



WP 34_09

Gary Koop

University of Strathclyde
and

The Rimini Centre for Economic Analysis, Italy

Dimitris Korobilis

University of Strathclyde
and

The Rimini Centre for Economic Analysis, Italy

FORECASTING INFLATION USING DYNAMIC MODEL AVERAGING

Copyright belongs to the author. Small sections of the text, not exceeding three paragraphs, can be used provided proper acknowledgement is given.

The *Rimini Centre for Economic Analysis* (RCEA) was established in March 2007. RCEA is a private, non-profit organization dedicated to independent research in Applied and Theoretical Economics and related fields. RCEA organizes seminars and workshops, sponsors a general interest journal *The Review of Economic Analysis*, and organizes a biennial conference: *Small Open Economies in the Globalized World* (SOEGW). Scientific work contributed by the RCEA Scholars is published in the RCEA Working Papers series.

The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Rimini Centre for Economic Analysis.

1 Introduction

Forecasting inflation is one of the more important, but difficult, exercises in macroeconomics. Many different approaches have been suggested. Perhaps the most popular are those based on extensions of the Phillips curve. This literature is too voluminous to survey here, but a few representative and influential papers include Ang, Bekaert and Wei (2007), Atkeson and Ohanian (2001), Groen, Paap and Ravazzolo (2008) and Stock and Watson (1999). The details of these papers differ, but the general framework involves a dependent variable such as inflation (or the change in inflation) and explanatory variables including lags of inflation, the unemployment rate and other predictors. Recursive, regression-based methods, have had some success. However, three issues arise when using such methods.

First, the coefficients on the predictors can change over time. For instance, it is commonly thought that the slope of the Phillips curve has changed over time. If so, the coefficients on the predictors that determine this slope will be changing. More broadly, there is a large literature in macroeconomics which documents structural breaks and other sorts of parameter change in many time series variables (see, among many others, Stock and Watson, 1996). Recursive methods are poorly designed to capture such parameter change. It is better to build models designed to capture it.

Second, the number of potential predictors can be large. For instance, Groen, Paap and Ravazzolo (2008) consider ten predictors. Researchers working with factor models such as Stock and Watson (1999) typically have many more than this. The existence of so many predictors can result in a huge number of models. For instance, if the set of models is defined by whether each of m potential predictors is included or excluded, then the researcher has 2^m models. This raises substantive statistical problems for model selection strategies. In light of this, many authors have turned to Bayesian methods, either to do Bayesian model averaging (BMA) or to automate the model selection process. Examples in macroeconomics and finance include Avramov (2002), Cremers (2002) and Koop and Potter (2004). Furthermore, computational demands can become daunting when the research is facing 2^m models.

Third, the model relevant for forecasting can potentially change over time. For instance, the set of predictors for inflation may have been different in the 1970s than now. Or some variables may predict well in recessions but not in expansions. This kind of issue further complicates an already difficult econometric exercise. That is, if the researcher has 2^m models and, at each point in time, a different forecasting model may apply, then the number of combinations of models which must be estimated in order to forecast at time τ

is $2^{m\tau}$. Even in relatively simple forecasting exercises, it can be computationally infeasible to forecast by simply going through all of these $2^{m\tau}$ combinations. For this reason, to our knowledge, there is no literature on forecasting inflation with many predictors where the coefficients on those predictors may change over time and where a different forecasting model might hold at each point in time. A purpose of this paper is to fill this gap.

In this paper, we consider a strategy developed by Raftery, Karny, Andrysek and Ettlter (2007) which they refer to as dynamic model averaging or DMA. Their approach can also be used for dynamic model selection or DMS where a single (potentially different) model can be used as the forecasting model at each point in time. DMA or DMS seem ideally suited for the problem of forecasting inflation since they allow for the forecasting model to change over time while, at the same time, allowing for coefficients in each model to evolve over time. They involve only standard econometric methods for state space models such as the Kalman filter but (via some empirically-sensible approximations) achieve vast gains in computational efficiency so as to allow DMA and DMS to be done in real time despite the computational problem described in the preceding paragraph.

