EUROPEAN CENTRAL BANK

EUROSYSTEM

Working Paper Series

Massimo Ferrari, Helena Le Mezo

Text-based recession probabilities

No 2516 / January 2021

Abstract

This paper proposes a new methodology based on textual analysis to forecast U.S. recessions. Specifically, the paper develops an index in the spirit of Baker et al. (2016) and Caldara and Iacoviello (2018) which tracks developments in U.S. real activity. When used in a standard recession probability model, the index outperforms the yield curve based forecast, a standard method to forecast recessions, at medium horizons, up to 8 months. Moreover, the index contains information not included in yield data that are useful to understand recession episodes. When included as an additional control to the slope of the yield curve, it improves the forecast accuracy by 5% to 30% depending on the horizon. These results are stable to a number of different robustness checks, including changes to the estimation method, the definition of recessions and controlling for asset purchases by major central banks. Yield and textual analysis data also outperform other popular leading indicators for the U.S. business cycle such as PMIs, consumers' surveys or employment data.

Keywords: U.S. recessions, forecast, textual analysis.

JEL Codes: E17, E47, E37, C25, C53.

Non-technical Summary

There are several empirical regularities that can be used to predict U.S. business cycle developments. Between them, the most commonly used indicator is the yield curve, i.e. the sequence of returns on government bonds with different residual maturities. From those data the so-called slope of the yield curve, i.e. the difference between the short and the long term interest rate, is extracted and used in recession probability models. Those models capture market participants expectations about future recessions. The intuition is that when agents expect a recession in the future they start rebalancing their asset portfolio, affecting the price (and hence returns) of government bonds. Several studies, Estrella et al. (2003) and Wright (2006), have shown that there is indeed a strong empirical relation between future recessions and the present slope of the yield curve. These historical regularities have been however challenged in the context of asset purchase programs (APP) by major central banks. APP purchases, in fact, induce changes in the yield curve structure that do not reflect agents' expectations but simply the purchases of long-term government securities by central banks. As a consequence, the relation between the slope of the yield curve and future recessions might have weakened.

There are also other popular modelling choices to predict future downturns which exploit commodity prices (in particular the copper price), surveys, stock market returns or even soft data such as consumers' purchases of households services.

This paper, instead, takes a novel approach to recession forecast. Instead of using financial market data, recessions are forecasted based on a novel index computed with textual analysis techniques and newspaper articles. Specifically, we construct an index based on more than 32,000 different sources reflecting the share of articles discussing a future U.S. recession published in a given day. This index has several practical advantages: first, it is easy to compute and interpret; second, it is available at high frequency and can be used to nowcast and monitor real-time economic developments; third, the index can be easily scaled to other countries as long as the news data coverage is sufficiently deep.

The text-based index is then used in a standard recession probability model to forecast

U.S business cycle developments. This exercise leads to three relevant conclusions. First, there is a strong empirical relation between the index and the prediction of future U.S. recessions at different horizons in the future. Second, the newspaper-based index significantly outperforms the yield curve model in forecasting recessions at short to medium horizons. Third, when the constructed indicator and the slope of the yield curve are simultaneously included in the model, the prediction accuracy is improved by 5 to 30% compared to other models. These evidences suggest that the newspaper-based index captures information that is not included in financial data and, hence, represent an important complement to already established metrics.

Those results are also robust to changes of the estimation method or different definitions of recessions. Most importantly, we account for possible APP distortions. In particular, the paper provides three robustness exercises which assess the impact of APP on U.S. recessions forecasts. First, the APP period is explicitly modelled, allowing the coefficients to vary after the start of APP by major central banks. Second, we account for APP distortions of the yield curve using an APP-neutral long term yield to compute the yield curve slope. Finally, APP might have increased the newspaper coverage of the yield curve and, considering the historical relation between the yield curve and recessions, that might have lead to more articles mentioning a U.S. recession. Consequently the index might also be biased in the latest part of the sample, distorting the empirical relation between the newspaper-based indicator and U.S. recessions. One way to account for that is to strip from the indicator all articles that mentions the yield curve. The main conclusions of the model remain unchanged after these additional tests.

1 Introduction

There are several empirical regularities in the U.S. business cycle that can be exploited to forecast U.S. recessions. Between them, a widely used indicator is the yield curve. Yield curve recession probability models exploit the slope of the yield curve, the difference between the short and the long term interest rate on Federal securities, to forecast U.S. recessions with standard binary regression methods. Estrella and Hardouvelis (1991), de Lindt and Stolin (2003) and more recently Benzoni et al. (2018), have shown that there are strong empirical regularities that make the slope of the yield curve a good recession indicator, particularly at medium term horizons. This approach has been revised in the context of asset purchase programs (APP) by major central banks. APP purchases generally target long-term security resulting in a compression of the long-end of the yield curve which does not proportionally translates in a reduction of short-term yields. As a consequence, the slope of the yield curve flattened because of policy actions and not in response to investors' believes on the future of the economy, leading to false recession signals. These concerns might be even more relevant in the U.S. for the role of global safe assets of U.S. treasuries.[1] Alternative popular modelling choices exploit commodity prices (in particular the copper price), stock market returns or even soft data such as consumers' purchases of households services to predict future economic slowdown.[2]

In this paper, we explore a different approach using information from newspaper articles. This approach is not new as, between others, newspaper data have already been used to measure geopolitical risk, Caldara and Iacoviello (2018), or trade policy uncertainty, Baker et al. (2016). This data source has the advantage of capturing a convolution of hard and soft information on the state of the economy at high frequency which might complement those provided by other more traditional indicators such as surveys, financial market and production inventories. Moreover, newspaper-based data are available at a daily frequency and can be updated fairly easily for a large set of economies.

[1]Moreover, Kearns et al. (2018) show that monetary policy spillovers on yields are stronger for economies which are more financially integrated, such as the US.

[2]See Liu and Moench (2016) and Hwang (2019) for a comparison of different indicators.

Relative to other applications, our paper innovates over three dimensions. First, we extend Baker et al. (2016) approach to study recessions and build an index that traces developments of U.S. real activity. The index correlates well with several U.S. business cycle indicators, such as the NBER recessions dummies or the IP growth. Second, we show that the constructed index outperforms the standard yield-curve model at medium horizons (from one to 8 months) in a formal -standard- regression framework. Finally, our results suggest that the text-based index provides complement information to the yield curve. Specifically, adding the index to a standard yield curve model specification, the forecast accuracy is improved by 5 to 30% depending on the forecast horizon. Our results are robust to a number of robustness exercises and to the comparison with other popular leading indicators.

2 Related literature

Our paper relates to two main strands of literature: recession probability modelling and the use of textual analysis tools in economics.

Recession probability modelling: The use of interest rates and their term structure as a predictor of recessions has been widely studied in the literature. Early studies such as Kessel (1965) and Fama (1986) found a strong link between the flattening of the yield curve (i.e the reduction of the differential between the short-term and the long-term rate) and periods of recessions or economic slowdowns. In times of economic uncertainty, market participants anticipate that central banks will cut the future policy rate to provide monetary policy accommodation. The expectation of lower future rates reduces longer-term rates, and this could result in an inverted yield curve.[3] To the extent that markets forecast correctly downturns of the economic activity, the flattening of the yield curve today can be used as a signal of higher probability of a future recession. Since then, there have been several empirical evidences that the term structure of interest rates have

[3]Another explanation would be that market participants expect a future recession and agents rebalance their portfolio acquiring safe assets. That leads to a rise in the price of bonds and a consequent fall of their remuneration.

greater predictive power than other leading economic and financial indices. In particular, Estrella and Mishkin (1996) compare the Treasury yield curve with other widely used indicators of future economic activity, finding that the differential between long and short-term interest rates outperforms many traditional forecasting measures, including the Commerce Department's index. In a similar spirit, more recently Gilchrist and Zakrajšek (2012) has proposed the "excess bond premium" as a leading indicator of the U.S. business cycle.

Yield curve information have been thus included in different forecasting frameworks for GDP, consumption and employment, see Stock and Watson (2003). A more recent approach is to use probit models to forecast the probability of a future recession captured by NBER recession dates.[4] Estrella and Hardouvelis (1991),Chauvet and Potter (2002) and Wright (2006) show that this framework delivers a more accurate forecast for recessions than continuous variable specifications. Another advantage of using a binary recession indicator variable is that it isolates the start and duration of recessions, whereas models forecasting continuous economic variables might suffer more from endogeniety issues. Specifically, Estrella and Hardouvelis (1991) and Wright (2006) , show that simple probit regressions give more precise recession forecasts than those produced by more sophisticated models, for 1 to 12-month ahead horizons. These types of models, however, have been recently criticized after the implementation of APP by major central banks. Asset purchases compress the long-end of the yield curve, possibly even more prominently for U.S. treasuries, distorting the historical relation between the yield curve and real activity, Engstrom and Sharpe (2018). The relevance of such distortion is still debated, Debortoli et al. (2019), for example, suggest that the impact should be modest.[5] In this paper we employ several robustness checks to account for the the potential distortions induced by the APP period.

[4]See Favara et al. (2016) and Federal Reserve Bank of St. Louis, NBER based Recession Indicators for the United States from the Period following the Peak through the Trough [USREC], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/USREC.

[5]There is a large literature trying to estimate the distortions generated by asset purchases on the yield curve and the implications for recession forecasts. Between the numerous works on the topic see: Engstrom and Sharpe (2018), Christensen (2018),Bauer and Mertens (2018b), Bauer and Mertens (2018a) and Gräb and Tizck (2020). At the same time Chauvet and Potter (2002) Chauvet and Potter (2005) find evidence of breaks in the relation between the yield curve and economic activity.

Textual analysis in economics: There is a rising strand of literature that uses textual analysis methods to construct economic indicators. For example, Hansen et al. (2018) study monetary policy decision-making by analysing the FOMC minutes via computation linguistic algorithms; Ke et al. (2019) construct an index of investors' sentiment by applying textual analysis methods to the Dow Jones Newswires; Gholampour and van Wincoop (2019) extract information from the tweets of professional traders to explain exchange rate movements; finally Minesso Ferrari et al. (2020) analyse trade-related announcements on social media to construct an indicator of trade tensions between the U.S. and China. Against this background, press articles are another important source of information as they reflect how agents perceive the state of the economy which can consequently alter their confidence and behaviour. Matsusaka and Sbordone (1995) and Batchelor and Dua (1998) have shown that consumer sentiment helps in predicting GDP, even after controlling for other relevant variables.[6] More recently, Kalamara et al. (2020) use data extracted from UK major newspapers to improve the prediction of an empirical business cycle model.

Textual analysis methods have also been widely used to measure uncertainty and estimate its role in business cycle fluctuations. Baker et al. (2016) and Caldara and Iacoviello (2018) create measures of uncertainty based on the share of articles discussing U.S. or global geopolitical tensions. They show that such indices can be used, in a VAR setting, to disentangle the role of "geopolitical tension" shocks. Caldara et al. (2020) use a similar approach to analyse the role of trade tensions. In this paper we use text data to create an high-frequency measure of the state of the U.S. business cycle. There are several benefits in using sentiment series derived from newspapers articles. First, they incorporate news and events that might take time to be reflected in more traditional, but low-frequency, macroeconomic variables. Moreover, this type of data provides useful signals on the present and future state of the economy summarising several "hard" and

[6] Taylor and McNabb (2007) investigate the role of consumer and business confidence for four European economies and find that sentiment variables indeed play a role in the prediction of recessions. More recently, Christiansen (2012) shows that sentiment variables improved significantly both in-sample and out-of-sample performance of recession prediction models, especially when paired with common recession predictors.

"soft" information that are discussed by economic commentators. Finally, the correlation between this indicator and the state of the economy is more likely to be stable over time, not being distorted by APP.

