The Value of News

Vegard Høghaug Larsen and Leif Anders Thorsrud

Norges Bank and Centre for Applied Macro- and Petroleum Economics (CAMP)

Berlin 2016

A larger research agenda

- Can "Big Data" and tell us something about the (macro) economy?
- Yes, it can - Using the same type of data as in what's presented today:
  1. We can predict asset prices on a daily frequency
  2. Construct high frequent (daily) measures of the business cycle
  3. Address the news driven view of the business cycle - today's topic

Why do economic activity fluctuate?

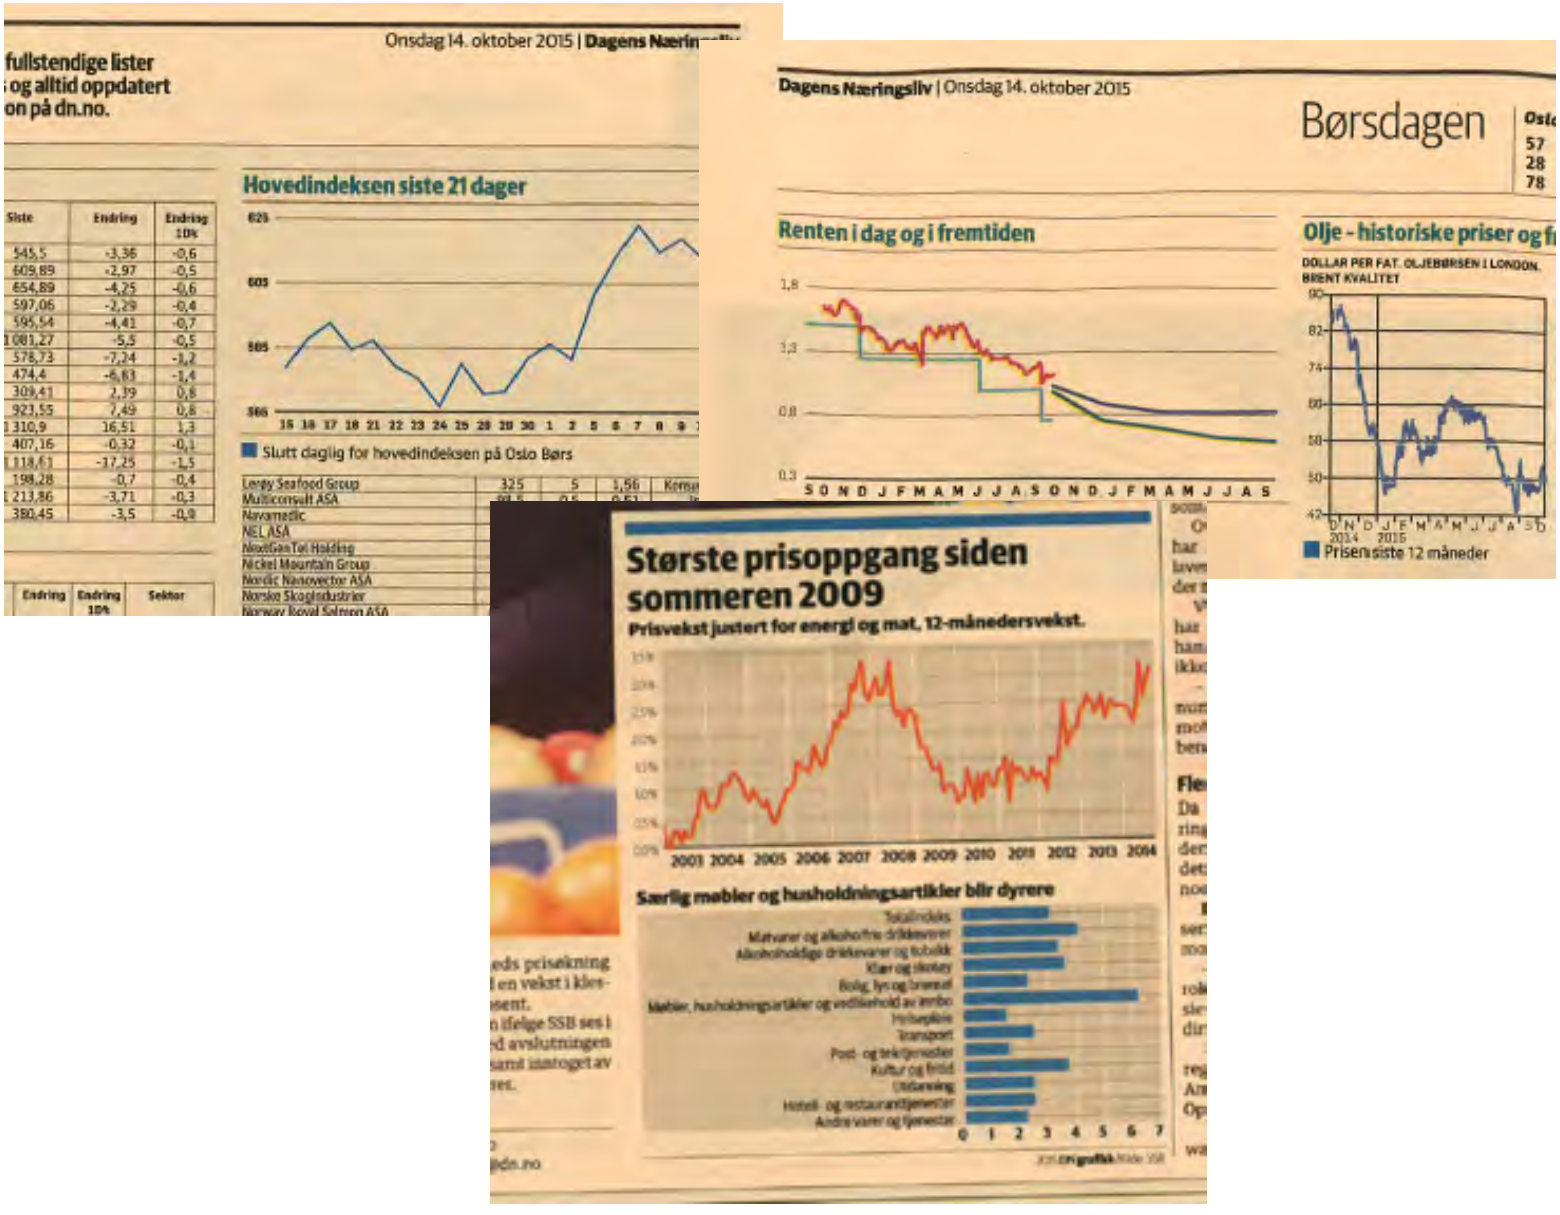

The News Driven Business Cycle view

- **Expectations** matter for explaining economic fluctuations
- Changes in expectations, due to **news**, is the key ingredient in the news-driven business cycle view, see e.g., Beaudry and Portier (2006, AER), Barsky and Sims (2012, AER), Blanchard et al. (2013, AER), Lorenzoni (2009, AER)

- What is news? Information received today that contain **fundamental information** on outcomes that might materialize in the future
- In theory models: Agents in the economy receive signals about future economic developments, and then face a signal extraction problem: **What is news and what is noise?**

Problem: Identifying news shock

- A problem with testing theories for news shocks is that the news is typically not observed
- One solution is to link news to innovations in asset prices, since asset prices should reflect all available information, Beaudry and Portier (2006, AER).
- But:
  - Asset prices change for a lot of reasons, likely containing both news and noise
  - News about what? Productivity, future policy, energy prices, etc.
  - The content of the news shock should matter for the structural interpretation

Our solution:

Use a primary source for news directly - the newspaper!