We use these methods in the context of a forecasting exercise with quarterly US data from 1959Q1 through 2008Q2. We use two measures of inflation and fifteen predictors and compare the forecasting performance of DMA and DMS to a wide variety of alternative forecasting procedures. DMA and DMS indicate that the set of good predictors for inflation changes substantially over time. Due to this, we find DMA and DMS to forecast very well (in terms of forecasting metrics such as log predictive likelihoods, MSFEs and MAFEs), in most cases leading to large improvements in forecast performance relative to alternative approaches.

2 Forecasting Inflation

2.1 Generalized Phillips curve models

Many forecasting models of inflation are based on the Phillips curve in which current inflation depends only on the unemployment rate and lags of inflation and unemployment. Authors such as Stock and Watson (1999) include additional predictors leading to the so-called generalized Phillips curve. We take as a starting point, on which all models used in this paper build, the following generalized Phillips curve:

$$y_t = \phi + x'_{t-1}\beta + \sum_{j=1}^p \gamma_j y_{t-j} + \varepsilon_t \quad (1)$$

where y_t is inflation which we define as $\ln\left(\frac{P_t}{P_{t-1}}\right)$, with P_t being a price index, and x_t a vector of predictors. This equation is relevant for forecasting at time t given information through time $t-1$. When forecasting $h > 1$ periods ahead, the direct method of forecasting can be used and y_t and ε_t are replaced by y_{t+h-1} and ε_{t+h-1} in (1).

In this paper we use quarterly data. We provide results for inflation as measured by the GDP deflator and by the consumer price index (CPI). As predictors, authors such as Stock and Watson (1999) consider measures of real activity including the unemployment rate. Various other predictors (e.g. cost variables, the growth of the money supply, the slope of term structure, etc.) are suggested by economic theory. Finally, authors such as Ang, Bekaert and Wei (2007) have found surveys of experts on their inflation expectations to be useful predictors. These considerations suggest the following list of potential predictors which we use in this paper. Precise definitions and sources are given in the Data Appendix.

- UNEMP: unemployment rate.
- CONS: the percentage change in real personal consumption expenditures.
- INV: the percentage change in private residential fixed investment.
- GDP: the percentage change in real GDP.
- HSTARTS: the log of housing starts (total new privately owned housing units).
- EMPLOY: the percentage change in employment (All Employees: Total Private Industries, seasonally adjusted).
- PMI: the change in the Institute of Supply Management (Manufacturing): Purchasing Manager's Composite Index.
- WAGE: the percentage change in average hourly earnings in manufacturing.
- TBILL: three month Treasury bill (secondary market) rate.
- SPREAD: the spread between the 10 year and 3 month Treasury bill rates.
- DJIA: the percentage change in the Dow Jones Industrial Average.
- MONEY: the percentage change in the money supply (M1).
- INFEXP: University of Michigan measure of inflation expectations.

- COMPRICE: the change in the commodities price index (NAPM commodities price index).
- VENDOR: the change in the NAPM vendor deliveries index.

This set of variables is a wide one reflecting the major theoretical explanations of inflation as well as variables which have found to be useful in forecasting inflation in other studies.

2.2 Time Varying Parameter Models

Research in empirical macroeconomics often uses time varying parameter (TVP) models which are estimated using state space methods such as the Kalman filter. A standard specification can be written, for $t = 1, \dots, T$, as

$$y_t = z_t \theta_t + \varepsilon_t \tag{2a}$$

$$\theta_t = \theta_{t-1} + \eta_t. \tag{2b}$$

In our case, y_t is inflation, $z_t = [1, x_{t-1}, y_{t-1}, \dots, y_{t-p}]$ is an $1 \times m$ vector of predictors for inflation (including an intercept and lags of inflation), $\theta_t = [\phi_{t-1}, \beta_{t-1}, \gamma_{t-1}, \dots, \gamma_{t-p}]$ is an $m \times 1$ vector of coefficients (states), $\varepsilon_t \stackrel{ind}{\sim} N(0, H_t)$ and $\eta_t \stackrel{ind}{\sim} N(0, Q_t)$. The errors, ε_t and η_t , are assumed to be mutually independent at all leads and lags. Examples of recent papers which use such models (or extensions thereof) in macroeconomics include Cogley and Sargent (2005), Cogley, Morozov and Sargent (2005), Groen, Paap and Ravazzolo (2008), Koop, Leon-Gonzalez and Strachan (2009), Korobilis (2009) and Primiceri (2005).

