DORTMUND CENTER
FOR DATA-BASED
MEDIA ANALYSIS

DoCMA Working Paper #3

January 2021

"For the times they are a-changin'"

Gauging Uncertainty Perception over Time

Henrik Müller, Nico Hornig and Jonas Rieger*

Cite as:
Müller, H., Hornig, N. & Rieger. J. (2021). "'For the times they are a-changin'. Gauging Uncertainty Perception over Time." *DoCMA Working Paper #3, Jan. 2021*
http://dx.doi.org/10.17877/DE290R-21878

Version 1.0, January 2021

*Prof. Dr. Henrik Müller is a professor for economic policy journalism at TU Dortmund University and Dortmund Center for data-based Media Analysis (DoCMA). Nico Hornig and Jonas Rieger are researchers at TU Dortmund University and DoCMA. The authors would like to thank Prof. Dr. Erich Schubert, Prof. Dr. Jörg Rahnenführer and Dr. Gerret von Nordheim for helpful comments.

Abstract:
This paper deals with the problem of deriving consistent time-series from newspaper content-based topic models. In the first part, we recapitulate a few our own failed attempts, in the second one, we show some results using a twin strategy, that we call *prototyping and seeding*. Given the popularity news-based indicators have assumed in econometric analyses in recent years, this seems to be a valuable exercise for researchers working on related issues.

Building on earlier writings, where we use the topic modelling approach Latent Dirichlet Allocation (LDA) to gauge economic uncertainty perception, we show the difficulties that arise when a number of one-shot LDAs, performed at different points in time, are used to produce something akin of a time-series. The models' topic structures differ considerably from computation to computation. Neither parameter variations nor the accumulation of several topics to broader categories of related content are able solve the problem of incompatibleness. It is not just the content that is added at each observation point, but the very properties of LDA itself: since it uses random initializations and conditional reassignments within the iterative process, fundamentally different models can emerge when the algorithm is executed several times, even if the data and the parameter settings are identical. To tame LDA's randomness, we apply a newish "prototyping" approach to the corpus, upon which our Uncertainty Perception Indicator (UPI) is built. Still, the outcomes vary considerably over time.

To get closer to our goal, we drop the notion that LDA models should be allowed to take various forms freely at each run. Instead, the topic structure is fixated, using a "seeding" technique that distributes incoming new data to our model's existing topic structure. This approach seems to work quite well, as our consistent and plausible results show, but it is bound to run into difficulties over time either.

Key words: uncertainty, economic policy, business cycles, Covid-19, Latent Dirichlet Allocation, Seeded LDA

1. Introduction: The Curse of Instability

In earlier writings (Müller et al. 2018), we put forward the proposition that the concept of *framing* (Entman, 1993), common in communication science, and the concept of the *narrative*, popular in economics (e.g. Shiller, 2017, 2019), were closely related. We offered a rather open definition, linking the two by suggesting that a media narrative consists of one or several media frames (as defined by Entman, 1993); plus: one or several protagonists – persons, institutions, or social groupings (nations, classes, etc.) –, whose relationships are (often) antagonistic and may change over time; plus: events, that are chronologically integrated and that are (often) assumed to constitute causal relationships. We held that "frame" was a rather static concept that applied during limited periods of time. In contrast "narrative" implied dynamic properties, i. e., the sorting of events, causes and effects over time, that explain how the current state of the world has come about, as stressed by Tenenboim-Weinblatt et al. (2016), also providing guesses about its future evolvement. Thus, the two concepts were complementary: a frame is to a narrative what a still photo is to a movie, the former showing more details and the latter providing a contextualization over time.

In terms of methodology, we showed that Latent Dirichlet Allocation (Blei et al., 2003) was a valuable tool for the measurement of media narratives. In Müller et al. (2018), we built a first version of our *Uncertainty Perception Indicator* (UPI), following the popular approach of Baker et al. (2016), but used a more refined method by pursuing a macro content analysis based on LDA. Thereby, we were able to identify certain drivers of economic uncertainty. Using data from a quarter-century of newspaper reporting, narratives could be detected that offered condensed versions of the evolving perceptions of economic policy.

Subsequent writings dug deeper into this approach. In Müller (2020), we broadened the underlying corpus of newspaper articles by including a set of leading media. Müller and Hornig (2020a) offered a taxonomy of economic uncertainty, highlighting the necessity for a more open query that refrains from pre-defining certain areas of economic policy, contrary to the approach of Baker et al. (2016). In Müller and Hornig (2020b), we further broadened the corpus, which now contains roughly three million newspaper articles, just for the German UPI.

Ever since, our aim has been to build time-series that may serve as suitable inputs into econometric models. By introducing newspaper content-based data, hitherto exogenous developments could be detected at an early stage before their full impact on economic variables could be felt; media narratives might also serve as a proxy for economic expectations. These properties have spurred hope to improve economic forecasts. However, building consistent time-series from topic models has proven to be a complex and at times frustrating undertaking.

One-shot LDAs may produce intriguingly plausible representations of public perceptions of economic developments, but are best at doing so with the benefit of hindsight. Gauging the cutting edge of economic perception on a recurrent basis is a more demanding exercise. That's what this paper is about. In the first part, we tell the story of our (failed) attempts to stabilize the UPI's LDA model over time. Extending the corpus from Q2 to Q3 of 2020 yielded non-identical topic structures. Neither changing the parameter settings to finer granulations ("zooming") nor merging topics to broader categories ("accumulating") nor LDA "prototyping" (Rieger et al. 2020a, b) could improve the results.

Referring to the film metaphor again, we weren’t able to produce a series of stills that could be assembled seamlessly to look like a movie. What we got, were pictures that looked somewhat familiar, but differed too much to lend themselves to the production of consistent time-series. The one-shot LDAs were good enough for large scale content analysis, but random in too many aspects. *The times they are a-changing*, as Nobel laureate Bob Dylan has taught the world – and, unfortunately, so are the topic models built to capture the evolving *zeitgeist.*

In the second part of this paper, we apply a different approach, a “seeded” LDA, that is set to reproduce a given topic structure at subsequent observation points. For now, this works pretty well on a quarter-to-quarter basis, but it is bound to yield erratic shifts over time, once the corpus’ topic structure changes to a degree that it clashes with the model’s stationarity. Therefore, a more flexible approach is needed; in the conclusions we sketch possible methodological enhancements we plan to pursue in future versions of our approach.