Step 1

We **generate a measure of news**, by decomposing a Norwegian business newspaper according to the topics it writes about using a **Latent Dirichlet Allocation (LDA)** model (Blei et al. (2003)):

- The LDA reduces a vast dataset of news articles into a much smaller dataset containing daily topic frequencies
- The topic frequencies are aggregated into time series - one time series for each topic

Key hypothesis

The more a newspaper write about a topic the more likely it is that this topic reflects something of importance for the economy's future needs and developments

We then ...

Step 2 **Construct an aggregate news series**

- **Evaluate the predictive power** of each news (topic) series - a "signal extraction" problem
- **Combine the news topics** - using the predictive results

Step 3 **Address the news driven view of the business cycle** using SVARs - commonly applied in the news literature

- News shocks identified using the aggregated news index

Contribution 1

- A growing number of studies (in finance in particular) use textual data to understand economic fluctuations. We build on this literature
- In contrast to word count procedures that have commonly been applied, the **LDA decomposition** offers **several advantages**:
  - We do not need to use "subjectively" chosen positive or negative words or a specific context
  - **Topics have content in their own right** - more likely to contain fundamental information than word counts (which is often associated with sentiment)
  - **Topics permit an interpretation**

Contribution 2

In contrast to previous work - in the news driven business cycle literature - we:

- **Use a primary source of new information - the newspaper**
- **Can explore** what type of news topics that are important for the aggregate economy, i.e., **what constitutes a news shock**?
- **Separate news from noise shocks**, i.e., if news shocks are innovations to our news index, what is unexpected innovations to asset prices? We call it noise

Main results:

- The LDA decomposition of the business newspaper delivers news topics that are easily classified and interpretable
- Many news topics predict macroeconomic outcomes. Especially noteworthy is the news topics ability to predict asset prices
- In line with theoretical predictions:
  - News shocks generate large and persistent aggregate fluctuations and a permanent increase in productivity and consumption
  - Noise shocks cause only an initial boom in consumption, and then a contraction
- Among the most important topics contributing to the news shocks are *Funding*, *IT/Startup*, *Oil production*, and *Monetary policy*

Step 1: Topic extraction

Data:

- Todays Business (TB), Norway's biggest business newspaper
- Retriever's "Atekst" database: 25 years of newspaper data (almost 500 000 articles, over 1 billion words)

The LDA takes a set of articles as input and return two sets of distributions:

- One set of distributions over words - one distribution for each topic
- One set of distributions over topics - one distribution for each article

Estimation: Bayesian MCMC...

