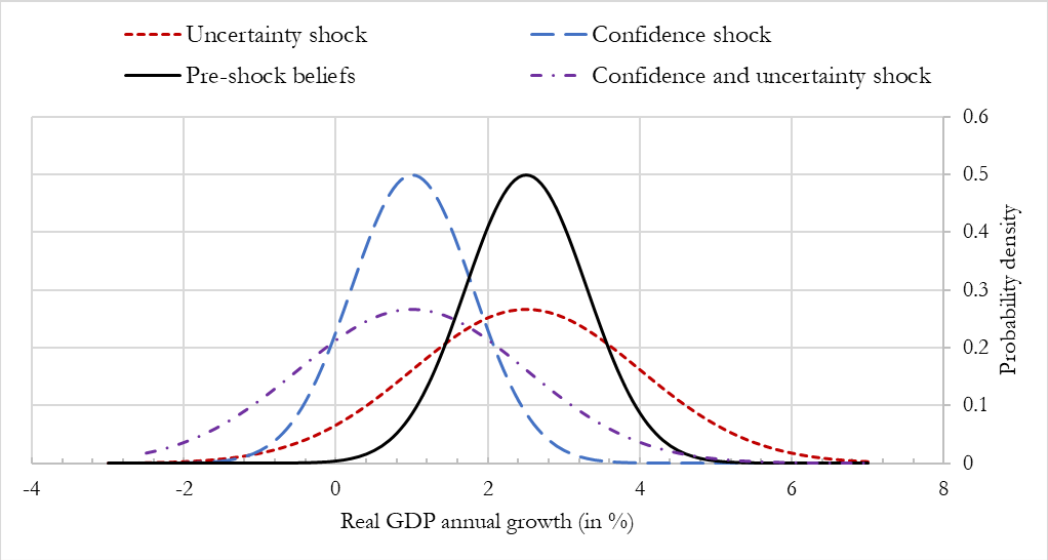
Second, one has to distinguish uncertainty from *risk*. In economic theory, the difference between these terms goes back to Knight (1921). Risk is defined as an event whose probability of occurrence can be determined by objective characteristics. For instance, the equity risk premium is calculated as the excess return over the risk-free rate that is available in the stock market. Risk can be furthermore reduced through various means, such as by grouping or through insurances. In contrast, uncertainty is an event whose possible states are unknowable or only measurable through subjective standards, such as intuition or personal judgment. Future economic outcomes are by nature uncertain both in probabilities and ranges. Hence, (rational) economic agents cannot make decisions based on pure mathematical calculations. In (Hicks 1931, pp. 170)'s words that implies: *'If uncertainty was removed, there would be no profit'*. [1]

[1] It is interesting to notice that there are terminological distinctions in the definition of uncertainty among economic schools. In neoclassical theory uncertainty is described as *epistemological* uncertainty, implying that the future *could* be determined, but due to the complexity of information available and the costs involved in identifying the relevant data, it is impossible for economic agents to derive an optimal decision. In Post-Keynesian theory, on the other hand, uncertainty is also called *fundamental* or *ontological* uncertainty, which refers to the view that the future *cannot* be determined, simply because predictable regularities between the past and the future are disrupted through the occurrence of innovations (North 2005) Therefore, it is further argued, the notion of 'error' for forecasts is in itself erroneous. In fact no 'correct' forecasts exist because economic markets are creative processes rather than statistical shadows of the past (Hausman 1994). See Bronk (2011) for a thorough analysis of the epistemological implications of the meaning of uncertainty in theory and practice.

Finally, uncertainty is not to be equalized with *pessimism* or low confidence. Bloom (2009) distinguishes between *first-moment* (levels) and *second-moment* (volatility) shocks and empirically showed that their economic impact is different. First-moment shocks typically have a persistent impact on economic activity from which the economy only slowly recovers. These can be productivity and/or demand shocks, which are closely linked to *shocks in confidence*. This is because after productivity or demand visibly changes, economic agents' expectations about the state of the economy consistently shifts towards more or less pessimistic beliefs. On the other hand, *shocks in uncertainty* have a temporary economic impact, where slowdown and bounce-back in economic activity both happens rapidly. The temporal effect of confidence shocks can be two to three years, whereas the effect of uncertainty shocks typically lasts for several months (Bloom 2009). In reality, uncertainty shocks rarely occur alone but usually coincide with confidence shocks, especially during times of crisis. Figure 1 illustrates the difference in the role of confidence versus uncertainty (shocks) in terms of future expectations. The probability density function (pdf) assigns different probabilities to a set of outcomes in GDP growth. The most likely outcome is the first-moment of the distribution and is represented by the mean, while the volatility in the outcome is the second-moment and corresponds to the width of the distribution. If we are given a current state, represented by the black line, and an economic recession occurs, a first moment shock shifts the curve to the left, i.e. decreases the average expected level of economic growth. This means that, on average, economic agents have more pessimistic expectations towards economic activity. If agents become more uncertain, on the other hand, the average expectation about future growth remains the same, but the range of possible outcomes becomes larger because agents assign greater probability to more extreme outcomes, i.e. higher and lower growth rates. Consequently, the width of the curve increases through this second moment shock, implying the *lack of* confidence in expectations about future economic developments (Haddow 2013).

For the application in this thesis, the most suitable definition of uncertainty lies in the initially defined con-

Figure 1 The Difference in Impact of Confidence and Uncertainty Shocks on Expectations.



Source: Haddow (2013)

cept of uncertainty *in beliefs*. It is aimed to approximate expressions of *uncertain cognitive states* by means of natural language. More precisely, an index will be created that measures *perceived* levels of uncertainty expressed in news. Regarding its distinction to risk it is noteworthy that language about uncertainty in the news will most likely also convey sentiments of risk. Hence, these two concepts cannot be clearly distinguished in the analysis of articles. Regarding the third differentiation, the evaluation section of this thesis furthermore investigates whether uncertainty can actually affect the first-moment of outcomes.

2.2 Evaluating Existing Measures of Uncertainty

Existing methods for measuring uncertainty can be categorized in *survey-based*, *market-based, forecast-based*, and *news-based* methods. Despite the fact that most of the literature developed uncertainty measures for autarkic economies like the US, regional aspects will be ignored in this evaluation. Moreover, it is noteworthy that different measures of uncertainty naturally influence the type of impact they capture. This impact heterogeneity of measures will be accounted for in the section thereafter.

Surveys are the most traditional method for identifying perceived uncertainty of economic outlooks, either by measuring the dispersion in expectations of firms regarding their employment, capital investment, and sales, as it is done for US firms by the Survey of Business Uncertainty (The Federal Reserve Bank of Atlanta 2020), or by estimating perceived uncertainty of consumers, where Leduc & Liu (2016) and The University of Michigan (2020) are common sources for US households. Nonetheless, surveys have strong limitations in their use as uncertainty proxies. Due to the verbal nature of surveys, it is hard to distinguish confidence and uncertainty when framing survey questions, which poses the risk of confounding the two concepts. Therefore, scholars have increasingly shifted their focus towards more indirect measures that use econometric tools to approximate uncertainty, discussed next.

Market-based indices typically use the cross-sectional degree in volatility in markets to approximate uncertainty. For instance, aggregate *realized volatility* (RV) of stock prices are a non-parametric method to forecast volatility on financial markets (Andersen et al. 2003). [2] Nonetheless, realized outcomes have only limited power for predicting future outcomes because they do not reveal any information about an agent's expectations about future market outcomes. On the other hand, *implied volatility* of forward looking derivatives contain information about the risk premia of stocks and are therefore useful for identifying investor preferences about the future. The most canonical example of this method is the *option-implied volatility index* (VIX) of the Chicago Board Options Exchange (C.B.O. Exchange 2003). However, despite their common use, volatility measures are commonly misinterpreted (Goldstein & Taleb 2007). The VIX was originally introduced to measure the expected 30-day volatility implied by the S&P 500 equity index in order to resolve the limitation of the Black-Scholes method of pricing options that assumes constant volatility over time. But nowadays it is oftentimes mistaken as a measure of expected future market volatility and therefore also referred to as *investors' fear gauge* (Triana 2009). However, in formal terms, implied volatility represents the *risk-neutral probability* – i.e. the probability adjusted for risks – of the volatility that itself depends on the realized stock market volatility. Hence, the index does not jump unless the market is actually in decline. Another problem is the difficulty to extrapolate uncertainty captured on financial markets to

[2] Alternatively, the higher-order moments (skewness) of distributions were shown to predict (slumps in) economic activity better than second-moment measures (variances or spreads) (Bloom 2009, Ferreira 2018)

uncertainty across the broader economy. Despite its drawback, the VIX counts among the first indices that takes into account the aspect of (investor) expectations in its measure. This aspect is therefore also incorporated into the measurement approach of this thesis. The difference in this paper is that the underlying data is not forward-looking prices but forward-looking *articles*.

A slightly newer strand of methods use *expert forecasts* on economic variables to identify uncertainty, such as the *degree of forecast disagreement*, based on the average dispersion across forecast densities or point forecasts (Abel et al. 2016). Another way is to use *forecast error variances* that account for the timely variation of the unforecastable component of a variable (i.e. the error term). Time-varying uncertainty around forecasts was shown to be particularly relevant for Swedish GDP growth (Ankargren 2020). Jurado et al. (2015) use the conditional variance of the errors of a collective number of US macroeconomic variables to capture macroeconomic uncertainty. The authors found that spikes in their index were less frequent compared to other uncertainty proxies, but showed larger magnitude, persistency and correlation with real economic activity. However, empirical studies showed that these forecast-based uncertainty measures are mostly reliable in economically stable periods (Lahiri & Sheng 2010), thus its usefulness is questionable given that uncertainty appears especially in times of instability and unexpected events.

The introduction of *news-based* uncertainty measures opened a whole new set of opportunities to approximate uncertainty, enabled through recent advancements in natural language processing and machine learning (see Gentzkow et al. 2019 for an extensive review). One of the most popular techniques is the *dictionary method*. These indices are typically constructed by counting the relative frequency of articles containing uncertainty-related keywords (Baker et al. 2016, Alexopoulos & Cohen 2009). The Economic Policy Uncertainty (EPU) index of Baker et al. (2016), based on articles mentioning keywords related to the economy, policy and uncertainty, is the most popular approach taken so far. One advantage of dictionary methods is its easy computation, interpretation, and replication, which is why it has been the most common approach in the literature. However, media outlets might intentionally stop using words such as uncertainty and instead convey uncertainty through more latent manners. [3] An alternative approach was taken by Manela & Moreira (2015) who applied support vector regression to identify a subset of terms on the front page of the Wall Street Journal whose occurrence correlate with variations in the VIX index. They found that the terms that drive most of the variation in market volatility refer to war and public policy. Measuring uncertainty with supervised machine learning has furthermore been advanced by Tobback et al. (2018) who used support vector machines with the purpose of identifying keywords that most likely predict articles labelled as conveying economic policy uncertainty.

An alternative strand of literature recently developed uncertainty measures with the use of *semantic methods* (Azqueta-Gavaldon 2017, Larsen 2017, Saltzman & Yung 2018). Azqueta-Gavaldon (2017) used EPU articles in several languages to identify topics with LDA and fount that the impact of topic-based uncertainty indices highly differed between Euro area countries. However, the selection of topics was supervised. Larsen (2017) further advanced this approach by identifying uncertainty-related topics in a broad set of business articles and using their orthogonal components for creating uncertainty indices. None of these scholars included the aspect of expec-

[3] Anecdotal reference can be found from former World Bank chief economist Joseph Stiglitz, who once explained that because words like 'depression' and 'recession' may cause feelings of anxiety and fear in people, the World Bank actively avoids their usage in its reports. In another instance, Alfred Kahn, former chief of the US administration's task force against inflation, used to substitute 'recession' with 'banana' during meetings with US president Carter.

tations in their measure.

Finally, an increasing number of literature refers to even more innovative data sources, such as Dzielinski (2012) and Castelnuovo & Tran (2017) who use the frequency of internet searches, measured with Google search trends data, in order to proxy uncertainty. These search-based measures are based on the assumption that individuals increase their information search in the occasion of uncertainty. Since no measure is a perfect proxy for uncertainty per se, some scholars also constructed composite indicators, as can be found in Charles et al. (2018) and Haddow (2013).

As can be seen, every measurement technique has its own peculiarities. For this thesis, a news-based approach was chosen given that the possibilities of analyzing latent uncertainty in news are still vast. Moreover, using news articles over other sources for the creation of economic indicators is furthermore especially advantageous for small economies like Sweden where the pool of hard economic variables is oftentimes restricted. News provide high-frequency and large-scale information, while cross-sectional or market based data is oftentimes available at infrequent time intervals only. My approach follows the same path as Azqueta-Gavaldon (2017) and Larsen (2017) in that it uses an unsupervised method to derive meaningful uncertainty indices. However, given that none of the existing news-based uncertainty indices incorporated the aspect of expectations into their measurement, I see it worthwhile to investigate whether expectations in news are as economically important as expectations in financial markets.

2.3 Why Uncertainty Shocks Matter

The central problem to which excessive levels of perceived uncertainty lead to is economic *inaction*, typically in form of postponing decisions in consumption and investment. This section presents how and to what extent these types of inaction reveal themselves in the macroeconomy. Empirical findings on the economic impact of uncertainty shocks can be divided into the different channels through which uncertainty affects economic activity, namely into *demand-side*, *supply-side* and *financial market* impact, as well as the role of *cross-country spillovers* for small and open economies like Sweden. In reality, these channels are not clear-cut because the impact in one channel can obviously affect the others. However, the differentiation between channels is useful for policy purposes since each channel requires different policy responses to shocks.

On the *demand side*, the empirical literature established the stylized fact that uncertainty inversely co-moves with consumption and savings (Castelnuovo et al. 2017, Carroll 1997, Romer 1990). Households have a risk-averse behaviour towards consumption, and uncertainty shocks can cause households to delay consumption in order to secure themselves against temporary income shocks. Moreover, uncertainty skews consumption patterns, where spending on perishable goods were shown to be preferred over durable goods (Romer 1990). The degree of uncertainty in one's labour income furthermore determines the level of *precautionary savings* that is accumulated (Caballero 1990). Households may have a target wealth-to-permanent-income ratio, and if wealth is below the target or if they face income uncertainty, they engage in so called *buffer-stock saving behaviour* (Carroll 1997). In addition, Fernández-Villaverde et al. (2011) showed that precautionary households in small and open economies reduce their holdings of foreign debt when faced with increased interest rate uncertainty in order to reduce fluctu-

ations in their expected marginal utility.

On the *supply-side*, empirical evidence prevails to confirm the stylized facts that uncertainty inversely co-moves with output, investment, returns, and employment (Bloom 2009, Jurado et al. 2015, Scotti 2016, Lazear & Spletzer 2012, Baker et al. 2016). Uncertainty can cause a rapid fall in productivity, output, and employment, from which they quickly recover. Investment and hiring levels tend to overshoot after the recovery period due to labour attrition and growth in business-conditions which causes firms to over-proportionally increase their net hiring and investment rates (Bloom 2009). A rise in uncertainty can furthermore cause large losses in aggregate productivity in the short-term because of inefficient labour allocation and skill-to-job matching, but to aggregate gains thereafter because low productivity firms leave the market and transfer resources to high productivity firms (Lazear & Spletzer 2012). On another dimension, the wait-and-see attitude of firms causes a reduction in resource reallocation given that productive (unproductive) plants pause expansion (contraction) (Bloom et al. 2018).

The demand- and supply-side impact of uncertainty was shown to be amplified through distortions in *financial markets*. On the demand-side for credit, uncertainty causes US firms to increase their holdings of cash and cut their debts to hedge against future shocks, which reinforces reductions in investment and hiring (Lin et al. 2017). On the supply-side of credit, the recessionary impact of exogenous uncertainty shocks on the US economy was shown to be up to five times larger during tightened credit regimes and accounted for 1% of the reduction in industrial production during the recession period between 2007 and 2009 (Caldara et al. 2016, Alessandri & Mumtaz 2019). After uncertainty shocks, banks tend to tighten credit conditions, which especially affects new and less-established firms. Since these firms typically drive economic innovation and growth, these restricted investment opportunities can be particularly economically harmful (Gilchrist et al. 2014). Alessandri & Bottero (2017) furthermore revealed that uncertainty reduces the likelihood that banks accept new credit applications. What is more, banks' lending behaviour becomes less responsive to changes in the short-term interest rate which implies that bank lending is a weak channel for the effectiveness of monetary policy during uncertain periods.

When it comes to *cross-country spillovers* of uncertainty shocks, the establishment of global value chains since the 20th century have led to spatial transmission in innovation, productivity, but also uncertainty (Feldman 1999, Davis 2016). Davis (2016) found that spikes of economic policy uncertainty in countries that account for two thirds of global output are related to global events, such as the Global Financial Crisis or the European immigration crisis. Especially for small open economies this spillover effect has important economic consequences. In Sweden, the spillover of foreign uncertainty shocks seems to have a delayed effect on the Swedish economy compared to domestic uncertainty shocks, which negatively impact economic activity in the same period (Armelius et al. 2019).

These findings have crucial implications for fiscal and monetary policy. Because the financial market and the supply side of the economy are less responsive to policy decisions during uncertain times, it is of highest relevance for policy makers to have useful effective measures at hand which identify periods of elevated uncertainty and with which they can estimate the effectiveness of their policies. With the enhancement of the Swedish news-based uncertainty index, this thesis seeks to enable exactly this.

3 Foundations of NLP for Macroeconomic Analysis

Natural Language Processing (NLP) is a sub field of statistical machine learning that combines the disciplines of computational linguistics and applied mathematics to derive methods for retrieving numerical information that is encoded in large-scale, unstructured textual data. While its use in the field of economics has become widespread just in the last couple of years, some scholars have already derived powerful applications for economic research, some of which are extensively reviewed in Mullainathan & Spiess (2017), and for NLP in particular in Gentzkow et al. (2019). A review of NLP applications for econometric analysis goes beyond the scope of this thesis. Therefore, this section focuses on a short introduction into the key concepts of NLP and its most important applications in macroeconomic research in order to facilitate the conceptual understanding of the methodology pursued in this thesis. The reader is encouraged to inspect Table B1 in the Appendix which provides definitions of the key terminologies that are used used in NLP analysis and in the subsequent section.

Common to all statistical text analysis applications is the conversion of a text corpus D with individual documents d into a numerical array $\mathbf{C}$. Each row, $\mathbf{c_d}$, represents a single observation d (i.e. a document) and the numerical vector representation of the frequency (or presence) of all its pre-defined tokens. For simplicity, assume that tokens refer to terms and call their numerical vector representation *term-frequency vector*, where the term frequency of term n in observation d is denoted $\text{tf}(d)^n$. The term frequencies define the vector space $V : d \rightarrow \mathbb{R}^{\dim(N)}$, where the dimension of N is determined by the set of all terms present in the corpus. Text data inherently differs from more conventional data sources in that $\mathbb{R}^{\dim(N)}$ is highly dimensional which can pose severe computational problems. Therefore, prior to mapping $\mathbf{C}$, it is necessary to reduce the dimensionality of N – i.e. the number of tokens – through well-established pre-processing techniques, which are also applied in this analysis and described in Appendix B. (Gentzkow et al. 2019, Kalamara et al. 2020)
The mapping of $\mathbf{C}$ is subsequently used to predict values $\hat{\mathbf{V}}$ of unknown outcomes $\mathbf{V}$ with statistical methods. These methods can be broadly aggregated into three different categories which differ from each other by the level on which they process text: *lexical or dictionary* methods, which is the analysis of token occurrences or frequencies in a corpus; *syntactic* methods, which refers to the analysis of ordered sequences of tokens within document sentences of a corpus; and *semantic* methods, i.e. the analysis of the conveyed meaning of tokens, usually by analyzing their co-occurrences to other tokens.

Applications in macroeconomic research are mainly focused on generating high-frequency and accurate predictions rather than causal inference of the mapping from $\mathbf{C}$ to $\hat{\mathbf{V}}$ (Gentzkow et al. 2019). For instance, a classic text analysis task in macroeconomic time series is the sentiment analysis of news in order to improve predictions of the real economy (see Doms & Morin 2004, Shapiro et al. 2020) or financial markets (see Nyman et al. 2018, Garcia 2012). The unknown variable $\mathbf{V}$ represents the true sentiment of an economy at time t, while prediction $\hat{\mathbf{V}}$ identifies the degree of positivity or negativity conveyed in the newspapers of the same economy at time t.
For instance, Kalamara et al. (2020) extensively evaluated different sentiment analysis methods applied on newspapers to improve more 'traditional' forecasts of GDP, inflation, and unemployment, and concluded that dictionary methods enhanced with supervised machine learning yield the best forecast improvement. Thorsrud (2020) and Larsen & Thorsrud (2019*a*) recently derived a high-frequency, nowcasting measure of Norwegian GDP out of latent topics in newspapers. Alternative high-frequency applications can be found in Baker & Fradkin (2011) who

use Google search queries to predict the job search activity of recipients of unemployment insurance. Their application is especially attractive to the more traditional survey methods in labour market research where people might refrain from honestly revealing their true actions and beliefs, while Google search queries reveal a less biased picture of labour markets.
The analysis of news has furthermore contributed to the better understanding of expectations and 'animal spirits'. Larsen et al. (2019) found that news topics have a predictive potential for both inflation and inflation expectations. Famously, Tetlock (2007) predicts stock market prices and trading volumes with the degree of pessimism in Wall Street Journal columns. In another paper, the authors found that share of negative words in financial news is predictive for low earnings and returns and capture fundamental aspects of firms that are otherwise difficult to quantify (Tetlock et al. 2008).

As can be seen, the majority of macroeconomic applications have been focused on using newspaper sentiment and narratives for the prediction of hard economic variables, such as GDP or stock returns. However, even though these studies found a relationship between news and economic activity, little is known about the channels through which news affect economic outcomes. One potential and plausible channel would be the market participants, i.e. firms and households, whose beliefs are changed by narratives in the news. A change in beliefs can lead to a change in behaviour which consequently affects economic activity. The analysis conducted in this study aims to fill this missing revelation in the literature by analyzing how uncertainty conveyed in news can affect the confidence of firms and households. Furthermore, it counts among the few applications that focuses on the role of news in a small and open economy, namely Sweden, and therefore contributes to the understanding of the role of news outside of autarkic economies.

4 Measurement of News-Based Expectation Uncertainty

This section of the thesis is concerned with measuring economic uncertainty on the basis of newspaper articles with the purpose of creating news-based uncertainty indices. The first sub section provides a description of the Swedish newspaper sample used for the index creation. Subsequently, I explain the two pursued methodologies in detail. The two methodologies I explore are the (supervised) dictionary method and the (unsupervised) semantic method. The first approach is pursued since it counts as the most standard method for uncertainty identification in news. It was used by Armelius et al. (2017) to create the Swedish Economic Policy Uncertainty (EPU) index, against which I compare the uncertainty indices developed in this thesis. However, one of the major limitation of the dictionary method is that dictionaries are manually crafted and relatively restricted in the selection of vocabulary they entail, thereby potentially causing noise in the selected article sample. Therefore, the second approach serves as an alternative investigation into whether the unsupervised identification of uncertainty – specifically uncertainty-related narratives – can help to more precisely capture uncertainty.
Upon reading this section, the reader will realize that the measurement procedure is an iterative process. Hence, the obtained uncertainty indices are not reported separately but are included and analyzed in this same section.

4.1 Data Sourcing & Sampling Strategy

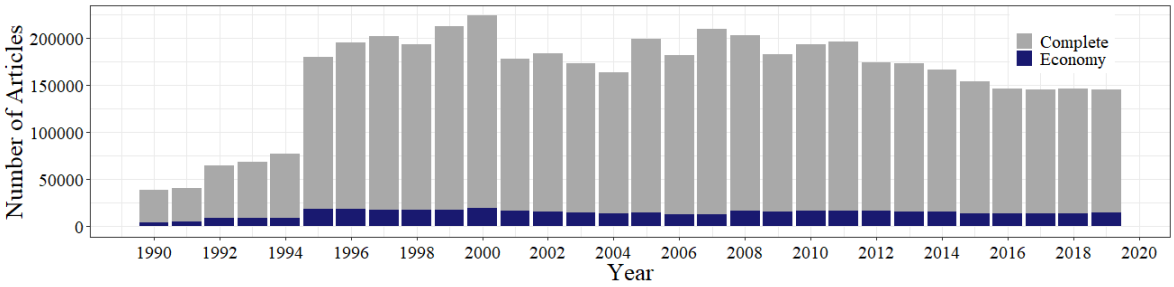
The basis for constructing the uncertainty index is laid by Swedish newspaper articles published in print between 01.01.1990 and 31.12.2019 by the following Swedish publishers: 'Dagens Nyheter', 'Dagens Industri', 'Svenska Dagbladet', 'Aftonbladet', and 'Expressen'. These are national publishers that count the highest annual number of published editions and have the largest readership in Sweden. Two things are important to consider in the selection of the news sources: First, it is assumed that it is news about nation-wide or international stories that will shift readers' beliefs about the economy by a considerable scale. These stories are likely to be more extensively reported in national rather than regional news. It might be true that households are more affected by news coverage related to specific events happening within their proximity. From a macroeconomic perspective, however, regional shocks to beliefs will most likely have a negligible impact on the Swedish economy overall. Moreover, if a regional event is relevant enough to have consequences on a national level, national news are very likely to pay attention and include them in their news coverage, so that these events are accounted for in this analysis. Second, Sweden has a variety of regional publishers which might be more frequently read but which are not included in the data. However, regional news about major economic events undoubtedly source from larger national news sources and it is assumed that national coverage about the stories of interest will spill over to regional coverage. Thereby, the stories covered this data set will sufficiently approximate the stories directly consumed by the majority of readers.

