Bank of Russia

The Central Bank of the Russian Federation

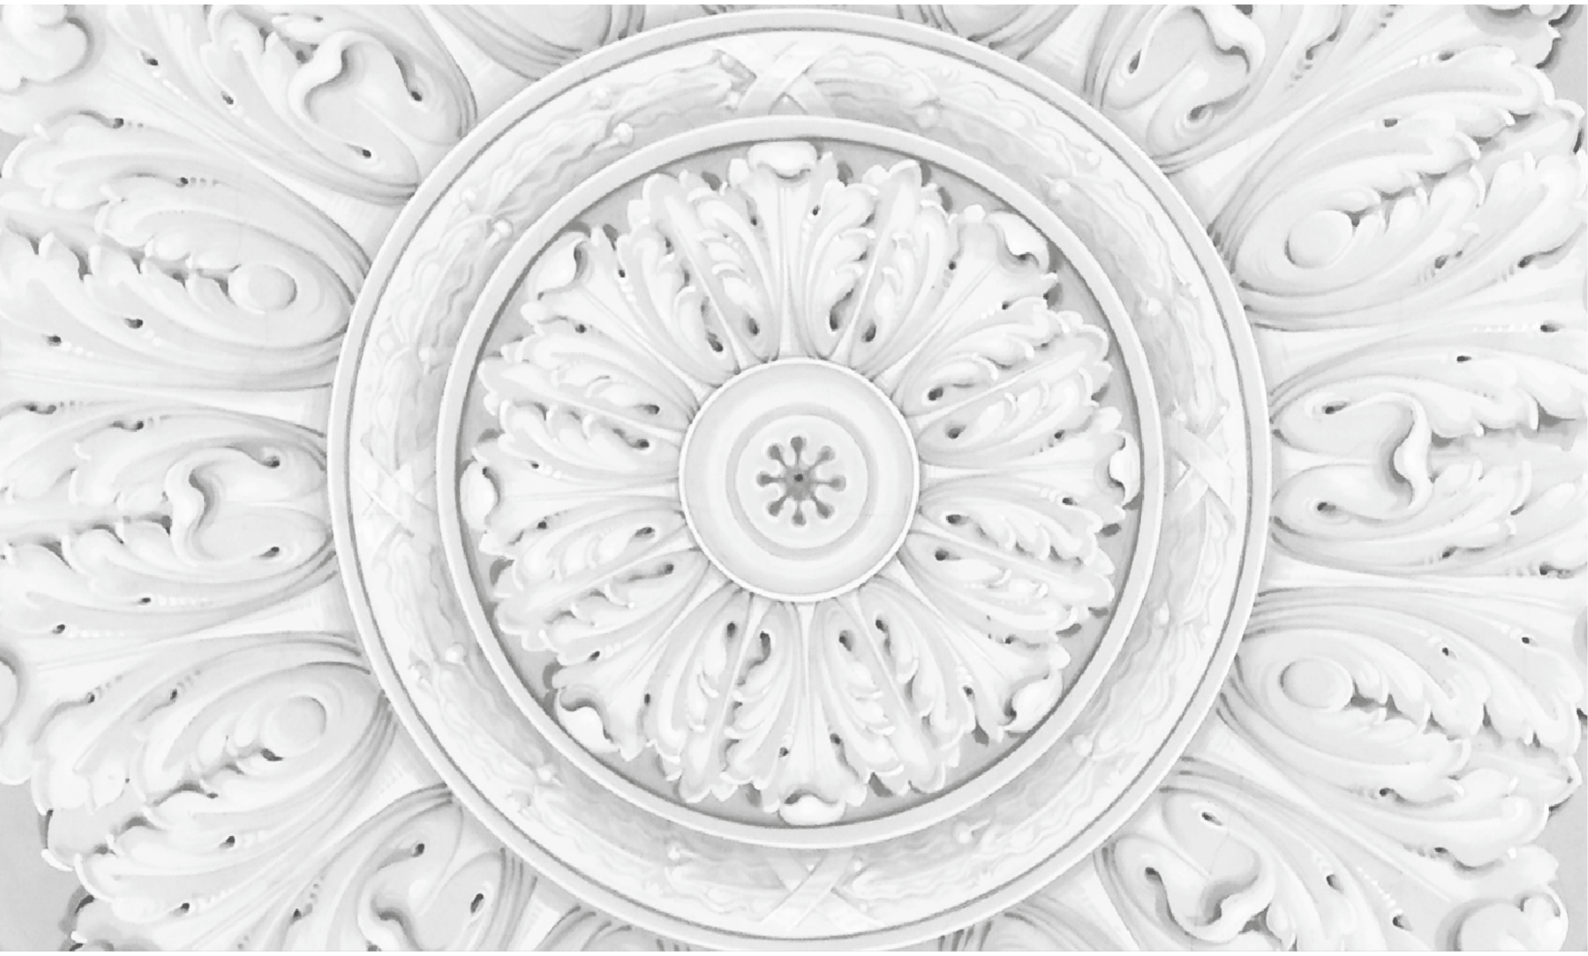

WORKING PAPER SERIES

Konstantin Styrin

Forecasting inflation in Russia by Dynamic Model Averaging

No. 39 / December 2018

Konstantin Styrin
Bank of Russia. E-mail: styrinka@cbr.ru

Address 12 Neglinnaya street, Moscow, 107016
Telephone +7 495 771-91-00, +7 495 621-64-65 (fax)
Website www.cbr.ru

Abstract

In this study, I forecast CPI inflation in Russia by the method of Dynamic Model Averaging (Raftery et al., 2010; Koop and Korobilis, 2012) pseudo out-of-sample on historical data. This method can be viewed as an extension of the Bayesian Model Averaging where the identity of a model that generates data and model parameters are allowed to change over time. The DMA is shown not to produce forecasts superior to simpler benchmarks even if a subset of individual predictors is pre-selected "with the benefit of hindsight" on the full sample. The two groups of predictors that feature the highest average values of the posterior inclusion probability are loans to non-financial firms and individuals along with actual and anticipated wages.

Keywords: Bayesian model averaging, model uncertainty, econometric modeling, high-dimension model, inflation forecast.

JEL classification: C5, C53, E37.

Forecasting Inflation in Russia by Dynamic Model Averaging

Konstantin Styrin*

December 14, 2018

Abstract

In this study, I forecast CPI inflation in Russia by the method of Dynamic Model Averaging (Raftery et al. (2010); Koop and Korobilis (2012)) pseudo out-of-sample on historical data. This method can be viewed as an extension of the Bayesian Model Averaging where the identity of a model that generates data and model parameters are allowed to change over time. The DMA is shown not to produce forecasts superior to simpler benchmarks even if a subset of individual predictors is pre-selected "with the benefit of hindsight" on the full sample. The two groups of predictors that feature the highest average values of the posterior inclusion probability are loans to non-financial firms and individuals along with actual and anticipated wages.

Keywords: Bayesian model averaging, model uncertainty, econometric modeling, high-dimension model, inflation forecast

JEL Codes: C5, C53, E37

*Bank of Russia, styrinka@cbr.ru

1 Introduction

In this study, I apply the method of Dynamic Model Averaging, DMA hereafter, to forecast monthly CPI inflation in Russia 1 through 6 months ahead. This method was developed by Raftery et al. (2010) and illustrated using an engineering application. Koop and Korobilis (2012) was the first paper that applied this methodology to forecast the US inflation. Koop and Korobilis (2011) use this method to forecast UK macroeconomic variables and compare its performance with other data-rich models. Byrne et al. (2018) is an application of DMA to forecasting exchange rates while Dangl and Halling (2012) to forecasting stock returns.

The DMA assumes that the data-generating process for inflation is not known with certainty. It states instead that, with some probability, inflation data can be generated by any of K alternative models, which differ in terms of a set of predictors employed. The DMA forecast on date t is obtained as a weighted average of forecasts produced by all alternative models with weights proportional to the predictive density of the respective individual model $k \in \{1, \ldots, K\}$ on date t. In that respect, the DMA is a generalization of the Bayesian Model Averaging, BMA hereafter, a conventional approach to dealing with model uncertainty in Bayesian econometrics – see, e.g., ch. 11 in Koop (2003) for a general discussion and Geweke and Whiteman (2006) for application to forecasting. Wright (2009) applies the BMA to forecast inflation in the US. Unlike BMA, the DMA method allows the identity of a model that generates data to randomly change over time. Furthermore, DMA feature time-varying parameters and stochastic volatility.

I conduct a standard pseudo-out-of-sample forecasting exercise for the DMA forecast and a set of benchmark forecast on data sample 2002M1-2017M9 using first 60 months of observations as an initial estimation sample. My findings suggest that the DMA does not produce forecasts superior to simpler benchmarks even if a subset of individual predictors is pre-selected "with the benefit of hindsight" on the full sample. The two groups of predictors that feature the highest average values of the posterior inclusion probability are loans to non-financial firms and individuals along with actual and anticipated wages.

The rest of the paper is organized as follows. Section 2 presents the methodology of the DMA and lays out details of the pseudo-out-of-sample forecasting exercise. Section 3 describes data. Section 4 exhibits and discusses empirical findings, and Section 5 concludes.

2 Methodology

2.1 Dynamic Model Averaging

The description of the DMA method follows closely Koop and Korobilis (2012). I denote by y_t the variable to be forecast. In this exercise, this variable is the monthly rate of CPI inflation. If the forecast horizon is h, then

$$y_{t+h} = 1200 \log \left(\frac{CPI_{t+h}}{CPI_{t+h-1}} \right),$$

i.e. y_t is the annualized growth rate of the Consumer Price Index (CPI) between dates $t+h-1$ and $t+h$. The vector of all potential predictors is z_t. It contains lagged values of monthly CPI inflation and other macroeconomic variables that potentially have some forecasting value with regard to future inflation. I consider all alternative forecasting models, each using a different subset of z_t as predictors. If the dimensionality of z_t is M, then the overall number of all subsets of z_t is $K = 2^M$, and so is the number of all alternative forecasting models that can be constructed from z_t. There is no prior knowledge of which particular combination of predictors yields the best forecast. Furthermore, the performance of each individual forecasting model can change over time.