3 Data

Our main source of data is Factiva Analytics which gathers articles from major newspapers around the globe since the early 80s.[7] Importantly, the platform allows to sort articles by topic, geographical location and language. In this paper, we consider only English articles published by US-based newspapers and covering domestic news. In this way we avoid contamination by news about foreign recessions or global shocks which do not affect the U.S. to focus only on factors relevant for the domestic economy.

Similarly to Caldara et al. (2020), Baker et al. (2016) and Caldara et al. (2020), we compute our indicator for the U.S. business cycle as the share of newspaper articles discussing a recession or a slowdown in the U.S. published in each day.[8] Intuitively, when a crisis is about to occur, newspapers devote more space to the discussion of economic slowdown compared to periods of expansion. It is also important to construct the index using the share of articles. In fact the total number of articles published has grown over time, reflecting the increase in the number of pages of published newspapers. That would naturally increase the number of articles discussing a contraction in the most recent part of the sample preventing a comparison across time. The share, instead, should be relatively unaffected by the rise in the media coverage and, thus, provides a more consistent indicator. Baker et al. (2016) and Caldara et al. (2020) use the same approach to construct their political and trade uncertainty indices.

Specifically, for each calendar day from January 1^{st} 1985 to September 2020 we compute the share of articles discussing a recession in the US economy relative to all articles

[7] Factiva collects news from 32,000 different sources including, among many others, The Wall Street Journal, the New York Times as well as Dow Jones articles.

[8] This methodology is indeed similar to the R-word index developed by The Economist. See *Gauging the gloom*, The Economist, Sep. 16^{th} 2011.

published that day. The standardization by the total number of articles is essential to account for the increase in the number of articles published on all topics since the 1980s. Our approach has the advantage to be simple and easily scalable to other countries. The resulting index can be computed on each calendar date, allowing to assess relatively high frequency changes in newspaper sentiment and, hence, being a good nowcasting tool. To formally compare the information extracted from newspapers to those provided by financial data, we aggregate the index at monthly frequency, which is the frequency most commonly used for the analysis of US recessions, see Wright (2006).

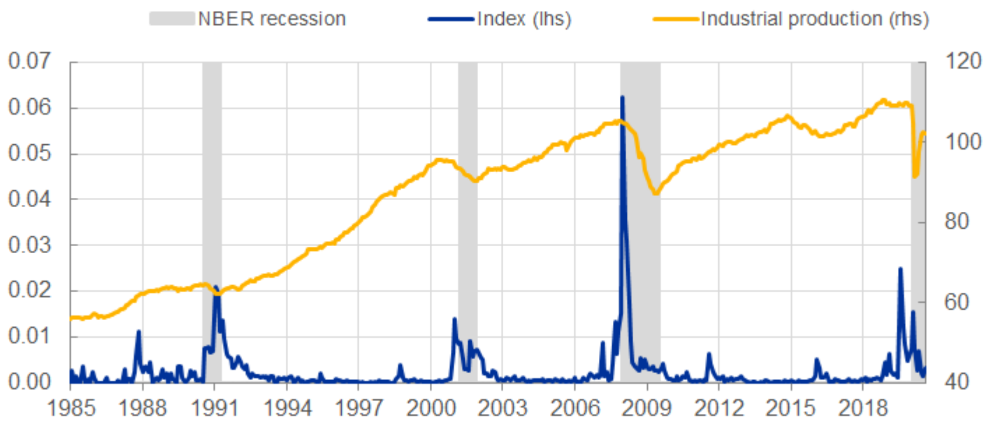

Figure 1: Newspaper-based index of US recessions.

Notes: the text-based index (left scale) is reported against NBER recession dummies (shaded areas) and US industrial production (right scale). The index is computed as the share of articles discussing a slowdown of the U.S. economy relative to all other articles published in the U.S.

Figure 1 plots the index against NBER recession dummies and U.S. industrial production. The index picks up well the 2001 and 2008 recessions and, more generally, correlates with contractions of U.S. output. Notice that the NBER defines a month to be in a recession with a lag of about two quarters.[9] Therefore, when the indicator picks up a recession in a given month, that month has not yet being defined as in a recession period. Figure 1 also suggests that the newspaper-based index could be a good coincidence indicator for U.S. real activity. In the following of the paper, we formally test its nowcasting and forecasting properties in a standard recession probability model.

[9]The NBER defines a recessions as "*two consecutive quarters of negative real gross domestic product (GDP) growth*", see The NBER's Business Cycle Dating Procedure.

Table 1: Descriptive Statistics

Variables	Mean	Standard Deviation	Time periods
Index	0.218	0.470	Jan 1985-Sep 2020
Spread	1.75	1.12	Jan 1985-Sep 2020

4 The information content of newspaper data

4.1 Forecasting U.S. contractions

To formally test the predictive power of the index, we include it in a recession probability model in the spirit of Ang et al. (2006) and Wright (2006).[10] Those frameworks are based on a probit regression where the probability of being in a recession in a future horizon $t+k$ is explained by the slope of the yield curve at t (the difference between short- and long-term rate).[11] In our specification, instead of exploiting the yield differential, we use the newspaper-based index to forecast U.S. recessions. More formally the model is specified as follows:

$$P\left(Recession_{t+k} = 1|t\right) = \Phi\left[\alpha_k + \beta_k Index_t\right] \tag{4.1}$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function. Recession periods are defined using NBER recession dummies and Equation (4.1) is estimated for $k \in [1, 16]$ at monthly frequency. By construction, the model conditions the prediction for periods $t+k$ only on the information available at time t.[12] In spirit, this framework is similar to the local projection by Jordá (2005) where the outcome variable in the future is conditional on the information set available in the present. Estimates of β_k are reported in Table 2 along with two fit statistics: the log-likelihood function and the pseudo-R^2. Overall, estimates show a strong and positive relation between the index and future recessions at any time horizon.

[10] Other relevant examples are: King et al. (2007), Chen et al. (2011) and Johansson and Meldrum (2018).

[11] The properties of these models are discussed in Estrella et al. (2003).

[12] Additionally agents do not know if they are entering a recession in period t as, generally, NBER recessions are announced, by construction, with a lag of 2 quarters. As a result, the NBER announcement of a recession period is likely to be exogenous to the index in t.

The predictive power of the index is statistically significant at all horizons, but the link with future recessions is stronger in the medium run (up to 8-month ahead). At longer horizons, larger than one-year ahead predictions, the index maintains a statistically significant relation with the probability of future recessions but explains a relatively small share of the volatility of the dependent variable. This most likely reflects the difficulty of producing accurate long-term forecasts of recessions with high frequency indicators. Long-term economic dynamics are likely to be captured by structural indicators or credit variables, which correlate more with long-term economic trends.

Turning to quantitative estimates, marginal effects describe the increase in the probability of a future recession for a 1% increase in the newspaper-based index. In this model, marginal effects are just the derivative of the dependent variable for one of the explanatory variable;[13] in Equation (4.1) the marginal effect is: $\frac{\partial P(Recession_{t+k}=1|t)}{\partial Index_t}$. Figure 2 plots the marginal effect of a 1% increase in the index (i.e. the share of articles discussing future recession rises by 1 p.p.). A 1% increase in the newspaper indicator leads to a 20% higher likelihood of recession in the next 4 months and about 10% higher in a 16 months horizons. To understand the economic significance of these numbers one can consider historical episodes. For example, in the run-up to the global financial crisis the indicator rose by about 1.5 to 2 p.p., implying a 30 to 40% probability of a crisis in the next year, see Figure 1.

Taken together, these evidences support the initial claim that newspaper-based information contain indeed relevant data to study business cycle fluctuations. To contrast these results against a standard recession model, we estimate Equation (4.1) using the slope of the U.S. yield curve instead of the newspaper index. Table A.4 and Figure A.1, reported in Appendix A, report the results of that estimate. Results suggest that our constructed indicator captures a different set of information compared to what is implied by yields. Specifically, the relation between the slope of the yield curve and future recession is not statistically significant at short horizons. Moreover, the newspaper-based index greatly outperforms the fit of the yield curve model up to 8-months in the future. At longer

[13] Given the non-linear nature of the estimator, they capture the change in the probability of a positive outcome for a unitary increase in that specific regressor (keeping the others constant).

horizons, instead, the explanatory power of the yield curve model (as defined by the pseudo-R^2 and the log-likelihood) starts to catch up with our baseline specification.

Table 2: Estimation results of Equation (4.1) from 1 to 16 months ahead

Months ahead	β_k	*s.e.*	Obs.	Log Likelihood	R^2
1	1.51	0.30	428	-102.27	0.24
2	1.46	0.29	427	-102.96	0.24
3	1.44	0.29	426	-103.74	0.23
4	1.50	0.30	425	-101.15	0.25
5	1.53	0.31	424	-100.13	0.26
6	1.65	0.30	423	-98.19	0.27
7	1.55	0.29	422	-102.06	0.24
8	1.36	0.25	421	-107.85	0.20
9	1.22	0.2	420	-111.58	0.17
10	1.16	0.22	419	-113.06	0.16
11	1.02	0.20	418	-116.74	0.13
12	0.96	0.19	417	-118.35	0.12
13	0.89	0.20	416	-120.92	0.10
14	0.77	0.18	415	-123.48	0.08
15	0.68	0.17	414	-125.21	0.06
16	0.51	0.15	413	-128.12	0.04

Notes: Coefficients and robust standard errors of Equation (4.1) estimated at different forecast horizons. The R^2 is the pseudo-R^2.

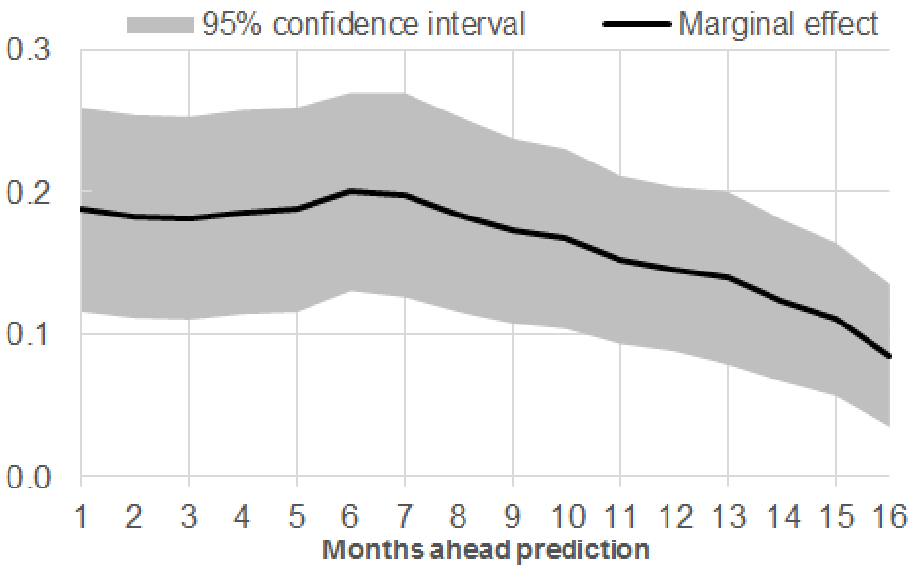

Figure 2: Marginal effects from Equation (4.1).

Notes: marginal effects $\left(\frac{\partial P(Recession_{t+k}=1|t)}{\partial Index_t}\right)$ from the probit regression for a 1% increase in the newspaper-based index (i.e. a 1% increase in the share of newspaper articles discussing a recession in the US). The grey shaded areas show 95% confidence intervals.