Within the given period, the total number of printed articles well exceeds 4 million publications. Figure 2 presents the total yearly article counts over the analyzed period. Between 1990 to 1994, the annual numbers amount to only around one third of articles in the rest of timeline, where it roughly ranges between 15,000 and 20,000 articles per year. I use a subsample of the entire corpus for all subsequent analysis by approximating articles related to economic subjects. Therefore, I select articles that mention the roots *ekonomi* or *ekonomisk* and exclude articles that include the keywords $\{$*sport, spel, fotboll, finalist, cup, familj, kvinn, barn, bok*$\}$. I furthermore filter out duplicates and end up with a corpus of 430,466 articles. Filtering articles with these keywords is reasonable because, first, it

improves the signal-to-noise ratio of the corpus. Even though I might miss articles that talk about economic expectations without explicit mentioning, I also avoid articles that are not necessarily economically meaningful, such as weather forecasts. Second, the subsample is only around 10% of the size of the total sample, which substantially reduces the computational burden for conducting all subsequent analyses.

The articles were provided by the Swedish media archive Retriever in a *tokenized* and *bag-of-words* format,

Figure 2 Number of Articles per Month.



meaning that each article is presented as a vector, $\mathbf{c_d}$, whose length is equal to the number of unigrams present in that article and whose elements c_{dn} are the absolute frequencies with which word n occurs in article d (see illustration in Figure A1 in the Appendix). This special format is due to copyright requirements of Swedish publishers. The raw texts, titles and publisher names of the articles are not identifiable i.e. the only metadata available is the date of publishing. This also implies that any syntax or paragraph analysis is not possible. However, this restriction does not pose any limitations on the principal methods pursued in this thesis. To test the validity of the sample, I replicate the existing Swedish news-based index – the EPU index – by following the method of Armelius et al. (2017).[4] Figure A2 in the Appendix compares the original and replicated EPU series which have a correlation of 0.71. The reason of no perfect correlation can be twofold: First, I include articles from 'Dagens Nyheter' in my sample, whereas the original paper does not. This will skew the number of articles used for the index calculation. Second, since I also include inflections of the keywords, their variation is likely not completely identical to the authors', which can result in more or less articles included. The high sensibility of news-based indices when it comes to their replication should be kept in mind for their validity in practice, which I further elaborate in §6. Having that said, I still regard the replication as successful since the trend between the two series is almost identical and the uncertainty peaks differ mostly in their magnitude. Therefore, I see the article sample as valid for all subsequent analyses.

4.2 Dictionary Method

Dictionary text analysis is concerned with the analysis of word or document occurrences based on pre-specified lists of keywords. The method pursued here is based on Baker et al. (2016), who introduced the Economic Policy Uncertainty (EPU) index. I identify the monthly number of articles in the sampled corpus that mention keywords

[4] I count the monthly number of articles where keywords of the three components **E**conomic ('ekonomi', 'ekonomisk'), **P**olicy ('riksbank', 'centralbank'. 'regering', 'departement', 'reglering'), and **U**ncertainty ('osäker','oro') co-occur. I also include words that use these keywords as roots. I then normalize the monthly count by the number of articles per month that mention the economic component and scale the series.

related to specific categories, thereby using three sets of dictionaries. The relative frequency of these articles over time consequently comprises the first obtained index.

Articles are counted if they mention at least one token from dictionary *economy*, $\mathbf{e} \in \{$*ekonomi, ekonomisk*$\}$, from dictionary *forecast*, $\mathbf{f} \in \{$*utsikt, förhandsberäkning, prognos, beräkning, förutsäga, förutsägelse, uppskattning, förväntning, förhoppning, förmodan, förväntan, framtid, spå, förutspå, bebåda, förkunna, förutse, prediktera, förutbestämma, hopp*$\}$, and at least one token from dictionary *uncertainty*, $\mathbf{u} \in \{$*osäkerhet, osäker, oro*$\}$. Whenever adequate, I also include the inflections of these keywords. The tokens in dictionary **e** and **u** are standard selections from the literature, while the tokens in dictionary **f** were identified through introspection and translated from English reference words to Swedish. Synonyms were found with the Sweden's People's Dictionary (*Folkets Lexicon*) and subsequently cross-evaluated by a native Swedish speaker. The combination of these three dictionaries aims to approximate articles about uncertain expectations about the economy in the news either in form of forecasts or beliefs (similar to Baker et al. (2016) who assume that articles with keywords about the economy, policy, and uncertainty approximate economic policy uncertainty). In total, 26,850 articles are counted. The resulting index is in the following called EFU_t (Economic Forecast Uncertainty).[5] Specifically, the calculation looks as follows:

$$EFU_t = \frac{\sum_{i=1}^{D_t} d_{i,t}(e=1, f=1, u=1)}{\sum_{i=1}^{D_t} d_{i,t}}, \tag{1}$$

where D_t is the total number of articles in the corpus published in month t, and $d(e=1, f=1, u=1)$ counts as one if article d mentions a word of each keyword lists e, f, u, and counts as zero otherwise. I normalize the indices by total article counts in order to account for time-varying structural factors, such as the rising number of articles published from 1995 onwards (see Figure 2).

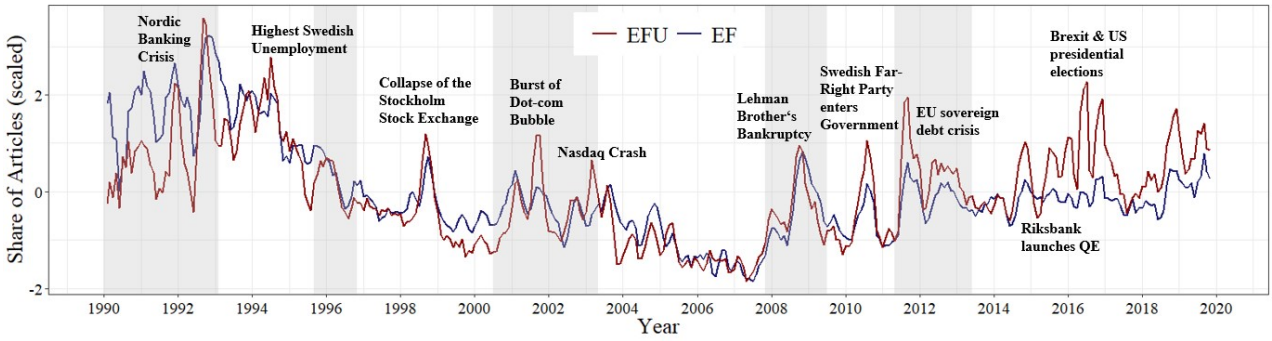
I create an additional index which accounts for the amount of news that convey *general* forward looking expectations about the economy, therefore called EF_t (for Economic Forecast). This is done by counting articles that mention tokens related to *economy* (dictionary **e**) and *forecast* (dictionary **f**), but *not* to uncertainty. The more general EF index will play a more important role in the evaluation of the uncertainty indices in §5. Looking at EFU articles alone ignores the fact that news about expectations in general rises during periods of economic stress and uncertainty, which might falsely suggest that it is news about uncertain expectations only that have an economic impact. In total, EF comprises 161,482 articles and is calculated as:

$$EF_t = \frac{\sum_{i=1}^{D_t} d_{i,t}(e=1, f=1)}{\sum_{i=1}^{D_t} d_{i,t}}, \tag{2}$$

Figure 3 plots the obtained time series of EFU and EF. It can be seen that the series follow a similar trend and seem to increase during periods of economic contractions (marked grey). However, the EFU index has more nuanced peaks compared to the more general EF index. These peaks are further detailed in Table B2 in the Appendix. The discrepancy between EF and EFU especially stands out between 1998 and 1999, where the Russian financial crisis and the collapse of the Stockholm Stock Exchange happened. Higher peaks are also found in late 2001 after the burst of the Dot-com bubble and after the Nasdaq's slump in October 2002. During the Brexit votum and the US presidential election of Donald Trump in 2016, the difference between the indices is again clearly visible. As

[5] In the following, *economic expectation uncertainty* and *economic forecast uncertainty* are used interchangeably.

Figure 3 EFU versus EF and Uncertainty-related Shocks..



Notes: Rolling means of the scaled series with a window of 3 months are presented in order to facilitate comparison. Grey segments indicate recession periods in the Swedish economy as defined by the OECD.

can be seen, uncertainty spikes might relate to both Swedish and international events. However, it is plausible to assume that external uncertainty elevates uncertainty in Sweden's economy to some degree given that it is a small and open economy.

Noticeable are also periods where EFU falls below EF. This was the case before the recession periods in 1995 and in 2000, and also in the early years before 2008. After the Great Recession, however, uncertainty played a more consistent role in news about economic expectations, as can be seen by the higher occurrence of EFU articles compared to EF articles even in periods of economic growth. The only time EFU falls below EF is in 2015. This might have been due to the aftermath effects of the European sovereign debt crisis when the European Central Bank launched its bond-purchasing program. This program ended in late 2014 when EFU started to rise again. Nonetheless, the Riksbank, the Swedish central bank, launched its quantitative easing program in February 2015, following central bank decisions in US and the EU, which might have caused EFU to slump.

These depictions are obviously anecdotal and little is yet known about what type of stories EFU articles cover and whether they are economically relevant. For example, it is not clear whether articles convey uncertainty about issues in specific sectors (e.g. financial or labour market) or regions (e.g. Sweden or elsewhere). For this purpose, it is useful to measure the latent thematic components of the articles over time. This will be done in the next section.

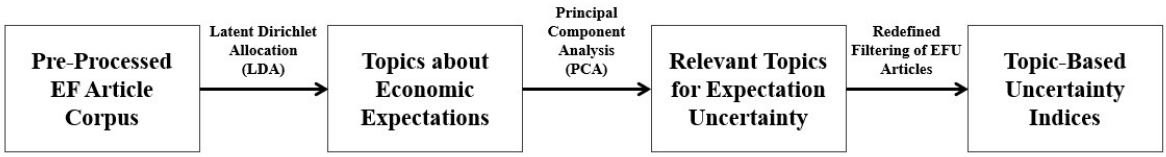
4.3 Semantic Method

The previous section used manually selected keywords in order to identify articles about uncertain economic expectations and counted their relative frequency over time in order to create an index. In contrast, the objective of this section is to identify latent *narratives* related to economic expectations that are relevant for uncertainty, i.e. diving into the semantic structure of articles.

The procedure is composed of three main steps. First, I identify topics in articles about economic expectations

with Latent Dirichlet Allocation (LDA). These topics aim to approximate the broader and more fuzzy concept of narratives. LDA is chosen primarily because it can account for the *intra*document latent structures of a text and the *inter*document latent structure of the entire corpus at the same time, therefore bypassing limitations of older methods for latent text classification in the field of Information Retrieval. [6] I use the more general corpus of economic forecast articles (EF) because I can improve the quality of the topic inferences with a larger sample size ($\sim$160,000 versus $\sim$ 26,000). As a second step, I use the inferred topic compositions and measure the timely occurrence of the topics in the EFU articles. Subsequently, I use Principal Component Analysis (PCA) to identify those components that explain most of the variation in the occurrence of topics over time. As a last step, I identify those topics that are most strongly connected to the principal components and create topic-based uncertainty indices by counting articles in the EFU corpus that entail these principal component topics with a high probability.
The outlined procedure allows me to identify topics related to uncertainty and to distinguish them from other less relevant topics. In the following sub sections, each of the steps – presented by the arrows in Figure 4 – will be explained in detail. Before starting the process, the LDA algorithm requires further pre-processing of the tokenized articles in order to reduce the dimensionality of the corpus. The steps undertaken are detailed in Appendix B.

Figure 4 Identification Process of the Semantic Method.



Latent Dirichlet Allocation

Latent Dirichlet allocation (LDA) was introduced by Blei et al. (2003) and Blei (2012), and serves as a powerful tool for finding hidden structures in highly heterogeneous data. This makes it especially attractive for identifying topics discussed in documents given the elaborate relationship among words in textual data. LDA is therefore also called *probabilistic topic modelling* and falls into the category of *generative probabilistic modelling*. *Generative* means that articles are generated from latent, i.e. hidden variables, which in this case are topics. The generative process is defined by a joint probability distribution over both the latent and the observed variables, the latter being the fixed vocabulary of the articles corpus. *Probabilistic modelling* means that data analysis makes use of this joint probability distribution to infer the conditional distribution of the latent topics given the vocabulary, also called the *posterior distribution*. For LDA, articles are composed of probability distributions over topics and topics are composed of probability distributions over words. The only assumption made is that LDA assigns positive probabilities for every topic in articles, but the degrees change depending on the article. The word order in the documents and the document order furthermore do not matter for computing the posterior distributions (i.e. bag-of-words assumption).

Formally, the following notation is used to describe the generative process: A topic, denoted as k, is an element of the total set of K topics. The set of topics can be expressed as $\phi = \{\phi_1, ..., \phi_k\}$, where each ϕ_k is a probability

[6] These are the *term-frequency-inverse-document-frequency* (td-idf) of tokens, *Latent Semantic Indexing* (LSI), and *probabilistic LSI* (pLSI). See Blei et al. (2003) for a thorough technical comparison of different text classifying methods.

distribution over the fixed vocabulary of the corpus and expresses the probability for each word to occur in topic k. On the other hand, an article is defined as $\theta_d = \{\theta_{d,1}, ..., \theta_{d,k}\}$, where $\theta_{d,k}$ expresses the probability with which topic k appears in article d. Next, we are given an article corpus D, where each article d is composed of a finite sequence of N unique words, denoted by $d = (w_1, w_2, \ldots, w_N)$, where w_n is the n-th word in the sequence. Finally, the probability that a word appears in topic k, and the probability of topic k to appear in article d is given by (Griffiths & Steyvers 2004, p. 5228):

$$p(w_n|topic = k) = \phi_{n,k} \quad \text{and} \quad p(k|article = d) = \theta_{d,k}. \tag{3}$$

Consequently, the joint probability distribution of the hidden and observed variables corresponds to (Blei 2012, pp. 80):

$$p(\phi, \theta, z, w) = \prod_{i=1}^{K} p(\phi_i|\beta) \prod_{d=1}^{D} p(\theta_d|\alpha) \left(\prod_{n=1}^{N} p(z_{d,n}|\theta_d) p(w_{d,n}|\phi, z_{d,n}) \right), \tag{4}$$

where $z_{d,n}$ is the topic assignment of a word n in article d and $w_{d,n}$ is the observed word n in article d. α and β are Dirichlet parameters, which are further detailed below. Given this notation, the LDA algorithm follows the subsequent process to generate each article d in corpus D (Blei et al. 2003, pp. 996):

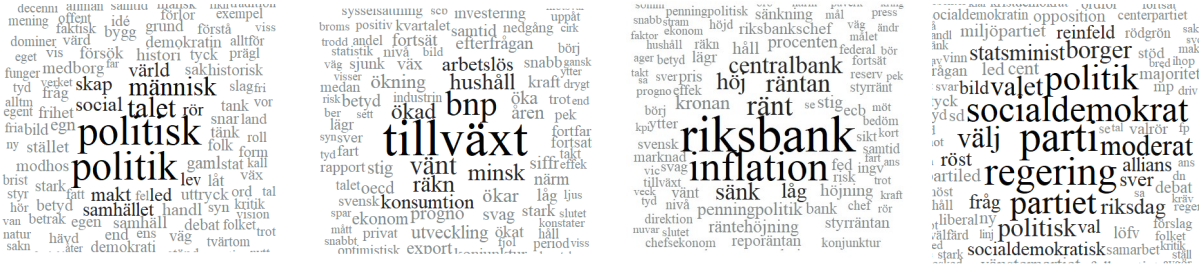
1. Choose the number of words, $N \sim Poisson(\xi)$.
2. Choose a topic distribution over a fixed set of topics, $\theta_d \sim Dir(\alpha)$.
3. For each of the N words $w_{d,n}$:
   (a) Randomly choose a topic $z_{d,n} \sim Multinomial(\theta)$ from the distribution obtained in Step 2.
   (b) Given that topic, randomly choose a word $w_{d,n}$ from $p(w_{d,n}|z_{d,n}, \beta)$, a multinomial probability conditioned on the topic.

Hence, LDA uses these steps to infer the underlying topic structure of the corpus, where each topic is a probability distribution over words (ϕ_k), and each article is a probability distribution over topics (θ_d). Knowing the algorithm for generating the documents, the distributions ϕ and θ can be estimated by computing the posterior distribution of the hidden variables. In order to preserve space, the inference procedure is detailed in Appendix B.

Before estimation, there are hyperparameters that need to be set manually, namely α, β, and K. The scalars α and β are Dirichlet priors which respectively determine the characteristics of the topic distributions over documents ($p(\theta|\alpha)$) and the word distributions over topics in a document ($p(\phi|\beta)$). For β, higher values result in a more specific word density over topics, i.e. topics will be characterized by few words with large probabilities and many words with low probabilities. Similarly, values of α determine the specificity of the topic density over documents.[7] A high value of α leads to more dense topic probabilities per article, which implies that the articles compositions are fairly uniform and that articles are therefore similar to each other. In contrast, a low α implies more sparse articles and skewed topic probabilities, so that articles are very different from each other and characterised by a few topics. Readers unfamiliar with the Dirichlet distribution can consult Appendix B which further explains the role of the Dirichlet priors and its influence on the inference procedure.

[7] Note that α is directly linked to the Dirichlet distribution in step 2 of the generative process, and hence is responsible for the concentration of the Dirichlet distribution from which θ is sampled.

Figure 5 News Topics of the EF Article Corpus in Form of Probability Distributions over Words.



(a) Politics & Society (b) Business Cycles (c) Central Bank (d) Government

The selection of these priors does not rely on principled techniques, so I determine their values based on assumptions about the corpus structure. Articles are more or less centered around news about economic expectations (by the nature of filtering). However, the filtering procedure was kept general, so there might still be substantive variety in the article themes. I therefore select $\alpha = 0.5$, so that the topic distributions in articles are spread out, although certain topics will be more dominant than others which enables me to better divide between relevant from irrelevant articles (e.g. an article might talk about the topic economic recession but even more about the topic of sports clubs facing financial difficulties). For determining the probability distribution of the topics, I set $\beta = 0.001$. This will imply that a small selection of words will be highly characterizing for topics and helps to find nuances among topics with a similar overall theme (e.g. topics 'Elections' and 'Government' fall under the overall theme of 'Politics' and can still be distinguished).[8]

The last parameter that needs to be determined is K, the total number of topics. It is clearly impossible to know the exact number of topics present in the corpus. In order to choose the topic number objectively, I follow the original authors' proposition to compare the maximum log-likelihood of topic densities on a held-out article corpus – in language modelling also referred to as *perplexity* (Blei et al. 2003). The exact procedure can be found in Appendix B. The optimal topic number derived is $K^{opt} = 43$.

Table B3 in Appendix B entails a list of the inferred topics together with the 20 tokens that have the highest probability of belonging to the respective topic. The only supervised step is the manual labelling of the topics, which serves to facilitate their interpretation. The figures in 5 list examples of topics that occur most frequently in the articles in form of word clouds. The size of the words in the cloud increases with the probability of belonging to that specific topic. Topics which occur frequently in the articles are intuitively related to the subject of economic expectations. One central subject of economic forecasts are business cycles, which is also highly reported in Swedish news (Figure 5b). Political themes are equally strongly related. Governments can detrimentally influence economic development through public policy decisions, especially in Sweden, given the strong influence of its public policy regimes (Figure 5a and 5d). However, as can be seen in Table B3 not all topics will be necessarily relevant for uncertainty, such as topics 'Sports'. The procedure for selecting uncertainty-relevant topics will be detailed in the subsequent sections.

[8] Similar studies, such as Larsen & Thorsrud (2019*a*) and Larsen (2017), used $\alpha = \frac{50}{K}$ and $\beta = \frac{200}{N}$

Principal Component Analysis

The LDA model provides us information about the underlying topics inherent in articles about economic expectations. In the subsequent step, the following questions are sought to be answered in order to derive uncertainty measures from the topics: First, what is the prevalence of the topics in articles about economic forecast uncertainty over time? Second, how much uncertainty do these topics convey at each time step? These questions are answered by translating the probability distribution of articles into meaningful time series, denoted as $T_{t,k}$. The conversion looks as follows: We are given m articles about economic forecast uncertainty (all EFU articles). For each topic k, the sum of its probability distribution in the articles (i.e. the sum of $\theta_{m,k}$) in month t is multiplied by the number of uncertainty tokens u_m appearing in t. The sum of $\theta_{m,k}$ approximates the timely prevalence of a topic, meaning that the more probable a topic is to appear in the sum of articles in a given month, the more relevant that topic is for that certain period for economic forecast uncertainty. The multiplication by the number of uncertainty tokens serves as a scaling purpose. The more uncertainty tokens are used in a given month, the more annotated are highly prevalent topics in those articles mentioning uncertainty. This procedure provides me with a raw measure of the uncertainty-adjusted prevalence of topics over time, which I then normalize by the total number of EFU articles at t:

$$T_{t,k} = \frac{\sum_{m \in t} u_m \theta_{m,k}}{\sum_{i=1}^{D_t^{EFU}} d_{i,t}^{EFU}} \tag{5}$$

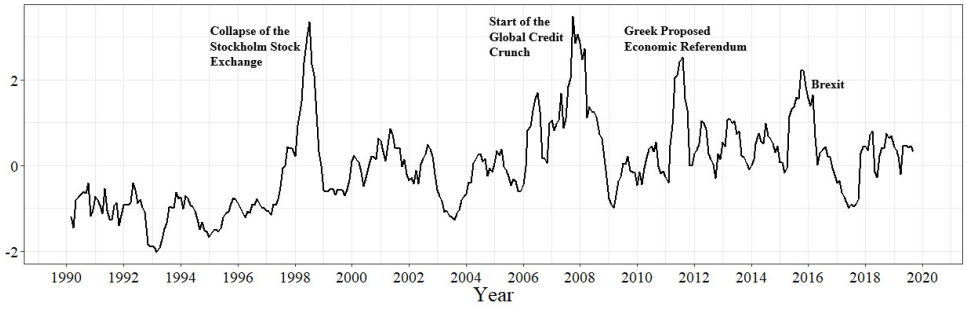
where $\theta_{m,k}$ is the probability that topic k appears in article m in month t. Normalization is necessary since the number of EFU articles, and therefore also the sum of u_m, increases over time and would otherwise inflate the prevalence measure. I end up with 43 different $T_{t,k}$. Examples can be examined in Figure 6.

The obtained measures could hypothetically already be used as proper uncertainty indices. However, there are several aspects that speak against it. First, little is yet known about which of these topic-based indices are the most important drivers of the EFU corpus. Neither the general prevalence nor the uncertainty-adjusted prevalence of topics directly tell how much of the total timely variation is captured by the individual topics. Second, many topics share common, domain-specific vocabulary which indicates that they are semantically related. It is extremely likely that they capture similar narratives and consequently also similar uncertainty-related events. An illustration is provided in Subfigure 6d in 6 between the topic time series of 'Central Bank' and 'Global Risk'. The latter topic is more prevalent during well known periods of international crises (the Russian and Asian Crises in the late 1990s and the European Sovereign Debt Crisis in the early 2010s). The 'Central Bank' topic reacts less to international shocks, which strongly suggests that the topic is related to Sweden's national central bank, i.e. the Riksbank. However, both topics frequently co-move which suggests that they still might capture an intersection of events that are relevant for narratives about both international risks and the Riksbank. For instance, the European Sovereign Debt Crisis in the early 2010s caused high levels of international uncertainty and risk due to the systemic nature of the Eurozone. At the same time, it was highly relevant for Sweden's (small and open) economy due to its proximity to Eurozone economies. It is therefore useful to look at certain topics in combination over time. Nonetheless, manual aggregation is burdensome and is prone to the risk of not capturing all relevant and related information.

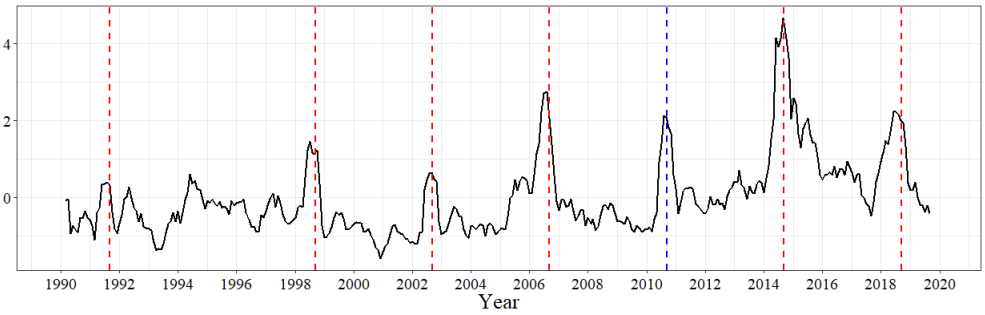
Fortunately, the temporal inter-topic relationship can be investigated more systematically. Instead of relying on the supervised aggregation of topic indices, I continue my procedure with an unsupervised approach by using Prin-

Figure 6 Examples of the Uncertainty-Adjusted Prevalence of Topics over Time.

