MASTER'S THESIS

MASTER IN ECONOMICS, FINANCE AND COMPUTER SCIENCE

Predicting Recessions: Consumer and Business Confidence versus Financial and Macroeconomic Indicators

Inmaculada Millán Díaz

Supervisors: Máximo Camacho (UMU) and
Manuel Ruíz (UPCT)

University of Huelva and International University of Andalusia

November 2018

MASTER'S THESIS

MASTER IN ECONOMICS, FINANCE AND COMPUTER SCIENCE

Predicting Recessions: Consumer and Business Confidence versus Financial and Macroeconomic Indicators

Abstract

In this paper we use Boosted Regression Trees (henceforth, BRT), as alternative to tradional Probit approaches, to revisit the usefulness of different groups of leading indicators –financial indicators, real economic variables and the *animal spirit* indexes– for predicting US recessions. In particular, three main hypotheses are tested: i) if the relative importance of different groups of indicators as predictors –and their marginal effects– of recessions at the US economy is different depending on the forecast horizon; ii) if the predictive power of some leading indicators has declined over time while others have gained in importance; and, iii) we also evaluate the consumer confidence and business climate indicators in comparison with other financial and non financial indicators, in order to ckeck if the *information view* hypothesis applies (Barsky and Sims, 2009), i.e. if movements in confidence reflect information about future economic prospects, with real economic effects. Results are consistent with previous literature, pointing again towards the importance of the interest rates and the interest rate spread as predictors of recession, but not at all forecast horizon. Consumer confidence and business climate indicators emerge as the most influential predictors of US recessions for a forecast horizon of three months.

Key words: Business cycle turning points; boosting regression trees; predicting recessions; US; animal spirits; business climate, consumer confidence.

Contents

1 Introduction

There is a large body of empirical studies aimed at comparing the accuracy of alternative models and economic leading indicators as predictors of the forecasts of recession. The benchmark model for forecasting the state of the business cycles is given by a model where the state of the business cycle denoted by y_t, —a dichotomous variable that is one if there is an official recession in the subsequent k quarters and zero otherwise— that is explained by a vector of N forward-looking economic indicators:

$$y_t = f(x_{t-k-1}, b)$$

where x_{t-k-1} is a set of predictors, b is the vector of parameters and k the forecast horizon.[1]

Early studies, by using simple limited dependent variable models —logit or probit models— explored the probability of recession conditioned on macroeconomic variables and financial indicators —specially interest rates, term spread and the yield curve— as in (Dueker, 1997), (Estrella and Mishkin, 1997) and (Estrella and Mishkin, 1998).

Later this literature evolved, at least, into three directions. One of them has been to explore the robustness and the accuracy of previous approaches, by using alternative econometric approaches. In this group we can include at least, the works of (Qi, 2001) who used neural networks, (Peláez, 2005) and (Wright, 2006), (Kauppi and Saikkonen, 2008), (Nyberg, 2010), (Kim and Swanson, 2014) and (Proaño and Theobald, 2014) who applied different versions of dynamic probit models for predicting recessions, or (Liu and Moench, 2016) who applied multivariate probit

[1] NBER recessions, in our case.

models, among others.[2]

Some research lines have focused to test the relative importance of different groups of leading indicators. The use of financial variables as predictors has been extensively explored, maybe due to the existence of a theory supporting their role potential for predicting recessions. Very shortly, the rationale for why term spreads might be a useful leading indicator is that under the expectations hypothesis the term spread (the short-term rates minus long-term yields) give us a rough measure of the difference between current short-term interest rates and the average of expected future short-term interest rates (Wright, 2006). The higher is the term spread, the more restrictive is current monetary policy, and the more likely is a recession over the subsequent quarters. Complementary to this, other sets of variables like macroeconomic and expectation predictors, were also explored. In that sense, one could argue that autonomous fluctuations in beliefs (i.e. changes in the consumer and business confidence measures) can have effects on real activity becoming a leading indicator of recessions. In addition, consumer and business confidence might be based on sets of information captured by consumers and entrepreneurs that does not otherwise show up in econometricians' information sets (Barsky and Sims, 2009). To test, the power of forecasting of these set of variables is the aim of this paper.

Finally, other researchers chose to check the stability of the forecast relationship to account for structural changes or other nonlinearities. The motivation of this type of works is given by the fact that the behavior of many predictors across the business cycle can change. Thus, the relative importance of these leading indicators can change over time and the existence of potential breakpoints in the

[2] This body of literature is surveyed by (Wheelock et al., 2009).

in-sample forecasting relationship should be taken into account —see, (Estrella et al., 2003), (Stock and W Watson, 2003), (Giacomini and Rossi, 2009) and (Benati and Goodhart, 2008) as works in which the existence of a time-varying relationship is explored—.

To face up the problems mentioned above, in a single framework, a handful of recent works have explored the use of machine learning algorithms for checking and (re)evaluate the usefulness of different groups of indicators as predictors of recessions. For instance, (Buchen and Wohlrabe, 2011) evaluated the forecast performance of boosting in comparison to traditional models.

(Ng, 2014) screens the forecasting performance of a set of 1500 leading economic indicators consisting of 132 real and financial time series plus their lags, by using boosting. In particular, she is interested in checking the usefulness of this approach as an alternative way to detect recessions. This analysis allows her not only to identify the relative predictive power of different predictors in different sub-periods but also how the accuracy of them depends on the forecast horizon.

(Berge, 2015) explores different models for predicting business cycle turning points. His results suggest that successful models of recessions condition on different subsets of economic indicators at different forecast horizons.[3] The results of (Berge, 2015) point out that leading indicators associated to real economic activity (especially indicators of the labor market) are the best predictors of recessions at very short horizons. He also finds that the yield curve is a predictor of turning points, but with the caveat of some recessions (like 2001 and 2007). Finally, he provides

[3] In order to evaluate the robustness he applies different strategies: weighted forecasts,a Bayesian model and the use of both linear and nonlinear boosting as a method for improving a *machine learning* algorithm. He only rejects the use of the *unweighted* model averages.

evidence about the accuracy of bond spreads and indicators of the housing market as predictors into the medium term.

A fourth piece of research complementary to (Ng, 2014) and (Berge, 2015) is the work of (Döpke et al., 2017) in which the approach is extended and applied to predict German recessions. In particular, they apply BRT to reexamine the usefulness of a set of leading indicators for predicting recessions.

