Irish Fiscal Advisory Council

Working paper series

Estimating Ireland's Probability of a Recession

Eddie Casey and Niall Conroy

Working Paper No. 18
February 2023

Suggested reference:

Casey, E. and N. Conroy, (2023). "Estimating Ireland's Probability of a Recession". Irish Fiscal Advisory Council Working Paper Series No. 18. Dublin. Available at: www.fiscalcouncil.ie/working-papers/

Estimating Ireland's Probability of a Recession

Eddie Casey and Niall Conroy[1]

February, 2023

Abstract

This paper explores ways to predict future recessions in Ireland. We make three key contributions. First, we assess multiple potential indicators, including news articles and Google searches. Second, we show that a mixture of domestic and international variables, such as the US Term Spread, perform well when attempting to predict recessions in Ireland. This involves assessing the pseudo real-time historical and out-of-sample performance of models, and the false positive and false negative rates at different thresholds. Third, we identify a useful indicator for predicting recessions that to our knowledge has not been used elsewhere: the surveyed expectations of employment prospects.

Keywords: Cycles, forecasting and simulation, recession, term spread, probit model

JEL No. C25, C53 E32, E37

This report can be downloaded at www.FiscalCouncil.ie

1 The authors are, respectively, the Chief Economist, and an Economist at the Irish Fiscal Advisory Council. Email: admin@fiscalcouncil.ie. The opinions expressed and arguments employed in this paper do not necessarily reflect the official views of the Fiscal Council. We would like to acknowledge the feedback and kind assistance received from members of the Council and Secretariat of the Fiscal Council.

1. Introduction

A key question in macroeconomic forecasting is how likely is it that a downturn will soon hit the economy?

Standard macroeconomic forecasting tools often rely heavily on recent momentum in the economy to predict future growth. Most of the time, this is fine. Future growth won't be remarkably different from recent growth. On average, such forecasting models tend to benefit from taking account of recent momentum, and there will be a strong predictive power attached to how variables have behaved recently.

However, by relying on recent momentum, these tools can tend to miss major turning points. That is, at times when we may care more than usual about what is happening in the economy — for instance, on the precipice of a major downturn — our forecast models will often place too much weight on recent growth as a predictor of future growth. This can lead users to mistakenly assume growth will continue more or less uninterrupted even though key developments might be fostering the beginning of a downturn.

Reflecting the need to develop more sophisticated ways of identifying turning points, there is a rich strand of literature specifically assessing ways to predict downturns. The difference between short- and long-term interest rates, the term spread, has received a lot attention as having a strong link with recessions (Kessel, 1965; Fama, 1986; and Benzoni et al., 2018).

For the US, the term spread has been particularly useful for predicting recessions. Large declines have preceded every US recession since 1953, and there has been only one occasion when the spread turned negative without a subsequent recession (Wheelock and Wohar, 2009). Probit models that use the term spread as a predictor of future recessions tend to show significant outperformance relative to models that rely on other variables when predicting U.S. recessions (Estrella and Hardouvelis, 1991; Estrella and Mishkin, 1998).

For Europe, there is similar success from using the term spread. Bernard and Gerlach (1998) find the term spread useful for forecasting recessions up to two years ahead in countries such as Belgium, France, Germany, the Netherlands, and the UK, while Moneta (2005) finds it useful for the euro area as a whole.

However, the (domestic) term spread has been found to be less useful as a predictor of recessions for Ireland. Stuart (2020) finds that Ireland's own term spread has no predictive power for recessions in Ireland over the full period 1972–2018 using monthly data. It only has weak significance during the Exchange Rate Mechanism period of 1979–1989 when the Irish Central Bank had some limited discretion over monetary policy, given that the currency was allowed float within a defined band relative to other currencies.[2]

Given the weakness of the domestic term spread, we might usefully consider using the US equivalent as a predictor of recessions in Ireland. Along these lines, Bernard and Gerlach (1998) find the US term spread useful for forecasting recessions in the UK – an understandable outcome given their trade linkages. However, for most countries they find that the domestic term spread is the best indicator, rather than US or German spreads.

Building on the international and domestic research, we look at ways to predict downturns in Ireland. In line with findings elsewhere, and given Ireland's openness and the important direct and indirect trade links between the US and Ireland, we assume that movements in the US term spread can give a useful insight into how the Irish economy will perform.[3] In effect, we are assuming that its capacity to signal downturns in the US economy will mean that it is also a useful indicator of recessions in Ireland. This follows from the view that a US recession would tend to herald subsequent downturns in Ireland and in other economies Ireland trades with.

We also consider additional predictors. A useful set of potential predictors is shown in OECD (2012). They highlight some of the desirable characteristics of any predictors considered. First, series should be economically relevant and cover a broad range of activity. Second, high frequency (say, for example monthly rather than quarterly or annual) series are preferred. Third, the series should ideally be released in a timely fashion. Fourth, the series should not be subject to subsequent revisions. This leads us to consider a number of variables such as the main

2 The less informative role for the term spread in small, open economies is partly judged to be due to the fact that central banks tend to fix, or manage heavily, their exchange rates, such that the domestic term spread only weakly reflects domestic economic conditions. Stuart also finds that the UK or German spread is not a significant predictor of recessions in Ireland.

3 As an alternative to the term spread, we could consider the federal funds rate also or a corporate credit spread for the US. Wright (2006) finds that including the federal funds rate as well as the term spread provides superior out-of-sample recession forecasts. King, Levin, and Perli (2007) find that including the corporate credit spread yields similar benefits.

economic indicators produced for Ireland's main trading partners by the OECD — composite measures of economic performance.

One useful indicator we assess is a survey-based measure of business expectations around future employment within various sectors and of consumer expectations for future unemployment available from the European Commission. This is available for the construction and industry sectors as well as for consumers from 1985, with services and retail being added from 1998 on.

Another useful set of predictors can be gleamed from Google searches and news articles.

Tkacz (2013) explores the early potential for the Google search terms such as "recession" and "jobs" to predict US recessions. Although the data are only available from 2004, they find some promise for their use as predictors, with movements in Google Trends data correlated with the 2008/2009 recession up to three months in advance.

In a similar vein, Minesso, Lebastard and Le Mezo (2022) consider a measure of US newspaper articles on economics that contain the words "recession" or "slowdown". Constructing an index based on these shares, they show a strong positive and statistically significant relation between the newspaper index and future recessions at various time horizons, but most strongly at the 8-month horizon.

Our contributions to the literature are three-fold. First, we assess the merit of multiple potential predictors of recessions for Ireland, including recent innovations that rely on textual analysis. Second, we show that a mixture of domestic and international predictors, including the US Term Spread, perform well when attempting to predict recessions in Ireland. This involves assessing the pseudo real-time historical and out-of-sample performance of models, and the false positive and false negative rates at different thresholds. Third, we identify a potentially very useful variable for predicting recessions that — to our knowledge — has not been used elsewhere. That is the surveyed expectations of businesses and consumers about future employment prospects.

2. Data and methodology

As in other literature, we define a recession as a change in the unemployment rate equal to or exceeding its 80th percentile. For Ireland, and over the sample period we assess, this equates to a rise in the unemployment rate over 12 months of at least 0.8 percentage points. We then define recessions as a binary variable (equal to one or zero) based on this variable, which we attempt to predict.

The focus on unemployment is useful considering the distortions to many of Ireland's aggregate national accounts measures arising from foreign-owned multinational enterprises. The approach is somewhat arbitrary, but it is a convention in the literature and nonetheless captures big increases in unemployment rates. [4]

One concern in using unemployment rates in an Irish context relates to the role of migration. Throughout Irelands economic history, economic downturns have been accompanied by significant outward migration (FitzGerald and Kearney, 1999). As a result, an increase in unemployment may be more modest than would otherwise be the case. As a robustness check, (shown in Annex D), we also consider using quarterly real modified domestic demand to define recession periods. We find no significant impact on results.

To identify suitable predictors of a recession, we draw on the characteristics recommended by the OECD (2012). The following characteristics are identified:

- **Higher frequency** indicators are preferred (example: monthly indicators are preferred to quarterly indicators)
- **Revisions should be avoided**: series that are not subject to significant revisions are preferred as more reliable and stable indicators
- **More timely indicators** are preferred: data should ideally be available very soon after the period to which they refer
- **Longer time series** with no breaks in the series are preferred

4 As a robustness check (shown in Annex D), we also consider using quarterly real modified domestic demand to define recession periods. We find no significant impact on results.

- **Economic justification is needed**: an economic justification for the relationship is needed

- **Broader coverage is better**: a series with a broader coverage of economic activity is preferable to a narrowly-defined series, with limited importance for the wider economy

With these characteristics in mind, the OECD (2012) also identify four types of indicators that may satisfy these conditions:

1) **Early-stage indicators**: indicators that measure the early stages of production, such as new orders, order books, and housing starts.

2) **Rapidly responsive indicators**: indicators that respond rapidly to changes in economic activity, such as profits and inventories.

3) **Expectation-sensitive indicators**: these are indicators that measure or are sensitive to expectations, such as stock prices, prices of raw materials, and business survey expectations related to production or the general economic climate.

4) **So-called "Prime Movers":** indicators relating to monetary policy and foreign economic developments such as money supply, terms of trade, and external demand.

Table 1 summarises the series we consider based on the desired characteristics set out above.

Table 1: Summary of the key predictors we assess

Indicator	Expected sign
Domestic indicators	
1) Early-stage indicators (PMI new orders, new business, housing commencements)	(–)
2) Employment expectations indicators (DG ECFIN sectoral surveys)	(–)
3) Input price indicators (PMI input prices)	(+)*
4) Expectation-sensitive indicators (ISEQ index, PMI business expectations)	(–)
5) Google searches	(+)
6) News articles (Irish Times)	(+)
External indicators	
7) Prime Movers: TED spread	(+)
8) Prime Movers: Term spread[5]	(–)
9) External demand indicators (OECD composite leading indicators: EA, US, UK)	(–)

Notes: * There may be some ambiguity around the expected sign on input prices. For a recession triggered by a negative supply shock, one would expect input prices to rise. By contrast, for a recession that is driven by falling demand, one might expect input prices to fall at or before the onset of the recession.

A key set of time series we consider is the employment expectations indicators.[6] These are based on regular surveys conducted by the European Commission of different sectors and of consumers. These are particularly useful as they give a timely sense of expectations around employment — something which should have a strong link to future changes in unemployment. They are typically released in the final days of the month to which they refer.

In addition, we also consider Google Trends data on searches for the terms "recession" and "unemployment" located in Ireland. Separately construct an index of the share of articles in the Irish Times covering the term "unemployment".

All series selected are available at a monthly or higher frequency, with results available relatively quickly — typically within a few days of the month ending. Exceptions are the OECD's Composite Leading Indicator, which tends to be published between 1–2 weeks after the reference month ends, and the housing

5 As a robustness check, we also repeated the analysis using the euro area term spread instead. We found similar results, however as the US term spread was more consistently significant, we use the US term spread in our final specifications.

6 The Commission's "Business and Consumer Surveys" are regular surveys conducted by the Directorate General for Economic and Financial Affairs for different sectors. The questions are of the form "how do you expect your firm's total employment to change over the next 3 months?". Responses are either "Increase", "Remain unchanged", or "Decrease". In the case of consumers, the question is "how do you expect the number of people unemployed in this country to change over the next 12 months?" Responses range from "Increase sharply" to "fall sharply". In each case, the responses are published as balances, based on the differences between positive and negative answers (in percentage points of total answers).

commencements indicator from the Department of Housing, Local Government and Heritage, which can take 3 weeks or more to be published.

We use seasonally adjusted series where there is an obvious need to do so. For instance, we use it for the survey indicators of employment expectations and composite leading indicators. However, we do not seasonally adjust other variables, such as the PMI indicators of new business, business expectations or new orders as there is no discernible seasonal pattern.

In the case of some categories of indicators where there are a large number of time series, we use the principal components method as a dimension-reducing technique. This allows us to identify a single time series that draws on the common features within specific categories of indicators we assess. In each case, we take the first principal component as an estimate of the common factors in the data.[7]

The initial dataset we consider spans the months between January 1998 and December 2022, matching the availability of the employment expectations data and the official monthly unemployment data. However, we also consider an extended sample period (starting in 1985) for our preferred estimation. Annex A sets out more detail on all the series we consider as predictors.

Estimation technique

To estimate the likelihood of a recession in the next *k* months (for example, 12 months), we use a probit model. This can be described as:

$$P(Recession_{t+k}) = \Phi[\alpha_k + \beta_{ki} X_{t_i}] \quad (1)$$

where $Recession_{t+k}$ is our k-month-ahead recession indicator based on unemployment rate changes, X_{t_i} is our vector of predictors, and Φ means that we are taking the standard normal cumulative distribution function of our variables.

7 The principal components are linear combinations of our original series weighted by their contribution to explaining the variance in a particular orthogonal dimension. The objective is dimension reduction. We basically wish to create one principal components from our larger number of initial variables, which share some characteristics in common. The first principal component accounts for as much of the variability in the initial time series as possible. This is applied to the categories 1–4 and category 9, the indicators of external demand.

The coefficients from a standard probit estimation would be unbiased, but the serial correlation in the dependent variable will cause the standard errors to be too small. We therefore report errors obtained after applying the Newey-West correction.[8]

While some previous studies have included current recession status ($Recession_t$) as a predictor of future recession probability ($P(Recession_{t+k})$), we have not done so. In practical terms, an early warning or predictor equation such as this would be used before the onset of recession is apparent. As a result, we feel it is best to not include current recession status when estimating such probit models.

