Bulletin
de la Banque de France
250/3 - JANUARY-FEBRUARY 2024

BANQUE DE FRANCE
EUROSYSTÈME

Economic research

What does an inversion of the yield curve tell us?

In June 2023, the yield on 10-year government bonds fell below the yield on 2-year bonds in the euro area. Inversions of the yield curve such as this have sometimes been perceived as a signal of recession, and certain assessments covering past events appear to confirm this. However, other factors may explain the inversion, particularly the asset purchase programme of the European Central Bank (ECB), whose accumulated stocks are driving down long-term rates. Therefore, the negative slope of the euro area yield curve tends to overstate the risk of recession. Besides, the ECB's staff economic projections for December 2023 estimated euro area GDP growth at 0.8% in 2024.

Matthieu Bussière, Stéphane Lhuissier
Monetary and Financial Analysis Directorate

JEL codes
C25, E37,
E43

The views expressed in this article are those of the authors and do not necessarily reflect those of the Banque de France or the Eurosystem. All errors and omissions are the responsibility of the authors.

-37 basis points
the spread between 10-year and 2-year government bond yields in Germany in November 2023

16.8%
the frequency of occurrence of yield curve inversion (negative slope) since 1970

8%
the one-year ahead probability of entering into recession, calculated using a statistical forecasting model based on the November 2023 yield curve, after taking into account financial conditions and the price of oil. Not taking into account these explanatory factors, the probability would be 22%

Slope of the yield curve in the euro area and Germany between January 1970 and November 2023
(percentage points)

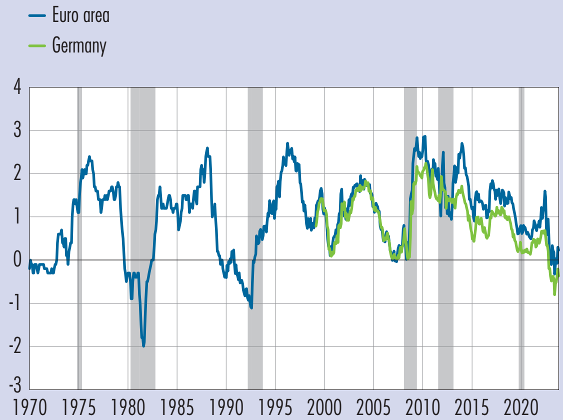

Sources: Bloomberg and authors' calculations.
Notes: The slope of the yield curve is calculated as the difference between 10-year and 2-year government bond yields. The shaded bands represent periods of recession in the euro area, as dated by the Centre for Economic Policy Research (CEPR).

In June 2023, the inversion of the yield curve (when 2-year government bond yields rise above the 10-year yields) stoked fears of a recession in the euro area. Despite a slight upturn in 2021 and 2022, the spread between long-term and short-term euro area bond yields, known as the "slope of the yield curve", has fallen into negative territory, and to its lowest level since the recession of 1992 (see Chart 1). This observation is particularly striking in the case of German bonds, which are more sensitive to key interest rate expectations and less susceptible to default risk. Indeed, the slope reached -37 basis points in November 2023.

The slope of the yield curve is highly volatile over time, but typically stays in positive territory because investors demand additional compensation for making long-term investments (a "term premium"). Therefore, a negative slope is relatively rare, but not unprecedented. It is conventionally perceived as a signal of recession. However, there is no consensus among economists as to the mechanisms behind this correlation, and the correlation itself is undermined by major exceptions, such as the sovereign debt crisis that rocked the euro area between 2010 and 2012 and the public health crisis of 2020: neither were preceded by an inverted yield curve. Furthermore, as the indicator is extremely volatile, assessing the risk of a recession becomes highly dependent on the date taken into consideration.