Following this strategy, this paper re-evaluates the predictive power of a number of commonly followed financial and macroeconomic variables, in the US economy, exploring the extent to which the business climate and the consumer confidence indexes are for forecasting the on set of recessions in the US. To this end, we use a BRT approach for allowing nonlinearity in the relationship between our predictors and recessions and looking for changes in the relative importance and marginal effects of these indicators depending on the forecast horizon. By using the general framework proposed by (Döpke et al., 2017) for forecasting recessions we explore the potential predictive power of different subsets of indicators at different forecast horizons.

In particular, three main hypotheses are tested: i) the relative importance of different groups of indicators as predictors –an their marignal effects– of recessions at the US economy is different depending on the forecast horizon; ii) the predictive power of some leading indicators has declined over time while others have gained in importance; and, iii) we also evaluate the consumer confidence and business climate indicators in comparison with other financial and non financial indicators.

Results are consistent with previous literature, pointing again towards the importance of the interest rates and the interest rate spread, but not at the shortest forecast horizon. At the shortest one, consumer confidence and business climate indicators emerge as the best predictors of US recessions.

The rest of the paper is structured as follows. The prediction model based on boosted regression trees is explained in Section 2. Section 3 is devoted to describing the data and sources and to report and discuss the prediction results. Finally Section 4 concludes.

2 Boosted Regression Trees

As we mentioned, we model recessions as a binary variable, $y_{t+k} \in \{0, 1\}$, where $y_t = 1$ denotes a recession according to the The National Bureau of Economic Research (NBER), $t = 1, \ldots,$ denotes a time index, and k denotes the forecast horizon. It will also be useful to define $\tilde{y}_{t+k} = 2y_{t+k} - 1$, such that $\tilde{y}_{t+k} \in \{-1, 1\}$ as (Döpke et al., 2017) and (Ng, 2014) do. The objetive is to model the links between recessions and a sets of leading indicators, $\boldsymbol{x_t} = (x_{t,1}, x_{t,2}, \ldots, x_{t,N})$, by means of a function $F(\boldsymbol{x_t})$, such as minimize the expected value of a loss function, $\mathcal{L}$.[4] (Friedman et al., 2000) suggest to minimize the exponential loss function

$$\mathcal{L}(F) = E(e^{-\tilde{y}_{t+k} F(\boldsymbol{x_t})}), \tag{1}$$

for selecting $F(\boldsymbol{x_t})$. The loss function increases when $\tilde{y}_{t+k}$ and $F(\boldsymbol{x_t})$ have different signs, and it decreases when $\tilde{y}_{t+k}$ and $F(\boldsymbol{x_t})$ have the same signs. In order words, the loss function decreases when $\boldsymbol{x_t}$ helps to classify recessions at forecast horizon k.

[4] E represents expectation over the joint distribution of $\tilde{y}_{t+k}$ and $\boldsymbol{x_t}$.

If we denote the conditional probability of recession as $P(\tilde{y}_{t+k} = 1 \mid \boldsymbol{x_t})$, the equation (1) can be rewritten as

$$E(e^{-\tilde{y}_{t+k}F(\boldsymbol{x_t})}) = P(\tilde{y}_{t+k} = 1 \mid \boldsymbol{x_t})e^{-F(\boldsymbol{x_t})} + P(\tilde{y}_{t+k} = -1 \mid \boldsymbol{x_t})e^{F(\boldsymbol{x_t})}. \tag{2}$$

The necessary condition for minimizing the loss function is given by

$$\frac{\partial E(e^{-\tilde{y}_{t+k}F(\boldsymbol{x_t})})}{\partial F(\boldsymbol{x_t})} = -P(\tilde{y}_{t+k} = 1 \mid \boldsymbol{x_t})e^{-F(\boldsymbol{x_t})} + P(\tilde{y}_{t+k} = -1 \mid \boldsymbol{x_t})e^{F(\boldsymbol{x_t})} = 0, \tag{3}$$

consequently $\mathcal{L}(F)$ is minimized by setting $F(\boldsymbol{x_t})$ to one-half of the log-odds-ratio:

$$F(\boldsymbol{x_t}) = \frac{1}{2}\log\frac{P(\tilde{y}_{t+k} = 1 \mid \boldsymbol{x_t})}{P(\tilde{y}_{t+k} = -1 \mid \boldsymbol{x_t})}. \tag{4}$$

Rearranging terms and given that $P(\tilde{y}_{t+k} = 1 \mid \boldsymbol{x_t}) = 1 - P(\tilde{y}_{t+k} = -1 \mid \boldsymbol{x_t})$ we have

$$P(\tilde{y}_{t+k} = 1 \mid \boldsymbol{x_t}) = \frac{e^{F(\boldsymbol{x_t})}}{e^{-F(\boldsymbol{x_t})} + e^{F(\boldsymbol{x_t})}} = p_+(\boldsymbol{x_t}), \tag{5}$$

$$P(\tilde{y}_{t+k} = -1 \mid \boldsymbol{x_t}) = \frac{e^{-F(\boldsymbol{x_t})}}{e^{-F(\boldsymbol{x_t})} + e^{F(\boldsymbol{x_t})}} = p_-(\boldsymbol{x_t}). \tag{6}$$

The basic idea of boosting is to divide a problem into simple problems such that the function to estimate, $F(\boldsymbol{x_t})$, is expressed in terms of simpler functions, $T(\boldsymbol{x_t})$. Then,

$$F(\boldsymbol{x_t}) = \sum_{m=0}^{M} T_m(\boldsymbol{x_t}), \tag{7}$$

where M denotes the number of functions that we want to consider.

In the machine-learning literature $F(\boldsymbol{x_t})$ is known as a strong learner and the simpler functions, $T(\boldsymbol{x_t})$, are known as weak learners.

In the algortigm the variables are included sequentially and no change is made to the coefficients of the variables already included. The size of the ensemble is determined by M.

We use a boosting algorithm known as gradient-descent boosting (Friedman, 2002). The algorithm is described below.