LDA - a graphical representation

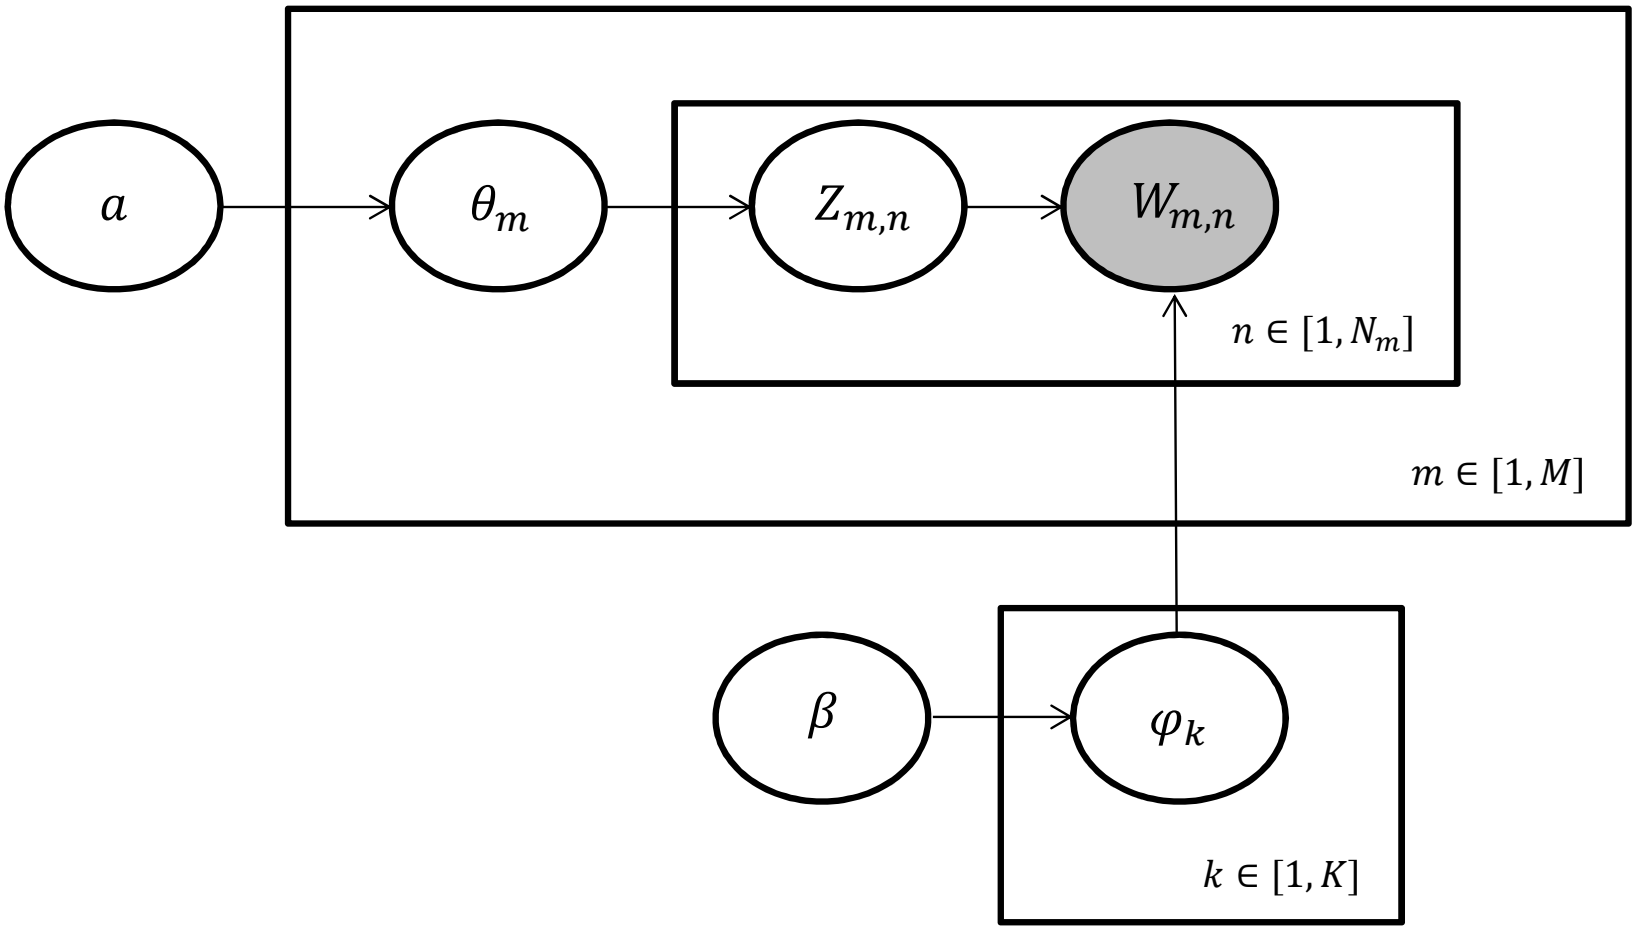

The generative process for the LDA model

1. Pick the overall theme of articles by randomly giving them a distribution over topics, i.e.: Choose $\theta_i \sim \mathrm{Dir}(\alpha)$, where $i \in \{1, \ldots, M\}$.
2. Pick the word distribution for each topic by giving them a distribution over words, i.e.: Choose $\varphi_k \sim \mathrm{Dir}(\beta)$, where $k \in \{1, \ldots, K\}$.
3. For each of the word positions i, j, where $j \in \{1, \ldots, N_i\}$, and $i \in \{1, \ldots, M\}$
   - 3.1 From the topic distribution chosen in 1., randomly pick one topic, i.e.: Choose a topic $z_{i,j} \sim \mathrm{Multinomial}(\theta_i)$.
   - 3.2 Given that topic, randomly choose a word from this topic, i.e.: Choose a word $w_{i,j} \sim \mathrm{Multinomial}(\varphi_{z_{i,j}})$.

