Recession Forecasting with Credit and Oil Price Channels

Frank Schmid
frank.schmid@frankschmid.com

This version: 14 November 2026
First version: 13 March 2026

Abstract. This paper develops a recession forecasting model that distinguishes between credit-channel and oil-channel recessions. The model is a multinomial logit that predicts the probability of each recession type one quarter ahead, with exclusion restrictions under which credit-channel recession risk is driven by the excess bond premium (EBP) of Gilchrist and Zakrajšek and the term spread, while oil-channel recession risk is driven by crude oil price changes. Every series is constructed from its primary source. Two model changes are evaluated separately on quarterly data spanning 1965:Q2–2026:Q2. The Kauppi–Saikkonen dynamic extension—adding the lagged recession indicator—raises in-sample discrimination, measured by the area under the receiver operating characteristic (ROC) curve (AUC), and supports within-recession signaling. Replacing the Baa–Aaa corporate bond spread with the EBP raises discrimination and raises the pre-onset signal for the 2001 recession, by isolating the financial intermediation component that the raw spread mixes with expected default risk. The model identifies the recession channel consistent with a narrative classification for the majority of the National Bureau of Economic Research (NBER)-dated postwar episodes; the exceptions are documented. The term spread terms do not reproduce the inversion mechanism of the yield curve literature once the EBP and recession persistence are conditioned on, and a term-spread probit benchmark shows little out-of-sample discrimination over the evaluation window. On the most recent predictor values the model's forecast is continued expansion, with the oil channel carrying the larger part of the response to the 2026 rise in crude oil prices. The forecast is also stated over the twelve-month window that published recession probabilities use, obtained by iterating the one-quarter model in closed form, and evaluated on that definition. Magnitudes and test statistics are reported in the body and tables.

JEL classification: E32, E37, C25, C35.

This paper was developed with assistance from generative AI tools across the design, analysis, coding, and drafting stages. The author supervised all stages, verified the results, and takes full responsibility for the content.

Contents

Abbreviations and Acronyms iii

Greek Letters iii

1 Introduction 1

2 Literature Review 2

2.1 The Yield Curve as a Recession Predictor . . . 2

2.2 Dynamic Probit Models . . . 2

2.3 Credit Spreads and the Excess Bond Premium . . . 2

2.4 Oil Prices and Recession Risk . . . 3

3 Data 3

3.1 Series Construction . . . 3

3.2 Regime Classification . . . 4

3.3 Predictors . . . 5

4 Methodology 5

4.1 Multinomial Logit with Exclusion Restrictions . . . 5

4.2 Kauppi–Saikkonen Dynamic Extension . . . 5

4.3 The Excess Bond Premium . . . 6

5 Model Progression 6

5.1 From Static to Dynamic . . . 6

5.2 From Baa–Aaa to the Excess Bond Premium . . . 7

6 Results 7

6.1 Final Model Coefficients . . . 7

6.2 Pre-Recession Signal Strength . . . 8

6.3 Within-Recession Signaling . . . 9

7 Out-of-Sample Evaluation 10

8 Current Forecast 11

8.1 One Quarter Ahead . . . 11

8.2 The Twelve-Month Horizon . . . 12

9 Conclusion 15

Bibliography **16**

Abbreviations and Acronyms

AIC	Akaike Information Criterion
AR	Autoregressive
AUC	Area under the receiver operating characteristic curve
BEA	Bureau of Economic Analysis
BIC	Bayesian Information Criterion
BLS	Bureau of Labor Statistics
COVID-19	Coronavirus disease 2019
EBP	Excess bond premium
FRED	Federal Reserve Economic Data
GDP	Gross Domestic Product
GFC	Global Financial Crisis
JEL	Journal of Economic Literature
NBER	National Bureau of Economic Research
OOS	Out-of-sample
OPEC	Organization of the Petroleum Exporting Countries
pp	Percentage points
PPI	Producer Price Index
ROC	Receiver operating characteristic
S&L	Savings and loan
TS	Term spread

Greek Letters

β	beta (pronounced BAY-tuh); coefficient vector
δ	delta (pronounced DEL-tuh); coefficient on the lagged recession indicator

1. Introduction

This paper develops a recession forecasting model that distinguishes between credit-channel and oil-channel recessions. The model is a multinomial logit that predicts the probability of each recession type one quarter ahead, with exclusion restrictions under which credit-channel recession risk is driven by the excess bond premium (EBP) of Gilchrist and Zakrajšek[1] while oil-channel recession risk is driven by crude oil price shocks. We extend the static specification with the dynamic component of Kauppi and Saikkonen,[2] which adds a lagged recession indicator to capture recession persistence.

The paper makes four contributions. First, we classify the eight National Bureau of Economic Research (NBER)-dated postwar recessions as either credit-channel or oil-channel on narrative grounds and show that the predictors of each type differ. Credit-channel recessions (1969–70, 1981–82, 1990–91, 2001, 2008–09) are preceded by deterioration in credit market conditions, measured by the EBP. Oil-channel recessions (1973–75, 1980) are preceded by oil price shocks; the 2020 episode is assigned to the oil category as the supply-side outcome class, with the caveats stated in Section 3.

Second, we impose exclusion restrictions on the multinomial logit: credit market variables enter only the credit-channel equation, and oil price variables enter only the oil-channel equation. This structure gives each coefficient a direct economic interpretation and prevents the model from conflating the two channels.

Third, we evaluate two model changes separately. The Kauppi–Saikkonen dynamic extension raises the area under the receiver operating characteristic (ROC) curve (AUC) and supports within-recession signaling. Replacing the Baa–Aaa corporate bond spread with the EBP raises the AUC further and concentrates its gain in the 2001 episode, where the EBP isolates the financial intermediation channel from expected default risk.

Fourth, every series in the model is constructed from its primary source, and the construction is documented in Section 3 at the level of the retrieved series, the aggregation rule, and the transformation.

This paper is the third in a series.[3] The first demonstrates that the Federal Reserve's reaction function cannot be identified from aggregate time-series data. The second shows that supply-side indicators forecast inflation when the output gap does not. The present paper extends the supply-side framework to recession forecasting, using oil prices as the supply-side indicator and the EBP as the credit-channel indicator.

The remainder of the paper proceeds as follows. Section 2 reviews the four strands of literature on which the model rests. Section 3 describes the construction of each series from its primary source, the recession coding, and the narrative channel classification. Section 4 specifies the model. Section 5, the paper's principal section, reports the model progression: the Kauppi–Saikkonen extension raises the in-sample AUC from 0.890 to 0.926 with the Baa–Aaa spread held in place, and the replacement of that spread with the EBP raises it to 0.935, with the gain concentrated in the 2001 episode, whose one-quarter-ahead signal rises from 4.7 percent to 45.0 percent. Section 6 reports coefficients, pre-recession signals, and within-recession signaling, and documents that the

[1] Simon Gilchrist and Egon Zakrajšek, "Credit Spreads and Business Cycle Fluctuations," *American Economic Review* 102, no. 4 (2012): 1692–1720.