---

8 Another approach could be to estimate a probit model with Markov-switching coefficient variation and a lagged dependent variable, which would allow for non-linearities (Dueker, 1997).

3. Results

Now we turn to the results of our initial estimation which is to assess the full set of predictors in the nine categories which were set out in Table 1.

We initially estimate the model for the period January 2005 to December 2022, with our recessions being predicted at a 3-, 6-, 9-, 12 and 15-month-ahead horizon. This sample period is chosen as it is a consistent period over which we have all indicators.[9] It is primarily limited by the availability of our Google Search term measure. Model 1 in Table 2 shows our initial estimation with all variables considered for a 15-month-ahead prediction horizon.

Table 2: Model selection at 15-month horizon

	Preferred model	Model 2	Model 1
Term Spread (-15)	-0.47	-0.52	-0.59
	(0.21)**	(0.28)*	(0.27)**
TED Spread (-15)	1.26	1.37	1.36
	(0.55)**	(0.59)**	(0.72)*
Employment Expectations (-15)	-0.48	-0.57	-0.41
	(0.15)***	(0.2)***	(0.17)**
New Activity (-15)		0.12	-0.26
		(0.24)	(0.21)
ISEQ (-15)		0.00	0.00
		(0)	(0)
Input Prices (-15)			-0.56
			(0.21)***
Main Economic Indicators (-15)			0.17
			(0.17)
Search Terms (-15)			-0.38
			(0.16)**
News (-15)			-0.30
			(0.11)***
Constant	-0.67	-0.54	0.22
	(0.52)	(1.41)	(1.5)
BIC	0.74	0.78	0.80
Pseudo R-Squared	0.41	0.41	0.51
Observations	224.00	224	212
Sample	2004M04 2022M11	2004M04 2022M11	2005M04 2022M11

Notes: Robust standard errors (Newey-West correction) are estimated. *, ** and *** denote statistical significance at a 10% ,5% and 1% level, respectively. The BIC is the Bayesian Information Criterion.

9 As a robustness check, we estimate the models excluding the Covid induced recession starting in March 2020. Given this recession was inherently unforecastable as it was caused by a global pandemic, it may be considered very different to other recessions experienced. In any event, we find no significant differences when excluding this recession.

Starting in Model 1, we find that the following four variables have the wrong sign on their coefficients: News, Google Search Term, Input Price, and External Demand measures. As a result, we drop these variables. The model is then re-estimated without these variables in Model 2. Looking at the results in model 2, the coefficients on the New Activity measure and the ISEQ index are found to be statistically insignificant. As a result, these two variables are dropped for our final, preferred specification.

We are left with three variables in a preferred specification for the 15-month horizon: the Term Spread, the TED Spread, and the Employment Expectations measure. They all have the expected sign and are statistically significant at least at the 5% level. Annex B shows the same sequential approach we take to selecting indicators in a preferred model for other prediction horizons. Following this approach leads us to also keep the New Activity measure as it proves significant at shorter prediction horizons (at the 9-month horizon).

Figure 1 shows each of the three indicators from our preferred specification for the 15-month horizon along with recession bands. The latter are based on our recession definition, which draws on the change in the unemployment rate relative to 12 months ago.

We can see that each of the variables in this preferred specification appears to have reasonable success in identifying each of the 1991, 2002 (Dot-Com Bubble) and 2008 (Financial Crisis) recessions. For the 2020 recession, the measures coincide with the initial March 2020 collapse in activity and employment, but they do not identify it in advance.[10]

10 An exception is the term spread, which deteriorated in 2019 to the extent that the yield curve had actually inverted. However, these were due to economic factors that were outside the bounds of what could explain the non-economic factors behind the pandemic-induced recession of 2020.

Figure 1: Key recession predictors following model selection

TED Spread

Term Spread

Labour Sentiment

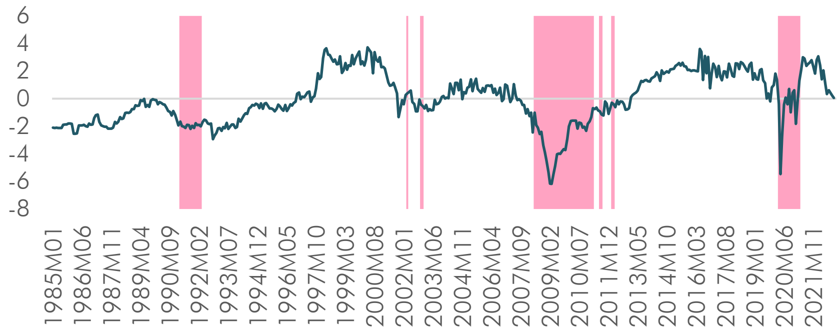

New Activity

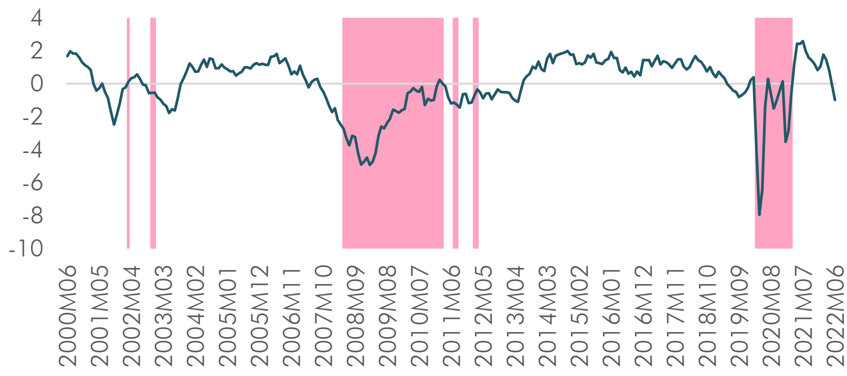

Notes: Recession bands in pink are shown for our main definition: a change in the unemployment rate greater than its 20th percentile, which, for the sample period considered, equates to a 0.8 percentage point increase vs 12 months ago.

Formal tests of predictive power

The formal tests of the predictive power of our key indicators are next assessed in a recession probability model that includes all four indicators over different prediction horizons. This corresponds to estimating equation (1), with k = 1 to 15 month horizons. More specifically, it is equivalent to estimating:

$$P(Recession_t) = \Phi[\alpha + \beta_1 * TED\ spread_{t-k} + \beta_2 * Term\ spread_{t-k} + \beta_3 * Labour\ confidence_{t-k} + \beta_{4*} New\ activity_{t-k}] \quad (2)$$

We start with a one-month-ahead prediction of our recession indicator and gradually increase the horizon by one month at a time to a 15-month-ahead prediction. The estimated coefficients and standard errors are shown for each indicator and at each prediction horizon in Table 3.

Table 3: Estimation results for our preferred model from 1 to 15 months ahead

	TED Spread		Term Spread		Labour sentiment		New activity				
Months ahead	β	S.E.	β	S.E.	β	S.E.	β	S.E.	R^2	AIC	BIC
1	-0.83	0.45*	-0.14	0.17	-0.88	0.22***	-0.42	0.18**	0.58	0.47	0.53
2	-0.63	0.40	-0.22	0.16	-0.88	0.24***	-0.36	0.18**	0.56	0.48	0.55
3	-0.26	0.39	-0.18	0.15	-0.78	0.22***	-0.25	0.18	0.52	0.53	0.60
4	-0.09	0.41	-0.17	0.15	-0.69	0.21***	-0.28	0.18	0.50	0.55	0.61
5	0.02	0.43	-0.23	0.17	-0.75	0.21***	-0.27	0.18	0.52	0.53	0.60
6	0.41	0.49	-0.18	0.19	-0.61	0.18***	-0.30	0.17*	0.50	0.55	0.62
7	0.53	0.51	-0.19	0.2	-0.57	0.18***	-0.29	0.17*	0.49	0.56	0.63
8	0.86	0.53	-0.17	0.21	-0.47	0.16***	-0.36	0.18**	0.49	0.56	0.63
9	1.27	0.63**	-0.19	0.22	-0.43	0.15***	-0.38	0.18**	0.50	0.56	0.62
10	1.17	0.63*	-0.26	0.20	-0.41	0.14***	-0.39	0.19**	0.48	0.58	0.64
11	1.27	0.63**	-0.25	0.18	-0.42	0.15***	-0.22	0.18	0.42	0.63	0.70
12	1.52	0.66**	-0.24	0.18	-0.38	0.15**	-0.08	0.20	0.38	0.68	0.75
13	1.05	0.59*	-0.39	0.17**	-0.50	0.16***	0.01	0.22	0.37	0.69	0.76
14	1.20	0.60**	-0.43	0.17**	-0.50	0.16***	0.01	0.22	0.38	0.68	0.75
15	1.16	0.58**	-0.49	0.18***	-0.52	0.18***	0.04	0.23	0.39	0.68	0.74

Notes: The coefficients are reported as marginal effects for a 1% increase in the predictor. Robust standard errors (Newey-West correction) are estimated. Both are shown for different forecast horizons. *, ** and *** denote statistical significance at a 10% ,5% and 1% level, respectively. The R^2 is the pseudo-R^2. The AIC and BIC are the Akaike and the Bayesian information criteria, respectively. Equations are estimated over the sample period 2001M9 to 2022M12.

We can see from Table 3 that the indicators are useful as predictors of recessions, but with each proving useful at different horizons. The labour sentiment index proves to be a useful indicator at all horizons considered. It achieves statistical significance at the 1% level for all but one of the forecasting horizons considered.

The TED Spread and the Term Spread are both statistically significant at longer horizons. Indeed, the TED spread only has the correct sign and statistical significance for a nine-month-ahead horizon up to the 15-months-ahead we assess. Similarly, the Term Spread only achieves statistical significance when a horizon of 13 months or more is considered.

By contrast, the New Activity measure is primarily significant at relatively shorter forecasting horizons, proving significant at the 5% level for horizons of 1 to 2 months ahead and 8 to 10 months ahead.

In terms of the overall fit of the models estimated, there is evidence that the models produce a reasonably good in-sample fit at all horizons considered but that they appear to lose their power over longer horizons. The pseudo R-squared ranges from 0.58 at the one-month horizon to 0.37 at the 13-month horizon. Similarly, the Bayesian information criterion (BIC) suggests that the models appear to lose their power over longer horizons (see Annex C for charts of AIC, BIC, pseudo-R-squared and Root Mean Squared Errors).

We next show the marginal effects over different prediction horizons for each of the key variables left in our preferred model in Figure 2. These charts can be read as, taking the first panel for example, a 1% deterioration in employment expectations leads to a 12% higher likelihood of a recession 4 months ahead, an around 8% higher likelihood 9 months ahead, and a 5% higher likelihood over a 15-month horizon. The employment expectations indicator is statistically significant at all forecast horizons.

Figure 2: Marginal effects of key indicators

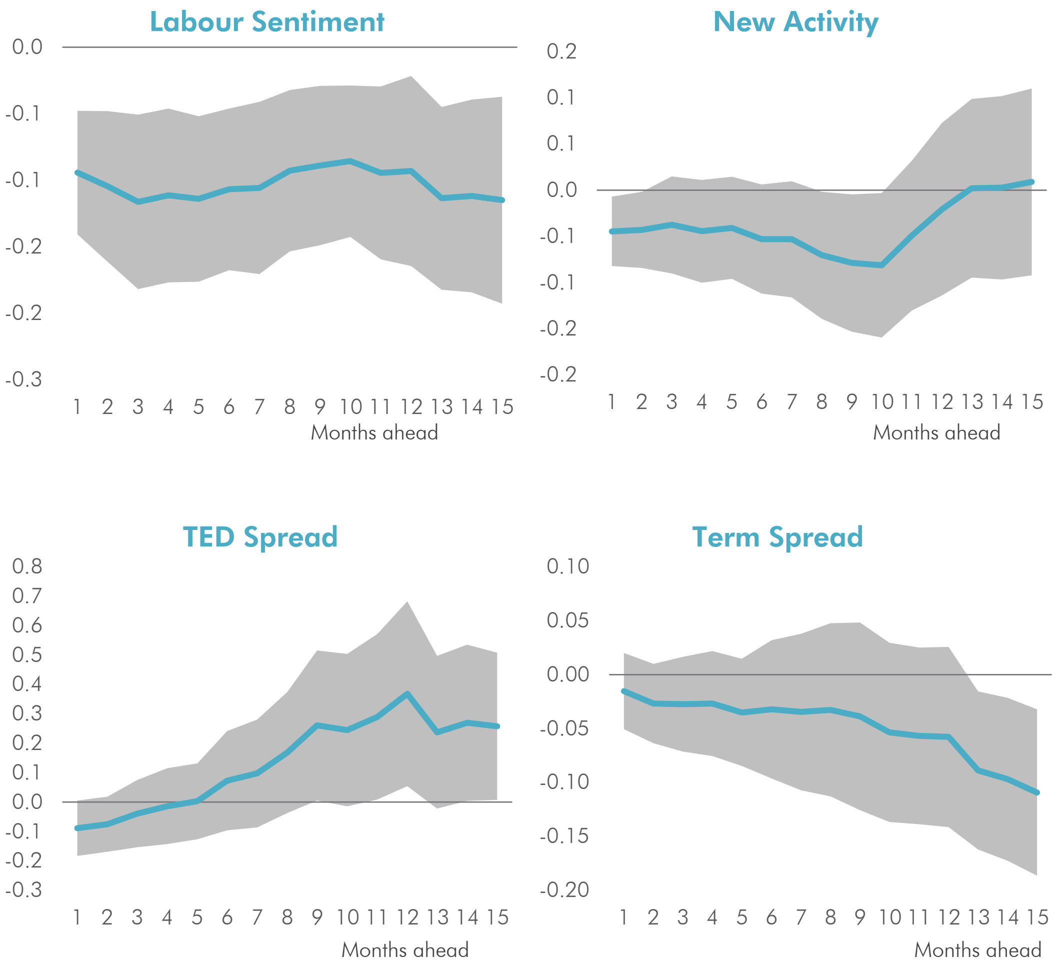

Notes: Marginal effects from the probit regression for a 1% increase in each predictor variable (for example, a 1% increase in the new activity index). Grey shaded areas report 95% confidence intervals. The estimation results are for models at various prediction horizons that include all four predictor variables.