Topic 14: A distribution over words

We call it the Funding topic

A topic representation of the whole corpus

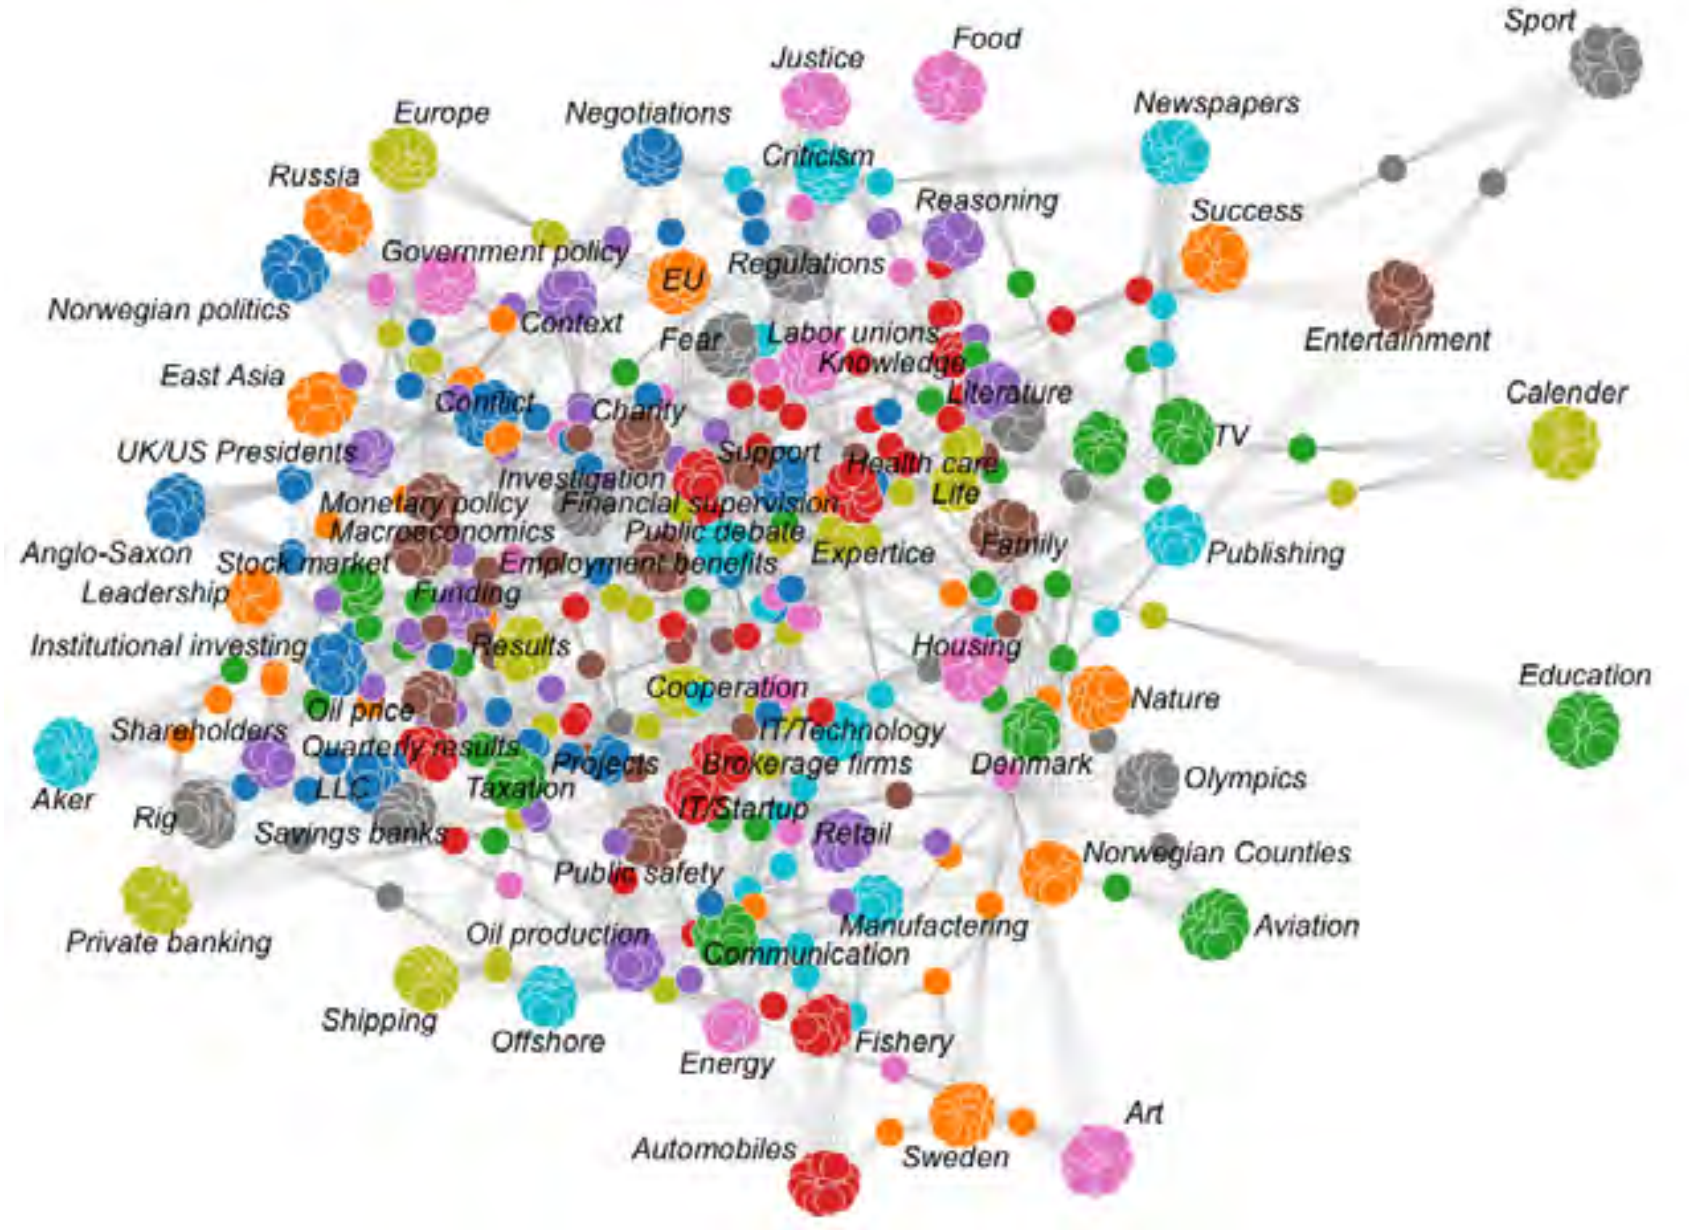

The *Funding* topic as a time series

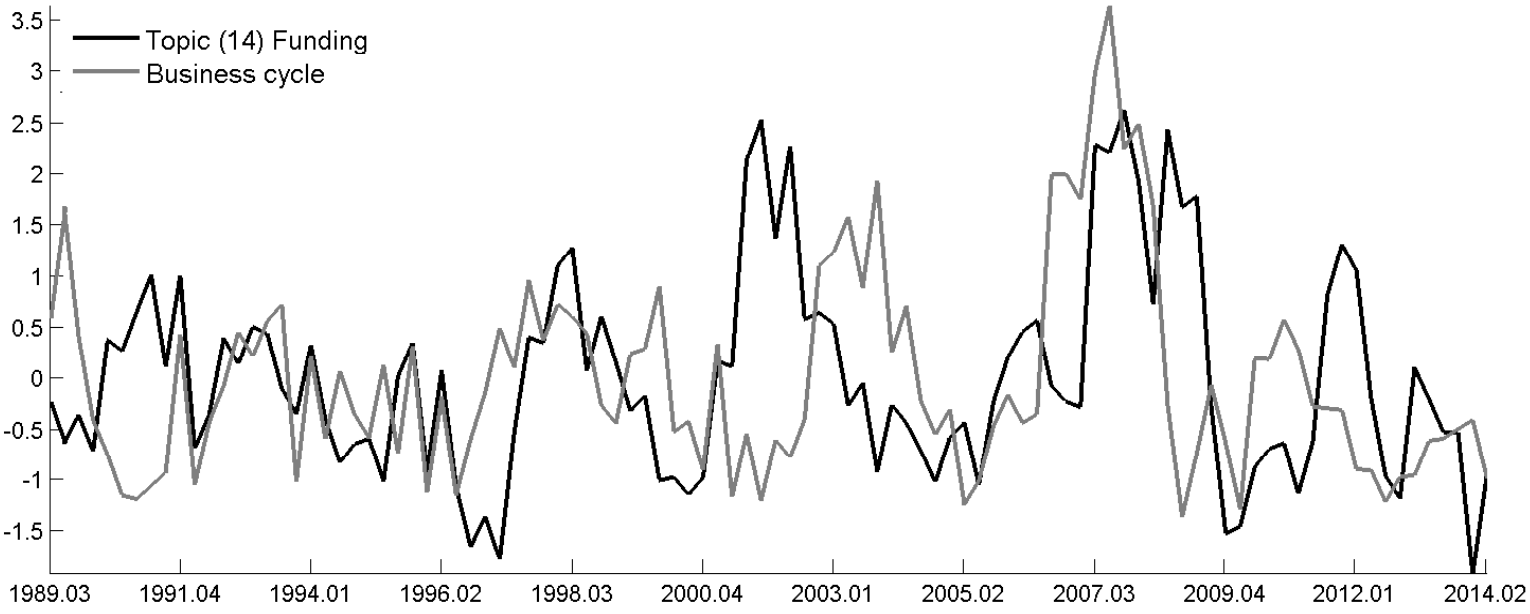

Note: Each individual topic time series is transformed to year-on-year growth rates and standardized. The Business cycle series is the standardized value of Hodrick Prescott (HP) filtered GDP, with the smoothing parameter set to 40000. Because the topics are not sign identified, the business cycle estimate is reported in absolute value.

Step 2: Identifying fundamental topics

- How do we link the frequency measures of topics to economic outcomes?
- We **run predictive regressions** of the form **ARX(p)**, where X represents the topics – one ARX for each topic – and **compare against** simple **AR(p)** specifications
- The predictive regressions are estimated as **Latent Threshold Models (LTMs)**, introduced by Nakajima and West (2013)
- Estimation and evaluation: Bayesian MCMC and posterior odds ratio...