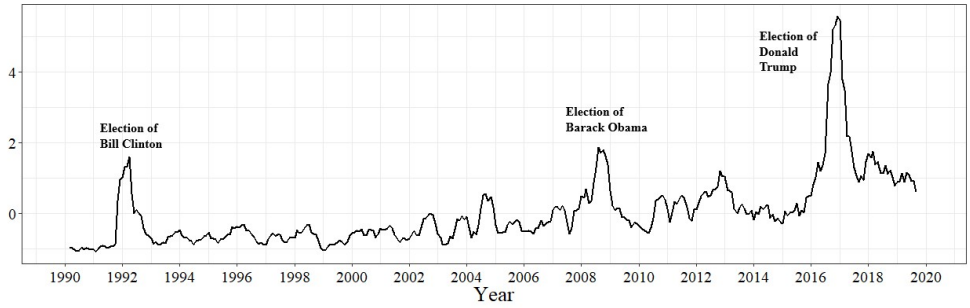
(a) Stock Exchange Uncertainty (Topic 35)



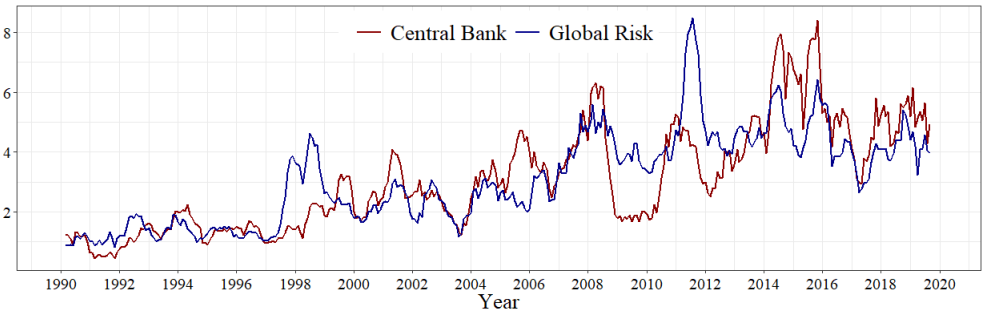
(b) Governmental Uncertainty (Topic 4)



(c) US Presidency Uncertainty (Topic 6)



(d) The Timely Co-movement of Topics 'Global Risk' (Topic 16) and 'Central Bank' (Topic 3)



Notes: 3-month rolling means are presented. In panel 6b, the dotted lines indicate election dates. Red dotted lines indicate that the majority of the government elected was left-leaning, while blue dotted line indicates that the government elected was right-leaning.

cipal Component Analysis (PCA). PCA is especially useful for data sets where a large number of features needs to be reduced into a smaller set of features that still capture most of the information in the data and are at the same time uncorrelated to each other, i.e. the so called *principal components* of the data. I use PCA in order 1) to determine which narratives capture common principal components over time in the EFU article corpus, and 2) to identify which narratives are the most relevant for economic forecast uncertainty in order to refine the identification of uncertainty for further analysis.

The procedure starts with a t x K matrix $\mathbf{X}$, with t individual observations (monthly intervals) and K-dimensional features, where K is the number of topics available. Each feature corresponds to the centered values of each T_k. PCA finds a reduced set of dimensions, where each dimension is a linear combination of the K features. The principal component i of the first time step is the linear combination of the features $X_1, X_2, ..., X_K$ with the largest sample variance (Hastie et al. 2009, pp. 375):

$$z_{i1} = \phi_{11} X_{i1} + \phi_{21} X_{i2} + ... + \phi_{p1} X_{iK}, \tag{6}$$

where $\phi_{11}, \phi_{21}, ..., \phi_{p1}$ are referred to as the *loadings* of the principal component and z_{i1} is the principal component score. The loadings are normalized to $\sum_{j=1}^{K} \phi_{j1}^2 = 1$, which prevents that they capture arbitrarily large variance. The combination of the former equations translates into an optimization problem for each principal component loading vector. For instance, the first principal component solves:

$$\underset{\phi_{11}, \phi_{21}, ..., \phi_{K1}}{\text{maximize}} \left\{ \frac{1}{t} \sum_{i=1}^{t} \left(\sum_{j=1}^{K} \phi_{j1} x_{ij} \right)^2 \right\} \quad \text{subject to} \quad \sum_{j=1}^{K} \phi_{j1}^2 = 1 \tag{7}$$

The number of components computed is $min(n-1, p)$. Hence, in this case, I end up with 43 components. Each component identified will explain a distinct proportion of variance, called PVE. The PVE is simply the proportion of variance explained by the m-th component relative to the total variance:

$$PVE = \frac{\frac{1}{t} \sum_{i=1}^{t} \left(\sum_{j=1}^{K} \phi_{jm} x_{ij} \right)^2}{\sum_{j=1}^{K} \frac{1}{t} \sum_{i=1}^{t} x_{ij}^2}, \tag{8}$$

where the numerator defines the variance explained by the m-th principal component and the denominator denotes the total variance of the centered data. Hence, the sum of PVEs of all components sum to one. For all subsequent steps, I select only the principal components that can explain 5% or more of the sample variation, which are two in total. These two components can explain around 72.8% of the entire sample variation. Figure B6 in the Appendix plots how much variation each of these two principal components captures over time.

It is important to note that the sign of the principal component scores at each time interval cannot be directly linked to uncertainty: whether a component explains more or less variation at a time does not translate to more or less uncertainty in the news. Rather, the narrative interpretation of the components is derived by looking at the feature loadings ϕ_k of each topic in each selected principal component. Table 1 reports the 5 topics with the highest absolute feature loading in each respective component. Although the topic loadings are relatively low in their absolute value, the identified topics can be seen as naturally relevant for articles about economic expectations and are also strongly linked between each other. In the following, a narrative description of the topic compositions in Table 1 is attempted.

The first principal component captures a substantial part of the total variation (PVE = 67%) and, more than clearly, relates to topics about the economy. Narratives about business cycles and global risks are inherently interlinked

Table 1 Top Principal Components - Narrative Description.

Component 1		**Component 2**	
(Number) Topic	Loading	(Number) Topic	Loading
(3) Central Bank	0.40	(27) Politics &Society	0.42
(35) Stock Exchange	0.36	(4) Government	0.36
(22) Business Cycles	0.33	(8) European Union	0.27
(25) Currencies	0.31	(3) Central Bank	-0.30
(15) Probabilities	0.28	(35) Stock Exchange	-0.33

Notes: The numbers on the left side of the topic labels correspond to the topic numbers in Table in B3 in Appendix B.

given that fluctuations in business cycles lead to higher (perceived) risks. For a small and open economy it is also plausible that the subject of currencies play an important role for communicating economic prospects. Finally, aspects of monetary policy seem to play a leading role for all of these topics (Topic 3 'Central Bank'). The 'Probabilities' topic does not provide a lot of room for interpretation. Nonetheless, narratives about the other four topics might use especially many probability-related words that express risk.

The second principal component explains much less variation (PVE = 6%) and its narratives are much more related to what the initial EPU index tries to capture. Judging from the word distributions, the topic with the highest loading is related to social and political narratives, closely followed by a topic about the government. The latter topic is most likely related to the national Swedish government because it includes the token 'riksdag' with a high probability. However, this component also captures narratives about international politics given that the third topic is clearly about the European Union. The relation between these three topics might indicate that narratives about policy issues drive the second highest timely variation in the EFU articles. The opposite signs of the last two loadings have a direct interpretation: News about politics, the government or the European Union move in opposite direction to news about the Riksbank and the stock exchange. This can be interpreted in relation to time, i.e. during times of macroeconomic or financial distress, news focus less on policy matters and more on narratives about the central bank and the stock market. Alternatively, it can also be interpreted in relation to newspaper outlets, i.e. newspapers covering economic and financial news are less likely to cover news about the other topics.

Topic-Based Uncertainty Indices

The previous section revealed which topics capture common components over time in articles about economic forecast uncertainty and how much variation they explain overall. In the following final step, topic-based uncertainty indices are derived based on the outcome of the PCA. Concretely, I construct two distinct uncertainty indices, each corresponding to the topics in the two principal components depicted in Table 1.

For each index and component, I identify the relative number of EFU articles where the topics assigned to the respective component are most probable to occur. That means that EFU articles are counted if at least one of the topics with a *positive* loading presented in Table 1 count among the *top three* most probable topics to appear in the articles (i.e. articles with the highest $\theta_{d,k}$ belonging to these topics). This procedure is based on two important assumptions: First, the probability distribution of an article's topic approximates the proportion of an article that

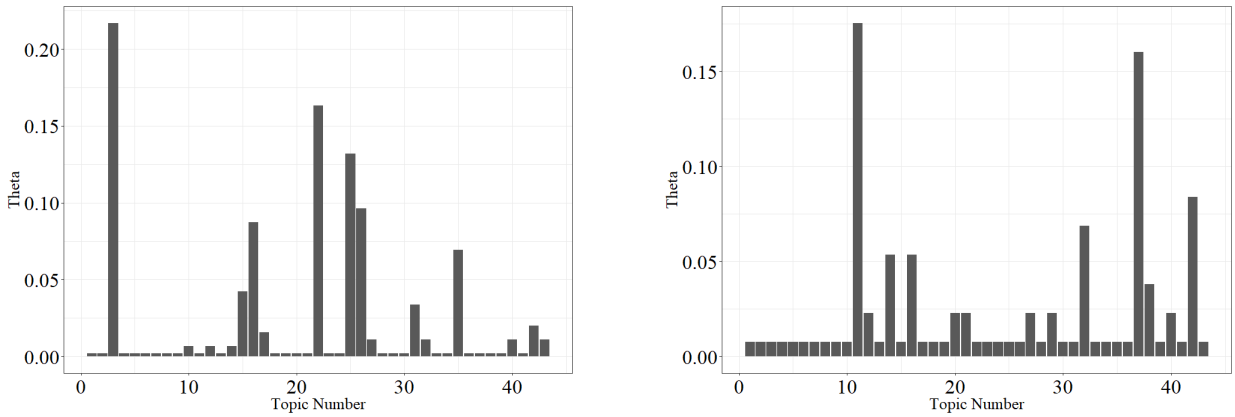
speaks about the respective topic (e.g. Suppose that the probability distribution of a topic k in article d is $\theta_{d,k}$ = 0.5. Then this is interpreted as 50% of this article covers narratives about topic k.). Second, three is chosen as a threshold based on the assumption that a reader mostly grasps and remembers at most three key narratives from reading an article. Consequently, an article will most strongly signal narratives that have the highest proportion in that article. Figure 7 provides an illustration of the selection procedure of articles. Finally, the raw number of articles identified for each component is then normalized by the total number of articles at each time step. The final outcome of these topic-based uncertainty indices can be examined in the figures in 8.[9]

The first of these topic-based uncertainty indices stands for *macroeconomic* expectation uncertainty and is in the following denominated as EFU-MACRO. It can be seen that this index clearly rises during periods of economic contractions happening in Sweden (early 1990s and 2000s, followed by 2008 and 2012) and abroad (Russian and Asian economic crisis in the late 1990s). Another more striking observation is that since 2008 the index steadily rises and keeps being elevated until the end of the analyzed period. This might indicate a structural break in macroeconomic uncertainty levels, where news coverage about this topic has been kept high after the Great Recession. The second index is related to *policy* expectation uncertainty, in the following called EFU-POL, and shows a more noisy picture. The difference to the other measures is that social policy uncertainty was extraordinary high in the mid 1990s, most likely being related to the rising Swedish unemployment rates, from which the Swedish labour market only slowly recovered. It furthermore shows peaks in the early 2000s, when the Swedish government held a referendum that determined whether Sweden will enter the Eurozone, and peaks higher during 2016 when Brexit and the election of Trump happened. Policy uncertainty also seems to show a structural break after 2008. However, compared to the first index, this break seems to have happened in the slope of this index rather than its mean. Whether news-based policy expectation uncertainty captures the same or different information from news-based macroeconomic expectation uncertainty is examined in the next section, where EFU, EFU-MACRO and EFU-POL are going to be evaluated.

[9] Although the principal components are orthogonal, the topic-based uncertainty indices might still partly correlate with each other because some articles contain topics from both principal components with a high probability. In total, 645 articles fall into this category, hence, this aspect is of minor importance. Moreover, the principal goal of this last step is to remove articles that contain topics that are *not* relevant, i.e. topics that do not drive most of the variation on the EFU articles, and are therefore regarded as noise.

Figure 7 Example of Article Selection Procedure for the Computation of EFU-MACRO.

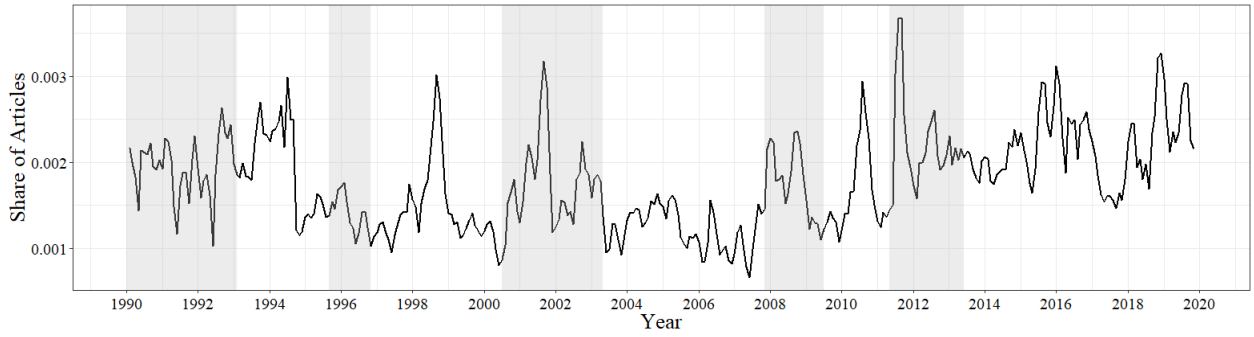
(a) Article counted. (b) Article not counted.



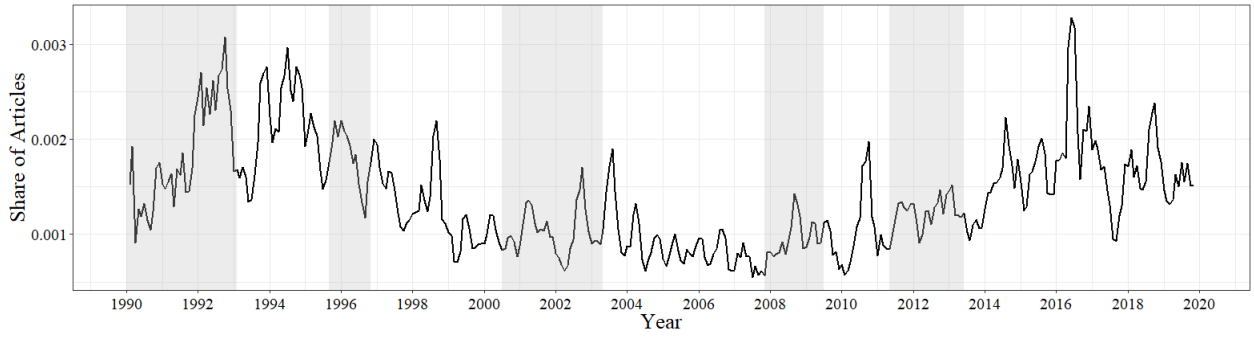
Notes: The plots present the topic probability distributions of two randomly drawn articles. The x-axes list the topics in sequential order and the y-axes illustrate how probable each topic is to appear in the respective articles ($\theta_{d,k}$). For instance, for creating EFU-MACRO we proceed in the following way: In the left graph we see that topic 3 (part of top 5 loadings of the first principal component as seen in Table 1) counts among the top three probability distributions of that random article. Hence, for EFU-MACRO this article is counted. On the right graph we see that none of the relevant topics count among the top three probability distributions of that random article. Therefore, this particular article is not counted.

Figure 8 Final Topic-Based Uncertainty Indices.

(a) Macroeconomic Uncertainty EFU-MACRO



(b) Policy Uncertainty EFU-POL



Notes: Rolling means of the scaled series with a window of 3 months are presented. Grey segments indicate recession periods in the Swedish economy as defined by the OECD. The EFU-MACRO index comprises 8,784 articles and the EFU-POL index comprises 4,997 articles.

5 Evaluation of News-Based Expectation Uncertainty

Having derived meaningful uncertainty indices from Swedish news articles in the previous section, this part seeks to evaluate if and how much these newly obtained news-based uncertainty indices impact the beliefs of economic agents, proxied by Swedish firm and consumer confidence indicators. Additionally, I also analyze the relationship between the existing Swedish Economic Policy Uncertainty (EPU) index and beliefs. This is done to compare if news-based expectation uncertainty, measured by EFU, EFU-MACRO, and EFU-POL, can give different insights compared to news-based economic policy uncertainty, measured by the EPU. This comparison lies in the focus of all subsequent analysis. To the best of my knowledge, the analysis in this section comprises at the same time the first systematic investigation into the relationship between news-based uncertainty and confidence in the literature.

Swedish confidence indicators are obtained from Sweden's National Institute of Economic Research (NIER). These are key confidence indicators which survey Swedish households and firms about their perceptions of the economy in the present and future (Konjunkturinstitutet 2020). *CI* summarizes the degree of confidence expressed by manufacturing, construction, retail and services, as well as consumers. CI_I aggregates the confidence of industries (defined as the aggregated confidence of manufacturing, construction, retail and services) and CI_C captures the confidence of consumers. CI_{micro} measures consumers' perception of their personal finances and CI_{macro} measures consumers' general view of the Swedish economy. Details about the questions, on which these indicators are based, are given in Table A2 in the Appendix. As can be seen each category covers questions about respondent's assessments of the present and future situation. Hence, they are regarded as appropriate proxies for measuring beliefs. The data spans monthly observations between 1996:07 and 2019:12 for *CI* and CI_I, while the data for CI_C, CI_{micro} and CI_{macro} spans from 1993:01 to 2019:12.

The first relevant question is whether (different types of) uncertainty in the news precede changes in confidence, thereby giving insights into its predictive potential. This timely relationship is investigated in §5.1 by analyzing the cross-correlation of news-based uncertainty and confidence, and subsequently by conducting Granger causality tests. Having this dynamic relationship in mind, the second question to answer is what effect a sudden rise in news-based uncertainty has on confidence over time, and vice versa. This question will be answered by analyzing the impulse responses of Swedish confidence to a shock in the news-based uncertainty in the framework of a structural VAR in §5.2.

5.1 Lead-versus-Lag Relationship with Swedish Confidence

Initially, I explore whether news-based uncertainty precedes changes in confidence. An initial visual inspection can be conducted with Figure 9, which plots the time series of general Swedish confidence (*CI*)[10] against the inverse values of EFU (Subfigure 9a) and the inverse values of EPU (Subfigure 9b).[11] Clearly, up until the early 2000s, news coverage about economic policy uncertainty increased after changes in confidence, while increased news coverage about economic forecast uncertainty preceded changes in confidence. In the time period afterwards

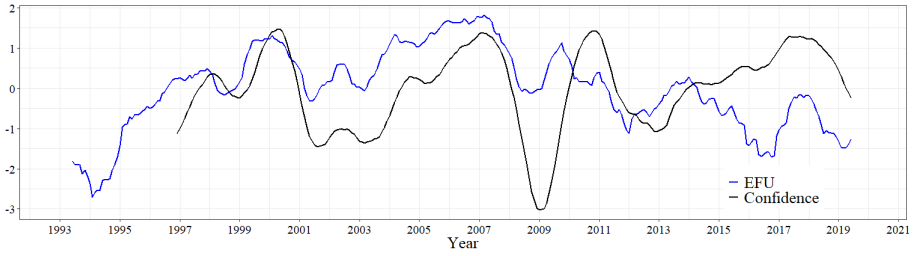
[10] Since most of the confidence indices have a similar visual pattern, only the general CI is presented here.

[11] Since uncertainty was shown to inversely co-move with consumption and investment (see §2.3) and confidence has a positive timely relationship with the latter two variables, the inverse values of uncertainty positively co-move with confidence and facilitate comparison.

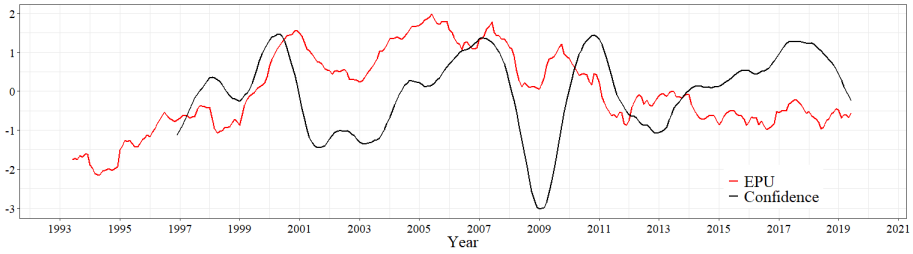
both EFU and EPU seem to show preceding movements, especially between 2009 and 2015. This relationship is investigated more systematically by examining the cross-correlation of news-based uncertainty with Swedish confidence. The cross-correlation (CCF) of two variables provides information on whether lag values of a predictor

Figure 9 Does News-based Uncertainty lead or lag Confidence? - Visual Inspection.

(a) EFU and Swedish Confidence (*CI*)



(b) EPU and Swedish Confidence (*CI*)



Notes: EFU and EPU values are plotted as 12-month rolling and inverse means. Source of the Confidence Index: Author's rendering of NIER data.

variable x at time $t-i$ correlate with 'current' values of a response variable y at time t, i.e. whether x leads y or whether the relationship is reverse. In (Enders 2014, p. 269), the cross-correlation $\rho_{yx}(i)$ between response variable y_t and predictor variable x_{t-i} is defined as follows:

$$\rho_{yx}(i) \equiv \frac{Cov(y_t, x_{t-i})}{\sigma_y, \sigma_x} = \frac{\frac{1}{n}\sum_{t=i+1}^{T}(y_t - \bar{y})(x_{t-i} - \bar{x})}{\sigma_y, \sigma_x} \tag{9}$$

where σ_y, σ_x are the time independent standard deviations of y_t and x_t, and $Cov(y_t, x_{t-i})$ is the cross-covariance between y_t and x_{t-i}. Note here that $\rho_{yx}(i) \neq \rho_{yx}(-i)$, but that $\rho_{yx}(i) = \rho_{xy}(-i)$. The figures in 10 compare the monthly cross-correlations of the dictionary indices EFU and EPU, as well as the topic-based indices EFU-MACRO and EFU-POL with Swedish confidence within a six-year horizon. The figures can be interpreted in the following way: On the x-axes, the intervals before (after) zero represent monthly lags the uncertainty index (the confidence index). The y-axes indicate the degree of cross-correlation. Hence, a negative first lag of an uncertainty index means that rising uncertainty in $t-1$ correlates with falling confidence in t by the indicated value.

Comparing the figures in 10, EFU has the strongest and most significant relationship to all types of confidence, both in lags and in magnitude. The difference is especially visible between the EFU and EPU. Lags in EFU significantly and negatively correlate with present confidence by at least 6 months, ranging from 0.25 to -0.7. In contrast, the EPU mostly cross-correlates the strongest in leads by magnitudes of 0.2 to -0.5. Among the topic-based uncertainty indices, macroeconomic uncertainty shows a similar but less consistent picture to the EFU, where it cross-correlates with general and macroeconomic confidence the strongest. Policy uncertainty shows the strongest

cross-correlation with consumer and microeconomic confidence. In terms of magnitude, topic-based uncertainties are weaker than the more general EFU.

All news indices indicate a positive cross-correlation with general and industry confidence both in their long-term lags and in their leads. The latter case implies that high levels of general and industry confidence are followed by high levels of uncertainty. This can especially hold true in cases of economic bubbles: High levels of excitement among market participants drastically fall after inconsistencies in markets are corrected, which then increases uncertainty about the expected path of the economy. The positive cross-correlation in the lags of uncertainty seem counter-intuitive but not unreasonable. According to the evidence in Bloom (2009), the economic impact of uncertainty is short-sighted and economic activity can bounce back quickly after uncertainty shocks. Hence, the fact that high uncertainty is followed by high confidence in the long-term seems to be plausible. In addition, it is worth to note that the EPU has also non-significant short-term positive lags which could potentially mean that policy uncertainty shocks in particular do not have a strong relation to confidence.

The long-term positive dynamics do not exist between uncertainty and the different consumer confidence measured ($CI_C, CI_{micro}, CI_{macro}$). Hence, uncertainty in the news might play a heterogeneous role for industry and households. All uncertainty measures precede and follow long periods of low confidence among consumers. This dynamic confirms the general intuition that private households are more negatively prone to news sentiment given that they are they source most of their information about the economy from news. Notably, here again EFU correlates stronger in lags, while EPU has the highest cross-correlation in its 6 to 12 month leads. The strongest difference between the EPU and all newly constructed indices is found for macroeconomic confidence, which corresponds to respondents' feelings in regard to the present and future economic situation. The EPU does not show any relevant cross-correlating relationship, while the other news indices lead macroeconomic confidence.

