



RUHR

ECONOMIC PAPERS

Boris Blagov
Henrik Müller
Carsten Jentsch
Torsten Schmidt

The Investment Narrative – Improving Private Investment Forecasts with Media Data

Imprint

Ruhr Economic Papers

Published by

RWI – Leibniz-Institut für Wirtschaftsforschung

Hohenzollernstr. 1-3, 45128 Essen, Germany

Ruhr-Universität Bochum (RUB), Department of Economics

Universitätsstr. 150, 44801 Bochum, Germany

Technische Universität Dortmund, Department of Economic and Social Sciences

Vogelpothsweg 87, 44227 Dortmund, Germany

Universität Duisburg-Essen, Department of Economics

Universitätsstr. 12, 45117 Essen, Germany

Editors

Prof. Dr. Thomas K. Bauer

RUB, Department of Economics, Empirical Economics

Phone: +49 (0) 234/3 22 83 41, e-mail: thomas.bauer@rub.de

Prof. Dr. Wolfgang Leininger

Technische Universität Dortmund, Department of Economic and Social Sciences

Economics – Microeconomics

Phone: +49 (0) 231/7 55-3297, e-mail: W.Leininger@tu-dortmund.de

Prof. Dr. Volker Clausen

University of Duisburg-Essen, Department of Economics

International Economics

Phone: +49 (0) 201/1 83-3655, e-mail: vclausen@vwl.uni-due.de

Prof. Dr. Ronald Bachmann, Prof. Dr. Manuel Frondel, Prof. Dr. Torsten Schmidt,

Prof. Dr. Ansgar Wübker

RWI, Phone: +49 (0) 201/81 49-213, e-mail: presse@rwi-essen.de

Editorial Office

Sabine Weiler

RWI, Phone: +49 (0) 201/81 49-213, e-mail: sabine.weiler@rwi-essen.de

Ruhr Economic Papers #921

Responsible Editor: Wolfgang Leininger

All rights reserved. Essen, Germany, 2021

ISSN 1864-4872 (online) – ISBN 978-3-96973-067-6

The working papers published in the series constitute work in progress circulated to stimulate discussion and critical comments. Views expressed represent exclusively the authors' own opinions and do not necessarily reflect those of the editors.

Ruhr Economic Papers #921

Boris Blagov, Henrik Müller, Carsten Jentsch, and Torsten Schmidt

The Investment Narrative – Improving Private Investment Forecasts with Media Data

Bibliografische Informationen der Deutschen Nationalbibliothek

The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie;
detailed bibliographic data are available on the Internet at <http://dnb.dnb.de>

RWI is funded by the Federal Government and the federal state of North Rhine-Westphalia.

<http://dx.doi.org/10.4419/96973067>

ISSN 1864-4872 (online)

ISBN 978-3-96973-067-6

Boris Blagov, Henrik Müller, Carsten Jentsch, and Torsten Schmidt¹

The Investment Narrative – Improving Private Investment Forecasts with Media Data

Abstract

Corporate investment in Germany has been relatively weak for a prolonged period after the financial crisis. This was remarkable given that interest rates and overall economic activity, important determinants of corporate investment, developed quite favourably during that time. These developments highlight the fact that the dynamics of business cycles varies over time: each cycle is somewhat different. A promising new line of research to identify the driving factors of business cycles is the use of narratives (Shiller 2017, 2020). Widely shared stories capture expectations and beliefs about the workings of the economy that may influence economic behavior, such as investment decisions. In this paper, we use Latent Dirichlet Allocation (LDA) to identify topics from news (text) data related to corporate investment in Germany and to construct suitable indicators. Furthermore, we focus on isolating those investment narratives that show the potential to lead to substantial improvement of the forecasting performance of econometric models. In our analysis, we demonstrate the benefit of using media-based indicators to improve econometric forecasts of business equipment investment. Newspaper data carries important information both on the future developments of investment (forecasting) as well as on current developments (nowcasting). Moreover, the identified investment narrative enables the researcher to improve her/his understanding of the investment process in general and allows to incorporate exogenous developments as well as economic sentiment, news and other relevant events to the analysis.

JEL-Code: C53, C82, E32

Keywords: Narrative economics; mixed-frequency; nowcasting via media data

September 2021

¹ Boris Blagov, RWI; Carsten Jentsch, TU Dortmund; Henrik Müller, TU Dortmund; Torsten Schmidt, RWI. – All correspondence to: Torsten Schmidt, RWI, Hohenzollernstr. 1/3, 45128 Essen Germany, e-mail: torsten.schmidt@rwi-essen.de

I. Introduction

Corporate investment in Germany has been relatively weak for a prolonged period after the financial crisis. This was remarkable given that interest rates and overall economic activity, important determinants of corporate investment, developed quite favourably during that time. However, the weak recovery in corporate investment was not an isolated development but can be observed in many other European economies. One possible explanation is that the consequences of the financial crisis continued to dampen the recovery for some time. Another one holds that the sovereign debt crisis in the euro area has contributed significantly to the slowdown in investment dynamics. These developments highlight the fact that business cycles vary considerably over time: each cycle is somewhat different. Time and time again new factors are at play that standard forecasting strategies struggle to incorporate. This requires the development of new econometric techniques and the use of non-standard data sources to improve the forecasting performance.

In fact, forecasts of corporate investment deteriorated in the 2010s. Standard approaches that use variables like financing costs and economic activity (CBO 2018; Rapach and Wohar 2007) proved less useful than in the past. To improve the forecasting performance for corporate investments several attempts have been made recently by using information e.g. from economic surveys (Driver and Meade 2019; Giordano et al. 2021). Driver and Meade (2019) argue that business surveys are able to incorporate economic developments that are not captured by commonly used economic time series, and therefore are not taken into account by forecasters. The flexible processing of information in particular with regard to the time-varying importance of different drivers of investment activity is clearly an advantage of this type of data. However, the informational content of the surveys, i.e. what respondents actually believe, still remains unknown.

A promising new line of research to identify the (possibly time-varying) driving factors of business cycles is the use of narratives (Shiller 2017, 2020). Widely shared stories capture expectations and beliefs about the workings of the economy that may influence economic behavior as e.g. investment decisions. Therefore, the inclusion of such narratives can be expected to improve economic forecasts. To make the narrative approach usable for forecasting, newspaper corpora are an available source of information. These text data need to be transformed into quantitative time series to make the contained information accessible for commonly used time series methods. Thorsrud (2016) and Larsen and Thorsrud (2019) used newspaper data for Norway and for the US, respectively, to construct daily indicators that track the development of GDP quite well. Ellingsen et al. (2020) use indicators based on newspaper texts to improve forecasts for macroeconomic time series. In addition, they apply unsupervised topic modelling approaches, such as Latent Dirichlet Allocation (LDA) proposed by Blei et al. (2003), which has become a standard tool for text mining purposes. Such text mining tools allow to semi-automatically form thematically related text categories ("topics"), assigning individual texts to these topics. One potential advantage of this approach is that it enables researchers to identify distinct patterns in newspaper coverage that may contain additional information concerning economic activity.

In this paper, we use LDA to construct indicators derived from news topics that are related to equipment investment in Germany. Compared to previous applications in econometric analyses, we combine the LDA method, which is unsupervised per se, with expert knowledge used to label the obtained LDA topics. Adding meaning to data facilitates better interpretations of the results,

as we focus on isolating investment narratives that are capable of enhancing the forecasting performance of econometric models significantly. LDA topics have been described as media frames (di Maggio et al. 2013), a concept pioneered by Entman (1993). Moreover, for longer periods of time, topics can also be interpreted as narratives (Müller et al. 2018, 2021), spanning a sequence of events and providing a quantitative specification of the concept of the economic narrative, as suggested by Shiller (2017, 2019, 2020).

