



Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Production, Manufacturing and Logistics

News-based forecasts of macroeconomic indicators: A semantic path model for interpretable predictions

Stefan Feuerriegel^{a,*}, Julius Gordon^b^aETH Zurich, Weinbergstr. 56/58, Zurich 8092, Switzerland^bUniversity of Freiburg, Platz der Alten Synagoge, Freiburg 79098, Germany

ARTICLE INFO

Article history:

Received 12 May 2017

Accepted 29 May 2018

Available online xxx

Keywords:

Forecasting

Text mining

Financial news

Macroeconomic indicators

Partial least squares

ABSTRACT

The macroeconomic climate influences operations with regard to, e.g., raw material prices, financing, supply chain utilization and demand quotas. In order to adapt to the economic environment, decision-makers across the public and private sectors require accurate forecasts of the economic outlook. Existing predictive frameworks base their forecasts primarily on time series analysis, as well as the judgments of experts. As a consequence, current approaches are often biased and prone to error. In order to reduce forecast errors, this paper presents an innovative methodology that extends lag variables with unstructured data in the form of financial news: (1) we apply a variety of models from machine learning to word counts as a high-dimensional input. However, this approach suffers from low interpretability and overfitting, motivating the following remedies. (2) We follow the intuition that the economic climate is driven by general sentiments and suggest a projection of words onto latent semantic structures as a means of feature engineering. (3) We propose a semantic path model, together with estimation technique based on regularization, in order to yield full interpretability of the forecasts. We demonstrate the predictive performance of our approach by utilizing 80,813 ad hoc announcements in order to make long-term forecasts of up to 24 months ahead regarding key macroeconomic indicators. Back-testing reveals a considerable reduction in forecast errors.

© 2018 Elsevier B.V. All rights reserved.

1. Introduction

Macroeconomic developments, such as cyclic downturns or the economic circumstances associated with the U. S. subprime crisis, affect firm operations in multiple ways and represent direct challenges to management (e. g. [Demyanyk & Hasan, 2010](#); [Goudie & Meeks, 1982](#)). Examples include changes in the price of goods and raw materials, as well as the impact on overall demand, supply chain utilization and even operational processes ([Xu, Pinedo, & Xue, 2017](#)). Therefore, firms are interested in foreseeing the future economic climate in order to manage operations accordingly and hedge potential risks. In this context, operational research (OR) has a long tradition of addressing such risks ([Demyanyk & Hasan, 2010](#)). Our discipline has thus contributed to anticipating a variety of developments at a macroeconomic level, including financial distress ([Geng, Bose, & Chen, 2015](#)), liquidity risks ([Shaik, 2015](#)), credit risks ([Akkoç, 2012](#); [Desai, Crook, & Overstreet, 1996](#)), finan-

cial crises ([Demyanyk & Hasan, 2010](#); [Huang, Kou, & Peng, 2017](#)), currency crises ([Sevim, Oztekin, Bali, Gumus, & Guresin, 2014](#)) and bankruptcy ([Du Jardin, 2015](#); [McKee & Lensberg, 2002](#); [Sun & Shenoy, 2007](#)), especially in the financial sector ([Tam & Kiang, 1992](#))

Future expectations regarding the macroeconomic environment play a critical role in the decision-making process for many organizations ([Xu et al., 2017](#)). Hence, decision-makers across all sectors must analyze the current economic environment and form accurate expectations about future economic trends in order to support the operational strategy of organization and long-term management. As a result, macroeconomic variables, and the accurate prediction thereof, form the basis for a wide array of OR models (e. g. [Calabrese, Degl'Innocenti, & Osmetti, 2017](#); [Fethi & Pasiouras, 2010](#); [Gutiérrez & Lozano, 2012](#)).

The importance of accurate long-term forecasts for firm operations has driven the extensive amount of research conducted with respect to macroeconomic predictions. Specific examples from the OR domain include short-term predictions of asset-related values, including government bonds ([Tay & Cao, 2001](#)) and stock

* Corresponding author.

E-mail addresses: sfeuerriegel@ethz.ch (S. Feuerriegel), juliusgordon89@gmail.com (J. Gordon).<https://doi.org/10.1016/j.ejor.2018.05.068>

0377-2217/© 2018 Elsevier B.V. All rights reserved.

indices (e. g. Huang, Nakamori, & Wang, 2005; Kung & Yu, 2008; Oztekin, Kizilaslan, Freund, & Iseri, 2016). Further research focuses on forecasting macroeconomic indicators of single countries (e. g. Mahmoud, Motwani, & Rice, 1990) or the relationship between countries (Sermpinis, Theofilatos, Karathanasopoulos, Georgopoulos, & Dunis, 2013). Other works propose agent-based simulations to study the behavior of human forecasters (Bovi & Cerqueti, 2016).

Previous efforts at forecasting macroeconomic indicators have made use of various input features and methodologies. Historic time series data is a staple input for macroeconomic forecasting, and has been applied to make both short- and long-term predictions (e. g. Jansen, Jin, & de Winter, 2016; Mahmoud et al., 1990; Sermpinis et al., 2013). A prevalent alternative is the subjective judgments of professional forecasters (Matsypura, Thompson, & Vasnev, 2018) such as those used by the European Central Bank. However, both time series models and subjective forecasts suffer from severe prediction errors (Jansen et al., 2016), possibly because they cannot benefit from predictors that better measure the state and outlook of the economy.

It is plausible that macroeconomic forecasting could also benefit from the advances of the big data era (Mortenson, Doherty, & Robinson, 2015) and potential improvements in predictive modeling. Especially unstructured data, such as news, promises further insights into global systems and the heterogeneous trends and events that occur within them. Hence, news represents a powerful source of information for financial forecasting and is thus likely to serve as an intriguing – yet currently rare – predictor for macroeconomic forecasts.

Our research therefore aims to improve macroeconomic forecasts using financial disclosures. Financial news, in particular, conveys rich information about expected firm performance that often goes beyond pure numeric data (e. g. Tsai & Wang, 2017). Text mining of financial news has been successfully used in conjunction with predictive analytics to infer profitable short-term investment decisions in capital markets (Feuerriegel & Prendinger, 2016; Nassirtoussi, Aghabozorgi, Wah, & Ngo, 2014). However, it is unclear whether financial filings – which should essentially mirror the current health of the firm – also provide long-term prognostic capacity.

To make long-term predictions of macroeconomic indicators, we construct an approach as follows: we use financial news and insert corresponding linguistic features into various machine learning methods. Our specific setting represents a major difference from previous news-based forecasts, in the sense that the outcome variable is reported in monthly or quarterly resolution. This results in fewer observations and thus extremely high-dimensional predictor matrices with severe risks of overfitting. To overcome this challenge, we develop a form of feature engineering based on semantic projections and, on top of that, propose a semantic path model that fulfills the demand of practitioners by being fully interpretable. Afterwards, we conduct an out-of-sample evaluation of the predictive performance. The results demonstrate that our approach is superior in terms of relative performance when compared to common time series models that serve as a benchmark.

The remainder of this paper is structured as follows. Section 2 reviews related research on macroeconomic forecasting in order to outline the aforementioned research gap. As a remedy, Section 3 explains our methodology for making news-based macroeconomic predictions, with a particular focus on the proposed projections to latent semantic structures. Section 4 reports the datasets, based on which Section 5 compares the predictive performance of both traditional time series models and news-based forecasts. A discussion of managerial implications follows in Section 6, while Section 7 concludes with an outlook on future research.

2. Background

2.1. Predictive models for macroeconomic forecasts

Beyond human predictions (Matsypura et al., 2018), common approaches to macroeconomic forecasting include various time series models (Allen & Morzuch, 2006), for instance, auto-regressive moving-average (ARMA) models and their multivariate variations, such as vector autoregression and vector error correction models. However, their use entails certain challenges, since they must be adapted to cope with structural changes in the underlying systems (Litterman, 1986). Nevertheless, time series models still appear to be the standard benchmark against which the predictive performance of alternative approaches is measured (de Gooijer & Hyndman, 2006).

Since vector autoregression and its variants are limited to a few variables, researchers have proposed alternative models, such as Bayesian vector autoregression (BVAR). The BVAR accounts more accurately for uncertainty regarding the structure of the economy and is utilized to handle instances with dozens of relevant variables (Litterman, 1986). Its use for forecasts requires extensive computational resources, especially when dealing with many predictors (Carriero, Clark, & Marcellino, 2015; Litterman, 1986).

Bridge equations have been suggested as an alternative means of incorporating relevant explanatory variables into predictive models. This approach is a form of linear regression that time-aggregates economic time series to bridge the information gap between low and high frequency indicators. It is a widely used method for forecasting macroeconomic variables (Jansen et al., 2016). However, these forecasts are generally limited to one or two periods ahead.

A different concept is represented by leading indicators. These are supposed to provide prognostic capabilities for future changes in the economic outlook. Examples include major stock indices, the Ifo Business Climate Index from Germany, or the University of Michigan Consumer Sentiment Index for the United States. These indices can yield important information regarding future cyclical fluctuations of the general economy (Allen & Morzuch, 2006) and they have thus been used to forecast imminent turning points in the business cycle (Layton & Smith, 2007). Predictive models using leading indicators have been applied, for instance, in single variable ARMA models, multivariate models, diffusion indices and BVAR (Allen & Morzuch, 2006; Stock & Watson, 2002). However, it is not clear whether these indices have a prognostic potential for future developments.

Since previous approaches to macroeconomic forecasting have primarily focused on using various forms of economic data as dynamic predictors, the use of qualitative or unstructured data, such as financial news, deserves the attention of researchers. However, the inclusion of larger numbers of predictors in traditional time series models can prove to be problematic, as such models typically struggle with such settings (Jansen et al., 2016; Litterman, 1986). The above literature review thus motivates our choice of models: we specifically decided to utilize different machine learning models that can cope with non-linear relationships and large numbers of textual features.

2.2. Predictive text mining with financial news

News-based predictions of macroeconomic indicators constitute an innovative area of research. Due to the scarcity of related publications, we decided to extend our literature overview and also include news-based forecasts of stock returns in addition to macroeconomic variables.

Text mining of financial news have been successfully used in conjunction with predictive analytics to infer profitable short-term

investment decisions in capital markets (Feuerriegel & Prendinger, 2016). Previous approaches have largely focused on predicting the response to financial news in the form of *short-term* reactions of *stock markets* (Nassirtoussi et al., 2014), whereas we develop a forecasting methodology for the general *macroeconomic environment* over *long-term forecast horizons*. Text mining of financial news for predictions in capital markets has been extensively studied in recent years; for a comprehensive review we refer to Nassirtoussi et al. (2014).