2. A series of one-shot LDAs does *not* produce a consistent time-series! A brief account of failed attempts

In Müller and Hornig (2020b), we found a somewhat different topic structure in Q2 than in Q1 (Müller and Hornig 2020a), even though the vast majority of topics were closely related, as measured by Hellinger distance. We attributed these changes in the model’s structure to the enlarged corpus, now including left-of-center *Süddeutsche Zeitung*, and to the extreme uncertainty events due to the Corona crisis, that were unfolding to full extend in Q2 of 2020. Furthermore, we found that the appropriate value of the LDA parameter *K*, i.e. the number of topics LDA is set to produce, increased with the bigger corpus. While in Q1 a value of *K*=12 yielded the most coherent results, in Q2 it was *K*=14, which came as little surprise, since the inclusion of *Süddeutsche* added a whole range of topics and angles that were more or less absent in the corpus’ other two papers, *Handelsblatt* and *Die Welt*. After all, *Süddeutsche*, Germany’s biggest broadsheet with a national circulation, added about 50 per cent to the combined three-newspaper corpus.

For this paper, we augmented the corpus by including the articles published in the three newspapers in Q3 of 2020. The overall share of texts related to economic uncertainty fell in July and August, but showed a slight uptick in September (fig. 1), as Corona infection rates were on the rise again.

Figure 1: UPI for Germany, Jan. 2000 to Sept. 2020 (monthly data), selected events*

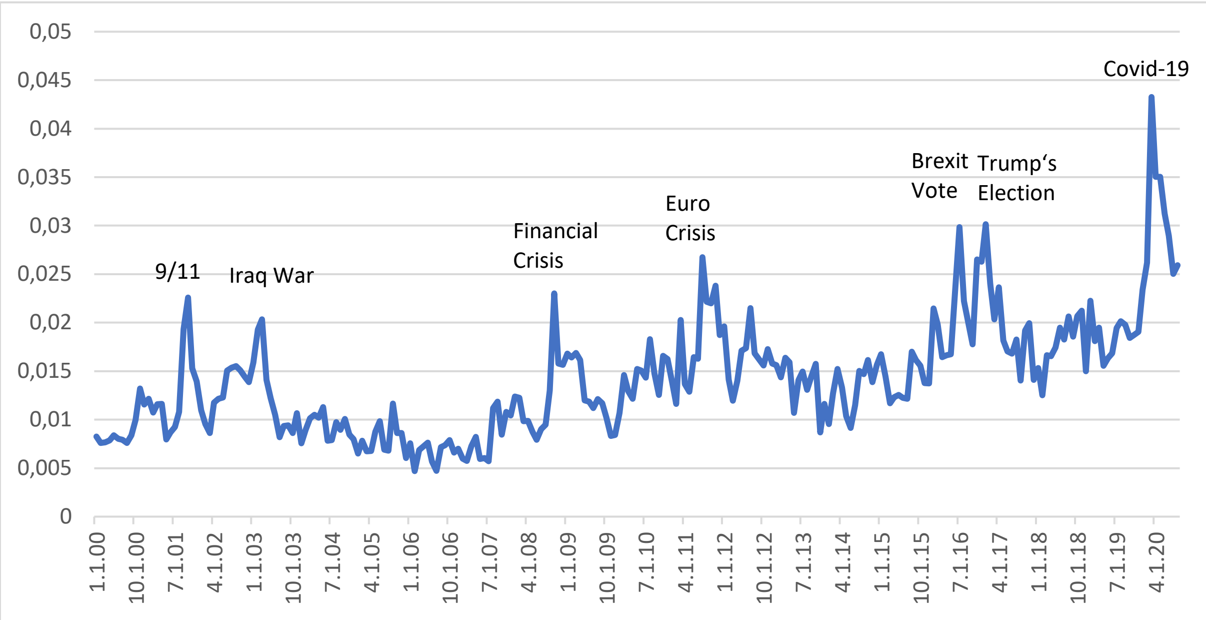

*analysis corpus share of entire corpus

Running the LDA algorithm on this corpus, all settings were the same as in the Q2 model. Unfortunately, the exercise produced different topics, and these were not as coherent as the ones found in Q2, that is, they could not be labelled quite as distinctly. To heal these insufficiencies, we tried three different strategies: zooming, accumulating, and prototyping.

Zooming

Due to the extreme rise of overall economic uncertainty in the first half of 2020, it seemed possible that the thematic distribution of topics in the corpus had changed. In order to capture these potential changes, we re-ran the algorithm with higher values of *K* (16, 18, 20) for three different corpora (all including *Handelsblatt*, *Süddeutsche,* and *Die Welt*), ending in Q1, Q2, and Q3 respectively. The finer granulation should make the models more sensitive to changes in the underlying structure, zooming into the subtler details of the corpus. A preliminary eyeballing analysis singled out *K*=18 as the most suitable parameter setting. Labelling the topics of each of the three models provided the following findings:

*Table 1: Topics of one-shot LDAs in Q1-Q3 2020 (K=18)**

Q1	Q2	Q3
Central Banks (13)	Central Banks (7)	Central Banks (11)
German Economy (9)	German Economy (6)	German Economy (8)
German Politics (10)	German Politics (15)	German Politics (18)
Energy & Climate Change Mitigation (5)	Energy & Climate Change Mitigation (1)	Energy & Climate Change Mitigation (15)
Big Business I (2)	Business II (12)	Big Business (1)
Financial Markets I (1)	Financial Markets I (11)	Financial Markets I (2)
Financial Markets II (focus on retail investors) (11)	Financial Markets II (focus on retail investors) (13)	Financial Markets II (focus on retail investors) (16)
Big Business II (focus on car industry) (17)	Big Business I (8)	Digitalization (3)
Miscellaneous (1)	Social Policy (Pensions Labor market regulation etc.) (2)	Energy (4)
Legal Risks (3)	≈ Society (3)	Geopolitics (USA) (5)
Miscellaneous (health, sports, nutrition) (6)	Euro Zone Tensions (4)	Miscellaneous (Arts, Culture, Sports) (6)
Geopolitics (7)	Miscellaneous (5)	Eurozone Tensions (focus on banks) (7)
Society (8)	Digitalization (9)	EU Conflicts I (9)
Social Policy (Pensions Labor market regulation etc.) (12)	Miscellaneous (health, student life, sports, nutrition…) (10)	EU Conflicts II (10)
Trade (focus on China) (14)	Trade (focus on China) (14)	Legal Risks (12)
Digitalization (15)	Legal Risks (16)	Society I (Health, Migration, Corona) (13)
Eurozone Tensions (16)	Geopolitics (17)	Emerging Market Troubles (14)
Banks (18)	EU Conflicts (18)	Society II (17)

*topics occurring in all LDAs are marked yellow; connections show content-related similarity; topic number in parentheses