Our research strategy is structured in three subsequent stages:

First, large text corpora consisting of newspaper articles from three major nation-wide German broadsheet newspapers – Süddeutsche Zeitung, Welt and Handelsblatt – spanning more than two decades of reporting are processed using several queries to construct interpretable sub-corpora. On these sub-corpora a host of LDA models with different parameter settings are computed and their content structure reviewed by human coders. Two models are chosen for further content analyses. In a further step, thematically closely related LDA topics are aggregated to what we call “Investment Factors” (IF).

Second, we explore the predictive power of these models’ topics and IFs regarding equipment investment. As single LDA topics and aggregated IFs can be continuously updated using newly arriving text data, so that they are available at a higher (monthly) frequency in comparison to commonly used quarterly macro variables, our approach allows to investigate their ability to enhance both forecasting and nowcasting. To explore the forecasting properties of the data, we aggregate the monthly series to a quarterly level and use both univariate and multivariate autoregressive setups to evaluate forecasts of next period’s investment growth. Thus, we use newspaper data only up to the last quarter in which investment figures are also available. This approach allows us to identify the most informative topics regarding investment. In the next step, we fully exploit the timeliness of media data and utilise a mixed-frequency vector autoregressive model for nowcasting, which also incorporates media data from the current quarter, for which investment figures are not yet available. The results single out several topics and IFs that indeed enhance the models’ predictive power.

Third, we explore the content of these best performing topics and IFs more closely. By applying a formula proposed by Müller et al. (2018) we extract investment narratives, i.e. underlying stories that drove public discourse and affected investment decisions according to our data over the period considered.

One advantage of this explorative multi-stage research strategy is that it is fundamentally data-driven. In contrast to other contributions to this field (e.g. Shiller 2019, 2020) we do not set out to find certain narratives in the data, but initially rely on the unsupervised LDA method to structure the text content, thereby taming researchers’ own convictions, which tend to induce biases when economic narratives are formulated.

This paper is organized as follows: In section II we describe the textual database and our approach to construct time series that capture different aspects of corporate investment dynamics. Section III describes the different model setups that we use to evaluate the informational content of the

newly constructed indicators for fore- and nowcasting. Section IV picks up these results and formulates the investment narratives. Section V concludes.

II. Empirical Specification

A. Text data selection

This section describes how time series containing information on the development of equipment investment are generated from texts. It is plausible to assume that a) economic developments are reflected in media reports *before* they become visible in national accounts or sentiment indicators (allowing "nowcasting"), b) economic policy and other exogenous factors of economic importance are captured early on by the media, and c) media content has an impact on the formation of expectations of economic subjects and may therefore influence decisions.

One important initial question concerns the selection of media. Since business cycle research requires long time series, preferably spanning several cycles, traditional daily newspapers, which have a long publication history and relatively constant criteria for issue selection, unlike online or social media content, are the obvious choice. Since our aim is mapping developments in the whole of Germany, leading broadsheet newspapers with a national circulation are preferable to regional media. They cater to an elite audience, i.e. to people who actually take investment decisions. Furthermore, they influence the reporting of other (follower) media and can thus be expected to capture broader public discourses. For this study, the content of three German mainstream newspapers with different leanings is used: Handelsblatt (business-focussed), Süddeutsche Zeitung (center-left), and Die Welt (center-right). The corpora were obtained from LexisNexis and the publishing houses. In total, they comprise about four million articles.

TABLE 1

Overview of the considered corpora

Media (Period)	Search Strategy	Query	Size (no. of articles)
<i>Corpora</i>			
I. Handelsblatt (1.1.94–31.12.19)		–	993.381
II. Süddeutsche Zeitung (1.1.94–31.12.19)		–	1.802.895
III. Handelsblatt, Süddeutsche Zeitung, Welt (1.1.00 – 31.12.19)		–	2.885.539
<i>Analysis Corpora</i>			
I.a HB (1.1.94–31.12.19)	explorative	Wirtschaft*+Investition* (1 mention each) ¹	52.049
I.b HB (1.1.94–31.12.19)	explorative	Wirtschaft*+Investition* (2 mentions each)	15.222
I.c HB (1.1.94–31.12.19)	theory-informed	list of investment-relevant terms (5 out of 20) ²	130.721
II.a Süddeutsche (1.1.94–31.12.19)	explorative	Wirtschaft*+Investition* (1 mention each)	30.817
II.b Süddeutsche (1.1.94–31.12.19)	explorative	Wirtschaft*+Investition* (2 mentions each)	6.156
II.c SZ (1.1.94–31.12.19)	theory-informed	list of investment-relevant terms (5 out of 20)	73.052
III.a HB, SZ, Welt (1.1.00–31.12.19)	explorative	Wirtschaft*+Investition* (1 mention each)	75.845

¹ * indicate truncations.

² Articles chosen include at least 5 of the following items: anbot, anlage, bank, entwicklung, export, import, investition, kapital, konjunktur, lohn, loehne, nachfrage, politik, regulierung, risiken, risiko, steuer, unsicher, wirtschaft, zins (supply, facility, bank, development, exports, imports, capital, business cycle, wage(s), demand, politics, regulation, risk(s), tax, uncertain, economy, interest).

To narrow down the corpora regarding our research interest, an explorative and a theory-informed query were tested to construt sub-corpora. First, two explorative approaches are used, where all articles in which the word stems "wirtschaft*" (economy) and "investment*" co-occur at least once or twice, respectively, are chosen. Second, using a theory-informed approach, all articles containing a subset of words from a list of pre-specified terms related to investment decisions are included. All three content selection strategies are applied separately to corpora from Handelsblatt and Süddeutsche Zeitung. Both were available from 1994. Additionally, we constructed a combined corpus, including close to 2.9 million articles from all three newspapers published between 2000 and 2019 (Die Welt was available from 2000), on which we deploy the first query

only. The resulting sub-corpora differ considerably in size (see Table 1 for an overview). With regard to investment activity, Handelsblatt proved to be particularly productive: as the only pure-play business newspaper in Germany, it provides broader and more detailed information on economic developments than other print media. However, the merged corpus from all three newspapers also represents a promising basis for new types of economic indicators; the inclusion of several media should provide greater thematic breadth and robustness against arbitrary editorial decisions. Therefore, we chose the two broadest corpora (I.c and III.a) for further analysis, relying on LDA to subsequently sort out content irrelevant to our research interest.

Figure 1: Frequency of the analysis corpora over time (shares relative to entire corpora, three-month moving average)




Figure 1 shows the frequency of the two corpora over time. III.a is characterized by three distinct peaks, associated with the financial crisis (2008/09), the investment slump following the European sovereign debt crisis (2014), and the manufacturing recession (2018/19) respectively. In addition, I.c follows a secular upward trend starting in 2014, a period when weak business investment activity rose in public awareness considerably.