Suppose that model k uses a subset $z_t^{(k)} \subseteq z_t$, $k = 1, \ldots, K$, as predictors. Then it can be written as

$$y_t = z_t^{(k)\prime} \theta_t^{(k)} + \epsilon_t^{(k)} \tag{1}$$

where $\theta_t^{(k)}$ is a vector of time-varying unobserved parameters that follows a vector random walk process

$$\theta_{t+1}^{(k)} = \theta_t^{(k)} + \eta_{t+1}^{(k)} \tag{2}$$

with a disturbance term being multivariate normal:

$$\eta_t^{(k)} \sim \text{i.i.d.}\mathcal{N}(0, Q_t^{(k)})$$

In the language of state-space models, equation (2) is called the law of motion for the unobserved state $\theta_t^{(k)}$ or, simply, state equation, whereas equation (1) is called measurement equation. The error term in the measurement equation (2) is assumed to be normal:

$$\epsilon_t^{(k)} \sim \text{i.i.d.}\mathcal{N}(0, H_t^{(k)})$$

Denote by y^t the history of observations of the CPI inflation up to date t, $y^t \equiv (y_t, y_{t-1}, \ldots, y_0)$, and by L_t the identity of the model that generates data on date t. Conditional on model k being the data-generating process (DGP hereafter) on date $t-1$ and given the history of observations up to date $t-1$, the vector of parameters has a multivariate normal distribution

$$\theta_{t-1} \,\big|\, L_{t-1} = k, y^{t-1} \sim \mathcal{N}\left(\hat{\theta}_{t-1}^{(k)}, \Sigma_{t-1|t-1}^{(k)}\right)$$

where, with some abuse of notation,

$$\hat{\theta}_t^{(k)} \equiv \mathbb{E}(\theta_t \,\big|\, L_t = k, y^t\,),$$

$$\Sigma_{t|t}^{(k)} \equiv \mathbb{E}\left[(\theta_t - \hat{\theta}_t^{(k)})(\theta_t - \hat{\theta}_t^{(k)})' \,\big|\, L_t = k, y^t\right]$$

Conditional on model k being the DGP on the next date, date t, a one-step ahead forecast of θ_t is

$$\theta_t \,\big|\, L_t = k, y^{t-1} \sim \mathcal{N}\left(\hat{\theta}_{t-1}^{(k)}, \Sigma_{t|t-1}^{(k)}\right)$$

where

$$\Sigma_{t|t-1}^{(k)} \equiv \mathbb{E}\left[(\theta_t - \hat{\theta}_{t-1}^{(k)})(\theta_t - \hat{\theta}_{t-1}^{(k)})' \,\middle|\, L_t = k, y^{t-1}\right]$$

Finally, an updated estimate of θ_t that exploits the information available on date t and is conditional on model k being the DGP on date t, i.e. the nowcast of θ_t, is

$$\theta_t \,\middle|\, L_t = k, y^t \sim \mathcal{N}\left(\hat{\theta}_t^{(k)}, \Sigma_{t|t}^{(k)}\right)$$

Given the law of motion for $\theta_t^{(k)}$, equation (2), the conditional variances of the one-step-ahead forecast and nowcast, both conditional on model k being the DGP on date t, are related as

$$\Sigma_{t|t-1}^{(k)} = \Sigma_{t-1|t-1}^{(k)} + Q_t^{(k)}$$

where the covariance matrix of the disturbance term in state equation (2), $Q_t^{(k)}$, is unobserved and, hence, has to be estimated. Following Raftery et al. (2010) and Koop and Korobilis (2012), I use an approximation by assuming that

$$\Sigma_{t|t-1}^{(k)} = \frac{1}{\lambda}\Sigma_{t-1|t-1}^{(k)} \tag{3}$$

where smoothing parameter λ is set very close to one from below, $0 < \lambda \lesssim 1$.

An application of standard Kalman filter formulas yields an updating equation for $\theta_t^{(k)}$ and $\Sigma_{t|t}^{(k)}$:

$$\hat{\theta}_t^{(k)} = \hat{\theta}_{t-1}^{(k)} + \Sigma_{t|t-1}^{(k)} z_t^{(k)} \left(H_t^{(k)} + z_t^{(k)\prime} \Sigma_{t|t-1}^{(k)} z_t^{(k)}\right)^{-1} \left(y_t - z_t^{(k)\prime} \hat{\theta}_{t-1}^{(k)}\right)$$

$$\Sigma_{t|t}^{(k)} = \Sigma_{t|t-1}^{(k)} - \Sigma_{t|t-1}^{(k)} z_t^{(k)} \left(H_t^{(k)} + z_t^{(k)\prime} \Sigma_{t|t-1}^{(k)} z_t^{(k)}\right)^{-1} z_t^{(k)\prime} \Sigma_{t|t-1}^{(k)}$$

The probability density of θ_{t-1} conditional on the history of observations up to date $t-1$ is obtained as a weighted average of model-specific nowcast densities:

$$p(\theta_{t-1} \,|\, y^{t-1}) = \sum_{k=1}^{K} p(\theta_{t-1} \,|\, L_{t-1} = k, y^{t-1}) \Pr(L_{t-1} = k, y^{t-1})$$

where

$$p(\theta_{t-1} \left| L_{t-1} = k, y^{t-1} \right.) \sim \mathcal{N}\left(\hat{\theta}^{(k)}_{t-1}, \Sigma^{(k)}_{t-1|t-1}\right)$$

Now I introduce some new notation for the probability of model l being the DGP on date t conditional on the history of observations up to date s, $s \leq t$:

$$\pi_{t|s,l} \equiv \Pr\left(L_t = l \left| y^s \right.\right)$$

It follows that

$$\Pr\left(L_{t-1} = k \left| y^{t-1} \right.\right) = \pi_{t-1|t-1,k}$$

The DMA framework assumes that the identity of the data-generating model can be different on different dates. In general, the process of switching between alternative data-generating models can be characterized by an unrestricted matrix of transition probabilities

$$\mathcal{P} = (p_{kl})^K_{k,l=1}$$

where p_{kl} is the probability of model k being the DGP on date t conditional on model l being the DGP on date $t-1$, and $\sum_{k=1}^K p_{kl} = 1$. The model prediction equation can then be written as

$$\pi_{t|t-1,k} = \sum_{l=1}^{K} \pi_{t-1|t-1,l}\, p_{kl} \tag{4}$$

The difficulty with the unrestricted matrix $\mathcal{P}$ is that it introduces a big number of additional parameters to be estimated. These parameters are transition probabilities p_{kl}. As a result, the estimation problem is difficult to handle even when the number of predictors in z_t is moderate. In order to circumvent this obstacle, Raftery et al. (2010) replace equation (4) by an approximation:

$$\pi_{t|t-1,k} = \frac{\pi^{\alpha}_{t-1|t-1,k}}{\sum_{l=1}^K \pi^{\alpha}_{t-1|t-1,l}}$$

where α is a forgetting factor that is set very close to one from below, $0 < \alpha \lesssim 1$.

Given $\pi_{t|t-1,k}$ and y^t, the updating equation for $\pi_{t|t,k}$ is

$$\pi_{t|t,k} = \frac{\pi_{t|t-1,k} p_k(y_t|y^{t-1})}{\sum_{l=1}^{K} \pi_{t|t-1,l} p_l(y_t|y^{t-1})}$$

where

$$y_t|L_t = k, y^{t-1} \sim p_k(y_t\left|y^{t-1}\right.) \equiv \mathcal{N}(z_t^{(k)\prime}\hat{\theta}_t^{(k)}, H_t^{(k)} + z_t^{(k)\prime}\Sigma_{t|t-1}^{(k)} z_t^{(k)})$$

The DMA recursive point forecast is then obtained by averaging of point forecasts of all individual models with weights equal to conditional probabilities of the respective model being the DGP on the date for which the forecast is made:

$$\mathbb{E}(y_t|y^{t-1}) = \sum_{k=1}^{K} \pi_{t|t-1,k} z_t^{(k)\prime}\hat{\theta}_{t-1}^{(k)} \tag{5}$$

Along with the DMA forecast, I also compute the DMS point forecast, which is the point forecast of the model with the highest conditional probability of being the DGP on the date for which the forecast is made:

$$k^* = \text{argmax}_k \pi_{t|t-1,k}$$

$$\mathbb{E}(y_t|y^{t-1}) = z_t^{(k^*)\prime}\hat{\theta}_{t-1}^{(k^*)}$$

It is straightforward to see that, given $H_t^{(k)}$, $\pi_{0|0,k}$ and $\theta_0^{(k)}$, $k = 1, \ldots, K$, one can apply derived analytical formulas to compute forecasts with no need to run MCMC posterior simulations, which are likely to be computationally demanding and time consuming in this case. This lightens the computational burden significantly.

Following Koop and Korobilis (2012), the variance of the error term in the measurement equation (1) is estimated as

$$\hat{H}_t^{(k)} = \kappa \hat{H}_{t-1}^{(k)} + (1-\kappa)\left(y_t - z_t^{(k)\prime}\hat{\theta}_t^{(k)}\right)^2$$

where κ is a decay factor, $0 < \kappa < 1$.

Dangl and Halling (2012) suggest one more layer of averaging – by forgetting factor λ in equation (3). A higher value of λ implies that the uncertainty introduced by an innovation in the error term of state equation (2), η_{t+1}, is small. This might be true in tranquil times but not in times of turbulence. In order to account for time variation of uncertainty caused by innovations to the vector of coefficients θ_t, Dangl and Halling (2012) allow the foregetting factor λ to vary over time. This is implemented as the following. They consider a grid of values for λ, say, $\lambda_j \in \{0.90, 0.91, \ldots, 1.00\}$. Conditional on λ_j, $j = 1, 2, \ldots, J$, a forecast is obtained by formula (5). Finally, all forecasts $\mathbb{E}\left(y_t \left| y^{t-1}, \lambda_j\right.\right)$ are averaged with weights equal to posterior probabilities $\Pr\left(\lambda_j \left| y^{t-1}\right.\right)$ to yield an ultimate forecast:

$$\mathbb{E}\left(y_t \left| y^{t-1}\right.\right) = \sum_{j=1}^{J} \Pr\left(\lambda_j \left| y^{t-1}\right.\right) \mathbb{E}\left(y_t \left| y^{t-1}, \lambda_j\right.\right)$$

2.2 Pseudo out-of-sample forecasting

With the data described in the next section, I evaluate the retrospective performance of forecast based on DMA and DMS and compare it with the performance of a few benchmark forecasts. The benchmark forecasts that I employ are those produced by Bayesian Model Averaging and Bayesian Model Selection (Koop and Potter (2004); Wright (2009)), Unobserved Components – Stochastic Volatility model (Stock and Watson (2007)), Bayesian autoregression of order 2, separately, with time-invariant and time-varying coefficients, DMA applied to a minimalist set of predictors that includes an intercept and two lags of inflation, and, finally, the "kitchen sink" model with time-varying parameters.

Taken literally, the DMA considers as an individual model any possible mix of predictors from the pool with the size varying from one – a time-varying intercept and one predictor with a time-varying coefficient – to M where all predictors from the pool along with the intercept are included (at least, in the implementation of Koop and Korobilis (2012)). DMA then averages out over those individual models with weights equal to posteriors probabilities of a respective model being the data-generating process on a given date. The overall number of individual models equals $K = 2^M$. A data set containing $M = 40$ time series is quite moderate by modern standards. For example, a popular data set composed by Stock and Watson (2012) contains about 200 quarterly U.S. macro series. If $M = 40$, i.e. five times fewer than the size of Stock and Watson (2012)'s dataset, the number of models $K = 2^M$ is of the order of 10^{12}. Handling such a big number of models requires quite a lot of computational power and prohibitive amount of time needed for computations on a research-purpose desktop or laptop computer with typical technical characteristics. For this project, I have employed the *R* package called *eDMA* (Catania and Nonejad (2018)), which features efficient algorithms and parallel computations. If $M = 20$ and the number of time observations is 189, my ASUS laptop with Intel Core i9 12-core processor and 64 GB of RAM completes all computations for the main exercise within about 10 minutes. Adding each extra predictor doubles the number of models and, hence, computation time. Once the number of predictors increases by 20, i.e. from $M = 20$ to $M = 40$, then the computation time rises by the factor of $2^{20} \approx 10^6$ and approaches $1/6 \times 10^6 \approx 19$ years, which is prohibitively high.

