March 4, 2016

Probabilities of the U.S. Economy Entering a Recession in the Coming Year

Travis Berge, Nitish Sinha, and Michael Smolyansky

Executive Summary

The increase in financial market volatility and softening in some economic indicators since the beginning of the year has led many market observers to question whether the U.S. economy is at a heightened risk of entering a recession in the near future. In recent weeks, Wall Street economists have circulated a number of statistical estimates of recession probabilities. These estimates are based on a wide range of indicators, both macroeconomic and financial. While many indicators are potentially informative, evidently, different specific indicators and analysis can yield materially different conclusions.

To better understand the empirical issue and highlight the uncertainties surrounding forecasts of this nature, we employ a statistical technique, Bayesian Model Averaging (BMA), able to incorporate information from a wide range of economic and financial variables to produce estimates of the probability that the U.S. economy transitions into a recession at some point over the next 12 months. BMA combines a large number of potential forecasting models to produce a forecast that is a weighted average of each individual model forecast. The weights given to each forecast model are related to that model's ability to explain previous recessions. In this way, economic indicators that have not anticipated past NBER-dated recessions are downweighted, given little or no weight in the final forecasting model.

A key advantage of this approach is that it allows for different sets of variables to be informative about recession risks at different forecast horizons. In particular, our analysis reveals that while certain indicators of real economic activity and some financial variables have forecasting power for predicting recessions at the 3-month horizon, real variables have considerably less information content at the 12-month horizon. In contrast, forward-looking financial variables, primarily corporate bond credit spreads and

the slope of the Treasury yield curve, maintain considerable predictive power about the risk of recession 12 months hence.

We estimate that the current probabilities of the U.S. economy being in a recession 3, 6, or 12 months from now are 13, 18 and 16 percent, respectively. At all horizons, the estimated recession probabilities are about in line with the unconditional probability of the economy being in recession, 15 percent since 1973. Notably, however, the probability of recession at each horizon has increased since December 2015. This increase in recession risk has been driven by a deterioration in financial indicators that have predictive power within our model – in particular, a flattening of the Treasury yield curve and a widening of corporate credit spreads. Similarly, the uncertainty surrounding the estimated recession probabilities has also increased since December 2015; for example, the 68-percent confidence interval around the 12-month ahead recession probability has increased to a range of 0 to 50 percent.

In the final section of the memo, we evaluate the performance of the BMA methodology by performing a pseudo out-of-sample forecasting exercise. Specifically, we end the model estimation prior to the onset of both the 2001 and 2008 recessions to determine what the model-implied recession probabilities were ahead of these recessions. Six months prior to the 2001 (2008) recession the model assigned about a 30 (25) percent probability that the economy would be in recession in 6 months—not remarkably strong signals, but noticeably higher than the current 6-month-ahead reading.

Modeling recession probabilities

Our objective is to predict a binary outcome: will the economy be expanding or contracting at a particular date in the future, given our knowledge of the world today? Clearly, there is no single indicator, or even fixed set of indicators, that contains comprehensive information about the state of the economy 3, 6, or 12 months from now. The economic outlook is usually mixed, with different indicators pointing in different directions. Further, there is no reason to expect any set of indicators to be equally predictive about the macroeconomic state prevailing at different forecast horizons.

For these reasons, we consider a fairly large number of possible recession indicators. Specifically, we consider a set of 17 monthly variables chosen to describe different aspects of the economy: broadly speaking, labor market indicators, measures of real economic activity, and forward-looking financial variables such as equity returns, credit spreads, the Treasury yield curve and indicators of financial market stress.[1] Our dataset begins in January 1973 and continues through February 2016.[2]

With these 17 indicators in hand, the econometric problem becomes choosing between a very large number of potential forecasting models. We use BMA to elicit the best model at each forecast horizon.[3] To be specific, the model is a weighted average of a suite of static probit regressions that use NBER recession dates as the dependent variable. In each probit regression, the dependent variable Y_t denotes the binary state of the economy: $Y_t = 1$ if the NBER has declared that month t falls in a recession, and $Y_t = 0$ if the NBER has declared that month t falls in an expansion instead. Accordingly, each model estimates the probability that month t will be declared a recession by the NBER with following equation:

$$\Pr(Y_t = 1 \mid x_{t-h}) = \Phi(\alpha + x'_{t-h}\beta) \qquad (1)$$

where Φ is the cumulative standard normal probability distribution.

