Nowcasting GDP in Real-Time with a Tone-Adjusted, Time-Varying Layered Topic Model*

Dorinth van Dijk **Jasper de Winter**

De Nederlandsche Bank
Economic Policy & Research
Econometrics & Modelling Department

EEA 2023, August 28, 2023

* Views expressed are those of the authors and do not necessarily reflect the position of De Nederlandsche Bank.

DeNederlandscheBank
EUROSYSTEEM

Introduction

Idea

- Can we use **newspaper articles** to track the business cycle and nowcast GDP growth?
- Extract **topics** from newspaper articles using unsupervised machine-learning model
- Extract **sentiment** using lexicon-based method
- Combine **topics** and **sentiment** in **tone-adjusted time-varying news topics**

Motivation

- Understand what drives business cycle fluctuations
- Nowcast current pace of economic growth to have point of departure medium term projections

Main contributions

- Analyze unique new source of Dutch newspaper texts
- Extend tone-adjusted topicmodel (LDA) with **time-variation** and **layering** in topics
- Analyze forecasting quality of tone-adjusted time-varying news topics in nowcasting model

Main idea paper. . .

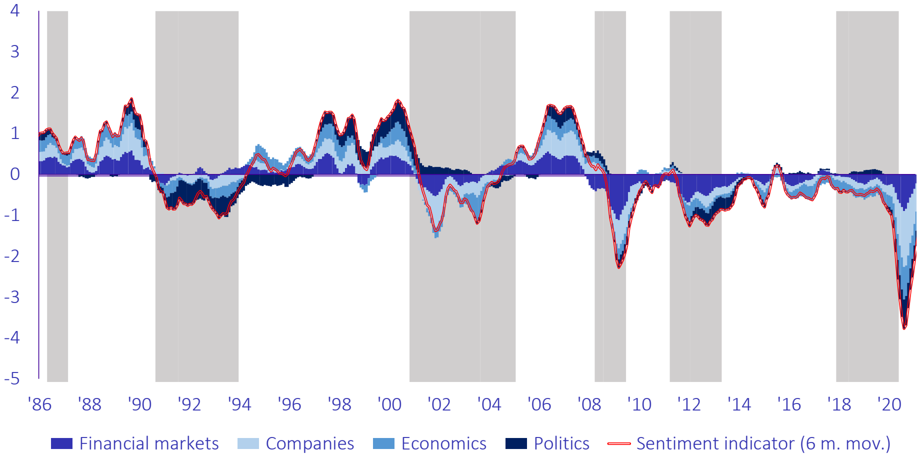

OECD recessions in grey

. . . & the four main takeaways

1. **Newspaper sentiment** is a good **coincident indicator** of the business cycle
2. **Tone-adjusted time varying layered topics** add "**story-telling**" layer to newspaper sentiment
3. **Tone-adjusted topics** embody information **not captured in other monthly indicators**, especially when nowcasting and forecasting
4. **Time-variation and layering** add (little) to forecasting power

Outline presentation

1. **Sentiment analysis**
2. **Topic Model**
3. **Nowcasting**

1. Sentiment analysis

Sentiment analysis

A short word on the data

- Complete **full-text archive** of Dutch “Financial Times” (Financieele Dagblad)
- Strong focus on financial markets, macro-economics and political-economic issues
- Analyzed period | 36 years | January 1^{st} 1985–January 1^{st} 2021
- “Standard” cleaning steps from the literature (e.g. Hansen et al., 2018 and Thorsrud, 2020)

Sentiment extraction

- Customized and extended Dutch version of **Loughran and McDonald (2011)**
- Check for **double negations** i.e.: deficit decreased, unemployment decreased
- Total list: 1,532 words | Positive: 468 | Negative: 1,063
- Sentiment score **per article** (see e.g. Tetlock, 2007 and Shapiro et al., 2020):
  (# positive words - # negative words)/(# words in article)