The model given by (2a) and (2b) is an attractive one that allows for empirical insights which are not available with traditional, constant coefficient models (even when the latter are estimated recursively). However, when forecasting, they have the potential drawback that the same set of explanatory variables is assumed to be relevant at all points in time. Furthermore, if the number of explanatory variables in z_t is large, such models can often over-fit in-sample and, thus, forecast poorly.

Popular extensions of (2a) and (2b) such as TVP-VARs also include the same set of explanatory variables at all times and suffer from the same problems. Even innovative extensions such as that of Groen, Paap and Ravazzolo (2008) involve only a partial treatment of predictor uncertainty. In an inflation forecasting exercise, they use a model which modifies the measurement equation to be:

$$y_t = \sum_{j=1}^m s_j \theta_{jt} z_{jt} + \varepsilon_t,$$

where θ_{jt} and z_{jt} denote the j^{th} elements of θ_t and z_t . The key addition to their model is $s_j \in \{0, 1\}$. Details of the exact model used for s_j are provided in Groen et al (2008). For present purposes, the important thing to note is that it allows for each predictor for inflation to either be included (if $s_j = 1$) or excluded (if $s_j = 0$), but that s_j does not vary over time. That is, this model either includes a predictor at all points in time or excludes it at all points in time. It does not allow for the set of predictors to vary over time. It is the treatment of this latter issue which is the key addition provided by DMA.

2.3 Dynamic Model Averaging

To define what we do this paper, suppose that we have a set of K models which are characterized by having different subsets of z_t as predictors. Denoting these by $z^{(k)}$ for $k = 1, \dots, K$, our set of models can be written as:

$$\begin{aligned} y_t &= z_t^{(k)} \theta_t^{(k)} + \varepsilon_t^{(k)} \\ \theta_{t+1}^{(k)} &= \theta_t^{(k)} + \eta_t^{(k)}, \end{aligned} \tag{3}$$

$\varepsilon_t^{(k)}$ is $N(0, H_t^{(k)})$ and $\eta_t^{(k)}$ is $N(0, Q_t^{(k)})$. Let $L_t \in \{1, 2, \dots, K\}$ denote which model applies at each time period, $\Theta_t = (\theta_t^{(1)'}, \dots, \theta_t^{(K)'})'$ and $y^t = (y_1, \dots, y_t)'$. The fact that we are letting different models hold at each point in time and will do model averaging justifies the terminology “dynamic model averaging”. To be precise, when forecasting time t variables using information through time $t-1$, DMA involves calculating $\Pr(L_t = k | y^{t-1})$ for $k = 1, \dots, K$ and averaging forecasts across models using these probabilities. DMS involves selecting the single model with the highest value for $\Pr(L_t = k | y^{t-1})$ and using this to forecast. Details on the calculation of $\Pr(L_t = k | y^{t-1})$ will be provided below.

Specifications such as (3) are potentially of great interest in empirical macroeconomics since they allow for the set of predictors for inflation to change over time as well as allowing the marginal effects of the predictors to change over time. The problems with such a framework are that many of the models can have a large number of parameters (and, hence, risk being over-parameterized) and the computational burden which arises when K is large implies that estimation can take a long time (a potentially serious drawback when forecasting in real time).

To understand the source and nature of these problems, consider how the researcher might complete the model given in (3). Some specification for how predictors enter/leave the model in real time is required. A simple way of doing this would be through a transition matrix, P , with elements $p_{ij} = \Pr(L_t = i | L_{t-1} = j)$ for $i, j = 1, \dots, K$. Bayesian inference in such a model is theoretically straightforward, but will be computationally infeasible since P will typically be an enormous matrix. Consider the case where we have m potential predictors and our models are defined according to whether each is included or excluded. Then we have $K = 2^m$ and P is a $K \times K$ matrix. Unless m is very small, P will have so many parameters that inference will be very imprecise and computation very slow.¹ Thus, a full Bayesian approach to DMA can be quite difficult. In this paper, we use approximations suggested by Raftery, Karny, Andrysek and Ettler (2007) in an industrial application. These approximations have the huge advantage that standard state space methods (e.g. involving the Kalman filter) can be used, allowing for fast real time forecasting.