Financial news contains fundamental and qualitative information that influences the expectations of market participants, and has thus driven the field of market prediction. A broad range of textual materials has been in the course of recent research, including newspaper articles from media sources and regulated filings, such as ad hoc announcements, 8-K filings and annual reports (Nassirtoussi et al., 2014). The common feature tying these previous efforts together is the focus on making short-term, usually same-day, predictions.

The challenge of working with textual data, such as financial news, lies in how to efficiently transform it into a machine-friendly form and then apply a classifier to it. Hence, text mining approaches commonly follow a similar procedure with the goal of extracting and selecting relevant features that best represent the original text (Nassirtoussi et al., 2014). In this regard, the bag-of-words approach breaks up the free-running text into single words (or n -grams) which can then be used as features. Afterwards, feature selection (sometimes joined with dimensionality reduction) aims to reduce the set of features while keeping relevant information. Common approaches in financial text mining (Nassirtoussi et al., 2014) include frequency-based statistics, i. e. term weighting. This then presents the input to machine learning classifiers. We draw upon the aforementioned approach utilizing high-dimensional input matrices, which results in one of our three text-based forecasting routines.

3. Methods

3.1. Overview

This section introduces the research framework for forecasting macroeconomic indicators based on financial news. This task entails two critical challenges: feature engineering and interpretability.

First, the input matrix is high-dimensional as the number of predictors exceeds the number of observations. For instance, the macroeconomic indicators with quarterly resolution are predicted based on a term-document matrix with 536 terms, while a single fold during cross-validation might consist of only 71 samples. We address this issue through unsupervised techniques for dimension reduction and, in addition, pioneer in replacing the statistical approaches to feature generation (e. g. Liang, Lu, Tsai, & Shih, 2016) with a supervised technique that incorporates domain-specific knowledge. More specifically, it aggregates terms according to a semantic classification with projections onto corresponding constructs, thereby yielding a low-dimensional input with only a couple of predictors.

Second, machine learning is often regarded as a black-box mechanism, whereas practitioners from operational management demand the potential for interpretation as a means to validate the computer-based forecasts with their own beliefs and expectations, as well as to gain trust in the automated decision support. Hence, we are specifically interested in models that attain a beneficial trade-off between predictive performance and explanatory power. As a remedy, we propose a semantic path model that can offer full accountability of the forecasts but yields a performance comparable to the best black-box models from machine learning.

Here we follow the intuition that the economic climate is linked to the perception (i. e. sentiments) of the involved stakeholders. Given a meaningful choice of dimensions, we can later decompose the overall forecasts in terms of the underlying constructs.

We adapt to the two aforementioned challenges by suggesting a variety of predictive methods for our forecasting task. The underlying approach for text-based forecasts consists of preprocessing of financial text, including feature extraction (cf. online appendix); optional feature engineering, in which we project individual words onto latent semantic structures; model training and selection; and multi-step ahead predictions using a test set. The individual approaches can be grouped as follows:

0. *Baseline: Time series models.* These build upon auto-regressive lags and combine them in linear and non-linear fashion.
 1. *Machine learning models with high-dimensional input.* A series of machine learning models is tested with high-dimensional input as a comparison. The models are later fed with an input matrix that is (optionally) subject to unsupervised dimension reduction techniques.
 2. *Models with semantic features.* We develop a supervised technique for feature engineering. This approach involves projections of terms onto semantic constructs that then serve as a low-dimensional feature space.
 3. *Semantic path model.* Similar to the semantic features, we develop a path model that operationalizes the semantic classification of terms into different categories relevant in a macroeconomic context in order to obtain interpretable forecasts. To circumvent overfitting from collinear constructs, we develop a novel estimation technique with an additional regularization.

We note that each of the text-based models is later trained in different variants, i. e. the models are optionally augmented by the same lags as the time series baseline.

3.2. Baseline: time series models

In order to evaluate the predictive performance of our forecasting methodology, we employ a number of benchmark models that forecast the macroeconomic time series purely on the basis of historic data points. These models include an auto-regressive (AR) process based on l lagged values. Formally, let i denote the current time period. We then predict the value in time step $i+h$, i. e. that is h steps ahead. This results in

$$Y_{i+h} = \alpha + \beta_1 Y_{i-1} + \dots + \beta_l Y_{i-l} + \varepsilon_i, \quad (1)$$

with coefficients $\alpha, \beta_1, \dots, \beta_l$, inputs Y_{i-1}, \dots, Y_{i-l} and where Y_{i+h} gives the predicted indicator. In addition, we incorporate non-linear relationships by applying a random forest (RF) to historic values for each of the macroeconomic indicators. This machine learning model frequently achieves a high predictive power with little need for tuning (Hastie, Tibshirani, & Friedman, 2013).

For both AR and RF models, we experiment with different specifications that vary in their lag order: on the one hand, we utilize only 1 lag and, on the other hand, all lags of order 1 to 6. By testing both variants, we also adhere to the selection made in previous literature (e. g. Litterman, 1986; Stock & Watson, 2002). We specifically decided to predict the raw values of each of the macroeconomic indicators, instead of applying additional transformation. This choice stems from the need of practitioners who require clear expectations in common and interpretable scales. As a comparison, we also predict a first-difference term as part of our evaluation.

3.3. Models for textual predictors

All subsequent models involving textual predictors are estimated in three different variants: on the one hand, we forecast

macroeconomic indicators based purely on the narrative materials, excluding the history of the outcome variable. On the other hand, we follow the above time series models and augment the predictor matrix by 1 and 6 auto-regressive lags of the to-be-predicted macroeconomic indicator. Thus the model can further learn from additional seasonal and trend characteristics.

3.3.1. Machine learning models for high-dimensional input

As noted earlier, the document-term matrices are wide, since the number of features (words) far exceeds the number of observations. We thus choose machine learning models that are known for their performance with wide data sets, as they reduce the risk of overfitting via feature selection, regularization and dimension reduction (Hastie et al., 2013). These are as follows: least squares absolute shrinkage operator (LASSO), ridge regression, elastic net (ENET), gradient boosting (GBM), principal component regression (PCR), random forest (RF) and partial least squares regression (PLS-R).

Following common conventions, the above models are estimated based on the term-document matrix. The entries are scaled by tf-idf weighting in order to reflect how frequently a term appears in a document in the corpus. Beyond that, the experiments are repeated with different unsupervised transformations for dimension reduction that are applied to the input X as part of a feature engineering step. More precisely, we apply a principal component analysis (PCA) and compute a latent semantic analysis (LSA), which are then fed into the different machine learning models. Altogether, we obtain a total of 21 different models (i. e. 7 classifiers, each with raw tf-idf, as well as 2 unsupervised dimension reduction techniques).

3.3.2. Models with semantic features

We now present an alternative approach to feature generation. More precisely, we replace the purely statistical approach behind unsupervised transformations for dimension reduction (i. e. PCA and LSA) with a supervised routine that encodes domain-specific knowledge.

Our approach assumes that the overall macroeconomic climate is driven by the perceptions (or sentiments) of individual stakeholders. For instance, a construct “macroeconomic uncertainty” could be indicated by a term “risk”, while “increase” links to a construct “positive outlook”. Hence, we identify a new set of features z_1, \dots, z_m which represent such meaningful constructs (i. e. perceptions). Each construct z_i is represented by a linear combination of terms x_1, \dots, x_n or, more specifically, a subset $\mathcal{I}_i \subseteq \{1, \dots, n\}$ thereof; i. e. $z_i = \sum_{j \in \mathcal{I}_i} \phi_{ij} x_j$ with coefficients ϕ_{ij} . The subset then incorporates our knowledge with regard to how terms correspond to the different constructs. Here merely a categorization \mathcal{I}_i of words is needed in the form of a simple wordlist that neither has weights nor is mutually exclusive. As a result, the high-dimensional input matrix is projected onto a low-dimensional feature space based on the the pre-specified semantic structure.

The weights ϕ_{ij} are computed in a supervised manner by making use of the response Y . That is, the weights ϕ_{ij} are chosen such that the z_i explain as much of the covariance between X and Y as possible. Mathematically, the computation draws upon the existing routine for projections on latent constructs (Hastie et al., 2013). This algorithm computes the coefficients for the each construct z_i similar to least squares, i. e. $\min_{\phi_{i1}, \dots, \phi_{im}} \|Y - \phi X\|_2^2$, with the difference that it involves the orthogonalization step from path modeling, as terms can appear in more than one construct (Tenenhaus, Vinzi, Chatelin, & Lauro, 2005). This highlights the differences from traditional PLS-R as a predictive model, in which the number of constructs m is subjected to tuning and *all* terms map on *all* features. Conversely, our semantic features impose a structure where

only a subset of terms maps to each feature, thereby incorporating domain knowledge.

We additionally utilize a Bayesian neural network as the underlying predictive model in order to infer Y from z_i , thereby essentially yielding a universal structure model, USM for short (Buckler & Hennig-Thurau, 2008). In fact, universal structure modeling has found to be effective in multiple applications from the fields of decision support and operations research (e. g. Oztekin, Kong, & Delen, 2011; Turkyilmaz, Oztekin, Zaim, & Fahrettin Demirel, 2013; Turkyilmaz, Temizer, & Oztekin, 2016).

3.3.3. Semantic path model

We now present the semantic path model that also involve projections onto low-dimensional feature spaces. These again build upon a meaningful representation of different perceptions of the macroeconomic climate as latent constructs. However, this approach attains full interpretability as the underlying combination construct is linear. Fig. 1 illustrates an example in which terms are mapped onto two constructs – *negative outlook* and *economic uncertainty* – which then yield the final prediction. Such path modeling techniques are common in the social sciences (cf. Aguirre-Urreta & Marakas, 2014; Aubert, Rivard, & Patry, 1996; McIntosh, Edwards, & Antonakis, 2014), where the constructs quantify different aspects of human behavior (Rigdon, Ringle, & Sarstedt, 2010; Tenenhaus et al., 2005). However, their adaptation to semantic structures for linguistic materials reveals an innovative application with considerable potential for interpreting language-based predictions. Mathematically, the model builds upon the prescribed paths of which terms $\mathcal{I}_i \subseteq \{1, \dots, n\}$ relate to construct i . Each construct is then computed by a linear combination $z_i = \sum_{j \in \mathcal{I}_i} \phi_{ij} x_j$, i. e. a projection of words onto latent semantic constructs. Then, the final prediction is an additional linear model $\tilde{Y} = \psi_0 + \sum_{i=1}^m \psi_i z_i$. Evidently, one can directly interpret the forecast and decompose its value into individual components $\psi_1 z_1, \psi_2 z_2$, etc. Hence, decision-makers can assess the underlying reasoning of the model and evaluate the contribution of the different semantic categories to the overall outcome.