Sentiment analysis

(a) Positive words

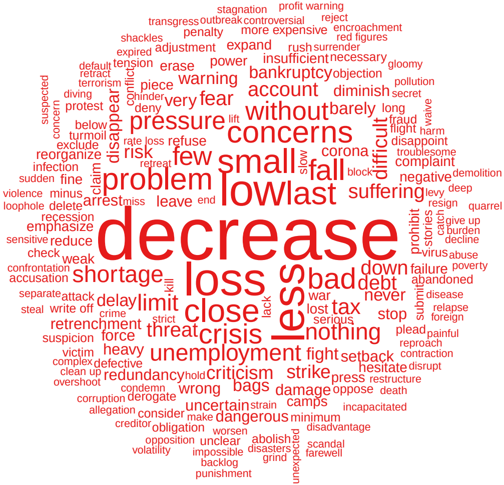

(b) Negative words

Sentiment: newspaper sentiment and GDP growth

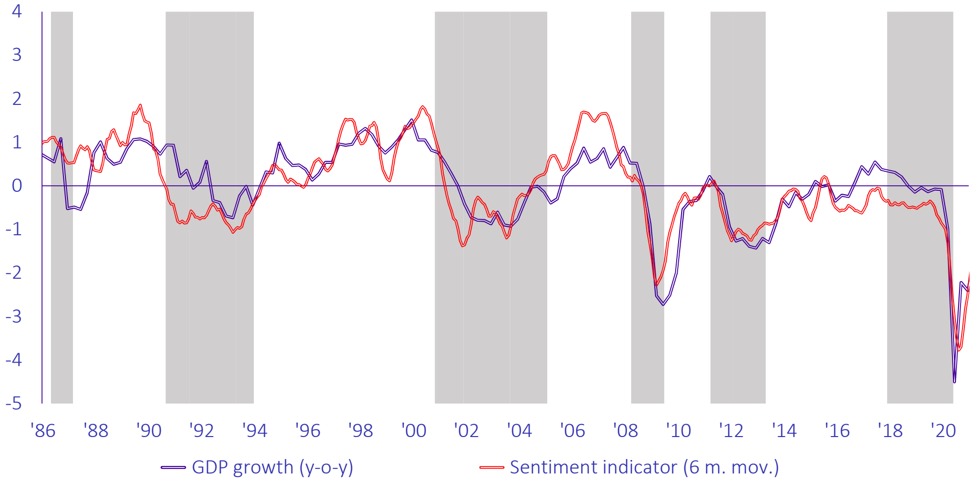

OECD recessions in grey

. . . Next part of presentation: sentiment→ per topic sentiment

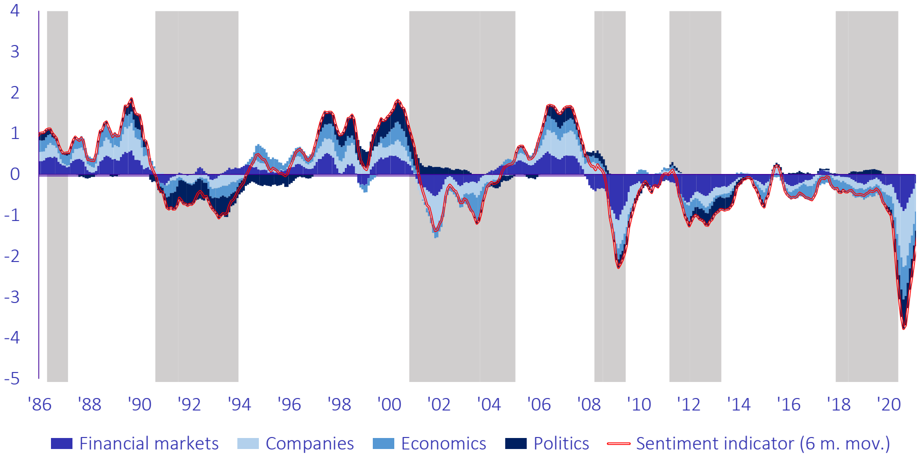

OECD recessions in grey

2. Topic Model

Model: intuition Latent Dirichlet Allocation

Model: intuition Latent Dirichlet Allocation

Step 1: Draw prob. distribution for words over topics ($\phi_{k=1/4}$)