Previous results suggest that the newspaper-based index has the potential to outperform yield curve models at medium horizons and to complement the information provided by financial data when forecasting the more distant future. To formally test this hypothesis, we run a third specification of the model where we include both the index and the slope of the yield curve:

$$P\left(Recession_{t+k} = 1|t\right) = \Phi\left[\alpha_k + \beta_k Index_t + \gamma_k\left(r_t^{3-month} - r_t^{10-year}\right)\right] \tag{4.2}$$

Table 3 and Figure 3 report the regression coefficients and the marginal effects for each forecast horizon of Equation (4.2). There are several interesting results from this exercise. First, only the index is significant at all forecast horizons. This suggests that indeed at short to medium-term the yield curve does not provide additional information compared to those already present in the newspaper-based indicator. At longer horizons, instead, Table 2 implies a somewhat weaker performance of the indicator compared to yield data. Notably, however, the index maintains its explanatory power even when yield-curve data are included in the model. This last result shows that the indicator contains some information that are useful for long-term forecasts that are not already captured by information present in spread between short- and long-term rates. Finally, Table 3 shows that the use of both indicators improves the overall fit of the model as captured by a simple metric like the log-likelihood or the pseudo-R^2. Those simple measures of fit suggest, again, that both variables provide valuable (and different) information on future developments of the U.S. business cycle. Finally, the elasticities of the recession probability to the index and the yield curve are very similar between the baseline model and Equation (4.2), see Figure 3. That, again, implies that the two variables contain complementary information on the future of the economy.

Table 3: Estimation results of Equation (4.2) from 1 to 16 months ahead

Months ahead	β_k	*s.e.*	γ_k	*s.e.*	Obs.	Log Likelihood	R^2
1	1.54	0.30	-0.05	0.08	428	-102.10	0.24
2	1.45	0.29	0.04	0.07	427	-102.87	0.24
3	1.40	0.29	0.11	0.07	426	-102.90	0.24
4	1.45	0.31	0.18	0.08	425	-99.01	0.27
5	1.46	0.32	0.26	0.09	424	-95.58	0.29
6	1.65	0.29	0.38	0.10	423	-89.07	0.34
7	1.61	0.29	0.49	0.11	422	-87.44	0.35
8	1.48	0.26	0.64	0.13	421	-85.62	0.36
9	1.39	0.24	0.77	0.14	420	-82.37	0.39
10	1.41	0.24	0.91	0.17	419	-77.57	0.42
11	1.29	0.24	1.00	0.19	418	-76.15	0.43
12	1.34	0.27	1.18	0.22	417	-71.01	0.47
13	1.41	0.30	1.38	0.24	416	-66.82	0.50
14	1.35	0.31	1.52	0.26	415	-64.72	0.52
15	1.16	0.29	1.55	0.25	414	-65.12	0.51
16	0.85	0.24	1.46	0.21	413	-69.72	0.48

Notes: Coefficients and robust standard errors of Equation (4.2) estimated at different forecast horizons. The R^2 is the pseudo-R^2.

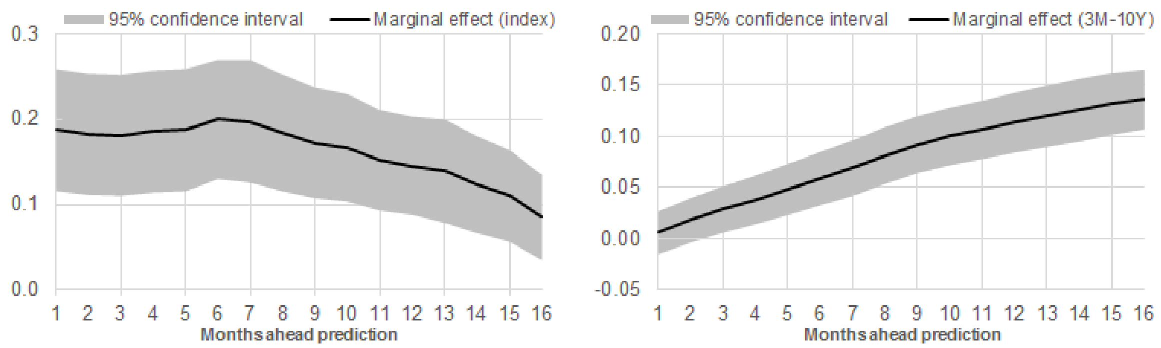

Figure 3: Marginal effects from Equation (4.1).
Notes: marginal effects $\left(\frac{\partial P(Recession_{t+k}=1|t)}{\partial Index_t}\right)$, $\left(\frac{\partial P(Recession_{t+k}=1|t)}{\partial r_t^{3-month} - r_t^{10-year}}\right)$ from the probit regression for a 1% increase in the newspaper index (left panel) and a 1% increase in the slope of the yield curve (right panel). The grey shaded areas show 95% confidence intervals.

4.2 A formal test for the goodness of fit

To formally test the predictive power of the index and compare it against the yield curve, we rely on the so-called *Receiver Operator Curves* (ROC)[14] analysis, also referred to as

[14]This statistics takes its name from radar controls. The ROC curve was first used during World War II in the context of the analysis of radar signals to quantify the ability of operators to correctly identify

the "area under the curve", which is the most commonly used method to evaluate forecasts of non-linear outcomes.

For a binary choice, this curve summarises the ability of a model to predict the correct binary outcome, i.e. a 0 (1) when a 0 (1) actually occurs. In a nutshell, the ROC is constructed by computing the true positive rate (i.e. the share of correct predictions of the model) against the false positive rate (i.e. the share of wrong predictions). Both rates are defined for a given threshold of the outcome variable, in this case a level of the predicted probability above (below) which a recession is considered to be predicted. The set of true positives and true negatives for different threshold values is then plotted as the ROC curve, showing the share of correct predictions of the model for any given threshold. The ROC curve is generally contrasted against the 45 degree line (the random model).[15] The more accurate is a model, the further away from the 45 degree line its ROC curve should lie, as the model should predict significantly better than a random choice. These curves are reported in Figure A.2 for the 1-month and 16-month ahead forecast.

The entire plot can also be summarized by the so-called ROC *statistic* defined as the ratio between the area below the ROC curve but above the 45 degree line and the total area above the 45 degree line. Intuitively, the ROC statistic captures the accuracy of the model as better models would have ROC curves further away from the random assignment curve and, hence, the distance between the two lines should be larger. The ROC statistic is naturally bounded between 0 (a model not different from coin-flipping) and 1 (the perfect model that always delivers correct predictions). Figure 4 reports the ROC statistics for the baseline model, Equation (4.1), the yield curve model[16] and the model including both variables, Equation (4.2).

incoming aircrafts. After the attack on Pearl Harbor, the United States army began new research to increase the accuracy of detection of Japanese aircrafts. Specifically, U.S. officials measured the ability of a radar receiver operator to make the distinction between enemy aircrafts (true positives) and noise (false positive) when reading radar signals on the screen. The resulting ratio between accurate and wrong calls was named the Receiver Operating Characteristic. Since then, the method has been widely adopted in statistics to quantify the level of true predictions delivered by models for non-linear outcomes.

[15]The 45 degree line represents the ROC curve for a random prediction model where both outcomes have the same probability of being chosen. For a binary model, in fact, the share of correct predictions of the random model (50%) does not depend on the value of the threshold.

[16]The specification of the yield curve model is reported in Equation (A.1) of Appendix A.

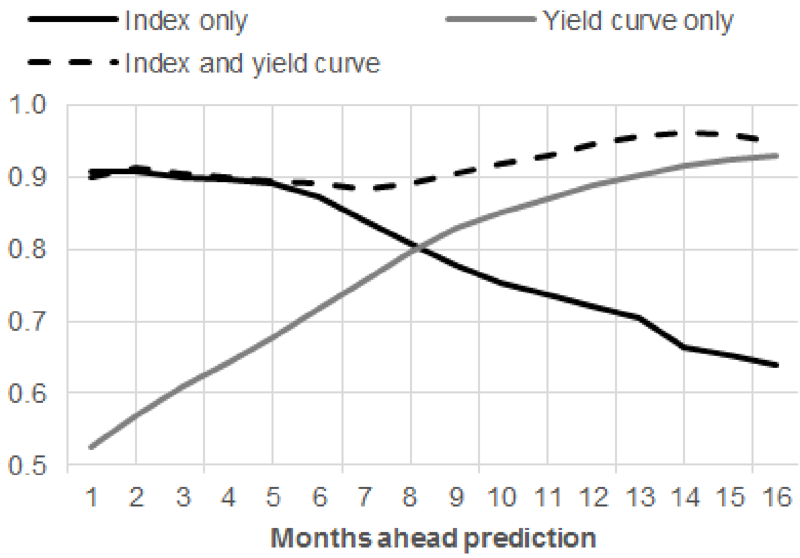

Figure 4: ROC statistic for the three models (index, yield curve, index and yield curve) at different forecast horizons.

Notes: the ROC statistic is computed as the ratio between the area below the ROC line but above the 45 degree line and the total area above the 45 degree line. It is naturally bounded between 0, the random assignment model, and 1, the model with all correct predictions.

Results reported in Figure 4 show several relevant properties of our index. First, Figure 4 confirms our previous results showing that the index outperforms the yield curve model at medium horizons, from 1 to 8 months. At those horizons predictions based on the index are significantly more accurate than those based on yields.[17] Second, although it is true that the index looses some degree of predicting power at longer horizons, after 9 months in the future, the share of accurate forecasts is still relevant up to 13 months in the future. Finally, the model that combines both variables gives "the best of both worlds", leveraging on the fact that the index and the yield curve contain complementary information useful for long-term forecasts. In particular, the ROC based on Equation (4.2) follows closely the accuracy of the index-only model at short horizons, leveraging on the strong forecasting properties of the newspaper-based index in the medium term. In the long run, Equation (4.2) exploits both indicators improving the ROC statistic relative to the yield curve model by about 15 p.p. (for the 8 to 11-month ahead prediction) to about 10 p.p. (for horizons longer than one year).

[17]This result confirms the preliminary evidences of Table 2 and Table A.4.

4.3 Predicted recessions

Equation (4.1) and Equation (4.2) can be used to compute implied recession probabilities at different horizons. Figure 5 shows the implied probabilities at the 1 to 12-month horizons for the three models considered before. Predictions of the index-only model are reported in the first row, those of the yield curve model in the second and those based on the model with both variables in the last one. As already suggested by Table 2, the index-based model forecasts well recession up to medium horizons; see, for example, the third chart in the upper panel of Figure 5. This result is quite remarkable considering that when the model picks up a recession in a given month the NBER has not yet defined the beginning of a recession period (the NBER requires 2 quarters, 6 months, of consecutive contractions to mark the beginning of a recession). However, at longer horizons, consider for example the fourth chart in the first panel of Figure 5, Equation (4.1) indeed delivers less accurate results generating more false-positives (i.e. predictions of a recession when no recession occurs). Forecasts based on the index are, however, on average more accurate than those based on the yield curve model. Contrast the first and the second panel of Figure 5. The yield curve model fails to pick up recessions at the 1 and 4-month ahead horizons. Moreover, at 8 and 12-month ahead the yield curve implies a significant share of false-positives, for example in the period 2003-2005, larger than what implied by the index-based model. Finally, the last row of Figure 5 shows the implied recession probabilities based on Equation (4.2). This model outperforms both specifications particularly at longer horizons. Consider the last chart of the lower panel in Figure 5 which shows the 12-month ahead prediction. Equation (4.2) sends much better signals ahead of the 2008 recession compared to both other specifications and does not show the same number of false positives implied by the yield curve model.