The above model can be estimated based on the algorithms for path modeling based on projections to latent structures (Tenenhaus et al., 2005), which we denote in the following:

- 1: Center X and Y
- 2: Initialize outer weights ϕ_{ij} with arbitrary values (e. g. $\phi_{ij} = 1$)
- 3: **repeat**
- 4: Compute the external approximation of latent constructs $z_i = \sum_{j \in \mathcal{I}_i} \phi_{ij} x_j$
- 5: Obtain inner weights and approximation of constructs
- 6: Calculate new outer weights ϕ_{ij}
- 7: **until** convergence of outer weights ϕ_{ij}
- 8: Obtain inner weights, i. e. path coefficients ψ_i , by ordinary least squares

$$\min_{\psi_0, \dots, \psi_m} \left\| \tilde{Y} - \psi_0 - \sum_{i=1}^m \psi_i z_i \right\|_2^2$$

Beyond the naïve estimation approach to path modeling based on projection to latent structures (Tenenhaus et al., 2005), we propose an alternative strategy: we argue that the potentially collinear relationships between the different constructs could lead to a low capacity for generalization and thus suboptimal forecasts. This differs from the majority of path models in social sciences, where constructs represent distinct concepts (e. g. Aubert et al., 1996) and thus mostly orthogonal vectors (Tenenhaus et al., 2005). However, in our case, a high negativity is likely to correspond to low positivity and vice versa. Hence, this can be addressed by replacing the ordinary least squares (OLS) estimator from the inner model in path modeling with regularized variants, namely, LASSO, ridge




Fig. 1. Illustrative depiction of our semantic path model. The dictionary-based assignment of words to semantic categories is shown by the presence of arrow-shaped paths. Here projections map the high-dimensional predictors onto latent semantic constructs.

regression and elastic net. This results in the general form

$$\min_{\psi_0, \dots, \psi_m} \left\| \tilde{Y} - \psi_0 - \sum_{i=1}^m \psi_i z_i \right\|_2^2 + \alpha_1 \|\psi\|_1 + \alpha_2 \|\psi\|_2^2, \quad (3)$$

where the coefficients α_1 and α_2 control the amount of shrinkage and are subject to hyperparameter tuning. This merely affects the magnitude of the coefficients ψ_i but maintains the linear relationship that lends itself to direct interpretation. In fact, the LASSO can set coefficients directly to zero and, accordingly, could even perform an implicit selection of relevant constructs. Closest to this are different approaches for shrinkage or the inclusion of sparsity in partial least squares regression (e. g. Chun & Keleş, 2010; Lê Cao, Rossouw, Robert-Granié, & Besse, 2008), but, to the best of our knowledge, penalization has not yet been incorporated in path modeling.

The semantic path model and PLS-R share similarities in terms of that both implement projections to latent structures. However, a clear difference becomes evident in how these mappings are defined. PLS-R follows a *data-driven* approach where all terms map onto a number of latent constructs that is subject to tuning. Thereby, the latent constructs lose their ability of interpretation and collinearity among features – such as in our high-dimensional setting – could impede generalizability. In contrast, our semantic path model focuses on a *knowledge-driven* methodology as it directly encodes a structure that lends to interpretation. The numerical results later indicate the potential performance improvements in our setting with extremely high-dimensional predictor matrices. Moreover, the proposed use of additional regularization can further bolster interpretation from more parsimonious structures and shrink the influence of collinear constructs, thereby reducing the risk of overfitting.

3.4. Model tuning

All of the previous models from both benchmarks and news-based predictors were tuned by following the default parameter search grid from Kuhn (2008). However, time series data does not lend itself to traditional random sampling used for cross-validation and we instead utilize a time-slice cross-validation with ten folds. This applies a rolling forecast origin, where the sample used for calibration contains only observations that occurred prior to the observations that form the validation set. The process is then repeated by rolling the origin of the validation set forward, thus incorporating the previous observations into the estimation procedure. The model is then retested using the next window of observations. Afterwards, the best configuration of hyperparameters is

retained, based on which the final model is fitted to the complete training data. All models are tuned with the goal of minimizing the root mean squared errors. All models preprocess the predictors by scaling and centering to zero mean and a variance of one.

4. Datasets

4.1. Financial news

To make long-term predictions of macroeconomic indicators, we use financial news in the form of ad hoc announcements. Ad hoc announcements represent regulatory filings that are obligatory for all firms listed on the German stock exchange (cf. German Securities Trade Act for details). This dataset offers a number of advantages. First, German regulations require listed companies to immediately publish any information which could potentially impact the price of their stock via this type of announcement. Second, the content of the announcement is authorized by the executive board and finally checked by the Federal Financial Supervisory Authority. The tight regulations reduce the likelihood of non-relevant content. Third, standardized distribution channels guarantee real-time access to disclosures. This allows one to derive a live picture of the industrial sector, since corporations self-report their current performance but also incorporate statements regarding their optimism, future expectations and legal risks. Altogether, we possess a rich source of information to help sense – potentially early – economic developments from a production-based view. Ad hoc announcements have been subject to extensive research revealing the benefits of their use in both practice and research, for instance, with regard to short-term investment decisions and risk predictions (e. g. Kraus & Feuerriegel, 2017).

Our corpus is composed of all 80,813 ad hoc announcements which we collected from the website of the EQS Group, the leading publisher of ad hoc announcements. Our filings are linked to publicly-traded companies and were published between July 1996 and April 2016. The dataset is split chronologically into a training set (60 %) and a test set (40 %). Announcements were selected such that our corpus includes those that are published in either English (our default analysis) or German, as well as a combined corpus of both. This facilitates later experiments that analyze the sensitivity of news-based forecasts with regard to the underlying language.

4.2. Macroeconomic variables

The ad hoc announcements are predominantly from German companies. Therefore, we selected a set of macroeconomic indica-

Table 1

List of predicted macroeconomic indicators.

Indicator	Region	Abbr.	Resolution (in use)	Source	Description	Obs.
Gross domestic product	Germany	GDP	Quarterly	OECD	Market value of all final goods and services produced	86
Gross domestic product	E. U.	GDP	Quarterly	ECB	Market value of all final goods and services produced	86
Unemployment rate	Germany	UR	Monthly	OECD	Number of unemployed people as a percentage of the labor force	260
Consumer price index	Germany	CPI	Monthly	OECD	Measures prices changes in a basket of goods and services	260
Industrial production index	Germany	IPI	Monthly	OECD	Measures the output in the industrial sector of the economy	260
Business confidence index	Germany	BCI	Monthly	OECD	Enterprises' assessment of production, orders and stocks	260
Consumer confidence index	Germany	CCI	Monthly	OECD	Households' expectations for the immediate future	260
Federal 10 year bond	Germany	Bund	Monthly	Bloomberg	German government bond denominated in euros	260

Table 2

Descriptive statistics of disclosures and macroeconomic variables. The table lists the average value of each time series by calendar year.

Year	Number disclosures	Average length	GDP: Germany	GDP: E. U.	UR	CPI	IPI	Bund	BCI	CCI
1996 (Jul–Dec)	424	277.38	481.58	1498.18	9.09	1.45	77.47	6.20	98.79	98.14
1997	1524	232.96	491.77	1534.49	10.05	1.94	80.06	5.65	100.20	98.37
1998	1911	281.13	504.56	1595.85	9.41	0.91	83.23	4.53	100.52	100.34
1999	3863	362.34	516.22	1669.71	8.58	0.59	84.16	4.55	99.51	100.85
2000	6954	402.03	529.12	1758.11	7.90	1.44	89.26	5.23	100.73	101.51
2001	8817	296.22	544.96	1839.15	8.01	1.98	89.64	4.83	99.21	100.56
2002	7351	288.17	552.32	1902.70	8.85	1.42	88.62	4.76	98.96	99.38
2003	4676	301.33	555.02	1957.47	9.84	1.03	88.80	4.09	99.25	98.25
2004	4095	309.71	567.65	2041.11	10.01	1.67	91.56	4.06	100.32	98.65
2005	4112	323.98	575.22	2114.99	11.29	1.55	94.93	3.36	100.11	99.01
2006	4221	345.60	598.31	2225.96	10.37	1.58	100.72	3.79	101.53	100.21
2007	4449	376.53	628.31	2350.20	8.80	2.30	107.59	4.24	101.94	101.74
2008	4012	382.68	640.43	2408.32	7.66	2.63	107.77	3.96	100.22	99.96
2009	3500	399.41	615.07	2322.03	7.85	0.31	89.23	3.26	97.28	97.28
2010	3054	432.50	645.01	2386.20	7.07	1.10	99.33	2.74	101.09	100.60
2011	2982	468.68	675.78	2449.54	5.91	2.07	107.89	2.65	101.95	101.74
2012	2903	477.53	689.57	2459.55	5.46	2.01	107.19	1.51	100.07	100.35
2013	3099	454.36	706.56	2484.55	5.31	1.50	107.32	1.62	100.34	100.44
2014	3031	469.42	730.98	2531.89	5.05	0.91	109.37	1.17	100.88	101.25
2015	2811	488.65	758.20	2613.64	4.69	0.23	109.76	0.51	100.71	101.05
2016 (Jan–Apr)	510	556.60	771.76	2645.19	4.31	0.17	111.06	0.09	100.62	100.47

tors with the aim of reflecting the general German economic climate. To account for the share of international firms, we also included a number of European indicators for comparison. Table 1 reports our macroeconomic indicators with both monthly and quarterly resolution.

4.3. Descriptive statistics

The average length of an ad hoc announcement over the period amounts to 366.15 words. Table 2 lists further descriptives per year. It further indicates an increasing length of the ad hoc announcements over the research period. The average number of ad hoc announcements published per year is to 4072 (for the period 1997–2015).

4.4. Semantic constructs

The proposed projections require a categorization of terms into different semantic structures. The method makes no a priori assumption regarding the structure and, hence, the final choice can be subject to the intended insights. In fact, the projections support the conventional wordlists from that are widely used in the finance domain or social sciences (the semantic path model merely requires a categorization with neither weights nor mutually-exclusive assignments). Other semantic categories might be beneficial when adapting our approach to other predictive settings, such as marketing.

In our computational experiments, the constructs should lend themselves to interpretation with regard to perceived sentiments; e. g. the valence of the outlook, the uncertainty of the economic development, and the associated confidence by the reporting companies. For this reason, we follow research on financial markets

(Nassirtoussi et al., 2014) and draw upon the Loughran–McDonald finance-specific dictionary¹, which categorizes words into different semantic categories that reflect six such perceptions. The *positive* and *negative* categories entail terms that refer to positive and negative assessments; *uncertainty* encodes terms used to signal vagueness or even risks with regard to macroeconomic development; *legal* refers to, e. g., potential liabilities in the future; *strong modal* and *tentative* (i. e. *weak modal*) measure the confidence with regard to statements, conclusions and firm developments. The assignment of words is not unique, as several words belong to more than one semantic category. For instance, all tentative (i. e. weak modal) terms are included in the uncertainty list. In total, 132 out of 2,362 entries appear in more than one construct. The choice of this dictionary is beneficial to our setting as the wordlists have been specifically designed to extract qualitative information from regulatory firm disclosures.