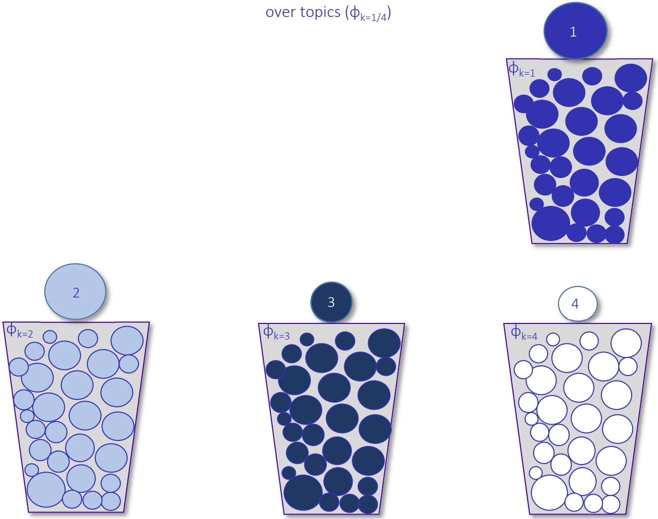

Model: intuition Latent Dirichlet Allocation

Step 2: Draw prob. distribution for topics over document for document d_1 ($\theta_{d=1}$)

Step 1: Draw prob. distribution for words over topics ($\phi_{k=1/4}$)