In our setup, we have many regressions that take the form of equation *(1)*, one model for every possible combination of the 17 indicators. BMA estimates the probability that the NBER will declare a month to be a recession from each regression, and then calculates a weighted average of these estimates. For each model, let $\hat{p}_{it}$ denote

[1] See the appendix for a complete list of indicators included in our data set.
[2] Many macroeconomic indicators lack February values at the time of writing, March 4, 2016 (though we were able to incorporate today's labor market data). For these indicators, we replace February values with Board staff estimates thereof.
[3] Leamer (1978) introduced Bayesian Model Averaging to the economics literature. More recently, many other applications have appeared in the economics literature. For example, Wright (2009) uses BMA to forecast future inflation; Piger and Morley (2008) use it to model trends and cycles in U.S. output; Faust, Gilchrist, Wright & Zakrajsek (2013) employ BMA using a large number of financial indicators to forecast real-time measures of economic activity. This memo draws heavily on the analysis of BMA forecasts for recession probabilities in Berge (2015).

the predicted probability of recession from model i, where $i = 1, \ldots, 2^k$.[4] The Bayesian model average forecast is the weighted sum:

$$\hat{p}_t = \sum_{i=1}^{2^k} \hat{p}_{it} \, Pr(M_i \mid D_{t-h}) \qquad (2)$$

where $\hat{p}_{it}$ denotes each individual forecast. The weights, $Pr(M_i \mid D_{t-h})$, denote the Bayesian posterior probability of model i given the data at time $t\text{-}h$. Combinations of variables that produce recession probabilities that match the actual NBER dates at each forecast horizon receive larger weight in the average of the model forecasts, while models that cannot anticipate recessions receive little or zero weight in the averaged forecast.

Model fit

Figure 1 plots the fit of the BMA model for forecast horizons of 3 and 12 months. The estimated recession probabilities generally rise during NBER-defined recessions, the shaded areas.[5] However, the model produces several false positives—periods where the estimated probability rises, but no recession occurs. For example, although the 3-month ahead recession probability (left panel) spikes immediately following the 1987 stock market crash, the economy continued to expand. The 12-month ahead recession probability (right panel) similarly rises somewhat in the late 1990s, well before the 2001 recession.

In addition, the final values in the plots of Figure 1 indicate that the estimated probability of recession has risen notably in the past few months. The forecast that February 2017 will be declared a recession by the NBER currently stands at just above 16 percent, essentially equal to the unconditional recession probability over this sample period (dashed blue line).

[4] We use the method of Raftery (1995) to perform BMA, which approximates the posterior likelihood of each model with a maximum likelihood estimate of its Bayesian Information Criterion.
[5] We evaluate the fit of each model more formally in the appendix.

Figure 1. Estimated probabilities from the BMA model, three- and twelve-month ahead forecasts.