B. Topic Modelling

To extract factors relevant to the analysis regarding investment, we use LDA (Blei et al. 2003) to form (unsupervised) thematically related clusters of texts ("topics") and to assign individual articles to these topics. If calibrated thoroughly, the topics' content is interpretable by human researchers who have the outcome of LDA models at their disposal, that is, characteristic words ("Top Words"), texts ("Top Texts"), and frequency distributions over time generated by the algorithm for each topic. In this regard, LDA thus becomes a *semi-supervised* method combining both quantitative and qualitative aspects. When running an LDA, a critical decision to be taken concerns the parameter K , which is the number of topics the LDA is set to form. The choice of K can be interpreted as using different zoom lenses: low K values produce overview-like, broad topics, similar to a wide-angle focal length; higher K values provide narrower, more detailed topics. For the choice of K , we have deliberately decided against using approaches for an automatic data-driven selection criterion. This is for two reasons: first, existing methods tend to lead to poorly delineated topics that are difficult to interpret (Chang et al. 2009) and require considerable additional computational effort. Second, it is essential for our further analyses to obtain well-interpretable and adequately tailored topics that allow answering the questions of interest. The content analysis was performed using the R package *tosca* (Koppers et al. 2020). After reviewing a number of different LDA models, we chose an analysis strategy that combines the broad queries discussed in A. with a high parameter value ($K=20$). The aim was to exclude as few relevant newspaper texts as possible by setting too narrow a search term and to allow detailed insights into content patterns at the same time. Topics that are irrelevant to the research goal can then be ignored in subsequent steps; thematically closely related topics can, in turn, be grouped according to four categories – politics, markets, finance, technology – we call "Investment Factors" (IF).

C. Empirical Results

Tables 2 and 3 provide an overview of both models' topics. The results are quite promising: 39 out of 40 topics obtained from both LDAs with $K=20$ topics each could be labelled unambiguously and were clearly distinguishable in their content and framing from one another. The vast majority of them were plausibly related to our research interest concerning business investment dynamics.

TABLE 2

Overview of LDA Model A

LDA Model A Corpus I.c, K=20				
Topic Label	Share (%)	Content	Relevant?	Investment Factor (IF)
German Politics	5.05	Party politics, Merkel, Schröder etc.; peaks during federal elections.	Yes	Politics
Miscellaneous	9.4	Diverse	No	–
Personal Finance I	6.35	Investor perspective, stock market focus	Yes	Finance
SME Financing	4.94	Credit to the German <i>Mittelstand</i> , 90s and 2000s topic	Yes	Finance
Personal Finance II	3.62	Focus on retirement	Yes	Finance
R & D	5.9	Research & Development, considerable rise after 2012	Yes	Technology
Energy	1.9	Oil and gas markets, topic is largely absent in the 2010s	Yes	Markets
German Economy	6.31	Domestic business cycle movements	Yes	Markets
International Conflicts	4.54	Political Conflicts involving China, Russia, India, Iran, Turkey, the US. Peaks: Crimea annexation (2014), Brexit, Trump's election (2016)	Yes	Politics
Financial Crisis	5.26	The Financial Crisis of 2008, its aftermath and fall-out, focus on banks' health	Yes	Finance
Stock Markets I	2.72	Developments at global bourses, focus on US and UK	Yes	Finance
Small German Banks	3.93	Health of small Banks, financing of SMEs; topic diminishes to zero after 2009	Yes	Finance
Globalization	7.56	Systemic discourse about openness and the perils of competition, several peaks (introduction of the Euro 1999, football world championship in Germany 2006, Euro Crisis 2012, Brexit Vote 2016, Trade War 2019)	Yes	Politics
Company Reports	6.94	Periodic reporting on corporate earnings	Yes	Markets
Central Banks	5.15	Monetary Policy, peaks: financial crisis, Euro crisis	Yes	Finance
Bond Markets	1.98	Yields and movements in fixed income markets	Yes	Finance
Stock Markets II	4.1	Developments at bourses including German ones	Yes	Finance
Energy Transition	4.27	Switching to renewables and EVs, pronounced increase after 2015	Yes	Technology
Foreign Direct Investment	4.31	Investing in non-domestic markets, 90s topic (Eastern Europe, Latin America etc., no China or India)	No	–
Structural Reforms	5.75	Domestic Reforms, topic is particularly strong in the 90s (<i>"Standortdebatte"</i>) and labor market reforms (<i>"Hartz Gesetze"</i>)	Yes	Politics


TABLE 3

Overview of LDA Model B

		LDA Model B III.a, K=20		
Topic Label	Share (%)	Content	Relevant?	Investment Factor (IF)
Financing SMEs	5.32	Credit to the German <i>Mittelstand</i> , peaks in the early 2000s.	Yes	Finance
Rankings of German Corporations	1.71	„Welt“-only topic Ranking	No	–
Sports & Culture	8.53	Professional Sports, Culture, Entertainment	No	–
EU	4.06	European dimension of investing, EU budget, Deficit surveillance	Yes	Politics
Labour Markets	4.63	Labour market developments and policies in Germany	Yes	Markets
Commercial Real Estate	4.92	Property Markets, Projects, Building activity	Yes	Markets
International Conflicts	3.53	Conflicts particularly in the Middle East with US involvement	Yes	Politics
IT Investments	3.73	Investments in Hardware, Software, Networks etc.	Yes	Technology
German Corp.	4.05	Investment by major German companies	Yes	Markets
Capitalism	10.43	Skeptical accounts of the capitalist social and economic model. discourse topic, pronounced rise starting in 2012	Yes	Politics
Company reports	5.3	Periodic reporting on quarterly and annual results	Yes	Markets
Energy	4.82	Energy markets, transition to renewables, peak in 2011 (Fukushima nuclear disaster)	Yes	Technology
Ressources Exporting Countries	3.77	Fokus on Russia, peaks: annexation of Crimea, election of Morales in Bolivia	No	–
German Fiscal Policy	7.76	Fiscal Stance, Taxation	Yes	Politics
Education	3.94	Training, Career	Yes	Markets
Private Investing	4.8	Private Asset Allocation, Wealth Management, financial market outlook (implicit information on financing conditions)	Yes	Finance
Monetary Policy	5.68	Central Banks' stance, peaks: financial crisis, Euro crisis	Yes	Finance
Investing in Asia	2.89	Investment dynamics in China, Taiwan, Japan, India...	No	–
Internat. Treaties	4.39	International investment pacts (TTIP, Ceta, etc.)	Yes	Politics
Business Cycle	5.78	Economic growth, Unemployment etc.	Yes	Markets

Figures 2 and 3 summarize the reporting patterns over the period considered. A couple of common threads of both models are noteworthy: politics have become more of an issue for investors in the 2010s, while market conditions have declined in relative relevance, as have financing conditions (though in the I.c-based model Finance is still the most important IF). Technology has increased markedly in importance in recent years, according to the analysis based on the Handelsblatt-only corpus (Figure 2), while it has stayed broadly stable, when the larger merged corpus is considered (Figure 3).

*Figure 2: Investment Factors Model A**




* Investment factors' shares relative to analysis corpus, three-month moving average

The rise of politics as an IF seems to be a result of populist politics taking center stage in Western countries, including the US and the UK, but also in Germany where the refugee crisis of 2015/16 fuelled the rise of right-wing populists that have rendered domestic politics and federal elections fiercely contested. Across the West, unpredictable politics have triggered unprecedented degrees of economic uncertainty (Baker et al. 2016; Davis 2019; Müller et al. 2021), until the Covid-19 pandemic, a natural disaster, hit in early 2021. On the other hand, the relative decline of real

(market) and financial influences on investment dynamics, as gauged by the corresponding IFs, may be explained by the fact that many German companies have diversified their sales and sourcing markets rapidly over the 2010s. A broader international market presence has impregnated firms against downturns in a single country or region. As far as external financing is concerned, this was a big issue in the 1990s and the early 2000s for many German companies, but has become less of a headache. Due to solid profit positions, firms on average have been able to finance their investment outlays internally; the German corporate sector as a whole has become a net saver.