The approximations used by Raftery et al (2007) involve two parameters, λ and α , which they refer to as *forgetting factors* and fix to numbers slightly below one. To explain the role of these forgetting factors, first consider the standard state space model in (2a) and (2b). For given values of H_t and Q_t , standard filtering and smoothing results can be used to carry out recursive estimation or forecasting. That is, Kalman filtering begins with the result that

$$\theta_{t-1} | y^{t-1} \sim N \left(\hat{\theta}_{t-1}, \Sigma_{t-1|t-1} \right) \quad (4)$$

where formulae for $\hat{\theta}_{t-1}$ and $\Sigma_{t-1|t-1}$ are standard (and are provided below for the case considered in this paper). Note here only that these formulae depend on H_t and Q_t . Then Kalman filtering proceeds using:

$$\theta_t | y^{t-1} \sim N \left(\hat{\theta}_{t-1}, \Sigma_{t|t-1} \right), \quad (5)$$

where

$$\Sigma_{t|t-1} = \Sigma_{t-1|t-1} + Q_t.$$

Raftery et al (2007) note that things simplify substantially if this latter equation is replaced by:

¹See, for instance, Chen and Liu (2000) who discuss related models and how computation time up to t typically involves mixing over K^t terms.

$$\Sigma_{t|t-1} = \frac{1}{\lambda} \Sigma_{t-1|t-1} \quad (6)$$

or, equivalently, $Q_t = (1 - \lambda^{-1}) \Sigma_{t-1|t-1}$ where $0 < \lambda \leq 1$. Such approaches have long been used in the state space literature going back to Fagin (1964) and Jazwinsky (1970). Raftery et al (2007) provide a detailed justification of this approximation and relate the resulting approach to statistical methods such as age-weighting and windowing and the reader is referred to their paper for details. The name “forgetting factor” is suggested by the fact that this specification implies that observations j periods in the past have weight λ^j . An alternative way of interpreting λ is to note that it implies an effective window size of $\frac{1}{1-\lambda}$. It is common to choose a value of λ near one, suggesting a gradual evolution of coefficients. Raftery et al (2007) set $\lambda = 0.99$. For quarterly macroeconomic data, this suggests observations five years ago receive approximately 80% as much weight as last period’s observation. This is the sort of value consistent with fairly stable models where coefficient change is gradual. With $\lambda = 0.95$, observations five years ago receive only about 35% as much weight as last period’s observations. This suggests substantial parameter instability where coefficient change is quite rapid. This seems to exhaust the range of reasonable values for λ and, accordingly, in our empirical work we consider $\lambda \in (0.95, 0.99)$. $\lambda = 0.99$ will be our benchmark choice and most of our empirical results will be reported for this (although we also include an analysis of the sensitivity to this choice).

An important point to note is that, with this simplification, we no longer have to estimate or simulate Q_t . Instead, all that is required (in addition to the Kalman filter) is a method for estimating or simulating H_t (something which we will discuss below).

Forecasting in the one model case is then completed by the updating equation:

$$\theta_t|y^t \sim N\left(\hat{\theta}_t, \Sigma_{t|t}\right), \quad (7)$$

where

$$\hat{\theta}_t = \hat{\theta}_{t-1} + \Sigma_{t|t-1} z_t \left(H_t + z_t \Sigma_{t|t-1} z_t' \right)^{-1} \left(y_t - z_t \hat{\theta}_{t-1} \right) \quad (8)$$

and

$$\Sigma_{t|t} = \Sigma_{t|t-1} - \Sigma_{t|t-1} z_t \left(H_t + z_t \Sigma_{t|t-1} z_t' \right)^{-1} z_t \Sigma_{t|t-1}. \quad (9)$$

Recursive forecasting is done using the predictive distribution

$$y_t|y^{t-1} \sim N\left(z_t\widehat{\theta}_{t-1}, H_t + z_t\Sigma_{t|t-1}z_t'\right). \quad (10)$$

We stress that, conditional on H_t , these results are all analytical and, thus, no Markov chain Monte Carlo (MCMC) algorithm is required. This greatly reduces the computational burden.