Table 2: Top words of the Topic "German Economy" of one-shot LDAs in Q1-Q3 2020

Q1	**Q2**	**Q3**
jahr	jahr	Jahr
wachstum	wachstum	Wachstum
wirtschaft	wirtschaft	Wirtschaft
quartal	quartal	Quartal
konjunktur	deutschen	Unternehmen
unternehmen	unternehmen	Konjunktur
deutschen	konjunktur	Deutschland
deutschland	deutschland	Deutschen
deutsche	deutsche	Rezession
rezession	rueckgang	Rueckgang
rueckgang	entwicklung	Aufschwung
erwartet	rezession	Oekonomen
entwicklung	industrie	Investitionen
aufschwung	aufschwung	Entwicklung
investitionen	oekonomen	Deutsche
oekonomen	Befragten	Industrie
industrie	Verbraucher	Erwartet
verbraucher	Erwartet	Prognose
konsum	Prognose	Bip
prognose	Investitionen	Arbeitsmarkt

While certain similarly labelled topics, sometimes split into two specific aspects, could be found in all of the three models (highlighted), even these topics were not completely identical. A case in point is “German Economy”, a topic that we’ve found in each and every uncertainty-related model we’ve run so far. It consists of articles about business cycle developments, forecasts, and early indicators; using the highly standardized language of national accounts, these topics tend to be pretty consistent over time. This is illustrated by table 2 showing the top 20 words of the “German Economy” topics. While the top four words are identical in each model, the rest differs mainly in terms of order.

The effect of these differences shows up in slightly varying frequency patterns (fig. 2).

*Figure 2: Topic "German Economy", Jan. 2000 to Sept. 2020 (monthly data)**

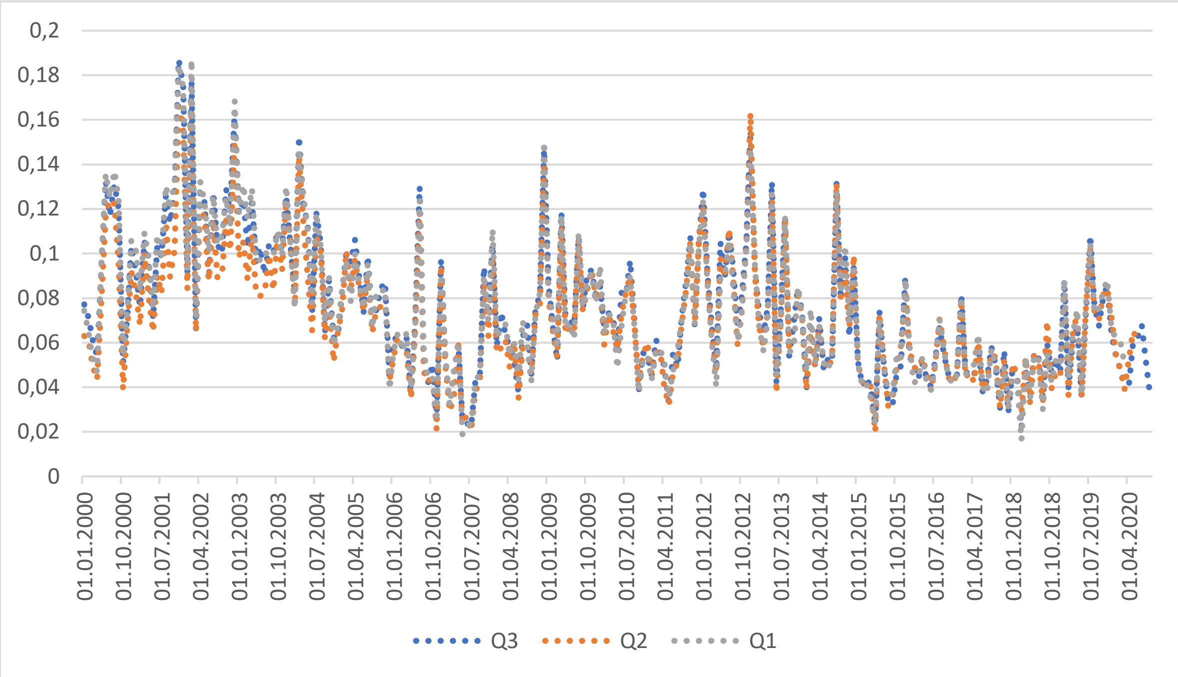

*shares in analysis corpuses

Accumulating

Could we overcome the problem of fuzziness if we combined the individual topics to broader categories? The rationale behind this exercise goes as follows: Even though the algorithm sorts certain events and themes into different categories in different models, these variations may be neutralized, if we combined these topics. In earlier versions of the UPI, we had already experimented with *Uncertainty Factors*; *UPI Politics* sums up domestic, European, and international developments in the realms of government, parliament, central bank, the judicial system, and civil society; *UPI Real Economy* consists of all the topics related to companies, sectors, and the business cycle; *UPI Financial Markets* captured uncertainty related to banks and bourses. Yet, again we found rather different patterns over time.

The patterns produced by the three topic models in Q1 to Q3 respectively were not entirely disparate; all of them show a rise of political uncertainty in the decade from 2008 (effects of the aftermath of the global financial crisis), followed by a rise of uncertainty in the real economy towards the end of the 2010s. But these patterns are far from identical either.

Prototyping

To be sure, LDA yields somewhat different models with each run, due to the probabilistic nature of the algorithm. Even using an exactly similar corpus and identical parameter settings, results are bound to differ, at least slightly. To overcome this problem, Rieger et al. (2020) have suggested a procedure they call "prototyping" (see next section for details), that yields a sort of mean LDA after a large number of re-runs (typically around 100). But applying this technique to our K=18 models of the 2020 Q1-Q3 corpora does not produce models consistent over time either. Even as we prototype the LDAs at each observation point (2020 Q1, Q2, Q3), the results differ considerably.

We conclude that a succession of one-shot LDAs is unable to result in a sufficiently consistent time-series. This outcome is highly unsatisfactory considering the purposes we have in mind for our uncertainty indicator. Imagine that changes in, say, consumer price levels were based on different baskets of goods each month, or that gross domestic product (GDP) was calculated using different definitions each quarter, statisticians would be unable to measure inflation or growth. Therefore, a different approach is needed to move closer to our goal of producing consistent time-series for text-based indicators such as the UPI.

3. Methodology: prototyping and seeding

As already mentioned, LDA has the far-reaching disadvantage, that random initialization and conditional reassignments within the iterative process of the Gibbs sampler (Griffiths and Steyvers, 2004) can result in fundamentally different models when executed several times on the same data and with identical parameter sets. This fact greatly limits the scientific reproducibility.