C1 Slope of the yield curve in the euro area and Germany between January 1970 and November 2023

(percentage points)

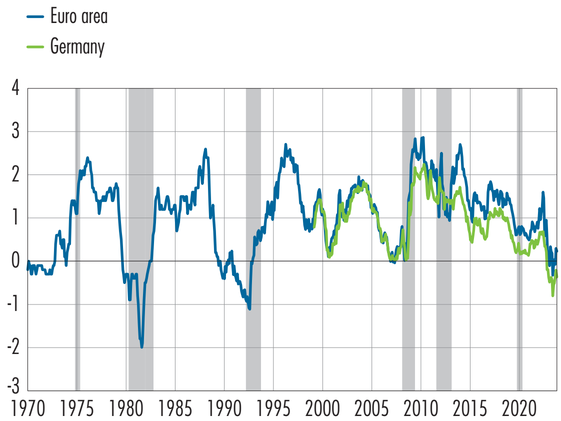

Sources: Bloomberg and authors' calculations.
Notes: The slope of the yield curve is calculated as the difference between 10-year and 2-year government bond yields. The shaded bands represent periods of recession in the euro area, as dated by the Centre for Economic Policy Research (CEPR).

For some time now, the academic literature has explored the link between the slope of the yield curve and the risk of recession using "probit" statistical models (Estrella and Hardouvelis, 1991, Estrella and Mishkin, 1998, Rudebusch and Williams, 2009, for the United States; Sabes and Sahuc, 2023, for the euro area). These studies show a high degree of statistical regularity between the yield curve and recession in the United States, but the link is not systematic for the euro area.

The yield curve inversions observed in the euro area since June 2023 also give us cause to re-examine this link. This bulletin therefore aims to review and explain forecasts of recession in the euro area based on the yield curve.

1 Breaking down the slope of the yield curve in accordance with expectations theory

There are different ways to measure the slope of the yield curve. Researchers generally use the difference between the 10-year government bond yield, which reflects investors' long-term outlook, and the 3-month bond yield, which offers a return close to the key interest rate set by the monetary authorities. Central bankers, on the other hand, prefer the difference between 10-year and 2-year government bond yields. Empirically, the two measures produce equivalent results in terms of forecasting recession (Bauer and Mertens, 2018).

It is important to remember that long-term interest rates primarily reflect the future path of short-term interest rates over the life of a bond expected by market participants, which naturally depends on their view of the coming trends in the business cycle and monetary policy. If the confidence of investors ebbs or if they fear a recession, they will probably expect the monetary authority to cut key interest rates to sustain demand and thus stabilise prices. Expectations of a downturn in future short-term interest rates therefore drive down long-term interest rates, which can result in an inversion of the yield curve. When market participants' predictions are accurate, the slope of the yield curve is associated with a higher probability of recession.

However, the slope of the yield curve is not determined by monetary policy expectations alone. It can also be affected by changes in market participants' attitudes to risk. To better understand its importance, we present a breakdown of a nominal bond yield, i_t^m for a given maturity m, into three components based on the expectations theory of the term structure of interest rates:

$$i_t^m = \underbrace{E_t \frac{1}{m}\left\{\sum_{j=0}^{m-1} r_{t+j}\right\}}_{\substack{\text{real interest}\\ \text{rate}\\ \text{expectations}\\ (r_A^m)}} + \underbrace{E_t \frac{1}{m}\left\{\sum_{j=0}^{m-1} \pi_{t+j}\right\}}_{\substack{\text{inflation}\\ \text{rate}\\ \text{expectations}\\ (\pi_A^m)}} + \underbrace{\Phi_t^m}_{\substack{\text{term}\\ \text{premium}}}$$

where E_t is the expectations operator at date t, r_t is the annualised real interest rate, π_t is the annualised inflation rate, and ϕ_t^m is the annualised term premium.

The first two parts of the breakdown show the expected path of the nominal interest rate over horizon m. Changes in long-term interest rates thus reflect the changes in the expected path of future real interest rates and inflation. The third part of the breakdown describes the risk premium, ϕ_t^m, which can also be further divided down into two subparts: the inflation risk premium $\Phi_{\pi,t}^m$ and the real risk premium $\Phi_{r,t}^m$, which reflect the additional return required by an investor to compensate for uncertainty with regard to inflation and real interest rate developments. Their sum captures the total compensation investors require to bear interest rate risk.