5. Results

This section summarizes the results of our out-of-sample predictive experiments. The predictive performance of the news-based models are evaluated relative to the benchmark models. Here we conducted two different prediction experiments. In the first experiment, the models were trained to predict the raw value of each indicator over multiple prediction horizons. In the second experiment, we looked to test the ability of our news-based approach to predict imminent turning points in the business cycle, where it predicts the change $\Delta Y_i = Y_{i+1} - Y_i$.

¹ See https://www3.nd.edu/~mcdonald/Word_Lists.html, last accessed January 12, 2018.

Table 3

Comparison of prediction performance across different quarterly macroeconomic indicators. The variable h indicates the number of time steps predicted ahead. Depending on the resolution of the predictive variable, the panels either report long-term forecasts of up to 8 quarters (such that both cases amount to 2 years). The column ΔY provides the results for the first difference or change in value from one period ahead to the current period. For each outcome and model type, only the performance of the best-in-breed model is reported (i. e. only the best-performing time series model, the best semantic path model, etc.). In each case, the table lists are the root mean squared error. The news-based models that equal or outperform the prediction error of the best benchmark model are in bold for each prediction horizon. In addition, the the P -value from the Diebold–Mariano test is stated in the parenthesis when comparing the best-in-breed model to the best baseline, as well as the corresponding model choice including lag structure and dimension reductions (if any).

Input/model	ΔY	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 8$
Predicted variable: quarterly Eurozone GDP						
Benchmark: lags	36.37	36.37	46.365	51.112	43.654	48.217
	AR6	AR6	AR6	AR6	AR1	AR1
Semantic path model	35.44 (0.261)	36.108 (0.466)	56.506 (0.902)	67.177 (0.897)	98.954 (0.998)	123.697 (0.999)
	OLS6	LASSO6	LASSO6	Ridge6	Ridge6	Ridge6
Semantic features	32.472 (0.130)	36.108 (0.466)	56.506 (0.902)	64.948 (0.879)	72.079 (0.993)	67.379 (0.974)
	RF6	LASSO6	LASSO6	PCR6	PCR6	PCR6
High-dimensional input	42.319 (0.846)	65.529 (0.995)	50.119 (0.681)	48.415 (0.366)	64.2 (0.970)	82.21 (0.994)
	LASSO	LASSO6	LASSO1	LASSO6	LASSO1	LASSO6
Sensitivity: lemmatization	33.693 (0.294)	65.19 (0.994)	48.375 (0.607)	47.968 (0.332)	51.121 (0.787)	101.873 (0.999)
	RF	LASSO6	LASSO1	LASSO6	LASSO1	LASSO6
Sensitivity: corpus	30.389 (0.101)	64.159 (0.989)	45.767 (0.472)	45.265 (0.189)	58.436 (0.93)	67.515 (0.972)
	Complete	German	Complete	German	Complete	Complete
Predicted variable: quarterly Germany GDP						
Benchmark: lags	12.264	12.264	17.381	19.354	22.265	30.398
	AR6	AR6	AR6	AR6	AR1	AR6
Semantic path model	13.412 (0.871)	33.725 (1.000)	34.699 (1.000)	57.075 (1.000)	91.805 (1.000)	92.378 (1.000)
	OLS6	LASSO6	Ridge6	Ridge6	Ridge6	Ridge6
Semantic features	12.47 (0.561)	28.078 (1.000)	26.325 (0.999)	31.74 (1.000)	46.205 (1.000)	52.095 (1.000)
	RF6	PCR6	PCR6	PCR6	PCR6	PCR6
High-dimensional input	9.921 (0.136)	23.36 (1.000)	26.7 (0.999)	44.149 (1.000)	59.875 (1.000)	63.918 (1.000)
	Elastic net	LASSO6	LASSO6	LASSO6	LASSO1	LASSO6
Sensitivity: lemmatization	9.143 (0.092)	22.86 (1.000)	25.299 (0.997)	40.808 (1.000)	57.024 (1.000)	65.838 (1.000)
	LASSO	LASSO6	LASSO6	LASSO6	LASSO1	LASSO6
Sensitivity: corpus	10.637 (0.216)	19.681 (0.999)	22.766 (0.995)	33.858 (1.000)	49.914 (1.000)	65.936 (1.000)
	German	Complete	German	German	German	German

Tables 3 and 4 compare the prediction performances as a best-of-breed measure. That is, we only report the best-performing time series baseline, the best result from semantic path models, etc. For each best-of-breed model, the performance is measured in terms of root mean squared error (RMSE). In addition, we validate the robustness of our results by employing the Diebold–Mariano test. The test statistic estimates whether our news-based models obtain statistically significant lower forecast errors compared to the benchmark models. We expect this value frequently to be low due to the small size of the test set and the overall difficulty of macroeconomic predictions. Later, the quality of the forecasts is compared across the different macroeconomic indicators; see Section 5.4. Here we provide the normalized RMSE as it facilitates comparison across different predicted variables.

5.1. Prediction performance across macroeconomic indicators

5.1.1. Eurozone GDP

Table 3 lists the RMSE values obtained from the predictive experiments undertaken for the Eurozone GDP. Several news-based models are able to surpass the accuracy of the benchmarks for a variety of the prediction horizons. First, in the case of the first-difference and 1 quarter ahead prediction horizon, both the se-

mantic path model and semantic features prove to be more accurate than both the high-dimensional input and benchmark models. Relative to best high-dimensional input, the RMSE of 42.319 and 65.529 represents 23 % and 44.8 % reduction, respectively.

When looking at the long-term predictions, the LASSO-trained high-dimensional input model outperforms the benchmark for the 3 quarters ahead prediction horizon, in this case achieving a RMSE of 48.415 which is a 5 % reduction relative to the baseline. Furthermore, the high-dimensional input prove to be more accurate over the medium and long-term prediction horizons than the semantic paths models.

5.1.2. German GDP

The prediction errors for the German GDP are as follows. We find the elastic net trained on the high-dimensional input the most accurate for first-difference prediction horizon, achieving a RMSE of 9.921. We observe improvements over the benchmark only for aforementioned horizons. This is especially interesting, since it differs from the prediction results for the whole Eurozone, where we achieved consistent improvements.

5.1.3. Unemployment

In terms of unemployment rate, the news-based models outperform the benchmark over the first-difference prediction. We find

Table 4

Comparison of prediction performance across different monthly macroeconomic indicators. The variable h indicates the number of time steps predicted ahead. Depending on the resolution of the predictive variable, the panels either report long-term forecasts of up to 24 months (such that both cases amount to 2 years). The column ΔY provides the results for the first difference or change in value from one period ahead to the current period. For each outcome and model type, only the performance of the best-in-breed model is reported (i. e. only the best-performing time series model, the best semantic path model, etc.). In each case, the table lists are the root mean squared error. The news-based models that equal or outperform the prediction error of the best benchmark model are in bold for each prediction horizon. In addition, the the P -value from the Diebold–Mariano test is stated in the parenthesis when comparing the best-in-breed model to the best baseline, as well as the corresponding model choice including lag structure and dimension reductions (if any).

Input/model	ΔY	$h = 1$	$h = 3$	$h = 6$	$h = 9$	$h = 12$	$h = 24$
Predicted variable: monthly unemployment							
Benchmark: lags	0.07	0.07	0.17	0.359	0.661	1.105	3.419
	AR1	AR1	AR1	AR1	AR1	AR1	AR1
Semantic path model	0.062 (0.016)	0.079 (0.985)	0.247 (1.000)	0.597 (1.000)	1.026 (1.000)	1.52 (1.000)	2.777 (0.000)
	LASSO6	OLS1	OLS6	OLS6	OLS6	OLS1	OLS
Semantic features	0.062 (0.009)	0.081 (0.996)	0.277 (1.000)	0.668 (1.000)	1.07 (1.000)	1.531 (1.000)	2.468 (0.000)
	LASSO6	LASSO1	LASSO1	LASSO1	LASSO1	LASSO1	PCR1
High-dimensional input	0.07 (0.601)	0.236 (1.000)	0.299 (1.000)	0.562 (1.000)	0.988 (1.000)	1.648 (1.000)	2.3 (0.000)
	GBM	LASSO1	LASSO1	LASSO1	LASSO1	PCA-ENET6	GBM
Sensitivity: lemmatization	0.071 (0.784)	0.236 (1.000)	0.293 (1.000)	0.516 (1.000)	0.817 (1.000)	1.45 (1.000)	2.375 (0.000)
	PCR	LASSO1	LASSO1	LASSO1	LASSO1	PCA-ENET6	ENET
Sensitivity: corpus	0.065 (0.083)	0.236 (1.000)	0.295 (1.000)	0.547 (1.000)	0.886 (1.000)	1.538 (1.000)	2.508 (0.000)
	Complete	German	German	German	German	German	Complete
Predicted variable: monthly inflation rate							
Benchmark: lags	0.29	0.29	0.497	0.656	0.711	0.793	0.749
	AR1	AR1	AR1	AR6	RF6	RF6	AR6
Semantic path model	0.331 (0.999)	0.371 (1.000)	0.738 (1.000)	0.99 (1.000)	1.156 (1.000)	1.227 (1.000)	0.787 (0.856)
	OLS6	LASSO6	LASSO6	Ridge6	LASSO6	OLS6	OLS6
Semantic features	0.323 (0.985)	0.369 (1.000)	0.669 (1.000)	0.956 (1.000)	1.082 (1.000)	1.175 (1.000)	0.752 (0.538)
	PCR	RF6	RF6	RF6	PLS-R6	PCR1	PCR6
High-dimensional input	0.282 (0.175)	0.37 (1.000)	0.622 (1.000)	0.904 (1.000)	0.91 (1.000)	0.902 (0.998)	0.694 (0.093)
	PCA-ENET6	LASSO6	LASSO1	LSA-PCR	LSA-PCR	LSA-PCR6	PCA-GBM1
Sensitivity: lemmatization	0.288 (0.406)	0.382 (1.000)	0.623 (1.000)	0.907 (1.000)	0.909 (1.000)	0.896 (0.998)	0.621 (0.005)
	PCA-RF	LASSO6	LASSO1	LSA-PCR	LSA-PCR6	LSA-PCR1	PCA-GBM6
Sensitivity: corpus	0.287 (0.368)	0.379 (1.000)	0.768 (1.000)	0.861 (1.000)	0.892 (1.000)	0.866 (0.954)	0.631 (0.021)
	Complete	German	German	German	German	German	German
Predicted variable: monthly industrial production rate							
Benchmark: lags	1.827	1.827	3.505	5.703	6.532	6.922	9.816
	AR1	AR1	AR6	AR6	AR6	AR1	AR6
Semantic path model	1.905 (0.947)	1.851 (0.739)	3.497 (0.439)	5.44 (0.024)	6.12 (0.169)	7.063 (0.663)	12.631 (1.000)
	Ridge1	LASSO1	LASSO6	OLS6	ENET6	Ridge6	Ridge
Semantic features	1.884 (0.879)	1.851 (0.739)	3.497 (0.439)	5.468 (0.024)	6.12 (0.169)	6.896 (0.464)	10.987 (0.999)
	PCR	LASSO1	LASSO6	LASSO6	ENET6	PLS-R6	PCR1
High-dimensional input	1.812 (0.128)	2.045 (0.984)	3.521 (0.546)	5.392 (0.268)	5.521 (0.020)	5.064 (0.000)	10.37 (0.946)
	LSA-LASSO6	LASSO1	LASSO1	PCA-LASSO6	PCA-GBM	PLS-R1	Ridge1
Sensitivity: lemmatization	1.817 (0.329)	2.035 (0.983)	3.522 (0.548)	5.364 (0.234)	5.403 (0.006)	4.978 (0.000)	10.415 (0.978)
	LSA-GBM	LASSO1	LASSO1	PCA-LASSO6	ENET1	PLS-R1	PCR6
Sensitivity: corpus	1.814 (0.107)	1.958 (0.928)	3.435 (0.323)	5.336 (0.125)	5.238 (0.003)	4.254 (0.000)	5.557 (0.000)
	German	German	Complete	Complete	Complete	German	German
Predicted variable: monthly business confidence index							
Benchmark: lags	0.057	0.057	0.318	0.791	0.96	0.955	1.023
	AR6	AR6	AR6	AR6	AR6	AR6	AR1
Semantic path model	0.061 (0.943)	0.057 (0.543)	0.334 (0.85)	0.799 (0.584)	0.878 (0.020)	0.847 (0.006)	1.146 (1.000)
	OLS6	LASSO6	OLS6	OLS6	LASSO6	Ridge6	OLS6
Semantic features	0.063 (0.989)	0.057 (0.543)	0.334 (0.858)	0.812 (0.71)	0.878 (0.020)	0.847 (0.006)	1.515 (1.000)
	LASSO6	LASSO6	LASSO6	LASSO6	LASSO6	Ridge6	USM1