Measuring the model's goodness of fit

A question worth asking at this stage is whether or not we should retain all four variables in a final model. With this in mind, a useful measure of fit for prediction models is the "Receiver Operator Characteristic" (ROC). This was originally used for assessing how well radar operators correctly identified incoming aircraft during World War II.

The ROC can be shown as a curve. It plots the true positive rate against the corresponding false positive rate for a given threshold. That is, if we were to arbitrarily set our threshold for identifying recessions as likely to occur 12 months out as being triggered by any likelihood prediction greater than 15%, then we can calculate the rate at which we get true positives (successful hits) and false positives (false alarms) for this chosen threshold.[11] The ROC curve shows this point but also all of the other corresponding true positive rates and false positive rates for any possible threshold that we might consider.

An ROC statistic, also called the "Area Under the Curve" or AUC can also be calculated as a useful single summary measure. This summary measure is the ratio of the area below the ROC curve, but above the 45 degree line, and the total area above the 45 degree line (see Figure 3A). A larger area suggests that the model performs significantly better than chance. The ROC statistic helps us assess the accuracy of a certain prediction model since better models would have a larger difference between the two areas. The ROC statistic is bounded between 0 and 1, with zero signifying a model no better than coin flip and one signifying the perfect model with flawless predictions.

Another important aspect is to consider threshold probabilities. This would be the probability if exceeded by the probit model, we would consider as signalling a likely recession. There is a trade-off here between true positives and false positives. Setting the threshold to a high level will minimize false positives, however this would lead to a low number of true positives (see the bottom left section of Figure 3A). Setting a lower threshold would lead to more true positives, but also more false positives (moving up and to the right in Figure 3A).

One approach to this trade off could be to maximise the ratio of true positives to false positives. In our case, a threshold of 19% would maximise this ratio. If one wanted to penalise false positives more, then one would set as higher threshold. Conversely, if wanted to reward true positives more, one would set a lower threshold.

11 The True Positive Rate (TPR) can be defined as $TPR = \frac{TP}{(TP+FN)}$ where TP is the total number of true positives we get for a given threshold and FN is the total number of false negatives we get. The False Positive Rate (FPR) can be defined as $FPR = \frac{FP}{(FP+TN)}$ where FP is the total number of false positives we get for a given threshold and TN is the total number of true negatives.

Figure 3: ROC Curve for 15-month-ahead prediction model

A. Including all four key indicators

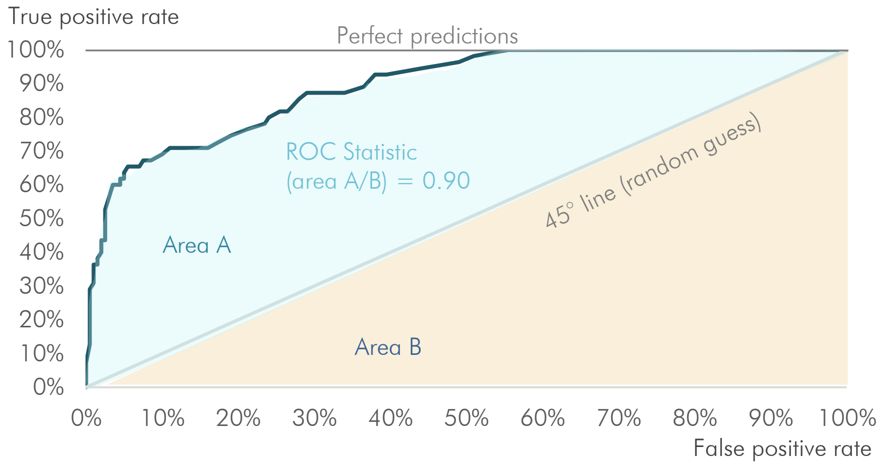

B. Assessing various combinations of our four key indicators

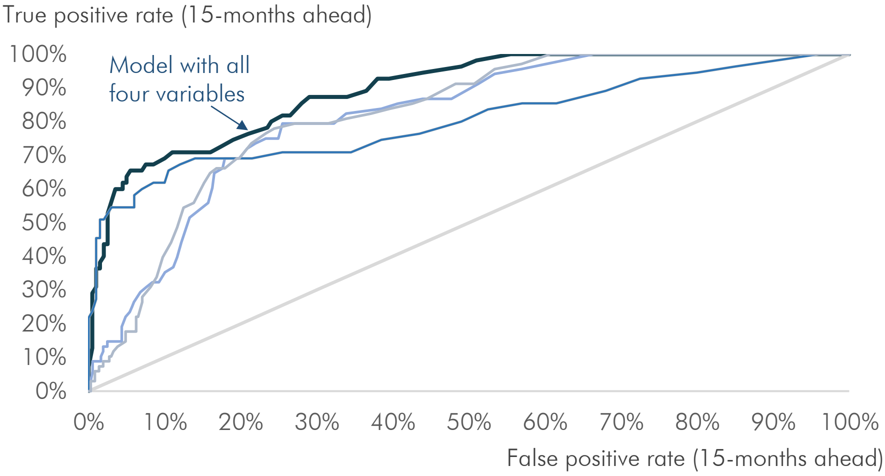

C. Assessing various combinations at different horizons

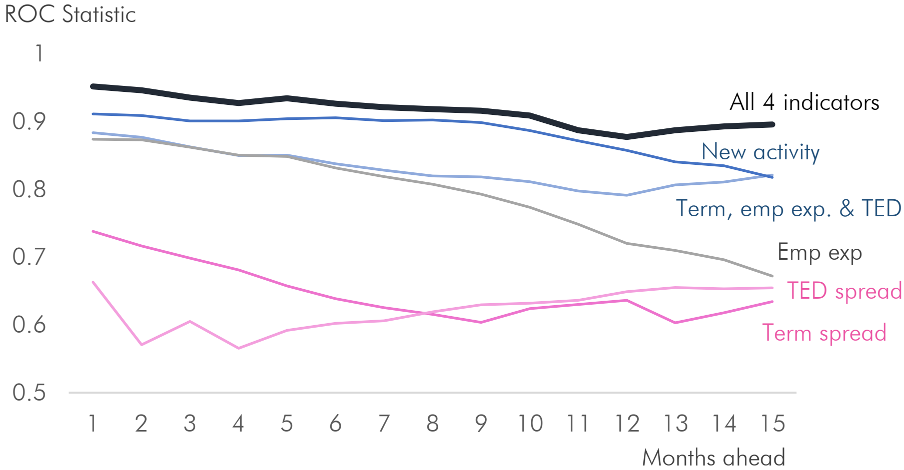

Notes: Panel B compares the model with all four predictors to other combinations using less than four predictors. In all cases it is 15 month ahead forecasts that are being assessed. Panel C shows different ROC statistics over the forecast horizons when using different combinations of variables as predictors. The model with all four predictor variables produces the highest ROC statistic at all forecast horizons.

We can use the ROC statistics to assess what combinations of our key variables would provide a better fit. Figure 3B shows the model specification with all four variables outperforming other combinations of models in this respect. We can see that the model with all 4 variables performs better than other alternatives where one or more of these variables are dropped.

We can also extend this analysis to compute just the ROC statistics for various models that include our four key indicators and for all lag lengths up to 15-months ahead. Figure 3C shows the results of this exercise. Again, the model with all four indicators outperforms other combinations at all forecast horizons. As expected, the ROC statistics is generally stronger (closer to one) at shorter prediction horizons. However, its fit appears to remain reasonably consistent for the 11 to 15 month horizon. We can see that the inclusion of the Term Spread and the TED Spread counteracts some of the loss in performance attributable to the other variables at longer horizons.

Examining an optimal lag length for each indicator

Before we turn to an assessment of the out-of-sample performance of our recession forecasting model, we can consider if it would be worth using different lag lengths for each of our individual indicators. Using the Aikaike Information Criterion (AIC), we can identify suitable lag lengths for each variable. Table 4 shows the AIC for each prediction horizon from 1 to 15 months. It also shows the estimated coefficients as, in the case of the Term Spread, the sign can switch direction and hence be dropped from our model selection.

Table 4 suggests that when producing forecasts of recessions, we should ideally use the most recently available data for employment expectations and new activity. By contrast, we should use older data for the TED Spread and the Term Spread, notably information from 12 and 15 months previously, respectively.

Table 4: Optimal lag lengths for models with individual variables

	TED Spread		Term Spread		Employment expectations		New Activity	
lag	Β	AIC	β	AIC	β	AIC	β	AIC
1	0.12	0.86	0.41	0.78	-0.60	0.59	-0.85	0.58
2	0.22	0.86	0.35	0.80	-0.58	0.59	-0.83	0.59
3	0.32	0.86	0.31	0.81	-0.54	0.61	-0.75	0.62
4	0.41	0.85	0.28	0.82	-0.52	0.63	-0.76	0.62
5	0.48	0.85	0.23	0.83	-0.52	0.63	-0.77	0.61
6	0.60	0.84	0.19	0.84	-0.48	0.65	-0.76	0.61
7	0.63	0.83	0.16	0.85	-0.45	0.66	-0.74	0.63
8	0.70	0.83	0.12	0.86	-0.42	0.68	-0.75	0.62
9	0.77	0.82	0.08	0.86	-0.40	0.70	-0.74	0.63
10	0.77	0.82	0.03	0.86	-0.36	0.72	-0.71	0.65
11	0.79	0.81	0.00	0.87	-0.32	0.74	-0.56	0.73
12	0.82	0.81	-0.05	0.86	-0.27	0.77	-0.44	0.80
13	0.76	0.82	-0.10	0.86	-0.26	0.78	-0.41	0.83
14	0.78	0.82	-0.13	0.86	-0.24	0.79	-0.40	0.84
15	0.77	0.82	-0.17	0.85	-0.22	0.80	-0.38	0.86

Notes: In each case, the coefficient and AIC is for a probit model where the selected variable (along with a constant) is the only explanatory variable

Out-of-sample forecasting

The previous sections look at recession probabilities based on in-sample forecasting. As in, we use the full sample of data to estimate fitted values. However, we are also interested in assessing how such models would perform when forecasting in real-time using data only available at certain points in time. To assess this, we examine forecasts of a future recession based on parameters estimated only on past data up to a certain point.

Given the limited number of recessions in our sample, we opt to extend our sample period as far back as possible. This allows us a sufficient window to, first, estimate our initial model, and second, to forecast beyond the sample used for estimation. We are able to extend our sample back to 1985 for our dependent variable using the Chow-Lin method to combine monthly Live Register data on unemployment with our annual unemployment data. We also have information on three of our key predictors back to 1985. As a result, for this exercise, we drop the New Activity measure from our probit models when constructing out of sample forecasts.

The approach taken here is to use an expanding estimation window, and then use that equation to estimate the probability of recession in 3, 6 or 12 months' time. To assess the out-of-sample performance, we need at least one episode of recession (so we can estimate an equation before forecasting out of sample). As a result, out-of-sample forecasting only begins after the first defined recession ends.[12] Thereafter we use the estimated equation to forecast 3, 6 or 12 months ahead. We then add one more recent observation to our estimation sample and forecast 3, 6 or 12 months ahead again from this new starting point. We repeat this process over and over until the we have out-of-sample forecasts running all the way from May 1992 to present.

Clearly, there are not a large number of recessionary episodes for us to assess the out of sample forecasting performance. In total there are 6 recessionary episodes of varying duration we identify after 1992.[13] Figure 4 shows the predicted recession probabilities compared to the recessionary periods we have defined. Unsurprisingly, the two short, mild recessions in 2002 are not deemed likely to occur 12 or months ahead. The onset of the global financial crisis is deemed to have a 15% chance of occurring 12 or months ahead of time according to our models.

The Covid-induced recession beginning in March 2020 was not deemed likely 12 or 6 months previously. This is hardly surprising as the Irish economy had been performing well at the time and this recession was caused by a once in a generation pandemic, with non-economic factors having a far greater bearing than any economic factors that we might control for as predictors. See Annex D for results excluding the Covid-induced recession, which are similar to the baseline results.