[2] Heikki Kauppi and Pentti Saikkonen, "Predicting U.S. Recessions with Dynamic Binary Response Models," *Review of Economics and Statistics* 90, no. 4 (2008): 777–791.

[3] Frank Schmid, "Observational Equivalence in Taylor Rule Estimation" (working paper, 2026); Frank Schmid, "Supply-Side Inflation Forecasting When the Phillips Curve Is Flat" (working paper, 2026).

term spread terms do not reproduce the inversion mechanism once the EBP and recession persistence are conditioned on. Section 7 reports the out-of-sample evaluation and its qualifications. Section 8 states the current forecast, and Section 9 concludes.

2. Literature Review

This section reviews four strands: the yield curve as a recession predictor, dynamic binary response models, credit spreads and the excess bond premium, and oil prices and recession risk.

2.1 The Yield Curve as a Recession Predictor

The work of Estrella and Hardouvelis and of Estrella and Mishkin established the term spread—the 10-year Treasury yield minus the 3-month Treasury bill rate—as a predictor of U.S. recessions.[4] The term spread's predictive content has weakened in recent decades. Johansson and Meldrum show that lower term premiums allow the yield curve to flatten without a corresponding increase in recession risk.[5] The 2007–09 Global Financial Crisis (GFC) was preceded by an inversion that had reversed by the onset quarter, and the persistent inversion of 2022–23 was not followed by an NBER-dated recession. Section 6 and Section 7 report where the term spread stands in the present specification.

2.2 Dynamic Probit Models

Dueker introduced the dynamic extension by including the lagged recession indicator.[6] Kauppi and Saikkonen developed a unified framework with four specifications: static, dynamic (adding the lagged recession indicator), autoregressive (adding the lagged model probability), and dynamic autoregressive (adding both).[7] They found that the dynamic model outperforms the static model for predicting recession duration. We adopt their dynamic specification (M2) as our preferred model.

2.3 Credit Spreads and the Excess Bond Premium

Gilchrist and Zakrajšek constructed a credit spread index from micro-level corporate bond data and decomposed it into an expected default component and a residual—the EBP.[8] They showed that the predictive content of credit spreads for future economic activity is due to the EBP, which reflects the risk-bearing capacity of the financial sector. Favara, Gilchrist, Lewis, and Zakrajšek used the EBP in a probit model to predict NBER-dated recessions.[9]

---

[4] Arturo Estrella and Gikas A. Hardouvelis, "The Term Structure as a Predictor of Real Economic Activity," *Journal of Finance* 46, no. 2 (1991): 555–576; Arturo Estrella and Frederic S. Mishkin, "Predicting U.S. Recessions: Financial Variables as Leading Indicators," *Review of Economics and Statistics* 80, no. 1 (1998): 45–61.

[5] Peter Johansson and Andrew Meldrum, "Predicting Recession Probabilities Using the Slope of the Yield Curve," *FEDS Notes* (Washington: Board of Governors of the Federal Reserve System, 2018).

[6] Michael J. Dueker, "Strengthening the Case for the Yield Curve as a Predictor of U.S. Recessions," *Federal Reserve Bank of St. Louis Review* 79, no. 2 (1997): 41–51.

[7] Kauppi and Saikkonen, "Predicting U.S. Recessions," 778–781.

[8] Gilchrist and Zakrajšek, "Credit Spreads," 1693–1699.

[9] Giovanni Favara, Simon Gilchrist, Kurt F. Lewis, and Egon Zakrajšek, "Recession Risk and the Excess Bond Premium," *FEDS Notes* (Washington: Board of Governors of the Federal Reserve System, 2016).

2.4 Oil Prices and Recession Risk

Hamilton documents that oil price increases preceded most postwar U.S. recessions and that the relation is nonlinear: large positive shocks carry disproportionate effects on economic activity.[10] We use the year-over-year percentage change in the Producer Price Index (PPI) for crude petroleum as the supply-side predictor. The series is available from Federal Reserve Economic Data (FRED) back to the 1940s and captures the energy-price channel of the 1973–75 and 1980 recessions.

3. Data

This section describes the construction of each series from its primary source, the sample, the dependent variable and its channel classification, and the predictors with their revision properties.

3.1 Series Construction

Every series is built from the primary release rather than from an intermediate record. Table 1 states the source, the retrieved series, and the transformation applied. Monthly series are aggregated to quarterly frequency by the arithmetic mean of the three months in the quarter; a quarter is retained only if all three months are present.

Table 1: Series Construction from Primary Sources

Variable	Source	Retrieved series	Transformation
Term spread	Board of Governors, H.15, via FRED	GS10, TB3MS	Quarterly means; 10-year Treasury yield minus 3-month Treasury bill rate
Excess bond premium	Board of Governors (Gilchrist–Zakrajšek)	`ebp_csv.csv`	Quarterly mean of monthly EBP; synthetic before 1973:Q1
Baa–Aaa spread	Moody's via FRED	BAA, AAA	Quarterly means; Baa yield minus Aaa yield
Crude oil price	BLS via FRED	WPU0561	Quarterly mean; year-over-year percentage change
Real GDP growth	BEA via FRED	A191RL1Q225SBEA	Quarterly, percent change at an annual rate
Recession regimes	NBER chronology	Business cycle chronology	Quarterly coding (Table 2)

FRED = Federal Reserve Economic Data; BLS = Bureau of Labor Statistics; BEA = Bureau of Economic Analysis. The NBER chronology is at `https://www.nber.org/research/data/us-business-cycle-expansions-and-contractions`*; series accessed 13 September 2026.*

The EBP of Gilchrist and Zakrajšek begins in 1973:Q1. For the earlier quarters we construct a synthetic EBP by standardizing the Moody's Baa–Aaa spread to the EBP distribution over the overlap period 1973:Q1–2026:Q2 (214 quarters). The Baa–Aaa spread is the Baa corporate bond

[10] James D. Hamilton, "Oil and the Macroeconomy Since World War II," *Journal of Political Economy* 91, no. 2 (1983): 228–248; James D. Hamilton, "What Is an Oil Shock?," *Journal of Econometrics* 113, no. 2 (2003): 363–398.

yield minus the Aaa yield; over the overlap it has mean 1.0566 percentage points (pp) and standard deviation 0.4313 pp. The EBP over the overlap has mean 0.0652 pp and standard deviation 0.5119 pp. The synthetic series shares the limitation that motivates the EBP in the first place—the Baa–Aaa spread mixes expected default risk with the intermediation component—so results for the 1969–70 episode rest on a proxy, and the synthetic values are marked where reported.

The estimation sample runs from 1965:Q2 through 2026:Q2 (245 quarterly observations): the target quarters for which a one-quarter-ahead forecast can be formed from predictors with complete four-quarter moving averages.

3.2 Regime Classification

The dependent variable is a three-category classification of the next quarter's regime: expansion (0), credit-channel recession (1), or oil-channel recession (2). Table 2 lists the recession quarters as coded.