Figure 3: Investment Factors Model B*



*Investment factors' shares relative to analysis corpus, three-month moving average


However plausible these observations may appear, it is still unclear whether and, if so, to what extent they contain additional, forecast-enhancing information on business investment dynamics. In the following sections, the time-series derived from our LDA models' topics and IFs are investigated with respect to their informational content and whether and to which extent they have the potential to improve forecasts in an econometric analysis.

III. Forecasting setup

To test the informational content of the media data, we will employ three different model setups with varying degrees of complexity. First, we will assess the predictive power of the data with respect to *forecasting* the future growth rate of equipment investment. Using univariate and multivariate autoregressive frameworks, we will test how well the information contained in the newspaper articles from the latest quarter, for which investment data is available, predicts the

future developments in equipment investment. Second, we will use a multivariate mixed-frequency autoregressive framework, to assess whether the investment narrative is informative for conducting *nowcasting* – whether the media data is informative regarding the current investment dynamics. Here we will utilise the timeliness of newspaper articles to its full potential, as they are available before most other indicators.

Figure 4: Growth rate of equipment investment – the target variable for forecasting



The target variables for the forecast will be the seasonally adjusted and calendar adjusted quarterly real growth rates of equipment investment, w_t , depicted in Figure 4. The forecast accuracy will be measured with the widely adopted Root Mean Squared Forecast Error (RMSFE), which compares the forecast, $\tilde{w}_t^x$, from some model x used for prediction with the realised data over a forecast horizon of length h . Given a dataset with T observations, the h -step ahead out-of-sample RMSFE is calculated as

$$RMSFE^x = \sqrt{\frac{\sum_{t=T+1}^h (w_t - \tilde{w}_t^x)^2}{h}}.$$

In this study we will focus on short-term forecasting, i.e. predicting the developments of the current quarter. Therefore, we set $h = 1$.

Univariate and bivariate models for forecasting

In forecasting comparisons, it is common to test the informational content of variables in the simplest framework possible. A standard approach in the forecasting literature is to begin with a naïve model – a model that does not require expert knowledge to use and create forecasts from. If

more sophisticated setup cannot produce any better forecasts, they are unwarranted. In the literature on forecasting macroeconomic aggregates, the standard naïve model is the univariate autoregressive (AR) model of order p ,

$$w_t = c + \sum_{j=1}^p a_j w_{t-j} + u_t ,$$

where w_t denotes the growth rate in our context, c is a constant (intercept), which reflects the long-run growth rate, a_j are the autoregressive coefficients which reflects how past values of the growth rate affect the current growth rate, and u_t is an error term. This auto-regression can be estimated by standard techniques such as the method of ordinary least squares. Then, based on the AR model fit, a forecast for the next quarterly growth rate is easily produced using the estimated parameters $\hat{c}$ and $\hat{a}_j$. Given data until time T , the one-step ahead forecast for time $T + 1$ is simply computed as

$$\tilde{w}_{T+1}^{AR} = \hat{c} + \sum_{j=1}^p \hat{a}_j w_{T+1-j}.$$

Thus, calculating the RMSFE from the AR model may serve as a benchmark to which the more sophisticated methods can be compared. Given this baseline forecast, the newly created indicators may be incorporated additionally in the model. A common approach is to use an auto-regressive distributed lag (ARDL) model, extending the standard AR model. Suppose that I_t is an indicator created from the new newspaper-based datasets. Then, an ARDL model of orders p and r is

$$w_t = c + \sum_{j=1}^p a_j w_{t-j} + \sum_{j=0}^r b_j I_{t-j} + u_t .$$

Similar to the AR framework, the model is used to generate forecasts for the target variable, which is then compared to the actual realisations. Finally, comparing the forecast errors of both models AR and ARDL leading to $RMSFE^{AR}$ and $RMSFE^{ARDL}$, respectively, is then revealing whether the indicator I_t has additional informational content regarding the target variable.

In the case of both the AR model and the ARDL model, the choice of lags p and r influences the forecasting performance. Using too few lags could miss important dynamics in the time series data, while too many could reduce the estimation precision of the model parameters – a problem called overfitting. In the forecasting literature, it is common to use information criteria (IC) to select the values of p and r . Common choices are either the Akaike information criteria (AIC), which tends to select more lags than might be necessary, or the Bayesian information criteria (BIC), which is on the other end of the spectrum and favours only few past values. In this analysis, all estimations have been performed using both ICs and the BIC turned out to generate lower RSMFE in all cases such that we focus only on the BIC results in what follows.

A drawback of the ARDL model is that it does not allow for general equilibrium effects as there is only one dependent variable. Thus, the iterative forecasts do not account for the full dynamics of the system. Therefore, in the second stage, the dynamics of the target variable and the indicator are considered jointly in a bivariate vector autoregression (VAR) of order p ,

$$Y_t = C + \sum_{j=1}^p A_j Y_{t-j} + U_t,$$

Where $Y_t = (w_t, I_t)'$ and C denote bivariate vectors, A_j are 2x2 coefficient matrices, and U_t are bivariate error terms. The VAR order p is chosen analogously to the univariate case using BIC. Also, the growth rate forecast, $\tilde{w}_{T+1}^{VAR}$, is generated as in the AR case by using the estimated $\hat{A}_j$ and $\hat{C}$ with the latest data.

Multivariate approach for nowcasting - a mixed-frequency VAR

The models presented in the previous subsections are used for forecasting next period's investment growth. It is well known that national accounts data is released with a significant delay. For example, the data for investment in the fourth quarter is published towards the end of next years' month of February. Forecasting the first quarter using the presented models then involves the aggregation of newspaper data from October to December. However, at the end of February, two full months of newspaper articles are already available, which are not taken into account when forecasting.

In contrast to forecasting, nowcasting intends to utilise all data available, including the two months from the respective quarter for which a forecast (nowcast) is produced. However, two challenges arise. The model has mismatching monthly and quarterly frequencies ("mixed frequencies") and the data is not "balanced" in the sense that data for some variables (in this case investment) is missing for the first quarter, while other variables (the media data) are available. Several approaches have been proposed in the literature to deal with these problems. Most notably – MIDAS and mixed frequency models (Ghysels and Marcellino 2018). The former is a set of regressions which intend to match monthly variables to quarterly observations via a lag polynomial, the latter utilising a state space approach to deal with the missing observations.

In this paper, we employ a Mixed-frequency Bayesian VAR model to incorporate the timely information from the media data. This class of models deals optimally with mismatching frequencies (in a mean-squared error sense) and allows for the inclusion of many variables simultaneously. The model is based on Schorfheide and Song (2015) and has become popular in recent years due to its excellent properties of combining monthly with quarterly data (Cross et al. 2020, Gefang et al. 2020, Koop et al. 2020). Döhrn et al. (2018) adapted it for the German economy and showed that it helps to improve the nowcast of GDP compared to quarterly models that use only data from the previous quarter.