Up to now, the so-called eye-balling method has been used in practice to select suitable results. There are solutions to average LDA runs (Nguyen et al., 2014), which has the disadvantage that no unique assignments can be obtained for individual tokens of the documents. A different method of objective and automated selection is perplexity optimization (Grün and Hornik, 2011). However, Chang et al. (2009) were able to show that selection mechanisms aiming for optimizing likelihood-based measures do not correspond to the human perception of a well-adapted model of text data. Instead, the authors propose a so-called intruder procedure based on human coding. The corresponding methodology is implemented in the software package tosca (Koppers et al., 2020).

The selection method LDAPrototype (Rieger et al., 2020a) determines a prototypical LDA by automated selection from a set of LDAs. The method improves the reliability of findings drawn from LDA results (Rieger et al., 2020b), which is achieved following a typical statistical approach. For a given combination of parameters, a number of models is calculated (usually about 100), from which that model is determined that is most similar to all other models in the set . For this purpose, pairwise model similarities are calculated using the S-CLOP measure (Similarity of Multiple Sets by Clustering with Local Pruning), which can be determined by a clustering procedure of the individual topic units based on topic similarities of the two LDA results considered. The methodology is implemented in the corresponding R package ldaPrototype (Rieger, 2020).

In addition to the LDAPrototype for initial estimates of the model, we use an implementation of LDA that uses the prototype result as initialization for subsequent quarters (https://github.com/JonasRieger/ldaGibbs). Therefore, we modify an existing implementation of LDA (Chang, 2015) by iterating the Collapsed Gibbs Sampler only over all the new data. This means that the topic assignments of all previously modeled articles remain stable and we obtain assignments for all new articles to the existing topics. This process of fitting new data to a predefined topic model is known as “seeding”.

4. Results

As a first step, we compute a LDAPrototype for the data up to 2020 Q2. Similar to the approach in Müller and Hornig (2020b), we use the parameter value K=14. The resulting model contains many of the familiar topics already presented in earlier writings (Müller and Hornig, 2020a, 2020b). "Big Business", "Financial Markets", "German Economy", "Central Banks", "German Politics", "Legal Risks", and "Society" appear to be somewhat stable topics in different LDAs conducted on continued data. Table 3 provides an overview of the model's topics.

*Table 3: Overview of Topics and Labels (LDAPrototype, Q2-2020, K=14)**

Topic No.	Label	Share in analysis corpus (per cent)	Content	Part of Uncertainty Factor…
1	Big Business	6.2	Developments at quoted international corporates	UPI Real Economy
2	Society	12.4	Debates on capitalism, globalization, democracy, populism, immigration, national identity	UPI Politics
3	German Politic	6.5	Political developments in Germany (national level)	UPI Politics
4	Geopolitics I	5.7	Geopolitical tensions	UPI Politics
5	Central banks	8.4	ECB, Fed etc. actions against crises	UPI Politics
6	Geopolitics II	4.8	Geopolitical economic tensions, trade, energy, resources, EMs (China)	UPI Politics
7	Legal Risks	6.2	Regulations and court rulings affecting businesses	UPI Politics
8	Energy & Climate Change Mitigation	3.9	Energy market developments, transition to sustainables etc., Fukushima disaster (2011) as focal event	UPI Real Economy
9	Human Resources	5.1	Education and knowledge, career issues, Corona implications, immigration, and skills market	UPI Real Economy
10	Companies & Markets	8.2	German corporates in trouble	UPI Real Economy
11	German Economy	7.9	Business cycle developments, forecasts, reports on survey data	UPI Real Economy
12	Financial Markets	7.5	Up and down at the bourses and how to invest	UPI Financial Markets
13	EU Conflicts	6.2	Brexit, Greece debt etc.	UPI Politics
14	Miscellaneous	11.2	Diverse	–

*we combine 3 and 7, 4 and 6, 1 and 10, due to their proximity.

However, the prototyping leads to the re-assignment of some issues compared to the models presented in our earlier papers. Most notably, this involves global energy and resources problems, that are part of "Energy and Climate Change Mitigation" (ECCM) in the older models; now, the LDAPrototype sorts them into an Emerging Markets-focused topic, that includes trade and other economic tensions as well as developments in markets for oil and other raw materials. China is one of the top protagonists in the topic. We label it "Geopolitics II" and combine it with "Geopolitics I", that focusses on strategic and military conflicts. As the two topics move fairly closely in tandem (fig. 3), driven by the same events, this procedure seems justified.