(continued on next page)

Table 4 (continued)

Input/model	ΔY	$h = 1$	$h = 3$	$h = 6$	$h = 9$	$h = 12$	$h = 24$
High-dimensional input	0.218 (1.000) LASSO6	0.28 (1.000) ENET1	0.784 (1.000) ENET1	1.291 (1.000) LASSO6	1.199 (0.986) PCA-RF6	0.976 (0.586) PCA-RF6	1.21 (0.987) PCA-PCR
Sensitivity: lemmatization	0.235 (1.000) GBM6	0.28 (1.000) ENET1	0.784 (1.000) ENET1	1.323 (1.000) LASSO6	1.286 (0.997) PCA-RF6	1.033 (0.811) PCA-GBM6	1.213 (0.988) PCA-PCR
Sensitivity: corpus	0.232 (1.000) German	0.281 (1.000) German	0.795 (1.000) Complete	1.272 (1.000) Complete	1.115 (0.971) Complete	0.996 (0.681) German	1.09 (0.81) Complete
Predicted variable: monthly consumer confidence index							
Benchmark: lags	0.085 AR6	0.085 AR6	0.454 AR6	0.998 AR6	1.212 AR6	1.206 AR6	1.396 AR1
Semantic path model	0.093 (0.996) OLS6	0.089 (0.824) LASSO6	0.516 (0.978) LASSO6	1.136 (0.984) LASSO6	1.263 (0.796) OLS6	1.074 (0.004) OLS6	1.841 (1.000) OLS6
Semantic features	0.094 (0.999) LASSO6	0.089 (0.824) LASSO6	0.516 (0.978) LASSO6	1.136 (0.984) LASSO6	1.226 (0.566) USM6	1.061 (0.002) USM6	2.003 (1.000) RF6
High-dimensional input	0.245 (1.000) RF6	0.293 (1.000) LASSO1	0.79 (1.000) LASSO1	1.217 (0.996) RF1	1.294 (0.946) PCA-Ridge1	1.229 (0.612) PCA-GBM1	1.375 (0.385) LSA-LASSO6
Sensitivity: lemmatization	0.246 (1.000) ENET6	0.293 (1.000) LASSO1	0.79 (1.000) LASSO1	1.235 (0.998) LASSO1	1.295 (0.916) PCA-ENET1	1.25 (0.767) PCA-LASSO1	1.389 (0.465) LSA-PCR6
Sensitivity: corpus	0.245 (1.000) Complete	0.294 (1.000) German	0.736 (1.000) German	1.188 (0.987) German	1.141 (0.054) Complete	1.25 (0.735) Complete	1.371 (0.359) Complete
Predicted variable: federal 10-year bond							
Benchmark: lags	0.23 RF6	0.266 AR1	0.62 AR6	1.076 AR1	1.399 AR1	1.723 AR6	2.207 AR6
Semantic path model	0.226 (0.303) LASSO1	0.444 (1.000) LASSO1	1.042 (1.000) OLS1	1.866 (1.000) OLS1	2.092 (1.000) OLS1	2.427 (1.000) Ridge1	2.484 (1.000) OLS
Semantic features	0.226 (0.303) LASSO1	0.392 (1.000) PLS-R6	1.048 (1.000) LASSO1	1.517 (1.000) PCR6	2.068 (1.000) PLS-R1	2.197 (1.000) PCR1	2.43 (1.000) PLS-R
High-dimensional input	0.209 (0.004) LSA-LASSO	0.434 (1.000) LASSO1	1.184 (1.000) LASSO1	1.942 (1.000) LASSO1	2.147 (1.000) LASSO1	2.144 (1.000) ENET1	2.121 (0.001) PCA-ENET
Sensitivity: lemmatization	0.212 (0.008) LSA-Ridge	0.441 (1.000) LASSO1	1.158 (1.000) LASSO1	2.035 (1.000) LASSO1	2.262 (1.000) PLS-R6	2.046 (1.000) ENET1	2.117 (0.000) PCA-ENET
Sensitivity: corpus	0.211 (0.004) Complete	0.422 (1.000) German	1.07 (1.000) German	1.771 (1.000) German	1.886 (1.000) German	1.868 (1.000) German	1.788 (0.000) Complete

both the semantic path model, as well as the semantic features, to be the most accurate with a RMSE of 0.062. Furthermore, the DM test results are highly significant for both models derived from semantic inputs.

We observe that news-based models cannot beat the benchmarks for medium-term predictions, while we again gain improvements from using news in the long-term. Here the news-based models consistently outperform the baseline for the 24 months ahead prediction horizon. The highest predictive accuracy stems from gradient boosting applied to high-dimensional input with a prediction error of 2.300. The best of the benchmarks reported only an RMSE of 3.419; this results in a forecast error reduction of 32.7%. Furthermore, the Diebold–Mariano test results are highly statistically significant for the 24 months prediction horizon for all news-based models. The semantic path model attains only an improvement of 18.8%, yet while being fully interpretable.

5.1.4. Inflation

We find support for the fact that the news-based models are superior to the benchmarks for the first difference prediction horizons. The high-dimensional input accounts for the lowest RMSE (0.209). Furthermore, the Diebold–Mariano test is statistically significant. For the prediction horizons 1–12 months ahead, our news-

based models attain similar levels of accuracy however are unable to best the leading benchmark. Further improvements in predictive performance are evident for the 24 months ahead prediction horizon. Here the high-dimensional input with PCA dimensionality reduction yields the lowest RMSE of 0.694, i. e. an improvement of 7.3% over the time series baseline. The observed results indicate that the financial news stemming from listed corporations can be utilized to capture the rate of change in consumer prices in the short-term, and that these, continue to be predictive over the long-term.

5.1.5. Industrial production

The news-based models outperform the benchmarks for a majority of tested prediction horizons. In particular, the news-based model achieve a lower forecast error for the mid-range prediction horizons. The experiments for the first-difference prediction reveal that the high-dimensional input with LSA preprocessing obtains a RMSE of 1.812 (i. e. an improvement of 0.8%. The news-based models outperform the benchmark for the 3–12 months ahead prediction horizons. Both the semantic path model and the semantic features attain the best RMSE (3.497), for the 3 months ahead prediction. The high-dimensional input is the most accurate for the 6–12 months ahead prediction horizon. The *p*-values from the individ-

ual DM tests indicate that predicted values from our news-based system are significantly more accurate over the before-mentioned prediction horizons.

5.1.6. Business confidence

The semantic path model, as well as semantic features, are able to outperform the benchmark and the high-dimensional input across a number of prediction horizons. Both methods register a the lowest RMSE for the 1 month (0.057), 9 month (0.878) and 12 month (0.847) ahead prediction horizons. The latter amounts to an improvement of 11.3%. Furthermore, these model significantly outperform the high-dimensional input across all predictions horizons. The business confidence index is calculated by conducting surveys of business leaders who know the current state of the corporate sector. The semantic path model can adapt at capturing the sentiment of the individual announcements; i. e. the same sentiment would be reflected in the index. The lack of long-term predictive power for the high-dimensional inputs could be due to the fact that the news content might already be factored into their survey answers, therefore reducing potential lag affects. Overall, the semantic models seems better suited for predictions of the business confidence index than the high-dimensional inputs.

5.1.7. Consumer confidence

We investigate the prognostic capability of financial news for the consumer confidence index. Overall, the predictive performance of the news-based models does not diverge greatly from those of the benchmarks. Of the news-based models, both semantic models prove to be more accurate than the high-dimensional input. The news-based models do narrowly outperform the benchmark for prediction horizons greater than 9 months ahead.

5.1.8. Government bonds

The news-based models perform best relative to the benchmarks for the first-difference prediction horizons. For instance, the high-dimensional input attains the lowest overall RMSE (0.209, i. e. a minus of 9.1% over the best baseline), closely followed by both the semantic models (0.226, i. e. a minus of 1.7%), while the RMSE of the best benchmark is 0.230. The high-dimensional input is also able to outperform best the benchmark for the 24 months ahead prediction horizon. The Diebold–Mariano test results for the high-dimensional input are all highly statistically significant.

5.2. Interpretation of latent semantic structures

Fig. 2 presents the results from the semantic path model and decomposes the forecasts into different constructs in order to facilitate interpretation. That is, we compute the value of each construct – given by predicted value $\psi_i z_i$ – and compare these figures across the different semantic categories according to the Loughran–McDonald finance-specific dictionary. These categories are: (i) positivity/negativity of economic outlook, (ii) uncertainty regarding economic development, (iii) legal risk for operations and (iv) confidence/modality of the disclosing company.