The mixed-frequency VAR model may be summarised with the following two equations. First, it assumes that the variables evolve on a monthly frequency

$$x_t = C + \sum_{j=1}^6 A_j x_{t-j} + U_t,$$

where x_t is the vector of the monthly values of *all* variables. Given that the monthly variables will be mapped on to quarterly observations, the lag-order of the monthly VAR has to be a multiple of 3. We follow Schorfheide and Song (2015) and choose a lag order of 6, which is equivalent to a quarterly VAR with two lags. Let y_t be the vector of quarterly values, e.g. for industrial production the quarterly value is the average of the three monthly values. Industrial production is observed every month of the quarter. Other variables, such as equipment investment is observed every third month (ignoring publication lags for the exposition). Moreover, the monthly values of equipment investment are never observed, only the quarterly average is. Thus, in this model, some variables are observed every month and every quarter, while others are observed only every quarter. This property can be formulated in a unified way by using the following second model equation,

$$y_t = M_t \Lambda_t x_t.$$

Here, M_t is a matrix indicating whether the monthly values of a variable are observed or not, while Λ_t indicates how the monthly values relate to the quarterly values (e.g., the quarter value should be the average of the three monthly values).

This model can be estimated using the Kalman filter (Kalman and Bucy 1961) to deal with the fact that some monthly values are never observed, and the overfitting problem is solved via Bayesian estimation. We follow Schorfheide and Song (2015) for the implementation of the estimation algorithm. The model is estimated via Gibbs sampling for which we let 12,000 iterations of which we discard the first 6000 and save the rest.

IV. Forecasting Results

In Section II the investment factors and topics from model A and model B were presented. The former consists only of the “Handelsblatt” newspaper and having a longer time series spanning from 1994, while the latter, starts at 2000, but incorporates more media outlets. In the estimation we differentiate between the two different samples, since the reduced number of quarterly observations of the target variable could increase estimation uncertainty.

For the forecast evaluation, the data set is divided into two periods. The first period – the training sample - is used for the initial estimation of the models. The second period, used for forecasting, is referred to as the evaluation sample.

There is a clear trade-off between the length of the training sample and the evaluation sample, respectively. A long evaluation sample ensures that the results are not a product of the selected data periods, but that the indices are actually informative for the target variable. However, it is also important that the models are estimated correctly - that there are enough observations to recover

the true relationships between the variables. In addition, it is important to use the same training and evaluation sample for model comparison.

The evaluation sample is chosen to begin in 2012Q1 and run through 2019Q4, for a total of 32 periods. That is, each model is initially estimated with data extending to 2011Q4. The training sample includes the global financial crisis purposely. During the crisis, investment declined considerably, hence RMSFEs from that period are extremely large. As a consequence, even a dataset with a small informational content could improve the forecast relative to the autoregressive model. When the forecast comparison is considered on average over the whole evaluation sample, just a few data points could erroneously suggest that a time series is informative.

The first forecast is made for 2012Q1. The 2012Q1 observation is then included in the estimation sample and a forecast for 2012Q2 is produced using the new parameters. Finally, the one-step ahead RMSFEs are collected and averaged over the evaluation window for a mean estimate of the RMSFE of the associated model, ie.

$$RMSFE_{mean}^x = \frac{RMSFE^x}{32}$$

Tables 4 and 5 present the forecasts for each bivariate regression, both the ARDL and the VAR setting for Model A and B, respectively. The first row reports the forecast error from the benchmark model, the naïve AR(1) model, which sits at 2.17. This will be the reference for all other models considered in this study. Below the first row, the first two columns of Table 3 show the absolute RMSFE of the ARDL and bivariate VAR models when an additional indicator is included. The last columns present the relative RMSFE, calculated by dividing the RMSFE from each model with the 2.17 RMSFE from the AR(1) case. It is evident that the ARDL specification has a higher forecast error than the VAR specification, highlighting the importance of the general equilibrium effects. When the VAR is extended with some of the newly created indices, the forecast of the target variables improves considerably. Among the 35 indicators based on newspaper data, 12 include additional information related to equipment investment.

Based on the long newspaper dataset (Model A), consisting of "Handelsblatt" data, the IFs politics and finance appear informative in relation to equipment investment and reduce the forecast error by about 10% (Table 4). The most interesting topics are business cycle, international conflicts and German politics.

TABLE 4

Newspaper-based indices and equipment investment for Model A

Investment Factor/Topic	absolute RMSFE			relative RMSFE	
	ARDL	VAR		ARDL	VAR
Naïve model	2.17	2.17		1	1
Markets	2.03	2.13		0.93	0.98
Business Cycle	2.11	2.03		0.97	0.93
Company Reports	2.44	2.61		1.12	1.20
Finance	2.21	1.96		1.02	0.90
Interest Rates*	2.16	2.27		0.99	1.05
International Bourses**	2.22	2.21		1.02	1.02
Personal Finance***	2.28	2.18		1.05	1.00
SME Finance	2.22	2.23		1.02	1.03
Small German Banks	2.17	2.18		1.00	1.00
Politics	2.18	1.92		1.01	0.88
Globalisation	2.25	2.19		1.03	1.01
Structural Reforms	2.15	2.14		0.99	0.98
International Conflicts	2.20	1.96		1.01	0.90
German Politics	2.17	1.95		1.00	0.90
Technology	2.33	2.31		1.07	1.06
R& D	2.28	2.23		1.05	1.03
Energy Change-over	2.17	2.17		1.00	1.00

Model A: Investment factors and topics based on a Handelsblatt text corpus. Average root mean-squared forecast errors (RMSFE) over an evaluation window from 2012Q1 to 2019Q4. First row shows the RMSFE of an AR(1) model for the growth rate of equipment investment. Each row thereafter corresponds to a model, where the respective topic has been included as a second regressor. The relative RMSFEs have been calculated with respect to the AR(1) model. Green indicates that the forecast error is at least 10 % lower and yellow that it is at least 5 % lower.

For the shorter newspaper sample (Model B), several financial and political topics again appear important: German fiscal policy, international treaties and SME financing (Table 5). On the technology side, the topic of IT Investments also significantly improves the forecast from equipment investment, by around 10%. An in-depth look at how and why these topics are important is presented in the Section V.

TABLE 5

Newspaper-based indices and equipment investment for Model B

Investment Factor/Topic	absolute RMSFE			relative RMSFE	
	ARDL	VAR		ARDL	VAR
Naïve model	2.17	2.17		1	1
Markets (Mod. B)	2.52	2.41		1.16	1.11
German Corporates	2.73	1.93		1.26	0.89
Business Cycle	2.16	2.25		0.99	1.04
Labour Market	2.53	2.29		1.16	1.05
Commercial Real Estate	2.21	1.98		1.02	0.91
Finance (Mod. B)	2.75	2.59		1.27	1.19
Personal Finance	2.49	2.49		1.15	1.15
Monetary Policy	2.93	3.06		1.35	1.41
SME Finance	2.19	1.95		1.01	0.90
Politics (Mod. B)	3.15	2.12		1.45	0.98
International Conflicts	2.20	2.48		1.01	1.14
EU	2.31	2.38		1.07	1.10
German Fiscal Policy	2.53	2.00		1.17	0.92
International Treaties	2.31	1.91		1.06	0.88
Globalisation	2.48	2.49		1.14	1.14
Technology (Mod. B)	2.40	2.06		1.11	0.95
Energy	2.66	2.43		1.22	1.12
IT Investments	2.57	1.95		1.18	0.90

Model B: Investment factors and topics based on text corpora featuring Süddeutsche Zeitung, Welt and Handelsblatt. Average root mean-squared forecast errors (RMSFE) over an evaluation window from 2012Q1 to 2019Q4. First row shows the RMSFE of an AR(1) model for the growth rate of equipment investment. Each row thereafter corresponds to a model, where the respective topic has been included as a second regressor. The relative RMSFEs have been calculated with respect to the AR(1) model. Green indicates that the forecast error is at least 10 % lower and yellow that it is at least 5 % lower.