Several approaches aiming to keep the number of individual models manageable are known in the literature. In an engineering application, Raftery et al. (2010) employ 5 predictors and consider only 17 their combinations as candidate DGPs, the selection being based on some external information about physical processes involved. In his inflation forecasting exercise by BMA, Wright (2009) considers only individual models each containing an intercept, lagged inflation, and one predictor thus limiting the number of models to $K = M$. Other papers impose an upper limit on the number of exogenous predictors: $M = 14$ in Koop and Korobilis (2012), and $M = 15$ in Groen et al. (2013). Both papers forecast inflation in the U.S. Onorante and Raftery (2016) nowcast Euro area GDP employing 30 predictors. The curse of model space dimensionality is circumvented by introducing stochastic search over the space of models. An important assumption needed for the approximation to be well-grounded is that DGP has to smoothly transit from one model to another (Catania and Nonejad (2018)). The inspection of posterior inclusion probabilities in Onorante and Raftery

(2016) and other papers, which feature occasional swings, suggests that this is unlikely to be the case. The interim conclusion seems to be that, given the present state of available computational power, the number of predictors that DMA is able to accommodate for practical purposes is unlikely to exceed 20-25.

In this study, I use a data set of 97 macro series for Russia that cover the time period January 2002 through September 2017, 189 time observations in total. In order to make the estimation manageable, I give my DMA forecast a certain "benefit of hindsight": for each forecast horizon, I pre-select a set of best performing predictors based on the so-called "hard thresholding" selection procedure (Ng (2013)). According to this method, I run 97 predictive regressions on the full sample with monthly inflation as the dependent variable and one variable at a time from the list of predictors as a regressor, in addition to an intercept. I then rank all predictors according to the absolute value of their t-statistic in a descending order. For two different forecast horizons $h = 1, 2, \ldots, 6$, the respective rankings need not coincide. I then pick 19 top predictors based on the absolute value of t-statistic for each h add the contemporaneous and lagged inflation y_{t-1} to it and consider only $K = 2^{19}$, each containing an intercept, contemporaneous inflation y_t, and up to 19 other predictors, the dependent variable being y_{t+h}. For this set of models, I then do a pseudo out-of-sample forecasting evaluation. The whole exercise is by no means *purely* pseudo out-of-sample forecasting since, as I already mentioned, the pre-selection of 19 out of 97 predictors is done on the full sample. If it turns out that, even with the benefit of hindsight, DMA does not yield considerable improvements over simpler benchmarks (which appears to be true), then this will make the case against DMA even stronger.

The design of the (post-pre-selection part of) the recursive pseudo-out-of-sample experiment is standard. First 60 time observations are reserved for initial estimation of the model. The first one-month ahead forecast is thus made for March 2007, two-month-ahead forecast for April 2007, etc. The last forecast is made for September 2017. Forecast horizons are $h = 1, 2, \ldots, 6$ months. After an h-month-ahead forecast from date t for date $t + h$

is obtained, date $t+1$ observation is added to the estimation sample, parameters of the forecasting model in hands are updated, and forecasts for date $t+1+h$ are produced from the perspective of date $t+1$. The procedure continues until the end of the sample period is reached. In order to avoid complications, I seasonally adjust all series on the full sample and do not address the "ragged edge" problem. As for the latter, it is well known that different series are updated with different delays. Most recent readings of financial series are available immediately whereas for data collected by national statistical agencies (Rosstat in Russia) or survey companies, it might take up to 2 months for the data to get prepared for publication. To make things simple, I pretend that February 2007 readings of all predictors are available when I am about to make a one-month-ahead forecast for March 2007, although, in reality, a February update of Rosstat data is released only about March 23. Seasonal adjustment on the full sample and ignorance of the "ragged edge" problem are thus another "benefit of hindsight" granted to my DMA forecast. Last but perhaps not least, my data set contains revised macro series that are different from a historical real-time data.

2.3 Alternative forecasting models

Below is the complete list of models that I employ in this study. In addition to DMA/DMS, it also includes benchmark forecasts against which I compare the out-of-sample forecast performance of DMA/DMS.

Dynamic Model Averaging (DMA). Following the literature, I set $\alpha = 0.99$ and $\kappa = 0.96$ whereas λ takes its values on a grid $(0.90, 0, 01, \ldots, 1.00)$.

Dynamic Model Selection (DMS). This is the same as DMA except that, instead of averaging over all models, only one model – with the highest posterior likelihood – is chosen for a forecast. The parametrization is similar to that of DMA.

Bayesian Model Averaging (BMA). DMA can be viewed as an extension of BMA. Similar to DMA, BMA assumes that the there is uncertainty with regard to the model that generates data and therefore considers several candidate models. Unlike DMA, though, BMA postulates that the identity of the data generating process is fixed but unknown and that model parameters do not change over times. BMA can be viewed as a special case of DMA with $\alpha = 1$, $\kappa = 1$, and $\lambda = 1$.

Bayesian Model Selection. Similar to BMA except that, instead of averaging over all models, only one model – with the highest posterior likelihood – is chosen for a forecast. The parametrization is similar to that of BMA.

Unobserved Components – Stochastic Volatility Model (UC-SV). This model gained a good reputation for forecasting inflation in the U.S. (Stock and Watson (2007)). It states that the inflation consists of persistent and transitory components:

$$y_t = \tau_t + \eta_t$$

$$\tau_t = \tau_{t-1} + \epsilon_t$$

where η_t and ϵ_t are two independent i.i.d. Gaussian processes with time-varying covariance matrices, respectively, Q_t and R_t where

$$\log(Q_t) = \log(Q_{t-1}) + u_t$$

$$\log(R_t) = \log(R_{t-1}) + v_t$$

with u_t and v_t being two independent i.i.d. Gaussian processes. The UC-SV forecast is $\mathbb{E}_t y_{t+h} = \mathbb{E}_t \tau_t$.

AR(2). Bayesian autoregression of order 2.

TVP-AR(2). Time-varying autoregression of order 2.

DMA-AR(2). This is a special case of the DMA model where only two lags of inflation are available as predictors.

TVP-KS. This is a "kitchen sink" regression that contains all predictors from the list with parameters allowed to change over time.

3 Data

In this study I employ monthly data covering the period from January 2002 through September 2017. The variable to be forecast is consumer price inflation. There are 97 exogenous predictors (as labeled in Koop and Korobilis (2012) – in the sense that neither of them is a lag of the variable to be forecast) that are listed and described in Table 1. As additional predictors, I also include a time-varying intercept and two lags of inflation. An intercept and two lags of inflation are set to be included in each individual model within DMA.

[TABLE 1 ABOUT HERE]

The macroeconomic variables that I consider as potential predictors are

- domestic prices of goods and services: PPI, cargo tariffs, etc.;
- money and credit: monetary aggregates and Bank of Russia's international reserves, credit to individuals and firms, etc.;
- labor market indicators: unemployment rate, number of employed;
- real economic activity indicators: retail sales, retail services, wholesale sales, investment, real disposable income, real wages, production of new houses, cargo shipments, etc.;
- financial market indicators: real and nominal exchange rates, stock market index, interest rates, etc.;
- survey indicators: various versions (composite, manufacturing, services, etc.) of the Purchase Management Institute Index (PMI) produced by Markit, a London-based consultancy firm; Russian Economic Barometer (REB) industry survey: current and anticipated prices for output and inputs, planned purchases of equipment, etc;
- commodity market indicators: international prices of oil, aluminum, wheat, etc.

- indicators of economic activity in systemic economies such as the U.S. and Euro area.

The rationale behind considering non-consumer prices as predictors is straightforward. Producer prices tend to be more flexible (or less sticky) empirically. They respond faster to macroeconomic shocks than consumer prices and can therefore, in theory at least, serve as a leading indicator with respect to the latter. A rise in transportation costs leads to an increase in all consumer prices, perhaps, with some delay. Lagged CPI inflation predicts well future CPI inflation, empirically. Furthermore, such a forecast, which is consistent with CPI inflation following a random walk, is very difficult to outperform with a more sophisticated model for the U.S. and other advanced economies (Atkeson and Ohanian (2001); Stock and Watson (2007); Faust and Wright (2013)). Activity variables convey information about the state of aggregate demand. Unusually strong growth in economic activity may signal about the ongoing rise in production costs and the resulting accumulation of inflationary pressures on prices. Exchange rates are potentially informative because of the incomplete and gradual exchange rate pass-through of their changes into retail prices of imported consumer goods, which is documented empirically (Burstein and Gopinath (2014)). Furthermore, along with other asset prices, exchange rates are essentially forward-looking variables that should respond to shifts in expected time path of future inflation and an anticipated response to it from monetary policy.

All variables except CPI inflation and financial market indicators – interest rates, stock market index, and exchange rates – were seasonally adjusted using the U.S. Bureau of Census X-13-ARIMA-SEATS seasonal filter as implemented in the *R* package *seasonal* (`http://www.seasonal.website`). The CPI inflation series was seasonally adjusted using the methodology adopted in the Bank of Russia (Sapova et al. (2018)). As it was already mentioned, seasonal adjustment was done on the full sample.

All predictors containing a unit root – quantities and prices in levels – were transformed to approximately stationary by log-differencing.

4 Findings

To evaluate the performance of forecasts, I employ two standard metrics: root mean squared forecast error (RMSFE) and mean absolute foecast error (MAFE). Tables 2 to 3 report RMSFE and MAFE for DMA/DMS forecasts and competitors.

[TABLE 2 ABOUT HERE]

Inspection of Table 2 suggests that the quality of DMA/DMS is not superior to benchmark forecasts. These two model outperform the rivals only at horizon three months with RMSFE equal to 5.1 percentage points (p.p.). For the forecast horizon $h = 1$ month, the best performer is the Bayesian Model Averaging (BMA) with RMSFE = 3.5 p.p. The precision of the DMA forecast is almost as good as that of BMA with RMSFE = 3.6 p.p. Surprisingly, for $h = 2$, the best forecast is delivered by the Time-Varying Parameter Kitchen Sink model where all predictors from the list are involved. This is unusual since, generally, kitchen sink regressions tend to produce inferior forecast due to a large number of parameters to be estimated. As in the previous case, the DMA/DMS forecast is roughly at par with the top performer yielding RMSFE equal to 4.7 and 4.8 p.p., respectively, against 4.6 of TVP-KS. For $h = 4, 5, 6$, the best forecast is generated by DMA-AR(2), which is lagging behind DFM/DMS for $h = 1, 2, 3$. Its RMSFE is steadily 5.3 p.p. against 5.4 p.p. of DMA and 5.7 to 6.0 p.p. of DMS. Overall, the DMA/DMS does not demonstrate a systematic advantage over other methods. Even for $h = 3$ where it beats its rivals, its advantage is marginal – just 0.3 p.p. relative to the closest follower. It is worth recalling that, in this exercise, DMA/DMS is given an important benefit of hindsight: for each h, the set of predictors involved are pre-selected on the full sample based on the hard thresholding method (Ng (2013)). Overall, one has to admit that the precision of DMA/DMS as well as their rivals is quite poor in absolute terms: the width of a 67% forecast interval is about 7 to 11 p.p., depending on the forecast horizon, which is far from satisfactory for practical purposes.

[TABLE 3] ABOUT HERE

Based on the MAFE metrics, the rankings almost do not change. DMA dominates the alternatives only for the forecast horizon $h = 3$. TVP-AR(2) is the best performer for $h = 1$ month, TVP-KS for $h = 2$ months, and DMA-AR(2) for $h = 4, 5, 6$ months. Again, DMA/DMS is almost at par with the winner for each forecast horizon with differences in MAFE being marginal.