In summary, these figures show that economic expectation uncertainty in news, in particular the more general EFU, might have a predictive potential for confidence in Swedish firms and households, while news-based economic policy uncertainty is either not leading or shows a weaker relation. It is evident that the cross-correlation does not yet tell us anything about the causal links between uncertainty and confidence since it does not account for confounding factors that can cause confidence to rise or fall through time. Therefore, I formally test the lead-or-lag relationship between the news indices and the confidence indicators by conducting Granger causality tests with the procedure of Toda & Yamamoto (1995). I start with a multivariate VAR, whose framework is specified as standard in the literature:

$$X_t = A_0 + \sum_{i=1}^{p} A_i X_{t-i} + \phi_1 Z_t + \phi_2 Z_{t-1} + \varepsilon_t, \tag{10}$$

$$\text{where} \quad X \equiv \begin{bmatrix} CI_* \\ Uncertainty \end{bmatrix}, \quad Z \equiv \begin{bmatrix} \Delta U \\ CPI(\%) \\ R(\%) \\ PPI(\%) \\ Prod(\%) \\ OMX - Vol \end{bmatrix}, \quad \text{and} \quad \varepsilon_t \sim V.W.N. \tag{11}$$

X_t is a vector of endogenous variables containing one of the confidence indices $CI_* \in \{CI, CI_I, CI_C, CI_{micro}, CI_{macro}\}$ and one of the news-based uncertainty indices $Uncertainty \in \{EFU, EPU, EFU - MACRO, EFU - POL\}$. Con-

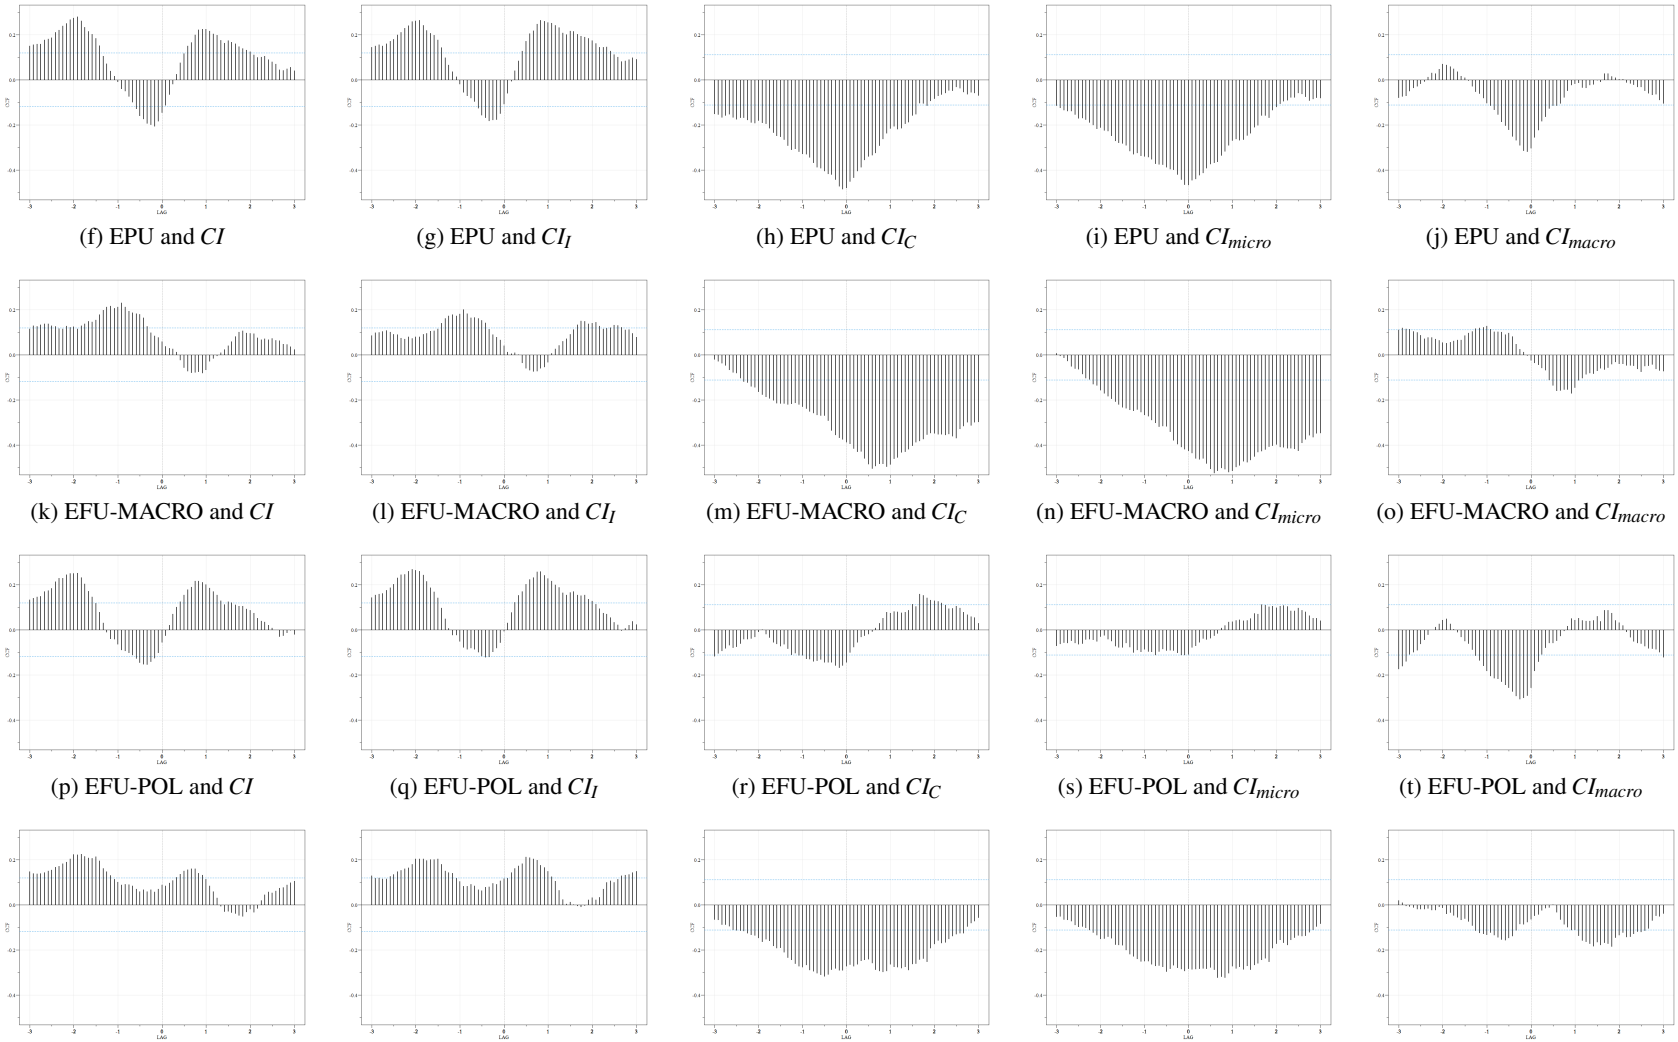
fidence and uncertainty measures are entered in levels. [12] The Z_t vector contains a set of (lagged) exogenous variables that control for other factors that might influence contemporaneous confidence: The change in the Swedish unemployment rate from t to $t-1$ ($\Delta U(\%)$) and the percent change in the Swedish consumer price index (CPI) are expected to influence household-level confidence. The Repo rate (R) aims to approximate short-term interest rates and the production growth rate ($Prod$) controls for potential shocks in production. Together with the percent change in the production price index (PPI), the three latter variables control for factors that are likely to shift firm-level confidence. Furthermore, fluctuations in the Swedish stock market are approximated by the 3-months rolling mean of the monthly standard deviation of the Swedish stock market index OMXS30 ($OMX-Vol$). All exogenous variables enter either as rates or first-differences, and are therefore stationary processes. It is expected that current confidence levels can be influenced by shocks in the previous period, hence, the first-order lags of exogenous variables are also included. It is important to note that most probably there remain other variables that influence contemporaneous levels of sentiment and news-based uncertainty, but which are hard to capture and therefore not contained in the model. These variables can be, for instance, shocks happening in foreign economies. It was also abstained from including additional lags in order to avoid low powered test statistics and potential overfitting of the VAR model, given that its order of integration is substantial (see Table C1 in the Appendix). The error term ε_t follows a white noise process and the null hypothesis of no serial correlation in the lags of the error term is formally tested with the Breusch-Godfrey test.[13] Monthly data is used, and the time spans from 1994:06 to 2019:12 for $CI_C, CI_{micro}, CI_{macro}$ and from 1996:07 to 2019:12 for CI, CI_I. The data sources are stated in Appendix A.
The VAR is estimated with a lag length of $p+d_{max}$, where d_{max} is supposed to be the maximal order of integration in the process. p was determined objectively with the Akaike Information Criterion (AIC), which is more likely to select the correct number of lags in small to medium sample sizes. The order of d_{max} was obtained with the Augmented Dickey-Fuller test. In cases where serial correlation between the error terms was found with lag length p, the lag length was increased until the null hypothesis of no serial correlation in the residuals could be rejected. Table C1 in the Appendix presents size of the estimated VARs as well as additional properties of the tests conducted. The procedure of Toda & Yamamoto (1995) tests the null hypothesis that the first p lags of the explanatory variable are jointly insignificant in the equation for the dependent variable using the standard asymptotic χ^2-statistics.

Tables 2 reports the statistics of testing Granger causality from the news-based uncertainty indices to the confidence indices, and Table 3 show the statistics of testing Granger causality in the opposite direction. Overall, there exists some evidence of Granger causality from news-based expectation uncertainty to Swedish confidence but less in the reverse direction. It is particularly macroeconomic expectations uncertainty that Granger causes all types of confidence the strongest. While EFU showed a stronger cross-correlation relationship with confidence, it significantly Granger causes only industry and consumer confidence and the EPU Granger causes consumer and micro confidence. On the other hand, there exists evidence of Granger causality from micro confidence to EPU, and from macro confidence to EFU. These findings suggest that news-based expectation uncertainty can be useful for predicting Swedish confidence, in addition to their use as a visual instrument.

[12] Obtaining unbiased estimates in a VAR usually requires all variables to follow a stationary process or otherwise to be cointegrated. By testing for unit roots with the Augmented-Dickey-Fuller test, the null hypothesis of unit roots could not be rejected for the EPU, EFU-POL, CI_C and C_{micro}. However, the Toda-Yamato procedure allows variables to enter in levels even if there is no cointegration relationship, and restrictions on the parameters can be tested because the standard asymptotic theory is still valid (Toda & Yamamoto 1995, pp. 227)

[13] This test was preferred over the more conventional Ljung-Box test because it was shown that the latter is biased in favor of maintaining the null hypothesis if not all regressors are strictly exogenous (Davies & Newbold 1979). This assumption does not hold in this framework since lagged dependent variables are used as regressors.

Figure 10 Does News-based Uncertainty lead or lag Confidence? - Cross-Correlation Functions.



(a) EFU and CI (b) EFU and CI_I (c) EFU and CI_C (d) EFU and CI_{micro} (e) EFU and CI_{macro}

(f) EPU and CI (g) EPU and CI_I (h) EPU and CI_C (i) EPU and CI_{micro} (j) EPU and CI_{macro}

(k) EFU-MACRO and CI (l) EFU-MACRO and CI_I (m) EFU-MACRO and CI_C (n) EFU-MACRO and CI_{micro} (o) EFU-MACRO and CI_{macro}

(p) EFU-POL and CI (q) EFU-POL and CI_I (r) EFU-POL and CI_C (s) EFU-POL and CI_{micro} (t) EFU-POL and CI_{macro}

Notes: The numbers in the x-axes present years and the y-axis corresponds to the cross-correlation of the news-index (varied by rows) with the confidence index (varied by columns). The blue dashed lines are the 95% lower and upper confidence intervals calculated as $-\frac{1}{n} \pm \frac{2}{\sqrt{n}}$.

Table 2 Granger Causality from News-Based Uncertainty to Confidence

Indices	CI	CI_I	CI_C	CI_{micro}	CI_{macro}
EFU	9.6	13.4	14.8	2.9	15.3
	(0.14)	(0.063*)	(0.011**)	(0.98)	(0.22)
EPU	4.9	1.5	14.6	16.0	8.1
	(0.56)	(0.98)	(0.023**)	(0.066*)	(0.23)
EFU-MACRO	13.7	17.0	21.0	16.1	35.1
	(0.018**)	(0.0093***)	(<0.01***)	(<0.01***)	(<0.01***)
EFU-POL	3.3	3.7	6.1	3.7	2.3
	(0.65)	(0.93)	(0.29)	(0.6)	(0.89)

Notes: χ^2 values are presented, with their respective p-values in brackets. *p<0.1, **p<0.05, ***p<0.01.

Table 3 Granger Causality from Confidence to News-Based Uncertainty

Indices	CI	CI_I	CI_C	CI_{micro}	CI_{macro}
EFU	6.4	5.8	5.2	13.7	29.8
	(0.38)	(0.56)	(0.4)	(0.19)	(0.003***)
EPU	8.4	11.6	3.8	14.9	4.6
	(0.21)	(0.11)	(0.7)	(0.093*)	(0.6)
EFU-MACRO	4.5	3.3	6.2	5.6	4.3
	(0.48)	(0.77)	(0.046**)	(0.13)	(0.23)
EFU-POL	8.1	18.0	6.0	8.5	11.4
	(0.15)	(0.036**)	(0.3)	(0.13)	(0.076*)

Notes: χ^2 values are presented, with their respective p-values in brackets. *p<0.1, **p<0.05, ***p<0.01.

5.2 Shock Response of Swedish Confidence

As could be seen in the plots of the new uncertainty indices, uncertainty-related events typically cause a sudden spike in news-based uncertainty. These sudden spikes represent *shocks* and it is useful to examine what impact these sudden shocks in news-based uncertainty have on Swedish confidence over time. This is done by analyzing the impulse response functions of the confidence indicators. Impulse response functions reveal the dynamic response of confidence in the occurrence of an identified shock (i.e. the sudden increase of news coverage about uncertain expectations). The identification of a news-based uncertainty shock relies on existing literature where it is identified based on a Cholesky decomposition in a structural VAR framework, where uncertainty is ordered first. The initially specified VAR model in §5.1 is therefore turned into a structural VAR model in order to be able to identify orthogonal shocks.

The figures in 11 plot the impulse responses of the respective confidence indices to a one standard deviation shock in the four different uncertainty indices. Overall, the figures show that it is especially economic and macroeconomic expectation uncertainty that affects confidence. Economic policy and policy expectation uncertainty both have a

non-significant or no impact, and the responses are most of the times delayed, consistent with the cross-correlation functions in Figure 10. EFU and EFU-MACRO consistently cause a long-term drop in Swedish confidence measures. The shocks are both statistically and economically significant, causing a drop in confidence of up to 0.25 standard deviations from which they recover only after up to 24 months. The response of general confidence is delayed by one period, while the response of the rest of the confidence indicators is contemporaneous. The negative impact of the EFU and EFU-MACRO seems to be especially persistent for consumer confidence, spanning over the 3-year time frame analyzed.

One concern can be that the articles in the EFU index do not convey only uncertainty but other, 'first-moment' sentiments, such as pessimism or 'bad news' in general. One way this is mitigated is by having the exogenous variable $OMX-Vol$ in the VAR. Stock markets are forward looking and volatility in stock market prices already incorporate the reactions of market participants to bad news about the future economy. Nonetheless, articles might include additional first-moment narratives that are not elsewhere captured. Therefore, the robustness of the uncertainty shocks is furthermore tested by controlling for shocks in news about (general) economic expectations. This is done by including the EF, created in section 4, as an additional endogenous variable into the sVAR. [14] EF is ordered first, one of either EFU, EFU-MACRO, EFU-POL second, followed by the confidence indices. The ordering is based on the assumption that news about general economic expectations will increase before news about uncertainty in economic expectations does. The robust shock response of the confidence indices can be examined in 12. The response to macroeconomic expectation uncertainty shocks remain robust, both economically and statistically. For EFU, the magnitude of the shocks prevails, but the responses are only significant for consumer and macroeconomic confidence.

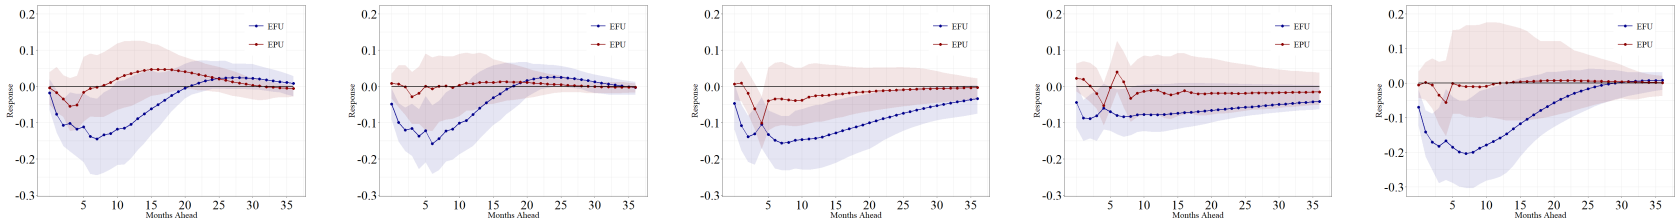
In Appendix C, I furthermore report the robustness of the impact of the EFU by ordering the confidence indices first in the Cholesky decomposition. The initial results are robust to this modification. EFU consistently causes a negative and significant drop in confidence, although its magnitude gets slightly reduced. The appendix furthermore illustrates the impulse response of EFU and EPU to a shock in confidence. Although confidence shocks (meaning an increase in confidence) can decrease news-based uncertainty by 0.1 standard deviations, this effect is not or only marginally significant for both EFU and EPU. Surprisingly, the EPU index even (non significantly) overshoots in the short-term after a positive confidence shock.

[14] It counts the relative frequency of articles that mention keywords related to the economy and expectations, but does *not* additionally filter by uncertainty keywords.

Figure 11 VAR Estimation of the Impact of a News-based Uncertainty Shock on Confidence.

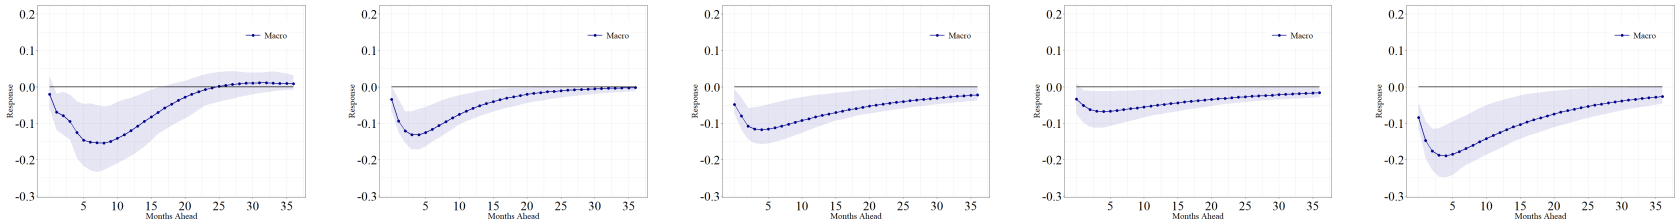
(a) **EFU versus EPU**

(b) *CI* (c) CI_I (d) CI_C (e) CI_{micro} (f) CI_{macro}

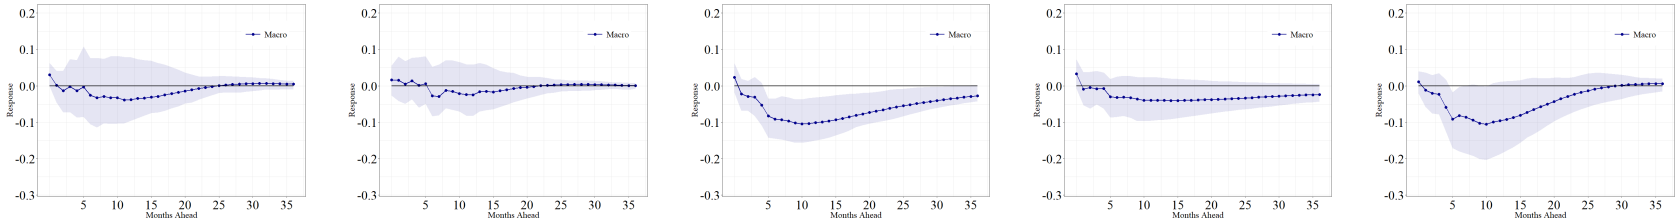


(g) **Topic-Based Uncertainty**

(h) EFU-MACRO and *CI* (i) EFU-MACRO and CI_I (j) EFU-MACRO and CI_C (k) EFU-MACRO and CI_{micro} (l) EFU-MACRO and CI_{macro}



(m) EFU-POL and *CI* (n) EFU-POL and CI_I (o) EFU-POL and CI_C (p) EFU-POL and CI_{micro} (q) EFU-POL and CI_{macro}

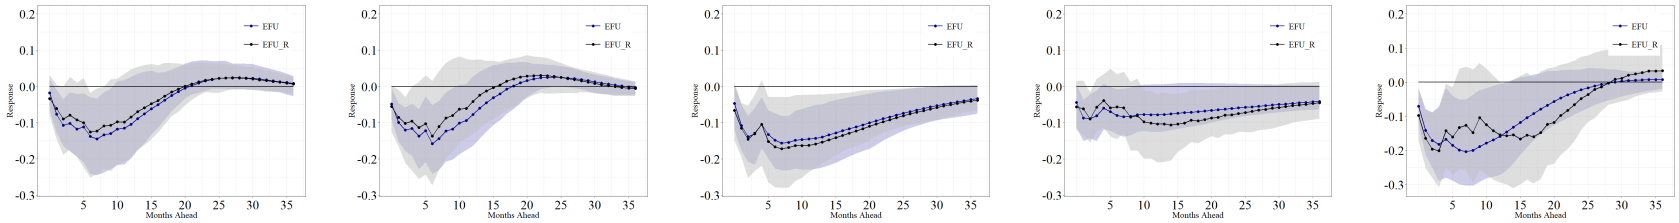


Notes: Shaded areas represent 95% standard residual-based confidence bounds.

Figure 12 Robust VAR Estimation of the Impact of a News-based Uncertainty Shock on Confidence of each shock response.

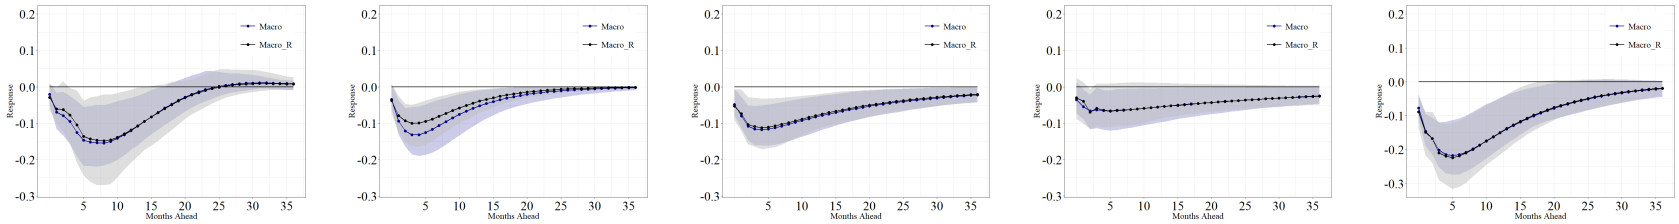
(a) **EFU**

(b) EFU and CI (c) EFU and CI_I (d) EFU and CI_C (e) EFU and CI_{micro} (f) EFU and CI_{macro}

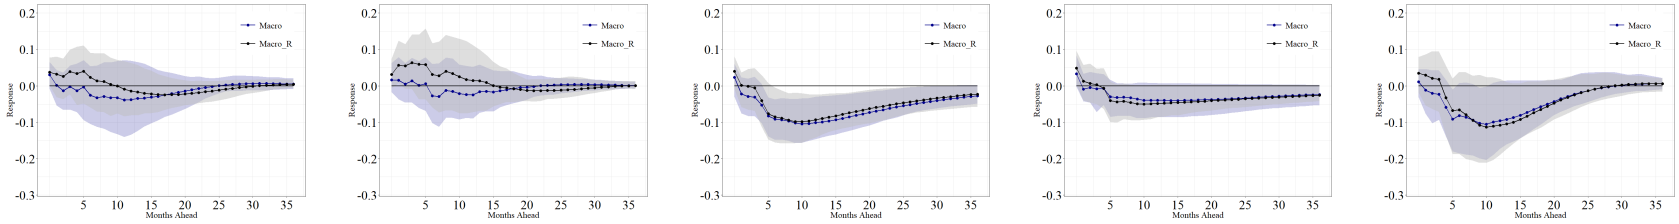


(g) **Topic-Based Uncertainty**

(h) EFU-MACRO and CI (i) EFU-MACRO and CI_I (j) EFU-MACRO and CI_C (k) EFU-MACRO and CI_{micro} (l) EFU-MACRO and CI_{macro}



(m) EFU-POL and CI (n) EFU-POL and CI_I (o) EFU-POL and CI_C (p) EFU-POL and CI_{micro} (q) EFU-POL and CI_{macro}



Notes: Shaded areas represent 95% standard residual-based confidence bounds of each shock response.