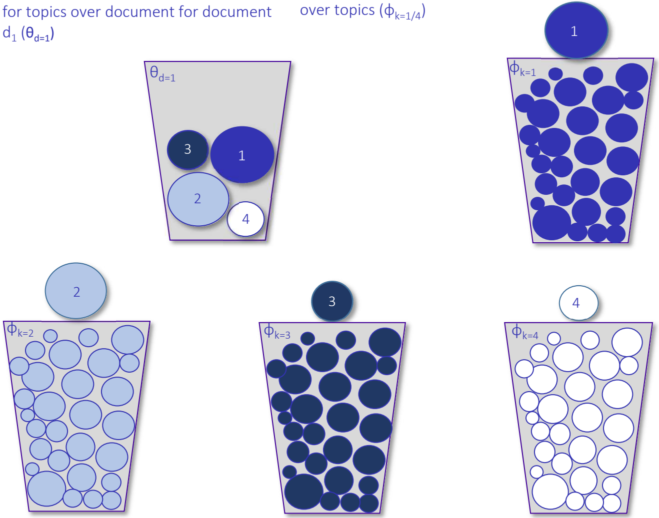

Model: intuition Latent Dirichlet Allocation

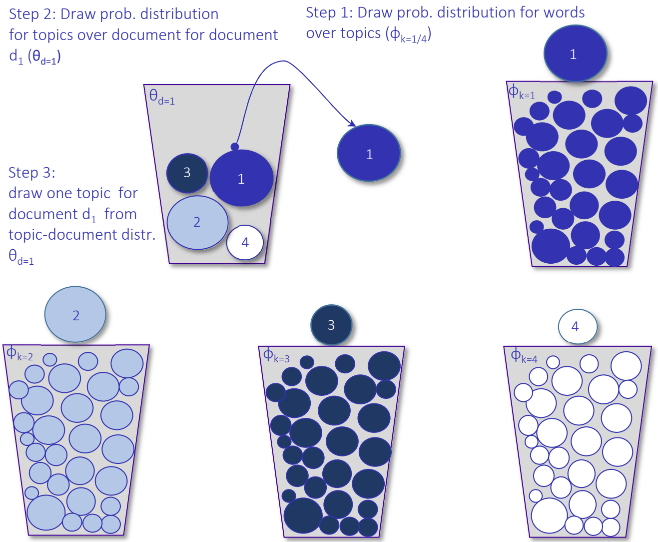

Model: intuition Latent Dirichlet Allocation

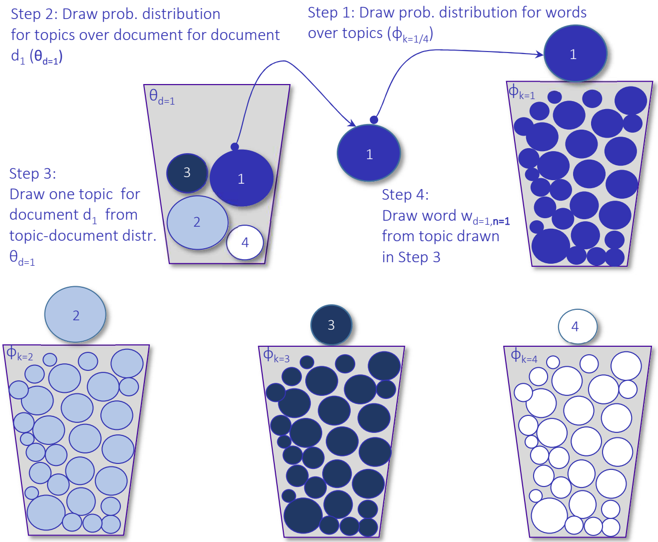

Model: intuition Latent Dirichlet Allocation

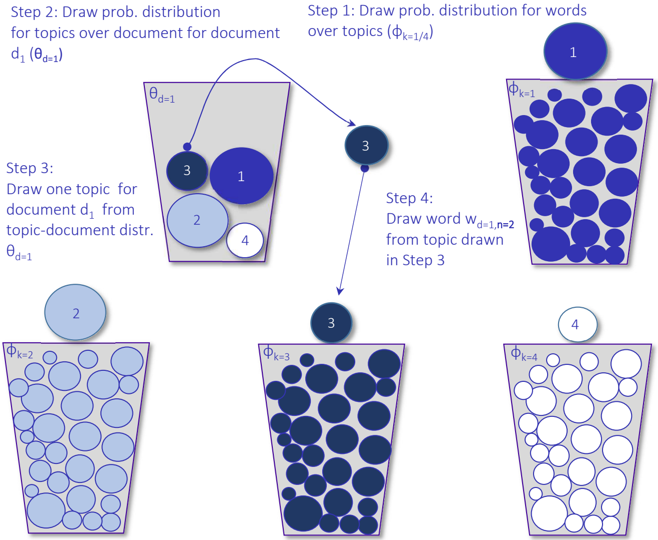

Repeat steps 3 and 4 until document d_1 is filled with N_1 words. **Repeat** Steps2/4 for d=2...d=D

Model: Inference of model parameters

Bayesian inference of model parameters

- Based on DGP just described, derive **joint distribution** of the document-topic distributions θ_d the topic-word distribution ϕ_k and the allocation of words $w_{d,n}$ to topics k in all documents.
- Bayesian inference via **Gibbs sampling** feasible but quite costly computationally, i.e.: $Pr(\boldsymbol{\phi}, \boldsymbol{\theta}, \boldsymbol{x}|\boldsymbol{w}, \alpha, \beta)$;
- **Collapsed Gibbs sampling** reduces computations to $Pr(\boldsymbol{x}|\boldsymbol{w}, \alpha, \beta)$
- **Essence Collapsed Gibbs sampler (CGS)**: With each pass of the CGS assign each word based on balance between how **likely** a word is for a topic and the **dominance** of a topic in a document based on the assignment of all **other words** to topics, i.e.:

$$Pr(x_{i,n} = K|\mathbf{x}_{-i}, w_i, d_i, .) \propto \textbf{``likeliness''} \times \textbf{``dominance''}$$

▸ Algorithm

Model: Extension layering

Intuition layering

Model: Extension layering

Intuition layering

- **Imposed hierarchy**, different from hierarchical topic models where hierarchy is based on correlations (Griffiths et al., 2003).
- In this paper:
  - Estimate 4 topics in first layer with ("**largest doll**");
  - Estimate 4 topics in second layer for **each of the 4 topic** in the first layer ("**second largest** doll");
  - Estimate 4 topics in third layer for **each of the 16 topic** in the second layer ("**third largest** doll");
  - **Estimation**: Use **words** assigned to topic predictive document-topic distribution;($\hat{\theta}_d$) of topic model in previous layer in slice t. Use random initialization using the Dirichlet priors for ($\hat{\theta}_d$) and ($\hat{\phi}_k$).

Model: Extension layering

Layer 1	Financial Markets	Firms	Economics	Politics
Layer 2	Markets	Infrastructure	Elections	Parliament
Layer3	1. Raw materials 2. Exchanges 3. International 4. Monetary Policy	17. Chemical & Pharma 18. Indices 19. Mobility 20. Company Results	33. Elections 34. Easten Europe 35. Africa & Asia 36. United States	49. Politics 50. Budgetary policy 51. Cabinets 52. Ministries
Layer 2	Financials	Multinationals	Indicators	National
Layer3	5. Corporate Finance 6. Financials 7. Banks 8. Insurance	21. Telecom 22. Customers 23. Big Tech 24. Media	37. International 38. Europe 39. Trading Partners 40. Fiscal Policy	53. Justice 54. Pensions & Healthcare 55. Supervision 56. Education & Research
Layer 2	News	Construction & Energy	Raw Materials	Lower Government
Layer3	9. Emissions 10. Takeovers 11. Trade 12. Insurers	25. Construction 26. Logistics 27. Energy 28. Industry	41. Asia 42. Oil & gas 43. Conflicts 44. Emerging Markets	57. Housing 58. Public-private 59. Agriculture & Fishery 60. Transport
Layer 2	Fin. Indices	Demography	European Union	Social Partners
Layer3	13. Stock Markets 14. Euronext 15. Analists 16. Results	29. Retail 30. Bankruptcies 31. Listed 32. International	45. Germany 46. European Union 47. Italy & Spain 48. France	61. Wage Negotiations 62. Labor Market 63. Entrepeneurs 64. Social Security & Pensions