Gradient boosting for minimizing $\mathcal{L}(F) = E(e^{-\tilde{y}_{t+k}F(\boldsymbol{x_t})})$:

1. Initialize the algorithm: $F_0 = T_0 = \frac{1}{2}\log\frac{P(\tilde{y}_{t+k}=1)}{P(\tilde{y}_{t+k}=-1)}$.

2. For $m = 1, \ldots, M$:

   a. Compute the negative gradient $\mathcal{L}'(F) = -\frac{\partial\mathcal{L}(F)}{\partial F}\ |_{F_{m-1}(\boldsymbol{x_t})}$.

   b. Let $T_m(\boldsymbol{x_t})$ be the best fit of $\mathcal{L}'(F)$ using predictor $\boldsymbol{x_t}$.

   c. Update $F_m(\boldsymbol{x_t}) = F_{m-1}(\boldsymbol{x_t}) + T_m(\boldsymbol{x_t})$.

3. When the recursion reaches $m = M$, the strong learner, $F_M(\boldsymbol{x_t})$, has been computed as the sum of the weak learners, $T_m(\boldsymbol{x_t})$, where $m = 0, \ldots, M$.

Step (2a) computes the adjusted response and step (2b) obtains the best model. Step (2c) then uses an approximation of the negative gradient vector to implement the update. Then gradient boosting amounts to repeatedly finding a predictor to fit the residuals not explained in the previous step. By introducing a parameter λ to slow the effect of $T_m(\boldsymbol{x_t})$ on $F(\boldsymbol{x_t})$, step (2c) can be modified to control the rate at which gradient descent takes place:

$$F_m(\boldsymbol{x_t}) = F_{m-1}(\boldsymbol{x_t}) + \lambda T_m(\boldsymbol{x_t}), \tag{8}$$

where $\lambda \in [0, 1]$ is known as shrinkage parameter. The parameters M and λ are related, a low value of the shrinkage parameter would necessitate a larger M.

At each iteration m, a regression tree partitions the space of the leading indicators into L disjoint regions, $\{R_{l,m}\}_1^L$, and predicts a separate constant value in each one. The regression tree that is integrated in step (2a) of the gradient-descent-boosting algorithm gives the optimal response for the given loss function such that (Friedman, 2002)

$$\gamma_{l,m} = \arg\min_{\gamma} \sum_{\boldsymbol{x_t} \in R_{l,m}} \mathcal{L}(F_{m-1}(\boldsymbol{x_t}) + \gamma). \tag{9}$$

The minimization problem specified in (9) can be solved using Newton's method (Friedman et al., 2000). The equation (7) can then be rewritten as

$$F(\boldsymbol{x_t}) = \sum_{m=0}^{M} \gamma_{l,m} \mathbf{1}_{\boldsymbol{x_t} \in R_{l,m}}, \tag{10}$$

where $\mathbf{1}$ denotes the indicator function. Combining the equations (8) and (10)

$$F_m(\boldsymbol{x_t}) = F_{m-1}(\boldsymbol{x_t}) + \lambda \gamma_{l,m} \mathbf{1}_{\boldsymbol{x_t} \in R_{l,m}}. \tag{11}$$

3 Empirical Analysis

3.1 Preliminary Analysis of Data

As a common practice in this body of literature, and given the scope of our empirical application, we define U.S. recessions following the definition adopted by the NBER Business Cycle Dating Committee and the "official" turning points determined by using the so-called "classical" approach. Thus recessions in our analysis refer to peak and trough dating as published by the NBER.[5] Table 1 reports the official peaks and trough provided by the U.S. Committee.

According the aim and scope of this paper, we include a subset of commonly used leading indicators (Marcellino, 2006) that could be grouped in three major categories: i) a variaety of financial indicators; ii) a group of variables for describing the real economy; and, iii) variables related to expectations (for both, consumers and firms). Table 2 summarises this set of variables and the source of data. In addition, in some cases, the original variables are transformed for resulting in stationary time series.[6] Finally, and regarding to the time span, the selection has been done establishiong a common range for the variables captured in our analysis: the period from 1960:2 until 2016:12.

3.2 Model Calibration

For estimating the BRT model, we use the package "gbm" (Generalized Boosted Regression Models) and an extended implementation of the Freud and Schapire's AdaBoost algorithm (Freund and Schapire, 1997) and Friedman's gradient boost-

[5]The algorithms of Bry and Boschan (1971) and Harding and Pagan (2003) are other methods for detecting business cycle turning points.

[6]Unit roots tests are not shown here for saving space but are available upon request.

ing machine (Friedman, 2001, 2002).[7] As in (Döpke et al., 2017), for each weak learner the tree depth is 5 (that is, number of divisions) and for each terminal node the minimum number of observations is 5. The shrinkage parameter, λ, mentioned in Section 2 takes a value of 0.005. We simulate the algorithm 1000 times where we use 70% of the data for training the model and the rest for testing (quasi out-of-sample model testing). We set a maximum number of 3000 weak learners but the average number of optimal weak learners for a forecast horizon of three months is 551 with a standard deviation of 112 and 513 with a standard deviation of 113 for a forecast horizon of 6 months. Finally, we use 30% of the data to train the data and to built the next weak learner in the out-ot-sample estimates.

3.3 In-Sample Results

Figure 1 shows the (in-sample) posterior probability of recessions denoted $P(y_t = 1 \mid \boldsymbol{x_{t-k}})$ along with the NBER recession dates. The estimated probabilities for the recessions clearly display spikes around the NBER recession dates. The 1980-1981 recessions are very close in time and with $k = 3$, the estimated probabilities shows some increases and decreases. The estimated probabilities when $k = 6$ fare better but the spikes are still not as pronounced as one had expected.

Following the strategy adopted by (Döpke et al., 2017) we also try to capture the relative importance of the predictor x_j. This is calculated measuring the variation that x_j causes in $F(\boldsymbol{x_t})$. (Breiman et al., 1984) proposed

$$\hat{I}_j^2(T) = \sum_{s=1}^{J-1} \hat{i}_s^2 1(\upsilon_s = j), \tag{12}$$

[7] We use the codes gently provided by (Döpke et al., 2017).

where s are the nonterminal nodes of the J-terminal nodes tree T, v_s is the splitting variable associated with node s, and $\hat{i}_s^2$ is the improvement in the square error as a result of the split.

For a collection of trees $\{T_m\}_1^M$ we average over all of the trees

$$\hat{I}_j^2 = \frac{1}{M} \sum_{m=1}^{M} \hat{I}_j^2(T_m). \tag{13}$$

The statistic (Friedman, 2001) is based on the number of times a variable is selected over the M steps, weighted by its improvement in squared error. The sum of $\hat{I}_j^2$ over j is 100. Higher values thus signifies that the associated variable is important (Ng, 2014).