6 Discussion

This section summarizes the main findings of measuring and evaluating news-based expectation uncertainty. §6.1 sheds light on the validity and limitations of both the proposed measurement method as well as the evaluation. In addition, §6.2 discusses avenues to use the newly derived uncertainty indices for future research and discusses possibilities to further extend natural language processing methods for measuring uncertainty in the news.

I create three different news-based uncertainty indices, where the first is the more general Economic Forecast Uncertainty (EFU) index, which represents the relative count of articles that convey uncertain expectations about the economy. The other two indices capture explicit narratives that are the most relevant for conveying expectation uncertainty. These narratives stem from latent topics in the articles about EFU which are identified with Latent Dirichlet Allocation (LDA) and Principal Component Analysis (PCA). I show that it is narratives about *macroeconomic* (EFU-MACRO index) and *policy* (EFU-POL index) expectation uncertainty that drive most of the variation in the articles about expectation uncertainty over time.

I evaluate the newly obtained indices by analyzing and comparing their impact on economic fundamentals and Swedish confidence to the currently available news-based uncertainty index, namely the Swedish Economic Policy Uncertainty (EPU) index. In order to analyze their effect on economic fundamentals, I use the benchmark multivariate VAR model of Armelius et al. (2017), the creators of the Swedish EPU. Based on this model, news-based expectation uncertainty and economic policy uncertainty have a similar impact on the aggregate economy. They cause a drop in GDP growth by up to 0.4 standard deviations which is significant at the 1%-level. While their impact on other macroeconomic variables is also negative, they are only marginally significant in statistic terms. The results of this analysis are reported in Appendix C.

The impact of news-based uncertain expectations on confidence highly differs from the impact of news-based economic policy uncertainty, despite having the same impact on economic fundamentals. First of all, by analyzing their cross-correlations, I show that all newly obtained indices lead Swedish confidence measures by a at least 6 months in advance, meaning that rising expectation uncertainty expressed in news 6 months ago correlates with falling firm and household confidence at present. In addition, I use a bivariate VAR framework where I account for exogenous shocks in other variables and find evidence that expectation uncertainty shows a stronger predictive potential for Swedish confidence. The EFU-MACRO index, obtained with the unsupervised method, Granger causes all types of confidence with a high statistical significance. On the other hand, the EPU can Granger cause consumer confidence.
Furthermore, impulse response analysis shows that it is shocks in news-based economic and macroeconomic expectation uncertainty that cause Swedish firm and consumer confidence to significantly fall in the long-term, from which they only slowly recover. News shocks in economic policy uncertainty does not affect the confidence of Swedish firms and households - the same being true for policy expectation uncertainty. Here again, shocks in macroeconomic expectation uncertainty, derived with the unsupervised method, show the strongest robustness to modifications in the VAR. The responses of consumer and macroeconomic confidence to news shocks in the EFU and EFU-MACRO are statistically and economically robust even after accounting for preceding shocks in news about general economic expectations. The consumer and macroeconomic confidence indicators capture responses

of households and it is plausible that they show the strongest response given that households typically use news as their main information source to update their beliefs and expectations about the economy, while firms use additional and potentially more important information sources, such as experts within their industries and their own experience.

There are two important learnings which can be drawn from these finding for the Swedish case. First, the results show that expectation uncertainty causes first-moment changes in the beliefs of Swedish consumers, opposed to the theory that uncertainty only causes changes in the second-moment (Bloom 2009). Particularly noteworthy is the impact on the macroeconomic confidence indicator which captures households' expectations about the future of the Swedish economy. This provides evidence to the argument of Shiller (2017) in that people use narratives, i.e. topics in the news, to update expectations, and that in this case narratives related to uncertainty deteriorate their expectations on the economic future. On the other hand, given that the impact of expectation uncertainty is less robust for the other confidence indicators, which focus more on the beliefs of firms and households at present, suggests that other first-moment sentiments in the news, such as pessimism or 'bad news' in general, impact the beliefs of firms and households regarding the current economic situation more strongly.
Second, it is shown that that policy-related uncertainty does not affect the confidence of Swedish firms and households, despite the evidence that news about expectation uncertainty and economic policy uncertainty both have the same impact on the aggregate Swedish economy. This suggests that each type of uncertainty plays a different role in news. Expectation uncertainty can affect economic activity through the channel of first-moment shocks in confidence. It therefore serves as a framing mechanism that changes readers' beliefs and expectations about the economy. Consequently, this change in beliefs also alters their behaviour. [15] In contrast, Swedish news about economic policy uncertainty rather serves as a *mirror* of the contemporaneous level of economic policy uncertainty present in the Swedish economy. It most likely affects economic activity through other channels which can be, for instance, resource inefficiencies in the public sector or the confidence of financial market participants. An alternative explanation can also be that news-based policy uncertainty might purely function as a second moment shock that happens only after shocks in confidence occur.

The fact that policy uncertainty and also policy expectation uncertainty does not affect Swedish firm and household confidence does not come from far away. As pointed out by Istiak & Serletis (2018), the impact of policy uncertainty is heterogeneous among different countries, where countries with stable regimes and trusted institutions will most likely experience less impact of policy uncertainty even in case of negative shocks. According to the Worldwide Governance Indicator by Kaufmann et al. (1999), Sweden counts among those countries with the best performing governance. [16] Therefore, it is no surprise to see that firms and households are less touched by news covering stories about policy uncertainty. Having that said, future research will certainly bear fruits in investigating additional channels through which news-based uncertainty affects economic activity in order to understand how uncertainty affects people's behaviour and eventually economic activity. Moreover, it is worthwhile to explore whether the same discrepancy between expectation uncertainty and policy uncertainty exists in other countries.

[15] The transmission channel from changes in confidence to economic fluctuations was not investigated in this thesis because it would require an additional research framework, which would have bursted the scope of this thesis. However, in previous studies it was shown that changes in consumer confidence are predictors of consumption spending in industrial economies (Dees & Brinca 2013, Ludvigson 2004).

[16] Measured by the following dimensions: Voice and Accountability, Political Stability and Absence of Violence, Government Effectiveness, Regulatory Quality, Rule of Law, Control of Corruption.

6.1 Validity and Limitations of Results

Validity of the Measurement

Common to all text-based indicators, the most consequential limitation of the newly created Economic Forecast Uncertainty (EFU) index and the topic-based indices are the large number of researcher degrees of freedom in the measurement design. Human language is highly flexible which leaves a lot of room for both the quantification and interpretation of meaning. This ranges from selecting appropriate keywords in dictionary methods up to selecting meaningful hyperparameters in unsupervised methods like in Latent Dirichlet Allocation. Only small changes can already detrimentally change the properties of the outcome, as was shown by Eckley (2015) who was able to increase the signal strength of the UK version of the news-based Economic Policy Uncertainty (EPU) index by 14% by only adding the plural form *uncertainties* to the dictionary. In this thesis, the aim was to base all supervised decisions on objective assumptions. These were explicitly mentioned throughout the process and serve to retrace all steps undertaken and facilitate the comparison of the proposed methods with methods of other authors.

The specific set and number of keywords chosen to initially filter articles about economic expectations might not be completely exhaustive, which might have resulted in missing some articles than convey uncertain expectations (false negatives). At the same time, the sampled corpus might entail articles that do not explicitly talk about uncertain expectations even though they mention the filtered keywords (false positives). Nonetheless, several factors provide evidence that these limitations do not pose severe restrictions on the internal validity of my results. First, the analysis of the principal components of the narratives of the EFU articles showed that topics that drive most of the timely variation in the topic prevalences are of either economic or political nature and therefore thematically relevant for economic expectations, therefore refuting the importance of false positives. It is true that there might still remain topics with a high prevalence, such as topics about sports or culture, that most likely do not carry high importance. However, it is likely that many of those 'exotic' articles still carry an uncertainty load, for example by covering narratives about sports clubs facing financial difficulties during recessions. The copyright restrictions of Swedish news articles unfortunately do not allow to inspect the encrypted articles, hence, this hypothesis remains an assumption.

In regard to false negatives, I do not account for articles that convey uncertainty without explicitly mentioning Swedish uncertainty keywords. While it is certainly worthwhile to explore more nuanced and less explicit uncertainty-related keywords for classifying uncertainty in articles, this would require careful semantic analysis. An earlier version of this thesis intended to create an uncertainty index by using keywords related to *modality*. However, the signal-to-noise ration of that index was particularly low, which is why it was abstained to use modal words for the uncertainty classification of EF articles. Thus, using explicit uncertainty keywords was the best option for the data available for this thesis. Moreover, Kalamara et al. (2020)'s comparative analysis on the economic impact of different UK-based uncertainty indices showed that dictionary-based uncertainty indices that account for more general uncertainty terms do not yield better results compared to the more popular approach followed here, which suggests that most of the variation in these indicators comes from explicit uncertainty terms.

Another potentially important concern is related to the timely relevance of the keywords used for filtering. Human

spoken and written language continuously transforms and adapts itself to changing environments and it is likely that the importance of words varies over time. One could argue that the word uncertainty has been increasingly used only after the Great Recession in 2008, thereby having become a *vogue term*, and that there might be a present bias with respect to the keywords used for filtering articles. While it could be the case that newspapers in the early 1990s used different vocabulary, it is most likely not of big concern given that the relative article count mentioning uncertainty in the 1990s is almost at the same level as it has been in the last 2010s, if not higher. Hence, the keywords are seen as relevant for the complete time period.

One could in addition argue that news about general economic conditions can also alter people's expectations about the future. This concern was inspected by evaluating an index that counted the share of general economic news over time (i.e. the relative count of articles initially filtered by the economy-related keywords). Despite yielding a very similar visual picture to the articles about general economic expectations (EF articles), the frequency of general economic news did not show any statistically significant impact on neither economic fundamentals nor economic confidence, which is why it was not included in the analysis. I leave the inspection of the role of general Swedish economic news, augmented with refined filtering of articles, for future research.

One last aspect to consider is the usability of the topic-based uncertainty indices and the added-value of analyzing the semantic context of articles with LDA. As can be seen from the process in §4, identifying narratives with LDA can be a cumbersome process which requires careful calibration and multiple steps. The topic-based uncertainty indices added especially strong value in the robustness of the uncertainty shocks. Firm and consumer confidence consistently deteriorated in cases of elevated news-based macroeconomic uncertainty. However, it remains questionable whether the improvement in the uncertainty index obtained with LDA is worth the extra steps. Based on the findings of this thesis, the supervised dictionary method is almost as effective in capturing the impact of news-based uncertainty as is the unsupervised semantic method.

Validity of the Evaluation

The evaluations pursued in this thesis are based on benchmark models from existing literature about economic uncertainty and news shocks, namely from Armelius et al. (2017) and Doms & Morin (2004). One drawback of the models in this thesis might be that the time frame analyzed is relatively small thereby reducing the statistical power of the test results. Several additional aspects could have been included into the VAR specifications, such as adding further lags for the control variables or additional controls. In the latter case, the most relevant control excluded is the aspect of the zero lower bound (ZLB) in US and Eurozone interest rates that was introduced by the FED and the ECB shorty after the Great Recession in 2008. It might be the case that the ZLB had an economic impact on Sweden as well, being a small and open economy close to the Eurozone, and potentially also on Swedish confidence, thereby influencing the impact of news-based uncertainty on confidence in the VAR model. However, some voices in the literature argue that the interpretation of the effects of the ZLB in a VAR is questionable, given that has been a highly non-linear event (Fernández-Villaverde et al. 2015). Furthermore, adding the ZLB would have meant to trim the time span to until 2008 for re-estimation, which would have resulted in too little observations in order to fit the higher-order VARs. Therefore, it was abstained to further modify the specifications with the aspect of ZLB.

Another aspect of doubt is the endogeneity of uncertainty for economic activity, which was imposed by the Cholesky decomposition. It is plausible to believe that uncertainty rises during recessions and drops during periods of an economic boom and this countercyclical behaviour was shown in several studies, such as Bloom et al. (2018), Jurado et al. (2015), Kehrig (2015). However, it might not be clear whether economic shocks caused uncertainty or the reverse. Baker & Bloom (2013) have tried to overcome the issue of endogeneity by including only exogenous disasters, such as natural disasters, terrorist attacks, or political protests, in their cross-country analysis of the economic impact of uncertainty. By estimating the influence of stock market volatility on GDP growth, they found that second moment shocks can explain at least half of the variation in economic growth after this selection of uncertain events. This approach was not adapted in this application given that it is highly questionable that there happened enough disasters in Sweden. Therefore, causal inferences about the impact of uncertainty are especially problematic in the context of policy and should be treated with high caution. Either way, the purpose of creating this index is not to find causality but to provide a tool for policy making in uncertain periods. This can be either through visual inspections merely or through including the index in a VAR model to estimate its economic impact, where the question of causality is secondary.

6.2 Perspectives for Future Research

Based on the evaluation results, the indices obtained in this thesis open new and attractive opportunities to further investigate the economic role of news-based uncertain expectations both in the Swedish and general context. In the US context, the EPU has been successfully used to predict market returns (Brogaard & Detzel 2015, Sum 2012, Ma et al. 2018), cross-country trade and investment (Tam 2018), and election uncertainty (Goodell et al. 2020). There is no doubt that all of these factors are detrimentally influenced by confidence or 'animal spirits'. Whether it is the bearishness or bullishness of financial markets, the confidence of a company owner about starting business relationships abroad, or the confidence of a citizen to vote for her favourite politician. Conversely, since news-based uncertainty was shown to impact confidence detrimentally stronger than news-based economic policy uncertainty, future studies should investigate the influence of the former in all these applications in further detail in order to understand whether it is really economic policy or rather expectation uncertainty that is driving people's beliefs and hence their behaviour.

The transmission of uncertainty through news can be extremely subtle, which makes it especially challenging to capture for economic evaluation. Based on the learnings derived from the methodology pursued in this thesis, however, I identify a number of promising possibilities to further extend upon the identification of news-based uncertainty both for the general and Swedish context. First of all, uncertainty classification should be attempted with less explicit keywords, which can equally transmit feelings of uncertainty or ambiguity through news. One option would be to better understand the role of *modality words* in economic news. Modality is a cross-language grammatical category that refers to the type of linguistic cues that authors use in order to convey beliefs, attitudes, and judgements. With this concept it is therefore possible to target the ambiguity in expectations about economic outcomes. The detection of uncertainty with modal expressions has already been well-researched in the field of natural language processing (Farkas et al. 2010, Auger & Roy 2008). Applications have been particularly prevalent in the biomedical sciences (Medlock 2008, Light et al. 2004) because the degree of uncertainty expressed in

scientific papers is regarded as extremely important for the reliability of scientific findings. (Light et al. 2004, p. 1) writes that *"Scientists are often not completely certain of a conclusion. This lack of definite belief is often reflected in the way [they] discuss their work"*. It is likely that this holds also true for journalists drawing conclusions about the expected path of the economy. However, the identification of modality-related uncertainty would benefit from the syntax and paragraph structure of articles, instead of using the limiting bag-of-words representation which is the only available format for this thesis. Tobback et al. (2018) has initiated the first attempt to use the count of modality expressions in articles about economic policy. Their method did not yield better forecasting performance than the current EPU-index, potentially because the exact position of modality words in the text is not accounted for in the author's dictionary method. Instead of using mere counts, it would be more insightful to detect the proximity between modality and economy related keywords (see also the suggestions for advanced sentiment classification by Pang et al. 2002). A very promising method in this field are the so called *continuous bag-of-words* or vector space models, especially the word embedding model introduced in Mikolov et al. (2013)'s seminal paper. In these models documents are represented as vectors and words can be analyzed based on their proximity to other word in the document and in the the entire document corpus. In such models, modality words would be given different weights, depending on their proximity to other words (e.g. 'probably' is given more weight if it precedes 'recession' rather than 'recovery').

For the Swedish context, it is worth exploring which news paper outlets transmit how much and what type of uncertainty. This thesis used a corpus with articles from 5 different national news sources, but the articles were not identified by the publishers. Hence, future research should account for the heterogeneity of different outlets with respect to their use in language, their narratives, and their economic impact. For examples, news-based uncertainty in boulevard journals that focus more strongly on lifestyle news might be economically less relevant than uncertainty conveyed in business and financial journals.

7 Conclusion

The agenda of policy makers and market participants is mostly concerned with current and future aggregate economic conditions. Uncertainty causes them a great deal of worry given that it leads to reduced consumption and investments which dampens economic growth. This thesis started with an overview about the measurement and impact of uncertainty, and shed light on the difficulty of capturing uncertainty due to the fact that it is not directly observable. News articles offer an attractive new source for measurement given their high-frequency and large-scale nature. In this thesis, three novel news-based uncertainty indices for the Swedish economy were developed that seek to augment the currently available Swedish news-based Economic Policy Uncertainty (EPU) index. I undertook modifications to its underlying subject of interest (focus on expectation uncertainty) and to its methodology (using unsupervised machine learning). The first index – the Economic Forecast Uncertainty (EFU) index – follows the dictionary method of the EPU. The other two indices were developed by identifying latent news topics relevant for expectation uncertainty – namely macroeconomic narratives (EFU-MACRO index) and policy narratives (EFU-POL index). Improvement was evaluated by analyzing first whether the new indices have a stronger influence on economic fundamentals. Second, their impact on beliefs, proxied by Swedish confidence indicators, was compared to the impact of the EPU. The second step furthermore initiated the first systematic, empirical investigation in the literature into the relation between news-based uncertainty and confidence, the latter being one of the major channels through which uncertainty – and news narratives – can affect the economy.

It was shown that both news-based economic policy uncertainty and economic expectation uncertainty have the same impact on the aggregate Swedish economy. However, the new indices added value to the existing Swedish EPU by showing that they have a predictive potential for Swedish confidence. First, the subjects of economic and macroeconomic expectation uncertainty in the news (EFU and EFU-MACRO) precede falling confidence in the long term. Second, shocks in these type of news cause a fall in confidence from which it only slowly recovers. The response is particularly strong for the confidence of consumers and their expectations about the economic future. Contrary to the theory in Haddow (2013), it was shown that news-based uncertainty can have a first-moment impact on beliefs. These empirical findings suggest that one of the channels through which uncertain economic expectations conveyed in news affect real activity can be the change in beliefs of economic agents. News about economic policy uncertainty and policy expectation uncertainty, however, do not influence beliefs. This might be due to the fact that countries like Sweden enjoy transparent governments. Hence, economic agents have trust in these institutions and policy uncertainty does not affect their confidence about the present and future. In contrast, the future outlook of the economy does not entirely lie in the hands of institutions but is obviously driven by uncontrollable factors or *animal spirits*. Therefore, news that convey uncertainty about the future economy will be more influential in shaping people's beliefs of the economy, and thereby also their behaviour.

Obviously, policy makers should not influence narratives propagated in the news for policy purposes. But the indices developed in this thesis can provide them useful new insights into how economic agents form their expectations and how these expectations can cause fluctuations in aggregate economic outcomes. By taking into account the type and amount of news about expectation uncertainty circulating in the economy, policy makers can improve their predictions about possible economic responses to their policy decisions and thereby enhance the calibration and effectiveness of their instruments during periods of high uncertainty.

References

Abel, J., Rich, R., Song, J. & Tracy, J. (2016), 'The Measurement and Behavior of Uncertainty: Evidence from the ECB Survey of Professional Forecasters', *Journal of Applied Econometrics* **31**(3), 533–550.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1002/jae.2430* 6

Alessandri, P. & Bottero, M. (2017), 'Bank Lending in Uncertain Times', *European Economic Review* **128**(C).
URL: *https://www.sciencedirect.com/science/article/abs/pii/S0014292120301343* 1, 8

Alessandri, P. & Mumtaz, H. (2019), 'Financial Regimes and Uncertainty Shocks', *Journal of Monetary Economics* **101**(C), 31–46.
URL: *https://econpapers.repec.org/paper/qmwqmwecw/wp729.htm* 8

Alexopoulos, M. & Cohen, J. (2009), 'Uncertain Times, Uncertain Measures', *University of Toronto, Department of Economics* .
URL: *https://econpapers.repec.org/paper/tortecipa/tecipa-352.htm* 6

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2003), 'Modeling and Forecasting Realized Volatility', *Econometrica* **71**(2), 579–625.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1111/1468-0262.00418* 5

Ankargren, S. (2020), 'Time-Varying Macroeconomic Forecast Uncertainty: Sweden and the Covid-19 Pandemic', *Swedish National Institute of Economic Research Working Paper* (153).
URL: *https://www.konj.se/english/publications/working-papers/working-paper/2020-09-16-time-varying-macroeconomic-forecast-uncertainty-sweden-and-the-covid-19-pandemic.html* 6

Armelius, H., Bertsch, C., Hull, I. & Zhang, X. (2019), 'Spread the Word: International Spillovers from Central Bank Communication', *Sveriges Riksbank Working Paper Series* (357).
URL: *https://www.bis.org/publ/work824.htm* 8

Armelius, H., Hull, I. & Stenbacka Köhler, H. (2017), 'The Timing of Uncertainty Shocks in a Small Open Economy', *Economics Letters* **155**, 31–34.
URL: *http://www.sciencedirect.com/science/article/pii/S016517651730109X* 11, 12, 33, 36, 61

Arun, R., Suresh, V., Veni Madhavan, C. E. & Narasimha Murthy, M. N. (2010), On Finding the Natural Number of Topics with Latent Dirichlet Allocation: Some Observations, *in* M. J. Zaki, J. X. Yu, B. Ravindran & V. Pudi, eds, 'Advances in Knowledge Discovery and Data Mining', Springer, Berlin, Heidelberg, pp. 391–402. 56

Auger, A. & Roy, J. (2008), Expression of Uncertainty in Linguistic Data, *in* '2008 11th International Conference on Information Fusion', pp. 1–8. 37

Azqueta-Gavaldon, A. (2017), 'Developing News-based Economic Policy Uncertainty Index with Unsupervised Machine Learning', *Economics Letters* **158**(C), 47–50.
URL: *https://econpapers.repec.org/article/eeeecolet/v3a158y3a2017i3ac3ap3a47-50.htm* 6, 7, 56

Baker, S. R. & Bloom, N. (2013), 'Does Uncertainty Reduce Growth? Using Disasters as Natural Experiments', *NBER Working Paper* (19475).
URL: *https://ideas.repec.org/p/nbr/nberwo/19475.html* 37

Baker, S. R., Bloom, N. & Davis, S. J. (2016), 'Measuring Economic Policy Uncertainty', *The Quarterly Journal of Economics* **131**(4), 1593–1636.
URL: *https://doi.org/10.1093/qje/qjw024* 1, 6, 8, 12, 13

Baker, S. R. & Fradkin, A. (2011), 'What Drives Job Search? Evidence from Google Search Data', *Stanford Institute for Economic Policy Research* .
URL: *https://ideas.repec.org/p/sip/dpaper/10-020.html* 9

Blei, D. M. (2012), 'Probabilistic Topic Models', *Communications of the ACM* **55**(4), 77.
URL: *http://dl.acm.org/citation.cfm?doid=2133806.2133826* 15, 16, 53

Blei, D., Ng, A. & Jordan, M. (2003), 'Latent Dirichlet Allocation', *Journal of Machine Learning Research* **3**, 993–1022. 15, 16, 17, 53, 54, 56

Bloom, N. (2009), 'The Impact of Uncertainty Shocks', *Econometrica* **77**(3), 623–685.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA6248* 1, 4, 5, 8, 26, 34

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I. & Terry, S. (2018), 'Really Uncertain Business Cycles', *Econometrica* **86**(3), 1031–1065.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA10927* 8, 37

Brogaard, J. & Detzel, A. (2015), 'The Asset-Pricing Implications of Government Economic Policy Uncertainty', *Management Science* **61**(1), 3–18.
URL: *https://pubsonline.informs.org/doi/10.1287/mnsc.2014.2044* 37

Bronk, R. (2011), Epistemological Difficulties with Neoclassical Economics, *in* 'Southern Economic Association 2011, 20th November 2011', LSE Research Online.
URL: *http://www.etnpconferences.net/sea/seaarchive/sea2011/index.php* 3

Caballero, R. J. (1990), 'Consumption Puzzles and Precautionary Savings', *Journal of Monetary Economics* **25**(1), 113–136.
URL: *http://www.sciencedirect.com/science/article/pii/0304393290900489* 7

Caldara, D., Fuentes-Albero, C., Gilchrist, S. & Zakrajšek, E. (2016), 'The Macroeconomic Impact of Financial and Uncertainty Shocks', *European Economic Review* **88**(C), 185–207.
URL: *https://ideas.repec.org/a/eee/eecrev/v88y2016icp185-207.html* 8

Cao, J., Xia, T., Li, J., Zhang, Y. & Tang, S. (2009), 'A density-based Method for Adaptive LDA Model Selection', *Neurocomputing* **72**(7-9), 1775–1781.
URL: *https://linkinghub.elsevier.com/retrieve/pii/S092523120800372X* 56

Carroll, C. D. (1997), 'Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis', *The Quarterly Journal of Economics* **112**(1), 1–55.
URL: *https://ideas.repec.org/p/nbr/nberwo/5788.html* 7

Castelnuovo, E., Lim, G. & Pellegrino, G. (2017), 'A Short Review of the Recent Literature on Uncertainty', *Australian Economic Review* **50**(1), 68–78.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-8462.12210* 7

Castelnuovo, E. & Tran, T. D. (2017), 'Google It Up! A Google Trends-based Uncertainty index for the United States and Australia', *Economics Letters* **161**, 149–153.
URL: *http://www.sciencedirect.com/science/article/pii/S0165176517303993* 7