The case with many models, (3), uses the previous approximation and an additional one. To explain this, we now switch to the notation for the multiple model case in (3) and let Θ_t denote the vector of all the coefficients. In the standard single model case, Kalman filtering is based on (4), (5) and (7). In the multi-model case, for model k , these three equations become:

$$\Theta_{t-1}|L_{t-1} = k, y^{t-1} \sim N\left(\widehat{\theta}_{t-1}^{(k)}, \Sigma_{t-1|t-1}^{(k)}\right), \quad (11)$$

$$\Theta_t|L_t = k, y^{t-1} \sim N\left(\widehat{\theta}_{t-1}^{(k)}, \Sigma_{t|t-1}^{(k)}\right) \quad (12)$$

and

$$\Theta_t|L_t = k, y^t \sim N\left(\widehat{\theta}_t^{(k)}, \Sigma_{t|t}^{(k)}\right), \quad (13)$$

where $\widehat{\theta}_t^{(k)}$, $\Sigma_{t|t}^{(k)}$ and $\Sigma_{t|t-1}^{(k)}$ are obtained via Kalman filtering in the usual way using (8), (9) and (6), except with (k) superscripts added to denote model k . To make clear the notation in these equations, note that, conditional on $L_t = k$, the prediction and updating equations will only provide information on $\theta_t^{(k)}$ and not the full vector Θ_t . Hence, we have only written (11), (12) and (13) in terms of the distributions which hold for $\theta_t^{(k)}$.

The previous results were all conditional on $L_t = k$, and we need a method for unconditional prediction (i.e. not conditional on a particular model). In theory, a nice way of doing this would be through specifying a transition matrix, P , such as that given above and using MCMC methods to obtain such unconditional results. However, for the reasons discussed previously, this will typically be computationally infeasible and empirically undesirable due to the resulting proliferation of parameters. In this paper, we follow the suggestion of Raftery et al (2007) involving a forgetting factor for the state equation for the models, α , comparable to the forgetting factor λ used with the state equation for the parameters. The derivation of Kalman filtering ideas begins with (4). The analogous result, when doing DMA, is

$$p(\Theta_{t-1}, L_{t-1} | y^{t-1}) = \sum_{k=1}^K p(\theta_{t-1}^{(k)} | L_{t-1} = k, y^{t-1}) \Pr(L_{t-1} = k | y^{t-1}), \quad (14)$$

where $p(\theta_{t-1}^{(k)} | L_{t-1} = k, y^{t-1})$ is given by (11). To simplify notation, let $\pi_{t|s,l} = \Pr(L_t = l | y^s)$ and thus, the final term on the right hand side of (14) is $\pi_{t-1|t-1,k}$.

If we were to use the unrestricted matrix of transition probabilities in P with elements p_{kl} then the model prediction equation would be:

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

but Raftery et al (2007) replace this by:

$$\pi_{t|t-1,k} = \frac{\pi_{t-1|t-1,k}^\alpha}{\sum_{l=1}^K \pi_{t-1|t-1,l}^\alpha}, \quad (15)$$

where $0 < \alpha \leq 1$ is set to a fixed value slightly less than one and is interpreted in a similar manner to λ . Raftery et al (2007) argue that this is an empirically sensible simplification and, in particular, is a type of multiparameter power steady model used elsewhere in the literature. See also Smith and Miller (1986) who work with a similar model and argue approximations such as (15) are sensible and not too restrictive.

The huge advantage of using the forgetting factor α in the model prediction equation is that we do not require an MCMC algorithm to draw transitions between models nor a simulation algorithm over model space.² Instead, simple evaluations comparable to those of the updating equation in the Kalman filter can be done. In particular, we have a model updating equation of:

$$\pi_{t|t,k} = \frac{\pi_{t|t-1,k} p_k(y_t | y^{t-1})}{\sum_{l=1}^K \pi_{t|t-1,l} p_l(y_t | y^{t-1})}, \quad (16)$$

where $p_l(y_t | y^{t-1})$ is the predictive density for model l (i.e. the Normal density in (10) with (l) superscripts added) evaluated at y_t .