Figure 2 shows that the indicators of business climate and consumer confidence are the most influential leading indicators when the forecast horizon is three months, whereas the relative importance of these indicators is different when the forecasting horizon considered is six. At this case the time series of the 5Y and 10Y spreads are the two most important predictors of reccessions for a six-month horizon. In addition, the relative importance of the shares prices exceeds 10% for a forecast horizon of three months and ranges between 5% to 10% for a forecast horizon of six months. Regarding the relative importance of real economic variables, the unemployed rate and the production growth ranges between 5% to 10% for a forecast horizon of three months and below the 5% for a forecast horizon of six months. The other leading indicators have a relative importance that is below 5%.

The marginal effects show the effect of a leading indicator on the probability of recession, where the effects of others leading indicators are controlled with a weighted-traversal technique given by (Friedman, 2001) where $\boldsymbol{z}_l$ is defined as a subset of size l of the input variables $\boldsymbol{x}$,

$$\boldsymbol{z}_l = \{z_1, \ldots, z_l\} \subset \{x_1, \ldots, x_N\} \tag{14}$$

and $\boldsymbol{z}_{\backslash l}$ as the complementary subset

$$\boldsymbol{z}_l \cup \boldsymbol{z}_{\backslash l} = \boldsymbol{x}. \tag{15}$$

For a tree and specific subset $\boldsymbol{z}_l$, a weighted traversal of the tree is performed. At the root of the tree, a weight value of 1 is assigned. For each nonterminal node visited, if its split variable is in $\boldsymbol{z}_l$, the weight is not modified. If the node's split variable is a member of $\boldsymbol{z}_{\backslash l}$, the current weight is multiplied by the fraction of training observations. When the tree traversal is complete, the dependence partial of the variables in $\boldsymbol{z}_l$ is the corresponding weighted average over the terminal nodes visited during the tree traversal. For a collection of M trees obtained through boosting the results for the individual trees are averaged.

Figure 3 shows the marginal effects for a forecast horizon of three months while Figure 4 shows the marginal effects for a six-month forecast horizon. The vertical axis shows the probabily of a recession measured in log-odds scale. The probability of recession is low for small values of the inflation rate but from values between 0 and 1 this increases and remains constant. The probability of recession, however, is low for small values of the discount rate but from values higher than 10 it increases notably. The probability of recession is constant for small values of the production growth, shares prices and short rate and from certain intermediate values the probability of recession decreases. The business climate and consumer opinion variables in a forecast horizon of three months are associated with a very high probability of recession for small values that decreases for higher values.

For a forecast horizon of six months, the graph of the marginal effects has the same form although the effects that cause in the probability of recession are not so important as for a forecast horizon of three months. The variables terms spread have the same form as the business climate and consumer opinion variables except that the marginal effects that they cause on the probability of recession are more important for a forecast horizon of six months than of three months. The probability of recession is rather insensitive to changes in the prices of oil, money growth, government yield and unemployed rate.

3.4 Out-Of-Sample Results

In this section, we develop an out-of-sample exercise, whose results are show in Figure 5. We start the experiment estimating the BRT model with data until 1976:12. We then use data for the next month to make a forecast with the estimated model. Then, we reestimate the model and use the reestimated model to make next month' forecast. We continue the recursive process until we reach the last period of the sample.

The recessions of 1980-1981 are very close in time and the model does not distinguish them well. The algorithm needs some more time to develop gradually in the recession of 1990-1991. The recession of 2001 has very short duration, therefore, is the worst estimate. The estimated probability of recession shows notable increases and decreases during the recession of 2007-2009.

Figure 6 shows how the relative importance of the leading indicators changes across the time. For a forecast horizon of three months we observe that the most influential variables are the business climate and the consumer opinion, and that this importance increases with time from approximately 10% to 25%-30%. In addition, the discount rate variable was very influential until 1980 but since then its relative importance decreases. For a forecast horizon of 6 months we observe that the term spread has a great relative importance of approximately 40% until 2000, decreasing up to 20%. However, it still remains the most influential variable for a forecast horizon of 6 months.

4 Conclusion

In this paper we applied the BRT approach on US data for asseing the relative importance of different groups of indicators and their marginal effects on the probability of recession, with special focus on consumer and business confidence variables.

Results show that both the term spread and consumer and climate confidence indexes are important leading indicators, but also that the confidence indexes have more predictive value at the shortest forecast horizon. Thus, our results are consistent with previous literature with respect to the predictive power of financial variables, but also provides support to the main hypothesis of this work, with regard the animal spirits. The approach also allow us to observe how the predictive power of some variables like the term spread has declined over time, while the consumer confidence index and the business climate have gained in importance. The results point to a strong relationship between the information set used for consumers and investors for forming their expectations/forecasts, at least in the short term. However term spread indicators keep their predictive power at 6 month horizon. Maybe the new ways, speed and forms of diffusion of information in the digital era, including the diffusion of information in digital social netweorks have a gowing importance in the way in which individuals make its own forecast and react. These phenomena might be behind the predictive power of these two indicators.

Then, one could argue that the potential research agenda should consider that the power of forecasting of different leading indicators change over time. Moreover the new availability of big amounts of data will allow to work with hundreds or even thousands of potential indicators potentially useful for predicting recessions, with no need of priors about the selection of indicators. In doing so, problems associated to the omission of relevant variables might be avoided.

References

Barsky, R. B. and Sims, E. R. (2009). Information, animal spirits, and the meaning of innovations in consumer confidence.

Benati, L. and Goodhart, C. (2008). Investigating time-variation in the marginal predictive power of the yield spread. Journal of Economic Dynamics and Control, 32(4):1236–1272.

Berge, T. J. (2015). Predicting recessions with leading indicators: Model averaging and selection over the business cycle. Journal of Forecasting, 34(6):455–471.

Breiman, L., Friedman, J. H., Olshen, R. A., and Stone, C. J. (1984). Classification and regression trees. Wadsworth International Group.

Buchen, T. and Wohlrabe, K. (2011). Forecasting with many predictors: Is boosting a viable alternative? Economics Letters, 113(1):16–18.

Döpke, J., Fritsche, U., and Pierdzioch, C. (2017). Predicting recessions with boosted regression trees. International Journal of Forecasting, 33(4):745–759.