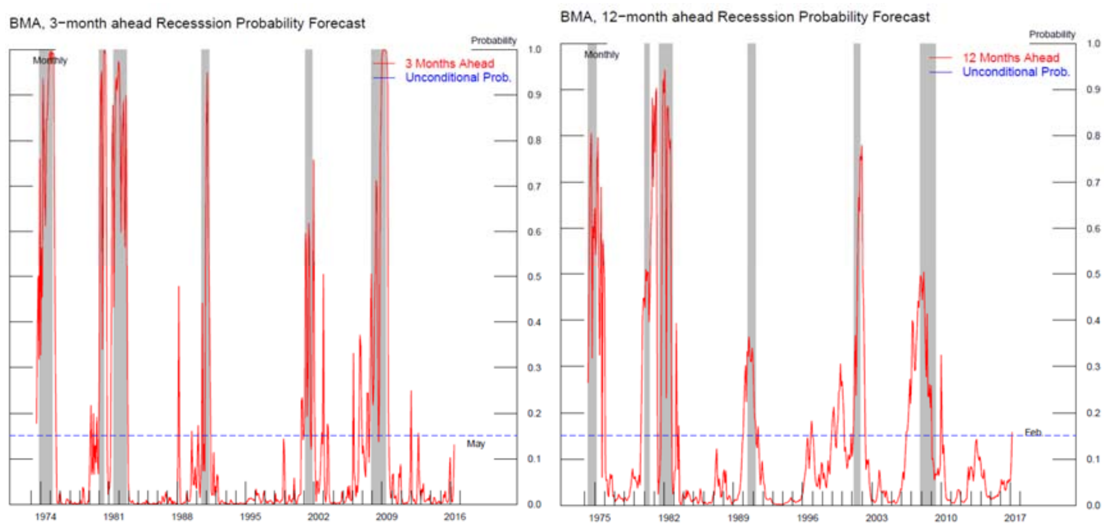

Note: Figures show $Pr(Y_t \mid x_{t-h})$, so that the date on the x-axis is the forecast that month *t* was declared recession given information *t-h* months prior. NBER recession dates shaded.

Table 1 gives the estimated regression coefficients produced by averaging across all models at each forecast horizon. The table shows variables for which the slope coefficient β, weighted across all models, is non-zero at the 90-percent confidence level, according to the *posterior inclusion probability*.[6] At the 3-month horizon, both real and financial variables are included as informative indicators of recession. The 3-month change in nonfarm payroll employment, real personal consumption expenditures, as well as the return on the S&P 500 index and the slope of the yield enter the model significantly with nonzero weight. However, as the forecast horizon lengthens, variables describing real economic activity drop out of the model, and financial indictors gain prominence. Indeed, the model that forecasts recession 12 months from now depends heavily on only two variables: the slope of the yield curve and the GZ credit spread index, a measure of credit market conditions that is described in detail in the companion memo by Favara, Lewis, and Suarez.

[6] The *posterior inclusion probability* can be thought of as the probability, given the data, that the "true" model includes a particular variable.

Table 1: BMA relies on different indicators at each forecast horizon.

Panel A:	3-month horizon		
	Posterior inclusion probability (%)	*Coef.*	*std. err.*
Change in payroll employment	100	-2.9	0.60
Slope of yield curve	100	-0.9	0.19
S&P 500 (3-month return)	100	-0.1	0.03
Real PCE	95	-1.0	0.39

Panel B:	6-month horizon		
	Posterior inclusion probability (%)	*Coef.*	*std. err.*
Slope of yield curve	100	-1.1	0.15
Change in payroll employment	100	-1.6	0.38
S&P 500 (3-month return)	100	-0.1	0.03

Panel C:	12-month horizon		
	Posterior inclusion probability (%)	*Coef.*	*std. error*
Slope of yield curve	100	-1.4	0.16
GZ Index	95	0.7	0.24

Note: Table shows only indicators with posterior inclusion probability greater than 90 percent.

Current recession risks

As indicated by the black line in the right panel of Figure 2, we currently estimate the probability of the U.S. economy being in recession 3, 6, or 12 months from now at 13, 18, and 16 percent, respectively. However, the recession probabilities produced by our model have moved higher since December, and the uncertainty around this forecast has widened considerably, as can be seen by comparing the current forecast to the left-side panel, which displays the forecast using data through December 2015.

Figure 2. Recession risks have increased since December 2015 but remain relatively subdued.