12 The first recession we identify starts in March 1991 and ends in April 1992.

13 March 2002 (duration 1 month), November 2002 (duration 2 months), May 2008 (long duration), July 2011 (duration 2 months), February 2012 (duration 2 months) and March 2020 (duration 12 months)

Figure 4: Out-of-sample forecasts

12-month-ahead prediction horizon, expanding estimation wind

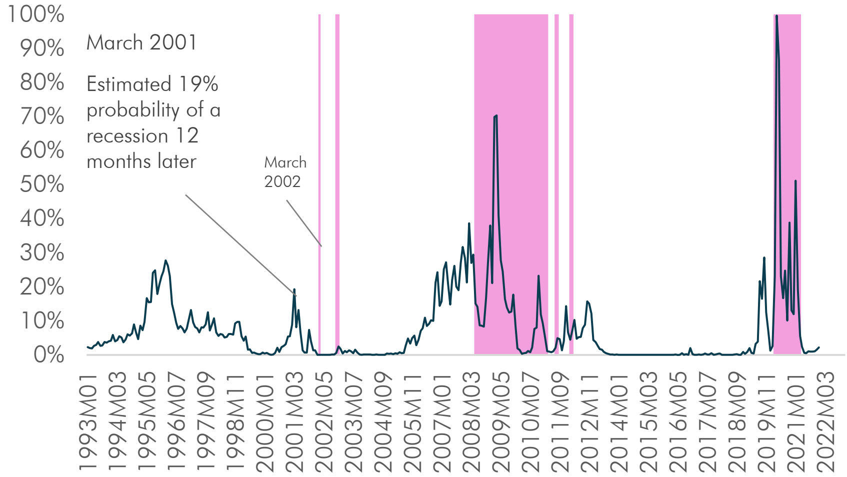

6-month-ahead prediction horizon, expanding estimation window

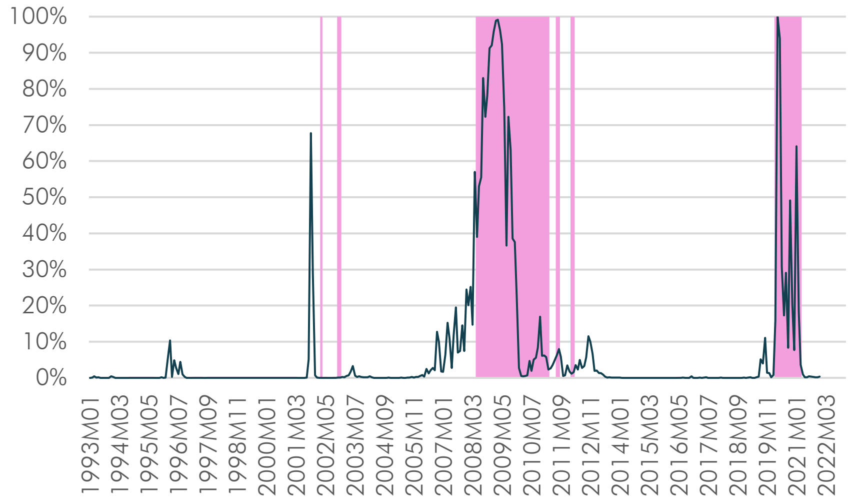

3-month-ahead prediction horizon, expanding estimation window

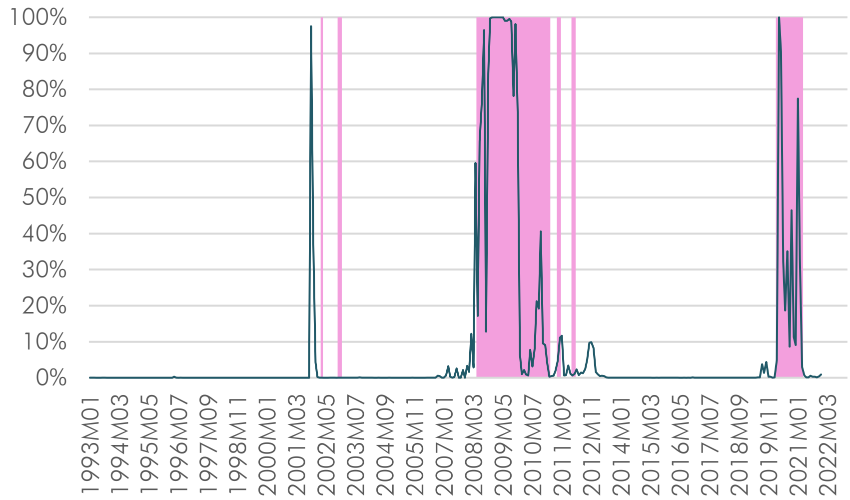

Finally, we examine how forecasting at different horizons performs relative to one another. Figure 5 shows the rate of true positives to false positives for each of the three forecast horizons considered for out of sample forecasting. We can see that 12 month ahead forecasts are significantly better than random guesses. However, forecasting performance is much improved if one is forecasting at a shorter horizon. Both 3 and 6 month ahead forecasts are much improved relative to 12 month ahead forecasts.

Figure 5: ROC curve for Out-of-sample forecasts

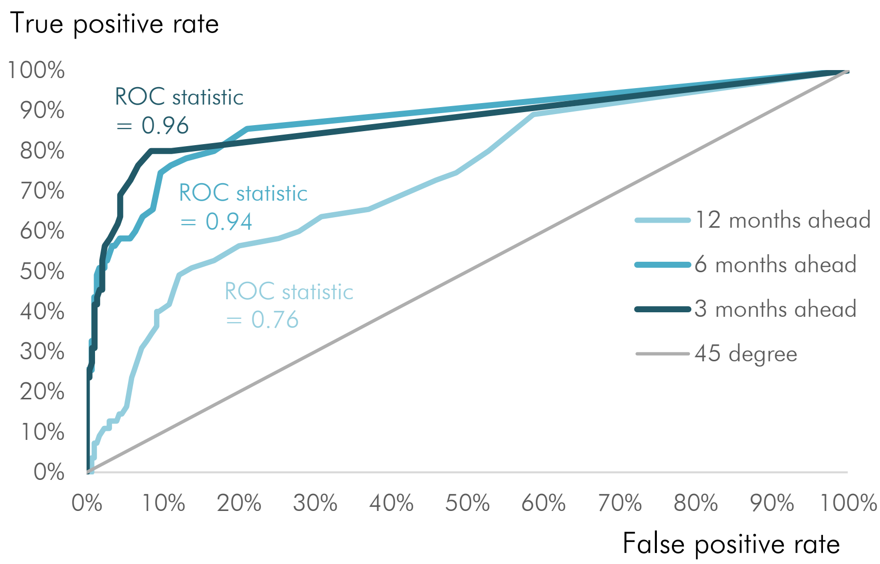

Note: Forecasts for 3, 6 and 12 months ahead are as shown in Figure 4.

Robustness checks

As a test to see whether our model is robust to other recession definitions, we explore a number of alternatives (Annex D). We find that the general results remain intact regardless of the recession definitions used. The forecasting performance diminishes somewhat in the case of recessions defined based on the technical recession definition (two consecutive quarters of contraction) for domestic demand measures. However, the performance is still notably better than chance predictions. Other alternative definitions based on the Bry-Boschan (1971) business cycle dating algorithms and based on a suite of output gap models result in predictions that are broadly as good, and in some cases better, than our baseline definition. We can also see that the models' forecasts are again more accurate at shorter horizons than at longer horizons.

4. Conclusions

We attempt to identify ways to predict Irish recessions based on available indicators that meet certain desirable criteria. While standard macroeconomic forecasting tools rely heavily on recent momentum, and so are less useful for identifying turning points, a rich literature has emerged using probit models and indicators, such as the Term Spread, to predict recessions.

We examine a large suite of potential predictors of recessions for Ireland, drawing on both domestic and international data. In addition, we build on recent innovations in textual analysis to develop an indicator based on news articles and references to unemployment. We also follow relatively recent literature for the US that uses Google Search data to help identify future recessions. In identifying a preferred model, we land on four key variables: the US Term Spread, the US TED Spread, a composite indicator of surveyed employment expectations for Ireland, and a composite measure of New Activity in Ireland.

To our knowledge, the use of surveyed expectations of businesses and consumers about future employment prospects is not a variable that has been used elsewhere for predicting recessions. However, we find that it is consistently useful as a predictor, at both short- and longer-term forecasting horizons.

Looking ahead, we can draw out some useful conclusions based on our findings. Employment expectations in Ireland have disimproved of late. The US Term Spread and TED spread both signal some concern about future activity. However, with the exception of the term spread, these indicators have not deteriorated substantially. The models therefore point to a relatively low likelihood of a recession in the coming year.

However, we would caution that there are many risks not captured by the model — some of which are still playing a major role in current developments. These include the impact on demand from rising prices, the weaker outlook overseas, and geopolitical risks outside the scope of the models, which rely on economic rather than non-economic factors to predict recessions.

This paper therefore provides a useful insight into how financial developments overseas, employment expectations at home, and a variety of other domestic leading indicators can provide some useful insight into the likelihood of a forthcoming recession. However, they do not provide a crystal ball that can capture factors that have had a big bearing on recent developments in economic activity, such as the Covid-19 pandemic and the cost-of-living crisis.

In terms of avenues for future research, our work could be developed further by exploring models that employ mixed frequencies of data, such as MIDAS models. Machine learning-type approaches, such as Lasso regression analysis, could also be usefully employed, particularly for the variable selection stage. Another interesting avenue to explore further would be additional text-based predictors, such as those that build on official transcripts, including those of central banks and other official bodies.

References

Benzoni, L., Chyruk, O., and D. Kelley., (2018). "Why does the yield-curve slope predict recessions?" Federal Reserve Bank of Chicago Working Paper, pp. 2018–15. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3298006

Bernard H. and S. Gerlach, (1998). "Does the Term Structure Predict Recessions? The International Evidence". No 1892, CEPR Discussion Papers. Available at: https://econpapers.repec.org/paper/cprceprdp/1892.htm

Bry, G. and C. Boschan (1971). "Cyclical Analysis of Time Series: Selected Procedures and Computer Programs". NBER, New York. Available at: https://www.nber.org/books-and-chapters/cyclical-analysis-time-series-selected-procedures-and-computer-programs

Estrella, A., and G. A. Hardouvelis, (1991). "Term structure as a predictor of real economic activity". Journal of Finance, 2(46). Available at: https://www.jstor.org/stable/2328836

Estrella, A., and F. S. Mishkin, (1998). "Predicting U.S. Recessions: Financial Variables as Leading Indicators". The Review of Economics and Statistics, 1998, 80 (1), pp. 45-61. Available at: https://www.jstor.org/stable/2646728

Fama, E.F., (1986). "Term premiums and default premiums in money markets". Journal of Financial Economics 17 (1), pp.175–196. Available at: https://www.sciencedirect.com/science/article/abs/pii/0304405X86900103

FitzGerald, J. and I.M Kearney, (1999). "Migration and the Irish Labour Market". ESRI Working Paper No 113. Available at: https://www.esri.ie/system/files?file=media/file-uploads/2015-07/WP113.pdf

Kessel, R. A., (1965). "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.

Minesso, M.F., L. Lebastard, and H. Le Mezo (2022). "Text-Based Recession Probabilities." IMF Econ Rev (2022). Available at: https://doi.org/10.1057/s41308-022-00177-5

Moneta, F., (2005). “Does the Yield Spread Predict Recessions in the Euro Area?” International Finance, Volume 8, Issue 2, Summer 2005, pp. 263-301. Available at: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-2362.2005.00159.x

OECD (2012). OECD System of Composite Leading Indicators. OECD Statistical Note. Available at: https://www.oecd.org/sdd/leading-indicators/41629509.pdf

Stuart, R. (2020). “Monetary regimes, the term structure and business cycles in Ireland, 1972-2018”. QUCEH Working Paper Series, No. 2020-03, Queen's University Centre for Economic History (QUCEH), Belfast. Available at: http://hdl.handle.net/10419/218817

Tkacz, G. (2013). “Predicting Recessions in Real-Time: Mining Google Trends and Electronic Payments Data for Clues”. C.D. Howe Institute, Commentary No. 387 September 2013 Financial Services. Available at: https://www.cdhowe.org/public-policy-research/predicting-recessions-real-time-mining-google-trends-and-electronic-payments-data-clues

Wheelock D.C. and M.E. Wohar (2009). “Can the term spread predict output growth and recessions? A survey of the literature”. Review, Federal Reserve Bank of St. Louis, vol. 91(Sep), pp. 419-440. Available at: https://research.stlouisfed.org/publications/review/2009/09/01/can-the-term-spread-predict-output-growth-and-recessions-a-survey-of-the-literature/

Annex A: Detail on Variables used

Table A1 Indicators used as predictors of recessions

Indicator	Justification	Expected sign
Domestic indicators		
1) Early Stage Indicators		(–)
Markit PMI Manufacturing New Orders	Weaker orders entail early signs of lower activity	
Markit PMI Construction New Orders	Weaker orders entail early signs of lower activity	
Markit PMI Services New Business	Weaker new business entails early signs of lower activity	
DoE Housing commencements	Weaker housing starts entails early signs of lower activity	
2) Employment expectations indicators		(–)
DG ECFIN Industrial Employment Expectations for the Months Ahead, Balance, SA	Weaker job expectations likely to mean increased unemployment	
DG ECFIN Construction Employment Expectations Over the Next 3 Months, Balance, SA	Weaker job expectations likely to mean increased unemployment	
DG ECFIN Services Expectations of the Employment Over the Next 3 Months, Balance, SA	Weaker job expectations likely to mean increased unemployment	
DG ECFIN Retail Employment Expectations Over the Next 3 Months, Balance, SA	Weaker job expectations likely to mean increased unemployment	
Eurostat Consumer Unemployment Expectations Over the Next 12 Months, SA, Inverse	Reduction in inverse likely associated with higher unemployment	
3) Input price indicators		(+)
Markit PMI Manufacturing Input Prices	Expectations sensitive, but also a driver of demand	
Markit PMI Construction Input Prices	Expectations sensitive, but also a driver of demand	
Markit PMI Services Input Prices	Expectations sensitive, but also a driver of demand	
4) Expectations-sensitive indicators		(–)
ISEQ Benchmark Overall Index Closing price	Weaker stock market signals expectations weaker demand	
Markit PMI Services Business Expectations	Weaker business expectations signals weaker demand	
Markit PMI Construction Business Expectations	Weaker business expectations signals weaker demand	
5) Google Searches		
Searches of terms “recession” and “unemployment”	Google searches of these terms could signal an increased expectation of a downturn among individuals	(+)
6) News Articles		
Share of articles containing the term “unemployment” in the Irish Times	News coverage could reflect an increased expectation of a downturn	(+)