Dueker, M. J. (1997). Strengthening the case for the yield curve as a predictor of us recessions. Federal Reserve Bank of St. Louis Review, 79(2):41.

Estrella, A. and Mishkin, F. S. (1997). The predictive power of the term structure of interest rates in europe and the united states: Implications for the european central bank. European economic review, 41(7):1375–1401.

Estrella, A. and Mishkin, F. S. (1998). Predicting us recessions: Financial variables as leading indicators. Review of Economics and Statistics, 80(1):45–61.

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

Freund, Y. and Schapire, R. E. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. Journal of computer and system sciences, 55(1):119–139.

Friedman, J., Hastie, T., Tibshirani, R., et al. (2000). Additive logistic regression: A statistical view of boosting. The Annals of Statistics, 28(2):337–407.

Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. Annals of Statistics, pages 1189–1232.

Friedman, J. H. (2002). Stochastic gradient boosting. Computational Statistics and Data Analysis, 38(4):367–378.

Giacomini, R. and Rossi, B. (2009). Detecting and predicting forecast breakdowns. The Review of Economic Studies, 76(2):669–705.

James, G., Witten, D., Hastie, T., and Tibshirani, R. (2013). An introduction to statistical learning, volume 112. Springer.

Kauppi, H. and Saikkonen, P. (2008). Predicting us recessions with dynamic binary response models. The Review of Economics and Statistics, 90(4):777–

791.

Kim, H. H. and Swanson, N. R. (2014). Forecasting financial and macroeconomic variables using data reduction methods: New empirical evidence. Journal of Econometrics, 178:352–367.

Liu, W. and Moench, E. (2016). What predicts us recessions? International Journal of Forecasting, 32(4):1138–1150.

Marcellino, M. (2006). Leading indicators. Handbook of Economic Forecasting, 1:879–960.

Ng, S. (2014). Boosting recessions. Canadian Journal of Economics/Revue Canadienne d'économique, 47(1):1–34.

Nyberg, H. (2010). Dynamic probit models and financial variables in recession forecasting. Journal of Forecasting, 29(1-2):215–230.

Peláez, R. F. (2005). Dating business-cycle turning points. Journal of economics and finance, 29(1):127–137.

Proaño, C. R. and Theobald, T. (2014). Predicting recessions with a composite real-time dynamic probit model. International Journal of Forecasting, 30(4):898–917.

Qi, M. (2001). Predicting us recessions with leading indicators via neural network models. International Journal of Forecasting, 17(3):383–401.

Stock, J. H. and W Watson, M. (2003). Forecasting output and inflation: The role of asset prices. Journal of Economic Literature, 41(3):788–829.

Wheelock, D. C., Wohar, M. E., et al. (2009). Can the term spread predict output growth and recessions? a survey of the literature. Federal Reserve Bank of St. Louis Review, 91(5 Part 1):419–440.

Wright, J. H. (2006). The yield curve and predicting recessions.

Table 1: Recessions: Peaks and Troughs (By the NBER)

	NBER Turning Points
Peak	1960: 4
Trough	1961: 2
Peak	1969:12
Trough	1970:11
Peak	1973:11
Trough	1975: 3
Peak	1980: 1
Trough	1980: 7
Peak	1981: 7
Trough	1982:11
Peak	1990: 7
Trough	1991: 3
Peak	2001: 3
Trough	2001:11
Peak	2007:12
Trough	2009: 6

Table 2: Data

Series	Source	Link	Transformation
Recession Phases	FRED	link	None
Industrial Production Index	FRED	link	Yes
Consumer Price Index	FRED	link	Yes
Total Share Prices	FRED	link	Yes
Consumer Opinion Surveys	FRED	link	None
Crude Oil Prices	FRED	link	Yes
Effective Federal Funds Rate	FRED	link	Yes
Discount Rate	FRED	link	None
Long-Term Government Yield (10 years)	FRED	link	Yes
Business Climate	FRED	link	None
Civilian Unemployment Rate	FRED	link	None
Money Supply M1	FRED	link	Yes
Money Supply M1, real	FRED	link	Yes
Money Supply M2	FRED	link	Yes
Money Supply M2, real	FRED	link	Yes
Term Spread (10 years)	FRED	link	None
Term Spread (5 years)	FRED	link	None

Note: The transformation made is the change over the previous year. FRED is the Federal Reserve Bank of St. Louis database. The data start in 1960/2 and the last forecast in made in 2016/12.