Altogether, we find the positive construct has the strongest impact on both Eurozone and German GDP and, to a lesser extent, on consumer confidence, industrial production and government bonds. The negative construct, on the other hand, affects the unemployment rate, inflation, industrial production, consumer confidence and government bonds. A rise in negative news could see firms downsize operations and cut investment spending, resulting in lower levels of employment. Reduced confidence in the economic outlook could see consumers also cut spending placing downward pressure on prices. The lag component is strongly predictive for all variables except business confidence. Potentially, a

change in perceptions at a corporate level would already be factored into the index by business leaders. The remaining indicators would be slower to adjust. We find the uncertainty construct to be predictive for inflation, the confidence indices and government bonds. Consumers, firms and financial markets all have an aversion to uncertainty. The legal construct seems to be predictive for industrial production, business and consumer confidence. A possible explanation is that firms facing legal difficulties would reduce output and other business investment impacting both the suppliers and consumers that function within the economic stream. Finally, we find the modal constructs to be a weak predictor across all of the chosen indicators.

5.3. Sensitivity analysis

Tables 3 and 4 entail additional sensitivity checks. First, we compare stemming as a default in contrast to lemmatization. While the latter adapts correctly to the semantic context, it suffers from larger predictor matrices and potentially higher risk of overfitting. However, our results suggest no clear indication concerning which approach is preferable with regard to our task. Second, we performed all analyses based on the English corpus and now compare this choice to both a corpus with German articles and a combined one. Again, we find no clear evidence that one consistently outperforms the other. A potential reason for this parity stems from the fact that the majority of firms publish in both languages and any potential differences therefore are relatively minor.

5.4. Comparison

While the previous results compared the use of time series models to a machine learning approach, we now turn our attention to the overall capacity of foreseen macroeconomic development. For this reason, we computed the normalized RMSE for all macroeconomic indicators as it facilitates comparisons between different outcome variables and their respective scales. Formally, it is defined by
$$\text{NRMSE} = \frac{\text{RMSE}}{Y_{\max} - Y_{\min}}$$

The normalized RMSE results are detailed in Table 5. We find the Eurozone GDP results to have lower prediction errors than German GDP. Potential reasons could be the export orientation of German companies, or the higher volatility of the Eurozone GDP, which can be better explained by production-related information from news disclosures. The prediction accuracy for the business confidence index relative to consumer confidence is also better. Given the direct link between the business leaders who are surveyed to estimate the index and the source of the ad hoc announcements, a higher rate of accuracy for the business confidence index could be expected.

In comparing the performance of both the semantic models relative to the high-dimensional input, we find the results varying depending on the indicator and prediction horizon. The semantic models consistently attain lower prediction errors for the confidence indices across all the prediction horizons except for the very long run. In fact, the semantic path model accomplishes the best forecast in 6 prediction experiments, while the semantic features do so in 8 cases and the high-dimensional in 12, yet the latter with hardly any interpretability. This is interesting in the light that our semantic path model enforces a linear structure among constructs and thus excludes potentially non-linear relationships as used by predictive models with semantic features or the universal structure model. While non-linearity introduces additional flexibility, our setting is extremely overspecified and the smaller parameter space of linear models can often be beneficial. Moreover, we see that traditional approaches to unsupervised dimensionality reduc-

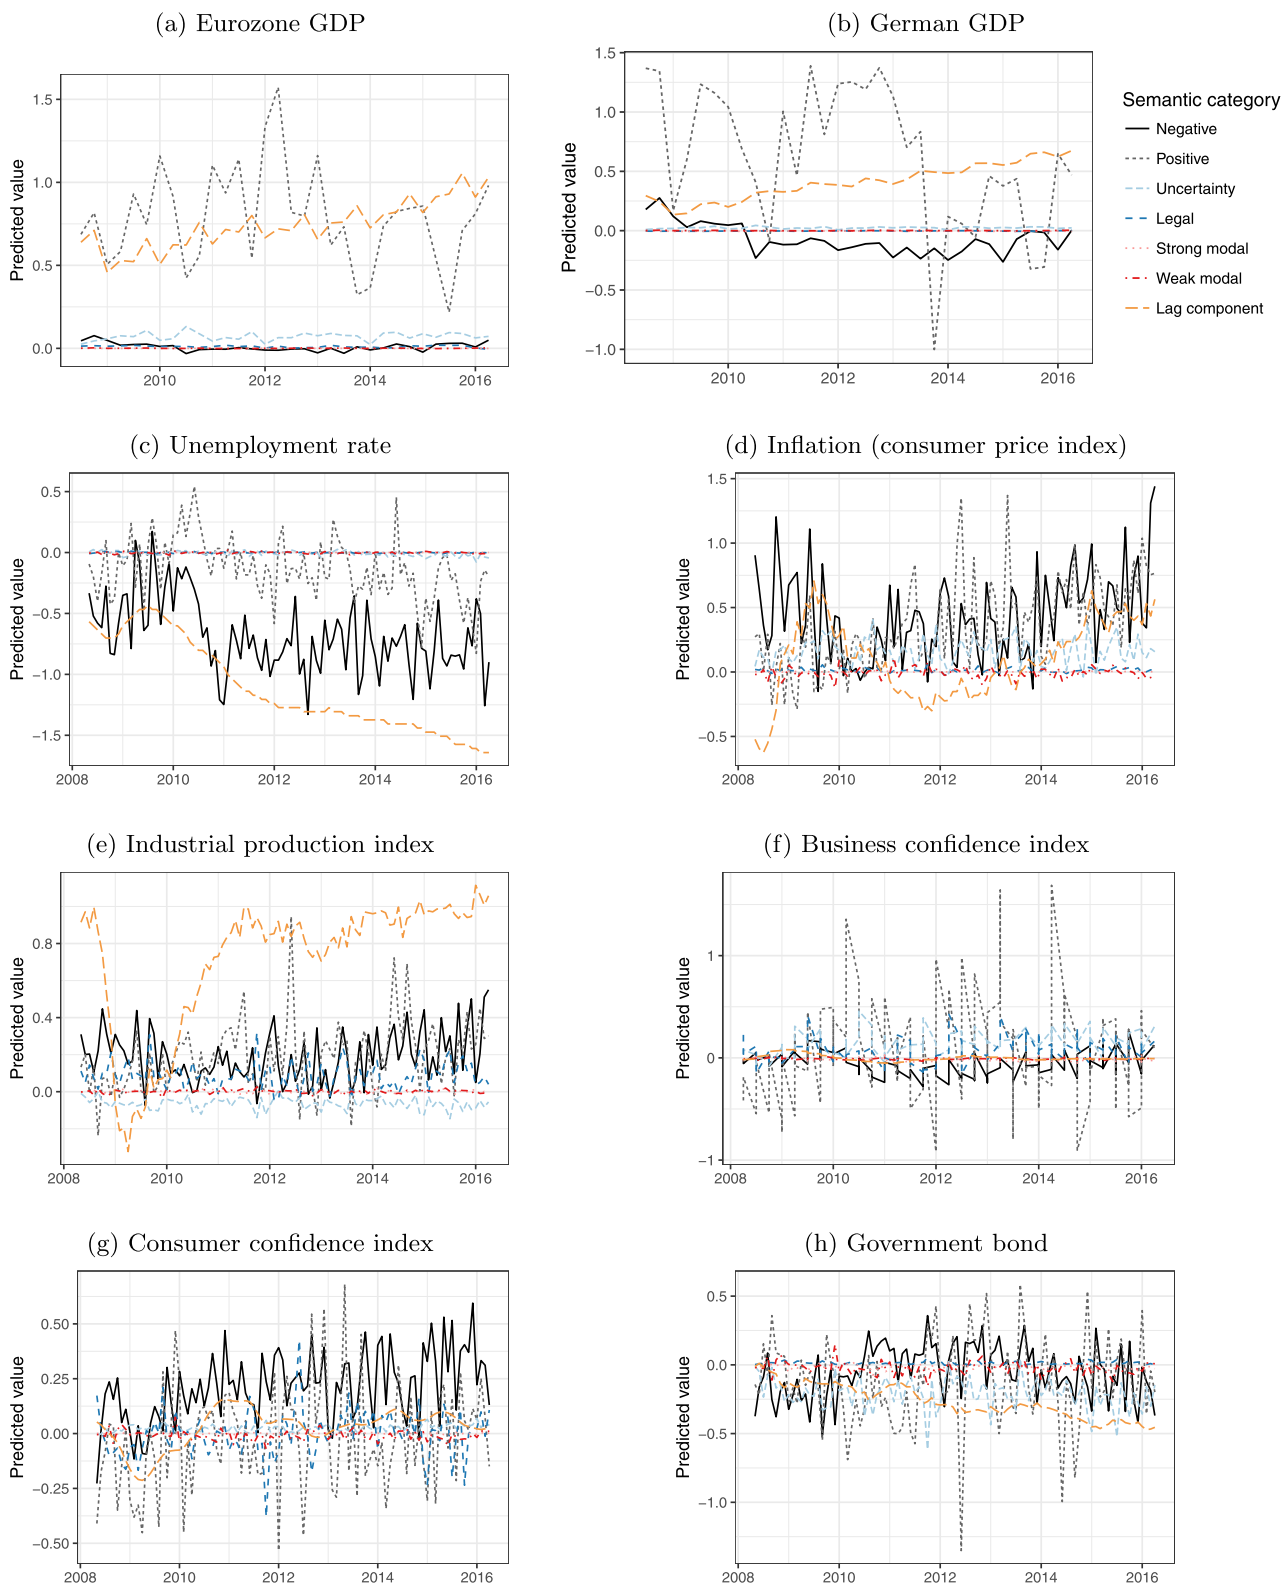


Fig. 2. Illustrated interpretability of semantic path models. The value of each construct – given by predicted value $\psi_i z_i$ – is plotted over the course of the test set. Note that the inner model from our semantic path model operates on normalized constructs z_i , which is even beneficial for the above comparisons. The latent structures directly lend to interpretation with respect to our semantic dimension. For reasons of comparability, we apply the same estimation technique to all macroeconomic indicators: the inner model is estimated via ordinary least squares for a prediction 1-year ahead based on the English corpus with 1 lag.

Table 5

Summary of text-based prediction performance across different macroeconomic indicators. Here the normalized RMSE is reported in order to facilitate comparison across outcomes (and their respective scales). The variable h indicates the number of time steps predicted ahead. Depending on the resolution of the predictive variable, the panels either report long-term forecasts of up to 8 quarters or 24 months (such that both cases amount to 2 years). The column ΔY provides the results for the first difference or change in value from one period ahead to the current period. For each outcome and model type, only the performance of the best-in-breed model is reported (i. e. only the best-performing time series model, the best semantic path model, etc.).