Model: Extension time-variation

Time variation enables trends in topics

- Standard LDA models will underestimate trends because the model is estimated over the whole time period 1985 – 2020. E.g. word "Brexit" has very low "likeliness" when measured over total sample but very high "likeliness" when measured in last time-slice;
- **Difference with other dynamic topic models**: We do **not** impose word-dynamics as in dynamic topic models but let the model decide (see e.g. Blei and Lafferty, 2006 and Bitterman and Rieger, 2022 for other approaches);
- **Fixed vocabulary** of $2,153$ words. time-invariant vocabulary based on total sample, so words can slowly gain importance in topics by appearing in **time slices**;
- **15 year rolling window** first slice 1985M1–2000M1, second slice 1985M2–2000M2. Window slides each month;
- **Estimation**: Use vocabulary of topic predictive document-topic distribution ($\hat{\theta}_d$) and topic-word distribution ($\hat{\phi}_k$) in time slice t to initialize collapsed Gibbs sampler in time slice $t+1$, repeat for time slice $[t+2,\ldots,T]$;
- **Stability of topic model is crucial**: We check/correct topic stability with Cosine Distance (Newman et al., 2010 and Aletras & Stevenson, 2014).

Model: Extension time-variation

Time variation in relevance "Brexit", within topic "Economics"

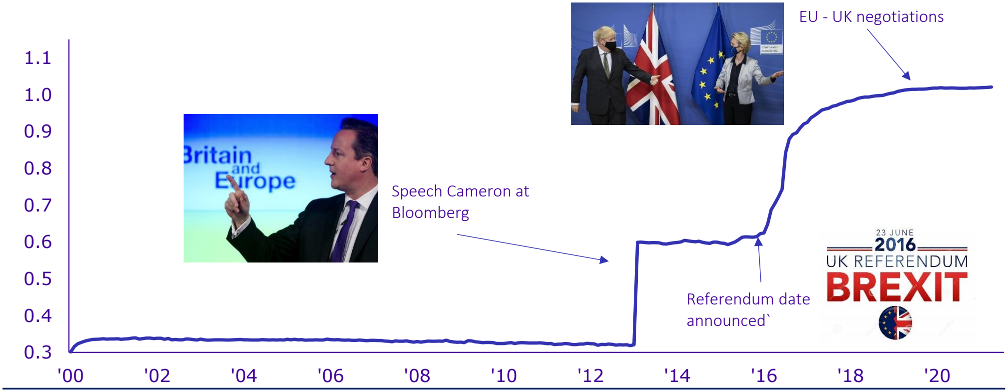

Higher score is more relevant, lambda = 0.6

Time variation in relevance top-20 words, within topic “Financial Markets”