C.B.O. Exchange (2003), 'Vix White Paper'. Retrieved: 01-10-2020.
URL: *http://www. cboe. com/micro/vix/vixwhite. pdf.* 5

Charles, A., Darné, O. & Tripier, F. (2018), 'Uncertainty and the Macroeconomy: Evidence from an Uncertainty Composite Indicator', *Applied Economics* **50**(10), 1093–1107.
URL: *http://www.tandfonline.com/doi/abs/10.1080/00036846.2017.1349294* 7

Davies, N. & Newbold, P. (1979), 'Some Power Studies of a Portmanteau Test of Time Series Model Specification', *Biometrika* **66**(1), 153–155.
URL: *https://academic.oup.com/biomet/article/66/1/153/223702* 27

Davis, S. (2016), 'An Index of Global Economic Policy Uncertainty', *NBER Working Paper* (22740).
URL: *https://econpapers.repec.org/paper/nbrnberwo/22740.htm* 8

Dees, S. & Brinca, P. S. (2013), 'Consumer Confidence as a Predictor of Consumption Spending: Evidence for the United States and the Euro Area', *International Economics* **134**, 1–14.
URL: *https://www.sciencedirect.com/science/article/pii/S2110701713000085?casa$_t$oken = nLoyNO − NZ9EAAAAA : X6JOFTYC$_s$TDLnybMJPBJzzLEpOL5nIyLeYnmAAWW2J9KHwvEMqqz5tslizEfuixShQq4jOibfA*34

Denny, M. J. & Spirling, A. (2018), 'Text Preprocessing For Unsupervised Learning: Why It Matters, When It Misleads, And What To Do About It', *Political Analysis* **26**(2), 168–189.
URL: *https://www.cambridge.org* 52

Deveaud, R., SanJuan, E. & Bellot, P. (2014), 'Accurate and Effective Latent Concept Modeling for ad hoc Information Retrieval', *Document Numérique, Lavoisier* **17**(1), 61–84.
URL: *http://dn.revuesonline.com/article.jsp?articleId=19419* 56

Doms, M. E. & Morin, N. J. (2004), 'Consumer Sentiment, the Economy, and the News Media', *FRB of San Francisco Working Paper* (09).
URL: *https://papers.ssrn.com/abstract=602763* 9, 36

Dzielinski, M. (2012), 'Measuring Economic Uncertainty and Its Impact on the Stock Market', *Finance Research Letters* **9**(3), 167–175.
URL: *http://www.sciencedirect.com/science/article/pii/S1544612311000468* 7

Eckley, P. (2015), 'Measuring Economic Uncertainty Using News-Media Textual Data', *University Library of Munich* . MPRA Paper 64874.
URL: *https://ideas.repec.org/p/pra/mprapa/64874.html* 3, 35

Enders, W. (2014), *Applied Econometric Time Series, 4th Edition*, 4. edition edn, Wiley, New York. 25

Farkas, R., Vincze, V., Móra, G., Csirik, J. & Szarvas, G. (2010), The CoNLL-2010 Shared Task: Learning to Detect Hedges and their Scope in Natural Language Text, *in* 'Proceedings of the Fourteenth Conference on Computational Natural Language Learning – Shared Task', Association for Computational Linguistics, Uppsala,

Sweden, pp. 1–12.
URL: *https://www.aclweb.org/anthology/W10-3001* 37

Feldman, M. P. (1999), 'The New Economics Of Innovation, Spillovers And Agglomeration: Areview Of Empirical Studies', *Economics of Innovation and New Technology* **8**(1-2), 5–25.
URL: *https://doi.org/10.1080/10438599900000002* 8

Fernández-Villaverde, J., Guerrón-Quintana, P., Kuester, K. & Rubio-Ramírez, J. (2015), 'Fiscal Volatility Shocks and Economic Activity', *American Economic Review* **105**(11), 3352–3384.
URL: *https://www.aeaweb.org/articles?id=10.1257/aer.20121236* 36

Fernández-Villaverde, J., Guerrón-Quintana, P., Rubio-Ramírez, J. F. & Uribe, M. (2011), 'Risk Matters: The Real Effects of Volatility Shocks', *The American Economic Review* **101**(6), 2530–2561.
URL: *https://www.jstor.org/stable/23045650* 7

Ferreira, T. R. T. (2018), 'Stock Market Cross-Sectional Skewness and Business Cycle Fluctuations', *FRB International Finance Discussion Paper* (1223).
URL: *https://ideas.repec.org/p/fip/fedgif/1223.html* 5

Garcia, D. (2012), 'Sentiment During Recessions', *Journal of Finance* **68**(3), 1267–1300.
URL: *https://papers.ssrn.com/abstract=1571101* 9

Gentzkow, M., Kelly, B. & Taddy, M. (2019), 'Text as Data', *Journal of Economic Literature* **57**(3), 535–574.
URL: *https://pubs.aeaweb.org/doi/10.1257/jel.20181020* 6, 9

Gilchrist, S., Sim, J. W. & Zakrajšek, E. (2014), 'Uncertainty, Financial Frictions, and Investment Dynamics', *FEDS Working Paper* (69).
URL: *https://ideas.repec.org/p/nbr/nberwo/20038.html* 8

Goldstein, D. G. & Taleb, N. N. (2007), 'We Don't Quite Know What We Are Talking About', *The Journal of Portfolio Management* **33**(4), 84–86.
URL: *https://jpm.pm-research.com/content/33/4/84* 5

Goodell, J. W., McGee, R. J. & McGroarty, F. (2020), 'Election Uncertainty, Economic Policy Uncertainty and Financial Market Uncertainty: A Prediction Market Analysis', *Journal of Banking & Finance* **110**.
URL: *http://www.sciencedirect.com/science/article/pii/S0378426619302584* 37

Griffiths, T. L. & Steyvers, M. (2004), 'Finding Scientific Topics', *Proceedings of the National Academy of Sciences* **101**(suppl 1), 5228–5235.
URL: *https://www.pnas.org/content/101/suppl*$_1$/5228 16, 56

Haddow, A. (2013), 'Macroeconomic Uncertainty: What is it, How can we measure it and Why does it matter?', *Bank of England Quarterly Bulletin* **53**(2), 10. 4, 7, 39

Hastie, T., Tibshirani, R. & Friedman, J. (2009), *The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition*, Springer Series in Statistics, 2 edn, Springer-Verlag, New York.
URL: *https://www.springer.com/de/book/9780387848570* 20

Hausman, D. M. (1994), *The Philosophy of Economics: An Anthology*, Cambridge University Press, Cambridge. 3

Hicks, J. R. (1931), 'The Theory of Uncertainty and Profit', *Economica* **32**, 170–189.
URL: *https://www.jstor.org/stable/2547922* 3

Istiak, K. & Serletis, A. (2018), 'Economic Policy Uncertainty and Real Output: Evidence from the G7 Countries', *Applied Economics* **50**(39), 4222–4233.
URL: *http://www.tandfonline.com/doi/abs/10.1080/00036846.2018.1441520* 34

Jensen, D. D. & Cohen, P. R. (2000), 'Multiple Comparisons in Induction Algorithms', *Machine Learning* **38**(3), 309–338.
URL: *https://doi.org/10.1023/A:1007631014630* 56

Jones, T. & Doane, W. (2019), 'textmineR: Functions for Text Mining and Topic Modeling'. Retrieved: 01-10-2020.
URL: *https://CRAN.R-project.org/package=textmineR* 57

Jurado, K., Ludvigson, S. C. & Ng, S. (2015), 'Measuring Uncertainty', *American Economic Review* **105**(3), 1177–1216.
URL: *https://www.jstor.org/stable/43495414* 6, 8, 37

Kalamara, E., Turrell, A., Redl, C., Kapetanios, G. & Kapadia, S. (2020), 'Making Text Count: Economic Forecasting Using Newspaper Text', *Bank of England Working Paper* (865).
URL: *https://papers.ssrn.com/abstract=3610770* 9, 35

Kaufmann, D., Kraay, A. & Zoido, P. (1999), 'Governance Matters', *World Bank Policy Research Working Paper Series* (2196).
URL: *https://papers.ssrn.com/abstract=188568* 34

Kehrig, M. (2015), 'The Cyclical Nature of the Productivity Distribution', *US Census Bureau Center for Economic Studies* (11).
URL: *https://papers.ssrn.com/abstract=1854401* 37

Keynes, J. M. (1936), *The General Theory of Employment, Interest, and Money*, Macmillan Cambridge University Press.
URL: *https://www.springer.com/de/book/9783319703435* 3

Knight, F. H. (1921), *Risk Uncertainty and Profit*, illustrated edition edn, Signalman Publishing, Kissimmee, Florida. 3

Konjunkturinstitutet (2020), 'Economic Tendency Survey'. Retrieved: 2020-11-01.
URL: *https://www.konj.se/english/publications/economic-tendency-survey.html* 24, 49

Lahiri, K. & Sheng, X. (2010), 'Measuring Forecast Uncertainty by Disagreement: The Missing Link', *Journal of Applied Econometrics* **25**(4), 514–538.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1002/jae.1167* 6

Larsen, V. (2017), 'Components of Uncertainty', *Centre for Applied Macro- and Petroleum economics (CAMP), BI Norwegian Business School* (4).
URL: *https://papers.ssrn.com/abstract=2955655* 6, 7, 17, 52, 56

Larsen, V. H. & Thorsrud, L. A. (2019*a*), 'Business Cycle Narratives', *CESifo Working Paper Series* (7468).
URL: *https://ideas.repec.org/p/ces/ceswps/*$_7$468.*html* 9, 17

Larsen, V. H. & Thorsrud, L. A. (2019*b*), 'The value of news for economic developments', *Journal of Econometrics* **210**(1), 203–218.
URL: *http://www.sciencedirect.com/science/article/pii/S0304407618302148* 52

Larsen, V. H., Thorsrud, L. A. & Zhulanova, J. (2019), 'News-driven Inflation Expectations and Information Rigidities', *Norges Bank* (5).
URL: *https://ideas.repec.org/p/bno/worpap/2019*$_0$5.*html* 10

Lazear, E. P. & Spletzer, J. R. (2012), 'Hiring, Churn, and the Business Cycle', *The American Economic Review* **102**(3), 575–579.
URL: *https://www.jstor.org/stable/23245596* 8

Leduc, S. & Liu, Z. (2016), 'Uncertainty Shocks are Aggregate Demand Shocks', *Journal of Monetary Economics* **82**(C), 20–35.
URL: *https://ideas.repec.org/a/eee/moneco/v82y2016icp20-35.html* 5

Light, M., Qiu, X. Y. & Srinivasan, P. (2004), The Language of Bioscience: Facts, Speculations, and Statements In Between, *in* 'HLT-NAACL 2004 Workshop: Linking Biological Literature, Ontologies and Databases', Association for Computational Linguistics, Boston, Massachusetts, USA, pp. 17–24.
URL: *https://www.aclweb.org/anthology/W04-3103* 37, 38

Lin, X., Bloom, N. & Alfaro, I. (2017), 'The Finance-Uncertainty Multiplier', *NBER Working Paper Series* (24571).
URL: *https://ideas.repec.org/p/red/sed017/887.html* 8

Ludvigson, S. C. (2004), 'Consumer Confidence and Consumer Spending', *Journal of Economic Perspectives* **18**(2), 29–50.
URL: *https://www.aeaweb.org/articles?id=10.1257/0895330041371222* 34

Ma, F., Wahab, M. I. M., Liu, J. & Liu, L. (2018), 'Is Economic Policy Uncertainty important to forecast the Realized Volatility of Crude Oil Futures?', *Applied Economics* **50**(18), 2087–2101.
URL: *http://www.tandfonline.com/doi/abs/10.1080/00036846.2017.1388909* 37

Manela, A. & Moreira, A. (2015), 'News Implied Volatility and Disaster Concerns', *Journal of Financial Economics* **123**(1), 137–162.
URL: *https://www.sciencedirect.com/science/article/pii/S0304405X16301751* 6

Medlock, B. (2008), 'Exploring Hedge Identification in Biomedical Literature', *Journal of Biomedical Informatics* **41**(4), 636–654.
URL: *http://www.sciencedirect.com/science/article/pii/S1532046408000087* 37

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S. & Dean, J. (2013), Distributed Representations of Words and Phrases and their Compositionality, *in* C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani & K. Q. Weinberger, eds, 'Advances in Neural Information Processing Systems 26', Curran Associates, Inc., pp. 3111–3119.
URL: *https://arxiv.org/abs/1310.4546* 38

Mullainathan, S. & Spiess, J. (2017), 'Machine Learning: An Applied Econometric Approach', *Journal of Economic Perspectives* **31**(2), 87–106.
URL: *http://pubs.aeaweb.org/doi/10.1257/jep.31.2.87* 9

North, D. C. (2005), *Understanding the Process of Economic Change*, Princeton University Press, Princeton.
URL: *https://www.jstor.org/stable/j.ctt7zvbxt* 3

Nyman, R., Kapadia, S., Tuckett, D., Gregory, D., Ormerod, P. & Smith, R. (2018), 'News and Narratives in Financial Systems: Exploiting Big Data for Systemic Risk Assessment', *Bank of England Working Paper* (704).
URL: *https://papers.ssrn.com/abstract=3098729* 9

Pang, B., Lee, L. & Vaithyanathan, S. (2002), Thumbs up? Sentiment Classification using Machine Learning Techniques, *in* 'Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing (EMNLP 2002)', Association for Computational Linguistics, pp. 79–86.
URL: *https://www.aclweb.org/anthology/W02-1011* 38

Porter, M. (1980), 'An Algorithm for Suffix Stripping', *Program* **14**(3), 130–137.
URL: *https://doi.org/10.1108/eb046814* 52

Romer, C. D. (1990), 'The Great Crash and the Onset of the Great Depression', *The Quarterly Journal of Economics* **105**(3), 597–624.
URL: *https://academic.oup.com/qje/article/105/3/597/1864581* 7

Saltzman, B. & Yung, J. (2018), 'A Machine Learning Approach to Identifying different types of Uncertainty', *Economics Letters* **171**, 58–62.
URL: *http://www.sciencedirect.com/science/article/pii/S016517651830257X* 6

Scotti, C. (2016), 'Surprise and Uncertainty Indexes: Real-time Aggregation of Real-Activity Macro-Surprises', *Journal of Monetary Economics* **82**(C), 1–19.
URL: *https://econpapers.repec.org/article/eeemoneco/v*$_3$*a*82$_3$*ay*$_3$*a*2016$_3$*ai*$_3$*ac*$_3$*ap*$_3$*a*1 − 19.*htm* 8

Shapiro, A. H., Sudhof, M. & Wilson, D. J. (2020), 'Measuring News Sentiment', *Federal Reserve Bank of San Francisco Working Paper Series* (1).
URL: *https://ideas.repec.org/p/fip/fedfwp/2017-01.html* 9

Shiller, R. J. (2017), 'Narrative Economics', *American Economic Review* **107**(4), 967–1004.
URL: *https://www.aeaweb.org/articles?id=10.1257/aer.107.4.967* 1, 2, 34

Sum, V. (2012), 'Economic Policy Uncertainty and Stock Market Performance: Evidence From the European Union, Croatia, Norway, Russia, Switzerland, Turkey and Ukraine', *Journal of Money, Investment and Banking* **25**, 99–104.
URL: *https://papers.ssrn.com/abstract=2094175* 37

Tam, P. S. (2018), 'Global Trade Flows and Economic Policy Uncertainty', *Applied Economics* **50**(34-35), 3718–3734. 37

Tetlock, P. C. (2007), 'Giving Content to Investor Sentiment: The Role of Media in the Stock Market', *The Journal of Finance* **62**(3).
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.2007.01232.x* 10

Tetlock, P. C., Saar-Tsechansky, M. & Macskassy, S. (2008), 'More Than Words: Quantifying Language to Measure Firms' Fundamentals', *The Journal of Finance* **63**(3).
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.2008.01362.x* 10

The Economist (2001), 'The R-Word'. Retrieved: 01-12-2020.
URL: *https://www.economist.com/finance-and-economics/2001/04/05/the-r-word* 1

The Federal Reserve Bank of Atlanta (2020), 'Survey of Business Uncertainty'. Retrieved: 01-10-2020.
URL: *https://www.frbatlanta.org:443/research/surveys/business-uncertainty* 5

The University of Michigan (2020), 'Surveys of consumers'. Retrieved: 01-10-2020.
URL: *https://data.sca.isr.umich.edu/charts.php* 5

Thorsrud, L. A. (2020), 'Words are the New Numbers: A Newsy Coincident Index of the Business Cycle', *Journal of Business & Economic Statistics* **38**(2), 393–409.
URL: *https://doi.org/10.1080/07350015.2018.1506344* 9

Tobback, E., Naudts, H., Daelemans, W., Junqué de Fortuny, E. & Martens, D. (2018), 'Belgian Economic Policy Uncertainty Index: Improvement through Text Mining', *International Journal of Forecasting* **34**(2), 355–365.
URL: *http://www.sciencedirect.com/science/article/pii/S016920701630098X* 6, 38

Toda, H. Y. & Yamamoto, T. (1995), 'Statistical Inference in Vector Autoregressions with Possibly Integrated Processes', *Journal of Econometrics* **66**(1), 225–250.
URL: *http://www.sciencedirect.com/science/article/pii/0304407694016168* 26, 27

Triana, P. (2009), *Lecturing Birds on Flying: Can Mathematical Theories Destroy the Financial Markets?*, John Wiley & Sons, New York. 5

A Appendix: Data

Article Format

Figure A1 Bag-of-Words Representation of Articles.

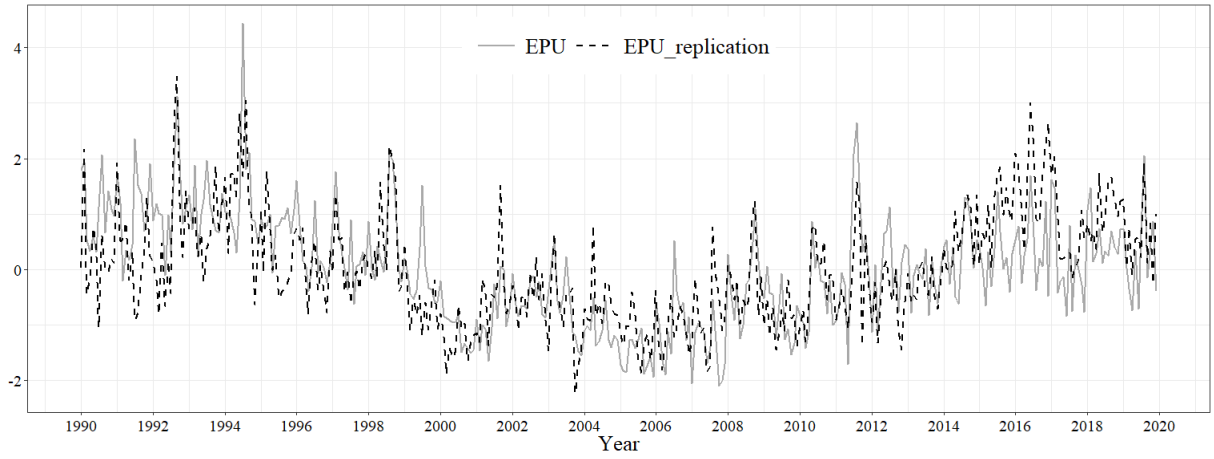
	docdate	sort_dict
	<date>	<chr>
1	2018-04-07	[('', 31), ('adoption', 1), ('aktiviteter', 1), ('aktörer', 1), ('aldrig', 2), ('alla', 3), ('allt', 5), ('andra', 2), ~
2	2018-04-06	[('', 656), ('&', 2), ('a', 68), ('adams', 2), ('adielsson', 2), ('alla', 2), ('allt', 1), ('andersson', 6), ('andra', ~
3	2018-04-06	[('', 1), ('att', 1), ('börjat', 2), ('då', 1), ('från', 1), ('fulham', 1), ('få', 1), ('för', 1), ('har', 2), ('i', 1)~
4	2018-04-05	[('', 9), ('agent', 1), ('aldrig', 1), ('allt', 1), ('alltså', 1), ('amatören', 1), ('andra', 1), ('aperfectdream', 1),~
5	2018-04-06	[('', 2), ('afpreutersttaftonbladet', 1), ('annan', 1), ('att', 2), ('av', 2), ('beskjutits', 1), ('beväpnad', 1), ('bl~
6	2018-04-06	[('', 14), ('akiva', 1), ('alltihop', 1), ('amerikansk', 1), ('anfaller', 1), ('annars', 1), ('artist', 1), ('att', 1),~
7	2018-04-05	[('', 9), ('alla', 2), ('alltid', 1), ('att', 5), ('av', 3), ('bland', 1), ('bli', 4), ('bobby', 2), ('bolt', 2), ('bor~
8	2018-04-05	[('', 34), ('aldrig', 1), ('allt', 3), ('alltid', 1), ('amorteringskrav', 1), ('amorteringskravet', 1), ('anges', 2), (~
9	2018-04-08	[('', 32), ('actor', 1), ('aftonbladet', 1), ('aktuell', 2), ('alla', 7), ('allra', 1), ('allt', 2), ('alltid', 2), ('a~
10	2018-04-05	[('', 1), ('annat', 1), ('annonseras', 1), ('att', 3), ('balar', 1), ('berättar', 1), ('bland', 1), ('både', 1), ('de',~

... with 24,978 more rows

Notes: Each row represents one article

Replication of The Swedish EPU-Index

Figure A2 Replication of the Swedish EPU-Index as outlined in §4.



Notes: The correlation between the original and the replicated EPU index is 0.71

Data Sources

Table A1 Sources of Macroeconomic Variables Used in the Evaluation.

Indicator Name	Abbreviation	Source	Interval	Adjusted
Confidence Indicators	CI_*	Swedish National Institute of Economic Research	Monthly	
OMXS30 Stock Price Index	$OMXS30$	Nasdaq database	Monthly	
3-Month Interbank Rate	R	FRED database	Monthly	
Unemployment Rate	U	FRED database	Monthly	
Consumper Price Index	CPI	FRED database	Monthly	Seasonally
Producer Price Index	PPI	FRED database	Monthly	Seasonally
Production of Total Industry	P	FRED database	Monthly	Seasonally
GDP Growth	GDP	FRED database	Quarterly	
Fixed capital formation	I	Statistics Sweden	Quarterly	Seasonally
Household Consumption Expenditure	C	Statistics Sweden	Quarterly	Seasonally

Notes: FRED stands for Federal Reserve Bank of St. Louis. Seasonal adjustment was conducted with X-13ARIMA-SEATS.

Table A2 Questions included in the Swedish Confidence Indicators

Indicator	Questions Included
Industry	*Manufacturing:* Order books (present situation assessment) – stock of finished goods (present situation assessment) + production volume (expectations) *Construction:* Order books (present situation assessment) + number of employed (expectations) *Trade:* Selling volume (outcome) – stock of goods (present situation assessment) + selling volume (expectations) *Private service sector:* The firm's business situation (outcome) + Demand for the firm's services (outcome) + Demand for the firm's services (expectations)
Consumers	Financial situation of the household now + Financial situation of the household within 12 months + Swedish economy now + Swedish economy within 12 months + good time to buy consumer durables now
Micro	Financial position of household (assessment of present situation) + financial position of household (expectations) + right time to make major purchases (assessment of present situation) + major purchases by household (expectations)
Macro	Swedish economy (assessment of present situation) + Swedish economy (expectations) - unemployment (expectations)

Notes: Each indicator is calculated as the mean of the net balances (seasonally adjusted and standardised) for the selected questions. *Source:* Konjunkturinstitutet (2020).

B Appendix: Measurement

NLP-Terminology

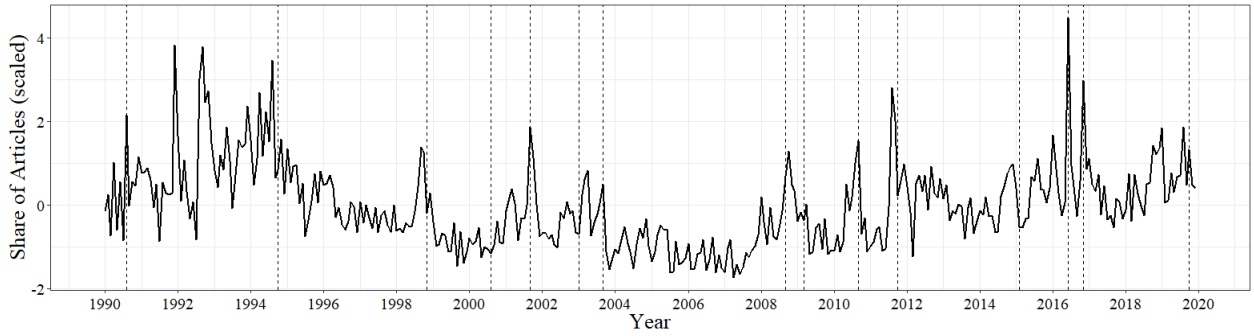
Table B1 Definition of Key NLP-Terminology.