*Figure 3: Geopolitical Topics before merging**

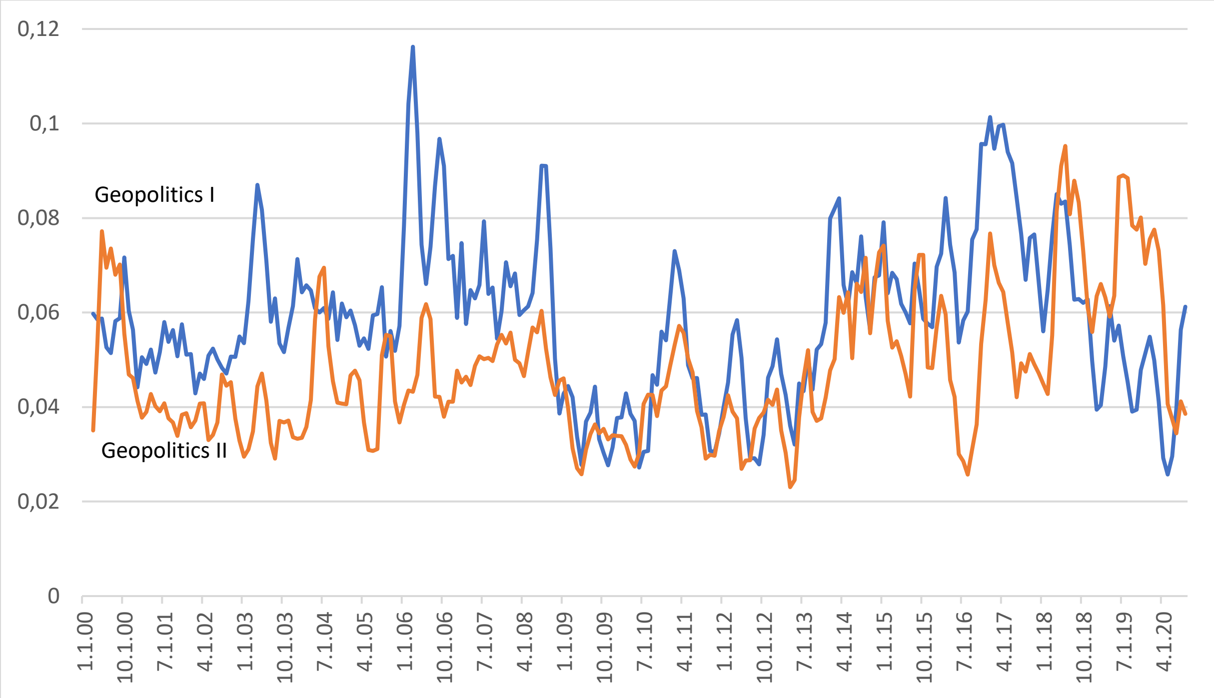

*three-month moving average

Traces of the Corona pandemic

A second difference involves the *Corona pandemic*, a new issue, unknown for most of the period covered, that our earlier models attached to "Central Banks" (in the Q1 model) and to "ECCM" (in the Q2 model) respectively. Now, the LDAPrototype attaches the Corona issue to a whole range of topics. Most pronouncedly, it appears in "Human Resources", where three of the 100 top-words are related to the pandemic. Among the topic's top articles, rank 29 is a piece from *Süddeutsche Zeitung*, published August 4 2020, that mentions Corona in a commentary, the headline reading *Bessere Absicherung für Gründer!* ("Better protection for founders!"), the text stressing the vulnerability of the self-employed during the Corona pandemic and the need for more government support for this group. Traces of the pandemic can also be found in other topics: in "Central Banks", one of the top 100 texts (in *Die Welt* on March 5 2020) about monetary policy decisions by the US Federal Reserve mentions Corona and the ensuing risks to the economic outlook ("Fed chief Jerome Powell explained the move, mentioning the economic risks that the spread of the new Corona virus represented"). Among the topics' top 100 articles, there are two mentions in "Legal Risks", one concerning the risk of unpaid rents for commercial property estate and another dealing with the possibility that residential property owners' annual assemblies cannot be held due to the pandemic. In "Companies & Markets", there are three mentions of Corona, in the context of VW, BMW, and SAP respectively. Two mentions about pandemic-related losses for stock markets and fixed income funds can be found in "Financial Markets". This is an interesting finding: the seeded LDAPrototype does not produce a distinct "Corona" topic, but traces of it can be found across the model.

General Patterns of Uncertainty Factors

The long-term pattern of economic uncertainty provides the familiar picture of a rise of relative uncertainty after the financial crisis of 2008, culminating in the advent of populist politics in 2016 and the following animosities disturbing the world trading system. After 2018, these political frictions start adding to real economic uncertainty. Over the long time-horizon of our analysis, the overall uncertainty shock due to the Corona virus is unique in its size (fig. 1), but the associated composition of uncertainty is not that extraordinary (fig. 4). While the *UPI Real Economy* is at its highest since 2000, in early 2020 the pandemic-induced rise is comparable to the shifts that occurred during the financial crisis and the effects of the trade war starting in 2018.

Figure 4: Uncertainty Factors – UPI Politics vs. UPI Real Economy and associated events*

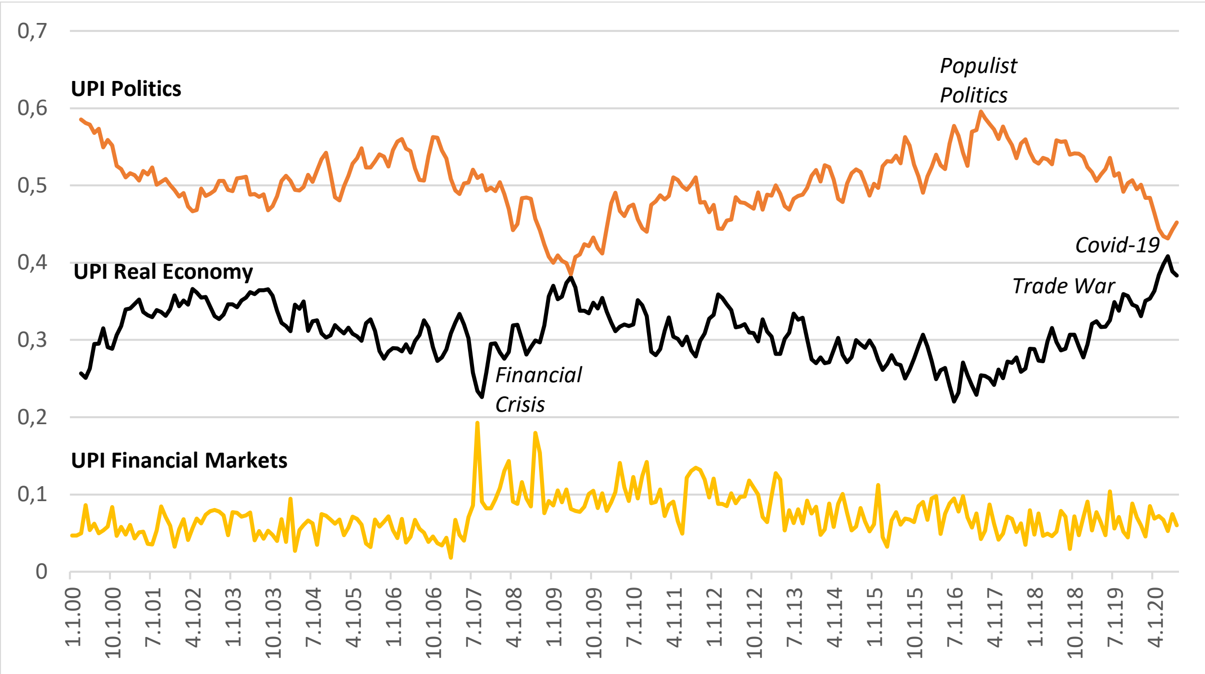

*share of analysis corpus relative to the entire corpus. Three-months moving averages. For the construction of the three categories see table 2.

International Drivers of Political Uncertainty

As in earlier exercises, political and real-economic uncertainty factors can be decomposed to show the developments of individual topics. Figures 5a and 5b display the components of the *UPI Politics.* Let's start with the international dimension of political uncertainty. In the early 2000s, there is plenty of uncertainty arising from central banks' actions; since we are dealing with German data these elevated levels can be attributed to the then-newish ECB, whose course of actions the public found hard to predict in its early years. From the second half of the 2000s onwards the *UPI Politics* is dominated by geopolitical tension, with escalating conflicts in the Middle East, the Financial Crisis, Russia's annexation of Crimea, Trump's election, and the following trade war as the most prominent events. Intra-EU conflicts that were largely absent in the 2000s – the sole exception being the rejection of the EU constitution by referendums in France and the Netherlands – gained prominence in the 2010s, that constituted a dismal decade for Europe. From the Euro Crisis to the Brexit vote, that settled Britain's subsequent departure from the bloc, the EU was a troubled and thus uncertainty-spreading environment for the Germany economy.