Recursive forecasting can be done by averaging over predictive results for every model using $\pi_{t|t-1,k}$. So, for instance, DMA point predictions are given by:

²Examples of simulation algorithms over model space include the Markov chain Monte Carlo model composition (MC³) algorithm of Madigan and York (1995) or the reversible jump MCMC algorithm of Green (1995).

$$E(y_t | y^{t-1}) = \sum_{k=1}^K \pi_{t|t-1,k} z_t^{(k)} \widehat{\theta}_{t-1}^{(k)}.$$

DMS proceeds by selecting the single model with the highest value for $\pi_{t|t-1,k}$ at each point in time and simply using it for forecasting.

To understand further how the forgetting factor α can be interpreted, note that this specification implies that the weight used in DMA which is attached to model k at time t is:

$$\begin{aligned} \pi_{t|t-1,k} &\propto [\pi_{t-1|t-2,k} p_k(y_{t-1} | y^{t-2})]^\alpha \\ &= \prod_{i=1}^{t-1} [p_k(y_{t-i} | y^{t-i-1})]^\alpha. \end{aligned}$$

Thus, model k will receive more weight at time t if it has forecast well in the recent past (where forecast performance is measured by the predictive density, $p_k(y_{t-i} | y^{t-i-1})$). The interpretation of “recent past” is controlled by the forgetting factor, α and we have the same exponential decay at the rate α^i for observations i periods ago as we had associated with λ . Thus, if $\alpha = 0.99$ (our benchmark value and also the value used by Raftery et al, 2007), forecast performance five years ago receives 80% as much weight as forecast performance last period (when using quarterly data). If $\alpha = 0.95$, then forecast performance five years ago receives only about 35% as much weight. These considerations suggest that, as with λ , we focus on the interval $\alpha \in (0.95, 0.99)$.

Note also that, if $\alpha = 1$, then $\pi_{t|t-1,k}$ is simply proportional to the marginal likelihood using data through time $t - 1$. This is what standard approaches to BMA would use. If we further set $\lambda = 1$, then we obtain BMA using conventional linear forecasting models with no time variation in coefficients. In our empirical work, we include BMA in our set of alternative forecasting procedures and implement this by setting $\alpha = \lambda = 1$.

We stress that, conditional on H_t , the estimation and forecasting strategy outlined above only involves evaluating formulae such as those in the Kalman filter. All the recursions above are started by choosing a prior for $\pi_{0|0,k}$ and $\theta_0^{(k)}$ for $k = 1, \dots, K$.

The preceding discussion is all conditional on H_t . Raftery et al (2007) recommend a simple plug in method where $H_t^{(k)} = H^{(k)}$ and is replaced with a consistent estimate. When forecasting inflation, however, it is likely that the error variance is changing over time. In theory, we could use a stochastic volatility or ARCH specification for $H_t^{(k)}$. However, to do this would greatly add to the computational burden. Thus, we prefer a

simple plug-in approach which is a rolling version of the recursive method of Raftery et al (2007). To be precise, let

$$\tilde{H}_t^{(k)} = \frac{1}{t^*} \sum_{j=t-t^*+1}^t \left[\left(y_t - z_t^{(k)} \hat{\theta}_{t-1}^{(k)} \right)^2 - z_t^{(k)} \Sigma_{t|t-1}^{(k)} z_t^{(k)'} \right].$$

Raftery et al (2007) uses this with $t^* = t$, but to allow for more substantial change in the error variances (e.g. due to the Great Moderation of the business cycle), we set $t^* = 20$ and, thus, use a rolling estimator based on five years of data. Following Raftery et al (2007), we can avoid the rare possibility that $\tilde{H}_t^{(k)} < 0$, by replacing $H_t^{(k)}$ by $\hat{H}_t^{(k)}$ where:

$$\hat{H}_t^{(k)} = \begin{cases} \tilde{H}_t^{(k)} & \text{if } \tilde{H}_t^{(k)} > 0 \\ \hat{H}_{t-1}^{(k)} & \text{otherwise} \end{cases}.$$