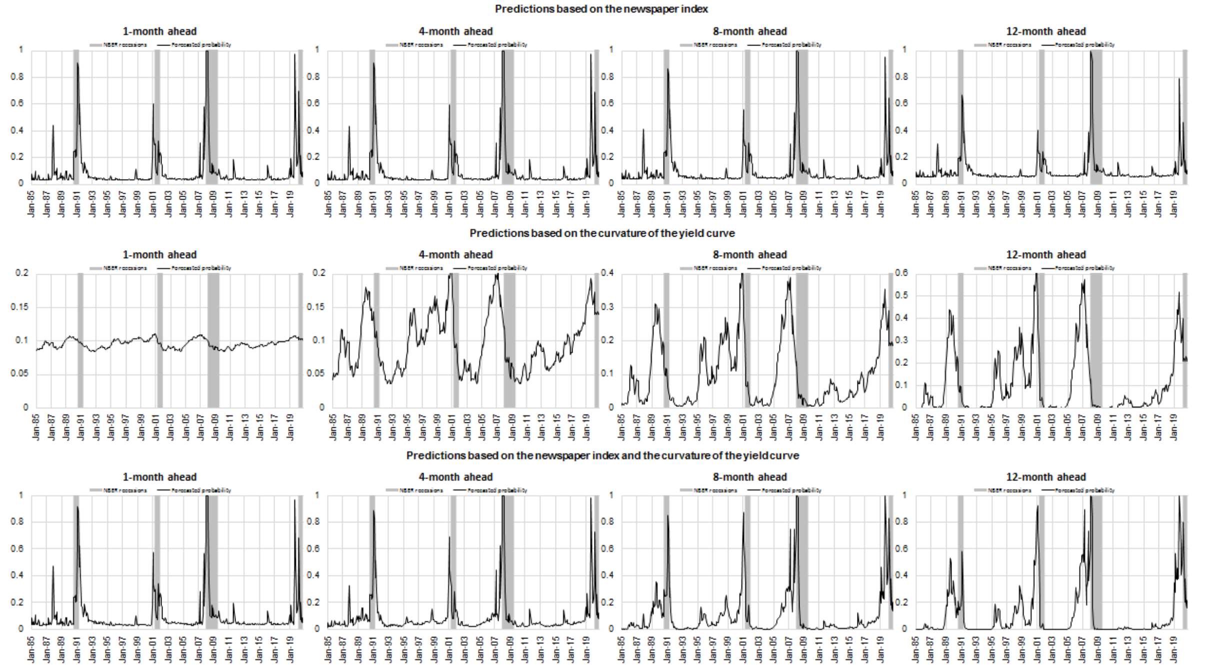

Figure 5: Predicted recession probabilities at different forecast horizons.
Notes: recession probabilities are computed based on Equation (4.1), for the index and the yield curve model, and Equation (4.2) for the model including both the yield curve and the index.

4.4 Robustness

Results reported in Section 4.1 and Section 4.2 are robust to different changes in the model.

Different estimation method. Our first robustness check is to change estimation method. We re-estimate Equation (4.1) and Equation (4.2) using a logit model instead of a probit. Logistic regressions have several practical advantages. Among them, residuals are not assumed to be normally distributed and homoscedasticity is not required. Under a logit framework, moreover, coefficients can be directly interpreted as marginal effects. These characteristics relax some of the assumptions imposed in our baseline specification. However, logistic models require that observations are independent from each other, hence they should not be repeated measures from the same data generating process (such as GDP). As such, logistic regressions are less suited for the analysis of time series data. For this reason, our preferred choice for Equation (4.1) is the probit specification. Our main results are, however, robust to the change of the estimation methods. Table A.2

and Table A.3 report estimated coefficients of the logistic regression which are broadly consistent with the baseline estimates presented earlier. Also ROC statistics, plotted in Figure 6 are in line with baseline results. Implied recession probabilities are also qualitatively similar to the baseline model, with the logistic estimation exhibiting higher uncertainty resulting in more volatile implied probability and more false-positives and negatives, see Figure A.3.

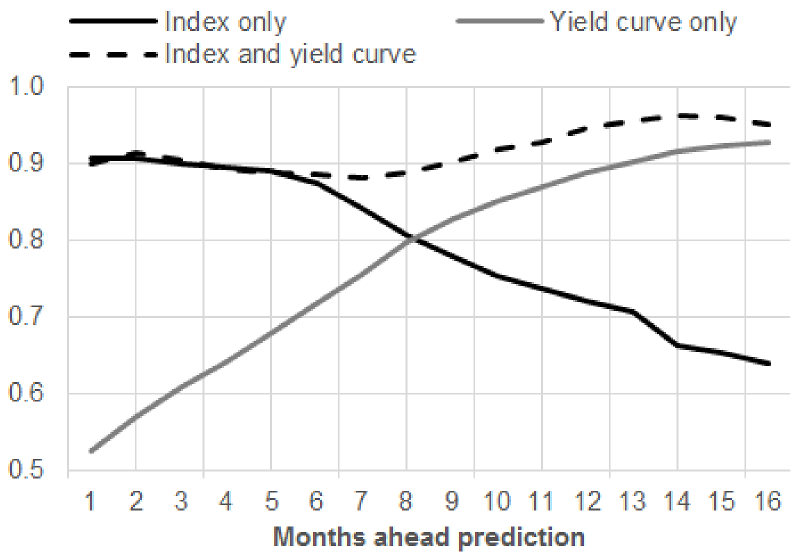

Figure 6: ROC statistic for the three models (index, yield curve, index and yield curve) at different forecast horizons estimated with a logit regression.

Notes: the ROC statistic is computed as the ratio between the area below the ROC line but above the 45 degree line and the area above the 45 degree line. It is naturally bounded between 0, the random assignment model, and 1, the model with all correct predictions.

Different definitions of recessions. As second robustness exercise, we changed the definition of recession. The NBER defines a recession as *"... a significant decline in economic activity spread across the economy, lasting more than a few months, normally visible in real GDP, real income, employment, industrial production, and wholesale-retail sales."*[18], thus weighting different indicators and measures of the business cycle. In most of the cases, the NBER requires at least 2 quarters of consecutive contractions to mark the beginning of a recession period. As robustness exercise, we change the definition of a recession to 4 or 8 consecutive months of IP contraction and estimate Equation (4.1) and Equation (4.2) with this dependent variable. These alternative definitions have also

[18]See The NBER's Business Cycle Dating Procedure: Frequently Asked Questions.

the practical advantages of increasing the number of recession periods in the sample.[19] Basing the definition of recession on IP contractions is broader than what considered by NBER dates and pics up more episodes which are related to cyclical slowdowns of U.S. real activity. Our main results are robust to this change of definition suggesting that the newspaper-based index indeed captures a rich set of information on the U.S. business cycle. Figure 7 plots the ROC statistics for these two models, confirming the baseline result that the index is superior to the yield curve up to medium horizons and that the combination of the two indicators improves the forecast at all horizons. Predicted recessions are consistent with the baseline estimates. Notably, the model pics up the 1991 recession despite it is not defined as such using this metric, see Figure A.6.

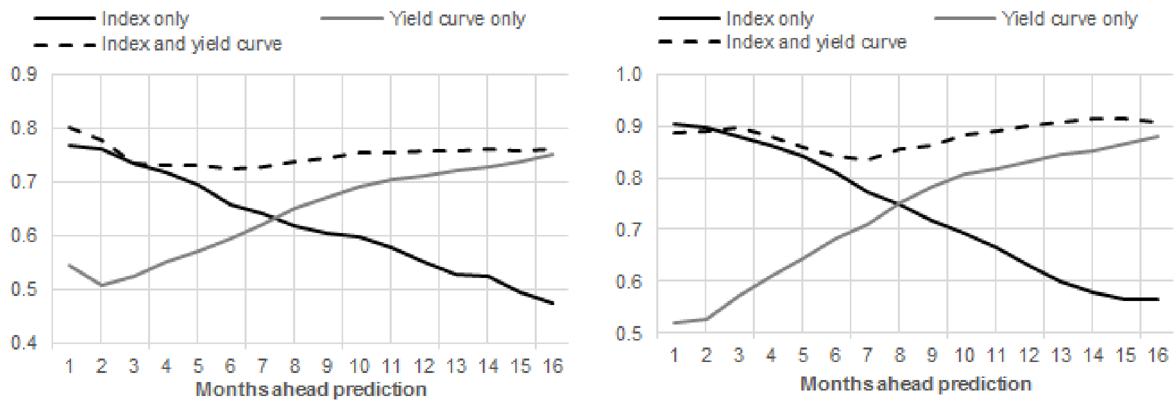

Figure 7: ROC statistic for the three models (index, yield curve, index and yield curve) at different forecast horizons and definition of recessions.

Notes: in the left panel a recession is defined as 4 consecutive months of IP contraction while in the right panel a recession is defined as 8 consecutive months of contraction. The ROC statistic is computed as the ratio between the area below the ROC line but above the 45 degree line and the area above the 45 degree line. It is naturally bounded between 0, the random assignment model, and 1, the model with all correct predictions.

Accounting for APP distortions. Since the global financial crisis all major central banks have started asset purchase programs, which have compressed global safe assets and particularly U.S. yields. This might bear important consequences for forecasting recessions. First of all, historical relations between the yield curve and recessions might be distorted, resulting in a less accurate forecasts based on treasury yields when the slope

[19]It has often been argued that binary models based on NBER recession dates suffer for the limited number of "positive" outcomes in the sample. Relaxing the recession perio definition allows to increase that number.

of the yield curve is compressed by central bank purchases rather than expectations. That could reduce the relevance of U.S. yields for forecasting the business cycle, Engstrom and Sharpe (2018).[20] As a result, also the comparison between the yield curve model and the index-based model might be biased (i.e. in the latest part of the sample, the yield curve model could have a higher false-positive ratio).

Moreover, there is another, more subtle, channel through which central bank interventions might affect our estimates. The historical relation between the flattening of the yield curve and recessions is well known. If in the more recent years the yield curve has flattened because of central banks' actions, economic commentator might have discussed that making reference to the relation between the yield curve and past U.S. recessions. As a result, our index (which is computed as the share of articles discussing U.S. recessions over all articles published) might be artificially inflated and bias results in the post global financial crisis period. To tackle these issues we propose three robustness exercises. First, we introduce a dummy for the APP period, which is interacted with other coefficients. Formally, we modify Equation (4.1) and Equation (4.2) as follows:

$$P\left(Recession_{t+k}=1|t\right)=\Phi\left[\alpha_k+\beta_k Index_t+\delta_k QE_t+\phi_k QE_t\times Index_t\right] \tag{4.3}$$

$$\begin{aligned} P\left(Recession_{t+k}=1|t\right)=\Phi\Big[&\alpha_k+\beta_k Index_t+\gamma_k\left(r_t^{3-month}-r_t^{10-year}\right)+\delta_k QE_t+\\ &+\phi_k QE_t\times Index_t+\vartheta_k QE_t\times\left(r_t^{3-month}-r_t^{10-year}\right)\Big]\end{aligned} \tag{4.4}$$

where QE_t is a dummy taking value of 1 for the period post November 2008 (when the FED first started quantitative easing interventions). This is a simple way to account for possible non-linearities in the APP period as the interaction terms should capture exactly the changed relation between the dependent variables and NBER recession dummies during the APP. Second, we correct the time series for the slope of the yield curve to account for APP effects and use the corrected series for the estimation. In this we follow Gräb et al. (2019).[21] Third, all articles discussing the yield curve are excluded from the

[20]There are however other studies suggesting that central banks' APP had limited impact on recession probability models based on the slope of the yield curve, see Bauer and Mertens (2018a), Bauer and Mertens (2018b), Christensen (2018) and Debortoli et al. (2019).

[21]Their procedure involves estimating an autoregressive distributed lag model for the U.S. term premium. In the model the term premium is regressed on foreign central bank holdings, domestic central

newspaper-based index. In this way all sources of interaction between the lower interest rate environment and recessions are mechanically excluded. Indeed, this choice eliminates several potentially relevant information on the U.S. business cycle. However, it has the advantage of producing very conservative estimates. If the main results hold with this definition of the index, they cannot be driven by any distortion induced by central bank interventions on the yield curve.

We report the ROC statistics for the three robustness exercises in Figure A.7, Figure A.8 and Figure A.9 respectively. The main results of the paper are robust to these three additional checks. Most importantly, the ROC statistics lead to the same conclusions as in the baseline specification: the index forecasts better than the yield curve at medium horizons and the model featuring both our index and the yield curve improves the statics between 5% to 30% depending on the horizon considered.