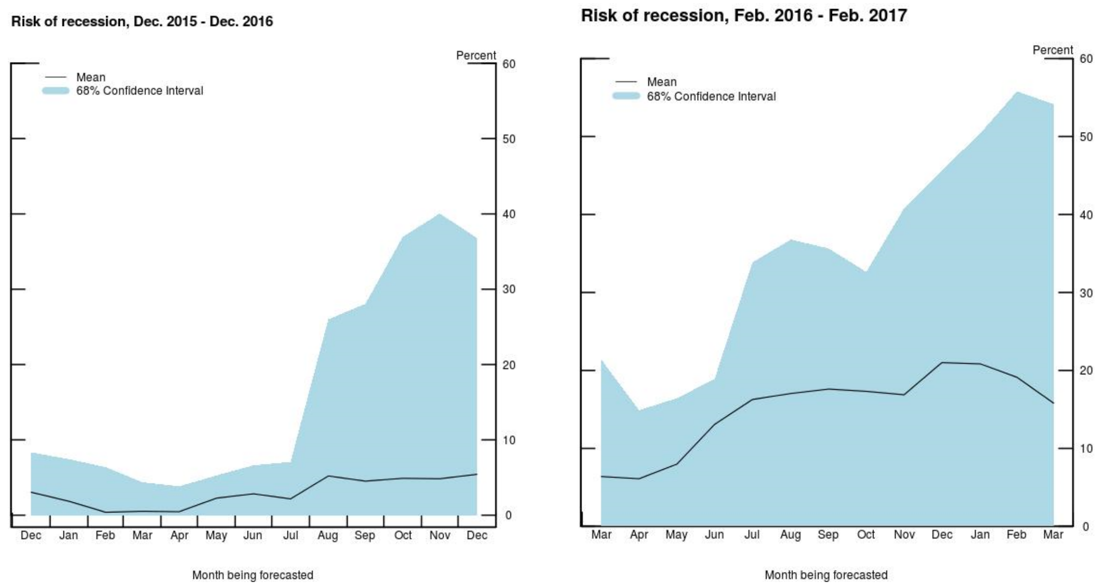

What accounts for the recent changes to the forecast? Consider first the 3-month horizon. In December, the 3-month ahead recession probability forecast stood at 1 percent, whereas by February it had risen to 13 percent. This increase was driven by the two financial indicators that are important predictors at this horizon (as detailed in Panel A of Table 1). Specifically, from December to February, the slope of the yield curve flattened from 2.2 percentage points to 1.5 percentage points, while the S&P 500 index also moved lower. And although the 3 month ahead forecast also depends on the growth in nonfarm payroll employment and real PCE, neither of these changed appreciably in recent months.

Turning to the 12-month ahead model, Panel C of Table 1 shows that the forecasted recession probability depends primarily on two variables, the GZ credit spread index and the slope of the yield curve. Incoming data in January and February have increased the model's view of the probability of recession 12 months hence from 5 percent to 16 percent. The increase has been driven by both a flattening of the yield curve and a widening of credit spreads, with the GZ index having increased from 2.5 percent to 2.9 percent. In sum, the model views the recent deterioration in financial indicators as signaling greater downside risk to the economy.

Finally, as forecasts of recession probabilities have increased since December 2015, the confidence intervals around these estimates have widened correspondingly. The blue shaded area in Figure 3 displays the 68 percent confidence intervals around the current BMA forecast. The current 68-percent confidence interval around the 12-month ahead forecast now ranges from 0 percent to 50 percent, whereas in December the range was 0 percent to 40 percent.[7] It is noteworthy that, despite the flexibility of the BMA approach, these confidence intervals span a very large range, emphasizing the limited ability of simple statistical models to forecast downturns.

An important caveat: Forecasting accuracy out-of-sample

While the BMA models were designed to fit the historical pattern of U.S. recessions, an additional issue is how they will likely perform out-of-sample. A natural question to ask, therefore, is how strong a signal might the BMA methodology have sent ahead of the 2001 and 2008 recessions? To shed light on this issue, we perform a pseudo out-of-sample forecasting exercise by, for example, using data only through the fall of 2000 to produce forecasts over the next 12 months to measure the extent to which the model could have anticipated the subsequent downturn.[8]

As shown in Table 2, six months prior to the 2001 recession, the model forecast a 35 percent probability that the U.S. economy would be in recession 12 months from that point in time, well above the unconditional average. The results indicate that by December, three months ahead of the NBER-dated recession, the model would have sent a fairly strong signal of the pending downturn.