Figure 5a: UPI Politics, international topics and associated events*

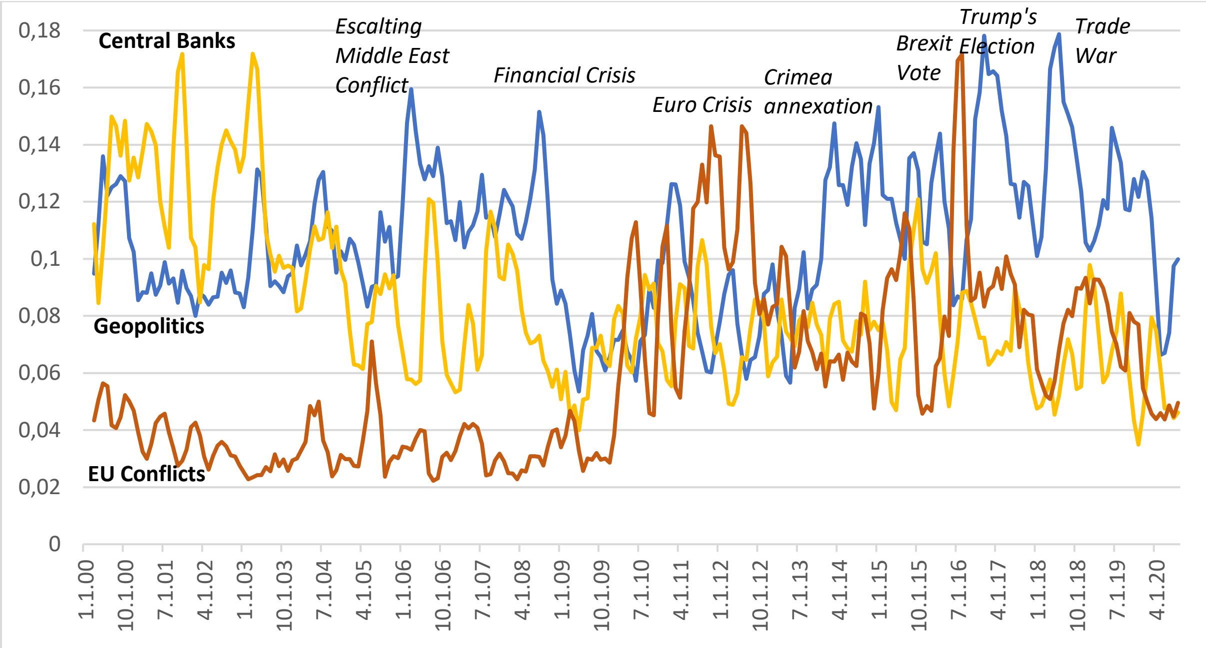

*share of analysis corpus relative to the entire corpus. Three-months moving averages.

Domestic Drivers of Political Uncertainty

In contrast to international uncertainty concerns, domestic politics proved rather predictable in the past one and a half decades. While in the early 2000s a wobbly reform-oriented coalition government between social democrats and greens kept political uncertainty at elevated levels, domestic politics calmed down considerably under the chancellorship of Angela Merkel, who assumed the top-job in the Fall of 2005. Only after the 2017 elections, when it was hard to form a new coalition government, the graph reaches a regional maximum.

Figure 5b: UPI Politics, domestic topics and associated events*

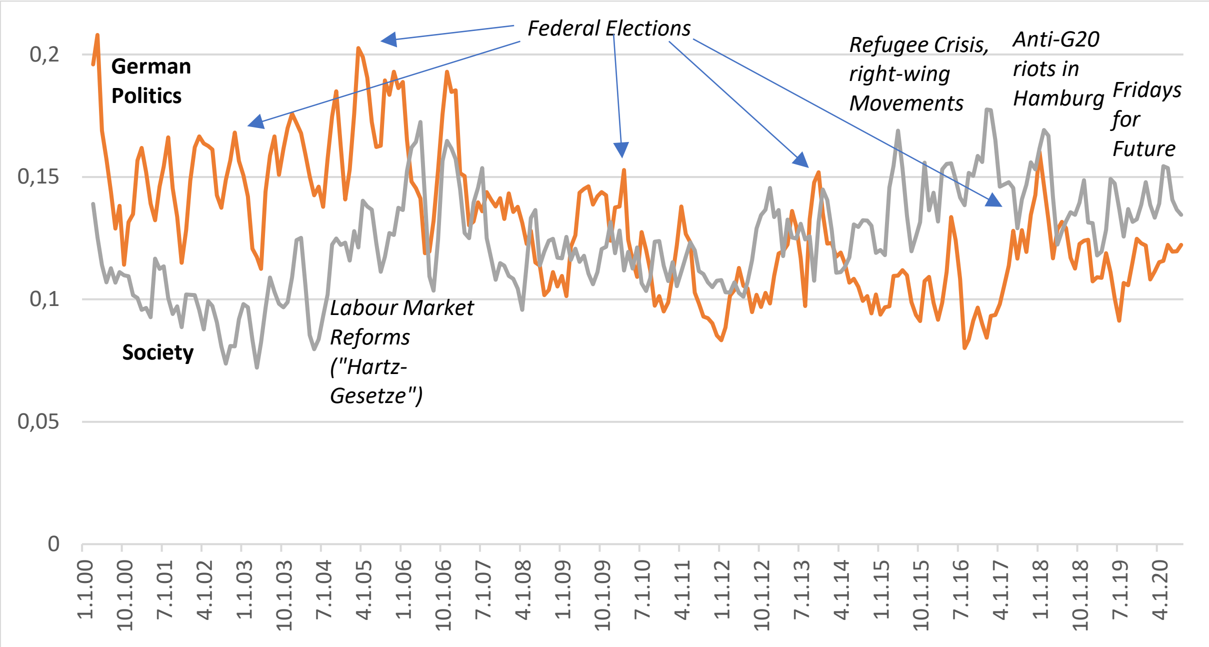

*share of analysis corpus relative to the entire corpus. Three-months moving averages.