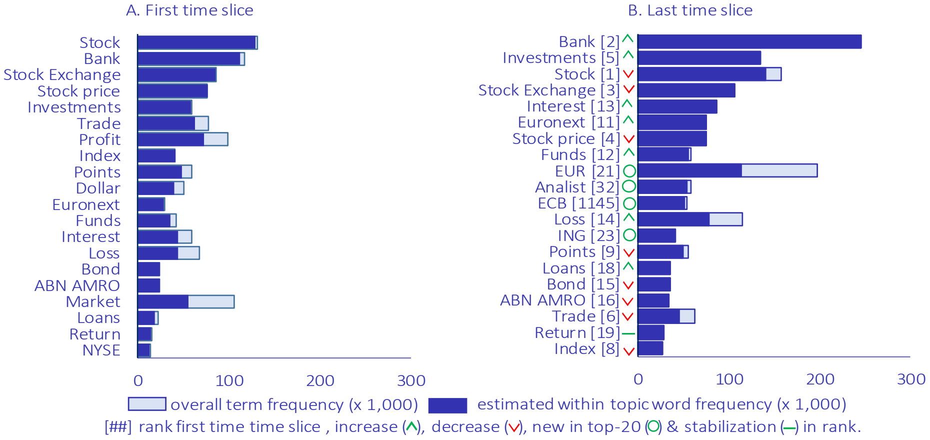

Other topics

Model: newspaper sentiment per topic

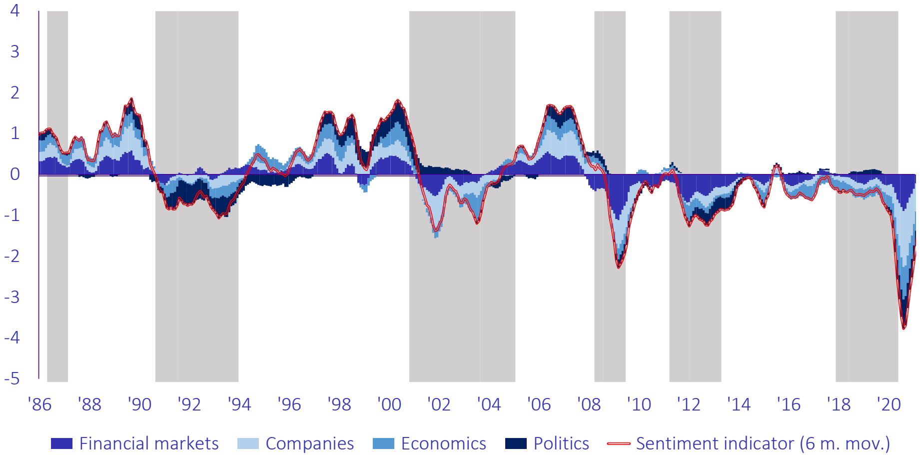

OECD recessions in grey

3. Nowcasting

Nowcasting GDP growth with(out) newspaper sentiment

Nowcasting horse-race

- **Pseudo real-time exercise**: using short-term forecasts of dynamic factor model **with(out) tone-adjusted topics** to see if newspaper sentiment adds to forecast accuracy;
- **Estimation Quasi maximum-likelihood** in a two-step procedure (Doz et al, 2012), taking into account **differences in frequencies** (GDP: quarterly, indicators: monthly) and **publication delays**. All series start in 1996M1
- **Evaluation**: 2003Q1–2020Q3 | 36 monthly macro-economic indicators | 64 newspaper sentiment indicators (1 month averages)
- Usual naming conventions: 2Q forecast, 1Q forecast, **nowcasting** (current quarter GDP forecast), **backcasting** (Previous quarter GDP forecast; GDP release 45 days after the quarter)

Outcome nowcasting exercise in a nutshell

Total sample

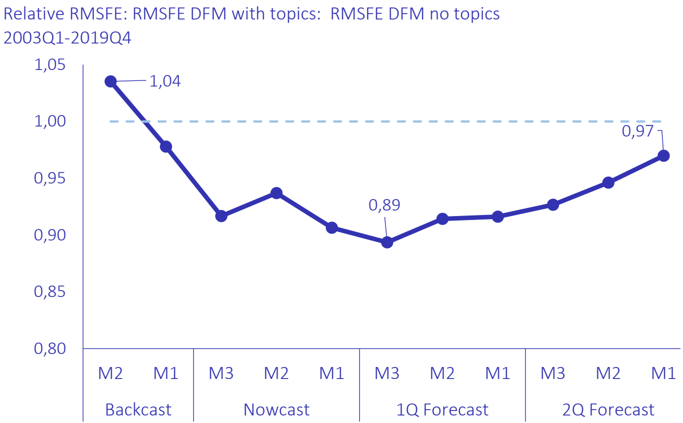

Sample excluding financial crisis

Relative RMSFE: RMSFE DFM with topics: RMSFE DFM no topics
2003Q1-2019Q4, excl. 2009Q1 and 2009Q2