Indicator	Justification	Expected sign
External indicators		
7) So-called "Prime movers"		Mixed
TED Spread	Higher US credit or default risk could signal global downturn	(+)
Term Spread (US)	Falling spread signals higher financial risk-taking before downturn	(–)
GBP per EUR exchange rate	Stronger euro could lead to weaker external demand	(+)
EUR per USD exchange rate	Stronger dollar could lead to stronger external demand	(–)
Ireland Import Prices Index	Higher import prices hurts demand	(+)
8) External demand indicators		(–)
OECD MEI Composite Leading Indicators, United States, Amplitude Adjusted, SA	Weaker external demand leading to higher likelihood of recession	
OECD MEI Composite Leading Indicators, United Kingdom, Amplitude Adjusted, SA	Weaker external demand leading to higher likelihood of recession	
OECD MEI Composite Leading Indicators, Euro Area 19, Amplitude Adjusted, SA	Weaker external demand leading to higher likelihood of recession	

Figure A1: Input variables

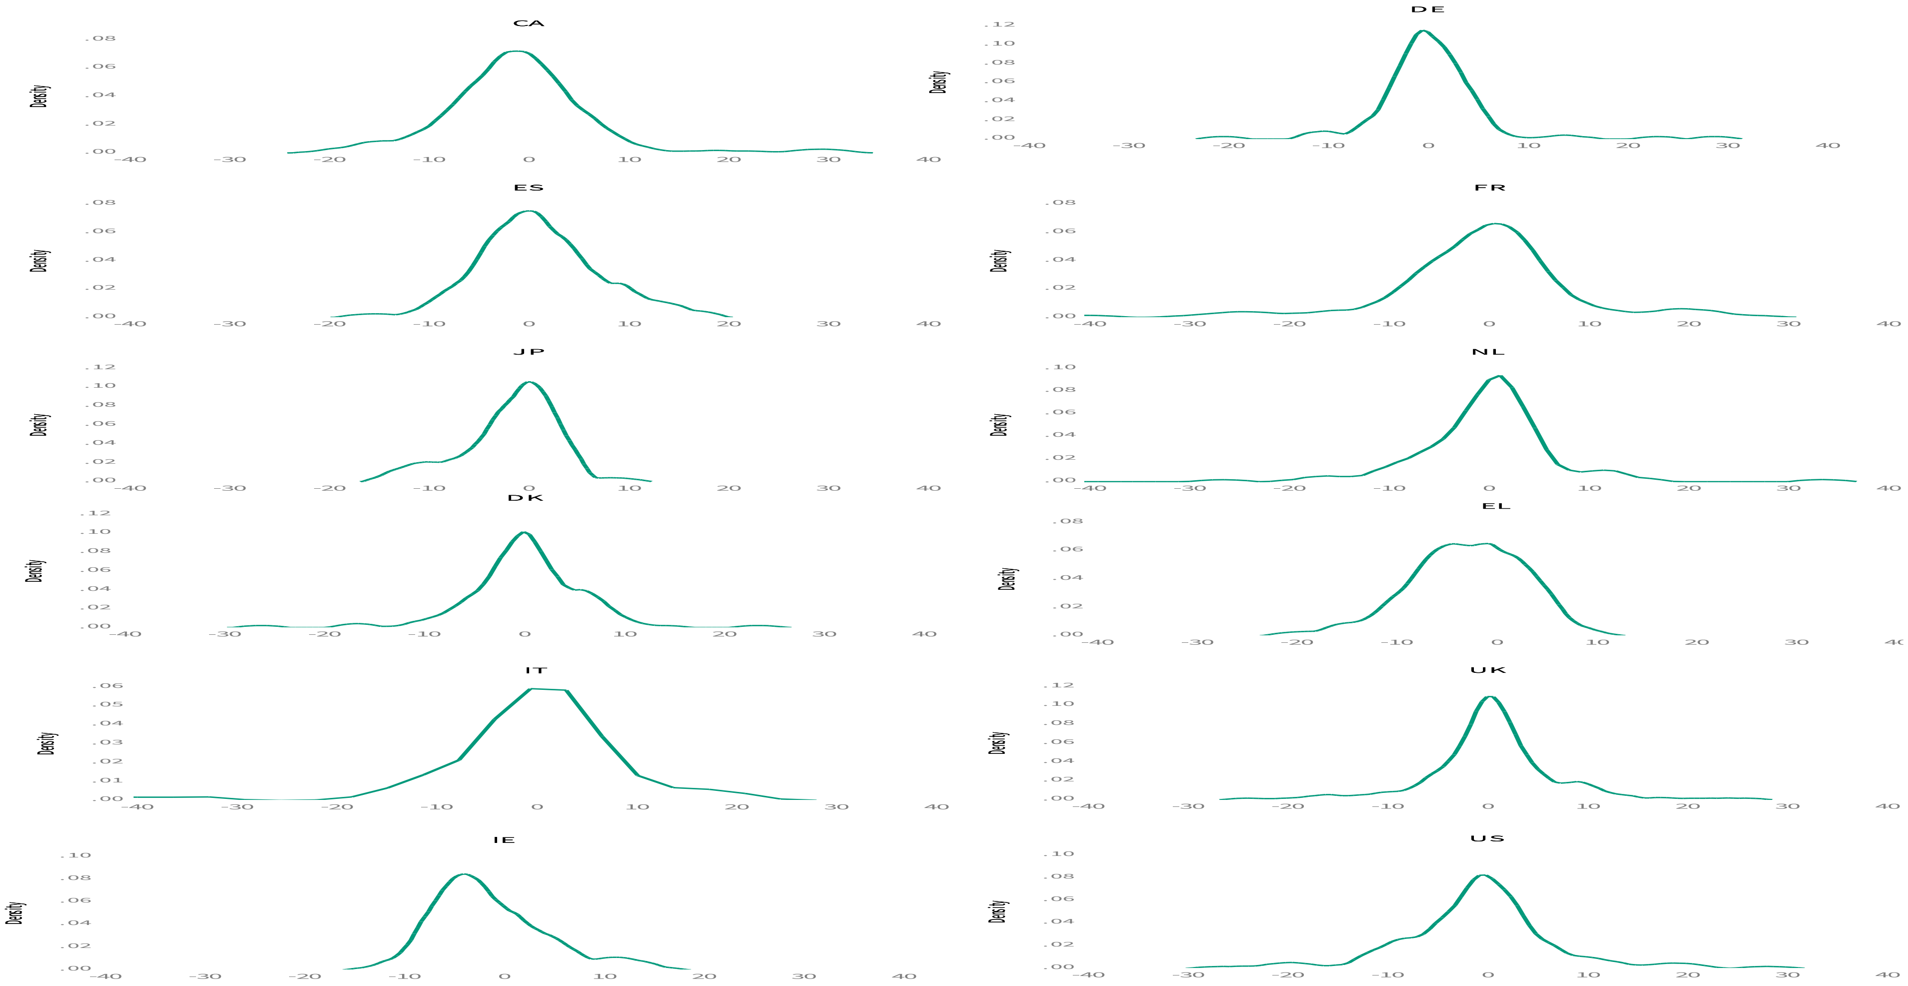

Figure A2: Distributions of indicators

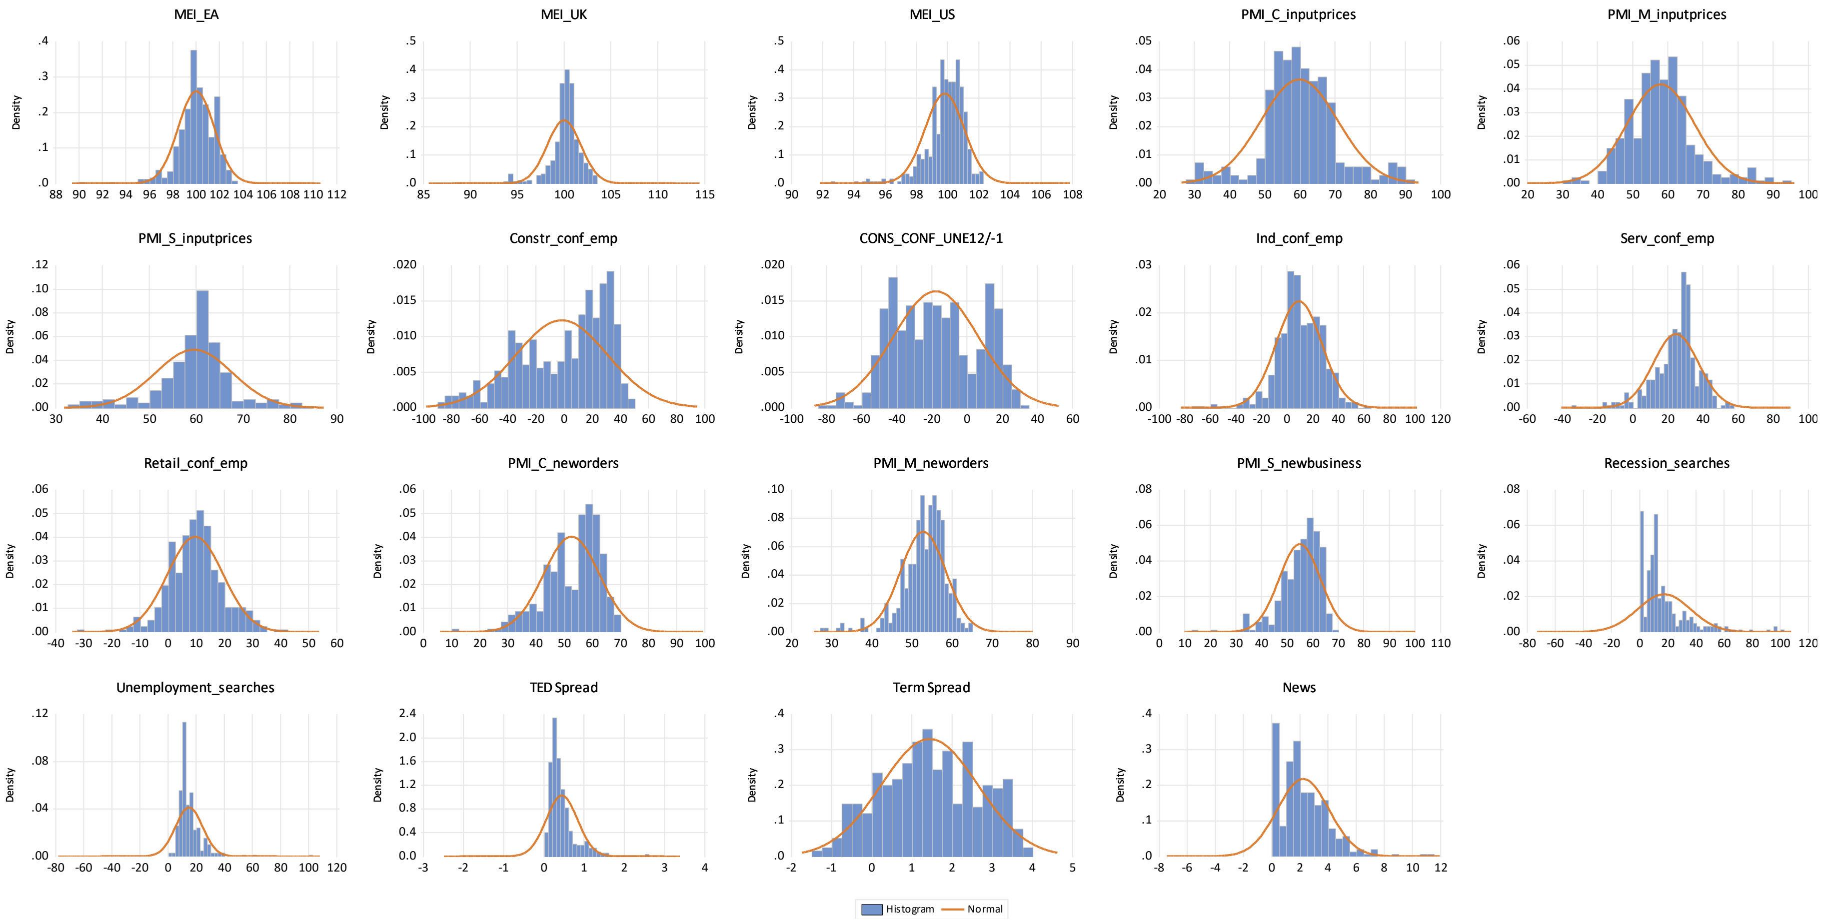

Figure A3: Principal components

Main economic indicators

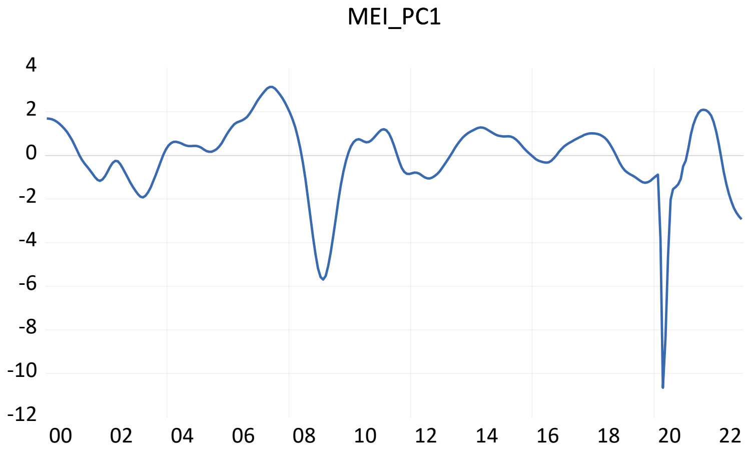

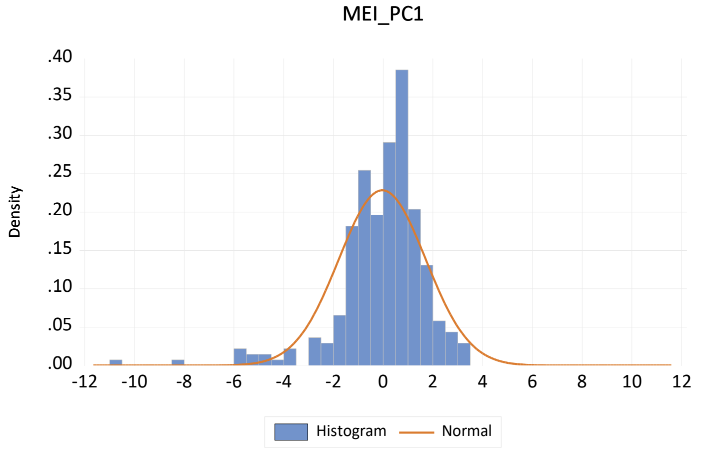

Input prices

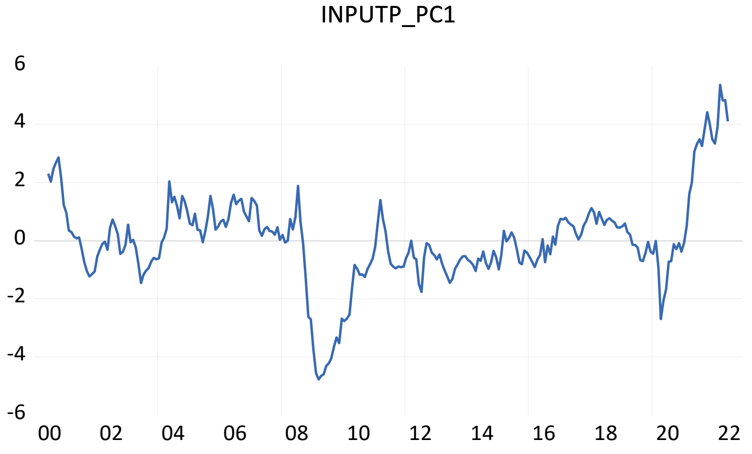

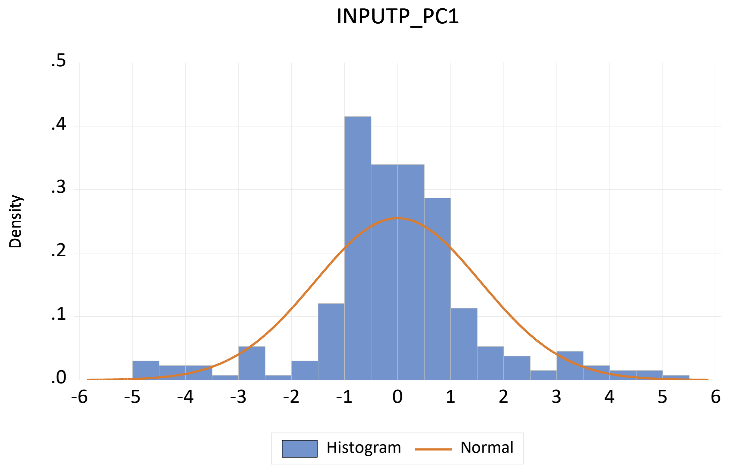

Employment expectations

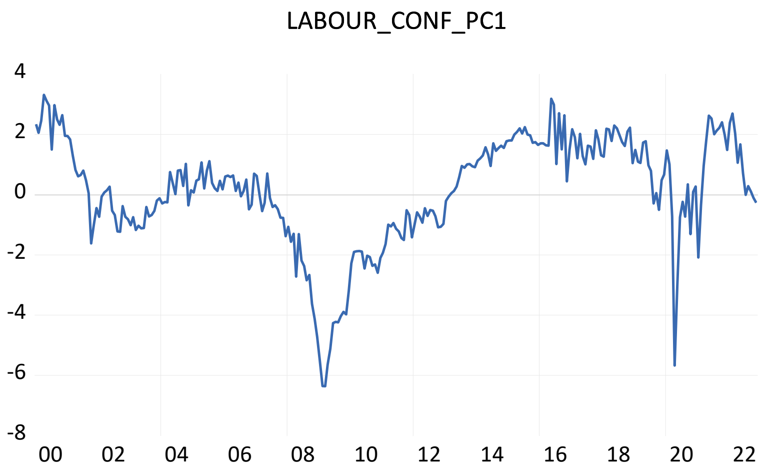

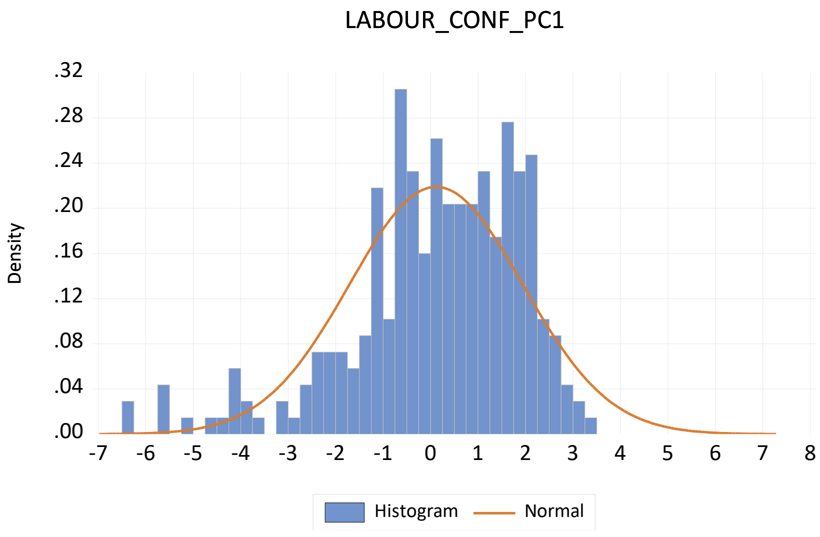

New activity

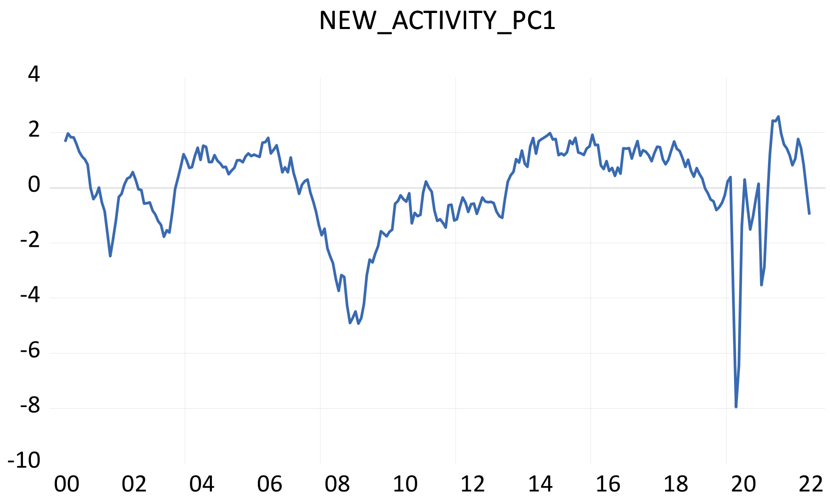

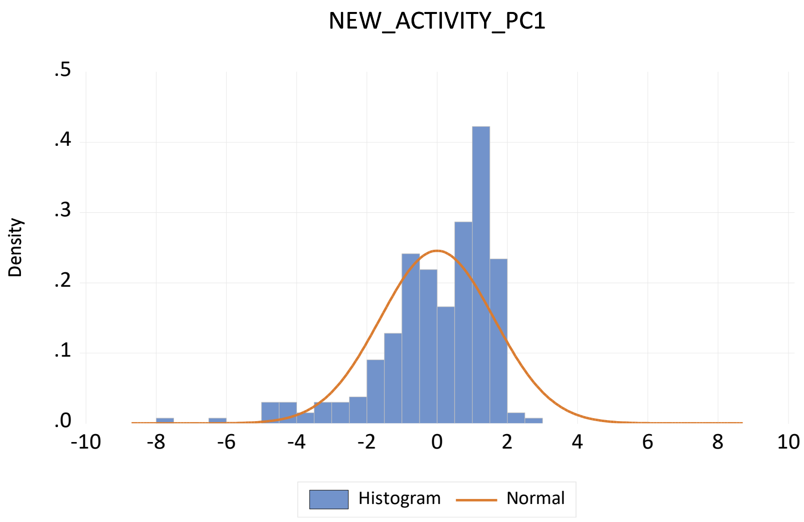

Annex B: Model Selection

Table B1: Model selection at 15-month horizon

	Preferred model	Model 5	Model 4	Model 3	Model 2	Model 1
TED Spread	-0.47	1.36	1.37	1.04	1.72	1.04
	(0.21)**	(0.64)**	(0.59)**	(0.49)**	(0.73)**	(0.67)
Term Spread	1.26	-0.50	-0.52	-0.62	-0.68	-0.59
	(0.55)**	(0.19)***	(0.28)*	(0.28)**	(0.31)**	(0.25)**
New activity		0.11	0.12	-0.15	-0.27	-0.28
		(0.24)	(0.24)	(0.22)	(0.21)*	(0.21)*
Hiring/employment expectations	-0.48	-0.57	-0.57	-0.49	-0.38	-0.41