Tables 4 to 9 and Figures 1 to 6 summarize findings related to the informativeness of each individual pre-selected predictor.

[TABLE 4] AND FIGURE 1 ABOUT HERE

For each pre-selected predictor, Tables 4 to 9 report the following statistics: (i) the sample mean of the coefficient on this predictor, (ii) the standard deviation of the coefficient, (iii) the sample mean of the posterior probability of the inclusion of this predictor to the data generating process, and (iv) the sample standard deviation of the inclusion probability. Depending on the forecast horizon, 3 to 9 pre-selected predictors have sample mean posterior inclusion probability of 0.25 or above.

As shown in Table 4, for $h = 1$ month, the most informative predictors are nominal effective exchange rate (f10), the output-to-input price ratio from the REB industrial survey (s36), the growth rate in credit to non-financial enterprises with maturity up to 1 year (f6), and the international price of wheat (f21). According to Figure 1, the informativeness of the top four predictors was not uniform over the sample period. At the beginning of the sample period, the inclusion probability for each of them was close to zero. The inclusion probability for the credit growth features a spike around 2010 and then goes below 0.30 quite fast whereas the three other predictors gain importance with the inclusion probability remaining almost about 0.9 for the nominal effective exchange rate after 2009, between 0.6-0.9 for the price of wheat starting from 2011, and in the range between 0.3-0.6 for the output-to-input price ratio in the industry after mid-2013.

[TABLE 5] AND FIGURE 2 ABOUT HERE

As it is shown on Table 5, for $h = 2$, the top predictors are the output-to-input price ratio (s36) and the anticipated growth rate of wages (s41) from the REB industrial survey and also the credit growth to non-financial enterprises (f6 and f7). According to Figure 2, the posterior probability of inclusion for loans, f6 and f7, tends to be high, 0.4 to 0.6, in the middle of the sample and declines toward the end of the sample. The inclusion probability of the two survey indicators, s36 and s41 instead grows over time reaching one in the case of the output-to-input price ratio (s36).

[TABLE 6] AND FIGURE 3 ABOUT HERE

Table 6 and Figure 3 summarize findings for $h = 3$. The most informative predictors are loans to individuals with maturity beyond one year (f9), overall loans to individuals (w4), monetary aggregate M0 (f13), and nominal wages (r24). The posterior inclusion probability for the two loan variables is moderate, below 0.4, between 2007 and 2015, but reaches one very fast afterwards. The predictive power of M0 and nominal wages is rather high, 0.6-0.7 and 0.5-0.6, respectively, before 2011 and then gradually declines.

[TABLE 7] AND FIGURE 4 ABOUT HERE

At the horizon $h = 4$ months, as Table 7 and Figure 4 illustrate, the most valuable predictors appear to be loans to non-financial enterprises and individuals with maturity longer than one year (f7 and f9, respectively), overall loans to individuals (w4), and expected wage inflation in manufacturing (s41). The posterior inclusion probability behaves quite unevenly for the three loan variables featuring spikes and abrupt declines. The expected growth of wages being uninformative before 2010 steadily gains predictive power afterwards: its posterior inclusion probability fluctuates between 0.5-0.9 on the second half of the sample period.

[TABLE 8] AND FIGURE 5 ABOUT HERE

Out-of-sample forecasting results for $h = 5$ are reported in Table 8 and Figure 5. The top predictors are loans to non-financial enterprises with maturity up to one year (f6), monetary aggregate M0 (f13), nominal wages (r24), and financial health of firms in manufacturing (s45). The predictive content of the loans and nominal wages decreases over time, with the inclusion probability declining from 0.4 at the beginning of the sample to 0.2 toward the end of the sample. M0 appears to be highly informative only occasionally, in 2015-2016, when its posterior inclusion probability experiences a spike to 0.6-0.7 and then drops to 0.3-0.4. The forecasting value of nominal wages also proved to be temporary: the inclusion probability is close to zero at the beginning of the sample period, it grows rapidly after 2010 reaching 0.6 in 2012 and then starts a steady descend approaching 0.1 at the end of the sample.

[TABLE 9] AND FIGURE 6 ABOUT HERE

Finally, Table 9 and Figure 6 contain forecast evaluation results for the longest forecast horizon we consider, $h = 6$ months. The leading predictors are nominal and real wages (r24 and r25), PMI index of input prices in services (s27) and expected purchases of equipment from REB industrial survey (s44). The posterior inclusion probability of wages wildly fluctuates over time in the range between 0.2 to 0.5. The PMI index remained almost uninformative until 2015 when its inclusion probability grew rapidly from 0.1 to almost 1.0 within a couple of months and then declined gradually to 0.6 during 2017.

To summarize, two groups of predictors tend to receive high posterior weights rather frequently. The first group is loans to non-financial firms and individuals. For all values of the forecast horizon, at least one member of this group is among four top performing predictors. Second, wages, either actual or anticipated, are among most informative predictors for forecast horizon of two months and longer. Occasional top predictors are the following. Monetary aggregate M0 appears to be helpful in forecasting inflation 3 and 5 months ahead. The output-to-input price ratio in manufacturing from the REB industrial survey is among best performers at horizons one and two months. International price of wheat and nominal effective exchange rate demonstrate a high predictive content in forecasting one month ahead. The latter might reflect a pass-through of changes in the ruble exchange rate into retail prices of imported goods although, as documented in the literature, the exchange rate pass-through effect is more pronounced at horizons longer than one month (Burstein and Gopinath (2014)). Finally, financial health of firms in manufacturing is a highly informative predictor at the horizon of 5 months whereas the PMI index of input prices in services along with expected purchases of equipment from REB industrial survey at the horizon of 6 months.

[TABLE 10 ABOUT HERE

Table 10 presents a decomposition of the variance of the forecast error into four components (Dangl and Halling (2012)). The first component, labeled as "Observations" is related to the variability of data. The second, "Coefficients" arises because the parameters of the model are estimated with error. The third component, labeled as "Model", accounts for the uncertainty about data generating process. The final component, "Time-varying parameters" reflects time variability of the coefficients in DMA. The variance decomposition is reported separately for DMA and BMA for six values of the forecast horizon. The inspection of Table 10 suggests that the variability of the data-generating process is an important factor of forecast errors, accounting for 10 to 17 percent of the forecast error variance. For BMA the importance of model uncertainty is much lower with the contribution not exceeding 4 percent. This is by construction though because, unlike DMA, this method assumes that the identity of the individual model that generates data is fixed but unknown. DMA, on the contrary, explicitly allows different models to generate data at different times with random switches among them. BMA also assumes that the parameters of the model are fixed, and therefore the last component of the variance decomposition, "Time-varying parameters," is always zero. For a similar reason, the contribution of estimation error for coefficient to the forecast error variance is much smaller for BMA than for DMA.

5 Conclusion

This paper applies the Dynamic Model Averaging method to forecasting monthly CPI inflation in Russia out of sample. Unlike superior performance of DMA documented in other studies, e.g., in Koop and Korobilis (2012) for the U.S. inflation, this method does not yield forecasts that would systematically beat simpler benchmarks in the case of Russia. Two groups of predictors feature the highest average values of the posterior inclusion probability. These are, first, loans to non-financial firms and individuals and, second, actual wages or anticipated wages from a survey. The former are likely to reflect inflation pressures originating from the aggregate demand whereas the latter are related to the cost of production. Among best performers, there is no single variable that remains evenly informative over the full sample period 2007-2018. A common pattern is that the posterior inclusion of a predictor is high over one subperiod and moderate to low over the rest of the sample. The forecast error variance decomposition suggests that the time variation in the identity of a model that generates data is a non-negligible source of forecast uncertainty.

References

Atkeson, A. and L. Ohanian (2001): "Are Phillips Curves Useful for Forecasting Inflation?" *Federal Reserve Bank of Minnesota Quarterly Review*, 25, 2–11.

Burstein, A. and G. Gopinath (2014): "International Prices and Exchange Rates," in *Handbook of International Economics, vol. 4*, ed. by G. Gopinath, E. Helpman, and K. Rogoff, North-Holland: Elsevier, 391–451.

Byrne, J., D. Korobilis, and P. Ribeiro (2018): "On the Sources of Uncertainty in Exchange Rate Predictability," *International Economic Review*, 59, 329–357.

Catania, L. and N. Nonejad (2018): "Dynamic Model Averaging for Practitioners in Economics and Finance: The eDMA Package," *Journal of Statistical Software*, 84, 1–39.

Dangl, T. and M. Halling (2012): "Predictive Regressions with Time-Varying Coefficients," *Journal of Financial Economics*, 106, 157–181.

Faust, J. and J. Wright (2013): "Forecasting Inflation," in *Handbook of Economic Forecasting, vol. 2*, ed. by G. Elliott and A. Timmerman, North-Holland: Elsevier.

Geweke, J. and C. Whiteman (2006): "Bayesian Forecasting," in *Handbook of Economic Forecasting, vol. 1*, ed. by G. Granger and A. Timmerman, North-Holland: Elsevier.

Groen, J., R. Paap, and F. Ravazzolo (2013): "Real-Time Inflation Forecasting in a Changing World," *Journal of Business and Economic Statistics*, 31, 29–44.

Koop, G. (2003): *Bayesian Econometrics*, Wiley.

Koop, G. and D. Korobilis (2011): "UK Macroeconomic Forecasting with Many Predictors: Which Models Forecast Best and When Do They Do So?" *Economic Modelling*, 28, 2307–2318.

——— (2012): "Forecasting Inflation Using Dynamic Model Averaging," *International Economic Review 53*, 867–886.

Koop, G. and S. Potter (2004): "Forecasting in Dynamic Factor Models Using Bayesian Model Averaging," *Econometrics Journal*, 7, 550–565.

Ng, S. (2013): "Variable Selection in Predictive Regressions," in *Handbook of Economic Forecasting, vol. 2*, ed. by G. Elliott and A. Timmermann, Elsevier, 752–789.

Onorante, L. and A. E. Raftery (2016): "Dynamic Model Averaging in Large Model Spaces Using Dynamic Occam's Window," *European Economic Review*, 81, 2–14.

Raftery, A., M. Kárný, and P. Ettler (2010): "Online Prediction under Model Uncertainty via Dynamic Model Averaging: Application to a Cold Rolling Mill," *Technometrics 52*, 52–66.

Sapova, A., A. Porshakov, A. Andreev, and E. Shatilo (2018): "Review of Methodological Specifics of Consumer Price Index Seasonal Adjustment in the Bank of Russia," Working paper no. 33, Bank of Russia.

Stock, J. H. and M. W. Watson (2007): "Has Inflation Become Harder to Forecast?" *Journal of Money, Credit, and Banking 39*, 3–34.

——— (2012): "Disentangling the Channels of the 2007-2009 Recession," *Brookings Papers on Economic Activity 42*, 81–135.

Wright, J. (2009): "Forecasting U.S. Inflation by Bayesian model averaging," *Journal of Forecasting 28*, 131–44.