Using the above equation, we can break down the slope of the yield curve (defined as the difference between the long-term interest rate i_t^m, and the short-term interest rate, i_t^n, with $m > n$) into two expectation components and one risk premium component:

$$\underbrace{i_t^m - i_t^n}_{\substack{\text{slope}\\ \text{of the yield}\\ \text{curve}}} = \underbrace{r_A^m - r_A^n}_{\substack{\text{slope}\\ \text{of real interest}\\ \text{rate expectations}}} + \underbrace{\pi_A^m - \pi_A^n}_{\substack{\text{slope}\\ \text{of inflation rate}\\ \text{expectations}}} + \underbrace{\Phi_t^m - \Phi_t^n}_{\substack{\text{slope}\\ \text{of the term}\\ \text{premium}}}$$

Movements in each of these three components (real interest rate and inflation rate expectations, and the term premium) can therefore affect the total slope of the yield curve, without necessarily implying a recession.

2 Recession probability models based on the yield curve

Building on previous studies, our recession probability model uses the yield curve as an explanatory variable. More specifically, we derive a probit model as follows:

$$Pr\,(RECESSION_{t,t+12} = 1) = \Phi(\beta_0 + \beta_1\, TS_t),$$

where $RECESSION_{t,t+12}$ is an indicator variable (0 or 1) that is equal to 1 if a recession in the euro area, as dated by the Centre for Economic Policy Research (CEPR), is observed in $t+12$ [1] based on information available in t; TS_t is the interest rate differential between 10-year and 2-year government bonds; Φ is the cumulative normal distribution function; β_0 and β_1, respectively, represent the constant and the elasticity of the probability of recession to the slope of the yield curve (TS_t – term spread). If the slope of the yield curve does not allow us to predict recessions, then $\beta_1 = 0$ and only the constant β_0 will determine the probabilities, which will therefore remain invariable over time.

The model is built on January 1970 to December 2009 data only, so that "out-of-sample" testing can be carried out over the subsequent period (excluding the interval when key interest rates reached their lower bound and the slope of the yield curve was deemed unreliable). We then take the model's estimated parameters to generate recession probabilities for each month up to November 2023, with a view to producing an out-of-sample forecast, generally considered more reliable and credible.

The results confirm the slope's capacity to predict numerous recessions (see Chart 2, left-hand panel). Indeed, before most of the recessions that materialised, the one year ahead probability of a recession calculated using the statistical model increases sharply. However, there is a downward trend in the model's predictive power after the Great Recession of 2008-09 (even if trends are difficult to identify based on a relatively limited number of observations). These results corroborate the work of Sabes and Sahuc (2023).

1 The academic literature generally considers a forecast horizon of one year. For longer horizons, this type of model's performance is vastly inferior.

C2 Probabilities of recession in the euro area, between January 1970 and November 2023