Terminology	Definition
Corpus	Large set of structured or unstructured text documents
Unigram	Smallest entity of a word with meaning in a given text (e.g. *debt*)
N-gram	Continuous sequence of *n* words with meaning in a given text (e.g. *household debt*)
Token	An individual unit of meaning in a text and the basic unit of discrete data for text analysis (both unigrams and n-grams can be tokens)
Stop Word	Most common tokens in a corpus that do not convey any semantic meaning
tf-idf (term frequency-inverse document frequency)	A score that estimates how characteristic a specific a token is for a document compared to the whole corpus, meaning that common tokens appearing in many document will have a low tf-idf score
Bag-of-Words	Simplified representation of the multiplicity of a text or corpus, where word order and grammar is disregarded
Document-term-matrix	Related to Bag-of-Words; matrix representation of a corpus, where each row is one document, each column is a token, and the matrix values measure how much and if each token appears in each document (word order and grammar again disregarded)
Tokenization	Converting a sequence of text into its individual units of meaning
Lemmatization	Extracting the linguistic root of a word that is common to all of its inflections (do have a meaning by themselves)
Stemming	Extracting the linguistic stem of a word that is common to all of its inflections (do *not* have a meaning by themselves)
Syntactic analysis	Analysis of the sentence or grammatical structure of the text
Semantic analysis	Analysis of the narrative meaning of a text

Notes: See §3 for a conceptual overview of NLP.

Description of Uncertainty-Related Shocks

Figure B1 *EFU* Index and Uncertainty-related Shocks.



Notes: Black dotted lines highlight uncertainty-related shocks, detailed in Table B2 in the Appendix.

Table B2 Major uncertainty-related Shocks as labelled in Figure B1.

Date	Description
1990 - 10	Gulf War I & Global Oil Price Shock
1995 - 09	Highest Swedish unemployment after the Swedish banking crisis
1997 - 11	Collapse of the Stockholm Stock Exchange, Asian Financial Crisis
1998 - 09	Russian Financial Crisis
2000 - 03	Nasdaq peak before burst of Dot-com bubble
2001 - 09	Burst of the Dot-com Bubble (approx.)
2002 - 10	Nasdaq bottom after burst of Dot-com bubble
2003 - 02	Gulf War II
2003 - 09	Swedish referendum against membership of the eurozone
2007 - 08	Start of the Credit Crunch and Global Financial Crisis
2008 - 09	Bankcrupcy of Lehman Brothers
2010 - 09	Swedish general elections and entering of far-right into government
2011 - 10	Greek proposed economy referendum & Nasdaq Crash
2015 - 02	Riksbank launched its bond purchase program
2016 - 06	Brexit Votum
2016 - 11	Election of Donald Trump as future US president
2019 - 10	Krona's weakest position against the US Dollar

Article Pre-Processing

The quality of a topic model is heavily sensitive to the type and order of pre-processing steps applied to the text corpus. Ideally, one would derive multiple pre-processed text corpora by varying number and order of pre-processing steps for each different corpus, and then compare the sensitivity of article proximities within each corpus (see Denny & Spirling 2018). This process, however, requires timely and computationally extensive resources which lie outside the scope of this thesis. Instead, I rely on personal experience and the approach of papers whose research framework is similar to mine, namely Larsen & Thorsrud (2019*b*) and Larsen (2017), in order to choose the sequence and type of the pre-processing conducted. First, all characters are transformed to lower case and extremely common and extremely rare characters were removed from articles. To the former belong all Swedish and English *stop words*, which are tokens that do not contribute any information relating to the subject of an article (words such as *och, av, är, etc*). The latter comprises tokens that appear less or equal to five times in the entire article corpus. I also remove tokens with which I filtered articles (dictionary **e**, **f** and **u**) since they will otherwise appear in many topics but do not convey any latent meaning in the texts. Next, the most common Swedish first and last names are eliminated.[17] I then transform each token into its linguistic *stem*, which is the part of a word that is common to all of its inflections (e.g. the stem of the tokens *osäkerhet* and *osäker* is *osäk*), via the standard stemming algorithm of Porter (1980). Furthermore, I only keep alphabetic characters in the corpus, meaning that all digits, white spaces, and punctuation characters are removed. The average number of tokens in the articles is 300. Finally, articles are removed if they count less than 50 tokens in order to avoid very short texts (around 1,000 articles). I exclude these articles because I want to focus on longer articles that are more likely to entail a sentimental or narrative load. Articles are also excluded if they count more than 10.000 tokens because I regard this length as non-standard for newspaper articles (around 30 articles).

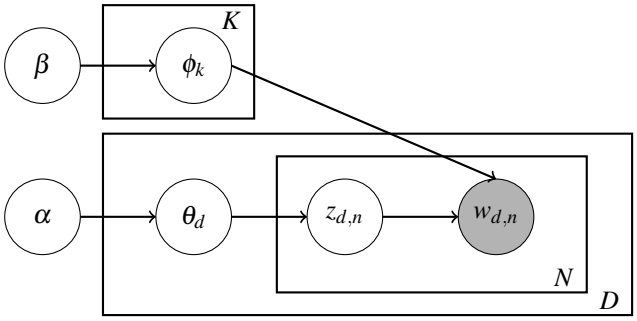
At this point it is important to note that the document-term-matrix, into which the corpus will be turned in order to compute the LDA model, is very sparse, meaning that the matrix entails a large number of word columns, but only very few cells with values other than zero (since most of the tokens will not appear in most of the texts). In order to increase the computing efficiency, I therefore reduce the sparseness of this matrix as a last step by removing tokens that are very distinct to articles. This means that I only keep the vocabulary that is widely shared across articles and thereby increase the share of non-zero values in the dtm-matrix. Sparse tokens are identified by the document-frequency (df) of tokens. A token i is removed if it appears in less than 0.0001% of all articles D, specifically if $df_i < D * 0.0001\%$ (that equals to at least 10 articles), or in other words if its sparseness is more than 99%. After these pre-processing steps, I remain with a corpus that consists of 100,179 articles and 15 million tokens (initially 16 million), out of which 61,924 tokens are unique (initially 816,029).

[17] I used peer-compiled stop word lists, which are publicly available under the following links:
Stop Words in Swedish: https://github.com/peterdalle/svensktext/tree/master/stoppord
Stop Words in English: http://snowball.tartarus.org/dist/snowball_all.tgz
Names: https://github.com/peterdalle/svensktext/tree/master/namn

Inference Procedure

The generative process of Latent Dirichelt Allocation (LDA) is illustrated graphically in Figure B2. The observed variables are the words of the articles (shaded grey) and all other variables are hidden (unshaded). The LDA model has three levels, where rectangles denote replications for each level: In the first level ϕ_k is sampled for each topic k. In the second level, θ is sampled for each document in D. Finally, z and w are sampled for every word in N of each document in D. α and β are hyperparameters, which are detailed below.

Figure B2 Probabilistic Graphical Model of LDA.



Notes: Observed variables are shaded grey and hidden variables are unshaded. Rectangles denote replications for each of the three levels of the LDA model. First, ϕ_k is sampled for each topic k. Second, θ is sampled for each document in D. Third, z and w are sampled for every word in N of each document in D. α and β are hyperparameters. *Source:* Blei et al. (2003).

Knowing the algorithm for generating the documents, the distributions ϕ and θ can be estimated by computing the posterior distribution of the hidden variables (Blei 2012, pp. 81):

$$p(\phi, \theta, z|w) = \frac{p(\phi, \theta, z, w)}{p(w)}, \tag{12}$$

where the numerator represents the joint distribution of all random variables (hidden and observed) and the denominator is the marginal probability of the observed words (i.e. the probability word n appearing in corpus D). Theoretically, the posterior has to be computed by summing the joint distribution over every possible combination of latent topic structures. But since the number of possible structures is exponentially large, the procedure is computationally impossible to implement. Therefore, the problem is solved by approximating the posterior distribution with inference algorithms that fit an alternative distribution over the hidden variable structure. The two most common approximations for the posterior distribution are i) *sampling-based* and ii) *variational* algorithms (Blei 2012). In the former case, a small sub-collection of samples is used to approximate the topic structure with maximal probability. Variational methods, on the other hand, transform the inference problem to an optimization problem by imposing a parameterized set of distributions over the hidden structure and identifying the member of that set that is closest to the posterior distribution. For memory efficiency reasons, I use a sampling-based method. Concretely Gibbs sampling with 1000 iterations is used, where a sequence of random variables are constructed, each dependent on the previous, and samples from the posterior are collected with the limiting distributions.

Hyperparameter Selection

The Dirichlet Distribution and Its Priors α, β

The probability model of LDA defines that the following variables follow a Dirichlet distribution:

$$\phi \sim Dirichlet(\beta) \quad \text{and} \quad \theta \sim Dirichlet(\alpha). \tag{13}$$

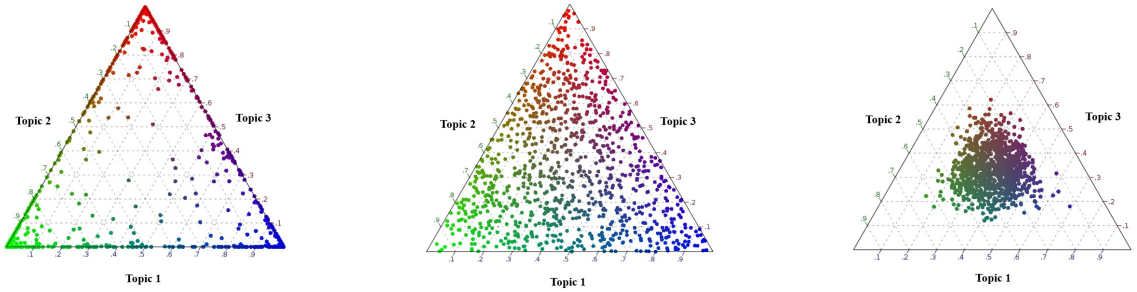
In the following, the relation between α and θ will be described in further detail, but note that the implications are similar for β and ϕ. The hyperparameter α is the parameter of the Dirichlet random variable $\theta_d = \{\theta_{d,1}, ..., \theta_{d,k}\}$, where $\theta_{d,k}$ is the probability that topic k appears in article d. Hence, θ_d is k-dimensional. It takes values in the (k-1)-simplex and is normalized to 1, i.e. $\sum_{k=1}^{K} \theta_{d,i} = 1$. The probability density on this simplex is (Blei et al. 2003, p. 996):

$$p(\theta|\alpha) = \frac{\Gamma(\sum_{i=1}^{k} \alpha_i)}{\prod_{i=1}^{k} \Gamma(\alpha_i)} \theta_1^{\alpha_1 - 1} \cdots \theta_k^{\alpha_k - 1}, \tag{14}$$

where α is a k-dimensional vector with $\alpha_i > 0$, and $\Gamma(x)$ is the Gamma function. Note that α was initially determined as a scalar, even though here it is defined as a k-vector. This means that all components of the k-vector are assigned the same value (which is the value of the scalar) and that the Dirichlet distribution is symmetrical. What does symmetry imply? As mentioned earlier, a high (low) value of α imposes dense (sparse) topic proportions to articles. A symmetrical Dirichlet assigns topic probabilities at random, which means that in dense articles, the topic probability distributions would be fairly uniform and in sparse articles the topic probabilities would be skewed towards certain sides randomly. On the other hand, an asymmetrical Dirichlet would also influence the distribution of the Dirichlet θ and assign higher probabilities to topics that were assigned a higher α_i. However, in almost all software applications of the LDA algorithm the Dirichlet distribution is set to be symmetrical.

The intuition of the relationship between the Dirichlet and α can be illustrated with a graphical example. If we are given a corpus of 1,000 articles ($D = 1,000$) and want to identify three latent topics in them ($K = 3$), we have a 3-dimensional vector $\alpha = (\alpha_1, \alpha_2, \alpha_3)$ which we have to set. For symmetric Dirichlet distributions, three scenarios are possible: $\alpha < 1, \alpha = 1, \alpha > 1$. In the figures in B3 and B4, examples of the symmetric and asymmetric Dirichlet distributions are respectively presented in two-dimensions. In the triangles, each dot represents $\theta_d = \{\theta_{d,1}, \theta_{d,2}, \theta_{d,3}\}$ of article $d \in (1, 1000)$ and the edges determine the probability distributions of each respective topic ($\theta_{d,k}$) in article d, where $k \in \{1, 2, 3\}$. For $\alpha < 1$, the articles are located close to the edges, implying that most articles contain only one topic or at most two topics with a high probability. On the other extreme, $\alpha > 1$ imposes very dense topic compositions on articles which means that each topic is about equally likely to appear in each article. For the asymmetric Dirichlet distributions in B4, each α_i can be arbitrarily set (although not applicable in LDA). In the scenario of low alphas, the distribution is uneven and also sparse, so that articles mostly contain only *one of two* topics with a high probability. Because sparseness is imposed, the less frequent topic can still have high probabilities in a few articles. Lastly, the scenario of high alphas make the distribution uneven but dense, hence, two topics are likely to occur *together* in the majority of articles and due to denseness the less frequent topic does rarely appear in any document.

Figure B3 Prior α and the Symmetric Dirichlet Distribution.



(a) $\alpha = 0.1$ (b) $\alpha = 1$ (c) $\alpha = 10$

Figure B4 Prior α and the Asymmetric Dirichlet Distribution.



(a) $\alpha = (0.1, 0.5, 0.5)$ (b) $\alpha = (5, 10, 10)$

Notes: Each dot represents $\theta_d = \{\theta_{d,1}, \theta_{d,2}, \theta_{d,3}\}$ of article $d \in (1, 1000)$. The edges of the triangles determine the probability distributions of each respective topic ($\theta_{d,k}$) in article d, where $k \in \{1, 2, 3\}$.

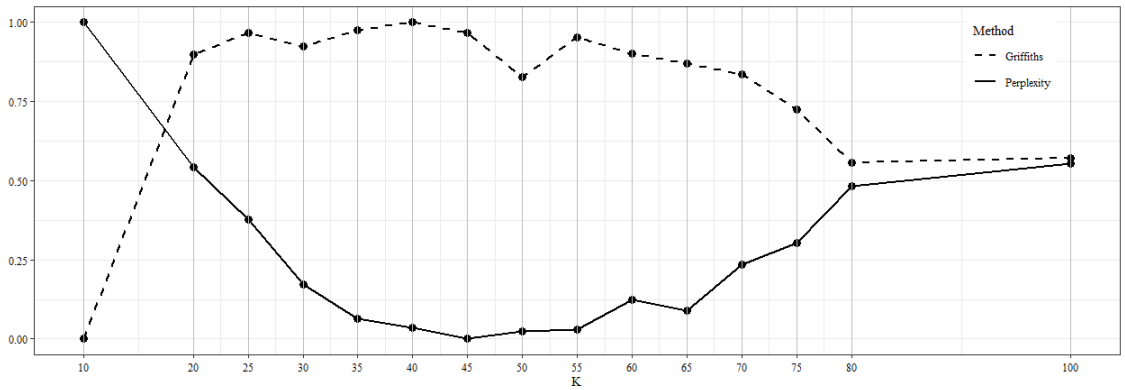
Optimal Number of Topics K

The NLP-literature provides several propositions to objectively approximate the most likely K (Blei et al. 2003, Griffiths & Steyvers 2004, Deveaud et al. 2014, Cao et al. 2009, Arun et al. 2010). I follow the original authors' proposition to compare the maximum log-likelihood of topic densities on a held-out article corpus - in language modelling also referred to as *perplexity*. The perplexity is monotonically decreasing in the likelihood of the held-out corpus, i.e. a lower perplexity score means better generalization performance. For efficiency reasons, I do not use the entire corpus for computing the perplexity scores but conduct a grid-search over a chosen set of topics $K \in \{10, 20, 25, 30, 35, 40, 45, 50, 60, 80, 100\}$ with 10-fold cross-validation.[18] This means that out of a random sample of 10,000 articles (10% of the total articles) I create 10 different random splits.[19] For each split I divide the articles into a training set (8000 articles) and a test set (2000 articles). For each K, I fit a topic model on each of the ten training sets and then calculate the perplexity score of the respective test set with the per-word likelihoods from the fitted model. Formally, for each K, perplexity is calculated as follows (Blei et al. 2003, p.1008):

$$perplexity(D^i_{test}) = \exp\left\{ -\frac{\sum_{d=1}^{D} \log p(\mathbf{w_d})}{\sum_{d=1}^{D} N_d} \right\}, \tag{15}$$

where i is the i-th split. Finally, I average the perplexity scores of the each K across all 10 train-test splits. For comparison, I also approximate the optimal number of K with the other most common method from Griffiths & Steyvers (2004). Here the log-likelihood maximization of $P(\mathbf{w}|K)$ is used, which is the likelihood of the probability that words are assigned to a given set of topics K. Similar to what we saw in (5), using $P(\mathbf{w}|K)$ poses computational problem since it requires to sum over all possible assignments of words to topics z, and so it is approximated by using the harmonic mean of a set of values of $P(w|z,K)$, where z is sampled from the posterior distribution $P(z|w,K)$. I calculate the harmonic mean of $P(w|z,K)$ for each K in the drawn sample of 10,000 articles. The scores for each method for each number of K is presented in Figure B5. K_{opt} seems to lie around 40 to 45 topics, which score slightly better than the rest of the topic numbers between 35 and 55. Therefore, a number in between the two optimum points, $K = 43$, is chosen.

Figure B5 Optimal K Approximation.



Notes: On the y-axis, the normalized values of the perplexity scores and the $P(w|z,K)$ (called *Griffiths*) are presented, respectively. A well fitted LDA model is indicated by a low perplexity score and a high Griffiths score.

[18] The selected range of topic numbers is based on previous studies that used LDA on similar article corpora: Azqueta-Gavaldon (2017) used $K_{opt} = 30$ and Larsen (2017) chose $K_{opt} = 80$.

[19] A split of 10 is regarded as the best practice in the literature of machine learning (Jensen & Cohen 2000)

Overview of Inferred Topics

Table B3 presents the list of inferred topics by the top 20 words that have the highest probability of belonging to the respective topic. The topics are sorted by the frequency with which they appear in the article corpus, also referred to as their *prevalence* score. Topic prevalence measures the marginal probability with which a topic appears in the whole corpus. For a topic k, prevalence P_k is the sum of the topic-per-document probabilities of all articles ($\theta_{d,k}$) normalized by the total number of articles:

$$P_k = \frac{\sum_d \theta_{d,k}}{D}. \tag{16}$$

The meaningfulness of topics can be furthermore evaluated with their *coherence* scores. Since articles were pre-filtered by keywords related to economic expectations, it is likely that we encounter a large set of vocabulary that is semantically related to these keywords and occurs frequently in many articles. Since these words are predominant, there will be most likely a large number of topics that entail these words with a high probability. However, topics where all these keywords co-occur are not necessarily useful since the words in question appear so frequently in the corpus as to be meaningless. In other words, frequent words are highly correlated with each other but they are statistically-independent of each other (Jones & Doane 2019). Having a large share of frequent words in a topic will consequently influence its quality. This quality can be measured with the concept of *probabilistic coherence*. It measures the correlation of word distributions within topics, controlling for statistical independence of words. Formally, for each pair of words a, b in the top I words in a topic, probabilistic coherence calculates $P(b|a)P(b)$, where a is more probable than b in the topic. In other words, if looking only at documents that contain the word a, then word b should be more more probable in those documents than words that are chosen at random from the corpus. $P(b|a)$ measures how probable b is only in documents containing a and $P(b)$ measures how probable b is in the entire corpus. If b is not more probable in documents containing a, then the difference $P(b|a)P(b)$ should be near zero. This difference is calculated across the top I words and the average gives the probabilistic coherence measure. The higher the coherence score of a topic, the more semantically meaningful it is in itself. The coherence and prevalence scores of each topic are presented in Table B3. Since the coherence measure is bounded between 0 and 1, I regard coherence above 0.1 as relatively good. There is only small set of articles, around 14 topics, that score below 0.1. All in all, I evaluate the topic output as useful given that most of the prevalent topics have coherence scores above 0.1, therefore being semantically meaningful and less noisy. Only a few of the highly prevalent topics has a coherence score below 0.1, such as the topic labelled as *Unknown*. It contains mostly propositions and verbs. Hence, its low coherence is intuitive because many articles use these auxiliary words, making the topic at the same time prevalent. Among the highest scoring coherent, it is intuitive that topics about sports and literature are highly coherent because of their specific vocabulary.

Table B3 Latent Topics in News about Economic Expectations.

Topic	Label	Coherence	Prevalence	Most Important Words
27	Politics & Society	0.09	5.8	politisk, politik, talet, människ, social, skap, makt, värld, samhället, rör, led, lev, gaml, egn, histori, försök, medborg, samhäll, handl, stället
22	Business Cycles	0.13	5.5	tillväxt, bnp, vänt, räkn, minsk, ökad, hushåll, konsumtion, arbetslös, ökning, ökar, progno, åren, stig, utveckling, öka, efterfrågan, väx, svag, oecd
31	Unknown	0.83	4.5	sver, bör, utveckling, kräv, svensk, skap, innebär, förutsättning, betyd, tyd, fråg, debat, mål, utveckl, förändring, uppgift, långsikt, samarbet, krav, lösning
43	Prepositions	0.37	4.0	gaml, börj, sa, såg, låt, nästan, bild, händ, värld, stod, satt, ord, kall,

				väg, ens, rikt, histori, trot, minn, plöts
3	Central Bank	0.30	3.8	riksbank, inflation, ränt, räntan, sänk, centralbank, höj, låg, kronan, penningpolitik, stig, procenten, räntehöjning, höjning, riksbankschef, fed, bank, reporäntan
10	Jobs	0.63	3.5	tyck, jobb, peng, känn, börj, tänk, väld, se, folk, person, berät, handl, sak, prat, fråg, arbet, hitt, håll, träff, försök
25	Currencies	0.19	3.4	steg, dollarn, punk, vänt, sjönk, doll, stärk, svensk, kronan, sa, räntan, euron, amerikansk, torsdag, kron, fredag, onsdag, ränt, medan, euro
35	Stock Exchange	0.14	3.2	börs, amerikansk, månad, uppgång, marknad, veckan, index, analytik, nedgång, steg, fortsat, kurs, föll, europ, positiv, vänt, börj, kraft, stark, veck
40	Reform	0.56	3.1	stat, regering, styr, förslag, verksam, utredning, ordför, ans, mynd, lag, regl, kritik, krav, frågan, förslaget, organisation, fråg, uppgift, ansv, uppdrag
17	Industry Growth	0.77	2.8	sver, företag, svensk, utveckling, investering, länd, skap, näringslivet, marknad, näringsliv, utländsk, land, väx, konkurrenskraft, internationell, tillväxt, invester, utveckl
8	European Union	0.17	2.8	eu, länd, europ, europeisk, storbritanni, kommission, brittisk, union, frankrik, parlamentet, fransk, bryssel, förhandling, folkomröstning, regering, premiärminist, medlemsländ
14	Global Conflict	0.19	2.7	militär, landet, irak, krig, kriget, fn, led, president, konflik, amerikansk, regim, värld, hot, stat, politisk, död, makt, sanktion, internationell, protest
32	Swedish Krona	0.12	2.4	kron, miljard, miljon, peng, cirk, drygt, räkn, kostnad, motsvar, stat, doll, betal, kost, beräkn, totalt, minsk, ytter, omkring, innebär, vänt
34	Job Market	0.11	2.3	jobb, arbetslös, arbet, arbetsmarknad, person, anställd, utbildning, antalet, lön, människ, sysselsättning, arbetskraft, arbetsgiv, företag, ung, minsk, ungdom, skap, grupp, personal
16	Global Risk	0.13	2.2	kris, imf, risk, finansiell, internationell, drabb, värld, finanskris, oro, landet, global, situation, tving, problem, talet, börj, snabbt, följd, åtgärd, värst
6	US Presidency	0.17	1.9	president, trump, röst, obam, amerikansk, valet, välj, bush, kongress, val, vinn, politisk, republikan, barack, demokrat, kandidat, konservativ, politik, huset, landet
13	Sports	0.12	1.8	spel, klubb, match, os, säsong, vm, lag, trän, laget, aik, svensk, djurgård, vinn, mål, tävling, göteborg, malmö, guld, poäng, miljon
5	Europe	0.14	1.8	emu, länd, sver, grekland, euron, valut, europeisk, euro, itali, spani, tyskland, europ, svensk, eu, euroländ, ecb, valutaunion, grekisk, valutan, kronan
37	Investment	0.21	1.8	akti, fond, bolag, börs, invester, köp, avkastning, placer, peng, risk, spar, sälj, förvalt, vinst, kapital, företag, värd, skandi, marknad, svensk
24	Research	0.13	1.8	lär, forskning, universitet, utbildning, forsk, student, skolan, kunskap, elev, skol, profess, studi, vetenskap, högskol, studer, arbet, stockholm, teknisk, läs
42	Financial Sector	0.13	1.7	bank, lån, ränt, norde, swedbank, hushåll, peng, köp, pris, rör, skuld, svensk, risk, finansinspektion, betal, bostadsmarknad, handelsbank, bolån, låg, utlåning
19	Culture	0.18	1.6	film, musik, konstnär, konst, spel, kultur, teat, publik, mod, scen, stockholm, utställning, låt, svensk, museet, samtid, festival, besök, artist, skap
1	Leisure	0.081	1.6	mat, bil, köp, fisk, pris, runt, sälj, butik, besök, vin, turist, vatt, djur, res, natur, bönd, somm, människ, sver, båt
2	Fiscal Policy	0.115	1.6	betal, pension, kron, skatt, inkomst, avgift, peng, spar, pensionär, lön, systemet, månad, tjän, mån, ersättning, försäkring, arbet, höj, lägr, välj
23	Energy	0.18	1.6	kärnkraft, utsläpp, energi, minsk, värld, vattenfall, bygg, använd, global, grön, miljön, klimatförändring, jord, miljö, bränsl, koldioxid, länd, teknik, ny, reaktor
36	Housing	0.07	1.5	stockholm, bygg, stad, lägen, hus, bostäd, flytt, lokal, förening, huset, gaml, byggnad, bor, städ, malmö, plan, kyrkan, hyr, området, kvadratmet
29	Emerging Economies	0.076	1.4	värld, länd, doll, global, indi, brasili, olj, oljepriset, priset, europ, pris, argentin, fat, landet, asi, amerikansk, mexiko, latinamerik, efterfrågan, kanad
33	Media	0.06	1.4	tv, tidning, se, läs, internet, medi, svd, digital, nyhet, kanal, använd, svt, nätet, radio, express, journalist, information, tjänst, program, titt
18	Community	0.067	1.3	kommun, stockholm, region, kommunal, län, landet, bygg, landsting, göteborg, skån, lokal, väg, trafik, sl, invån, landstinget, regional, stad, peng, tåg
39	Asia	0.01	1.3	japan, kinesisk, kin, japansk, landet, värld, hongkong, asi, länd, peking,