The Latent Threshold model

$$
\begin{aligned}
y_t &= x_{t-1}' b_t + u_t & u_t &\sim N(0, \sigma_u^2) && \text{(1a)}\\
b_t &= \beta_t \varsigma_t & \varsigma_t &= I(|\beta_t| \geq d) && \text{(1b)}\\
\beta_t &= \beta_{t-1} + e_t & e_t &\sim N(0, \Sigma_e) && \text{(1c)}
\end{aligned}
$$

- The **time-varying parameter** in the LTM allows us to **identify the sign** of the topic relative to an outcome variable
- The **threshold dynamics** allow for topics to be important for parts of our estimation sample and also avoid overfitting

Model evaluation

The predictive content of a model is evaluated based on the posterior odds ratio

$$PO_{ij} = \frac{p(y|M_i)}{p(y|M_j)} = \frac{p(y|ARX)}{p(y|AR)}$$

Relative marginal likelihood – ln $PO_{ij} > 2$

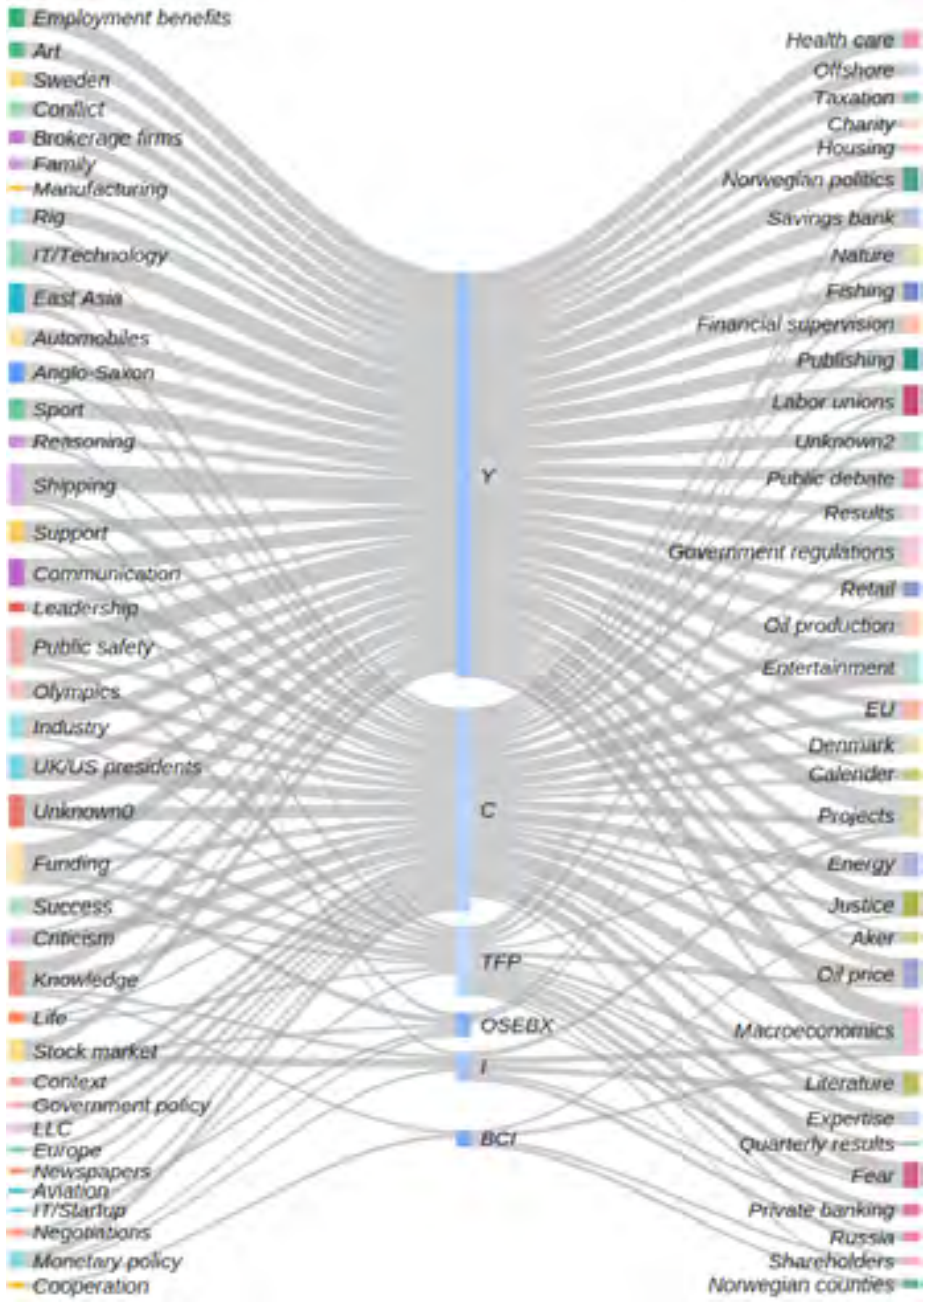

The aggregate news index

The aggregate news index is constructed as follows

$$NI_t = \sum_{i=1}^{T} w_i b_{i,t} n_{i,t-1}, \quad \text{where} \quad w_i = \frac{p(y|M_i)}{\sum_{i=1}^{T} p(y|M_i)}$$

We base our news index on the predictive results for asset prices:

- Not revised
- Forward looking, and should (according to theory) contain all fundamental information
- Commonly used

The aggregate news index cont'd

Step 3: News and business cycle fluctuations

The news view - in theory:

- Agents in the economy receive signals about future economic developments, and then face a signal extraction problem: **What is news and what is noise?**
- Changes in expectations about the future, due to **news**, leads to permanent increases in productivity, consumption, and output
- Noise, or wrongly filtered news signals, leads to a boom and bust pattern

The empirical strategy

- Specify standard **Structural Vector Autoregressions (SVARs)** (commonly applied in the literature)
- Identify shocks using a recursive ordering, i.e., Cholesky decomposition (commonly applied in the literature)
- But:
  - **Identify news shocks as unexpected innovations to our suggested news index**
  - **Identify noise shocks as unexpected innovations to asset prices (orthogonal to news innovations)**

Issues

- The nonfundamentalness problem: If economic agents can't separate between true news and noise in the data, how can the econometricians?
  - We can because the predictive filtering in Step 2 works as a signal extraction mechanism for (true) news
- The interpretation of news shocks: News shocks are anticipated fundamental shocks
  - We order the news index below productivity (fundamental) in the SVAR, i.e., news shocks can only affect productivity with a lag

Formally

- The VAR can be written as follows:

$$y_t = B_1 y_{t-1} + B_2 y_{t-2} + \cdots B_h y_{t-h} + u_t \sim N(0, \Omega) \qquad (2)$$

- We identify structural shocks using the Cholesky decomposition such that:

$$\begin{bmatrix} u_t^{TFP^a} \\ u_t^{\text{NI}} \\ u_t^{OSEBX} \\ u_t^{C} \\ u_t^{\pi} \\ u_t^{\text{R3MR}} \end{bmatrix} = \begin{bmatrix} a_{11} & 0 & 0 & 0 & 0 & 0 \\ a_{21} & a_{22} & 0 & 0 & 0 & 0 \\ a_{31} & a_{32} & a_{33} & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ a_{n1} & a_{n2} & \dots & \dots & \dots & a_{nn} \end{bmatrix} \begin{bmatrix} \vdots \\ e_t^{\text{News}} \\ e_t^{\text{Noise}} \\ \vdots \\ \vdots \\ \vdots \end{bmatrix}$$