Tables

Table 1: Description of predictors

Notation	Description	Source
f1	Interest rate on short-term ruble-denominated loans to individuals	BoR
f2	Interest rate on long-term ruble-denominated loans to individuals	BoR
f3	Interest rate on short-term ruble-denominated loans to non-financial enterprises	BoR
f4	Interest rate on long-term ruble-denominated loans to non-financial enterprises	BoR
f5	Interest rate on overnight interbank loans	BoR
f6	Loans to non-financial enterprises up to 1 year, % MoM	BoR
f7	Loans to non-financial enterprises beyond 1 year, % MoM	BoR
f8	Loans to individuals up to 1 year, % MoM	BoR
f9	Loans to individuals beyond 1 year, % MoM	BoR
f10	Nominal effective exchange rate, % MoM	BoR
f11	Real effective exchange rate, % MoM	BoR
f12	Monetary aggregate M2, % MoM	BoR
f13	Monetary aggregate M0, % MoM	BoR
f14	Bank of Russia's international reserves	BoR
f15	Interest rate on short-term deposits of individuals	BoR
f16	Interest rate on short-term deposits of non-financial enterprises	BoR
f17	Interest rate on long-term deposits of individuals	BoR
f18	Interest rate on long-term deposits of non-financial enterprises	BoR
f19	Moscow Exchange Stock Market Index, % MoM	MoEx
f20	International price of oil, % MoM	Bloomberg
f21	International price of wheat, % MoM	Bloomberg
f22	International price of natural gas, % MoM	Bloomberg
f23	International price of aluminum, % MoM	Bloomberg
f24	International price of nickel, % MoM	Bloomberg
f25	U.S. industrial production, % MoM	Bloomberg
f26	E.U. total orders in manufacturing	Bloomberg
f27	European Commission Manufacturing Confidence Index	Bloomberg
f28	U.S. ISM Manufacturing PMI SA	Bloomberg
f29	Industrial production in euro zone	Bloomberg
f30	Investment goods price deflator	Bloomberg

r1	Overall exports	Rosstat
r2	Exports to CIS	Rosstat
r3	Exports outside CIS	Rosstat
r4	OKVED industrial production index	Rosstat
r5	Industrial Production: Mining and Quarrying (NSA, 2005=100)	Rosstat
r6	Industrial Production: Manufacturing (NSA, 2005=100)	Rosstat
r7	Industrial Production: Electricity, Gas, and Water Supply (NSA, 2005=100)	Rosstat
r8	Industrial Production: Metallurgical Production and Finished Metalware (NSA, 2005=100)	Rosstat
r9	Industrial Production: Pulp, Paper, Publishing, and Printing (NSA, 2005=100)	Rosstat
r10	Industrial Production: Chemicals (NSA, 2005=100)	Rosstat
r11	Industrial Production: Coke and Petroleum Products (NSA, 2005=100)	Rosstat
r12	Industrial Production: Electrical and Optical Equipment (NSA, 2005=100)	Rosstat
r13	Industrial Production: Food, Beverages, and Tobacco (NSA, 2005=100)	Rosstat
r14	Industrial Production: Leather and Leather Products (NSA, 2005=100)	Rosstat
r15	Industrial Production: Other Nonmetallic Mineral Products (NSA, 2005=100)	Rosstat
r16	Industrial Production: Manufacture of Textiles (NSA, 2005=100)	Rosstat
r17	Industrial Production: Rubber and Plastic Products (NSA, 2005=100)	Rosstat
r18	Industrial Production: Transport Equipment (NSA, 2005=100)	Rosstat
r19	Industrial Production: Wood and Wood Products (NSA, 2005=100)	Rosstat
r20	Industrial Production: Machinery and Equipment n.e.c. (NSA, 2005=100)	Rosstat
r21	Output in Agriculture (NSA, % MoM)	Rosstat
r22	Output in Constructure	Rosstat
r23	Output of dwellings, % MoM	Rosstat
r24	Nominal wages	Rosstat
r25	Real wages	Rosstat
r26	Real Disposable Income (NSA, % MoM)	Rosstat

r27	Real retirement benefits paid	Rosstat
r28	Retail trade	Rosstat
r29	Retail trade: food, beverages, and tobacco	Rosstat
r30	Retail trade: non-food items	Rosstat
r31	Retail services	Rosstat
r32	Cargo shipments % MoM	Rosstat
r33	Cargo shipments by railroad	Rosstat
r34	Total Output – 5 Basic Indicators (NSA, % MoM)	Rosstat
r35	Unemployment rate, %	Rosstat
r36	Number of employed, % MoM	Rosstat
ro4	Index of industrial production	Rosstat
ro5	Industrial production: Mining and quarrying	Rosstat
ro6	Industrial production: manufacturing	Rosstat
ro7	Industrial production: supply of electric power, natural gas, and steam; air conditioning	Rosstat
ro8	Industrial production: supply of water; sewage; waste disposal and recycling	Rosstat
ro9	Metallurgical output	Rosstat
ro10	Output of finished metal items except machines and equipment	Rosstat
ro11	Output of paper products	Rosstat
ro12	Output of publishing products	Rosstat
ro13	Output of chemicals	Rosstat
ro14	Output of coal and petrochemicals	Rosstat
ro15	Output of electrical equipment	Rosstat
ro16	Output of food	Rosstat
ro17	Output of beverages	Rosstat
ro18	Output of tobacco products	Rosstat
ro19	Output of leather and leather products	Rosstat
ro20	Output of other mineral products	Rosstat
ro21	Output of textile products	Rosstat
ro22	Output of clothing	Rosstat
ro23	Output of rubber and plastic products	Rosstat
ro24	Output of means of transportation	Rosstat
ro25	Output of other transportation equipment	Rosstat
ro26	Output of timber and wood products	Rosstat
ro27	Output of other machines and equipment	Rosstat

s1	Index of business sentiment: Mining and quarrying	Rosstat
s2	Index of business sentiment: Manufacturing	Rosstat
s3	Index of business sentiment: Supply of electric power, natural gas, and water	Rosstat
s4	PMI: Composite – Output, SA	Markit
s5	PMI: Composite – New orders, SA	Markit
s6	PMI: Composite – Input prices, SA	Markit
s7	PMI: Composite – Output prices, SA	Markit
s8	PMI: Composite – Employments, SA	Markit
s9	PMI: Composite – Work backlog, SA	Markit
s10	PMI: Manufacturing, SA	Markit
s11	PMI: Manufacturing – Output, SA	Markit
s12	PMI: Manufacturing – New orders, SA	Markit
s13	PMI: Manufacturing – New export orders, SA	Markit
s14	PMI: Manufacturing – Finished goods, SA	Markit
s15	PMI: Manufacturing – Employment, SA	Markit
s16	PMI: Manufacturing – Stocks of purchase, SA	Markit
s17	PMI: Manufacturing – Quantity of purchase, SA	Markit
s18	PMI: Manufacturing – Input prices, SA	Markit
s19	PMI: Manufacturing – Output prices, SA	Markit
s20	PMI: Manufacturing – Delivery times, SA	Markit
s21	PMI: Manufacturing – Work backlogs, SA	Markit
s22	PMI: Services – Business activity, SA	Markit
s23	PMI: Services – New business, SA	Markit
s24	PMI: Services – Outstanding business, SA	Markit
s25	PMI: Services – Employment, SA	Markit
s26	PMI: Services – Prices charged, SA	Markit
s27	PMI: Services – Input prices, SA	Markit
s28	REB industry survey: Indebtness, anticipated	REB
s29	REB industry survey: Output prices	REB
s30	REB industry survey: Input prices	REB
s31	REB industry survey: Wages	REB
s32	REB industry survey: Employment	REB
s33	REB industry survey: Output	REB
s34	REB industry survey: Orders	REB
s35	REB industry survey: Stock of final output	REB

s36	REB industry survey: Output-to-input price ratio	REB
s37	REB industry survey: Purchases of new equipment	REB
s38	PMI: Services – Business expectations	Markit
s39	REB industry survey: Output prices, anticipated	REB
s40	REB industry survey: Input prices, anticipated	REB
s41	REB industry survey: Wages, anticipated	REB
s42	REB industry survey: Employment, anticipated	REB
s43	REB industry survey: Output, anticipated	REB
s44	REB industry survey: Purchases of equipment, anticipated	REB
s45	REB industry survey: Financial health, anticipated	REB
s46	REB industry survey: Stock of final output, anticipated	REB
s47	REB industry survey: Capacity utilization rate (100=normal)	REB
s48	REB industry survey: Labor utilization rate (100=normal)	REB
s49	REB industry survey: Stock of final output (100=normal)	REB
s50	REB industry survey: Orders (100=normal)	REB
s51	REB industry survey: Fraction of enterprises in good financial health, %	REB
w1	Imports	Rosstat
w2	Imports from CIS	Rosstat
w3	Imports from outside of CIS	Rosstat
w4	Ruble-denominated loans to households, % MoM	BoR
w5	Harmonized index of industrial production	BoR

Table 2: Pseudo-out-of-sample performance of DMA in comparison with benchmark forecasts as measured by RMSFE

Model	Horizon, months					
	1	2	3	4	5	6
DMA	3.6	4.7	5.1	5.4	5.4	5.4
DMS	3.8	4.8	5.1	5.7	5.9	6.0
BMA	3.5	5.1	5.6	5.6	5.7	5.7
BMS	3.6	5.2	5.7	5.6	5.8	5.8
UC-SV	5.6	6.4	6.7	6.9	6.9	6.8
AR(2)	3.8	5.3	5.7	5.8	5.8	5.8
TVP-AR(2)	4.0	5.3	5.6	5.6	5.6	5.8
DMA-AR(2)	4.0	5.6	5.4	5.3	5.3	5.3
TVP-KS	3.8	4.6	6.0	5.6	6.0	5.4

Notes: Entries are values of the root mean squared forecast error (RMSFE) obtained through the pseudo out-of-sample forecasting procedure and evaluated on the sample 2002m1 – 2017m9. 60 monthly observations are used for an initial forecast. DMA – Dynamic Model Averaging; DMS – Dynamic Model Sselection; BMA – Bayesian Model Averaging; BMS – Bayesian Model Selection; UC-SV – Unobserved Components-Stochastic Volatility Model; AR(2) – Bayesian autoregression of order 2; TVP-AR(2) – Bayesian autoregression of order 2 with time-varying parameters; DMA-AR(2) – DMA on contemporaneous inflation and its first lag; TVP-KS – Kitchen Sink model with time-varying parameters.

Table 3: Pseudo-out-of-sample performance of DMA in comparison with benchmark forecasts as measured by MAFE

Model	Horizon, months					
	1	2	3	4	5	6
DMA	2.7	3.5	3.5	3.9	4.0	3.8
DMS	2.8	3.5	3.7	4.0	4.3	4.2
BMA	2.7	3.6	3.8	3.8	3.9	3.8
BMS	2.8	3.7	4.0	3.8	4.0	3.9
UC-SV	3.6	4.0	4.1	4.2	4.3	4.5
AR(2)	2.7	3.8	3.9	4.0	4.0	4.1
TVP-AR(2)	2.6	3.7	3.8	3.7	3.8	4.0
DMA-AR(2)	2.6	3.8	3.8	3.7	3.7	3.7
TVP-KS	2.8	3.4	4.1	4.2	4.3	4.0

Notes: Entries are values of the mean absolute forecast error (MAFE) obtained through the pseudo out-of-sample forecasting procedure and evaluated on the sample 2002m1 – 2017m9. 60 monthly observations are used for an initial forecast. DMA – Dynamic Model Averaging; DMS – Dynamic Model Sselection; BMA – Bayesian Model Averaging; BMS – Bayesian Model Selection; UC-SV – Unobserved Components-Stochastic Volatility Model; AR(2) – Bayesian autoregression of order 2; TVP-AR(2) – Bayesian autoregression of order 2 with time-varying parameters; DMA-AR(2) – DMA on contemporaneous inflation and its first lag; TVP-KS – Kitchen Sink model with time-varying parameters.