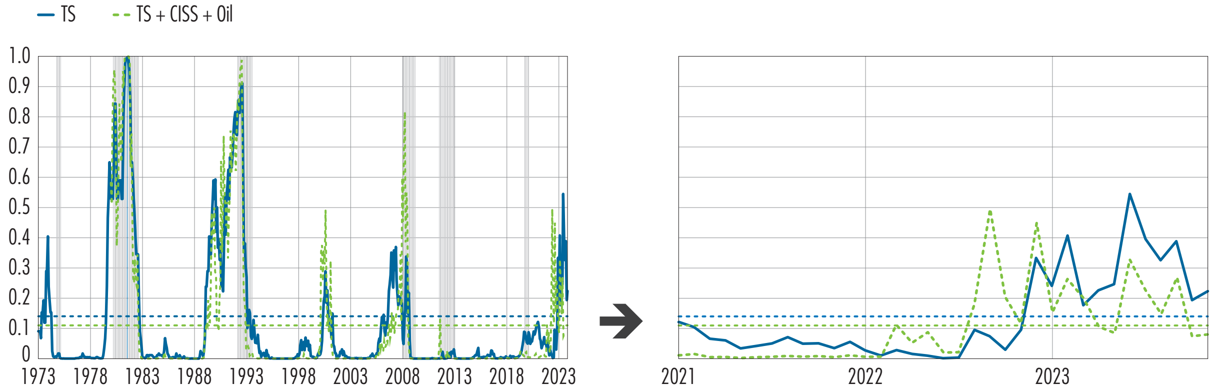

Sources: Bloomberg and authors' calculations.
Notes: Results from the probit model. Model 1 includes the slope of the yield curve (TS, term spread) only as an explanatory variable, while model 2 incorporates the price of oil (Oil) and a measure of financial conditions (CISS, the composite indicator of systemic stress). The shaded area represents periods of recession, as dated by the Centre for Economic Policy Research (CEPR). The dotted lines indicate the alert thresholds for each model (14% for model 1 and 11% for model 2).

A close look at the 2021-23 period reveals a sharp rise in the probability of recession from the end of 2022, from 0% to around 22% in November 2023 (see Chart 2, right-hand panel).

Nevertheless, this model remains far from perfect. On the one hand, the mechanisms that may explain the causal link between yield curve inversion and recession are still poorly understood. The generally accepted interpretation is based on economic agents' expectations: if they expect a recession one year ahead, they also expect key interest rates to fall sharply at the same moment, which results in long-term rates falling relative to short-term rates. This mechanism risks being a self-fulfilling prophecy in itself, whereby the inversion of the yield curve becomes a simple reflection of economic agents' expectations. On the other hand, an abundance of literature on recession risk estimation shows that incorporating additional information, such as financial conditions or the price of oil, into the model significantly improves its performance (see, for example, Gilchrist and Zakrajšek, 2012, and Favara et al., 2016). Our augmented model draws on Fonseca et al. (2023) to become:

$$Pr\,(RECESSION_{t,t+12} = 1) = \Phi(\beta_0 + \beta_1\,TS_t + \beta_2\,CISS_t + \beta_3\,OIL_t),$$

where $CISS_t$ (composite indicator of systematic stress) is a measure of financial conditions developed by Holló et al. (2012) and OIL_t is the Brent crude oil price in US dollars.

When financial conditions and the price of oil are taken into account, the estimated one-year ahead probability of recession reduces significantly compared to the previous model, decreasing by 14 percentage points to 8%. Overall, this result reflects the downward trend in the price of oil and the improvement in financial conditions seen over the past few months, which are boosting economic activity and thus reduce the risk of recession.

3 The yield curve: an important indicator, with caveats

We now compare the predictive performance of the two versions of the model described above using the receiver operating characteristic (ROC) curve. This metric evaluates the probability of detecting a recession – a "true positive" – against the probability of a false signal, or "false positive" (see appendix for further details).

Applying the ROC curve to both models (see Chart 3 below) allows us to visualise the potential distributions between true and false positives. The greater the area under the curve, the better the predictive capacity of the model. Unsurprisingly, the model that incorporates the slope of the yield curve, financial conditions and the price of oil is superior to the model that incorporates the slope of the yield curve as the only explanatory variable. And indeed, the area under the curve is around 0.83 for the first model, compared with 0.76 for the second. Estimates of probabilities of recession based solely on the yield curve

C3 ROC curves

(x-axis: rate of false positives; y-axis: rate of true positives; in %)