Table 2: Classification of NBER Recessions

Episode	Recession quarters	Duration	Channel	Primary driver
1969–70	1969:Q4–1970:Q4	5	Credit	Monetary tightening
1973–75	1973:Q4–1975:Q1	6	Oil	OPEC oil embargo
1980	1980:Q1–1980:Q3	3	Oil	Oil price shock following the Iranian revolution
1981–82	1981:Q3–1982:Q4	6	Credit	Volcker disinflation
1990–91	1990:Q3–1991:Q1	3	Credit	S&L crisis; Gulf War
2001	2001:Q1–2001:Q4	4	Credit	Dot-com bust; corporate credit deterioration
2008–09	2008:Q1–2009:Q2	6	Credit	Global Financial Crisis
2020	2020:Q1–2020:Q2	2	Oil	COVID-19 pandemic (assigned; see text)

Duration in quarters. OPEC = Organization of the Petroleum Exporting Countries; S&L = savings and loan; COVID-19 = coronavirus disease 2019. Quarterly dating follows the NBER chronology (peaks 1969:Q4, 1973:Q4, 1980:Q1, 1981:Q3, 1990:Q3, 2001:Q1, 2007:Q4, 2019:Q4; troughs as listed). The coding includes the peak quarter for the first six episodes and begins in the quarter following the peak for the 2008–09 and 2020 episodes. The coded quarters constitute the estimation data (35 recession quarters in total). The NBER has dated no peak after 2019:Q4, so every quarter from 2020:Q3 through 2026:Q2 is coded as expansion. Sources: Board of Governors of the Federal Reserve System; Moody's and the Bureau of Labor Statistics via Federal Reserve Economic Data; Bureau of Economic Analysis; National Bureau of Economic Research; accessed 13 September 2026; author's calculations.

The channel classification is assigned on narrative grounds—the proximate cause of each downturn as documented in the business cycle literature—rather than from the predictor values, so that the model's channel identification in Section 6.2 is evaluated against labels that do not embed the predictors. Hamilton documents the oil-price origins of the 1973–75 and 1980 episodes.[11] The credit-channel labels rest on the monetary tightening of 1969, the Volcker disinflation, the savings and loan (S&L) crisis, the dot-com bust with the associated corporate credit deterioration, and the 2008–09 financial crisis. Two assignments carry caveats. The 1990–91 episode followed both the S&L crisis and the Gulf War oil price increase; the credit label rests on the narrative weight of the

[11] Hamilton, "Oil and the Macroeconomy," 228–230; Hamilton, "What Is an Oil Shock?," 363–365.

S&L crisis. The 2020 episode fits neither channel: the pandemic shock was exogenous to both credit conditions and oil prices, the EBP was negative before onset, and oil prices were falling. It is assigned to the oil category as the supply-side outcome class. The assignment determines which equation absorbs the two 2020 quarters in estimation; the model identifies neither channel in advance for this episode (Table 5).

3.3 Predictors

The predictors, all dated one quarter before the outcome, are as follows. Credit-channel equation: (i) term spread inverted portion, $\min(TS, 0)$; (ii) GDP growth negative portion, $\min(GDP, 0)$; (iii) term spread 4-quarter moving average; (iv) EBP positive portion, $\max(EBP, 0)$; (v) EBP 4-quarter moving average; (vi) lagged recession indicator (Kauppi–Saikkonen dynamic component). Oil-channel equation: (i) GDP growth negative portion; (ii) oil price year-over-year positive portion, $\max(Oil, 0)$; (iii) oil year-over-year 4-quarter moving average; (iv) lagged recession indicator.

The predictors are available at the forecast date; their revision properties differ. The term spread is constructed from market yields and is not revised. The PPI is subject to one scheduled revision four months after initial publication. The EBP is re-estimated when the Federal Reserve Board updates the underlying credit spread decomposition, so historical values change across releases. GDP growth is revised repeatedly; the estimation uses the most recent vintage, and Section 7 states the implication for the out-of-sample evaluation.

4. Methodology

This section specifies the static multinomial logit, adds the dynamic extension, and describes the credit variable.

4.1 Multinomial Logit with Exclusion Restrictions

Let y_{t+1} denote the regime in quarter $t+1$, taking values 0 (expansion), 1 (credit-channel recession), or 2 (oil-channel recession). The model specifies

$$P(y_{t+1} = j \mid X_t) = \frac{\exp(z_j(t))}{1 + \exp(z_1(t)) + \exp(z_2(t))}, \tag{1}$$

where $z_0(t) = 0$ (expansion is the base category), $z_1(t) = X_D(t)\beta_D$ is the credit-channel index function, and $z_2(t) = X_S(t)\beta_S$ is the oil-channel index function. The exclusion restrictions are enforced by construction: X_D contains only credit market and term spread variables, while X_S contains only oil price variables. No credit variable appears in the oil equation, and no oil variable appears in the credit equation.

4.2 Kauppi–Saikkonen Dynamic Extension

Adding the lagged recession indicator y_t captures recession persistence:

$$z_j(t) = X_j(t)\beta_j + \delta_j\, y_t \qquad \text{for } j = 1, 2, \tag{2}$$

where $y_t = 1$ if the economy is in either recession regime in quarter t, and 0 otherwise. The forecast for quarter $t+1$ thereby conditions on the regime in quarter t. Section 7 discusses the availability of y_t in real time.

4.3 The Excess Bond Premium

The Gilchrist–Zakrajšek credit spread index is constructed from micro-level corporate bond data. The credit spread is decomposed into an expected default component—estimated from a Merton distance-to-default model[12]—and a residual, the EBP. The EBP captures the component of credit spreads that reflects the risk-bearing capacity of the financial sector. The EBP reached +3.27 pp in 2008:Q4 and +1.59 pp in 2000:Q4, one quarter before the 2001 recession, but was negative before the 1973–75, 1980, and 2020 episodes—a pattern consistent with the exclusion restrictions.

5. Model Progression

This section evaluates the two model changes in sequence: from static to dynamic, and from the Baa–Aaa spread to the EBP.

5.1 From Static to Dynamic

The static multinomial logit with exclusion restrictions and the Baa–Aaa credit spread (M1) achieves an in-sample AUC of 0.890, an F1 score of 0.582, and a Brier score of 0.090. Adding the Kauppi–Saikkonen lagged recession indicator to both equations (M2), holding the credit variable at the Baa–Aaa spread, raises the AUC to 0.926, raises F1 to 0.783, and lowers the Brier score to 0.053. The likelihood ratio test statistic for the two added parameters is 49.82 on two degrees of freedom ($p < 10^{-10}$). The dynamic autoregressive extension (M4) adds the lagged model-implied recession probability; its likelihood ratio statistic relative to M2 is 3.85 on two degrees of freedom ($p = 0.146$), and both information criteria select M2. The same ordering holds with the EBP in place of the Baa–Aaa spread: the dynamic extension carries a likelihood ratio statistic of 46.45 on two degrees of freedom ($p < 10^{-10}$), and the dynamic autoregressive extension adds 2.61 on two degrees of freedom ($p = 0.272$). Table 3 reports the progression.