Table 4: One-month-ahead DMA forecast

Predictor	$\mathbb{E}[\theta_t]$	Std.Dev.$[\theta_t]$	$\mathbb{E}[\mathbb{P}(\theta_t)]$	Std.Dev.$[\mathbb{P}(\theta_t)]$
Intercept	3.11	1.28	1.00	0.00
infl	0.47	0.09	1.00	0.00
infl(-1)	-0.03	0.04	0.19	0.13
s26	0.00	0.01	0.15	0.12
s46	0.00	0.00	0.12	0.09
s36	-0.01	0.00	0.25	0.19
s7	0.00	0.00	0.13	0.06
s39	0.00	0.00	0.13	0.07
s41	0.00	0.00	0.14	0.09
s27	0.00	0.00	0.13	0.05
f10	-3.77	1.45	0.81	0.20
f17	0.00	0.00	0.15	0.07
s45	0.00	0.00	0.08	0.06
f21	0.51	0.34	0.58	0.26
s6	0.00	0.00	0.13	0.10
f6	0.24	0.24	0.26	0.05
s43	0.00	0.00	0.08	0.03
s18	0.00	0.00	0.15	0.11
s44	0.00	0.00	0.07	0.05
s40	0.00	0.00	0.08	0.03
f18	0.00	0.01	0.17	0.09

Notes: The dependent variable is a monthly rate of CPI inflation one month ahead, infl(+1). Each individual model is bound to contain an intercept and contemporaneous inflation. $\mathbb{E}[\theta_t]$ is sample mean of the estimated coefficient of a regressor, and Std.Dev.$[\theta_t]$ is its standard deviation. $\mathbb{E}[\mathbb{P}(\theta_t)]$ is sample mean inclusion probability for a regressor, and Std.Dev.$[\mathbb{P}(\theta_t)]$ is its standard deviation. See Table 1 for a description of predictors.

Table 5: Two-month-ahead DMA forecast

Predictor	$\mathbb{E}[\theta_t]$	Std.Dev.$[\theta_t]$	$\mathbb{E}[\mathbb{P}(\theta_t)]$	Std.Dev.$[\mathbb{P}(\theta_t)]$
Intercept	0.62	3.02	1.00	0.00
infl	0.00	0.26	1.00	0.00
infl(-1)	0.01	0.01	0.09	0.04
s46	0.00	0.00	0.05	0.02
s41	0.01	0.01	0.42	0.27
s45	0.00	0.00	0.09	0.05
s43	0.00	0.00	0.11	0.08
s44	0.01	0.01	0.20	0.30
f15	0.00	0.02	0.14	0.05
s39	0.00	0.00	0.09	0.08
s40	0.00	0.00	0.05	0.01
s36	-0.03	0.04	0.36	0.40
s27	-0.01	0.01	0.22	0.15
f1	0.00	0.01	0.09	0.06
r21	-0.06	0.15	0.12	0.03
s3	0.00	0.00	0.09	0.04
f6	0.49	0.67	0.23	0.07
s6	0.01	0.02	0.16	0.17
f7	0.36	2.18	0.29	0.09
s26	0.00	0.01	0.15	0.08
s25	0.00	0.00	0.10	0.06

Notes: The dependent variable is a monthly rate of CPI inflation two months ahead, infl(+2). Each individual model is bound to contain an intercept and contemporaneous inflation. $\mathbb{E}[\theta_t]$ is sample mean of the estimated coefficient of a regressor, and Std.Dev.$[\theta_t]$ is its standard deviation. $\mathbb{E}[\mathbb{P}(\theta_t)]$ is sample mean inclusion probability for a regressor, and Std.Dev.$[\mathbb{P}(\theta_t)]$ is its standard deviation. See Table 1 for a description of predictors.

Table 6: Three-month-ahead DMA forecast

Predictor	$\mathbb{E}[\theta_t]$	Std.Dev.$[\theta_t]$	$\mathbb{E}[\mathbb{P}(\theta_t)]$	Std.Dev.$[\mathbb{P}(\theta_t)]$
Intercept	10.61	11.63	1.00	0.00
infl	0.10	0.12	1.00	0.00
infl(-1)	0.03	0.03	0.17	0.11
s40	0.00	0.00	0.05	0.02
s41	0.01	0.01	0.30	0.30
s44	0.00	0.00	0.06	0.04
s43	0.00	0.00	0.06	0.03
s39	0.00	0.00	0.15	0.15
s46	0.00	0.00	0.04	0.01
s3	0.00	0.01	0.07	0.03
f6	0.32	0.47	0.20	0.11
r24	-1.60	2.11	0.31	0.18
f15	-0.01	0.07	0.27	0.22
f1	0.00	0.00	0.06	0.02
f9	-72.32	129.08	0.50	0.29
s20	0.00	0.01	0.08	0.03
w4	65.57	119.76	0.52	0.28
f7	0.27	0.33	0.20	0.07
s25	0.00	0.00	0.07	0.04
f13	-2.56	2.65	0.38	0.18
s37	0.00	0.00	0.09	0.06

Notes: The dependent variable is a monthly rate of CPI inflation three months ahead, infl(+3). Each individual model is bound to contain an intercept and contemporaneous inflation. $\mathbb{E}[\theta_t]$ is sample mean of the estimated coefficient of a regressor, and Std.Dev.$[\theta_t]$ is its standard deviation. $\mathbb{E}[\mathbb{P}(\theta_t)]$ is sample mean inclusion probability for a regressor, and Std.Dev.$[\mathbb{P}(\theta_t)]$ is its standard deviation. See Table 1 for a description of predictors.

Table 7: Four-month-ahead DMA forecast

Predictor	$\mathbb{E}[\theta_t]$	Std.Dev.$[\theta_t]$	$\mathbb{E}[\mathbb{P}(\theta_t)]$	Std.Dev.$[\mathbb{P}(\theta_t)]$
Intercept	6.11	13.98	1.00	0.00
infl	0.20	0.13	1.00	0.00
infl(-1)	0.02	0.04	0.20	0.07
s20	0.00	0.01	0.08	0.06
s7	0.00	0.01	0.09	0.06
s44	0.00	0.00	0.08	0.04
s26	-0.01	0.01	0.14	0.09
s3	-0.01	0.01	0.09	0.05
s27	-0.01	0.01	0.14	0.10
f9	-8.62	20.96	0.34	0.11
s13	0.00	0.00	0.07	0.07
s41	0.02	0.02	0.50	0.32
s6	0.00	0.00	0.11	0.05
s46	0.00	0.00	0.10	0.06
s31	0.00	0.00	0.15	0.11
s35	0.00	0.00	0.06	0.04
w4	3.20	12.92	0.39	0.09
s43	0.00	0.00	0.11	0.07
f1	-0.01	0.01	0.10	0.10
f7	0.20	1.21	0.25	0.07
s40	0.00	0.00	0.10	0.07

Notes: The dependent variable is a monthly rate of CPI inflation four months ahead, infl(+4). Each individual model is bound to contain an intercept and contemporaneous inflation. $\mathbb{E}[\theta_t]$ is sample mean of the estimated coefficient of a regressor, and Std.Dev.$[\theta_t]$ is its standard deviation. $\mathbb{E}[\mathbb{P}(\theta_t)]$ is sample mean inclusion probability for a regressor, and Std.Dev.$[\mathbb{P}(\theta_t)]$ is its standard deviation. See Table 1 for a description of predictors.

Table 8: Five-month-ahead DMA forecast

Predictor	$\mathbb{E}[\theta_t]$	Std.Dev.$[\theta_t]$	$\mathbb{E}[\mathbb{P}(\theta_t)]$	Std.Dev.$[\mathbb{P}(\theta_t)]$
Intercept	0.41	6.37	1.00	0.00
infl	0.15	0.21	1.00	0.00
infl(-1)	0.00	0.01	0.06	0.02
s27	-0.02	0.03	0.18	0.18
s44	0.00	0.00	0.09	0.07
r29	0.50	0.83	0.19	0.07
r2	-0.06	0.09	0.10	0.05
s6	0.00	0.01	0.12	0.07
r24	-0.90	1.99	0.26	0.07
s37	0.00	0.00	0.05	0.04
f7	-0.74	3.53	0.20	0.07
f13	2.43	6.78	0.27	0.14
s13	-0.02	0.04	0.18	0.25
s7	0.00	0.00	0.07	0.03
s46	0.00	0.00	0.17	0.13
s25	0.00	0.00	0.06	0.02
s41	0.00	0.00	0.17	0.14
s3	-0.02	0.07	0.12	0.21
s8	0.00	0.01	0.08	0.06
s45	0.00	0.01	0.24	0.17
f6	0.13	1.26	0.30	0.11

Notes: The dependent variable is a monthly rate of CPI inflation five months ahead, infl(+5). Each individual model is bound to contain an intercept and contemporaneous inflation. $\mathbb{E}[\theta_t]$ is sample mean of the estimated coefficient of a regressor, and Std.Dev.$[\theta_t]$ is its standard deviation. $\mathbb{E}[\mathbb{P}(\theta_t)]$ is sample mean inclusion probability for a regressor, and Std.Dev.$[\mathbb{P}(\theta_t)]$ is its standard deviation. See Table 1 for a description of predictors.

Table 9: Six-month-ahead DMA forecast

Predictor	$\mathbb{E}[\theta_t]$	Std.Dev.$[\theta_t]$	$\mathbb{E}[\mathbb{P}(\theta_t)]$	Std.Dev.$[\mathbb{P}(\theta_t)]$
Intercept	0.07	9.33	1.00	0.00
infl	0.13	0.14	1.00	0.00
infl(-1)	0.00	0.03	0.07	0.03
r25	-5.29	14.27	0.35	0.07
r31	0.90	1.94	0.25	0.08
r24	7.33	15.14	0.36	0.08
r29	-0.44	1.50	0.20	0.06
s37	0.00	0.01	0.12	0.09
f13	-0.14	1.22	0.21	0.06
f24	0.02	0.08	0.10	0.02
s44	0.00	0.01	0.31	0.22
s23	0.00	0.00	0.05	0.02
f12	0.56	1.76	0.29	0.11
s25	0.00	0.01	0.13	0.08
s22	0.00	0.01	0.10	0.06
s5	0.00	0.00	0.05	0.03
s13	-0.03	0.06	0.25	0.34
s27	-0.04	0.07	0.30	0.37
f6	0.12	0.60	0.22	0.08
f7	1.20	2.00	0.28	0.07
f15	0.00	0.05	0.31	0.19

Notes: The dependent variable is a monthly rate of CPI inflation six months ahead, infl(+6). Each individual model is bound to contain an intercept and contemporaneous inflation. $\mathbb{E}[\theta_t]$ is sample mean of the estimated coefficient of a regressor, and Std.Dev.$[\theta_t]$ is its standard deviation. $\mathbb{E}[\mathbb{P}(\theta_t)]$ is sample mean inclusion probability for a regressor, and Std.Dev.$[\mathbb{P}(\theta_t)]$ is its standard deviation. See Table 1 for a description of predictors.