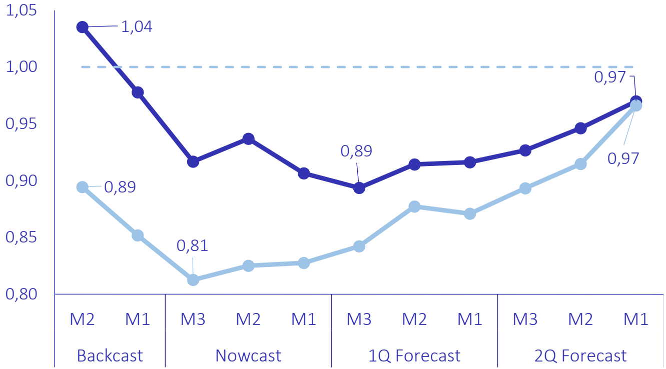

Outcome nowcasting exercise in a nutshell

Plain vanilla topic model versus layering and time- variation

LDA: plain vanilla topicmodel, L-LDA: Layered topicmodel, TVL-LDA: time varying layered topic model

topicmodel type		LDA	L-LDA	TVL-LDA
# topics per layer (X)		64	4 X 4 X 4	4 X 4 X 4
		(absolute RMSFE)	(relative RMSFE against LDA)	
Backcast	M2	*0,53*	0,98	1,00
	M1	*0,55*	0,99	0,99
Nowcast	M3	*0,57*	0,99	0,98
	M2	*0,59*	0,99	0,98
	M1	*0,62*	0,99	0,97
1Q Forecast	M3	*0,64*	1,00	0,98
	M2	*0,67*	1,01	0,99
	M1	*0,69*	1,01	0,99
2Q Forecast	M3	*0,71*	1,02	1,00
	M2	*0,71*	1,01	0,99
	M1	*0,71*	1,01	1,00

RMSFE: Root Mean Squared Forecast Error, 2003Q3-2019Q4

Outcome nowcasting exercise in a nutshell

Plain vanilla topic model versus layering and time-variation

LDA: plain vanilla topicmodel, L-LDA: Layered topicmodel, TVL-LDA: time varying layered topic model

topicmodel type		LDA	L-LDA	TVL-LDA
# topics per layer (X)		64	4 X 4 X 4	4 X 4 X 4
		(absolute RMSFE)	(relative RMSFE against LDA)	
Backcast	M2	*0,39*	0,99	0,97
	M1	*0,40*	1,00	0,97
Nowcast	M3	*0,41*	1,00	0,96
	M2	*0,44*	1,01	0,95
	M1	*0,46*	0,99	0,95
1Q Forecast	M3	*0,48*	1,00	0,96
	M2	*0,49*	1,01	0,99
	M1	*0,49*	1,03	1,00
2Q Forecast	M3	*0,50*	1,05	1,01
	M2	*0,50*	1,04	1,00
	M1	*0,51*	1,03	1,00

RMSFE: Root Mean Squared Forecast Error, 2003Q3-2019Q4

Concluding remarks

Four main takeaways from our research

1. **Newspaper sentiment** is a good **coincident indicator** of the business cycle
2. **Tone-adjusted time varying layered topics** add "**story-telling**" layer to newspaper sentiment
3. **Tone-adjusted topics** embody information **not captured in other monthly indicators**, especially when nowcasting and forecasting
4. **Time-variation and layering** add (little) to forecasting power

Thank you for your attention!

j.m.de.winter@dnb.nl

Algorithm LDA collapsed Gibbs sampler

Draws from the posterior distribution $Pr(x|w)$ are obtained by sampling from

$$Pr(x_i = K|\mathbf{x}_{-i}, w_i, d_i, .) \propto \underbrace{\frac{n^{(j)}_{-i,K} + \beta}{n^{(.)}_{-i,K} + V\beta}}_{\text{“likeliness”}} \times \underbrace{\frac{n^{(d_i)}_{-i,K} + \alpha}{n^{(d_i)}_{-i,.} + k\alpha}}_{\text{“dominance”}}$$

V	number of words in the vocabulary;
(j)	indicates w_i is equal to the jth term in the vocabulary, $j = [1 \ldots, V]$;
$n^{(j)}_{-i,K}$	freq. of the jth term assignment to topic K without the ith word;
d_i	document in the corpus to which word w_i belongs;
x_{-i}	vector of current topic membership of all words without the ith word w_i;
$(.)$	summation over index;
α	prior of the Dirichlet distribution of the topic-word distribution (ϕ_k);
β	prior of the Dirichlet distribution of the document-topic distribution (θ_d).

One draw for all words in the corpus equals one iteration of the Gibbs sampler. Based on topic-assignments you can calculate estimated predictive document-topic distribution $(\hat{\theta}_d)$ and topic-word distribution $(\hat{\phi}_k)$.