Table 3: Model Progression: Dynamics and the Credit Variable

Model	Credit variable	LL	k	AIC	BIC	AUC	F1	Brier
M1: Static	Baa–Aaa spread	−93.2	10	206.4	241.4	0.890	0.582	0.090
M2: Dynamic	Baa–Aaa spread	−68.3	12	160.6	202.6	0.926	0.783	0.053
M4: Dynamic AR	Baa–Aaa spread	−66.4	14	160.7	209.7	0.927	0.783	0.053
M1: Static	Excess bond premium	−87.1	10	194.1	229.1	0.897	0.622	0.082
M2: Dynamic (preferred)	Excess bond premium	−63.8	12	151.7	193.7	0.935	0.813	0.048
M4: Dynamic AR	Excess bond premium	−62.5	14	153.1	202.1	0.941	0.813	0.048

LL = log-likelihood; k = number of parameters; AIC = Akaike Information Criterion ($2k - 2\,\mathrm{LL}$); BIC = Bayesian Information Criterion ($k \ln 245 - 2\,\mathrm{LL}$); AUC = area under the ROC curve, computed as the rank statistic with ties averaged; F1 is evaluated at the threshold that maximizes it in sample; Brier = Brier score (lower is better). Rows with different credit variables are non-nested; no likelihood ratio test applies across them. Sources: Board of Governors of the Federal Reserve System; Moody's and the Bureau of Labor Statistics via Federal Reserve Economic Data; Bureau of Economic Analysis; National Bureau of Economic Research; accessed 13 September 2026; author's calculations.

[12] Robert C. Merton, "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance* 29, no. 2 (1974): 449–470.

5.2 From Baa–Aaa to the Excess Bond Premium

Replacing the Baa–Aaa spread with the EBP in the dynamic specification raises the log-likelihood from −68.3 to −63.8 at an unchanged parameter count. The two specifications are non-nested—one credit variable replaces another—so no likelihood ratio test applies; the information criteria order them (AIC 160.6 versus 151.7; BIC 202.6 versus 193.7), and every performance metric moves in the same direction: AUC from 0.926 to 0.935, F1 from 0.783 to 0.813, Brier from 0.053 to 0.048. The recession-by-recession picture locates the gain in the 2001 episode: the one-quarter-ahead signal rises from 4.7 percent to 45.0 percent. The EBP stood at +1.59 pp in 2000:Q4, reflecting the deterioration in financial intermediation after the dot-com bust, while the Baa–Aaa spread stood at 0.81 pp: the expected-default component was masking the intermediation signal.

6. Results

This section reports the final model coefficients, the pre-recession signal strength, and within-recession signaling.

6.1 Final Model Coefficients

Table 4 reports the coefficient estimates of the preferred specification.

Table 4: Final Model Coefficients (M2 with EBP)

Predictor	Credit channel	Oil channel
Constant	−2.879	−4.849
TS inverted portion	1.464	—
GDP negative portion	0.227	0.014
TS 4-quarter average	−0.975	—
EBP positive portion	2.309	—
EBP 4-quarter average	−0.182	—
Oil Y/Y positive portion (/100)	—	2.105
Oil Y/Y 4-quarter average (/100)	—	0.318
y_{lag} (dynamic component)	3.840	4.252

"—" indicates a coefficient constrained to zero by the exclusion restrictions. 12 parameters in total. TS = term spread; EBP = excess bond premium; Oil Y/Y = crude oil price year-over-year change. Sources: Board of Governors of the Federal Reserve System; Moody's and the Bureau of Labor Statistics via Federal Reserve Economic Data; Bureau of Economic Analysis; National Bureau of Economic Research; accessed 13 September 2026; author's calculations.

In the credit equation, the EBP enters the current positive portion with a coefficient of 2.309: elevated excess bond premiums raise the credit-channel index. The 4-quarter average of the EBP enters with a small negative coefficient (−0.182), so the contemporaneous reading rather than the trailing average carries the credit signal. The lagged recession indicator enters with a coefficient of 3.840, the largest term in the equation, which is the persistence mechanism of Section 6.3.

The two term spread terms enter with opposite signs. The inverted portion carries 1.464 and the 4-quarter average carries −0.975; because $\min(TS, 0)$ is nonpositive, a sustained inversion of 1 pp changes the credit index by −0.49, that is, it lowers the index. Conditional on the EBP and

on recession persistence, the term spread terms therefore do not reproduce the inversion mechanism of the yield curve literature reviewed in Section 2.1. Section 7 reports the corresponding out-of-sample result for a term-spread probit, which shows little discrimination over the evaluation window. The two findings are consistent with each other and with the term premium argument of Johansson and Meldrum; together they bound the contribution that the term spread makes to this specification.

The negative portion of GDP growth enters with small coefficients in both equations (0.227 credit, 0.014 oil); because $\min(GDP, 0)$ is nonpositive, the positive credit-equation coefficient implies that realized growth shortfalls lower the credit index by a small amount. Onset signals derive from the credit and oil price variables rather than from realized growth.

In the oil equation, oil price changes enter with coefficients of 2.105 on the current positive portion and 0.318 on the 4-quarter average, both scaled by 1/100: the contemporaneous component of oil price increases carries the larger weight in the index.

6.2 Pre-Recession Signal Strength

Table 5 reports the one-quarter-ahead forecast for each onset quarter, formed in the preceding quarter.

Table 5: Pre-Recession Signal (One Quarter Before Onset)

Episode	P(any)	P(Credit)	P(Oil)	EBP, pp	Oil Y/Y	Channel identified
1969–70	5.0%	4.1%	0.8%	−0.18†	+5.1%	Yes
1973–75	1.5%	0.5%	1.0%	−0.31	+12.9%	Yes
1980	3.5%	1.3%	2.2%	−0.13	+46.0%	Yes
1981–82	5.0%	2.5%	2.6%	+0.50	+54.0%	No
1990–91	4.9%	4.1%	0.8%	+0.15	−12.3%	Yes
2001	45.0%	43.8%	1.2%	+1.59	+36.3%	Yes
2008–09	10.1%	7.9%	2.2%	+0.30	+52.8%	Yes
2020	5.7%	4.9%	0.7%	−0.22	−6.7%	No

P(any), P(Credit), and P(Oil) are one-quarter-ahead forecasts for the onset quarter, formed in the preceding quarter; EBP and Oil Y/Y are the predictor values in that preceding quarter. Components may not sum to P(any) because of rounding. pp = percentage points. "Channel identified" states whether the larger of P(Credit) and P(Oil) matches the Table 2 classification. †Synthetic EBP (Baa–Aaa calibrated; Section 3.1). Sources: Board of Governors of the Federal Reserve System; Moody's and the Bureau of Labor Statistics via Federal Reserve Economic Data; Bureau of Economic Analysis; National Bureau of Economic Research; accessed 13 September 2026; author's calculations.

The model identifies the recession channel consistent with the narrative classification in six of eight episodes. The two exceptions are the 1981–82 episode and the 2020 episode. For 1981–82 the two channel probabilities are within 0.1 percentage point of each other (2.5 percent credit against 2.6 percent oil): oil prices were up 54.0 percent year over year in 1981:Q2 while the EBP stood at +0.50 pp, so both indicators carried signal and the oil equation held the marginally larger share. For 2020 neither indicator carried a signal: the EBP stood at −0.22 pp and oil prices were down 6.7 percent.