Alternative business cycle indicators. There are alternative business cycle indicators that have shown robust forecasting properties. Hwang (2019), for example, compares the forecast accuracy of financial data, employment statistics and surveys. The "excess bond premium" by Gilchrist and Zakrajšek (2012) has also proved to be a good leading indicator for the U.S. business cycle.[22] To gauge the robustness of our findings and compare forecasts based on Equation (4.2) against those delivered by alternative indicators we estimate the following model:

$$P\left(Recession_{t+k} = 1|t\right) = \Phi\left[\alpha_k + X_t \Gamma_k'\right] \tag{4.5}$$

where X_t is a matrix of alternative indicators and Γ_k a vector of loadings. We consider four sets of leading indicators: the S&P500 and the 3-month Libor-TBill spread; unemployment rate and weekly hours worked; the PMI and the Michigan consumers' confidence survey; the excess bond premium and the Gilchrist-Zakrajšek credit spread. The ROC statistics based on these models are compared against those implied by Equation (4.2)

bank holdings and country-specific fundamentals, including the level and volatility of growth and inflation. The corrected series are the fitted values from the regression. A similar method is used by Gräb and Tizck (2020).

[22] For a review of results based on Gilchrist and Zakrajšek (2012) see Favara et al. (2016).

in Figure A.10, where the dependent variable are NBER dummies, and Figure A.11, where recessions are defined as 4 months of consecutive IP contraction. Using NBER recession dummies our preferred model outperforms all other models at long horizons and is dominated by financial variables up to 8 months in the future. Results are similar when recessions are defined based on IP. Under this specification, the Gilchrist-Zakrajšek perform better than our model at short horizons. Overall, Figure A.10 and Figure A.11 show that the newspaper-based index performs well even when compared to some of the most used U.S. leading indicators.

5 Conclusion

This paper develops a text-based indicator to track U.S. economic activity slowdowns. The index is based on the share of newspaper articles that discuss a recession in the U.S. in each day. This approach, which is essentially an extension of Baker et al. (2016) to real variables, has several advantages: first, it is easy to replicate and does not require involved computational techniques; second, the indicator can be easily scaled to any country for which newspaper information are available; third, the indicator can be updated on a daily basis, to be used for nowcasting and monitoring high frequency developments; fourth the index summarizes a rich set of economic information which are not necessary present in financial market data.

We test the properties of the index in a standard recession probability model, showing that the index outperforms the standard yield curve model at medium horizons (up to 8 months). Moreover, we show that the text-based index adds relevant information on top of what is implied by the slope of the yield curve at all forecast horizons. This is formally tested with the use of ROC statistics, which show that the model featuring both the index and the yield curve outperforms other models by 5% to 30% depending on the forecast horizon. Our results are robust to a broad range of robustness checks: different estimation methods, different definition of recessions and, most importantly, accounting for distortions induced by asset purchase programs. Finally, we show that the text-based index performs well when compared to other U.S. leading indicators.

This paper, in conclusion, shows the potentials of the use of high-frequency textual analysis methods to forecast real economic activity. Such methods include a broad range of economic information, ranging from hard production data to consumers surveys, in one single indicator which complements those provided by financial data. Results from our estimations are promising; we thus leave to future research the integration of our index in more complex forecasting models based, for example, on smooth transitions or Markov chains.

References

A. Ang, M. Piazzesi, and M. Wei. What does the yield curve tell us about gdp growth? *Journal of Econometrics*, 131(1-2):359–403, 2006.

R. S. Baker, N. Bloom, and S. J. Davis. Measuring economic policy uncertainty. *Quarterly Journal of Economics*, 131(4):1593–1636, 2016.

R. Batchelor and P. Dua. Improving macro-economic forecasts: The role of consumer confidence. *International Journal of Forecasting*, 14:71–81, 1998.

M. D. Bauer and T. M. Mertens. Economic forecasts with the yield curve. *FRBSF Economic Letter*, 2018a.

M. D. Bauer and T. M. Mertens. Information in the yield curve about future recessions. *FRBSF Economic Letter*, 2018b.

L. Benzoni, O. Chyruk, and D. Kelley. Why does the yield-curve slope predict recessions? *Federal Reserve Bank of Chicago Working Paper*, pages 2018–15, 2018.

Berge, Travis J.and Jordà. Evaluating the classification of economic activity into recessions and expansions. *American Economic Journal: Macroeconomics*, 3(2):246–277, 2011.

F. Bianchi, H. Kung, and T. Kind. Threats to central bank independence: High-frequency identification with twitter. *NBER Working Papers*, 26308, 2019.

D. Caldara and M. Iacoviello. Measuring geopolitical risk. *International Finance Discussion Papers, Board of Governors of the Federal Reserve System (U.S.)*, 1222, 2018.

D. Caldara, M. Iacoviello, P. Molligo, A. Prestipino, and A. Raffo. The economic effects of trade policy uncertainty. *Journal of Monetary Economics*, 109:38–59, 2020.

M. Chauvet and S. Potter. Predicting a recession: evidence from the yield curve in the presence of structural breaks. *Economics Letters*, 77(2):245–253, October 2002.

M. Chauvet and S. Potter. Forecasting recessions using the yield curve. *Journal of Forecasting*, 24(2):77–103, 2005.

Z. Chen, A. Iqbal, and H. Lai. Forecasting the probability of us recessions: a probit and dynamic factor modelling approach. *Canadian Journal of Economics*, 44(2):651–672, 2011.

J. H. E. Christensen. The slope of the yield curve and the near-term outlook. *FRBSF Economic Letter*, 2018.

C. Christiansen. Predicting severe simultaneous recessions using yield spreads as leading indicators. *International Journal of Forecasting*, 32:1032–1043, 2012.

R. C. de Lindt and D. Stolin. Measuring economic policy uncertainty. *Journal of Monetary Economics*, 50:1603-1622, 2003.

D. Debortoli, J. Galí, and L. Gambetti. On the empirical (ir)relevance of the zero lower bound constraint. In *NBER Macroeconomics Annual 2019, volume 34*, NBER Chapters, pages 141–170. National Bureau of Economic Research, Inc, May 2019.

E. Engstrom and S. A. Sharpe. The near-term forward yield spread as a leading indicator : A less distorted mirror. *Finance and Economics Discussion Series, Board of Governors of the Federal Reserve System (U.S.)*, (2018-055), 2018.

A. Estrella and G. A. Hardouvelis. Term structure as a predictor of real economic activity. *Journal of Finance*, 2(46), 1991.

A. Estrella and F. S. Mishkin. Tthe yield curve as a predictor of u.s recessoins. *Current Issues in Economics and Finance, Federal Reserve Bank of New York*, 2, 1996.

A. Estrella, A. Rodrigues, and S. Schich. How stable is the predictive power of the yield curve? evidence from germany and the united states. *The Review of Economics and Statistics*, 85(3):629–644, 2003.

E. F. Fama. Term premiums and default premiums in money markets. *Journal of Financial Economics*, 17(1):175–196, 1986.

G. Favara, S. Gilchrist, K. F. Lewis, and E. Zakrajšek. Updating the recession risk and the excess bond premium. *FEDS Notes*, 2016.

V. Gholampour and E. van Wincoop. Exchange rate disconnect and private information: What can we learn from euro-dollar tweets? *Journal of International Economics*, 119 (C):111–132, 2019.

S. Gilchrist and E. Zakrajšek. Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4):1692–1720, 2012.

J. Gräb and S. Tizck. Estimating recession probabilities with a less distorted term spread. *ECB Mimeo*, 2020.

J. Gräb, T. Kostka, and D. Quint. Quantifying the "exorbitant privilege" - potential benefits from a stronger international role of the euro. In *The international role of the euro 2019*. European Central Bank, 2019.

S. Hansen, M. McMahon, and A. Prat. Transparency and deliberation within the fomc: A computational linguistics approach. *The Quarterly Journal of Economics*, 133(2): 801–870, 2018.

Y. Hwang. Forecasting recessions with time-varying models. *Journal of Macroeconomics*, 62, 2019.

P. Johansson and A. C. Meldrum. Predicting recession probabilities using the slope of the yield curve. *FEDS Notes*, 2018.

O. Jordá. Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182, 2005.

E. Kalamara, A. Turrell, C. Redl, G. Kapetanios, and S. Kapadia. Making text count: economic forecasting using newspaper text. *Bank of England working papers*, 865, 2020.

Z. T. Ke, B. T. Kelly, and D. Xiu. Predicting returns with text data. *NBER Working Papers*, 26186, 2019.

J. Kearns, A. Schrimpf, and D. Xia. Explaining monetary spillovers: The matrix reloaded. *BIS Working Papers*, (757), 2018.

R. A. Kessel. Explanations of the term structure of interest rates. In R. A. Kessel, editor, *The Cyclical Behavior of the Term Structure of Interest Rates*. National Bureau of Economic Research, 1965.

T. B. King, A. T. Levin, and R. Perli. Financial market perceptions of recession risk. *Finance and Economics Discussion Series*, 2007-53, 2007.

W. Liu and E. Moench. What predicts us recessions? *International Journal of Forecasting*, 32(4), 2016.

J. G. Matsusaka and A. M. Sbordone. Consumer confidence and economic fluctuations. *Economic Inquiry*, 33:296–318, 1995.

M. Minesso Ferrari and H. Le Mezo. Using information in newspaper articles as an indicator of real economic activity. *Economic Bulletin Boxes, European Central Bank*, 2, 2020.

M. Minesso Ferrari, F. Kurcz, and M. S. Pagliari. Do words hurt more than actions? *ECB Working Papers*, 2020.

J. H. Stock and M. W. Watson. Forecasting output and inflation: the role of asset prices. *Journal of Economic Literature*, 41(1):788–829, 2003.

K. Taylor and R. McNabb. Business cycles and the role of confidence:evidence for europa. *Oxford Bulletin of Economics and Statistics*, 2018-15(2):185–208, 2007.

J. H. Wright. The yield curve and predicting recessions. *Finance and Economics Discussion Serie, Board of Governors of the Federal Reserve System (U.S.)*, 2006-07, 2006.

A Appendix

A.1 Tables

Table A.1: Estimation results of Equation (4.1) from 1 to 16 months ahead using the slope of the yield curve as explainatory variable:

$$P\left(Recession_{t+k} = 1|t\right) = \Phi\left[\alpha_k + \beta_k\left(r_t^{3-month} - r_t^{10-year}\right)\right] \tag{A.1}$$

Months ahead	β_k	$s.e.$	Obs.	Log Likelihood	R^2
1	0.03	0.06	428	-135.04	0.00
2	0.11	0.06	427	-134.07	0.01
3	0.17	0.07	426	-132.45	0.02
4	0.23	0.07	425	-130.39	0.03
5	0.30	0.08	424	-127.41	0.05
6	0.37	0.08	423	-123.38	0.08
7	0.45	0.09	422	-118.60	0.12
8	0.56	0.10	421	-112.22	0.17
9	0.66	0.11	420	-106.12	0.21
10	0.74	0.13	419	-101.03	0.25
11	0.82	0.14	418	-96.67	0.28
12	0.91	0.16	417	-91.74	0.32
13	1.00	0.18	416	-87.61	0.35
14	1.10	0.20	415	-83.79	0.37
15	1.21	0.21	414	-80.22	0.40
16	1.30	0.18	413	-78.22	0.41

Notes: Coefficients and robust standard errors of Equation (4.1) estimated at different forecast horizons using the slope of the yield curve (i.e. the difference between the 3-month rate and the 10-year rate on treasury bills) as explanatory variable. The R^2 is the pseudo-R^2.

Table A.2: Estimation of Equation (4.1) and Equation (A.1)with a logistic model.