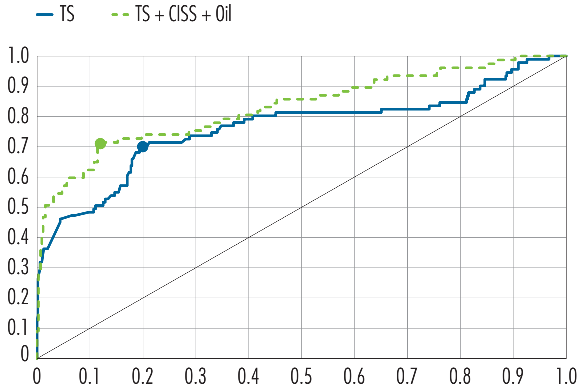

Source: Authors' calculations.
Notes: ROC – receiver operating characteristic.
This chart shows the ROC curve for version 1 of the model, which includes the slope of the yield curve (TS, term spread) as the only explanatory variable, and for version 2 of the model, which also incorporates the price of oil (Oil) and a measure of financial conditions (CISS, the composite indicator of systemic stress). The dots indicate the alert threshold. The estimation period stretches from January 1970 to December 2009 and the forecast period runs to November 2023.

remain inaccurate, as other factors, such as the price of oil and financial conditions, also play a part.

We also determine the optimal alert threshold (i.e. the point at which the ROC curve is closest to the upper left-hand corner of the chart). At this point, the ratio between the rate of true positives and the rate of false positives is maximised, i.e. it is the indicator's alert or trigger threshold. In the model incorporating all available information, this optimal value of 0.11 exceeds the probability of recession predicted for November 2023 of 0.08 (see Chart 2 above) suggesting that the model raises a weak signal of risk of a future recession. Nonetheless, the indicator is quite volatile and remains close to the trigger threshold.

In order to compare the predictive performance of the yield curve with that of other alternative economic indicators, we consider models that incorporate the following explanatory variables: (i) a measure of the growth rate of broad money (M3); (ii) the slope of the yield curve using German bond yields only; (iii) the price of oil; (iv) financial conditions; (v) the price of oil and financial conditions.

We measure the area under the ROC curve for each of the estimated models (see table). Models that include a yield curve among their explanatory variables produce the best predictive performances. The predictive capacities remain

Areas under the ROC curve

(between 0 and 1)

Forecast model	Area under the ROC curve
Slope of the yield curve (TS)	0.7639
TS + CISS +Price of oil in US dollars (Oil)	0.8324
Growth rate of the broad monetary aggregate (M3)	0.6042
Slope of the yield curve (TS) based on German bond yields	0.7661
Price of oil in US dollars (Oil)	0.6615
Indicator of financial conditions (CISS)	0.6574
Oil + CISS	0.6941

Source: Authors' calculations.
Notes: TS, term spread; CISS, composite indicator of systemic stress.
This table shows the area under the ROC curve derived from different probit models. The higher the value, the greater the predictive capacities of the model.

identical whether the curve is calculated for euro area sovereign bonds or for German sovereign bonds. Models that include financial conditions and the price of oil as the only explanatory variables also produce good results.

Although the yield curve offers the best predictive performance, its predictive capacities have probably declined in recent years due to the ECB's asset purchase programme conducted since 2015. This programme involves purchases of government, corporate, asset-backed and covered bonds to boost economic activity and inflation, and helped to substantially reduce long-term interest rates by narrowing risk premium. Despite ongoing reductions in this asset portfolio, the ECB still held a large quantity of government bonds in 2023, exerting a downward pressure on the term premium and therefore on the yield curve. Therefore, yield-curve based models tend to overstate the risk of recession in the euro area.

⁂

Due to its predictive capacity, the slope of the yield curve is often interpreted as a signal of recession. Our assessment for the euro area since 1970 demonstrates a statistical link between this slope and the risk of recession, but the link has probably weakened in recent years due to the ECB's asset purchase programme, which has had a significant impact on the term premium and therefore on long-term rates. Therefore, the negative slope of the yield curve observed in recent months in the euro area tends to be exaggerated, and consequently overstates the risk of recession.

References

Bauer (M. D.) and Mertens (Th. M.) (2018)
"Information in the yield curve about future recessions", *FRBSF Economic Letter*, No. 2018-20, Federal Reserve Bank of San Francisco, August.

Christiansen (Ch.), Nygaard Eriksen (J.) and Vinther Møller (S.) (2014)
"Forecasting US recessions: the role of sentiment", *Journal of Banking & Finance*, Vol. 49(C), December, pp. 459-468.