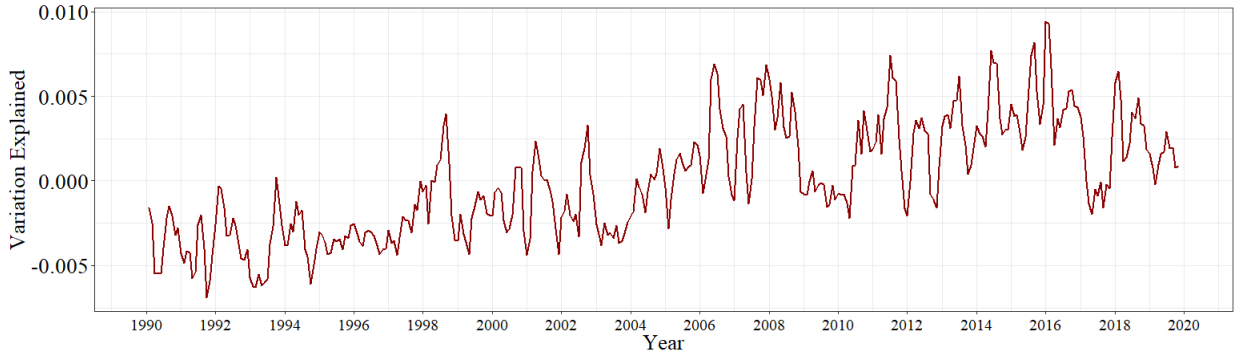
				sydkore, nordkore, wto, led, tull, asiatisk, utländsk, kines, amerikansk, taiwan
20	Automobile	0.14	1.2	bil, volvo, saab, tillverk, köp, fabrik, sälj, företaget, produktion, försäljning, scani, produk, modell, gm, marknad, ford, lastbil, bygg, ny, europ
7	Immigration	0.11	1.1	turkiet, landet, fn, flykting, människ, afrik, regering, bosni, kosovo, fatt, serbi, land, turkisk, miljon, befolkning, invandr, jugoslavi, länd, bistånd, etnisk
11	Electronic	0.07	1.05	sas, plan, noki, flyg, teli, bolaget, mobil, flygbolag, ericsson, passager, mobiltelefon, flygplats, island, konkurrent, flygplan, bolag, telefon, operatör
12	Trade Unions	0.1	1.0	avtal, lön, fack, part, avtalet, förhandling, medlemm, löneökning, arbetsgiv, ordför, facket, lönebildning, krav, förhandl, strejk, industrin, tco, avtalsrör, saf, metall
4	Government	0.02	3.78	parti, regering, socialdemokrat, politik, partiet, välj, moderat, valet, politisk, statsminist, riksdag, sver, röst, allians, socialdemokratisk, val, reinfeld
38	Financial Reporting	0.24	3.25	vd, bolaget, bolag, kron, akti, vinst, mkr, företaget, försäljning, omsättning, result, kvartalet, verksam, marknad, miljon, affär, anställd, miljard, ägar, rapport
15	Probabilities	0.53	3.23	se, tyck, påverk, svårt, betyd, frågan, händ, risk, väld, tyd, oro, gansk, handl, viss, effek, svd, sett, faktisk, exempel, särskilt
21	Public Policy	0.11	2.66	regering, skatt, sänk, offent, stat, finansminist, budget, politik, utgift, reform, förslag, skattesänkning, höj, åtgärd, finans, arbetslös, sver, statsskuld
41	Crime	0.82	1.81	polis, person, brott, uppgift, domstol, lämn, misstänk, lag, betal, avslöj, brotts, mynd, peng, express, svensk, dn, berät, skad, anklag, åklag
30	Eastern Europe	0.2	1.43	ryssland, rysk, ukrain, putin, pol, president, jeltsin, nato, moskv, landet, estland, väst, lettland, stat, ryss, baltisk, sovjetunion, ung, litau, premiärminist
28	Public Health	0.14	1.33	kvinn, vård, sjukvård, patient, läk, män, sjukhus, landsting, sjukdom, äldr, sjuk, behandling, personal, människ, läkemedel, omsorg, medicinsk, privat, medicin, kvalitet
26	Northern Neughbours	0.21	1.32	norg, dansk, finland, sver, svensk, norsk, danmark, nordisk, finsk, nord, förat, finländsk, bl, sidan, tvåår, detär, oslo, ans, norrmän, köpenhamn
9	Central Europe	0.15	0.74	tyskland,europ, frankrik, merkel, fransk, kohl, förbundskansl,österrik, bundesbank, cdu, östr, mark, östtysk, spd, ddr, deutsch, frankfurt, bonn

Notes: Topics sorted by the frequency with which they appear in the article corpus, also referred to as their *prevalence* score. Most Important Words refer to the tokens that were assigned the highest probability of belonging to the topic in question.

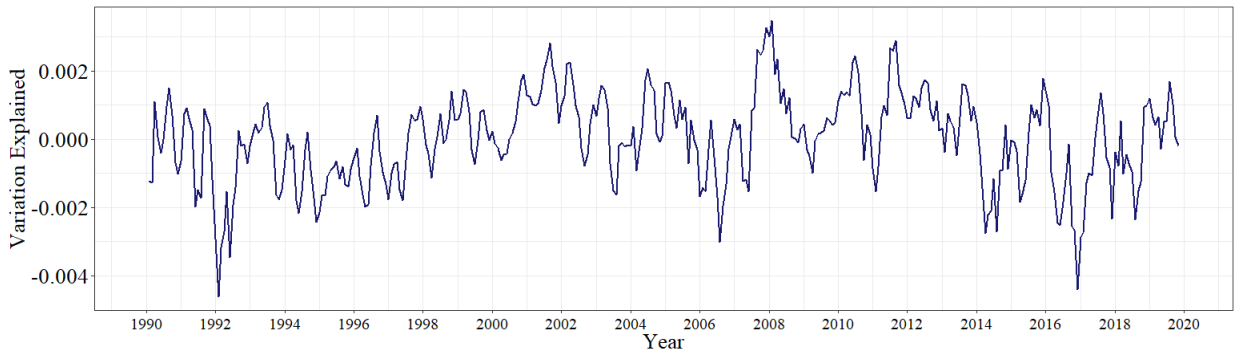
Principal Component Analysis

Figure B6 Timely Variation Captured by Top Principal Components.

(a) Principal Component 1: 66.7% of variation explained



(b) Principal Component 2: 6% of variation explained



Notes: The y-axis represents the principal component score of each component at time t (z_{ti}).

C Appendix: Additional Results

Impact on Economic Fundamentals

The multivariate SVAR specification of Armelius et al. (2017) serves as a benchmark model to estimate the impact EFU, EFU-MACRO and EFU-POL on the Swedish macroeconomy and compare their results with the impact generated by the EPU. Specifically, the SVAR analyzes the per-quarter effect of one of the newly created news-based uncertainty indices on Swedish GDP growth ($GDP_t(\%)$), household savings rate (as a percentage of total income) ($\Delta S_t(\%)$), gross fixed capital formation ($I_t(\%)$), the Repo rate ($R_t(\%)$), and the unemployment rate ($U_t(\%)$) (specified in this same order by the Cholesky decomposition). This also implies that uncertainty does not react to shocks of other variables in the first period. The analyzed period ranges from 1994:02 until 2019:04 and overall comprises 10 additional quarters for analysis compared to the original specification in Armelius et al. (2017). The number of lags amounts to 4. The uncertainty indices are averaged over quarters GDP, investment, and unemployment enter in quarterly growth rates, while savings enters in first-difference and the Repo rate in (percent) levels. The specification looks as follows:

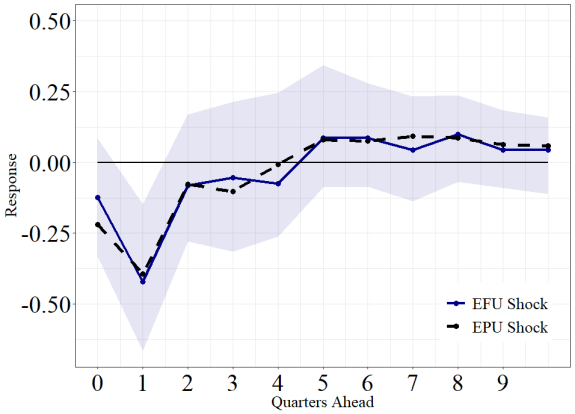
$$\mathbf{X_t} = \mathbf{A_0} + \sum_j \mathbf{A_j}\mathbf{X_{t-j}} + e_t, \quad \text{where} \quad \mathbf{X_t} \equiv \begin{bmatrix} Uncertainty_t \\ GDP_t(\%) \\ \Delta S_t(\%) \\ I_t(\%) \\ R_t(\%) \\ U_t(\%) \end{bmatrix} \quad \text{and} \quad e_t \sim W.N. \tag{17}$$

First I compare the impact of dictionary based uncertainty indices EFU and EPU. The first figure in C1 compares the impulse response functions of quarterly GDP growth to a shock in EFU and EPU respectively, and the figures below compare the impulse responses of the other variables to a shock in either EFU and EPU. Both economic forecast uncertainty and economic policy uncertainty cause a drop in GDP growth in the first two quarters by almost the same magnitude. The impact of the first lag of both uncertainty indices is statistically significant at the 1%-level. The impact of EFU and EPU shocks on the other variables in the SVAR are only marginally or not statistically significant, but still seem to move their responses into the right direction.

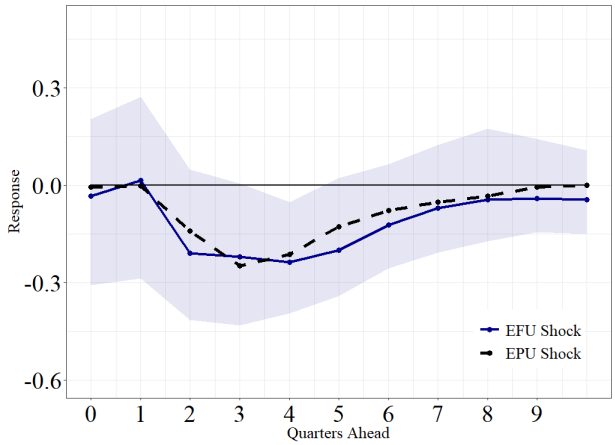
Next, I move to the topic-based uncertainty indices and use the same specification to estimate the impact of a shock in EFU-MACRO and EFU-POL. As can be seen from the impulse response functions in C2b, a shock in the topic-based uncertainty indices overall send a similar impulse to the macro economy as does the EFU index. The shock in EFU-MACRO is significant at a 1%-level for GDP growth but shocks to the other variables again mostly non-significant. Moreover, shocks in policy uncertainty measured with EFU-POL do not show any significant impact at all.

Figure C1 VAR Estimation of the Impact of a EFU versus EPU Shock on the Swedish Aggregate Economy.

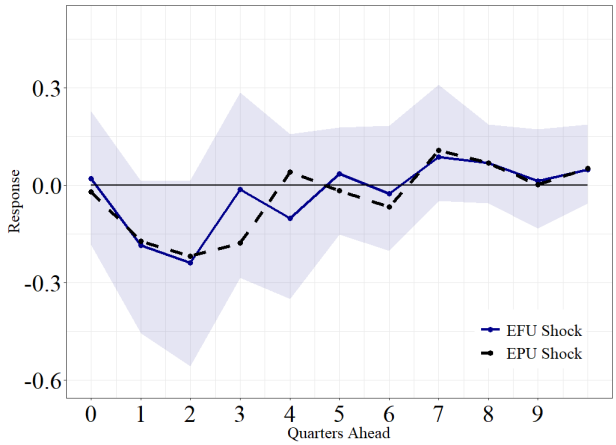
(a) GDP Growth (%)



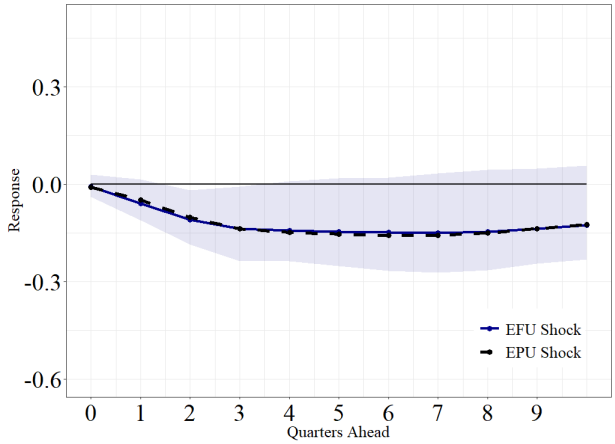
(b) ΔSavings Rate (%)



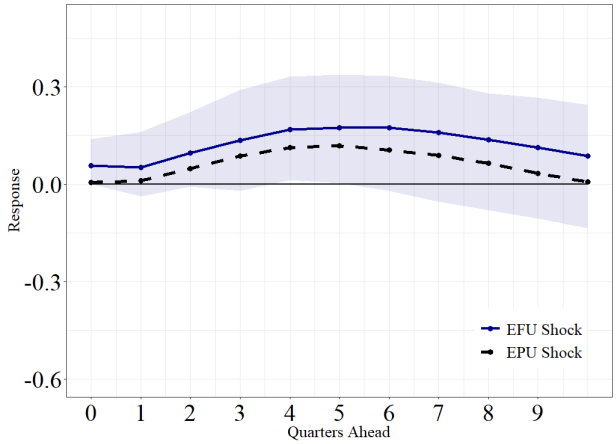
(c) Investment Rate (%)



(d) Repo Rate (%)



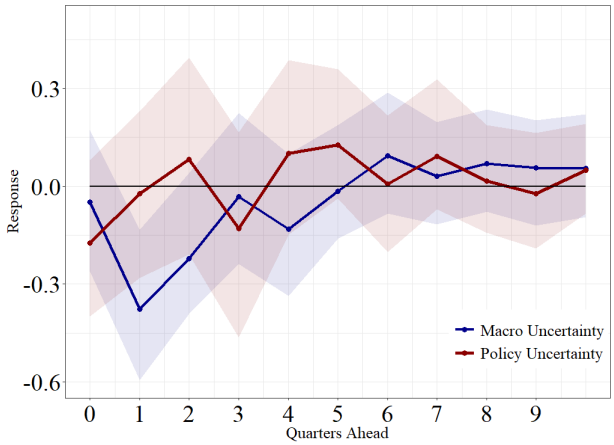
(e) Unemployment Rate (%)



Notes: Shaded areas represent 95% standard residual-based confidence bounds of a shock in the EFU-index.

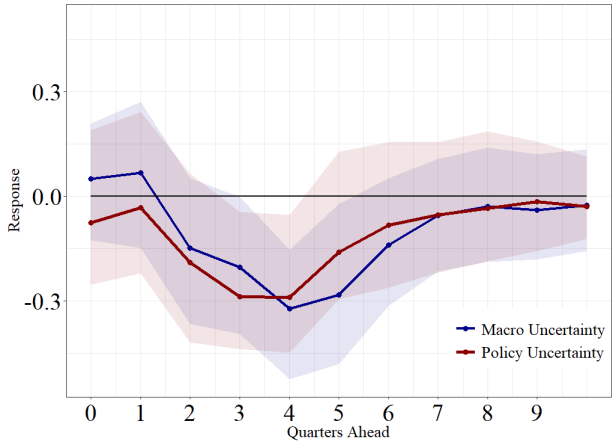
Figure C2 VAR Estimation of the Impact of a EFU-MACRO and EFU-POL Shock on the Swedish Aggregate Economy.

(a) GDP Growth (%)

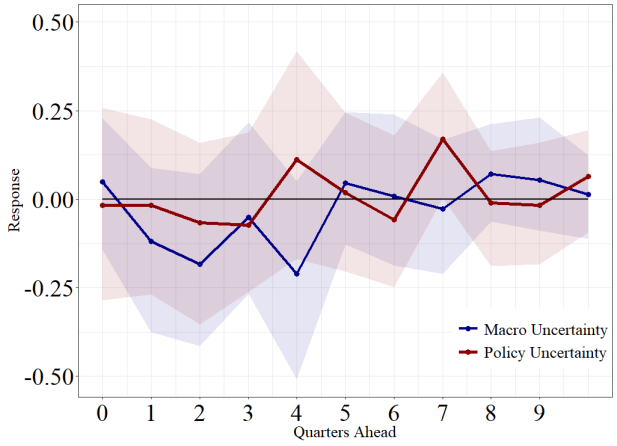


(b)

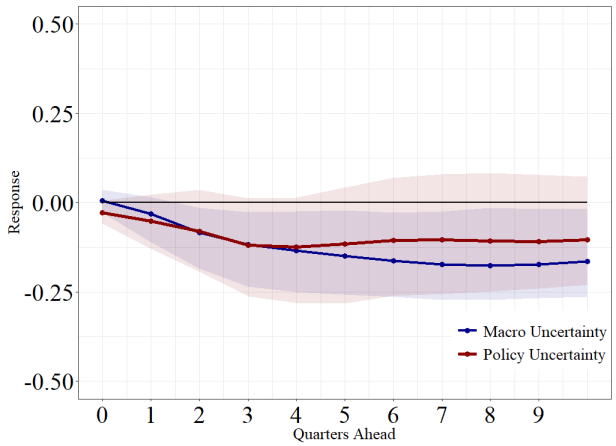
(c) ΔSavings Rate (%)



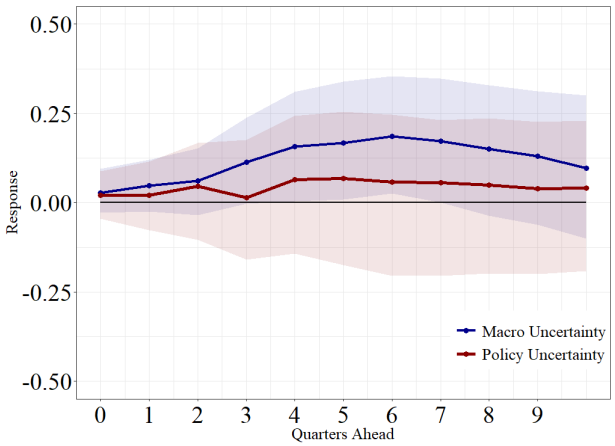
(d) Investment Rate (%)



(e) Repo Rate (%)



(f) Unemployment Rate (%)



Notes: Shaded areas represent 95% standard residual-based confidence bounds of a shock in EFU-MACRO and EFU-POL respectively.

Impact on Economic Confidence

Properties of the Toda-Yamamoto Procedure

Table C1 Propoerties of the VARs and Breusch-Godfrey Test Results for Serially Correlated Errors.

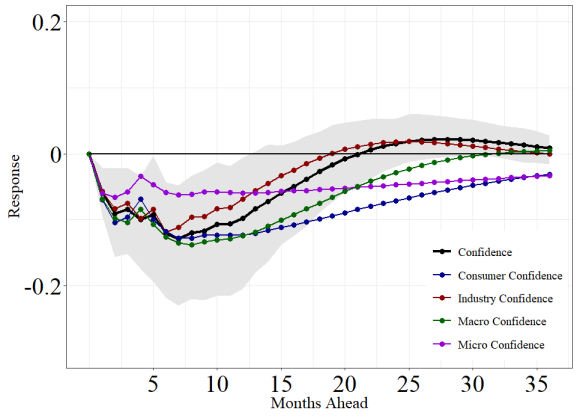
Variables	AIC-selected lags (p)	Lags chosen ($p+d_{max}$)	χ^2 (BG-LM)	p-value (BG-LM)	d.f.
$EFU-CI$	6	6 + 7	27.95	0.11	20
$EFU-CI_I$	7	7 + 7	22.05	0.34	20
$EFU-CI_C$	5	5 + 7	27.54	0.12	20
$EFU-CI_{Micro}$	5	10 + 7	20.09	0.45	20
$EFU-CI_{Macro}$	5	12 + 7	19.94	0.46	20
$EPU-CI$	6	6 + 7	24.23	0.23	20
$EPU-CI_I$	7	7 + 7	26.15	0.16	20
$EPU-CI_C$	6	9 + 7	24.89	0.21	20
$EPU-CI_{Micro}$	6	9 + 7	24.89	0.21	20
$EPU-CI_{Macro}$	6	6 + 7	19.43	0.49	20
$EFUMACRO-CI$	2	5 + 7	20.75	0.19	20
$EFUMACRO-CI_I$	2	2 + 7	56.72	0.45	20
$EFUMACRO-CI_C$	2	2 + 7	66.0	0.17	20
$EFUMACRO-CI_{Cmicro}$	3	3 + 7	60.202	0.20	20
$EFUMACRO-CI_{Cmacro}$	3	3 + 7	60.67	0.19	20
$EFUPOL-CI$	6	6 + 7	45.19	0.26	20
$EFUPOL-CI_I$	8	8 + 7	36.85	0.25	20
$EFUPOL-CI_C$	5	5 + 7	51.40	0.21	20
$EFUPOL-CI_{Cmicro}$	5	5 + 7	54.43	0.13	20
$EFUPOL-CI_{Cmacro}$	6	6 + 7	49.5	0.14	20

Notes: BG-LM refers to the Breusch-Godfrey LM statistic. d_{max} is the maximal order of integration in the process identified with the Augmented Dickey-Fuller Test. *p<0.1, **p<0.05, ***p<0.01.

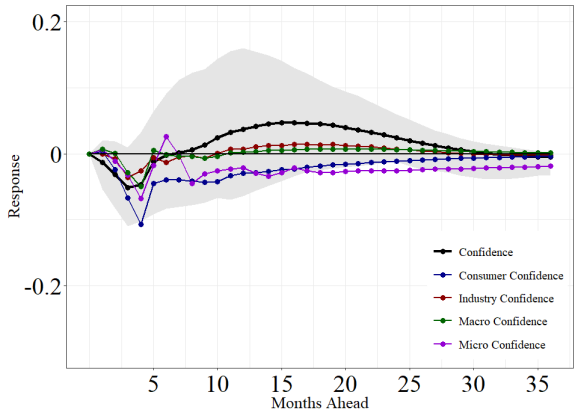
Robustness Checks

Figure C3 VAR Estimation of the Impact of a News-based Uncertainty Shock on Confidence - Confidence ordered first.

(a) Impulse Response of on EFU



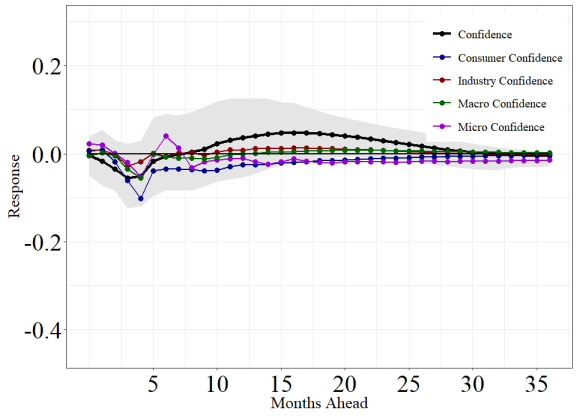
(b) Impulse Response on EPU



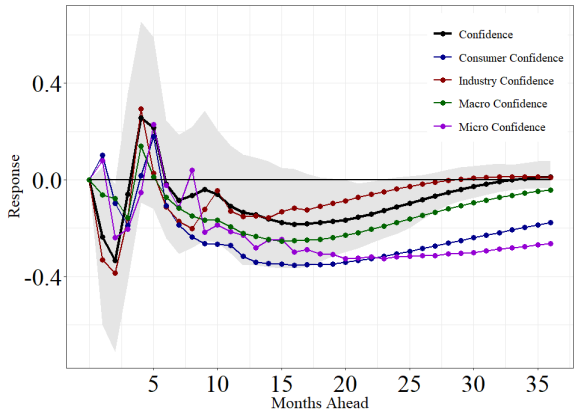
Notes: Grey segments represent the area between the upper and lower 1 standard-error bands of the shock response of general confidence (black line).

Figure C4 VAR Estimation of the Impact of a Confidence Shock on News-based Uncertainty.

(a) Impulse Response of EFU



(b) Impulse Response of EPU



Notes: Grey segments represent the area between the upper and lower 1 standard-error bands of the shock in general confidence (black line).