The pre-onset probability is below 11 percent for seven of the eight episodes. The model is a concurrent-plus-persistence predictor rather than a leading indicator; the exception is 2001, where the EBP reading of +1.59 pp one quarter before onset produces a signal of 45.0 percent.

Figure 1 plots the full history of the channel probabilities.

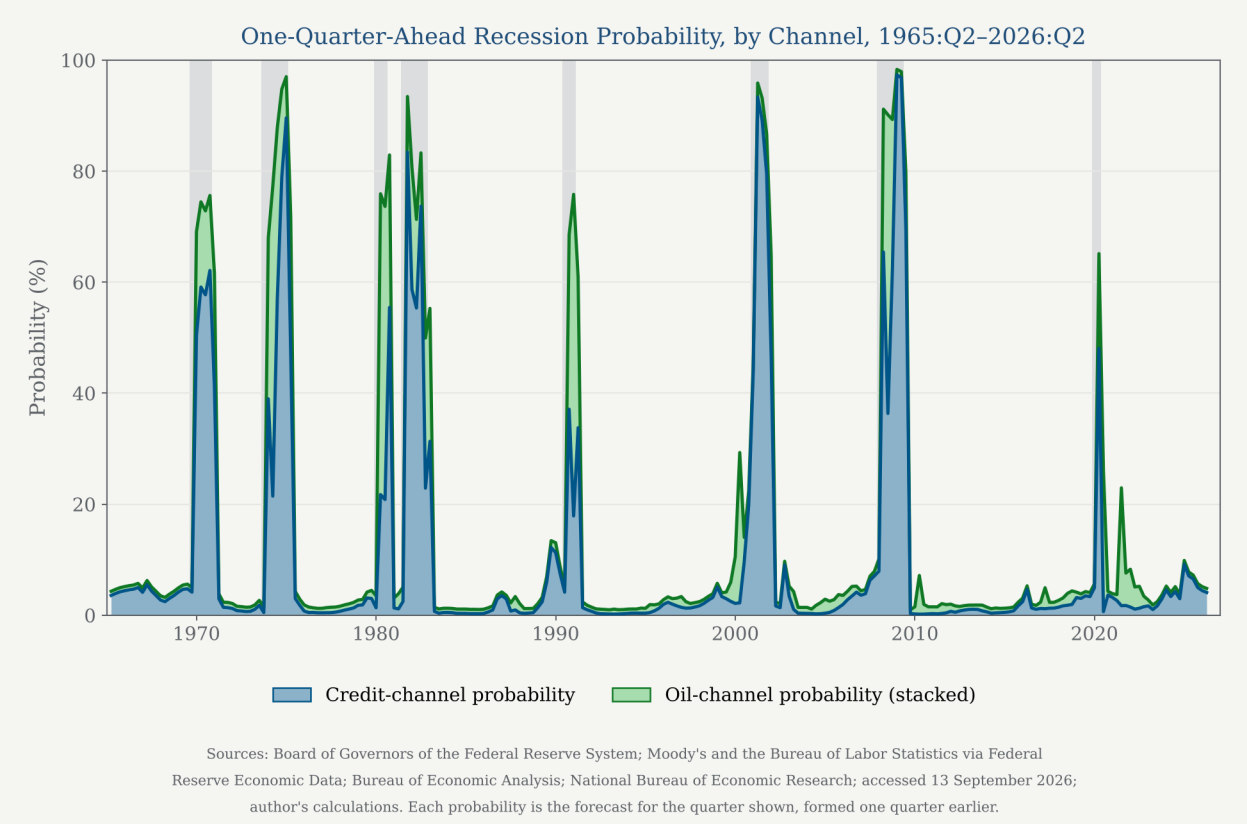


Figure 1: One-quarter-ahead recession probability by channel, 1965:Q2–2026:Q2. The horizontal axis is the forecast target quarter; the vertical axis is probability in percent. The lower band (blue fill with dark blue boundary) is the credit-channel probability; the upper band (green fill with dark green boundary) is the oil-channel probability, stacked on the credit band, so the upper boundary traces the total recession probability. Each probability is the model forecast for the quarter shown, formed one quarter earlier from the specification in Table 4. Gray vertical bands mark the recession quarters coded in Table 2. Sources: Board of Governors of the Federal Reserve System; Moody's and the Bureau of Labor Statistics via Federal Reserve Economic Data; Bureau of Economic Analysis; National Bureau of Economic Research; accessed 13 September 2026; author's calculations.

6.3 Within-Recession Signaling

The Kauppi–Saikkonen component supports within-recession signaling. For each of the 35 recession quarters t, we record the one-quarter-ahead probability formed in quarter t (the forecast for quarter $t+1$). This probability exceeds 50 percent in 33 of 35 recession quarters (94 percent; 32 of 33, or 97 percent, excluding the 2020 episode). The two quarters below the threshold are 1982:Q3, at 49.9 percent, and 2020:Q2, at 26.0 percent; the latter is the trough quarter of the two-quarter pandemic episode.

The 2008–09 sequence illustrates the mechanism. The forecast for the onset quarter, formed in 2007:Q4, stands at 10.1 percent. The forecasts formed in the six recession quarters are 91.2, 90.2, 89.3, 98.3, 97.9, and 79.8 percent (channel components may not sum to the totals because of rounding). The channel attribution within the episode tracks the indicators: the forecast formed in 2008:Q2—the quarter in which oil prices were still rising—assigns the larger share to the oil

channel (53.8 percent oil against 36.3 percent credit), while the forecasts formed in 2008:Q4 and 2009:Q1, with the EBP at +3.27 pp and +2.92 pp, assign 97.4 percent and 96.8 percent to the credit channel. The forecast formed in 2009:Q3—the first expansion quarter, for which the lagged indicator returns to zero—declines to 1.0 percent: the model leaves the recession state within one quarter of the trough.

7. Out-of-Sample Evaluation

We evaluate the model in a pseudo-out-of-sample (OOS) expanding-window exercise. The initial training window covers 1965:Q2–1989:Q4. Starting in 1990:Q1, we re-estimate the model each quarter and generate a one-quarter-ahead forecast. The exercise covers 146 predictions and four NBER recessions: 1990–91, 2001, 2008–09, and 2020.

Table 6: Out-of-Sample Performance (1990:Q1–2026:Q2)

Model	AUC	F1	Brier
Multinomial logit (M2 with EBP)	0.925	0.710	0.062
Probit: term spread only	0.547	0.219	0.104
GDP < 0 rule	—	0.516	0.103

The OOS period contains 146 predictions with 15 recession quarters (10.3%). F1 scores for the two probability models are computed at thresholds chosen ex post to maximize F1 over the OOS predictions; the GDP < 0 rule involves no threshold choice, so the F1 comparison favors the probability models by construction. AUC and the Brier score involve no threshold choice. "—" marks an AUC that is undefined: a deterministic rule produces no ranking. AUC, F1, and the Brier score carry the definitions given under Table 3. Sources: Board of Governors of the Federal Reserve System; Moody's and the Bureau of Labor Statistics via Federal Reserve Economic Data; Bureau of Economic Analysis; National Bureau of Economic Research; accessed 13 September 2026; author's calculations.