3 Empirical Work

Our empirical work is divided into three sub-sections. The first two of these sub-sections present results using DMA and DMS, implemented in our preferred way. This involves setting $\alpha = 0.99$, $\lambda = 0.99$, a noninformative prior over the models (i.e. $\pi_{0|0,k} = \frac{1}{K}$ for $k = 1, \dots, K$ so that, initially, all models are equally likely) and a relatively diffuse prior on the initial conditions of the states: $\theta_0^{(k)} \sim N(0, 100)$ for $k = 1, \dots, K$. The first sub-section presents evidence on which variables are good for predicting inflation. The second sub-section investigates forecast performance by comparing DMA forecasts to those produced by several alternative forecasting strategies. The third sub-section presents evidence on the sensitivity of our results to the choice of the forgetting factors. We present results for short-term ($h = 1$), medium-term ($h = 4$) and long-term ($h = 8$) forecast horizons for two measures of inflation: one based on the CPI, the other based on the GDP deflator. The list of potential predictors (which specifies the transformation used on each variable) is given in sub-section 2.1 (see also the Data Appendix). All of our models include an intercept two lags of the dependent variable.³

3.1 Which Variables are Good Predictors for Inflation?

In theory, DMA has a large potential benefit over other forecasting approaches in that it allows the forecasting model to change over time. Of course, in a particular empirical application, this benefit may be small if the forecasting model does not change much over

³Preliminary experimentation with lag lengths up to four indicated two lags leads to the best forecast performance for both our measures of inflation.

time. Accordingly, we begin by presenting evidence that, when forecasting inflation, the forecasting model is changing over time.

One striking feature of all of our empirical results is that, although we have 15 potential predictors (and, thus, tens of thousands of models), most probability is attached to very parsimonious models with only a few predictors. If we let $Size_k$ be the number of predictors in model k (note that $Size_k$ does not include the intercept plus two lags of the dependent variable which are common to all models), then

$$E(Size_t) = \sum_{k=1}^K \pi_{t|t-1,k} Size_k$$

can be interpreted as the expected or average number of predictors used in DMA at time t . Figure 1 plots this for our six empirical exercises (i.e. two definitions of inflation and three forecast horizons).

For the short forecast horizon ($h = 1$), the shrinkage of DMA is particularly striking. It consistently includes (in an expected value sense) a single predictor for both our definitions of inflation. For GDP deflator inflation at horizons $h = 4$ and $h = 8$, slightly more predictors are included (i.e. roughly 2 predictors are included in the early 1970s, but the number of predictors increases to 3 or 4 by the end of the sample). It is only for CPI based inflation at longer horizons that DMA chooses larger numbers of predictors. For instance, for $h = 8$ the expected number of predictors gradually increases from about two in 1970 to about eight by 2000. But even this least parsimonious case (which is still very parsimonious before 1990) excludes (in an expected value sense) half of the potential predictors.

Figure 1 shows clear evidence that DMA will shrink forecasts and provides some evidence that the way this shrinkage is done changes over time. But it does not tell us which predictors are important and how the predictors are changing over time. It is to these issues we now turn.