Figure 1: In-sample Performance

Forecast Horizon: 3 Months

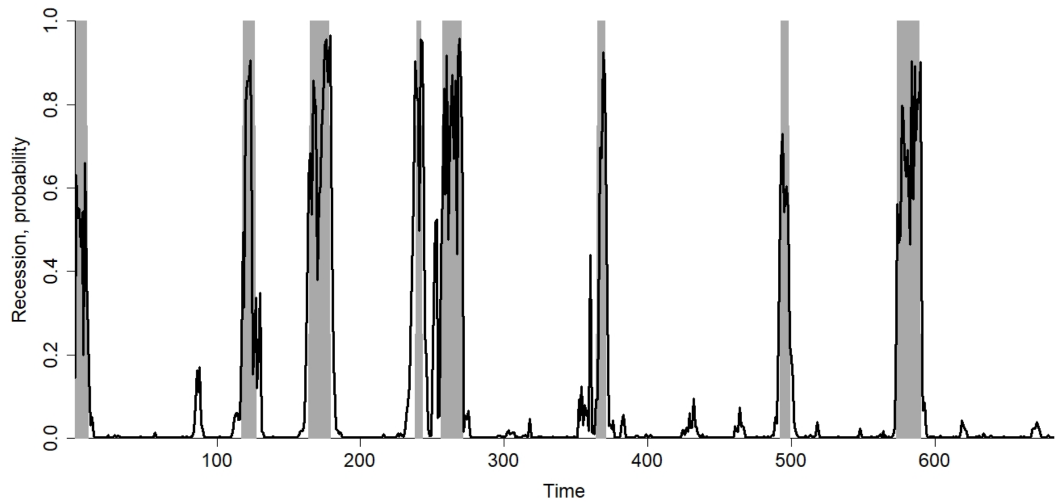

Forecast Horizon: 6 Months

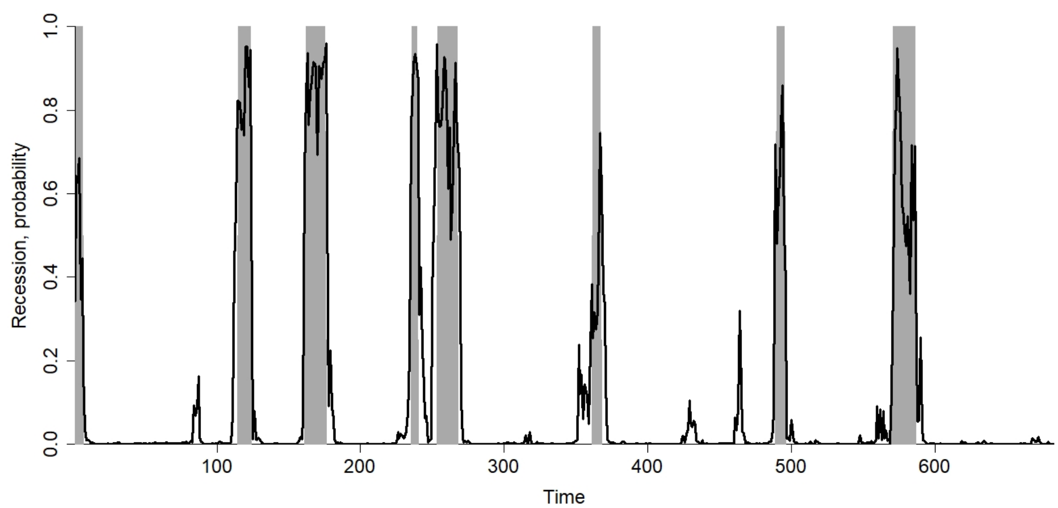

Note: Shaded areas indicate recessions. The horizontal axis shows time where 1 is the period 1960:2 and 683 the period 2016:12.

Figure 2: Relative Importance of Leading Indicators

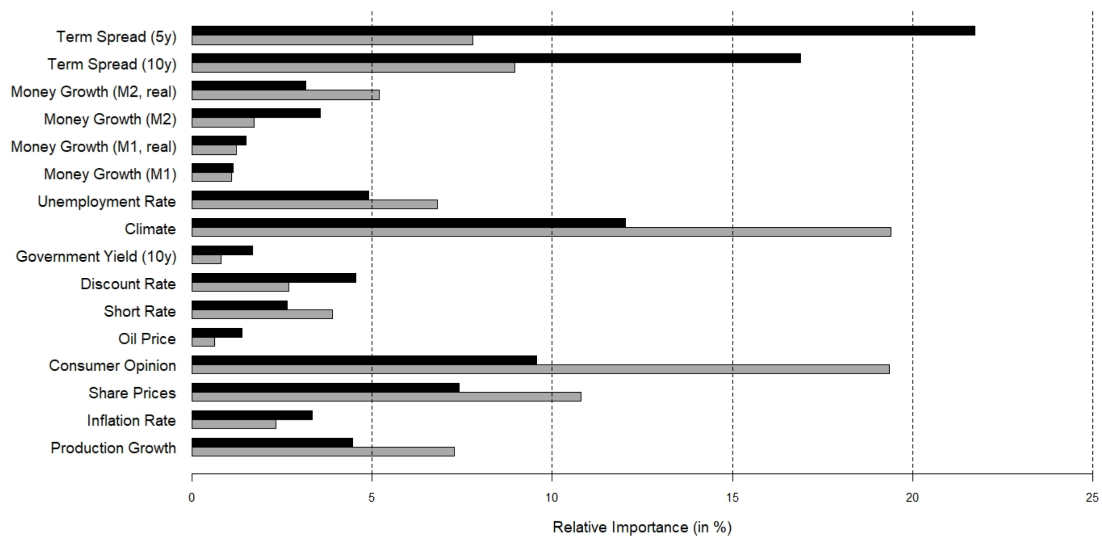

Note: The gray bar shows the relative importance for a forecast horizon of three months. The black bar shows the relative importance for a forecast horizon of six months.

Figure 3: Marginal Effects (Forecast Horizon: 3 Months)