The multinomial logit retains discriminative power out of sample: the AUC declines from 0.935 in sample to 0.925. The term-spread probit reaches an AUC of 0.547 over the same window, close to the 0.5 of an uninformative ranking, and its Brier score of 0.104 exceeds that of the GDP rule. This is the out-of-sample counterpart of the coefficient finding in Section 6: over this window the term spread on its own does not order recession risk. The episodes documented in Section 2.1 bear on the result—the 2008–09 recession followed an inversion that had reversed before onset, and the 2022–23 inversion was followed by expansion.

The AUC reported here, like the in-sample one, is computed over every prediction in the window, the recession quarters among them. Section 8.2 reports the corresponding figure over expansion origins alone, which is the quantity a forecaster looking forward from an expansion faces, and it is lower.

The 2008–09 episode is the most severe recession in the OOS window. The first OOS forecast formed during the recession—at 2008:Q1, for 2008:Q2—assigns an 89.4 percent probability, and the forecasts formed in subsequent recession quarters range from 90.4 through 99.3 percent, with the channel taking the larger share from the forecast formed in 2008:Q4 onward. The term-spread probit remains between 5 and 12 percent throughout the episode.

Three qualifications bound the exercise. First, predictors enter at their most recent vintage; a fully real-time exercise would use the GDP vintage available at each forecast date. Because the GDP terms enter only through the negative portion, the distinction matters in quarters in which the initial and revised releases disagree in sign. Second, the lagged recession indicator is taken

from the final NBER chronology; NBER turning-point announcements trail the turning points by several months to more than a year, so the indicator's value for the most recent quarter is not always known at the forecast date. Third, as noted under Table 6, the F1 thresholds are selected ex post. All three features imply that the reported OOS performance bounds real-time performance from above.

8. Current Forecast

8.1 One Quarter Ahead

As of 2026:Q2, the model forecasts the probabilities in Table 7 for 2026:Q3. The configuration of the indicators differs from the preceding quarters in one respect: the PPI for crude petroleum stands 53.3 percent above its year-earlier level, a reading comparable to those that preceded the 1980 and 1981–82 episodes. The EBP is negative (−0.29 pp): corporate credit markets show no intermediation stress. The term spread, the 10-year Treasury yield minus the 3-month Treasury bill rate, is positive at 0.80 pp and has steepened over the past four quarters. The lagged recession indicator is zero.

Table 7: Current Forecast (2026:Q2 Predictors)

Variable	Value
P(Expansion, 2026:Q3)	94.2%
P(Credit-channel recession)	3.5%
P(Oil-channel recession)	2.3%
Excess bond premium, pp	−0.29
Term spread, pp	0.80
Oil Y/Y	+53.3%
Oil Y/Y, 4-quarter average	+7.1%
GDP growth	+1.5%
Lagged recession indicator	0

Predictor values dated 2026:Q2, computed from the sources of Table 1; pp = percentage points. Probabilities are computed from the coefficients in Table 4 and those predictor values. Sources: Board of Governors of the Federal Reserve System; Moody's and the Bureau of Labor Statistics via Federal Reserve Economic Data; Bureau of Economic Analysis; National Bureau of Economic Research; accessed 13 September 2026; author's calculations.

The oil-channel probability of 2.3 percent is the model's response to the oil price increase. It remains small for two reasons that the specification makes explicit. The oil equation weights the 4-quarter average as well as the contemporaneous reading, and that average stands at 7.1 percent because the three preceding quarters were flat or negative; and the lagged recession indicator, the largest term in either equation, is zero. The configurations that produced oil-channel probabilities above 20 percent before the 1980 episode combined a contemporaneous increase with a sustained one. Figure 2 places the current readings in the historical record.

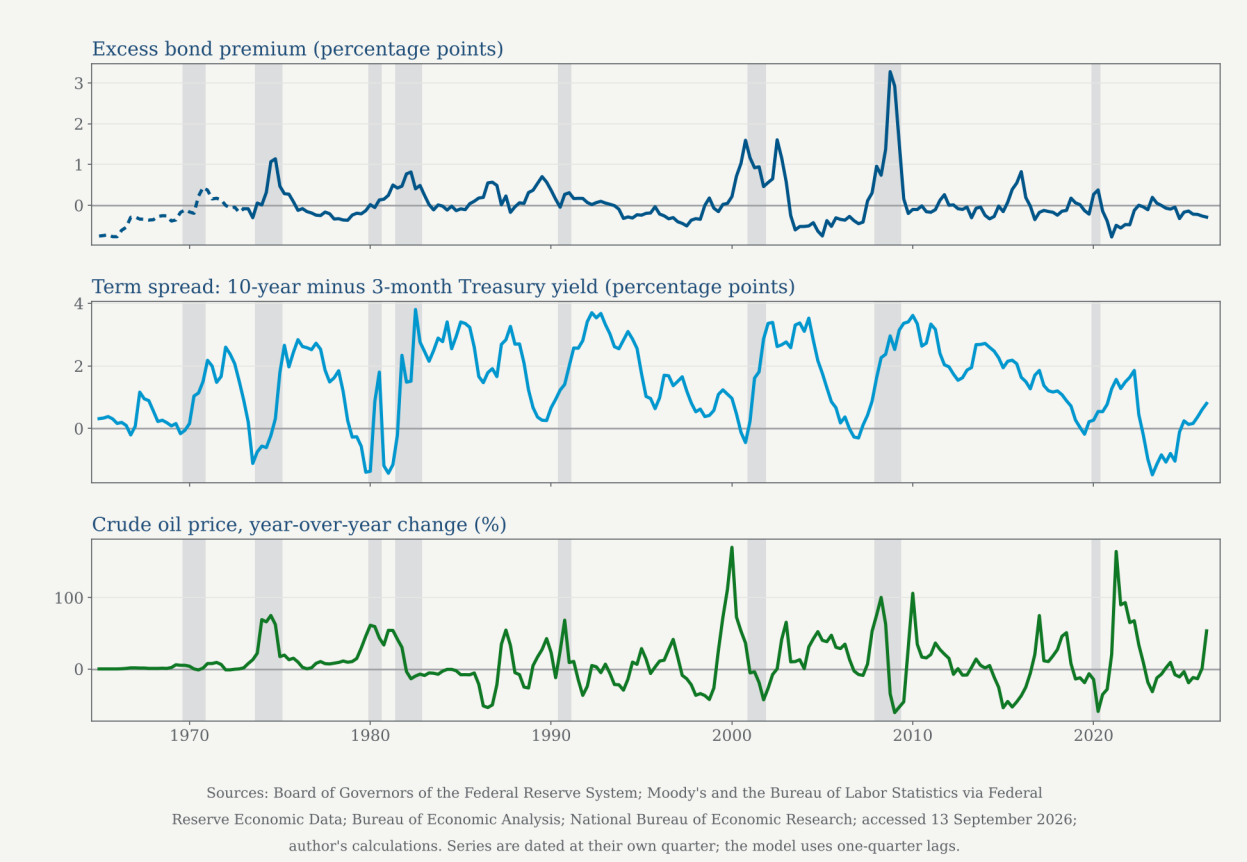


Figure 2: Model indicators, 1965:Q1–2026:Q2. Top panel: excess bond premium in percentage points (dark blue line); the dashed segment before 1973:Q1 is the synthetic series constructed from the Baa–Aaa spread (Section 3.1). Middle panel: term spread, the 10-year Treasury yield minus the 3-month Treasury bill rate, in percentage points (light blue line). Bottom panel: crude oil price, year-over-year change in percent (dark green line). Series are dated at their own quarter; the model uses one-quarter lags of each series. The horizontal reference line in each panel marks zero. Gray vertical bands mark the recession quarters coded in Table 2. Sources: Board of Governors of the Federal Reserve System; Moody's and the Bureau of Labor Statistics via Federal Reserve Economic Data; National Bureau of Economic Research; accessed 13 September 2026; author's calculations.