The second domestic politics topic has been labeled "Society". Different kinds of political movements fall into this category. As fig. 5b displays, this non-institutionalized part of political life first assumes considerable recognition in national media in the first half of the 2000s, when the extra-parliamentary opposition to the federal government's labor market reforms becomes a vocal part of domestic politics. In the 2010s, social movements rise in importance as a source of uncertainty: fueled by the refugee crisis, right-wing protests (such as Dresden-based "Pegida") gain traction; later in the decade, the hard-left riots against the G20 summit in Hamburg in 2017 and the Fridays for Future demonstrations are events that temporarily raise perceived uncertainty.

Uncertainty and the real economy

Over the entire period from 2000 to 2020 Q3, uncertainty associated with the real economy is mainly driven by developments in individual companies and sectors ("Companies & Markets"). This topic shows three pronounced peak-periods: the first during the early 2000s years, when the economy was rattled by poor performance ("German Disease"), insolvencies, and credit restrains; the second during the financial crisis and the ensuing global recession, the third in the course of the trade war and the sectoral recession of Germany's large manufacturing sector (fig 6).

Figure 6: UPI Real Economy, individual topics and associated events*

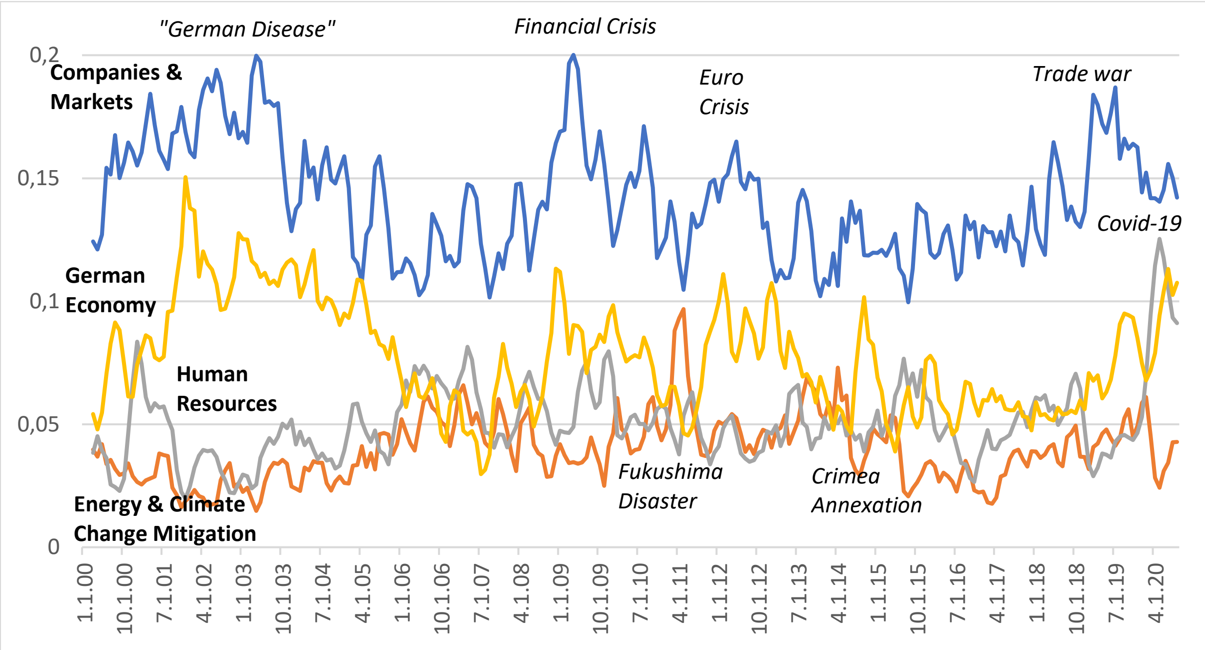

*share of analysis corpus relative to the entire corpus. Three-months moving averages.

The at the time apparently precarious state of the "German Economy" is prominent in the early 2000s, but wanes thereafter, never again reaching the levels of the German Disease period, not even during the Financial Crisis, the Euro Crisis, or the Covid-19 recession, underlining the resilience of the German economy.

Interestingly, we found two other themes connected with the state of the real economy: "Human Resources", a topic dealing predominantly with shortages in the markets for (skilled) labor. The rise of this topic in the early 2000s is associated with the then-perceived lack of IT experts as the internet revolution reached the German economy, followed by a significant drop as the bleak outlook for the German economy took hold. During the boom years of 2006 and 2007 the skills shortages theme gains some prominence, and again after 2014 as the output gap has turned positive once more (Sachverständigenrat, 2019, p. 55-56). During the Covid-19 crisis this topic reaches a maximum, when the shift to home office work and the issue of potentially lasting structural changes in the demand for skills becomes a major issue.

Furthermore, our LDA model produces an "Energy and Climate Change Mitigation" topic, that contains articles about energy markets, but also about climate change and resulting mitigation strategies. The topic reaches its maximum in 2011, when, following the disaster at the nuclear plant in Fukushima (Japan), the German government decides to phase-out nuclear power for good. Another, albeit smaller, peak in 2014 is associated with Russia's annexation of Crimea, that raises concerns about energy security; after all, Russia is a major supplier of oil and natural gas to the German economy.

5. Conclusion: The Curse of Stability?

In this paper, we have focused on the difficulties to produce consistent time-series from probabilistic topic models such as LDA. Attempts to reproduce our *Uncertainty Perception Indicator* (UPI) on a quarterly basis were haunted by profound shifts in the model's structure.

Not even a complete set of corresponding topics could be found at each observation point. Even topics carved out by subsequent models, that appeared to be pretty much identical at first sight, like “German Economy”, were found to differ somewhat at close inspection. We tried several strategies to deal with these problems, particularly *zooming* (raising the value of the LDA parameter *K)*, *accumulating* (grouping sets of topics), and *prototyping* (calculating a sort of mean LDA – as proposed by Rieger et al., 2020 – each quarter to increase the robustness of results)*,* all of which failed to yield the model’s stability we strove for.

Reluctantly, we tried a different approach, which can be classified as a "seeded" LDA, a field that encompasses a variety of methods that represent modifications of LDA using predetermined topics. In our case, this means that the topic structure is fixated at a certain point in time. Subsequently included new data are then distributed to this pre-defined structure. We chose Q2 2020 as the reference date for our UPI model; i.e. the “seed” model is based on data from Jan. 2000 to June 2020. It is then “prototyped” to minimize the randomness of its results. In a further step we added data for Q3 and examined the results.