Table 10: Variance contribution

Model	Horizon, months					
	1	2	3	4	5	6
	DMA					
Observations	71.42	58.64	61.78	69.15	63.40	55.59
Coefficients	16.07	23.13	25.54	15.40	18.06	24.53
Model	10.67	16.02	10.57	11.02	16.07	17.29
Time-varying parameters	1.83	2.21	2.11	4.44	2.47	2.59
	BMA					
Observations	91.07	93.04	91.30	94.95	94.45	93.60
Coefficients	5.97	3.66	5.15	2.42	3.41	4.05
Model	2.96	3.30	3.55	2.63	2.13	2.35
Time-varying parameters	0.00	0.00	0.00	0.00	0.00	0.00

Notes: Entries are fraction of variances of the forecast error explained by the error term in the measurement equation for inflation, uncertainty due to imprecisely estimated coefficients, uncertainty about a model that generates data, and time variation of parameters. Out-of-sample forecasts produced by DMA and BMA were evaluated on the sample 2002m1 – 2017m9. 60 monthly observations are used for an initial forecast.

Figures

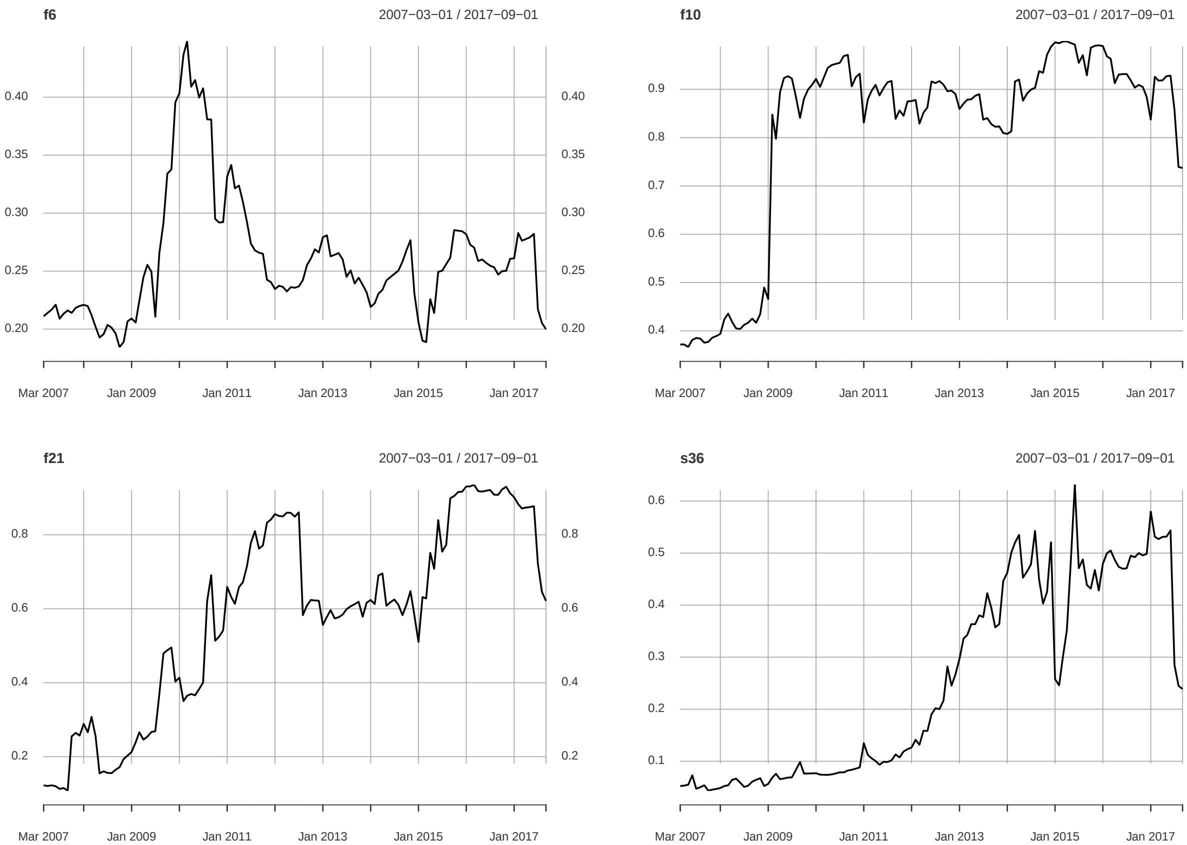

Figure 1: Inclusion probabilities of top one-month-ahead predictors: loans to non-financial enterprises up to 1 year (f6, upper left), nominal effective exchange rate (f10, upper right), international price of wheat (f21, lower left), and output-to-input price ratio in manufacturing (s36, lower right)

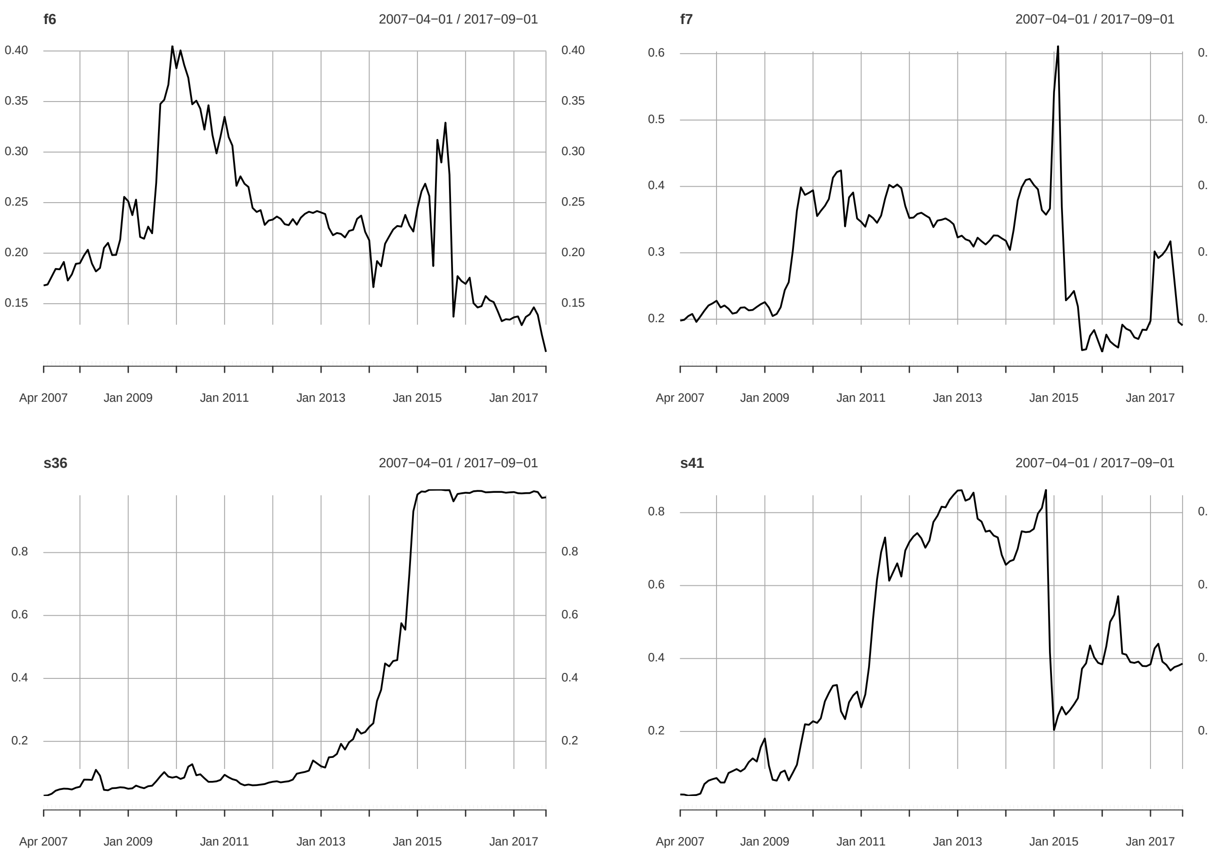

Figure 2: Inclusion probabilities of top two-month-ahead predictors: loans to non-financial enterprises up to 1 year (f6, upper left), loans to non-financial enterprises beyond 1 year (f7, upper right), output-to-input price ratio in manufacturing (s36, lower left), and expected wage inflation in manufacturing (s41, lower right)

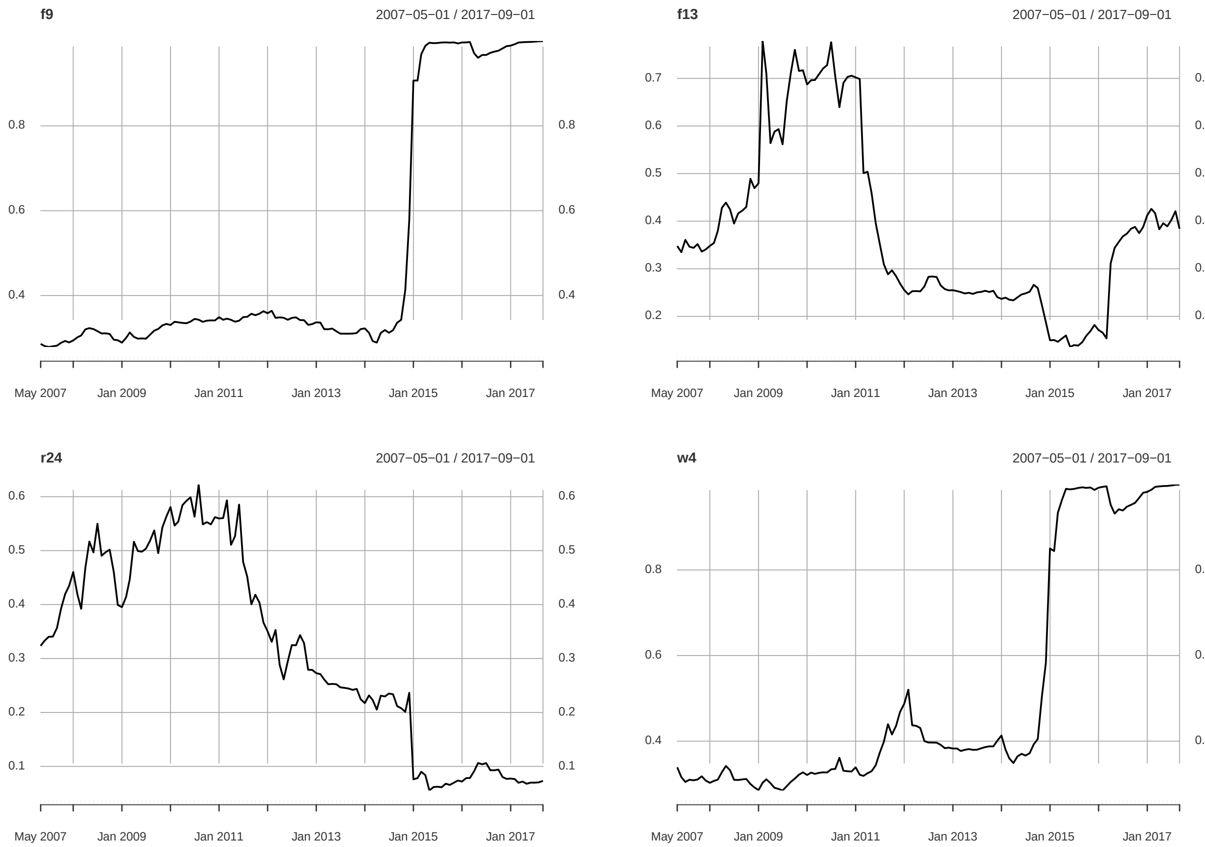

Figure 3: Inclusion probabilities of top three-month-ahead predictors: loans to individuals beyond 1 year (f9, upper left), monetary aggregate M0 (f13, upper right), nominal wage (r24, lower left), and loans to individuals (w4, lower right)