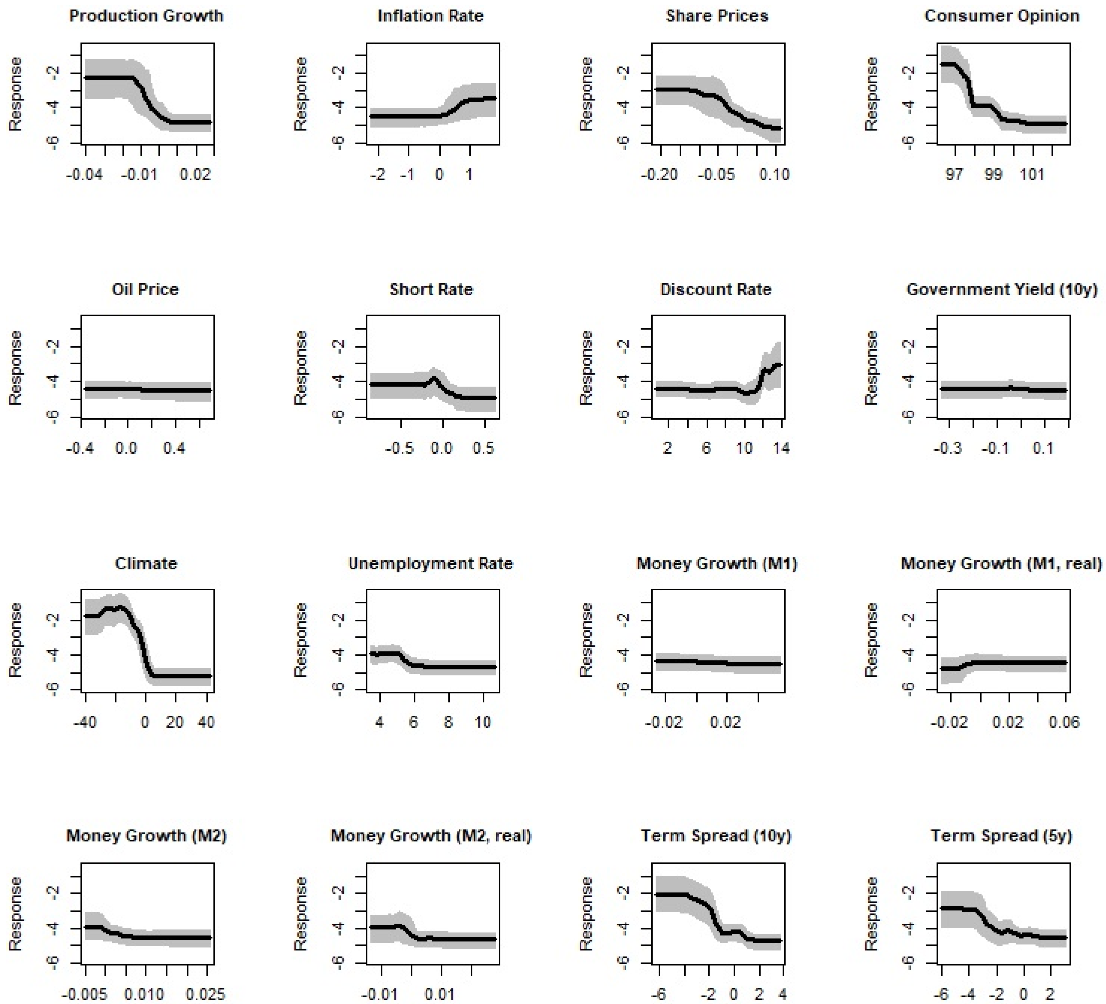

Note: The horizontal axis shows the leading indicators. The vertical axis shows the probabily of a recession measured in log-odds scale. The thick line shows the mean of the marginal effects computed across all simulation runs. The gray area denotes the 95% confidence interval.

Figure 4: Marginal Effects (Forecast Horizon: 6 Months)

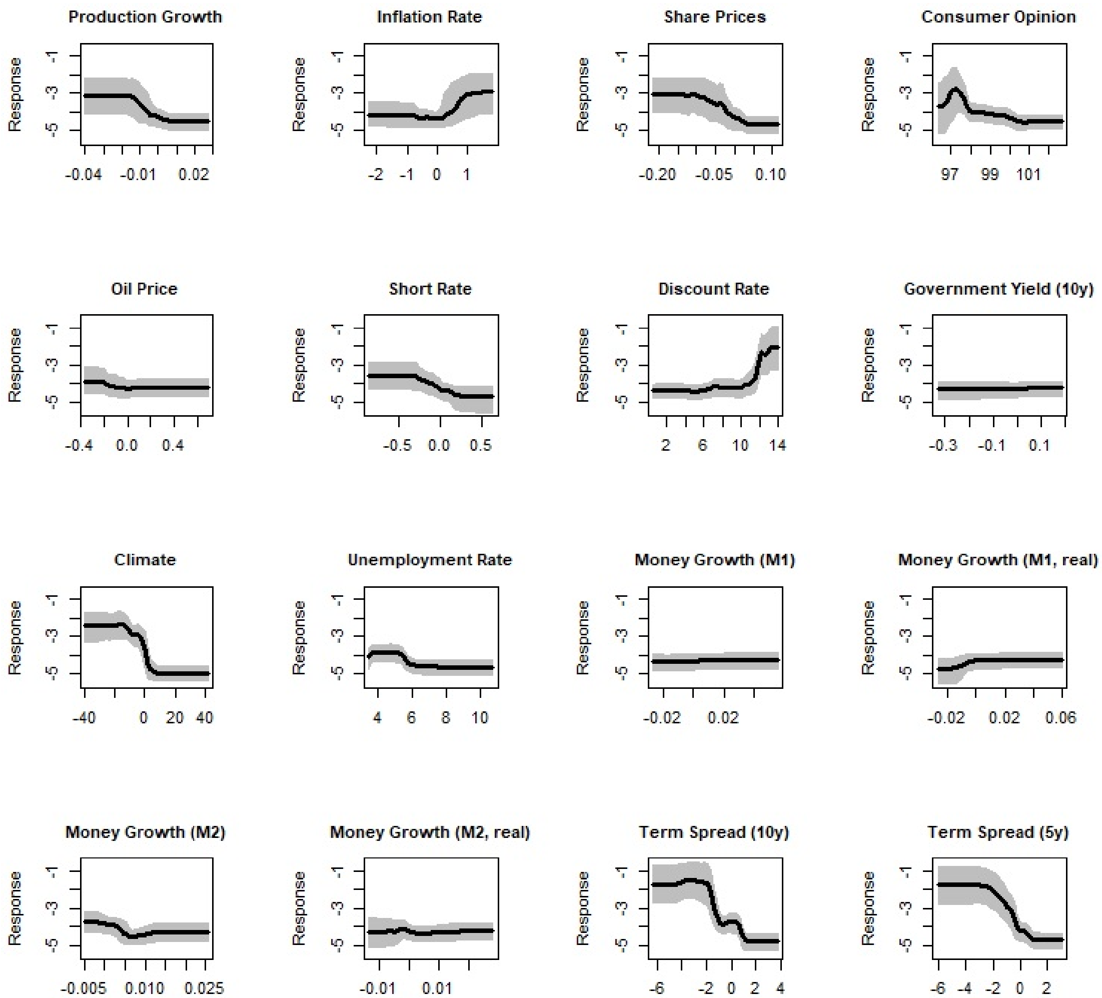

Note: The horizontal axis shows the leading indicators. The vertical axis shows the probabily of a recession measured in log-odds scale. The thick line shows the mean of the marginal effects computed across all simulation runs. The gray area denotes the 95% confidence interval.