Predicted variable	Model	Normalized RMSE						
		Quarterly resolution						
ΔY		$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 8$		
Eurozone GDP	Semantic path model	12.0	8.0	12.5	16.2	23.9	37.9	
	Semantic features	11.0	8.0	12.5	15.7	17.4	20.6	
	High-dimensional input	14.3	14.5	11.1	11.7	15.5	25.2	
German GDP	Semantic path model	21.1	18.5	19.0	32.0	57.3	77.1	
	Semantic features	19.6	15.4	14.4	17.8	28.8	43.5	
	High-dimensional input	15.6	12.8	14.6	24.7	37.4	53.4	
		Monthly resolution						
ΔY		$h = 1$	$h = 3$	$h = 6$	$h = 9$	$h = 12$	$h = 24$	
Unemployment	Semantic path model	15.4	2.2	6.7	16.1	27.7	41.1	99.2
	Semantic features	15.4	2.2	7.5	18.1	28.9	41.4	88.1
	High-dimensional input	17.6	6.5	8.1	15.2	26.7	44.5	82.1
Inflation	Semantic path model	18.9	9.7	20.4	34.1	39.9	42.3	29.3
	Semantic features	18.5	9.6	18.5	33	37.3	40.5	28.0
	High-dimensional input	16.1	9.7	17.2	31.2	31.4	31.1	25.9
Industrial production	Semantic path model	16.6	6.6	12.4	19.4	21.8	28.7	92.9
	Semantic features	16.4	6.6	12.4	19.5	21.8	28.0	80.8
	High-dimensional input	15.8	7.3	12.5	19.2	19.6	20.6	76.2
Business confidence	Semantic path model	4.6	0.9	5.1	12.3	13.5	13.7	32.6
	Semantic features	4.7	0.9	5.2	12.5	13.5	13.7	43.1
	High-dimensional input	16.4	4.3	12.1	19.9	18.5	15.7	34.4
Consumer confidence	Semantic path model	6.6	1.4	8.1	17.8	19.8	17.3	61.7
	Semantic features	6.6	1.4	8.1	17.8	19.2	17.0	67.1
	High-dimensional input	17.3	4.6	12.4	19.0	20.2	19.8	46.1
Government bonds	Semantic path model	21.5	9.8	24.2	50.2	56.2	65.3	71.3
	Semantic features	21.5	8.7	24.3	40.8	55.6	59.1	69.7
	High-dimensional input	19.8	9.6	27.5	52.2	57.7	57.7	60.9

tion rarely yield the best-of-breed model; conversely, our domain knowledge can outperform them frequently.

The overall performance of the news-based models varies depending on the macroeconomic variable and prediction horizon. Over the short- and medium term, the benchmark models tend to be superior in terms of consumer-centered variables, such as unemployment rate or inflation. However, financial news reflects both the current health and future expectations of corporations. As a consequence, news-based models appear especially useful for product-related indicators, such as the actual industrial product rate and confidence indices. Finally, in terms of long-term predictions, the news-based models tend to outperform the benchmark models. Altogether, this provides ample evidence of the strength of the predictive properties of our news-based inputs.

6. Managerial implications

The economic climate has long-ranging effects on firm operations. As a consequence, managers and policymakers alike are required to form expectations, hedge potential operational risks from economic downturns and adapt their operations accordingly. Our method thus assists managers in making such predictions and provides a cost-effective alternative to expert forecasts. Firm managers can utilize our predictions in order to, among other things, set long-term strategic plans, decide on investment levels, and manage human resources. Beyond that, managers in public sector institutions may benefit from our long-term predictions when setting fiscal policy, such as government spending, borrowing and taxation, and monetary policy concerning to control of the supply of money within the economy. Our innovative, news-based methodology also provides investment managers a greater ability to antici-

pate the impact of short-term economic fluctuations on their portfolios and, if necessary, take the necessary precautions to hedge against losses.

Professional macroeconomic forecasts, such as those from central banks, usually stem from both quantitative predictions and the judgment of economic experts (Matsypura et al., 2018). For example, the German Bundesbank combines factor models, demand- and supply-side bridge equations and time series models for industrial production. Expert knowledge is then applied to these quantitative forecasts to determine the final projections for key macroeconomic indicators.² The European Central Bank has a similar process and their projections “may incorporate a fair amount of judgment” (European Central Bank, 2016). Their questionnaire concludes that “respondents consider their predictions to be 40% judgment-based”.³ Hence, it is likely that these forecasts

Numerous studies have analyzed professional forecasts for their predictive performance and ability to identify potential biases (Blanc & Setzer, 2015; Mostard, Teunter, & de Koster, 2011). In the field of macroeconomic forecasting, a recent study (Dovern & Weisser, 2011) finds that the distribution of a forecasts accuracy varies significantly across indicators, forecasters and nations. A further study reveals that consensus estimates are poor predictors of gross domestic product as compared with statistical models (Jansen et al., 2016). However, the consensus estimates tend to per-

² German Bundesbank. *Forecasting models in short-term business cycle analysis: A workshop report*. URL: https://www.bundesbank.de/Redaktion/EN/Downloads/Publications/Monthly_Report_Articles/2013/2013_09_forecasting.pdf, last accessed on January 12, 2018.

³ European Central Bank. *Results of a special questionnaire for participants in the ECB Survey of Professional Forecasters*. URL: https://www.ecb.europa.eu/stats/prices/indic/forecast/shared/files/quest_summary.pdf, last accessed on January 12, 2018.

form better in periods of crisis. An intriguing approach would be to combine expert estimations and statistical forecasts (Lessmann, Sung, Johnson, & Ma, 2012).

Predicting macroeconomic variables is difficult, requiring time-intensive collection of economic data, which, as a result, is often out of date. However, our news source provides free and real-time access. Additionally, time series models are excellent at picking up trends, but perform poorly in situations of volatility (Jansen et al., 2016). In contrast, our approach can incorporate additional predictors besides trend, as it extracts qualitative information, as well as overall sentiments, regarding the economic outlook from financial news. At the same time, our semantic path model accounts for interpretable results and can thus be further assessed by decision-makers.

Our innovative, news-based methodology adds to previous literature and entails a number of important implications for researchers, as well as for the public and private sectors, by providing long-term forecasts of the macroeconomic climate. Here big data is once again shifting the frontiers of research (Mortenson et al., 2015) and is likely to change the way in which operations adapt to economic developments. As demonstrated by our previously discussed evaluations, the system elicits a high predictive accuracy, when compared with a number of benchmark models. This contributes to a growing body of literature that replaces man-made forecasts with computerized and fine-tuned predictive models. Our proposed methodology could be extremely valuable for small enterprises or countries in which macroeconomic forecasts are scarce due to the associated costs.

7. Conclusion

The volatile nature of the economy affects the operation of firms in key dimensions, such as prices for goods and services, financial risk, utilization of supply chains and customer demand. Hence, it is important for managers to obtain an accurate prognosis of developments in the economy in order to direct operations accordingly. Hitherto, macroeconomic forecasts have been dominated by expert opinions and rather simple time series models, while the big data era has created new opportunities to enhance the predictive power of macroeconomic forecasts based on unstructured data sources.

This work contributes a news-based methodology for predicting macroeconomic indicators. First, we experiment with conventional models from machine learning in order to predict macroeconomic outcomes on the basis of word occurrences and historic lags. Back-testing shows that this method outperforms the benchmark models in an out-of-sample evaluation for various prediction horizons and macroeconomic indicators. Yet text-based predictions represent a challenging undertaking due to potential overfitting for a high-dimensional input. Instead, we suggest an alternative approach whereby the words from different semantic categories are mapped onto latent structures as a form of feature engineering. This shrinks the feature space considerably, thereby achieving superior out-of-sample performance in multiple experiments. Beyond that, we propose semantically-structured variant of partial least squares that reaches a comparable accuracy, but while fulfilling the demand of practitioners in being fully interpretable. This contributes to a greater understanding of how qualitative information can be used to make long-term predictions for key macroeconomic indicators.

This paper opens avenues for future research in a number of directions. First, the proposed semantic path model can facilitate interpretable predictions in a variety of settings involving natural language. Second, the approach can be seamlessly extended to include alternative data sources, including additional news sources and time series. Third, one could weigh corporate disclosures ac-

ording to their market capitalization, since larger firms presumably have a great impact on the general economy than smaller firms.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at [10.1016/j.ejor.2018.05.068](https://doi.org/10.1016/j.ejor.2018.05.068).