	Newspaper index					Yiled curve				
Months ahead	β_k	$s.e.$	Obs.	Log Likelihood	R^2	β_k	$s.e.$	Obs.	Log Likelihood	R^2
1	2.71	0.59	428	-104.17	0.23	0.06	0.12	428	-135.05	0.00
2	2.65	0.64	427	-104.96	0.22	0.20	0.12	427	-134.13	0.01
3	2.59	0.63	426	-105.75	0.22	0.32	0.13	426	-132.57	0.02
4	2.78	0.68	425	-102.88	0.24	0.43	0.14	425	-130.55	0.03
5	2.86	0.71	424	-101.62	0.25	0.57	0.15	424	-127.53	0.05
6	3.01	0.62	423	-99.44	0.26	0.73	0.17	423	-123.32	0.08
7	2.84	0.61	422	-102.93	0.23	0.91	0.19	422	-118.20	0.12
8	2.44	0.51	421	-108.49	0.19	1.13	0.21	421	-111.52	0.17
9	2.18	0.45	420	-111.99	0.17	1.35	0.24	420	-105.15	0.22
10	2.07	0.43	419	-113.33	0.16	1.53	0.27	419	-99.71	0.26
11	1.80	0.38	418	-116.96	0.13	1.70	0.30	4418	-94.99	0.29
12	1.68	0.35	417	-118.52	0.12	1.89	0.34	417	-89.93	0.33
13	1.57	0.36	416	-121.08	0.10	2.07	0.37	416	-85.87	0.36
14	1.36	0.32	415	-123.52	0.08	2.24	0.39	415	-82.22	0.39
15	1.18	0.28	414	-125.25	0.06	2.37	0.38	414	-79.56	0.40
16	0.88	0.26	413	-128.22	0.04	2.40	0.34	413	-78.91	0.41

Notes: Coefficients and robust standard errors Equation (4.1) and Equation (A.1) estimated at different forecast horizons using the newspaper based index and the slope of the yield curve (i.e. the difference between the 3-month rate and the 10-year rate on treasury bills) as explanatory variables. The R^2 is the pseudo-R^2.

Table A.3: Estimation results of Equation (4.2) from 1 to 16 months ahead with a logistic model.

Months ahead	β_k	*s.e.*	γ_k	*s.e.*	Obs.	Log Likelihood	R^2
1	2.76	0.59	-0.09	0.15	428	-104.01	0.23
2	2.62	0.64	0.08	0.14	427	-104.85	0.22
3	2.52	0.61	0.23	0.15	426	-104.74	0.22
4	2.69	0.66	0.37	0.16	425	-100.39	0.26
5	2.76	0.69	0.55	0.18	424	-96.50	0.28
6	2.96	0.61	0.77	0.21	423	-90.12	0.33
7	2.90	0.58	1.02	0.24	422	-87.68	0.30
8	2.63	0.50	1.31	0.27	421	-85.64	0.35
9	2.48	0.46	1.58	0.31	420	-82.09	0.36
10	2.50	0.45	1.84	0.34	419	-77.27	0.39
11	2.29	0.45	2.02	0.37	418	-75.92	0.42
12	2.41	0.52	2.33	0.44	417	-71.07	0.43
13	2.53	0.56	2.64	0.49	416	-67.28	0.47
14	2.50	0.59	2.88	0.53	415	-65.19	0.50
15	2.24	0.58	2.92	0.51	414	-65.46	0.51
16	1.60	0.43	2.69	0.42	413	-70.46	0.47

Notes: Coefficients and robust standard errors of Equation (4.2) estimated at different forecast horizons. The R^2 is the pseudo-R^2.

Table A.4: Estimation results of Equation (4.1) from 1 to 16 months excluding articles discussing the yield curve from the calculation of the index

Months ahead	β_k	*s.e.*	Obs.	Log Likelihood	R^2
1	1.72	0.30	428	-99.73	0.26
2	1.63	0.30	427	-101.01	0.25
3	1.59	0.30	426	-102.02	0.24
4	1.66	0.31	425	-99.68	0.26
5	1.67	0.32	424	-99.13	0.26
6	1.40	0.32	423	-98.72	0.27
7	1.56	0.30	422	-104.08	0.23
8	1.36	0.27	421	-109.40	0.19
9	1.22	0.24	420	-113.09	0.16
10	1.15	0.23	419	-114.62	0.15
11	1.00	0.21	418	-118.25	0.12
12	0.92	0.19	417	-120.03	0.10
13	0.84	0.19	416	-122.33	0.09
14	0.74	0.18	415	-124.26	0.07
15	0.63	0.16	414	-126.10	0.06
16	0.47	0.14	413	-128.78	0.04

Notes: Coefficients and robust standard errors of Equation (4.1) estimated at different forecast horizons. The R^2 is the pseudo-R^2.

Table A.5: Estimation results of Equation (4.2) from 1 to 16 months excluding articles discussing the yield curve from the calculation of the index.

Months ahead	β_k	$s.e.$	γ_k	$s.e.$	Obs.	Log Likelihood	R^2
1	1.72	0.30	-0.02	0.08	428	-99.69	0.26
2	1.61	0.30	0.06	0.08	427	-100.75	0.25
3	1.56	0.30	0.14	0.08	426	-100.75	0.25
4	1.62	0.32	0.21	0.08	425	-96.82	0.28
5	1.64	0.32	0.29	0.09	424	-93.52	0.31
6	1.75	0.32	0.41	0.10	423	-88.29	0.34
7	1.66	0.31	0.52	0.11	422	-88.02	0.35
8	1.53	0.28	0.66	0.13	421	-85.75	0.36
9	1.44	0.26	0.79	0.15	420	-82.42	0.39
10	1.46	0.26	0.93	0.17	419	-77.59	0.42
11	1.32	0.26	1.02	0.19	418	-76.26	0.43
12	1.35	0.27	1.19	0.22	417	-71.48	0.47
13	1.41	0.30	1.39	0.24	416	-67.30	0.50
14	1.35	0.31	1.52	0.26	415	-65.14	0.51
15	1.14	0.28	1.55	0.25	414	-65.83	0.51
16	0.82	0.24	1.46	0.21	413	-70.26	0.47

Notes: Coefficients and robust standard errors of Equation (4.2) estimated at different forecast horizons. The R^2 is the pseudo-R^2.

Table A.6: Estimation of Equation (4.1) and Equation (A.1) whith the QE-adjusted slope of the yield curve.

	Newspaper index					Yiled curve				
Months ahead	β_k	*s.e.*	Obs.	Log Likelihood	R^2	β_k	*s.e.*	Obs.	Log Likelihood	R^2
1	1.76	0.33	428	-100.89	0.25	-0.01	0.06	428	-134.57	0.00
2	1.58	0.33	427	-102.55	0.24	0.06	0.06	427	-133.59	0.01
3	1.50	0.32	426	-103.66	0.23	0.12	0.07	426	-131.92	0.02
4	1.56	0.33	425	-101.02	0.25	0.17	0.07	425	-129.665	0.04
5	1.49	0.32	424	-99.80	0.26	0.23	0.08	424	-126.32	0.06
6	1.35	0.29	423	-93.51	0.31	0.30	0.08	423	-121.77	0.10
7	1.21	0.27	422	-92.74	0.31	0.36	0.09	422	-115.94	0.14
8	1.06	0.24	421	-100.43	0.25	0.42	0.10	421	-102.68	0.24
9	0.94	0.22	420	-105.08	0.22	0.50	0.10	420	-96.35	0.28
10	0.91	0.21	419	-107.71	0.20	0.58	0.11	419	-90.23	0.33
11	0.79	0.19	418	-112.02	0.16	0.64	0.12	418	-83.22	0.38
12	0.74	0.18	417	-113.38	0.15	0.73	0.14	417	-79.72	0.41
13	0.72	0.18	416	-116.73	0.13	0.81	0.16	416	-76.50	0.43
14	0.64	0.17	415	-119.62	0.11	0.90	0.19	415	-73.34	0.45
15	0.53	0.15	414	-118.86	0.11	1.01	0.20	414	-70.38	0.47
16	0.39	0.14	413	-121.79	0.09	1.12	0.19	413	-67.95	0.49

Notes: Coefficients and robust standard of Equation (4.1) and Equation (A.1) estimated at different forecast horizons using the newspaper based index and the QE-adjusted slope of the yield curve (i.e. the difference between the 3-month rate and the 10-year rate on treasury bills) as explanatory variables. The R^2 is the pseudo-R^2.

Table A.7: Estimation results of Equation (4.2) from 1 to 16 months whith the QE-adjusted slope of the yield curve.

Months ahead	β_k	*s.e.*	γ_k	*s.e.*	Obs.	Log Likelihood	R^2
1	1.76	0.33	0.00	0.08	428	-100.79	0.25
2	1.59	0.33	0.09	0.08	427	-102.19	0.24
3	1.53	0.32	0.17	0.08	426	-102.36	0.24
4	1.62	0.33	0.24	0.08	425	-98.23	0.27
5	1.58	0.33	0.32	0.09	424	-94.79	0.30
6	1.50	0.31	0.40	0.10	423	-85.76	0.36
7	1.40	0.28	0.47	0.11	422	-82.00	0.39
8	1.28	0.25	0.54	0.12	421	-81.25	0.40
9	1.22	0.24	0.65	0.14	420	-77.81	0.42
10	1.27	0.24	0.78	0.16	419	-71.83	0.46
11	1.17	0.25	0.86	0.18	418	-67.64	0.50
12	1.21	0.26	1.01	0.21	417	-63.76	0.52
13	1.29	0.28	1.20	0.24	416	-59.52	0.56
14	1.25	0.29	1.34	0.26	415	-57.45	0.57
15	1.05	0.28	1.37	0.25	414	-58.01	0.57
16	0.77	0.24	1.32	0.23	413	-60.79	0.55

Notes: Coefficients and standard errors from the estimation of Equation (4.2) for different forecast horizons.

A.2 Figures

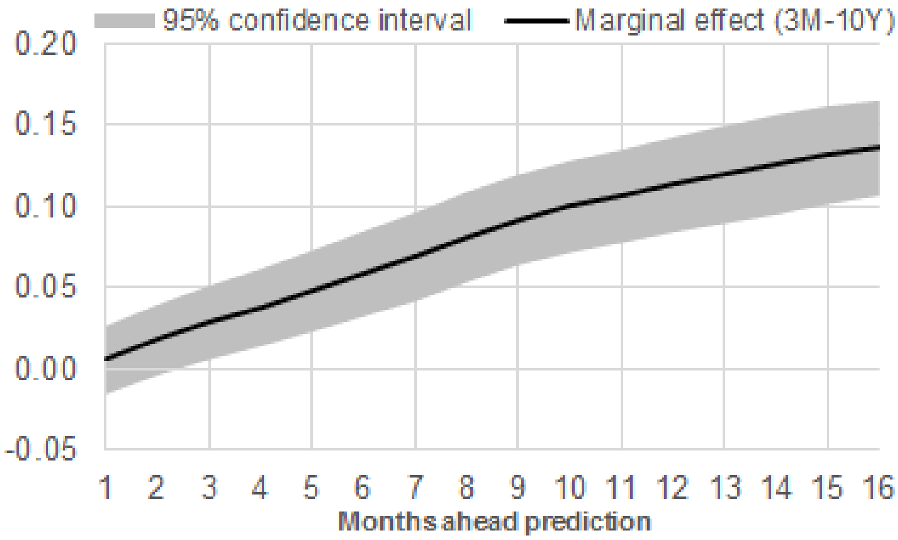

Figure A.1: Marginal effects from Equation (4.1) using the slope of the yield curve as explainatory variable.

Notes: marginal effects $\left(\frac{\partial P(Recession_{t+k}=1|t)}{\partial\left(r_t^{3-month}-r_t^{10-year}\right)}\right)$ from the probit regression for a 1% increase in the slope of the yield curve. The grey shaded areas report 95% confidence intervals.