The effects of a News shock

News shock

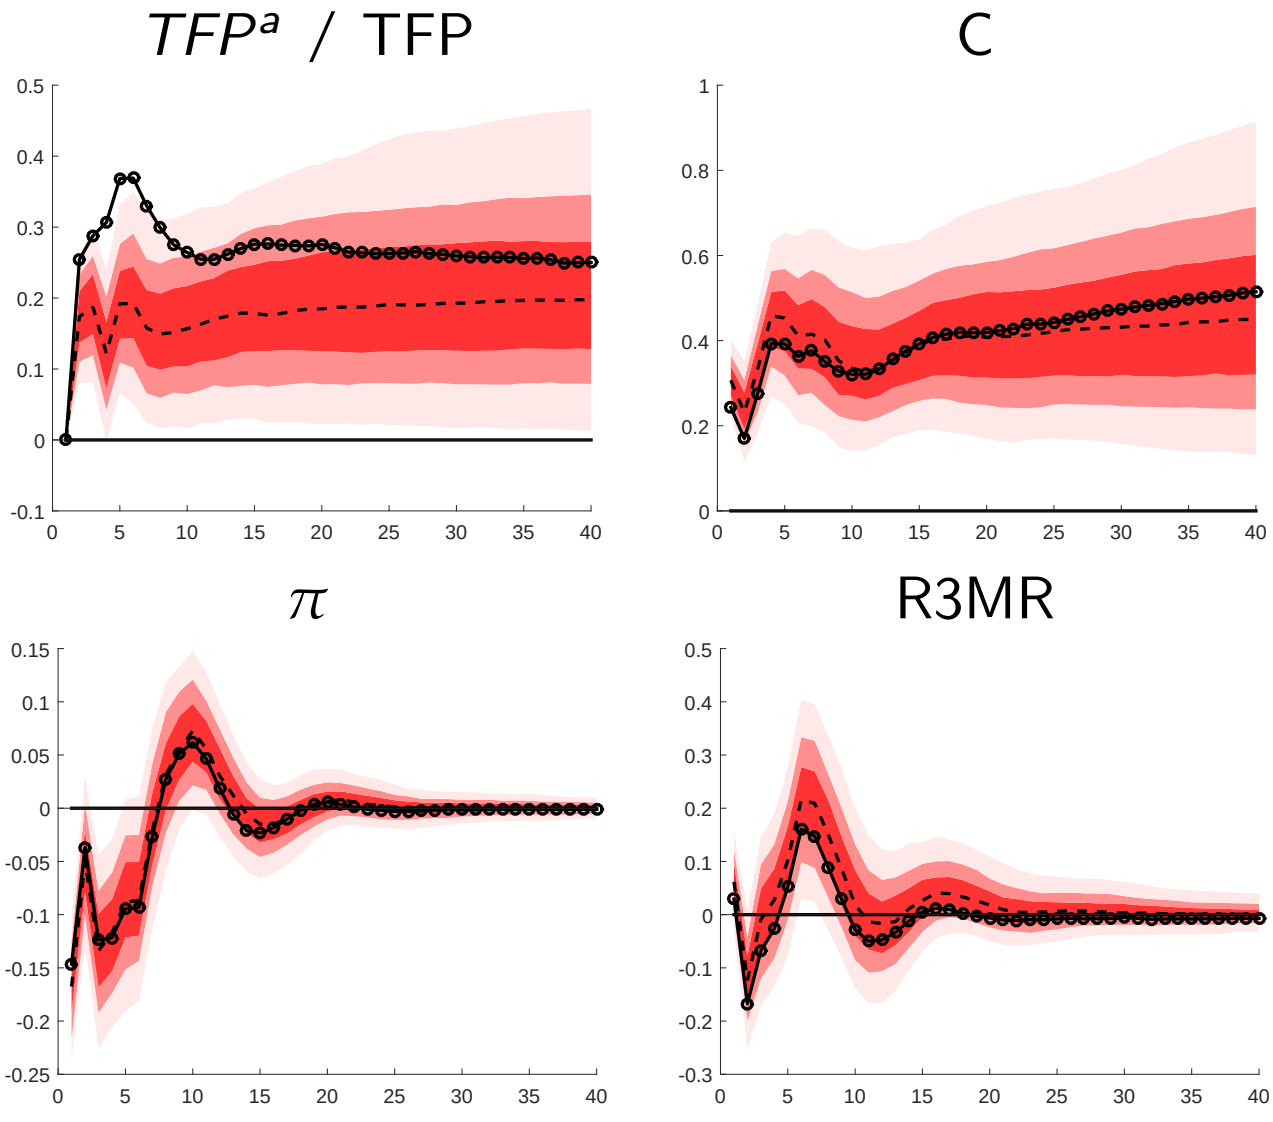

The effects of a Noise shock

Noise Shock

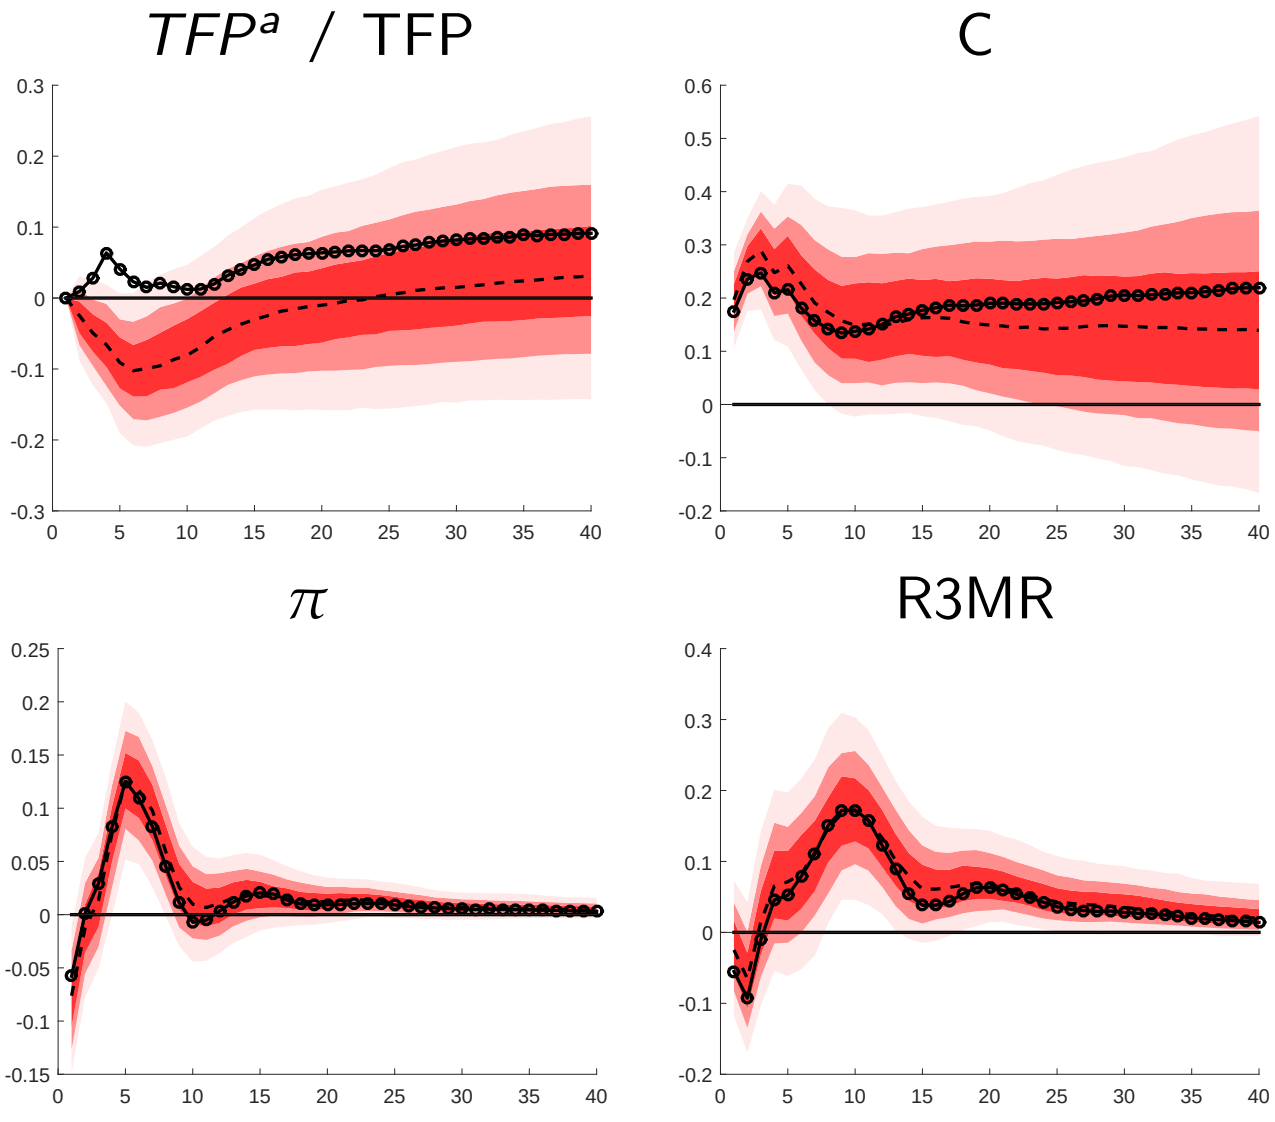

What constitutes a news shock?