Estrella (A.) and Hardouvelis (G. A.) (1991)
"The term structure as a predictor of real economic activity", *The Journal of Finance*, Vol. 46, No. 2, June, pp. 555-576.

Estrella (A.) and Mishkin (F. S.) (1998)
"Predicting U.S. recessions: financial variables as leading indicators", *The Review of Economics and Statistics*, Vol. 80, No. 1, February, pp. 45-61.

Favara (G.), Gilchrist (S.), Lewis (K. F.) and Zakrajšek (E.) (2016)
"Recession risk and the excess bond premium", *FEDS Notes*, US Federal Reserve, April.

Fonseca (L.), McQuade (P.), Van Robays (I.) and Liliana Vladu (A.) (2023)
"The inversion of the yield curve and its information content in the euro area and the United States", *European Central Bank Economic Bulletin*, No. 7/2023, November, pp. 42-46.
Download the document

Gilchrist (S.) and Zakrajšek (E.) (2009)
"Credit spreads and business cycle fluctuations", *American Economic Review*, Vol. 102, No. 4, June, pp. 1692-1720.

Holló (D.), Kremer (M.) and Lo Duca (M.) (2012)
"CISS – A composite indicator of systemic stress in the financial system", *Working Paper Series*, No. 1426, European Central Bank, March.

Miller (D. S.) (2019)
"There is no single best predictor of recessions", *FEDS Notes*, US Federal Reserve, May.

Rudebusch (G. D.) and Williams (J. C.) (2009)
"Forecasting recessions: the puzzle of the enduring power of the yield curve", *Journal of Business & Economic Statistics*, Vol. 27, No. 4, June, pp. 492-503.

Sabes (D.) and Sahuc (J.-G.) (2023)
"Do yield curve inversions predict recessions in the euro area?", *Finance Research Letters*, Vol. 52, No. C, March.

Appendix
The ROC curve

Assessing recession probability models is complicated by the fact that the predicted variable is binary, whereas the value predicted by the models is a continuous variable. No model is perfect: type I errors (the model wrongly predicts a crisis) and type II errors (the model fails to predict a crisis present in the sample) can arise. This leads us to question the threshold at which the signal sent by the model (the probability of recession) should be taken seriously.

One way to do so is to use the receiver operating function, more commonly known as the "ROC (receiver operating characteristic) curve”. Christiansen, Nygaard Eriksen and Vinther Møller (2014), Miller (2019), and Sabes and Sahuc (2023) in particular have used this metric. It is also applied in medical research, for example, to measure the capacity of a biological sample in detecting a pathology: above a certain threshold, the diagnosis of pathology is in principle confirmed; below it, the patient is considered “healthy”.

The ROC curve can be used to determine the capacity of a model to categorise recessions and expansions appropriately, by calculating the rates of “false positives” (the proportion of events that the model wrongly characterises as recessions, i.e. the risk of type I errors in probabilities) and “true positives” (the proportion of recessions that are accurately detected). Represented graphically, the value of the area under the curve is used to estimate the accuracy of binary events, such as recessions and expansions. More specifically, the curve assesses the spectrum of different area thresholds (between 0 and 1) for determining a recession, rather than assessing the predictive power at a given threshold. A model that takes historical data and makes a perfectly accurate classification between recession and expansion would only have true positives (i.e. 0 on the x-axis and 1 on the y-axis – top left of the chart). In this perfect classification, the area under the ROC curve would be equal to one. In contrast, a model making random guesses (the equivalent of flipping a coin) would on average result in an equal number of true and false positives, with an area under the ROC curve of 0.5.

Published by
Banque de France

Managing Editor
Claude Piot

Editor-in-Chief
Claude Cornélis

Editor
Didier Névonnic

Translator/English Editor
Scott Oldale

Technical production
Studio Creation
Press and Communication

ISSN 1952-4382

To subscribe to the Banque de France’s publications
https://www.banque-france.fr/en/alertes/abonnements