	(0.15)***	(0.2)***	(0.2)***	(0.2)**	(0.17)**	(0.18)**
Input prices					-0.59	-0.58
					(0.23)**	(0.22)***
External demand				0.21	0.27	0.21
				(0.13)	(0.14)*	(0.17)*
ISEQ			0.00	0.00	0.00	0.00
			(0)	(0)	(0)	(0)
Google search terms					-0.39	-0.41
					(0.16)**	(0.16)**
Newspaper articles						-0.45
						(0.18)**
Constant	-0.67	-0.74	0.10	-0.54	-0.54	0.97
	(0.52)	(1.41)	(1.38)	(1.41)	(1.41)	(1.61)
BIC	0.74	0.76	0.78	0.79	0.80	0.80
Pseudo R-Squared	0.41	0.41	0.41	0.43	0.48	0.51
Observations	224	224	224	224	212	212
Sample	2004M04 2022M11	2004M04 2022M11	2004M04 2022M11	2004M04 2022M11	2004M04 2022M11	2005M04 2022M11

Notes: Robust standard errors (Newey-West correction) are estimated. *, ** and *** denote statistical significance at a 10% ,5% and 1% level, respectively. The BIC is the Bayesian Information Criterion.

Table B2: Model selection at 12-month horizon

	Preferred model	Model 5	Model 4	Model 3	Model 2	Model 1
TED Spread	1.87	1.49	1.47	1.29	1.64	0.69
	(0.56)***	(0.64)**	(0.68)**	(0.62)**	(0.72)**	(0.65)
Term Spread		-0.29	-0.28	-0.45	-0.48	-0.53
		(0.2)	(0.18)	(0.27)*	(0.32)	(0.25)**
New activity			-0.03	-0.20	-0.24	-0.59
			(0.21)	(0.19)	(0.19)	(0.25)**
Hiring/employment expectations	-0.35	-0.45	-0.42	-0.36	-0.24	-0.08
	(0.12)***	(0.14)***	(0.16)**	(0.18)**	(0.17)	(0.21)*
Input prices				0.18	-0.49	-0.87
				(0.16)	(0.29)*	(0.36)**
External demand					0.26	0.33
					(0.18)	(0.18)*
ISEQ				0.00	0.00	0.00
				(0)	(0)	(0)**
Google search terms					-0.11	-0.15
					(0.15)	(0.14)
Newspaper articles						-0.99
						(0.23)***
Constant	1.87	1.49	1.47	1.29	1.64	0.69
	(0.56)***	(0.64)**	(0.68)**	(0.62)**	(0.72)**	(0.65)
BIC	0.77	0.75	0.78	0.81	0.84	0.72
Pseudo R-Squared	0.36	0.39	0.39	0.41	0.45	0.57
Observations	224	224	224	224	215	215
Sample	2004M04 2022M11	2004M04 2022M11	2004M04 2022M11	2004M04 2022M11	2005M01 2022M11	2005M01 2022M11

Notes: Robust standard errors (Newey-West correction) are estimated. *, ** and *** denote statistical significance at a 10% ,5% and 1% level, respectively. The BIC is the Bayesian Information Criterion.