As noted, all models rely only on past information to forecast equipment investment, while one of the main advantages of the indices produced in this feasibility study is their timely availability. To exploit this advantage, the indicators are tested in the MF-VAR framework next.

In the second month of each quarter, the detailed national accounts figures for the previous quarter are published. Thus, a nowcast for the current quarter can be performed as this data is released (Nowcast 1). As an example, Table 6 summarizes the availability of data for the first quarter following a National Accounts release in February. It includes additional monthly indicators. For Nowcast 1, no value of industrial production or new orders is available for the current quarter. However, all newspaper-based indexes for January are known for Nowcast 1.³ A second nowcast (Nowcast 2) will be conducted at the end of the following month, when industrial data for January becomes available.

TABLE 6

Example of data availability for various monthly indicators

	Nowcast 1- February			Nowcast 2 - March			Nowcast 3 - April		
	Indicator published for			Indicator published for			Indicator published for		
	Jan	Feb	Mar	Jan	Feb	Mar	Jan	Feb	Mar
Industrial production	0	0	0	1	0	0	1	1	0
New orders	0	0	0	1	0	0	1	1	0
Media outlets data	1	0	0	1	1	0	1	1	1
Interest rates	1	0	0	1	1	0	1	1	1

Data availability at different points in time when nowcasts are created – towards the end of February, March and April. A “1” indicates that the release of the monthly indicator for the corresponding month is available. For example, at the end of March the January industrial production is available.

Next, we estimate several MF-VAR models with different specifications of the dataset. Present in all models is the quarterly level of equipment investment.⁴ We test an array of control variables such as industrial production, new orders, interest rates, and survey data (capacity utilization and ifo business climate) and find that industrial production is the most informative regarding equipment investment. We then introduce the IFs and topics that were most informative from the previous analysis – those are Business cycle, Finance and Politics.

³ One could wait until the last day of the month to incorporate the newspaper articles for February. Typically, however, nowcasts on the day of the release of national accounts are most impactful. In that case, another option is to include data only up to the day of the release, say 25th of February in the example above. To keep the implementation simple, we opt for using only media outlets data for January in Nowcast 1.

⁴ To facilitate the comparison to the other results, after the estimation in levels we calculate the growth rates of equipment investment and use those for the RMSFE.

The mixed-frequency model relies on the disaggregation of the quarterly series in monthly values, which is achieved via the Kalman filter. The relationship between these two sets of variables is part of the model parameters space and therefore estimated. Therefore, to reduce the parameter uncertainty, in this section we work with the longer dataset (Model A).

The results are presented in Table 7. We find that, even one month of newspaper data improves the nowcast of investment equipment by at least 5%, while the business cycle topic achieves more than 10% improvement. These numbers are compared to the forecast that can be done using more standard approaches with balanced data – using newspaper data only until the last quarter, for which investment equipment data is also available. The second month leads to higher improvement, however, Nowcast 2 is the first time when industrial production data is also available (see Table 6). Hence, the gain cannot be completely attributed to the newspaper data. The final month of data does not appear to be as informative, as the first two. This suggests a lag structure, where newspaper articles lead the investment data by about two months.

TABLE 7

Nowcasts for equipment investment, using a mixed frequency VAR

	absolute RMSFE			relative RMSFE		
	Nowcast 1	Nowcast 2	Nowcast 3	Nowcast 1	Nowcast 2	Nowcast 3
Naïve model	2,17			1		
Business cycle	1.93	1.79	1.77	0.89	0.83	0.82
Finance	2.06	1.91	1.92	0.95	0.88	0.89
Politics	2.04	1.80	1.99	0.94	0.83	0.92

Nowcasts for the equipment investment produced towards the end of the first, second, and third month after the release of the national accounts data. Green indicates that the forecast error is at least 10% lower and yellow that it is at least 5% lower.

Having singled out the news-based time series with the best forecasting properties, it is worth asking whether the related content may be subsumed to an Investment Narrative. If we have an idea about the story beneath the data, we may be able to catch a glimpse of economic agents' convictions that influence their decisions – of the “stories that offer interpretations of economic events, or morals, of hints of theories about the economy” present in a society, to quote Shiller (2020, p. 792). In our case, isolating an Investment Narrative may help to solve the puzzle cited at the outset of this paper: why was corporate investment rather weak for a prolonged period of time? Note that the direction of causation can be bi-directional. A narrative may contribute to low investment levels, but it can also be driven by weak investment dynamics, thereby reinforcing other dampening factors.

V. Investment Narratives

As the estimates in the previous section have shown, several IFs derived from LDA Model A are able to improve the forecasts for equipment investment considerably, while the predictive power of Model B is non-existent for most of the topics and just weak for IF Technology. Model A is based on the Handelsblatt-only corpus using a theory-informed query that required at least five investment-related terms from a pre-determined list consisting of 20 terms to co-occur in a text for a newspaper article to be included into the analysis corpus. Table 7 concludes that the topic “Business Cycle” as well as the IFs “Politics” and “Finance” have some potential for nowcasting.

In this section, we analyse the content from which these time-series originate more closely. Furthermore, we relate the results to the concept of the economic narrative, for which we propose a confined yet flexible definition that lends itself to quantification by LDA. It considerably transcends Shiller’s (2020) intuitive deduction.

A. *Properties of a Narrative*


As noted above, definitions of economic narratives tend to lack coherence. To sharpen the concept, a definition has been proposed (Müller et al. 2018) that resorts to the concept of the media frame (Entman 1993). A media frame typically contains four elements – a problem definition, a problem diagnosis, moral judgements and possible remedies. A media narrative, according to our definition, consists of one or several frames plus one or several protagonists (persons, institutions, or social groupings such as nations, classes, etc. whose relationships are (often) antagonistic and may change over time; plus certain events, that are chronologically integrated and that are (often) assumed to constitute causal relationships. The results of LDA models can guide researchers to formulate media narratives, thereby capturing exogenous dynamics. By isolating the underlying meaning of content-derived time-series, economists get an impression of what they have actually measured. The results may also contain hints regarding potentially fertile additional variables for their forecasts. LDA in its original form is well-suited for the identification of media frames (di Maggio et al., 2013). Frame being an inherently static concept and LDA being a static method, they fit together well over limited time-horizons and for thematically limited text corpora. Over longer time-horizons, however, the correspondence between research object and method is less obvious (Müller et al., 2021): problem definitions may change due to enhanced knowledge, as do moral judgements; new protagonists may enter the plot; new events may prompt observers to look at former ones in a different light. LDA produces lists of Top Words and Top Articles that contain the four elements of a frame as well as protagonists; peaks of frequency distributions over time of certain topics indicate key events. This open concept is applied to the best performing time-series derived from the investment analysis (sec. IV).

B. *The best-performing Topics*

As Tables 4 to 7 show, political factors are among the most influential in our empirical analysis. This applies particularly to the topics *German Politics* and *International Conflicts*, the former dealing with the quarrels of parties and politicians in Germany and, to some degree, at the EU level, while the latter captures global as well as European developments. Figure 5 depicts both topics’ frequencies and highlights their peaks that are associated with selected events. Over the quarter-century covered by our data, both topics have trended upwards. Starting soon after the turn

of the millennium, the topics' frequencies suggest that political influences, at home and abroad, have become more important to business investment.