[7] Note that this widening of uncertainty is in a sense mechanical—uncertainty in the probit model is reflected in the linear combination from equation (1), $\alpha + x'_{t-h}\beta$. However, because the transformation from this combination to a probability is nonlinear, small movements in the covariates can produce large changes in their associated probability and confidence interval.

[8] Importantly, owing to data constraints, we use current-vintage data for this exercise. Any judgement on the model's success based on forecasts using revised data should be viewed as an upper-bound on the model capability.

Table 2: Recession probabilities prior to 2001 recession

	Forecast made using data through:	
	Sept-2000	*Dec-2000*
Current-month	6	4
Three-months hence	28	67
Six months hence	32	57
Twelve months hence	35	39

Note: NBER peak dated March, 2001; NBER trough is November, 2001.

Similarly, with data six months prior to the December 2007 peak in economic activity, June 2007, the model forecast a 25 percent probability that the U.S. economy would be in recession 12 months from that point in time. Three months later, with data through September 2007 in hand, the model forecast of a recession 12 months had risen to 40 percent. Most models do not send very strong signals of recession very far ahead of time.[9]

Table 3: Recession probabilities prior to 2007 recession

	Forecast made using data through:	
	June-2007	*Sept-2007*
Current-month	15	22
Three-months hence	26	22
Six months hence	26	41
Twelve months hence	25	41

Note: NBER peak dated December, 2007; NBER trough is June, 2009.

[9] See, for example, Chauvet and Potter (2008), Hamilton (2011), and Berge (2015) for discussions of the real-time performance of several different classes of recession models.

Appendix

Here we present a more formal evaluation of the model's in-sample fit. Figure A1 plots so-called Receiver Operating Characteristic (ROC) curve. The figure depicts the tradeoff associated with achieving a particular true positive rate—that is, the likelihood the model correctly forecasts recessions that actually occur—against the model's corresponding false positive rate—the likelihood it predicts high recession odds when the economy actually continues to expand.[10]

Figure A1. Tradeoff between true positive and false positive in-sample predictions in the BMA model

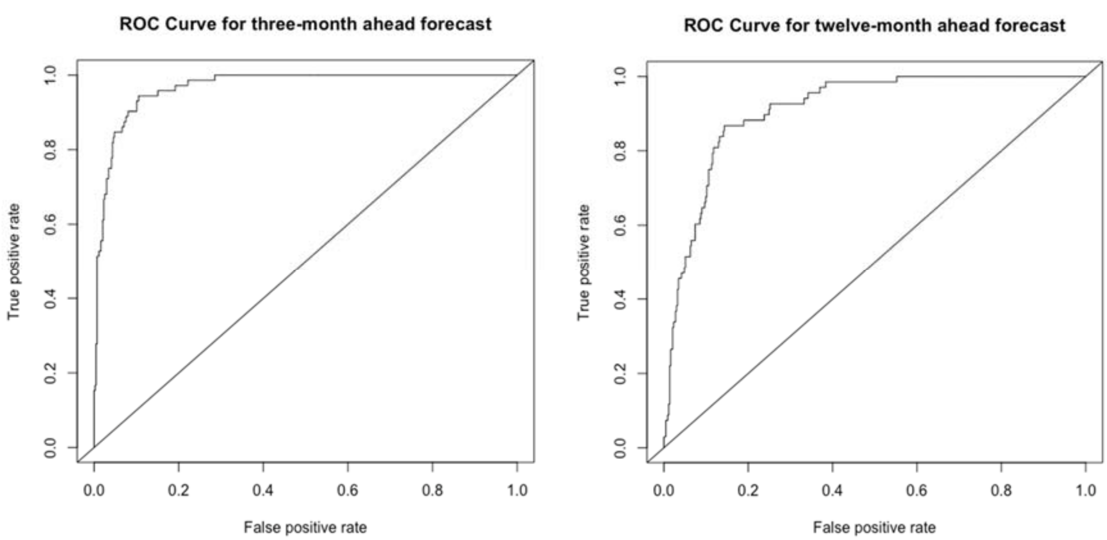

Focusing first on the left-side panel, at the 3-month horizon, to obtain a true positive rate of 85 percent—that is, the fraction of times the model signals recession and a recession is actually realized—one has to bear a false positive rate—periods when the model signals recession but an expansion occurs instead—of just 5 percent. According to the right panel, the 12-month ahead forecast is able to classify NBER-defined recessions quite well in sample: to obtain the same true positive rate of 85 percent, the corresponding false positive rate is only 15 percent, quite an impressive performance.