8.2 The Twelve-Month Horizon

Published recession probabilities are generally stated over a twelve-month window that includes the present rather than for a single quarter. The Wall Street Journal's Economic Forecasting Survey, conducted quarterly among a panel of more than 70 academic, business, and financial economists, asks for the probability that the economy is in recession at some point in the next twelve months, including the present, and publishes the panel average.[13] A one-quarter-ahead probability answers a different question over a quarter of the horizon, and the two are not comparable as printed. This section states the model's forecast on the twelve-month definition and evaluates it there.

[13] Anthony DeBarros, "About The Wall Street Journal Economic Forecasting Survey," *Wall Street Journal*, updated October 11, 2026.

Because the lagged recession indicator is the only dynamic term and enters as an indicator of either recession regime, iterating the one-quarter model over four quarters is a two-state chain whose transition probabilities are those of Equation (1). The cumulative probability has a closed form: no simulation is involved, no parameter is added, and the coefficients are those of Table 4. Writing $p_0(t)$ for the expansion probability of the first step, taken from the regime observed at t, and $q_0(t)$ for that of a later step, which is taken from expansion and so carries $y = 0$,

$$P(\text{no recession quarter in } t{+}1, \ldots, t{+}4 \mid X_t) = p_0(t)\, q_0(t)^3, \tag{3}$$

and the complement is the reported probability. The channel split is the first-passage decomposition: the probability that the first recession quarter of the window is a credit-channel one, and the same for the oil channel. The two sum to the cumulative probability by construction.

Equation (3) holds the predictors at their values at t over the window. This is the implementable rule: a forecaster at t knows them and knows nothing later. Reading the predictors at each later quarter from the record instead is a diagnostic no forecaster can use; over the 242 origins where both are defined it raises the mean reading from 23.1 percent to 24.4 percent, and the largest single disagreement is 30.1 percentage points, at the origin of 2000:Q2, where the held rule gives 45.3 percent and the realized path 75.5 percent. The values reported below are those of the held rule.

Table 8: Twelve-Month Recession Probability (2026:Q2 Predictors)

Variable	2026:Q2	2026:Q1
P(recession quarter within four quarters)	21.2%	17.9%
first entry through the credit channel	12.7%	15.2%
first entry through the oil channel	8.5%	2.7%
P(recession quarter in the next quarter alone)	5.8%	4.8%

Computed from Equation (3) and the coefficients of Table 4, with the predictors held at the values of the quarter named. Components sum to the total by construction and may not add in the displayed rounding. The 2026:Q1 column is the preceding quarter's reading on the same definition, reported beside the current one because a probability of this kind is read against its own last value. Sources: Board of Governors of the Federal Reserve System; Moody's and the Bureau of Labor Statistics via Federal Reserve Economic Data; Bureau of Economic Analysis; National Bureau of Economic Research; accessed 13 September 2026; author's calculations.

On the twelve-month definition the model's reading for the four quarters following 2026:Q2 is 21 percent, against 18 percent a quarter earlier. The composition moved more than the total: the credit-channel share fell from 15.2 to 12.7 percent as the EBP stayed negative, and the oil-channel share rose from 2.7 to 8.5 percent on the crude oil increase of Section 8.1. For comparison, the panel of the Wall Street Journal survey of 2 through 7 October 2026 put the probability at 24 percent, against 25 percent in its July survey.[14] The comparison is between a model conditioned on four indicators and a panel average over 71 forecasters, and it is offered as a placement of the model's reading rather than as a test of either.

Figure 3 plots the twelve-month probability over the sample.

[14] Matt Grossman and Anthony DeBarros, "Why Economists Think Higher Bond Yields Are Here to Stay," *Wall Street Journal*, October 11, 2026.