Figure 5: Topics German Politics and International Conflicts, selected events*



*share of analysis corpus, three-month moving average

German Politics

Protagonists: The topic's characteristic words provide an impression of its dominant players, namely "Merkel", "SPD", "Government", "CDU", "Party", "Schröder", "Finance Minister", "European Commission", as well as the locations where the narrative unfolds, i.e. "Berlin", "Brussels".

Events: As the frequency graph in Figure 5 shows, domestic politics are prominent in the investment discourse, with federal elections representing recurrent focal points.

Framing: Besides dealing with the usual disputes over government top positions, the Top Articles highlight debates about specific government policies. An article, published on 4 February 2003 is dealing with labour market and welfare state reforms and the conservatives' acclaimed willingness to cooperate with the at the left-leaning government. Another, published on 26 June 1999, reports that the federal government plans to prop up municipal investment. Several articles are concerned with enhancing the EU's efficiency. In November 2019, a story deals with reactions to the governing coalition's plans to introduce a "minimum pension" scheme, a piece of redistributive social policy that its opponents considered as economically detrimental. The common thread of this topic could be summarized as follows: private (and government) investment is sluggish in

Germany (*problem definition*), due to a lack of willingness to reform (*problem diagnosis*), thereby depriving German citizens of possibly higher levels of welfare (*moral judgement*); *potential remedies* include business-friendly reforms and a more coherent governance structure at the EU level.

International Conflicts

Protagonists: The actors in this topic are countries and their leading figures, as can be found in the Top Words, i.e. “China”, “Russia”, “USA”, “Trump”, “Putin”, “Turkey”, “Iran”, Hongkong”. Locations where events of the narrative take place are also detectable, namely “Beijing”, “Moscow”, “Washington”, “Shanghai”.

Events: The topic’s frequency is characterized by peaks that are driven by distinct events. The Iraq war and the ensuing transatlantic tensions, the global financial crisis and the Euro crisis, the Russian annexation of Crimea, the Brexit vote and the election of Donald Trump as well as the trade war he triggered drive this narrative.

Framing: The Top Texts of the topic deal with geostrategic tensions of various kinds. On 3 February 2002 a top article’s headline reads: “Bush’s Taiwan policy angers China”. Another one deals with China’s influence in Africa and in Myanmar. Yet other ones concern the conflict between Iran and Israel after the US withdrawal from the anti-nuclear weapons treaty with the Mullah state, or the situation in Crimea after Russia annexed it respectively. In a broad sense, the topic International Conflicts captures the weakening of the Western-dominated international order (*problem definition*). Differences between the USA, Germany, France and other European NATO members over the Iraq campaign, and later during the Trump presidency, laid bare deeper tensions within the West that could be interpreted as a relative decline of the West’s power (*problem diagnosis*). The Global Financial Crisis showed the flaws in US-style capitalism, the Euro Crisis did the same for the European model (*problem diagnosis*). The surge in populist politics that culminated in the Brexit referendum and the election of Donald Trump in 2016 can be considered to be further steps in the decline of the West. The subsequent rise of uncertainty proves detrimental to business (*moral judgement*). More international cooperation and more predictable politics could ease the tensions (*potential remedies*).


Other forecast-enhancing time-series

Two more time-series were found to improve the forecast. Their content is more straightforward as far as investment is concerned: reporting on financial conditions and the business cycle.

The *IF Finance* stays at rather high levels over the 1990s and early 2000s, before succumbing to a somewhat more relaxed mode in the years preceding the global financial crisis. Afterwards it increases, peaking in 2011 at the height of the Euro crisis, and declines markedly in the following years (Figure 2). Since only the combined IF, rather than its individual topics, contributes to

forecasting accuracy, we consult its topic composition (Figure 6).⁵ As the graphs indicate, German business investment has become less dependent on credit financing, traditionally provided by small German banks, as international financial markets have come to play a more important role in company financing (market developments are partly captured in terms of Personal Finance). Interest rates, particularly those set by central banks, played an important role during years of financial stress in Germany in 2003 and during the Euro crisis, but have largely vanished in importance since the ECB started quantitative easing in 2015 and interest rates across the duration spectrum stayed persistently at ultra-low levels.

Figure 6.: IF Finance – frequencies of individual topics*




*shares relative to analysis corpus, three-month moving average

The topic *Business Cycle* contains reports about short to medium-term growth perspectives of the German economy, such as news on early indicators, survey results and such. The topic's frequency distribution shows a peak in the early 2000s, when anxiety about Germany being the "sick man of

⁵ The graphs are partly derived from combined topics that are closely related. "Personal Finance" combines both topics with this label (see above, Table 2), "International Bourses" encompasses two topics on international stock market developments. "Interest Rates" is a combination of the topics "Central Banks" and "Bond Markets".

Europe” dominated the national mood, before setting off to a secular decline (Figure 7), indicating the greater openness of the German economy where investment is no longer closely related to domestic demand, but rather influenced by global factors.

*Figure 7: Topic Business Cycle**



*share of analysis corpus, three-month moving average

C. A German Investment Narrative

The examination of the topics and IFs that came out of our econometric exercise can be interpreted as several strands of a single investment narrative. This economic story is set in an outside world where forces are at play that act upon an inside world inhabited by companies, investors, and financial institutions. The outside world evolves according to two narrative threads: one about domestic politics and another one about international relations. Both are afflicted by conflicts whose recurrent culminations (international crises, federal elections) structure the narrative; before the financial crisis (Brexit, the election of Angela Merkel, the Trump presidency and so on) the world was somewhat different than afterwards. What's more, these culminations shift the conditions under which the inside world operates. The inhabitants of the inside world, for their part, adapt to changing circumstances. Traces of these changes can be found in Figures 6 and 7;

e.g. companies strive to become less dependent on credit and on domestic demand, the banking sector consolidates etc.

As the narrative evolves, *events* take place and fade, *protagonists* come and go. However, the basic *frame* seems to be fairly stable over time. It asserts that detrimental forces hinder companies to live up to their full potential.

Subsuming the empirical evidence, a German Investment Narrative can be formulated as follows:

Over the past quarter-century business investment in Germany has become increasingly driven by political developments, both at the international and at the domestic level, that have caused subpar levels of business investment (*problem definition*). At the domestic level, subsequent governments have refrained from tackling deep-seated structural problems in earnest; at the international level the decline of Western dominance has brought about increased uncertainty (*problem diagnosis*), epitomized in a series of financial and security crises (*events*). To blame are quarrels between domestic government and opposition leaders as well as between international leaders, such as Donald Trump and Wladimir Putin (*protagonists*), all of whom tend to pursue their individual political interests rather than the common good (*moral judgement*). *Possible remedies* include a more efficient EU and welfare state reforms at home.

It does not really matter if, or to what extent, this story is true. It does not even matter if the newspaper's readers themselves believe it. What matters is that it is the underlying thread of reported reality, a latent narrative, that media users *believe others to believe*. This “third-person effect” (e.g. Lischka 2016) influences expectations and economic decision-making and thereby real economic activity.

VI. Conclusions

In this paper, we demonstrated that media-based indicators have the propensity to improve econometric forecasts of business equipment investment considerably. Newspaper data carries important information both for future developments of investment (forecasting) as well as for current developments (nowcasting). Regarding the latter, we find that even one month of media-based indicators improves the nowcasting performance with respect to RMSFE by at least 5%.