References

- Aguirre-Urreta, M. I., & Marakas, G. M. (2014). Research note—partial least squares and models with formatively specified endogenous constructs: A cautionary note. *Information Systems Research*, 25(4), 761–778. doi:10.1287/isre.2013.0493.
- Akkoç, S. (2012). An empirical comparison of conventional techniques, neural networks and the three stage hybrid adaptive neuro fuzzy inference system (ANFIS) model for credit scoring analysis: The case of Turkish credit card data. *European Journal of Operational Research*, 222(1), 168–178. doi:10.1016/j.ejor.2012.04.009.
- Allen, P. G., & Morzuch, B. J. (2006). Twenty-five years of progress, problems, and conflicting evidence in econometric forecasting: What about the next 25 years? *International Journal of Forecasting*, 22(3), 475–492. doi:10.1016/j.ijforecast.2006.03.003.
- Aubert, B. A., Rivard, S., & Patry, M. (1996). Development of measures to assess dimensions of is operation transactions. *Omega*, 24(6), 661–680. doi:10.1016/S0305-0483(96)00019-9.
- Blanc, S. M., & Setzer, T. (2015). Analytical debiasing of corporate cash flow forecasts. *European Journal of Operational Research*, 243(3), 1004–1015. doi:10.1016/j.ejor.2014.12.035.
- Bovi, M., & Cerqueti, R. (2016). Forecasting macroeconomic fundamentals in economic crises. *Annals of Operations Research*, 247(2), 451–469. doi:10.1007/s10479-015-1879-4.
- Buckler, F., & Hennig-Thurau, T. (2008). Identifying hidden structures in marketing's structural models through universal structure modeling. *Marketing Journal of Research and Management*, 4(2), 49–68.
- Calabrese, R., Degl'Innocenti, M., & Osmetti, S. A. (2017). The effectiveness of tarp-cpp on the us banking industry: A new copula-based approach. *European Journal of Operational Research*, 256(3), 1029–1037. doi:10.1016/j.ejor.2016.07.046.
- Carriero, A., Clark, T. E., & Marcellino, M. (2015). Bayesian VARs: Specification choices and forecast accuracy. *Journal of Applied Econometrics*, 30(1), 46–73. doi:10.1002/jae.2315.
- Chun, H., & Keleş, S. (2010). Sparse partial least squares regression for simultaneous dimension reduction and variable selection. *Journal of the Royal Statistical Society (Series B)*, 72(1), 3–25. doi:10.1111/j.1467-9868.2009.00723.x.
- Demyanyk, Y., & Hasan, I. (2010). Financial crises and bank failures: A review of prediction methods. *Omega*, 38(5), 315–324. doi:10.1016/j.omega.2009.09.007.
- Desai, V. S., Crook, J. N., & Overstreet, G. A. (1996). A comparison of neural networks and linear scoring models in the credit union environment. *European Journal of Operational Research*, 95(1), 24–37. doi:10.1016/0377-2217(95)00246-4.
- Dovern, J., & Weisser, J. (2011). Accuracy, unbiasedness and efficiency of professional macroeconomic forecasts: An empirical comparison for the G7. *International Journal of Forecasting*, 27(2), 452–465. doi:10.1016/j.ijforecast.2010.05.016.
- Du Jardin, P. (2015). Bankruptcy prediction using terminal failure processes. *European Journal of Operational Research*, 242(1), 286–303. doi:10.1016/j.ejor.2014.09.059.
- European Central Bank (2016). *A guide to the Eurosystem/ECB staff macroeconomic projection exercises*. Frankfurt am Main: European Central Bank.
- Fethi, M. D., & Pasiouras, F. (2010). Assessing bank efficiency and performance with operational research and artificial intelligence techniques: A survey. *European Journal of Operational Research*, 204(2), 189–198. doi:10.1016/j.ejor.2009.08.003.
- Feuerriegel, S., & Prendinger, H. (2016). News-based trading strategies. *Decision Support Systems*, 90, 65–74. doi:10.1016/j.dss.2016.06.020.
- Geng, R., Bose, I., & Chen, X. (2015). Prediction of financial distress: An empirical study of listed Chinese companies using data mining. *European Journal of Operational Research*, 241(1), 236–247. doi:10.1016/j.ejor.2014.08.016.
- de Gooijer, J. G., & Hyndman, R. J. (2006). 25 years of time series forecasting. *International Journal of Forecasting*, 22(3), 443–473. doi:10.1016/j.ijforecast.2006.01.001.
- Goudie, A. W., & Meeks, G. (1982). The effects of macroeconomic developments on individual companies' flows of funds. *Omega*, 10(4), 361–371. doi:10.1016/0305-0483(82)90015-9.
- Gutiérrez, E., & Lozano, S. (2012). A competing risks analysis of the duration of federal target funds rates. *Computers & Operations Research*, 39(4), 785–791. doi:10.1016/j.cor.2010.10.002.
- Hastie, T. J., Tibshirani, R. J., & Friedman, J. H. (2013). *The elements of statistical learning: Data mining, inference, and prediction*. Springer Series in Statistics (2nd). New York, NY: Springer.
- Huang, W., Nakamori, Y., & Wang, S.-Y. (2005). Forecasting stock market movement direction with support vector machine. *Computers & Operations Research*, 32(10), 2513–2522. doi:10.1016/j.cor.2004.03.016.
- Huang, Y., Kou, G., & Peng, Y. (2017). Nonlinear manifold learning for early warnings in financial markets. *European Journal of Operational Research*, 258(2), 692–702. doi:10.1016/j.ejor.2016.08.058.

- Jansen, W. J., Jin, X., & de Winter, J. M. (2016). Forecasting and nowcasting real gdp: Comparing statistical models and subjective forecasts. *International Journal of Forecasting*, 32(2), 411–436. doi:10.1016/j.ijforecast.2015.05.008.
- Kraus, M., & Feuerriegel, S. (2017). Decision support from financial disclosures with deep neural networks and transfer learning. *Decision Support Systems*, 104, 38–48. doi:10.1016/j.dss.2017.10.001.
- Kuhn, M. (2008). Building predictive models in R using the caret package. *Journal of Statistical Software*, 28(5), 1–26. doi:10.18637/jss.v028.i05.
- Kung, L.-M., & Yu, S.-W. (2008). Prediction of index futures returns and the analysis of financial spillovers: A comparison between GARCH and the grey theorem. *European Journal of Operational Research*, 186(3), 1184–1200. doi:10.1016/j.ejor.2007.02.046.
- Layton, A. P., & Smith, D. R. (2007). Business cycle dynamics with duration dependence and leading indicators. *Journal of Macroeconomics*, 29(4), 855–875. doi:10.1016/j.jmacro.2006.02.003.
- Lê Cao, K.-A., Rossouw, D., Robert-Granié, C., & Besse, P. (2008). A sparse PLS for variable selection when integrating omics data. *Statistical Applications in Genetics and Molecular Biology*, 7(1), 35. doi:10.2202/1544-6115.1390.
- Lessmann, S., Sung, M.-C., Johnson, J. E., & Ma, T. (2012). A new methodology for generating and combining statistical forecasting models to enhance competitive event prediction. *European Journal of Operational Research*, 218(1), 163–174. doi:10.1016/j.ejor.2011.10.032.
- Liang, D., Lu, C.-C., Tsai, C.-F., & Shih, G.-A. (2016). Financial ratios and corporate governance indicators in bankruptcy prediction: A comprehensive study. *European Journal of Operational Research*, 252(2), 561–572. doi:10.1016/j.ejor.2016.01.012.
- Litterman, R. B. (1986). Forecasting with Bayesian vector autoregressions: Five years of experience. *Journal of Business & Economic Statistics*, 4(1), 25–38. doi:10.1080/07350015.1986.10509491.
- Mahmoud, E., Motwani, J., & Rice, G. (1990). Forecasting us exports: An illustration using time series and econometric models. *Omega*, 18(4), 375–382. doi:10.1016/0305-0483(90)90027-7.
- Matsypura, D., Thompson, R., & Vasnev, A. L. (2018). Optimal selection of expert forecasts with integer programming. *Omega*, 78, 165–175. doi:10.1016/j.omega.2017.06.010.
- McIntosh, C. N., Edwards, J. R., & Antonakis, J. (2014). Reflections on partial least squares path modeling. *Organizational Research Methods*, 17(2), 210–251. doi:10.1177/1094428114529165.
- McKee, T. E., & Lensberg, T. (2002). Genetic programming and rough sets: A hybrid approach to bankruptcy classification. *European Journal of Operational Research*, 138(2), 436–451. doi:10.1016/S0377-2217(01)00130-8.
- Mortenson, M. J., Doherty, N. F., & Robinson, S. (2015). Operational research from Taylorism to Terabytes: A research agenda for the analytics age. *European Journal of Operational Research*, 241(3), 583–595. doi:10.1016/j.ejor.2014.08.029.
- Mostard, J., Teunter, R., & de Koster, R. (2011). Forecasting demand for single-period products: A case study in the apparel industry. *European Journal of Operational Research*, 211(1), 139–147. doi:10.1016/j.ejor.2010.11.001.
- Nassirtoussi, A. K., Aghabozorgi, S., Wah, T. Y., & Ngo, D. C. L. (2014). Text mining for market prediction: A systematic review. *Expert Systems with Applications*, 41(16), 7653–7670. doi:10.1016/j.eswa.2014.06.009.
- Oztekin, A., Kizilaslan, R., Freund, S., & Iseri, A. (2016). A data analytic approach to forecasting daily stock returns in an emerging market. *European Journal of Operational Research*, 253(3), 697–710. doi:10.1016/j.ejor.2016.02.056.
- Oztekin, A., Kong, Z. J., & Delen, D. (2011). Development of a structural equation modeling-based decision tree methodology for the analysis of lung transplantations. *Decision Support Systems*, 51(1), 155–166. doi:10.1016/j.dss.2010.12.004.
- Rigdon, E. E., Ringle, C. M., & Sarstedt, M. (2010). Structural modeling of heterogeneous data with partial least squares. *Review of Marketing Research*, 7, 255–296.
- Sermpinis, G., Theofilatos, K., Karathanasopoulos, A., Georgopoulos, E. F., & Dunis, C. (2013). Forecasting foreign exchange rates with adaptive neural networks using radial-basis functions and particle swarm optimization. *European Journal of Operational Research*, 225(3), 528–540. doi:10.1016/j.ejor.2012.10.020.
- Sevim, C., Oztekin, A., Bali, O., Gumus, S., & Guresen, E. (2014). Developing an early warning system to predict currency crises. *European Journal of Operational Research*, 237(3), 1095–1104. doi:10.1016/j.ejor.2014.02.047.
- Shaik, S. (2015). Impact of liquidity risk on variations in efficiency and productivity: A panel gamma simulated maximum likelihood estimation. *European Journal of Operational Research*, 245(2), 463–469. doi:10.1016/j.ejor.2015.03.018.
- Stock, J. H., & Watson, M. W. (2002). Macroeconomic forecasting using diffusion indexes. *Journal of Business & Economic Statistics*, 20(2), 147–162. doi:10.1198/073500102317351921.
- Sun, L., & Shenoy, P. P. (2007). Using Bayesian networks for bankruptcy prediction: Some methodological issues. *European Journal of Operational Research*, 180(2), 738–753. doi:10.1016/j.ejor.2006.04.019.
- Tam, K. Y., & Kiang, M. Y. (1992). Managerial applications of neural networks: The case of bank failure predictions. *Management Science*, 38(7), 926–947. doi:10.1287/mnsc.38.7.926.
- Tay, F. E., & Cao, L. (2001). Application of support vector machines in financial time series forecasting. *Omega*, 29(4), 309–317. doi:10.1016/S0305-0483(01)00026-3.
- Tenenhaus, M., Vinzi, V. E., Chatelin, Y.-M., & Lauro, C. (2005). PLS path modeling. *Computational Statistics & Data Analysis*, 48(1), 159–205. doi:10.1016/j.csda.2004.03.005.
- Tsai, M.-F., & Wang, C.-J. (2017). On the risk prediction and analysis of soft information in finance reports. *European Journal of Operational Research*, 257(1), 243–250. doi:10.1016/j.ejor.2016.06.069.
- Turkyilmaz, A., Oztekin, A., Zaim, S., & Fahrettin Demirel, O. (2013). Universal structure modeling approach to customer satisfaction index. *Industrial Management & Data Systems*, 113(7), 932–949. doi:10.1108/IMDS-12-2012-0444.
- Turkyilmaz, A., Temizer, L., & Oztekin, A. (2016). A causal analytic approach to student satisfaction index modeling. *Annals of Operations Research*, 7(7), 209. doi:10.1007/s10479-016-2245-x.
- Xu, Y., Pinedo, M., & Xue, M. (2017). Operational risk in financial services: A review and new research opportunities. *Production and Operations Management*, 26(3), 426–445. doi:10.1111/poms.12652.