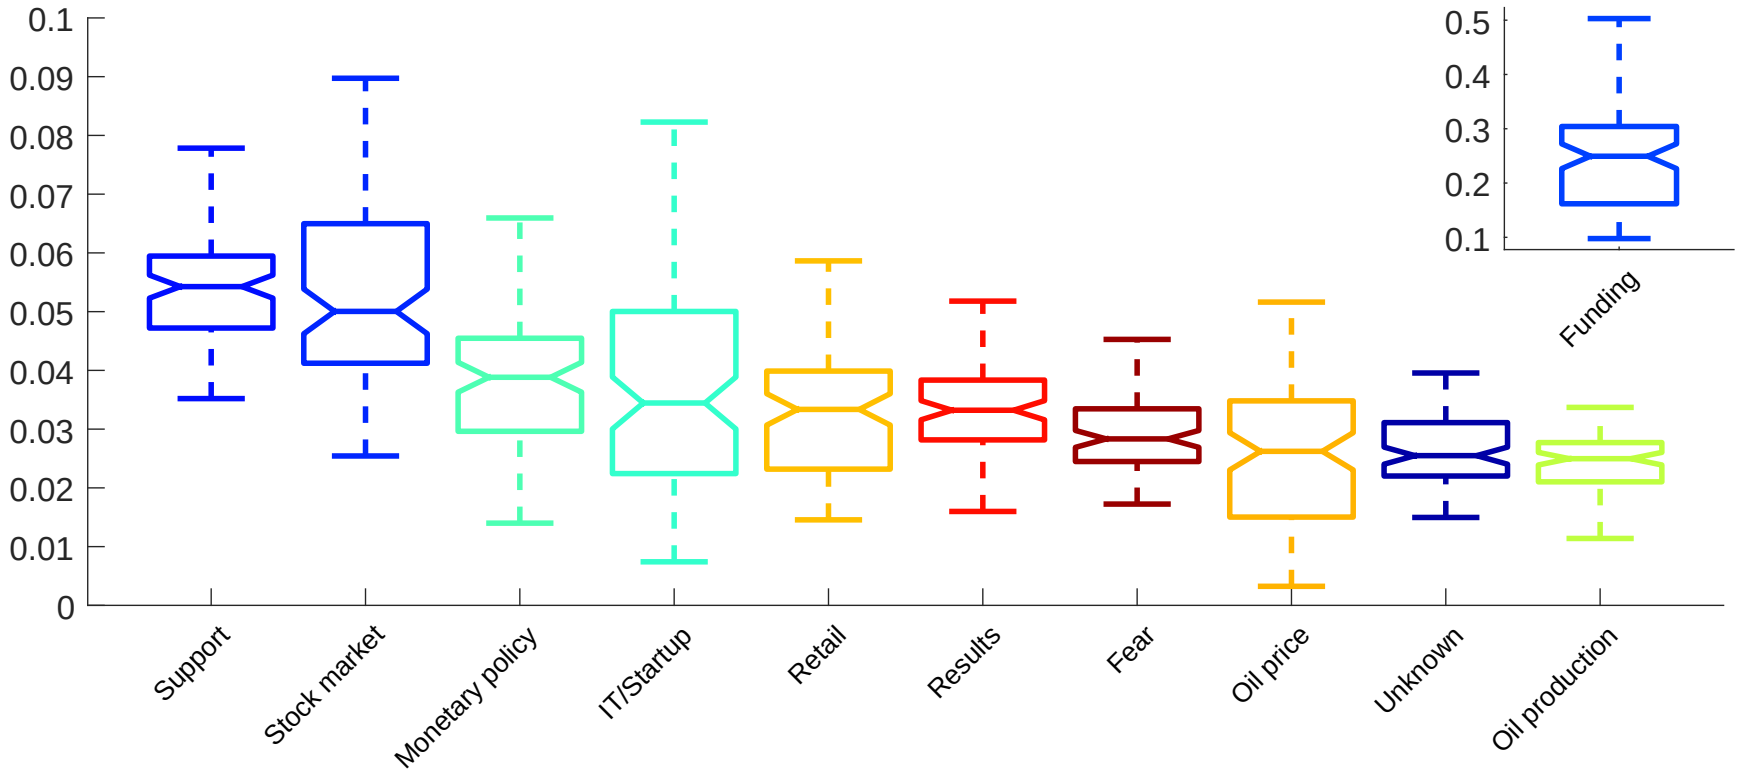

Average contribution to the aggregate news shocks

Variance decompositions

		Variable					
Shock	**Horizon**	TFP^a	**NI**	**Osebx**	**C**	π	**R**
News	1	0.00 (0.00)	1.00 (0.95)	0.50 (0.41)	0.10 (0.07)	0.06 (0.05)	0.01 (0.00)
	20	0.07 (0.13)	0.76 (0.63)	0.39 (0.30)	0.17 (0.13)	0.12 (0.11)	0.08 (0.06)
	40	0.07 (0.11)	0.73 (0.61)	0.39 (0.29)	0.13 (0.12)	0.12 (0.11)	0.07 (0.06)
Noise	1	0.00 (0.00)	0.00 (0.00)	0.48 (0.48)	0.04 (0.03)	0.01 (0.01)	0.00 (0.01)
	20	0.02 (0.02)	0.04 (0.04)	0.31 (0.31)	0.05 (0.04)	0.07 (0.06)	0.06 (0.06)
	40	0.02 (0.02)	0.04 (0.04)	0.30 (0.30)	0.03 (0.03)	0.07 (0.07)	0.06 (0.06)

Robustness/additional results

- Identifying news shocks using asset prices, as in Beaudry and Portier (2006, AER), leads to mixed results - using the news index does not
- Constructing the news index using the news topics that best predict productivity growth leads to qualitatively similar SVAR results - but news shocks explain more of the variation in productivity
- Replacing consumption (in the SVAR) with output, hours worked, or employment tells the same story

Implications

- Effects of news and noise shocks resemble those in theory models where agents face a signal extraction problem and the shocks affect productivity directly (see e.g., Barsky and Sims (2012, AER) and Blanchard et al. (2013, AER)):