One contribution of this paper is the exploratory inductive approach concerning the media data, which differs considerably from Shiller's (2020) intuitive deduction. At the outset of our quest, we tested a host of different media as well as strategies to construct sub-corpora to be analysed as well as different calibrations of LDA models. Time series constructed using these models were implemented in econometric models. The ones that yielded the best results were inspected more closely regarding their content. We conducted a narrative analysis guided by a suitable definition (Müller et al. 2018, 2021). Here, we specified it to capture a business investment narrative.

The narrative we found is indeed capable of enhancing forecasts and improving researchers' understanding of the investment process in general, as it introduces exogenous developments as well as economic sentiment, news and relevant events to the analysis.

Our formula enables researchers to conduct macro content analyses. After all, by isolating the underlying meaning of content-derived time-series, economists get an idea of what they have actually measured. However, as in other studies in this vein, a clear understanding of the microeconomic foundations is still missing. Future research may draw from behavioral economics or psychology, but also from communication science to better fathom the effects of narratives on human behavior.

News content-based time series are a promising way to produce indicators to be used in business cycle forecasting. However, LDA in its original form has the property of yielding slightly different results with each run, even when it is conducted using the same data and parameter settings. This is particularly annoying when LDA models need to be recalculated on a regular, e.g. monthly, basis. Adding new data leads to variations that hinder the development of consistent time series. These obstacles can be overcome by applying a novel method called *RollingLDA* (Rieger, Jentsch and Rahnenführer 2021). Contrary to the classic static LDA used in this paper, it allows for alterations in topics over time, a technique that seems capable of capturing the fluid nature of media narratives and collective memory, as shown in Müller et al., (2021). This approach has two major advantages: it produces consistent time-series over long periods, and it markedly reduces the arbitrariness of traditional LDA. The latter achievement is particularly due to the combination of the *RollingLDA* method with *LDAPrototype* (Rieger, Jentsch and Rahnenführer 2020) at various stages of the process.

Further research in this vein should focus on developing reliable narrative indicators for business investment and other variables relevant to business-cycle forecasts such as private consumption and exports.

References

- Baker, S. R., N. Bloom and S.J. Davis (2016). Measuring Economic Policy Uncertainty. *The Quarterly Journal of Economics*, 131(4), 1593–1636.
- Blei, D. M., A. Y. Ng, and M. I. Jordan (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3, 993–1022.
- Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica*, 77(3), 623–685.
<https://doi.org/10.3982/ECTA6248>
- Chang, J., J. L. Boyd-Graber, S. Gerrish, C. Wang and D. M. Blei (2009). Reading Tea Leaves: How Humans Interpret Topic Models. *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>
- Cross, J. L., C. Hou, and A. Poon (2020). Macroeconomic Forecasting with Large Bayesian VARs: Global-Local Priors and the Illusion of Sparsity. *International Journal of Forecasting*, 36(3), 899-915.
- Davis, S. J. (2019). Rising Policy Uncertainty. *NBER Working Paper* 26243,
<http://www.nber.org/papers/w26243>
- DiMaggio, P., M. Nag and D. M. Blei (2013). Exploiting affinities between topic modeling and the sociological perspective on culture: Application to newspaper coverage of U.S. government arts funding. *Poetics*, 41(6), 570 – 606.
- Driver, C., N. Meade (2019). Enhancing survey-based investment forecasts. *Journal of Forecasting* 38, 236– 255. <https://doi.org/10.1002/for.2567>
- Döhrn, R., G. Barabas, B. Blagov, A. Fuest, P. Jäger, R. Jessen, M. Micheli und S. Rujin (2018). Die wirtschaftliche Entwicklung im Inland: Aufschwung setzt sich fort, Gefährdungen nehmen zu. *RWI Konjunkturberichte* 69 (3), 21-56.
- Entman, R. M. (1993). Framing: Toward Clarification of a Fractured Paradigm. *Journal of Communication* 43(4), 51–58.
- Ellingsen, J., V.H. Larsen and L. A. Thorsrud (2020). News media vs. FRED-MD for macroeconomic forecasting. *BI Norwegian Business School CAMP, Working Paper Series*; 08/2020.
- Gefang, Deborah, G. Koop, and A. Poon (2020). Computationally Efficient Inference in Large Bayesian Mixed Frequency VARs. *Economics Letters* 191 (June 1, 2020), 109120.
<https://doi.org/10.1016/j.econlet.2020.109120>.

- Giordano, C., M. Marinucci and A. Silvestrini (2021), Forecasting corporate capital accumulation in Italy: the role of survey-based information. *Banca d'Italia Occasional Papers* 596.
- Ghysels, E. and M. Marcelliono (2018), Applied Economic Forecasting using Time Series Methods. *Oxford University Press*, ISBN: 978-0190622015.
- Kálmán, R., and R. Bucy (1961), New Results in Linear Filtering and Prediction Theory. *Journal of Fluid Engineering* 83 (1), 95-108. <https://doi.org/10.1115/1.3658902>.
- Koop, G., St. McIntyre, J. Mitchell, and A. Poon (2020). Regional Output Growth in the United Kingdom: More Timely and Higher Frequency Estimates from 1970. *Journal of Applied Econometrics* 35, no. 2 (2020), 176–97. <https://doi.org/10.1002/jae.2748>.
- Lasky, M. (2018). CBO's Model for Forecasting Business Investment. *CBO Working Paper* 2018-09.
- Larsen, V. H. and L.A. Thorsrud (2019). Business Cycle Narratives. *CESifo Working Paper Series* 7468, CESifo.
- Lischka, J. A. (2016). Economic News, Sentiment, and Behavior: How Economic and Business News Affects the Economy. *VS Verlag für Sozialwissenschaften*. <https://doi.org/10.1007/978-3-658-11541-8>
- Müller, H., G. von Nordheim, K. Boczek, L. Koppers and J. Rahnenführer (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., J. Rieger and N. Hornig (2021). 'We're rolling'. Our Uncertainty Perception Indicator (UPI) in Q4 2020: introducing RollingLDA, a new Method for the Measurement of Evolving Economic Narratives. *DoCMA Working Paper* #6, Mar 2021. <http://dx.doi.org/10.17877/DE290R-21974>
- Rapach, D.E. and M.E. Wohar (2007). Forecasting the recent behavior of US business fixed investment spending: an analysis of competing models. *Journal of Forecasting* 26, 33-51. <https://doi.org/10.1002/for.1010>
- Rieger, J., J. Rahnenführer, and C. Jentsch (2020). Improving Latent Dirichlet Allocation: On Reliability of the Novel Method LDAPrototype. *Natural Language Processing and Information Systems, NLDB 2020*, 118–125. Springer International Publishing. https://doi.org/10.1007/978-3-030-51310-8_11
- Schorfheide, F. and D. Song (2015). Real-Time Forecasting with a Mixed-Frequency VAR Model. *Journal of Business & Economic Statistics* 33, no. 3, 366-380. <https://doi.org/10.1080/07350015.2014.954707>.

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. NJ: *Princeton University Press*.

Shiller, R. J. (2020). Popular economic narratives advancing the longest U.S. expansion 2009–2019. *Journal of Policy Modeling* 42 (2020) 791–798

Thorsrud, L.A. (2016). Words are the new numbers: A newsy coincident index of business cycles. *Centre for Applied Macro- and Petroleum economics, BI Norwegian Business School, Working Paper 44*