Table B3: Model selection at 9-month horizon

	Preferred model	Model 5	Model 4	Model 3	Model 2	Model 1
TED Spread	1.50	1.52	1.12	1.90	2.06	1.68
	(0.57)***	(0.56)***	(0.67)*	(0.83)**	(0.81)**	(0.85)**
Term Spread			-0.30	-0.28	-0.11	-0.17
			(0.22)	(0.23)	(0.31)	(0.3)
New activity	-0.84	-0.55	-0.50	-0.33	-0.26	-0.34
	(0.19)***	(0.26)**	(0.25)**	(0.26)	(0.23)	(0.28)
Hiring/employment expectations		-0.29	-0.49	-0.47	-0.50	-0.55
		(0.23)	(0.16)***	(0.16)***	(0.16)***	(0.18)***
Input prices					-0.68	-0.87
					(0.31)**	(0.33)***
External demand				-0.28	-0.35	-0.45
				(0.24)	(0.23)	(0.23)*
ISEQ					0.00	0.00
					(0)	(0)
Google search terms				0.29	0.06	0.24
				(0.38)	(0.36)	(0.3)
Newspaper articles						-0.51
						(0.22)**
Constant	-1.57	-1.52	-1.01	-1.20	-3.08	-1.07
	(0.32)***	(0.33)***	(0.55)*	(0.55)**	(1.42)**	(1.53)
BIC	0.58	0.58	0.59	0.62	0.63	0.62
Pseudo R-Squared	0.53	0.55	0.57	0.59	0.63	0.66
Observations	224	224	224	218	218	218
Sample	2004M04 2022M11	2004M04 2022M11	2004M04 2022M11	2004M10 2022M11	2004M10 2022M11	2004M10 2022M11

Notes: Robust standard errors (Newey-West correction) are estimated. *, ** and *** denote statistical significance at a 10% ,5% and 1% level, respectively. The BIC is the Bayesian Information Criterion.

Table B4: Model selection at 6-month horizon

	Preferred model	Model 5	Model 4	Model 3	Model 2	Model 1
TED Spread	1.53	1.22	1.11	1.30	1.56	1.55
	(0.61)**	(0.63)*	(0.7)	(0.67)*	(0.64)**	(0.69)**
Term Spread		-0.19	-0.17	0.09	0.16	0.26
		(0.18)	(0.18)	(0.22)	(0.24)	(0.27)
New activity			-0.12	-0.02	0.05	0.09
			(0.24)	(0.25)	(0.24)	(0.25)
Hiring/employment expectations	-0.77	-0.88	-0.80	-1.04	-0.98	-1.05
	(0.2)***	(0.17)***	(0.21)***	(0.32)***	(0.31)***	(0.37)***
Input prices					-0.65	-0.72
					(0.3)**	(0.29)**
External demand	-0.65	-0.59		-0.80	-0.79	-0.92
	(0.25)***	(0.21)***		(0.23)***	(0.23)***	(0.28)***
ISEQ				0.00	0.00	0.00
				(0)	(0)**	(0)*
Google search terms						0.00
						(0.24)
Newspaper articles						-0.15
						(0.18)*
Constant	-1.43	-1.08	-1.09	-3.00	-4.13	-4.08
	(0.32)***	(0.48)**	(0.48)**	(1.34)**	(1.45)***	(1.82)**
BIC	0.53	0.55	0.57	0.58	0.57	0.61
Pseudo R-Squared	0.50	0.60	0.61	0.62	0.65	0.66
Observations	224	224	224	224	224	221
Sample	2004M04 2022M11	2004M04 2022M11	2004M04 2022M11	2004M04 2022M11	2004M04 2022M11	2004M07 2022M11

Notes: Robust standard errors (Newey-West correction) are estimated. *, ** and *** denote statistical significance at a 10% ,5% and 1% level, respectively. The BIC is the Bayesian Information Criterion.

Table B5: Model selection at 3-month horizon

	Preferred model	Model 5	Model 4	Model 3	Model 2	Model 1
TED Spread			0.16	0.22	0.27	0.89
			(0.52)	(0.64)	(0.71)	(0.85)
Term Spread					0.14	0.39
					(0.22)	(0.3)
New activity				0.03	0.04	0.10
				(0.29)	(0.29)	(0.31)
Hiring/employment expectations	-0.86	-0.76	-0.93	-1.01	-1.04	-0.91
	(0.2)***	(0.22)***	(0.28)***	(0.33)***	(0.35)***	(0.32)***
Input prices						-0.59
						(0.22)***
External demand	-0.49	-0.40	-0.48	-0.64	-0.72	-0.79
	(0.15)***	(0.16)**	(0.18)***	(0.24)***	(0.25)***	(0.27)***
ISEQ			0.00	0.00	0.00	0.00
			(0)	(0)*	(0)*	(0)**
Google search terms		0.43	0.46	0.47	0.38	0.49
		(0.28)	(0.27)*	(0.27)*	(0.24)	(0.23)**
Newspaper articles						0.05
						(0.2)*
Constant	-0.89	-0.90	-1.84	-2.12	-2.66	-4.48
	(0.23)***	(0.24)***	(0.67)***	(0.68)***	(1.26)**	(2.05)**
BIC	0.51	0.52	0.55	0.57	0.59	0.60
Pseudo R-Squared	0.59	0.61	0.62	0.64	0.64	0.67
Observations	224	224	224	222	222	222
Sample	2004M04 2022M11	2004M04 2022M11	2004M04 2022M11	2004M04 2022M09	2004M04 2022M09	2004M04 2022M09

Notes: Robust standard errors (Newey-West correction) are estimated. *, ** and *** denote statistical significance at a 10% ,5% and 1% level, respectively. The BIC is the Bayesian Information Criterion.

Annex C: Model Performance at different forecast horizons

Figure C1: Root Mean Squared Error for k = 1 month to 15 months, baseline specification

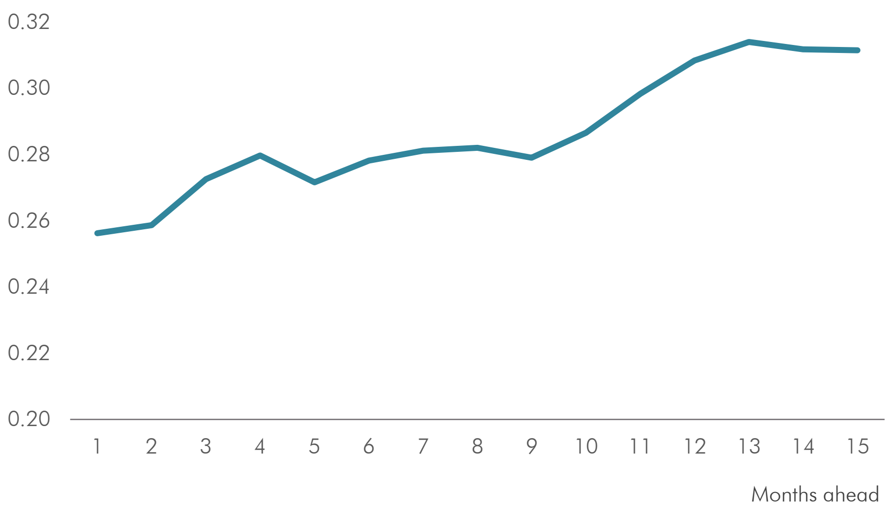

Notes: The root mean squared error shown is from using the baseline specification to forecast k months ahead,

where k ranges from one to fifteen months. These are within sample forecasts. The baseline specification is given by equation (2).

Figure C2: AIC for k = 1 month to 15 months, baseline specification

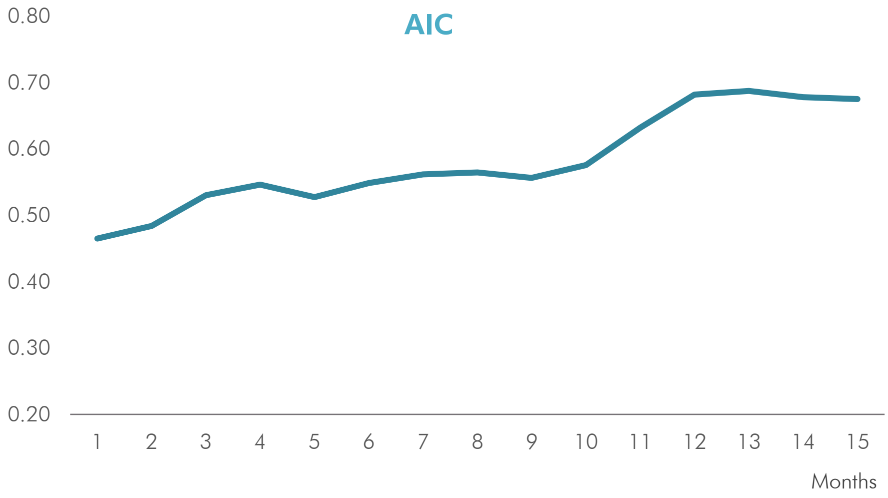

Notes: The Akaike Information Criterion (AIC) statistic shown is from estimating the baseline specification with independent variables at a lag of k months, where k ranges from one to fifteen months. Lower values of AIC are favoured. The baseline specification given by equation (2).

Figure C3: BIC for k = 1 month to 15 months, baseline specification

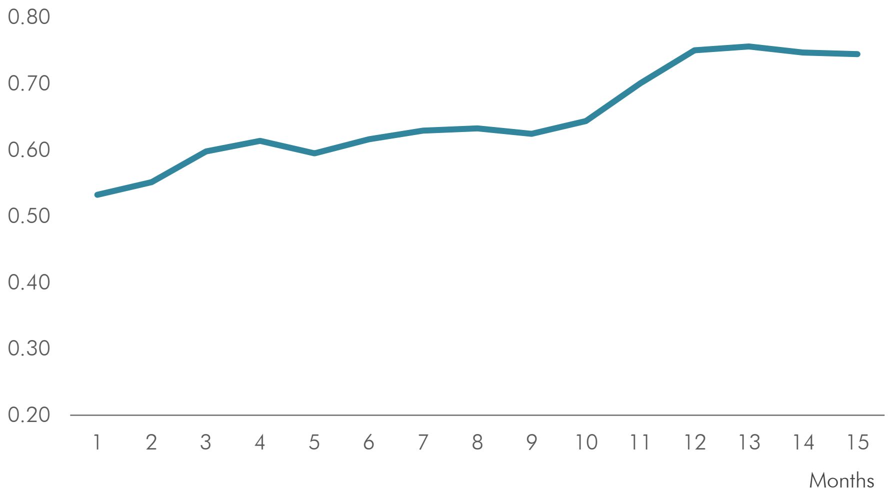

Notes: The Bayesian Information Criterion (BIC) statistic shown is from estimating the baseline specification with independent variables at a lag

of k months, where k ranges from one to fifteen months. Lower values of BIC are favoured. The baseline specification is given by equation (2).

Figure C4: Pseudo R^2 for k = 1 month to 15 months, baseline specification

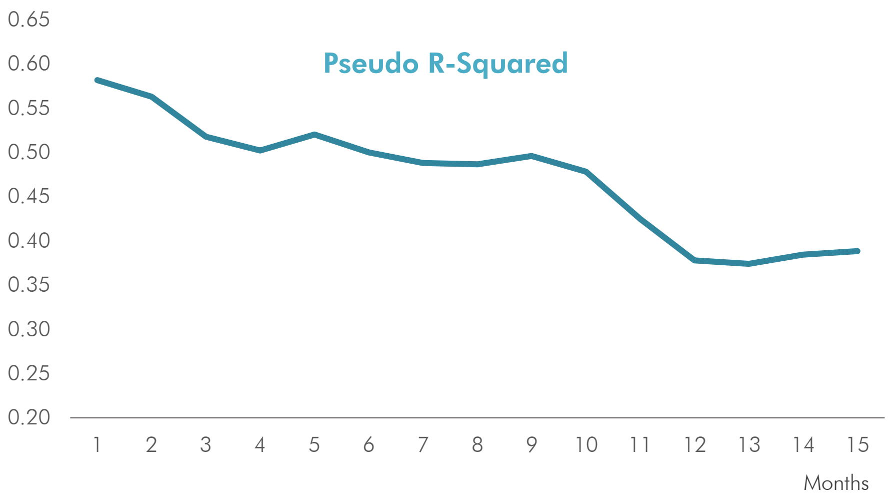

Notes: The Pseudo R^2 statistic shown is from estimating the baseline specification with independent variables at a lag

of k months, where k ranges from one to fifteen months. The baseline specification is given by equation (2).

Annex D: Robustness Checks

As a robustness check, we assess the out-of-sample forecasting performance of our model using different methods to identify recessions.

We use two alternative definitions based on quarterly modified domestic demand to identify our recession periods. First, we use the standard two quarter rule (a recession is defined as a period where demand contracts for two consecutive quarters on a seasonally adjusted basis).[14] Second, we apply the Bry and Boschan (1971) dating algorithm to our modified domestic demand series in log levels — a technique that is commonly applied to examine business cycle peaks and troughs.[15] Third, we construct recession dates using the output gap estimates of Casey (2019), which are based on a multivariate suite of models' approach to estimating potential output. These are derived from the output gap series, with a) falls in the output gap of at least two consecutive quarters required to indicate a recession, and b) an override for this rule depending on the scale of the contraction. The Bry and Boschan method only identifies 4 recessions, the first of which starts in 2008Q1. As a result, there is a very short period over which to assess out of sample forecasting performance using this method.

Figure D1 below shows the ROC curves for our out-of-sample predictions using these alternative recession dates relative to those shown in the main text. The forecasting performance diminishes somewhat in the case of recessions defined based on the technical recession definition (two consecutive quarters of contraction) for modified domestic demand measures. However, the performance is still notably better than chance predictions. Other alternative definitions based on the Bry-Boschan (1971) business cycle dating algorithms and based on a suite of output gap models result in predictions that are broadly as good, and in some cases better, than our baseline definition. We can also see that the models' forecasts are again more accurate at shorter horizons (Figures D2 and D3).

14 For the quarters which are deemed to be in recession, we assume all three months of that quarter are in recession. As modified domestic demand data only goes back to 1995Q1, we have a shorter sample. In addition, the first recession identified by this method begins in 2007Q2, so we have a much smaller number of out of sample recessions to assess forecasting performance.