[10] See Berge and Jorda (2011) for a formal introduction to ROC curves and their use in evaluating forecasts of NBER recession dates.

An additional statistic that emerges from the ROC curves is the area underneath the curve, or AUC. Values of the AUC vary between 0.5, indicating predictive performance akin to a coin-toss, and 1, perfect classification ability. The AUC's of the BMA models at the 3 and 12 month horizons are 97 and 91 percent, respectively. Of course, out-of-sample, the model's performance would deteriorate. Berge (2015) provides evidence from a pseudo out-of-sample forecasting exercise that the AUCs of very similar models that forecast three and twelve months hence obtain AUC statistics of 90 and 85 percent, respectively.

Appendix Table 1: Variables included in forecasting models

Variable	Definition/notes	Transformation
Financial Variables		
Slope of yield curve	10-year Treasury less 3-month yield	
Curvature of yield curve	2 x 2-year minus 3-month and 10-year	
GZ index	Gilchrist and Zakrajsek (AER, 2012)	
TED spread	3-month ED less 3-month Treasury yield	
BBB corporate spread	BBB less 10-year Treasury yield	
S&P 500, 1-month return		1-month log diff.
S&P 500, 3-month return		3-month log diff.
Trade-weighted dollar		3-month log diff.
VIX	CBOE and extended following Bloom	
Macroeconomic Indicators		
Real personal consumption expend.		3-month log diff.
Real disposable personal income		3-month log diff.
Industrial production		3-month log diff.
Housing permits		3-month log diff.
Nonfarm payroll employment		3-month log diff.
Initial claims		3-month log diff.
Weekly hours, manufacturing	4-week moving average	3-month log diff.
Purchasing managers index		3-month log diff.

Note: Treasury yields from Gurkaynak, Swanson and Wright (2007).

References

Berge, T.J. and O. Jorda (2011) "Evaluating the Classification of Economic Activity into Recessions and Expansions," *American Economic Journal: Macroeconomics* 3(2), 246-277, April.

Berge, T.J. (2015) "Predicting Recessions with Leading Indicators: Model Averaging and Selection over the Business Cycle," *Journal of Forecasting* 34(6): 455-471.

Chauvet, M. and J. Piger (2008) "A comparison of the real-time performance of business cycle dating methods," *Journal of Business and Economic Statistics* 26: 42–49.

Faust, J., S. Gilchrist, J.H. Wright, and E. Zakrajšek (2013), "Credit Spreads as Predictors of Real-Time Economic Activity: A Bayesian Model-Averaging Approach," *Review of Economics and Statistics* 95(5): 1501–1519.

Gilchrist, S., and E. Zakrajšek (2012), "Credit Spreads and the Business Cycle Fluctuations," *American Economic Review* 102(4): 1692-1720.

Gurkaynak, R., B. Sack and J.H. Wright (2006) "The U.S. Treasury Yield Curve: 1961 to the Present," FEDS Working Paper Series 2006-28.

Hamilton J.D. (2011) "Calling recessions in real time," *International Journal of Forecasting* 27(4): 1006–1026.

Leamer, E.E. (1978), *Specification Searches: Ad Hoc Inference with Nonexperimental Data* (New York: Wiley).

Piger, J., and J.C. Morely (2008) "Trend-Cycle Decomposition of Regime-Switching Processes," *Journal of Econometrics* 146: 220-226.

Raftery, A.E. (1995) "Bayesian Model Selection in Social Research," *Sociological Methodology* 25: 111-163.

Wright, J.H. (2009) "Forecasting U.S. Inflation by Bayesian Model Averaging", *Journal of Forecasting*, 28: 131-144.