Go Back

Model: Extension time-variation

Time variation in relevance top-20 words, within topic "Economics"

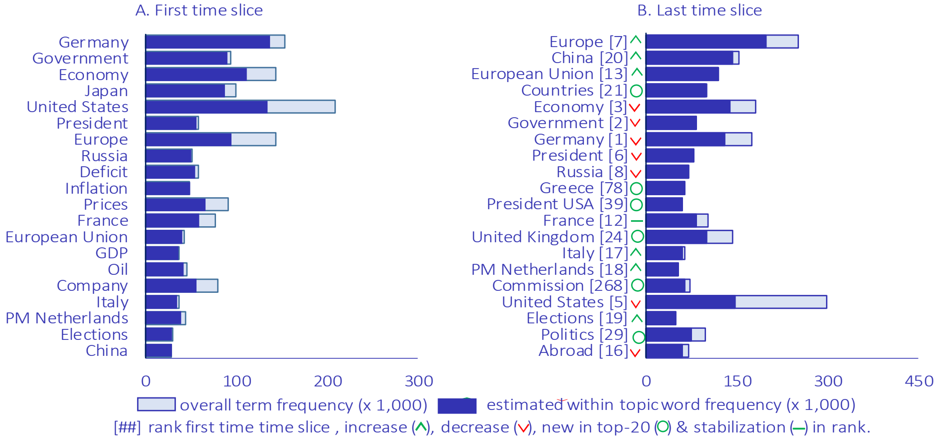

Go Back

Time variation in relevance top-20 words, within topic "Firms"

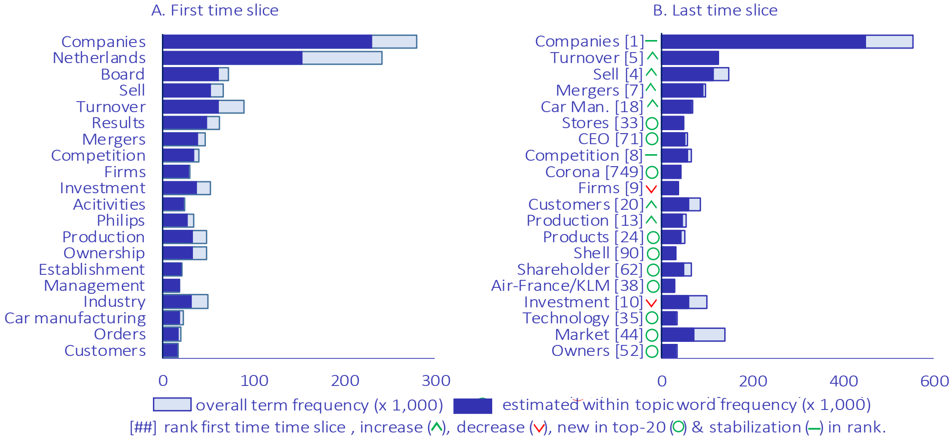

Go Back

Model: Extension time-variation

Time variation in relevance top-20 words, within topic “Politics”

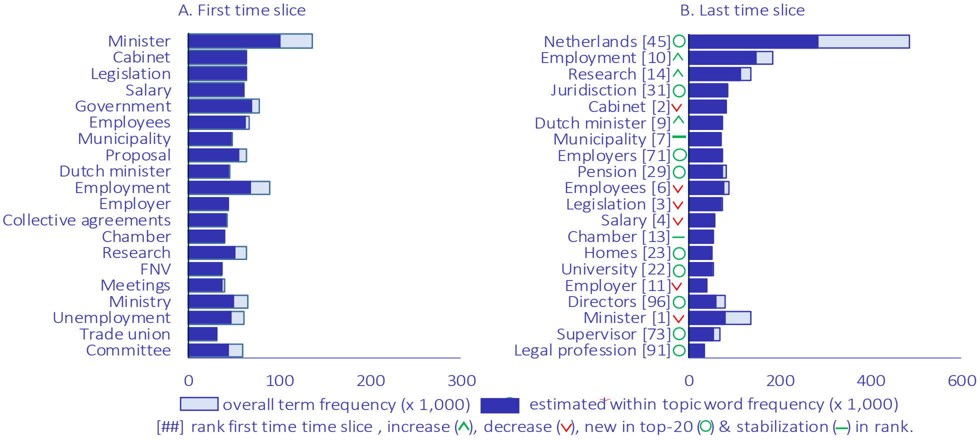

Go Back