  - A broad range of news topics contribute significantly to news shocks questioning the validity of such a restriction, but could be suggestive for future work on how news shocks theoretically can be an important driver of economic fluctuations
- Asset prices are almost fully explained by news and noise shocks together:
  - Using asset prices to identify news shocks likely mixes together the effects of news and noise shocks

Conclusions

We have:

- Constructed a novel measure of news and evaluated it's merits in explaining economic fluctuations
- Shown that news - from newspaper data - adds marginal predictive power for a range of economic variables
- Shown that news shocks - identified as unexpected innovation in the proposed news index - create economic responses in accordance with theory
- Provided better insights regarding what a news shock constitute
- Shown that with the news index at hand, we can also identify noise shocks

Other applications - consistent results

Predicting daily asset returns

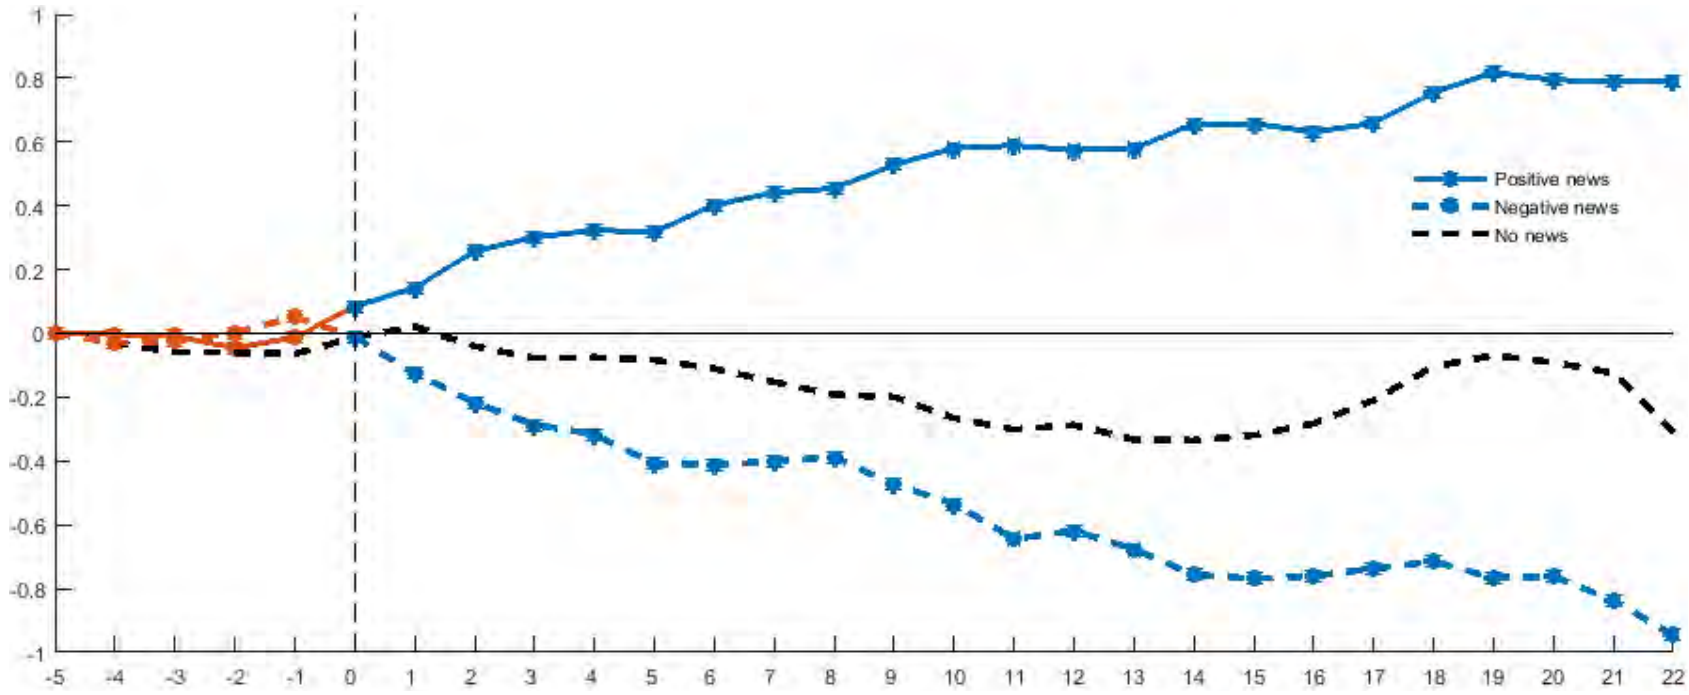

Other applications - consistent results cont'd

A daily business cycle measure