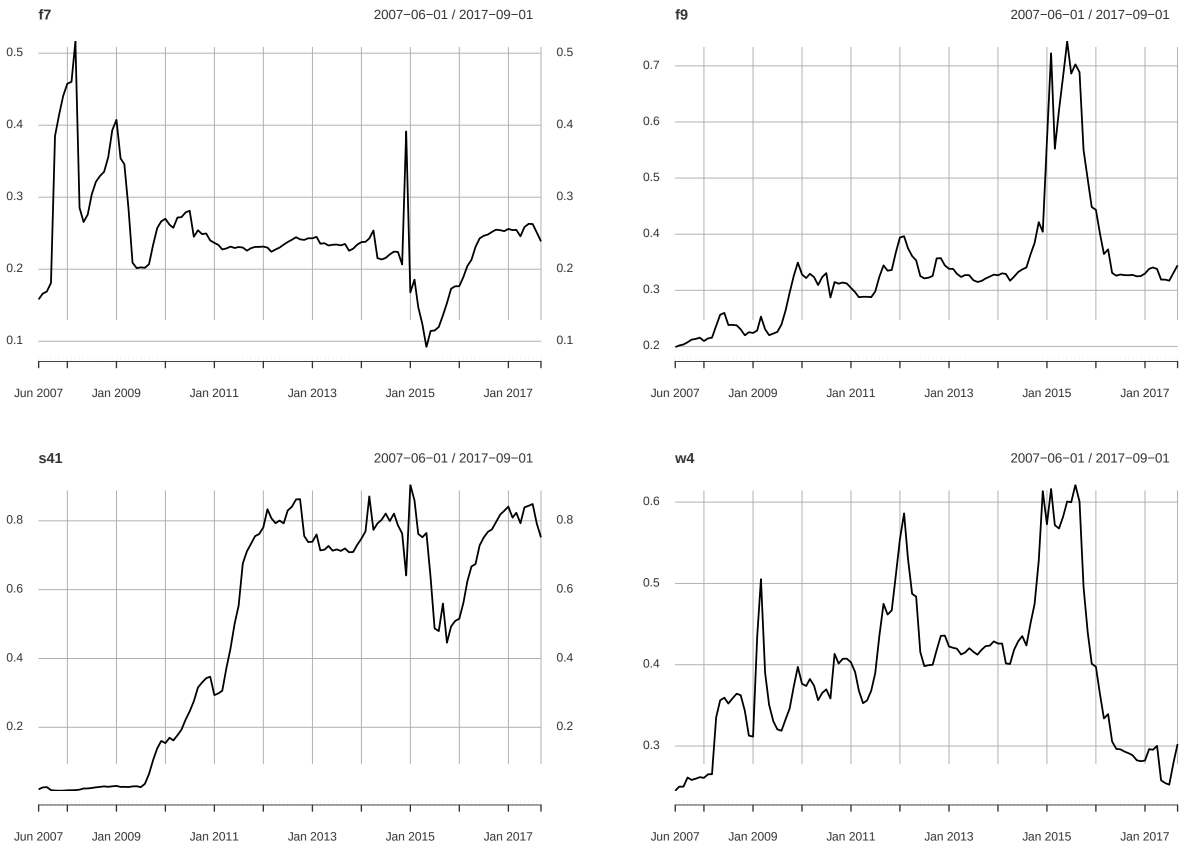

Figure 4: Inclusion probabilities of top four-month-ahead predictors: loans to non-financial enterprises beyond 1 year (f7, upper left), loans to individuals beyond 1 year (f9, upper right), expected wage inflation in manufacturing (s41, lower left), and loans to individuals (w4, lower right)

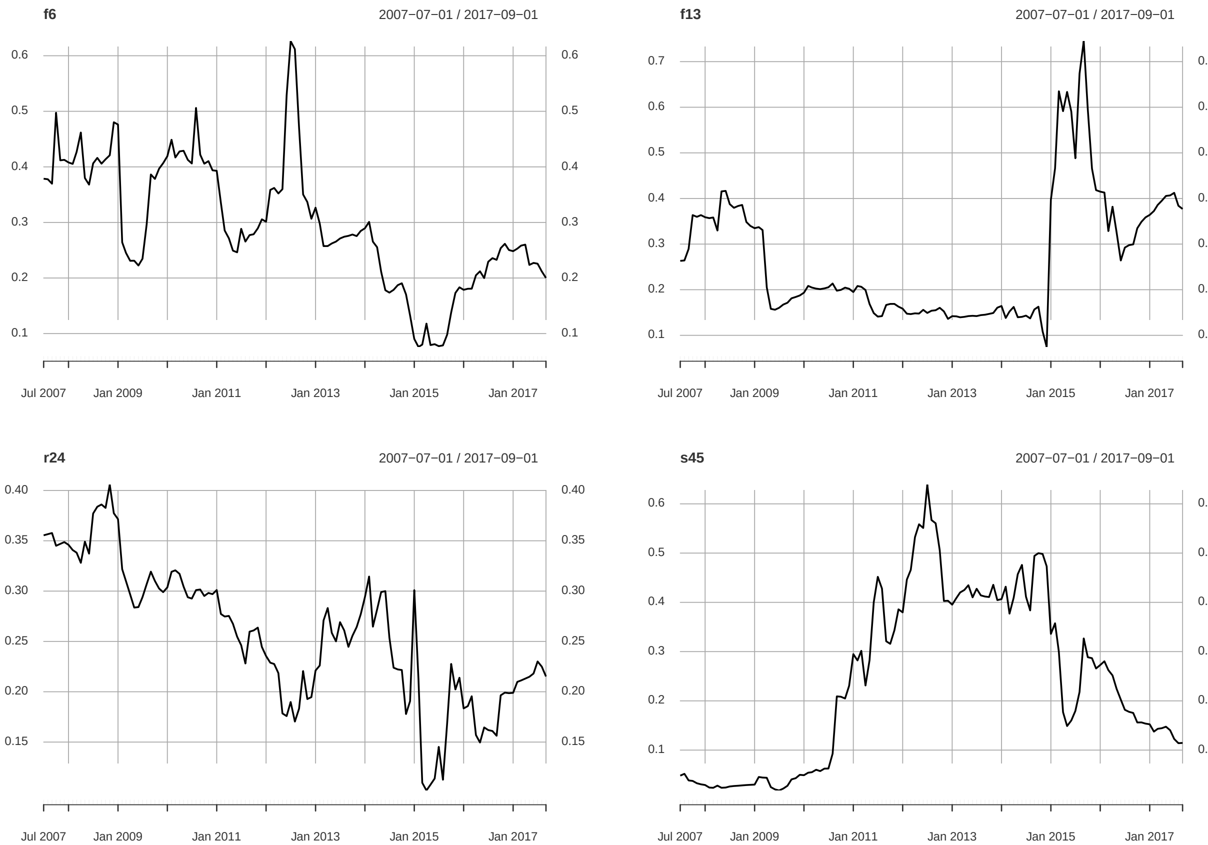

Figure 5: Inclusion probabilities of top five-month-ahead predictors: loans to non-financial enterprises up to 1 year (f6, upper left), monetary aggregate M0 (f13, upper right), nominal wage (r24, lower left), and financial health of firms in manufacturing (s45, lower right)

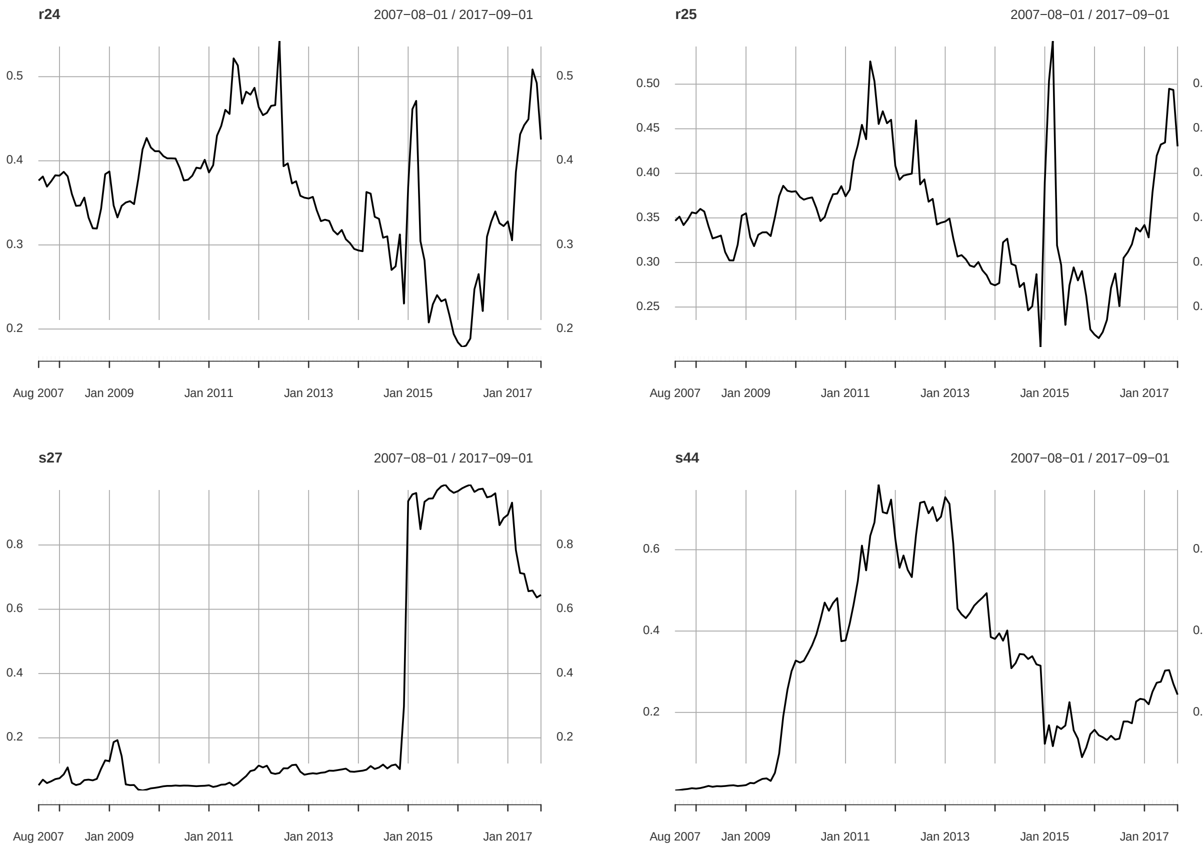

Figure 6: Inclusion probabilities of top three-month-ahead predictors: nominal wage (r24, upper left), real wage (r25, upper right), PMI input prices in services (s27, lower left), and expected purchases of equipment in manufacturing (s44, lower right)