Figure 5: Out-of-Sample Performance

Forecast Horizon: 3 Months

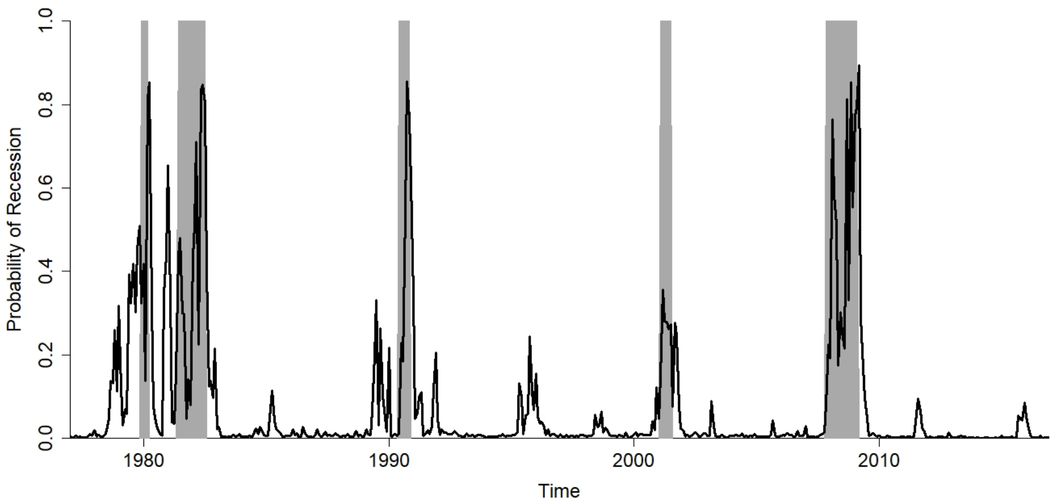

Forecast Horizon: 6 Months

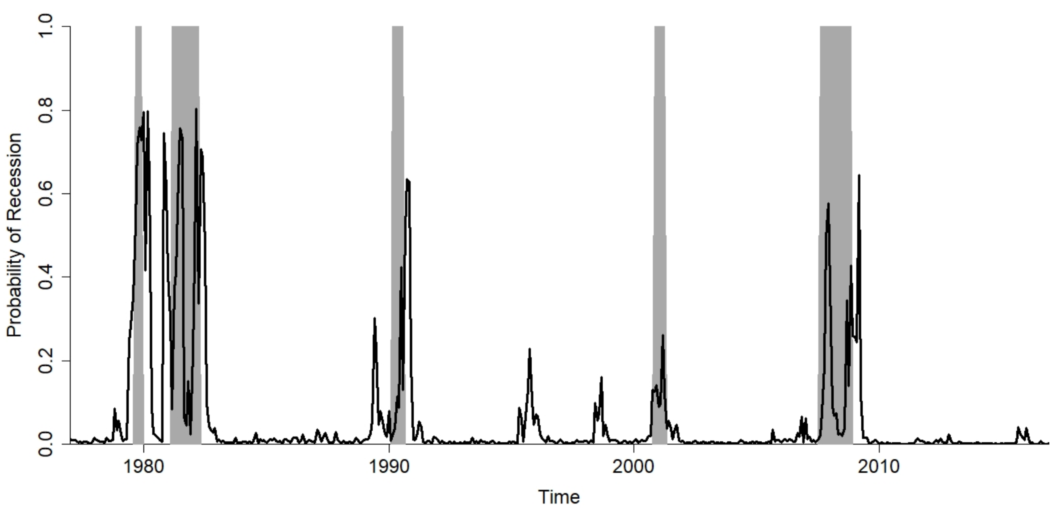

Note: Shaded areas indicate recessions

Figure 6: Changing Relative Importance of Leading Indicators

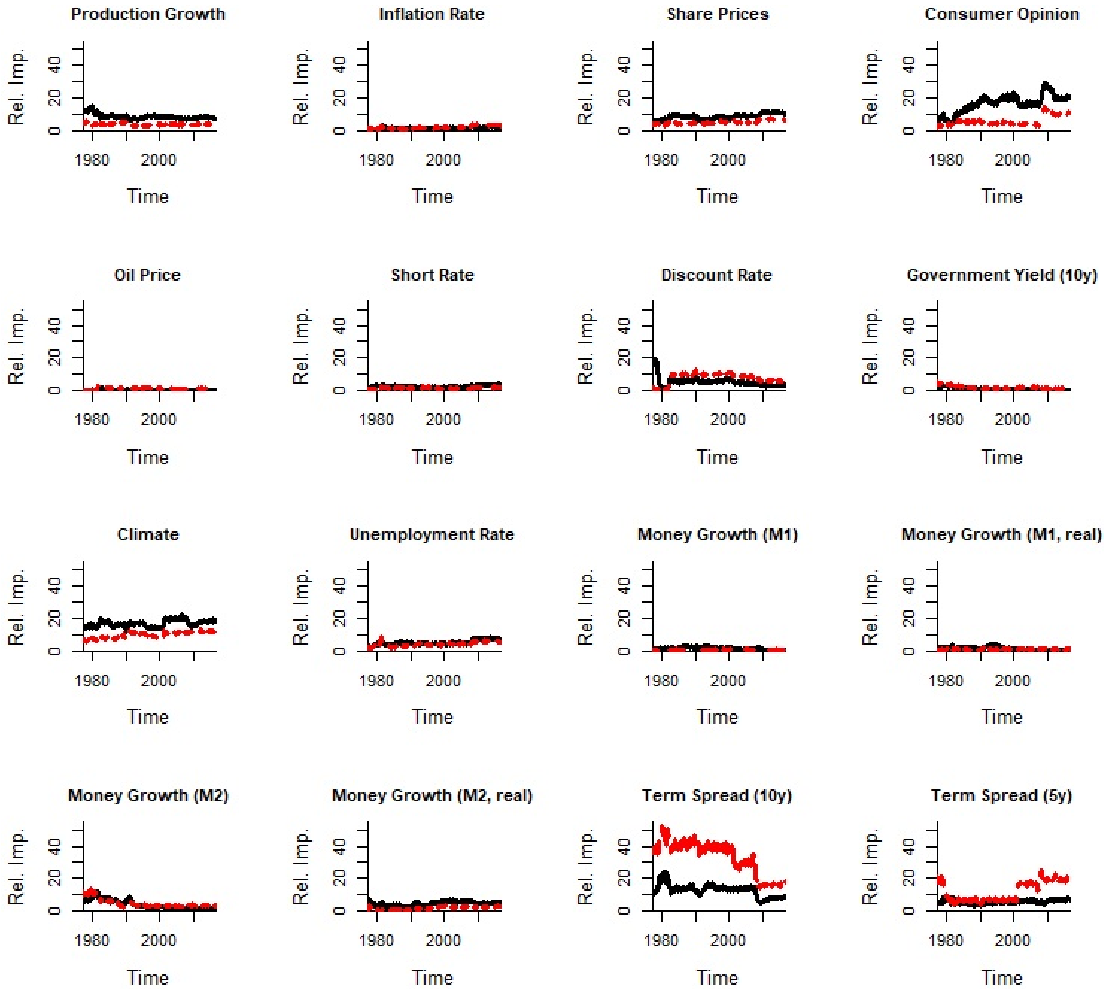

Note: Black (red) solid (dashed) lines: Forecast horizon 3 (6) months.