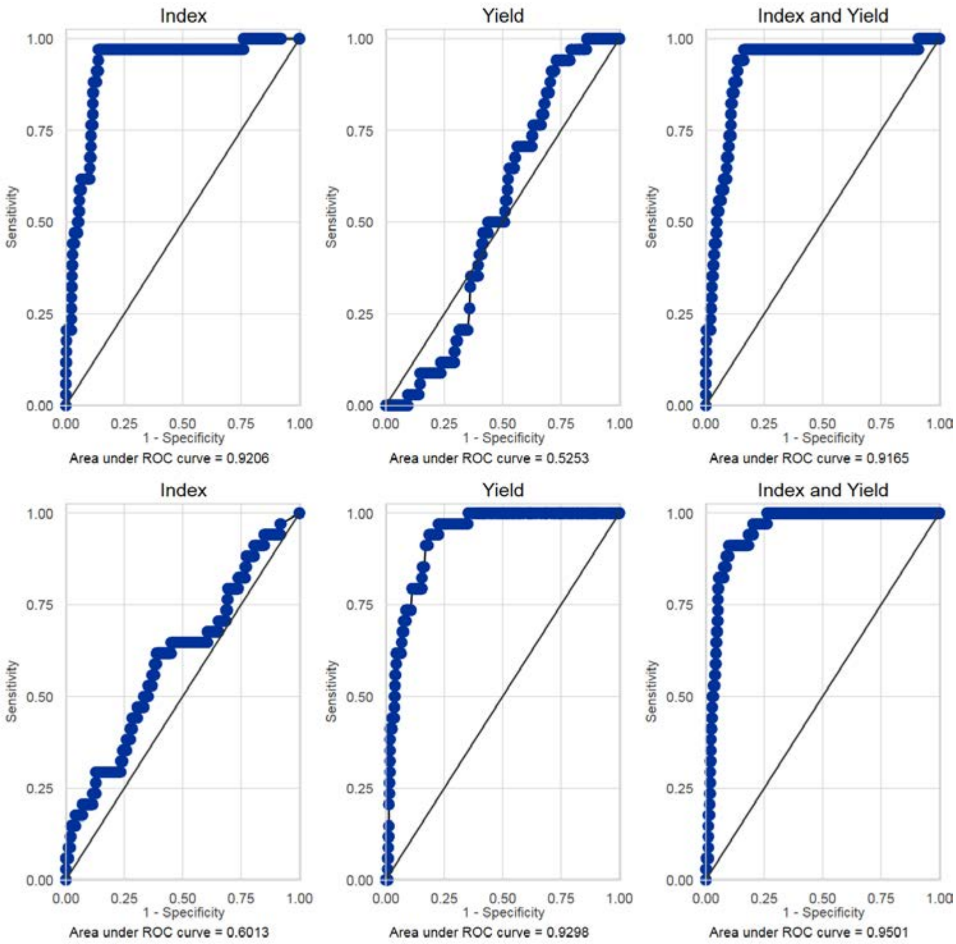

Figure A.2: Receiver Operator Curves for 1-month ahead predictions (upper panel) and 16-months ahead predictions (lower panel) of (i) the index only model Equation (4.1); (ii) the yield curve model Equation (A.1) and (iii) the index and yield curve model Equation (4.2).

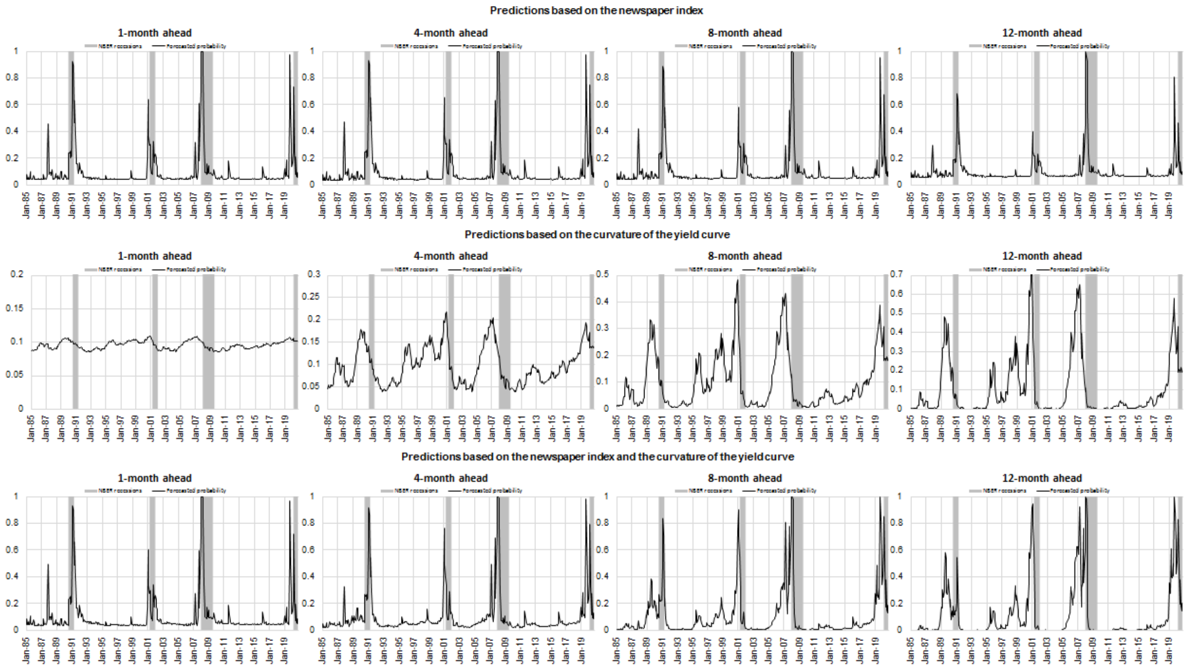

Figure A.3: Predicted recession probabilities at different forecast horizons based on the logit regression.

Notes: recession probabilities are computed based on Equation (4.1), for the index and the yield curve model Equation (A.1), and on Equation (4.2) for the model including both the yield curve and the index between regressors.

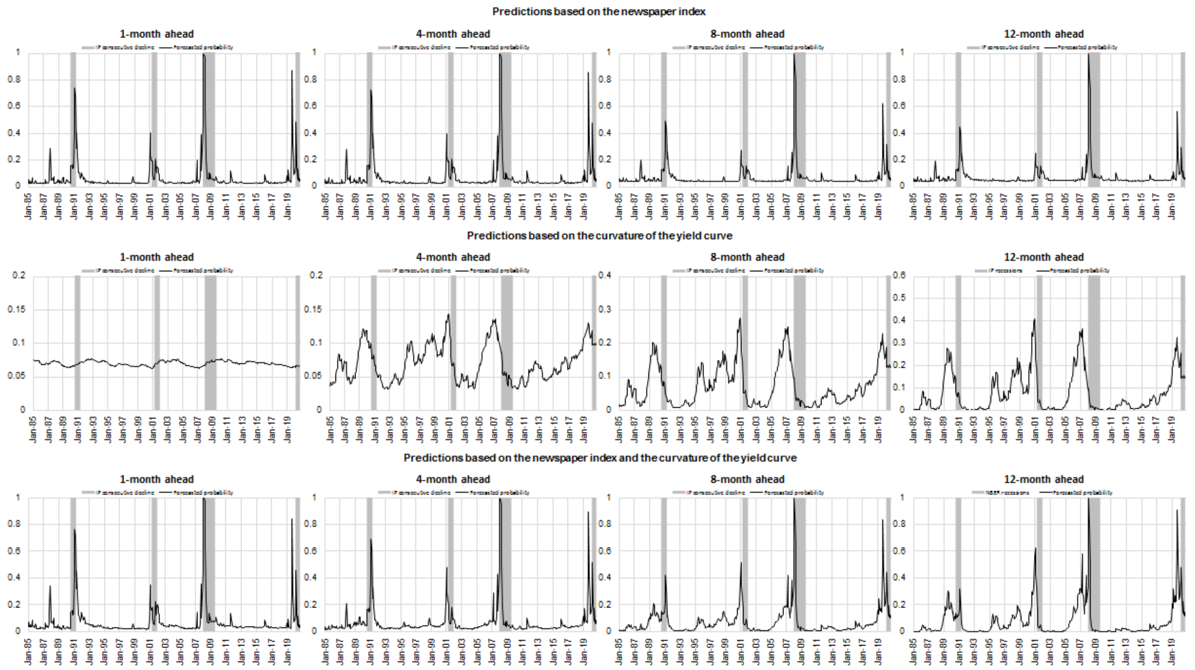

Figure A.4: Predicted recession probabilities at different forecast horizons when a recession is defined as 8 consecutive months of IP contraction.

Notes: recession probabilities are computed based on Equation (4.1), for the index and the yield curve model Equation (A.1), and on Equation (4.2) for the model including both the yield curve and the index between regressors.

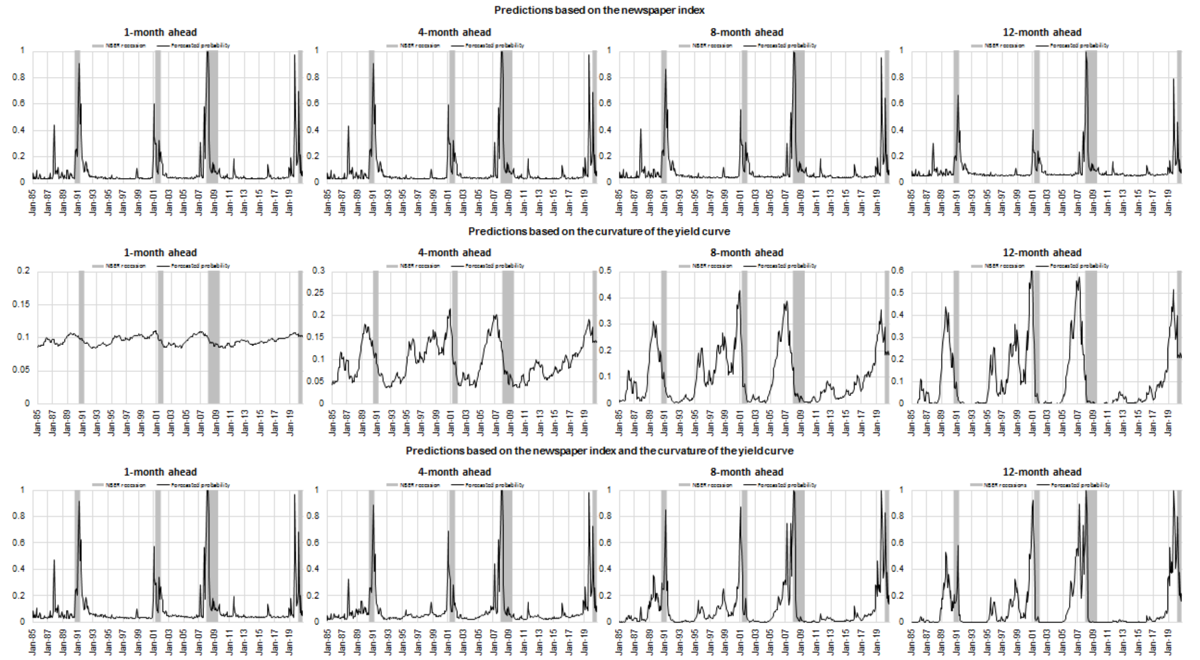

Figure A.5: Predicted recession probabilities at different forecast horizons based on probit regressions adjusting the yield curve for UMP periods.

Notes: recession probabilities are computed based on Equation (4.1), for the index and the yield curve model corrected by APP periods, and on Equation (4.2) for the model including both the yield curve and the index between regressors.