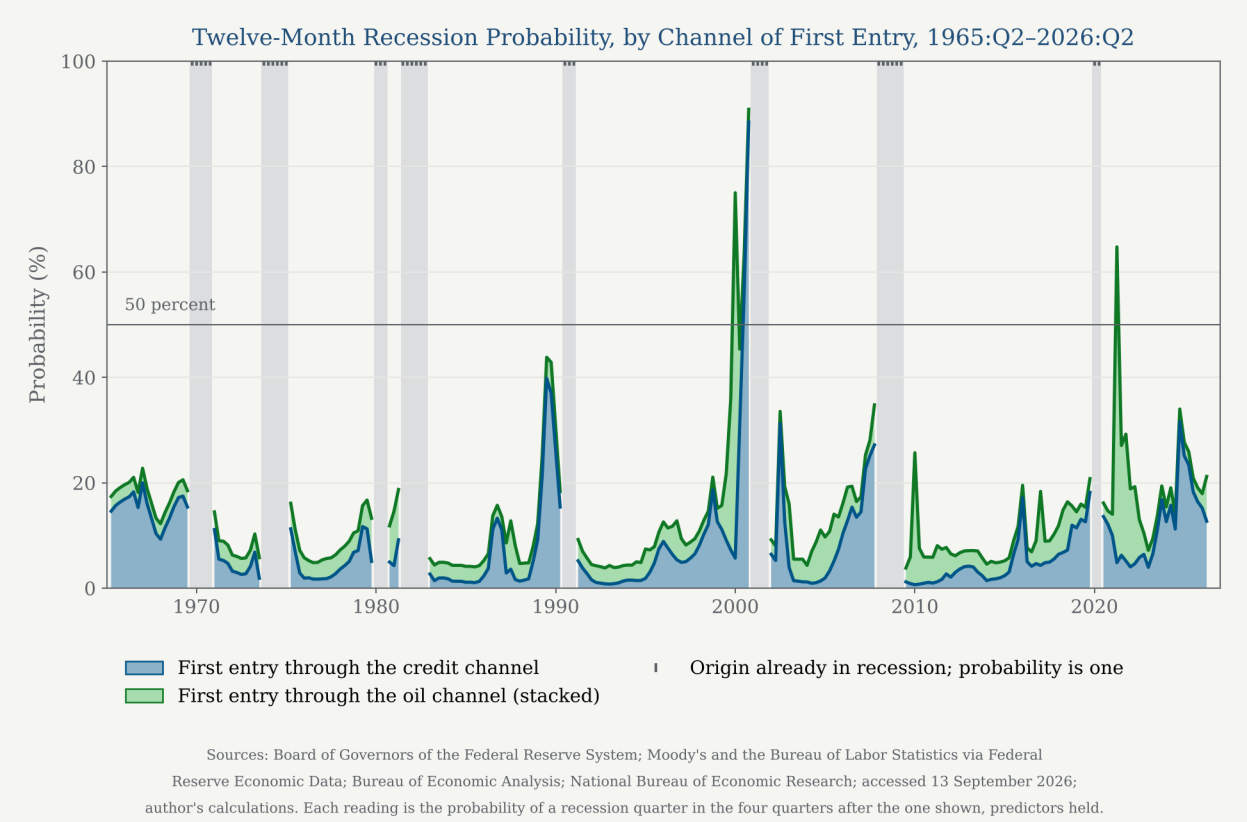


Figure 3: Twelve-month recession probability by channel of first entry, 1965:Q2–2026:Q2. The horizontal axis is the origin quarter; each reading is the probability that a recession quarter occurs in the four quarters following it, with the predictors held at the origin's values. The lower band (blue fill with dark blue boundary) is the probability that the first recession quarter of the window is a credit-channel one; the upper band (green fill with dark green boundary) is the same for the oil channel, stacked, so the upper boundary traces the cumulative probability. The horizontal reference line marks 50 percent, the threshold of Section 6.3. Tick marks at 100 percent mark origins already inside a recession, where the quantity is one by construction and the band is not drawn. Gray vertical bands mark the recession quarters coded in Table 2. Sources: Board of Governors of the Federal Reserve System; Moody's and the Bureau of Labor Statistics via Federal Reserve Economic Data; Bureau of Economic Analysis; National Bureau of Economic Research; accessed 13 September 2026; author's calculations.

The twelve-month figure is evaluated on its own definition over the 206 origins that are expansion quarters and whose four-quarter window is fully observed; 31 of those windows, 15.0 percent, contain a recession quarter. The AUC is 0.802 and the Brier score 0.105. Over the same 206 origins the one-quarter-ahead probability of Table 4's specification reaches an AUC of 0.772. Both are below the in-sample AUC of 0.935 reported in Table 3, and the difference is the subject of Section 6.3: that figure is computed over all 245 observations, the 35 recession quarters included, where the lagged recession indicator carries the discrimination. Restricted to the forecasting problem—an expansion quarter, looking forward—the model orders recession risk at 0.80 rather than at 0.94. The largest reading not followed by a recession quarter is 64.7 percent, at the origin of 2021:Q2, where crude oil prices stood above their pandemic-depressed level of a year earlier; 1999:Q4 at 36.0 percent and 2024:Q4 at 34.0 percent follow it.

9. Conclusion

This paper develops a recession model that decomposes recession risk into credit-channel and oil-channel components. The credit-channel equation, driven by the Gilchrist–Zakrajšek EBP and the term spread, captures recessions associated with financial intermediation stress and monetary policy tightening. The oil-channel equation, driven by crude oil prices, captures recessions associated with energy price shocks. Exclusion restrictions prevent conflation of the two channels, so that every coefficient has a direct economic interpretation.

The analysis yields four findings. First, the Kauppi–Saikkonen dynamic extension raises the AUC from 0.890 to 0.926 with the Baa–Aaa spread held in place (likelihood ratio statistic 49.82, $p < 10^{-10}$), and supports within-recession signaling: the one-quarter-ahead probability formed during recession quarters exceeds 50 percent in 33 of 35 quarters. Second, replacing the Baa–Aaa spread with the EBP raises the AUC further, to 0.935, and raises the 2001 pre-onset signal from 4.7 percent to 45.0 percent by isolating the financial intermediation channel from expected default risk. Third, the model's channel identification matches the narrative classification in six of eight episodes; the exceptions are the 1981–82 onset, where the two channel probabilities differ by 0.1 percentage point, and the 2020 episode, an exogenous pandemic shock outside the scope of financial and oil price indicators. Fourth, the term spread terms enter with signs that do not reproduce the inversion mechanism once the EBP and recession persistence are conditioned on, and a term-spread probit reaches an out-of-sample AUC of 0.547 over 1990:Q1–2026:Q2; the credit-channel signal in this specification is carried by the EBP.

Stated on the twelve-month definition that published recession probabilities use—a recession quarter at some point in the next four, the current quarter included—the model's reading for the four quarters after 2026:Q2 is 21 percent, against 18 percent a quarter earlier, with the oil channel accounting for 8.5 of the 21.2 percentage points against 2.7 of 17.9. On that definition, over the 206 expansion origins whose window is fully observed, the model reaches an AUC of 0.802 against a base rate of 15.0 percent.

The model's limitations follow from its structure. The pre-onset signal one quarter before onset is below 11 percent for seven of the eight episodes; the model is a concurrent-plus-persistence predictor rather than a leading indicator. The persistence mechanism conditions on the lagged NBER indicator, which is announced with a delay of several months to more than a year, and the evaluation uses final-vintage predictor data; both features imply that reported performance bounds real-time performance from above. The channel classification is a narrative assignment; the 1990–91 and 2020 labels carry the caveats stated in Section 3.2.

Bibliography

Dueker, Michael J. "Strengthening the Case for the Yield Curve as a Predictor of U.S. Recessions." *Federal Reserve Bank of St. Louis Review* 79, no. 2 (1997): 41–51.

Estrella, Arturo, and Gikas A. Hardouvelis. "The Term Structure as a Predictor of Real Economic Activity." *Journal of Finance* 46, no. 2 (1991): 555–576.

Estrella, Arturo, and Frederic S. Mishkin. "Predicting U.S. Recessions: Financial Variables as Leading Indicators." *Review of Economics and Statistics* 80, no. 1 (1998): 45–61.

Favara, Giovanni, Simon Gilchrist, Kurt F. Lewis, and Egon Zakrajšek. "Recession Risk and the Excess Bond Premium." *FEDS Notes*. Washington: Board of Governors of the Federal Reserve System, 2016.

Gilchrist, Simon, and Egon Zakrajšek. "Credit Spreads and Business Cycle Fluctuations." *American Economic Review* 102, no. 4 (2012): 1692–1720.

Hamilton, James D. "Oil and the Macroeconomy Since World War II." *Journal of Political Economy* 91, no. 2 (1983): 228–248.

Hamilton, James D. "What Is an Oil Shock?" *Journal of Econometrics* 113, no. 2 (2003): 363–398.

Johansson, Peter, and Andrew Meldrum. "Predicting Recession Probabilities Using the Slope of the Yield Curve." *FEDS Notes*. Washington: Board of Governors of the Federal Reserve System, 2018.

Kauppi, Heikki, and Pentti Saikkonen. "Predicting U.S. Recessions with Dynamic Binary Response Models." *Review of Economics and Statistics* 90, no. 4 (2008): 777–791.

Merton, Robert C. "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates." *Journal of Finance* 29, no. 2 (1974): 449–470.

Schmid, Frank. "Observational Equivalence in Taylor Rule Estimation." Working paper, 2026.

Schmid, Frank. "Supply-Side Inflation Forecasting When the Phillips Curve Is Flat." Working paper, 2026.