The thus calculated UPI and its components are plausible and broadly in line with earlier findings. The topics are clear-cut and driven by certain key events – as one would expect of a model based on newspaper content – so that they can be unambiguously labeled. For instance, geopolitical and European developments are distinctly separated into different topics. Traces of the economic fallout of the Corona pandemic can be found in a range of topics, and rightly so, considering the virus’ wide-reaching consequences. This feature suggests that the “seeded” model might actually be able to capture future developments seamlessly as well.

However, while the model’s stability over time clearly is an advantage, it comes with certain drawbacks, explaining our initial reluctance to implement the “seeded” LDA approach. After all, the aim of the UPI is to discover not just *known unknowns*, but also *surprising unknowns* coming from unexpected directions (Müller and Hornig, 2020a, p. 7). The pre-defined model structure, though, may be a hinderance in achieving this goal. Furthermore, time-series computed this way are bound to shifts erratically at some point in time, namely when the corpus’ topic structure changes to a degree that it clashes with the model’s stationarity. Therefore, a more flexible approach is needed, that reconciles the objective of flexibility with the need for stability.

For this purpose, we are currently developing a method that, in a first step, models a reliable model on an appreciable portion of data (for example, up to 2010) using our LDA prototyping method. In subsequent steps, newer articles are distributed according to the existing model. In the process, the assignments of the already modeled articles become constant, so that the associated time series also remains unchanged. To ensure appropriate flexibility, we introduce a parameter that controls after how many quarters the model forgets the information of the assignments of articles and can thus rearrange the topics. For example, if we choose the parameter to take the value of three quarters, new articles will always be initialized – or “seeded” – based only on the assignments of the articles of the last three quarters. This method is intended to combine the aim of a consistent time series with the flexibility needed to detect new topics and will be implemented in a subsequent paper.

6. References

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

Blei, D.M., Ng, A.Y., & Jordan, M.I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3, 993–1022. http://www.jmlr.org/papers/volume3/blei03a/blei03a.pdf

Chang, J., Boyd-Graber, J. L., Gerrish, S., Wang, C., & Blei, D. M. (2009). Reading Tea Leaves: How Humans Interpret Topic Models. In Y. Bengio, D. Schuurmans, & J. D. Lafferty (Eds.), 24 Council of Economic Advisors to the Federal Government. Annual Report 2020/2021, November 2020 Advances in Neural Information Processing Systems 22 (pp. 288–296). http://papers.nips.cc/paper/3700-reading-tea-leaves-how-humans-interpret-topic- models.pdf

Chang, J. (2015). Collapsed Gibbs Sampling Methods for Topic Models. LDA package. https://cran.r-project.org/web/packages/lda/lda.pdf

Entman, R. M. (1993). Framing: toward clarification of a fractured paradigm. *Journal of Communication*, 43(4), 51–58. https://doi.org/10.1111/j.1460-2466.1993.tb01304.x

Griffiths, T. L. & Steyvers, M. (2004). Finding scientific topics. Proceedings of the National Academy of Sciences of the United States of America 101 Suppl 1, 5228–5235. https://doi.org/10.1073/pnas.0307752101

Grün, B., Hornik, K. (2011). topicmodels: An R Package for Fitting Topic Models. *Journal of Statistical Software*. https://www.doi.org/10.18637/jss.v040.i13

Koppers, L., Rieger, J., Boczek, K., & Nordheim, G. von. (2020). *tosca: Tools for Statistical Content Analysis* (Version R package version 0.2-0) [Computer software]. https://doi.org/10.5281/zenodo.3591068

Müller, H., von Nordheim, G., Boczek, K., Koppers, L., & Rahnenführer, J. (2018). Der Wert der Worte – Wie digitale Methoden helfen, Kommunikations- und Wirtschaftswissenschaft zu verknüpfen. *Publizistik*, 63(4), 557–582. https://doi.org/10.1007/s11616-018-0461-x

Müller, H. (2020). Donald Trump als wirtschaftspolitischer Unsicherheitsfaktor. *Ifo Schnelldienst*, 73(1/2020), 26–29. https://www.ifo.de/DocDL/sd-2020-01-2020-01-22_1.pdf

Müller, H., & Hornig, N. (2020a). Expecting the Unexpected: A new Uncertainty Perception Indicator (UPI) – concept and first results. DoCMA Working Paper #1-2020. http://dx.doi.org/10.17877/DE290R-21089

Müller, H., & Hornig, N. (2020b). "I heard the News today, oh Boy. An updated Version of our Uncertainty Perception Indicator (UPI) – and some general thoughts on news-based economic indicators. *DoCMA Working Paper #2-2020*. http://dx.doi.org/10.17877/DE290R-21669

Nguyen, V.-A., Boyd-Graber, J., Resnik, P. & Chang, J. (2014). Learning a Concept Hierarchy from Multi-labeled Documents. Neural Information Processing Systems, 2014, 9 pages. http://users.umiacs.umd.edu/~jbg/docs/2014_nips_l2h.pdf

Rieger, J. (2020). ldaPrototype: A method in R to get a Prototype of multiple Latent Dirichlet Allocations. *Journal of Open Source Software*, 5(51), 2181. https://doi.org/10.21105/joss.02181

Rieger, J., Rahnenführer, J. & Jentsch, C. (2020a). Improving Latent Dirichlet Allocation: On Reliability of the Novel Method LDAPrototype. Natural Language Processing and Information Systems, NLDB 2020. LNCS 12089, pp. 118-125.

Rieger, J., Koppers, L., Jentsch, C., & Rahnenführer, J. (2020b). Improving Reliability of Latent Dirichlet Allocation by Assessing Its Stability Using Clustering Techniques on Replicated Runs. ArXiv:2003.04980 [Cs, Stat]. http://arxiv.org/abs/2003.04980

Sachverständigenrat zur Begutachtung der gesamtwirtschaftlichen Entwicklung (2019). Den Strukturwandel meistern. Jahresgutachten 2019/20, Jahresgutachten, No. 2019/20, ISBN 978-3-8246-1089-1, Sachverständigenrat zur Begutachtung der Gesamtwirtschaftlichen Entwicklung, Wiesbaden.

Shiller, R. J. (2017). Narrative economics. *American Economic Review*, 107(4), 967–1004. https://doi.org/10.1257/aer.107.4.967

Shiller, R. J. (2019). Narrative Economics: How Stories Go Viral and Drive Major Economic Events, Princeton University Press.

Tenenboim-Weinblatt, K., Hanitzsch, T., & Nagar, R. (2016). Beyond peace journalism: reclassifying conflict narratives in the Israeli news media. *Journal of Peace Research*, 53(2), 151–165. https://doi.org/10.1177%2F0022343315609091