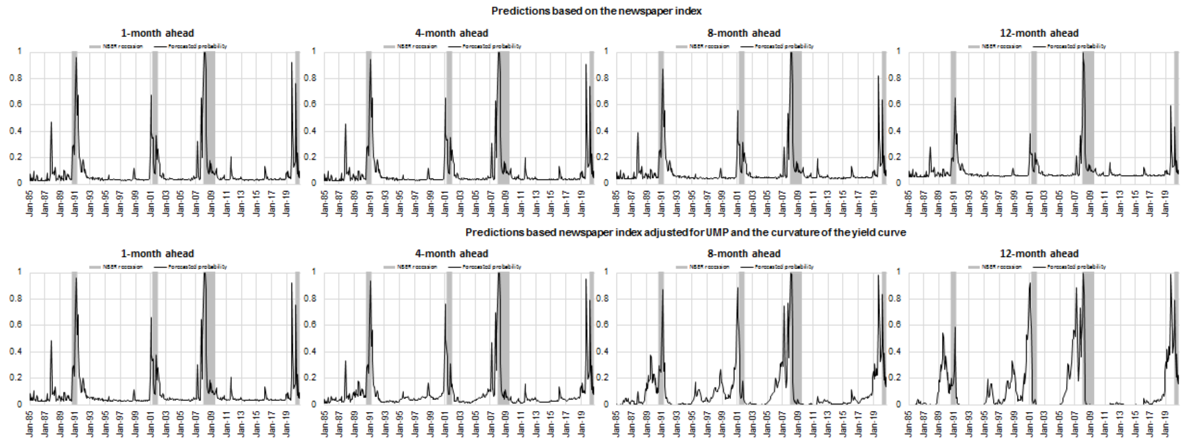

Figure A.6: Predicted recession probabilities at different forecast horizons based on probit regressions when correcting our index by UMP effects.

Notes: recession probabilities are computed based on Equation (4.1) when correcting the index by APP affects and on Equation (4.2) for the model including both the yield curve and the index corrected for APP effects between regressors.

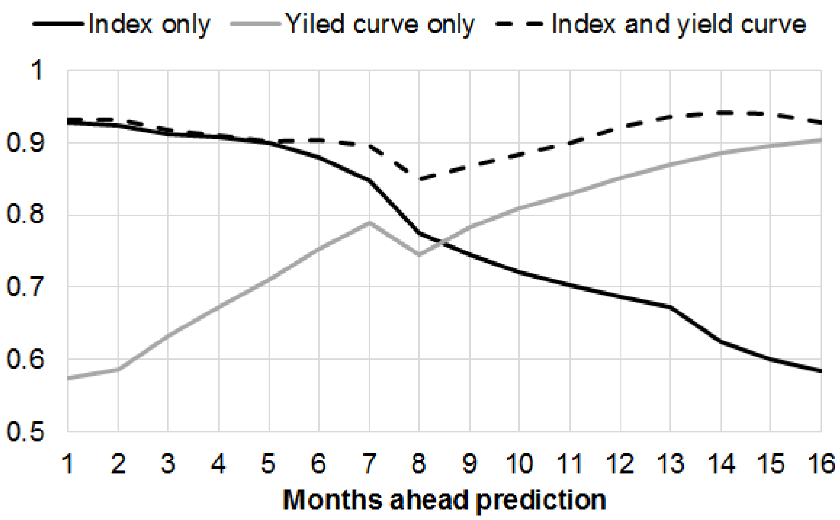

Figure A.7: ROC statistic for the models described by Equation (4.3) and Equation (4.4) at different forecast horizons estimated with a probit regression.

Notes: The ROC statistic is computed as the ratio between the area below the ROC line but above the 45 degree line and the area above the 45 degree line.

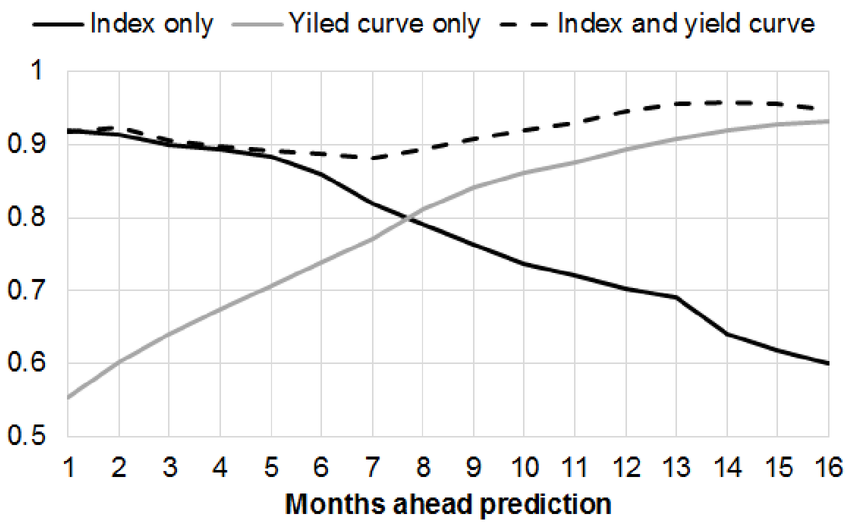

Figure A.8: ROC statistic for the models described by Equation (4.1) and Equation (4.2) at different forecast horizons estimated with a probit regression and QE-corrected yields.

Notes: The ROC statistic is computed as the ratio between the area below the ROC line but above the 45 degree line and the area above the 45 degree line.

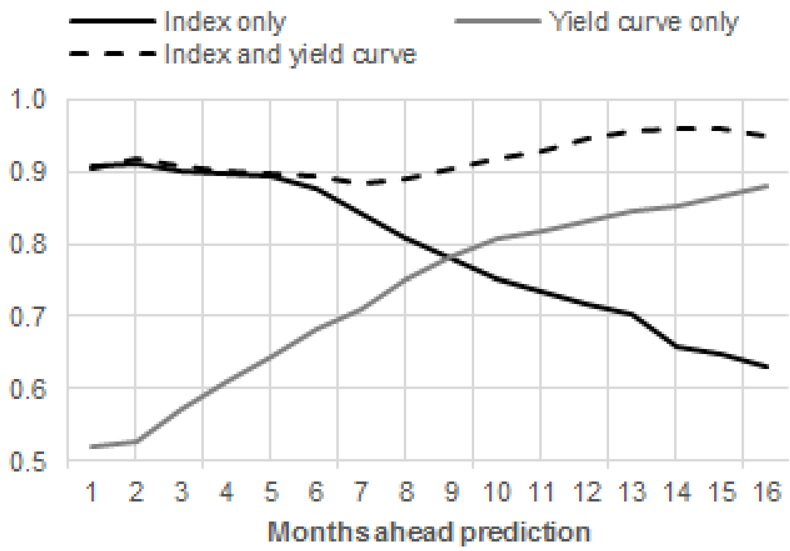

Figure A.9: ROC statistic for the models described by Equation (4.1) and Equation (4.2) at different forecast horizons estimated with a probit regression and when articles discussing the yield curve are excluded from the indicator.

Notes: The ROC statistic is computed as the ratio between the area below the ROC line but above the 45 degree line and the area above the 45 degree line.

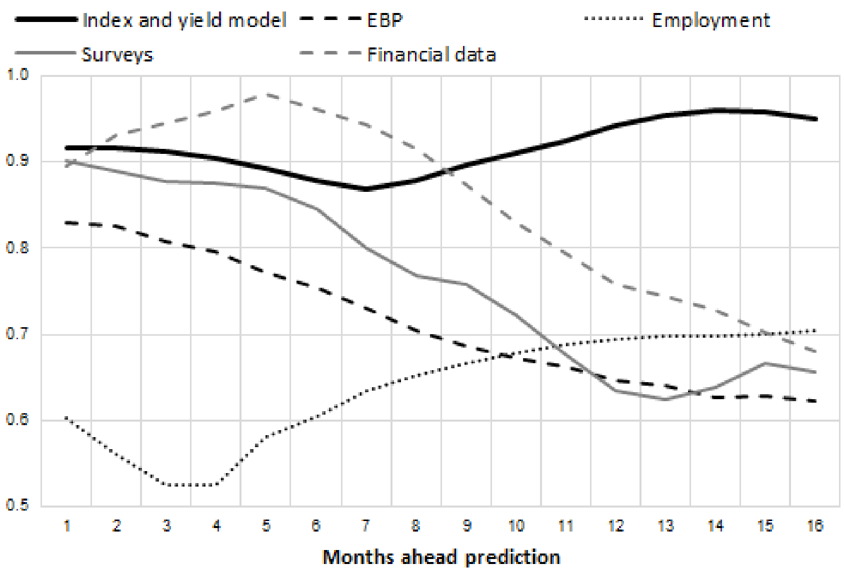

Figure A.10: ROC statistic for the models described by Equation (4.2) and Equation (4.5) at different forecast horizons estimated with a probit regression.

Notes: The ROC statistic is computed as the ratio between the area below the ROC line but above the 45 degree line and the area above the 45 degree line. Financial variables include the S&P500 and the 3-month Libor-TBill spread; employment variables include the unemployment rate and weekly hours worked; surveys include the PMI index and the Michigan consumers' confidence survey; EBP includes the excess bond premium and the terms spread by Gilchrist and Zakrajšek (2012).

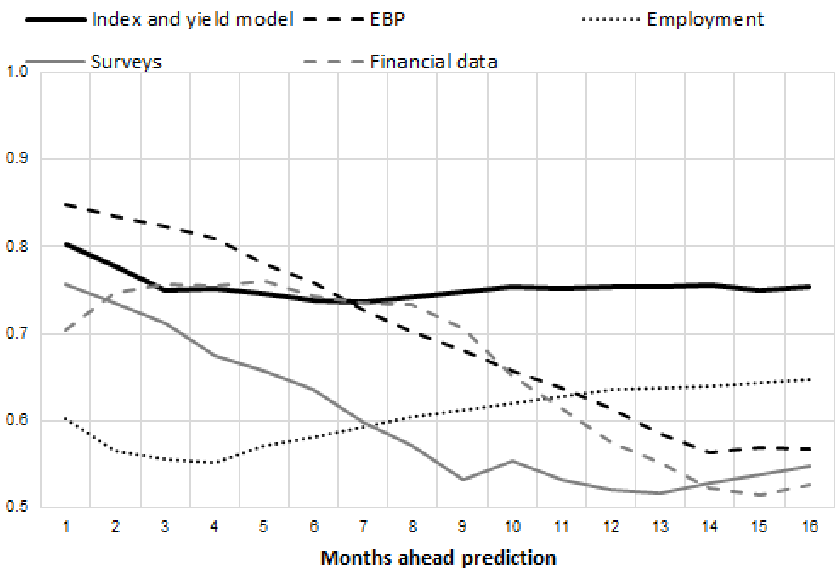

Figure A.11: ROC statistic for the models described by Equation (4.2) and Equation (4.5) at different forecast horizons estimated with a probit regression and using 4 consectuve months of IP contraction as definition of recession.

Notes: The ROC statistic is computed as the ratio between the area below the ROC line but above the 45 degree line and the area above the 45 degree line. Financial variables include the S&P500 and the 3-month Libor-TBill spread; employment variables include the unemployment rate and weekly hours worked; surveys include the PMI index and the Michigan consumers' confidence survey; EBP includes the excess bond premium and the terms spread by Gilchrist and Zakrajšek (2012).

Acknowledgements

We would like to thank M. Ca' Zorzi, J. Gräb, M. Jarocinski, M. Lenza, D. Lodge, A. Mehl, M.S. Pagliari, G. Rünstler, B. Schnatz, L. Stracca and participants to 2020 COMPTEXT conference as well as internal seminars for useful comments and discussions. The views expressed in the paper are those of the authors only and do not necessarily reflect those of the ECB or of the ESCB.

Massimo Ferrari
European Central Bank, Frankfurt am Main, Germany; email: massimo.ferrari1@ecb.europa.eu

Helena Le Mezo
European Central Bank, Frankfurt am Main, Germany; email: helena.le_mezo@ecb.europa.eu

Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.ecb.europa.eu

This paper can be downloaded without charge from www.ecb.europa.eu, from the Social Science Research Network electronic library or from RePEc: Research Papers in Economics. Information on all of the papers published in the ECB Working Paper Series can be found on the ECB's website.

PDF ISBN 978-92-899-4516-5 ISSN 1725-2806 doi:10.2866/864540 QB-AR-21-007-EN-N