15 The algorithm is based on parameters being set for a) the minimum phase for (2 quarters) expansions and contractions; b) the minimum cycle length for a complete contraction plus expansion (five quarters); c) an overrule for a) if, say, the contraction is especially large. We set these parameters in line with the standard rules used by the NBER.

Figure D1: ROC curve for 12-month-ahead forecasts

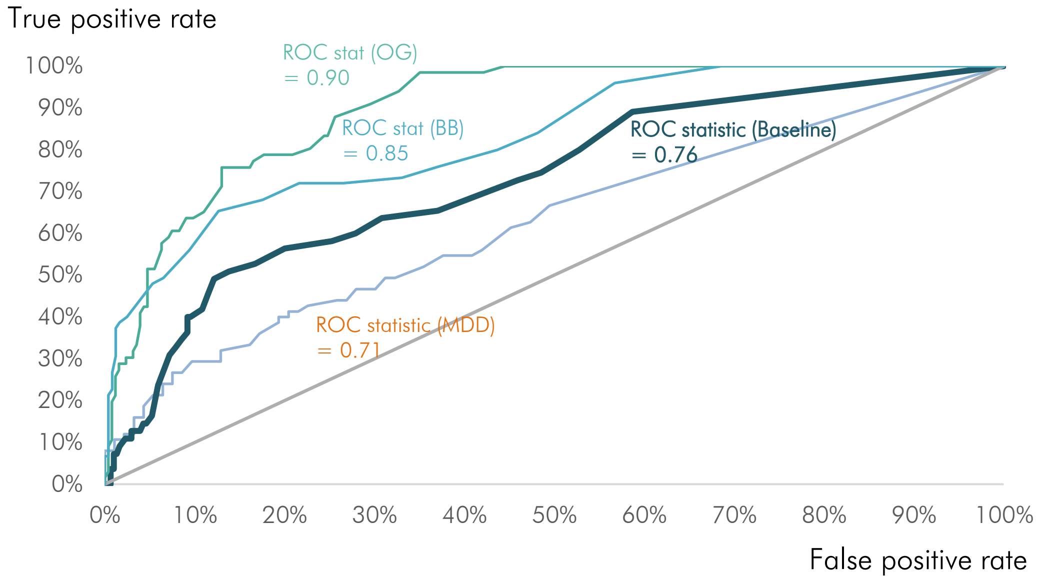

Notes: Baseline refers to the results using recession dates based on the monthly unemployment rate as used in the main text. OG refers to output gap, BB refers to the Bry and Boschan dating algorithm for identifying recession dates. MDD refers to modified domestic demand, where two consecutive quarters of declining seasonally adjusted modified domestic demand is deemed to signal the onset of a recession.

Figure D2: ROC curve for 6-month-ahead forecasts

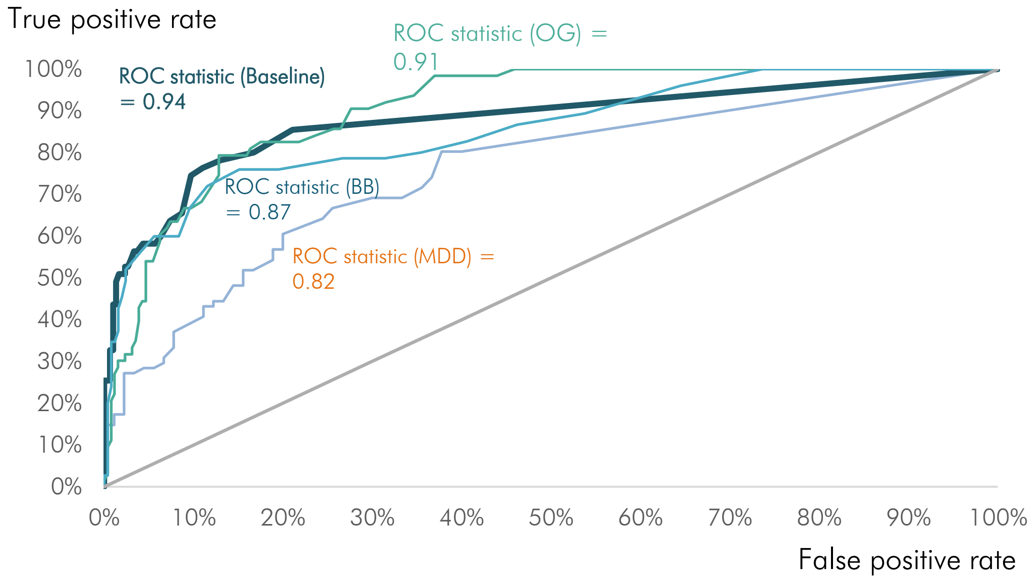

Notes: Baseline refers to the results using recession dates based on the monthly unemployment rate as used in the main text. OG refers to output gap, BB refers to the Bry and Boschan dating algorithm for identifying recession dates. MDD refers to modified domestic demand, where two consecutive quarters of declining seasonally adjusted modified domestic demand is deemed to signal the onset of a recession.

Figure D3: ROC curve for 3-month-ahead forecasts

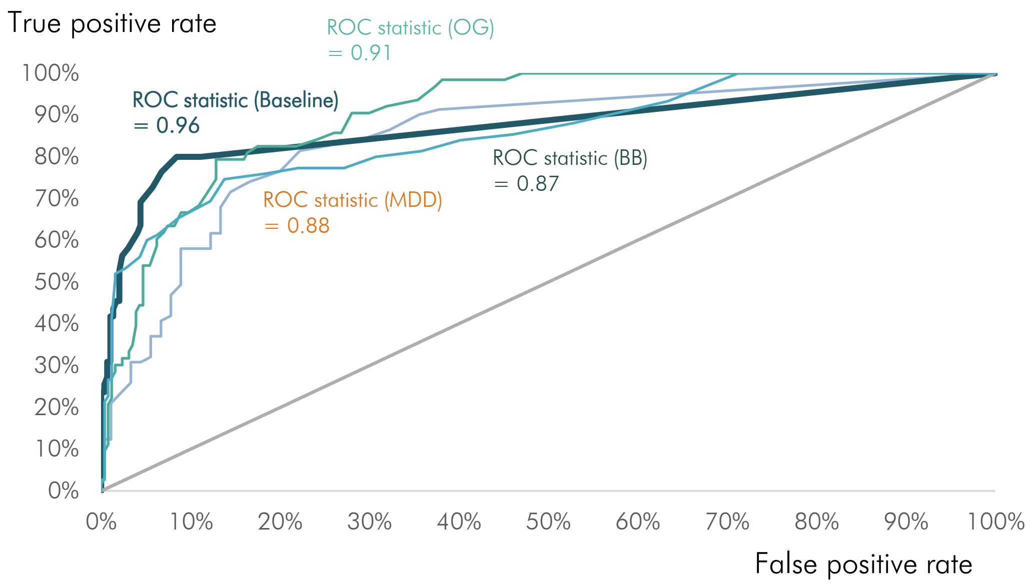

Notes: Baseline refers to the results using recession dates based on the monthly unemployment rate as used in the main text. OG refers to output gap, BB refers to the Bry and Boschan dating algorithm for identifying recession dates. MDD refers to modified domestic demand, where two consecutive quarters of declining seasonally adjusted modified domestic demand is deemed to signal the onset of a recession.

As a final robustness check, we assess the out-of-sample forecasting performance of our model excluding the Pandemic induced recession which started in March 2020. For both estimating our models and performing out of sample forecasts, we cut off our data in December 2019.

Figure D4, D5 and D6 show the ROC curves for our out-of-sample predictions for 12, 6 and 3 month ahead. These are shown for the baseline (including Covid) and when the Covid induced recession is excluded. We can see that forecasting performance is quite similar. In each case, forecasting performance is improved by omitting the 2020 recession. This is to be expected as that recession was caused by non-economic factors. However, the difference in forecasting performance is very small.

Figure D4: ROC curve for 12-month-ahead forecasts, excluding Covid

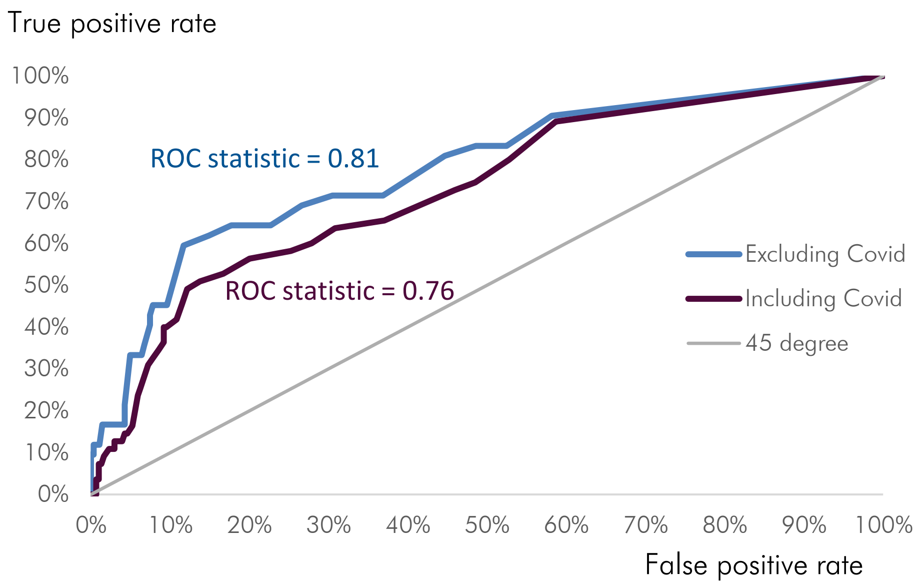

Notes: Baseline refers to the results using recession dates based on the monthly unemployment rate as used in the main text. Excluding Covid means the sample ends in December 2019.

Figure D5: ROC curve for 6-month-ahead forecasts, excluding Covid

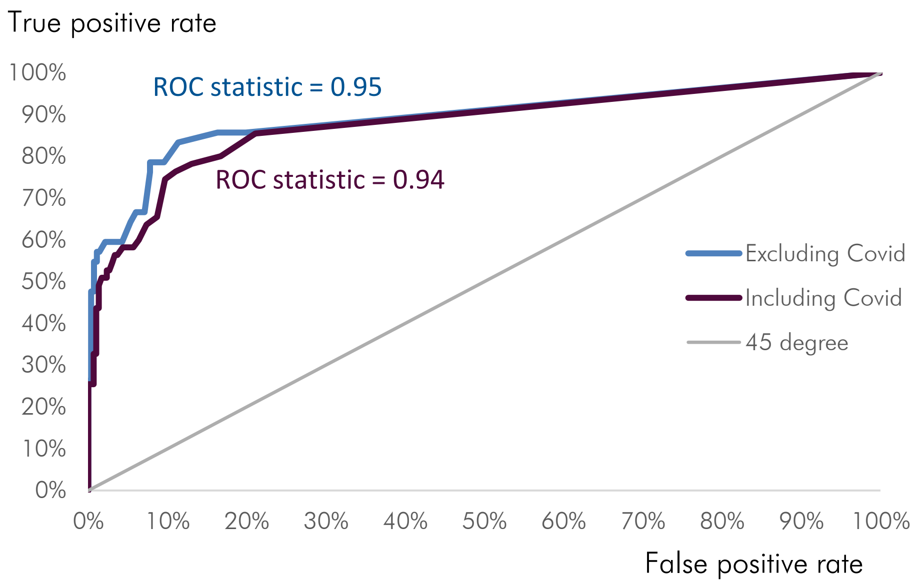

Notes: Baseline refers to the results using recession dates based on the monthly unemployment rate as used in the main text. Excluding Covid means the sample ends in December 2019.

Figure D6: ROC curve for 3-month-ahead forecasts, excluding Covid

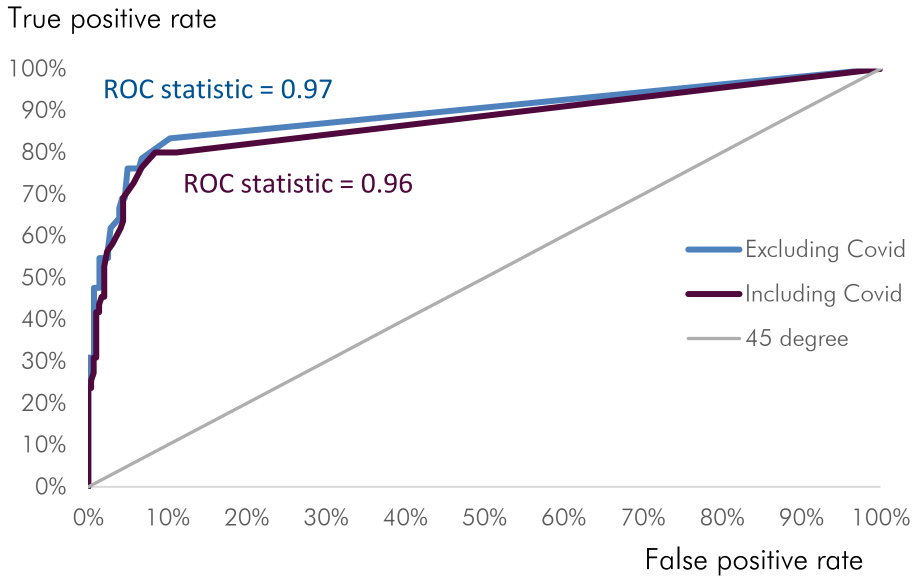

Notes: Baseline refers to the results using recession dates based on the monthly unemployment rate as used in the main text. Excluding Covid means the sample ends in December 2019.