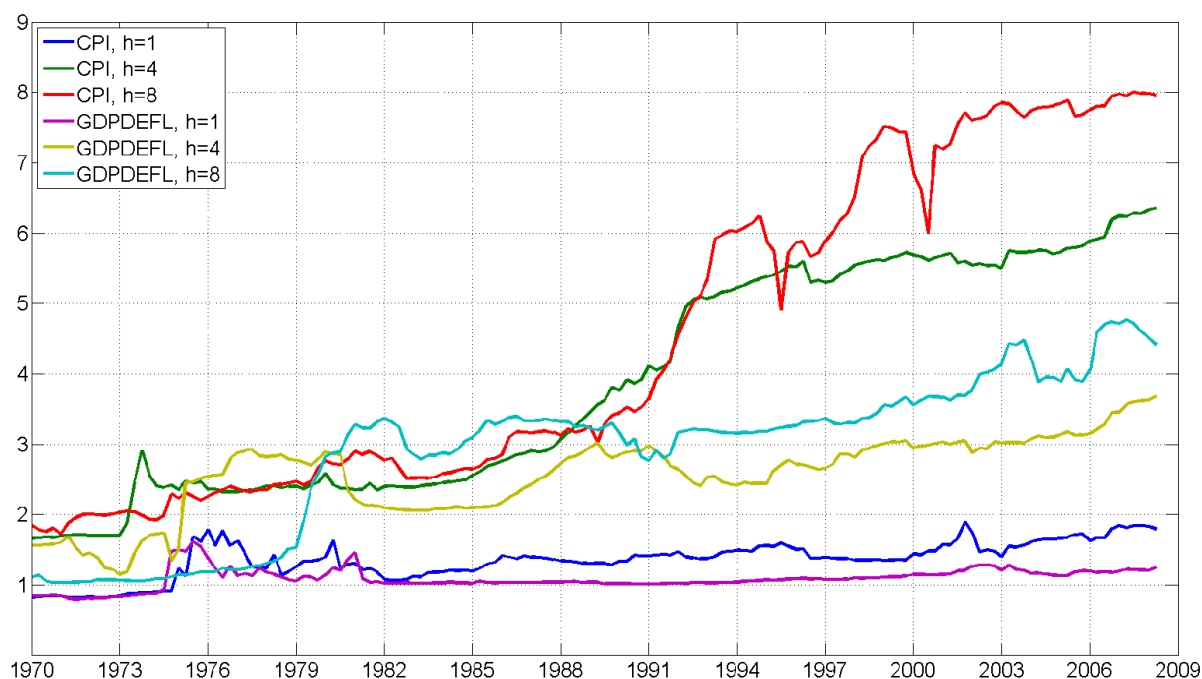


Figure 1: Expected Number of Predictors in Each Forecasting Exercise

Figures 2 through 7 shed light on which predictors are important at each point in time for each of our six empirical exercises. These graphs contain posterior inclusion probabilities. That is, they are the probability that a predictor is useful for forecasting at time t . Equivalently, they are the weight used by DMA attached to models which include a predictor. To keep the figures readable, we only present posterior inclusion probabilities for predictors which are important at least one point in time. To be precise, any predictor where the inclusion probability is never above 0.5 is excluded from the appropriate figure.

These figures confirm that DMS is almost always choosing parsimonious models and the weights in DMA heavily reflect parsimonious models. That is, with the partial exception of $h = 8$, it is rare for DMS to choose a model with more than two or three predictors.

Another important result is that for both measures of inflation and for all forecast horizons, we are finding strong evidence of model change. That is, the set of predictors in the forecasting model is changing over time.

Results for CPI inflation for $h = 1$ are particularly striking. Before 1975, no predictors come through strongly. Between 1975 and 1985 money is the only predictor. After 1985 the measure of inflation expectations comes through strongly. With regards to the inflation expectations variable, similar patterns are observed for $h = 4$ and $h = 8$. Before the mid- to late- 1980s there is little or no evidence that it is a useful predictor for inflation. But after this, it often is a useful predictor. To a lesser extent, the same pattern holds

with GDP deflator inflation. For instance, with $h = 1$ very few predictors are included, with money being an important predictor near the beginning of the sample and inflation expectations being important near the end. However, for GDP deflator inflation with $h = 1$, the predictor reflecting earnings (WAGE) comes through as being the strongest predictor after 1980 (this variable was not found to be an important predictor for CPI inflation).

Housing starts is another variable which often has strong predictive power for both measures of inflation. But, interestingly, only at medium or long horizons. For $h = 1$, there is no evidence at all that housing starts have predictive power for inflation.

The interested reader can examine Figures 2 through 7 for any particular variable of interest. Most of our potential explanatory variables come through as being important at some time, for some forecast horizon for some measure of inflation. Only CONS, DJIA, COMPRICE and PMI never appear in Figures 2 through 7. But it is clearly the case that there is a large variation over time, over forecast horizons and over measures of inflation in what is a good predictor for inflation. We stress that the great benefit of DMA and DMS is that they will pick up good predictors automatically as the forecasting model evolves over time.




Figure 2: Posterior Probability of Inclusion of Main Predictors (CPI inflation, $h = 1$)




Figure 3: Posterior Probability of Inclusion of Main Predictors (CPI inflation, $h = 4$)




Figure 4: Posterior Probability of Inclusion of Main Predictors (CPI inflation, $h = 8$)

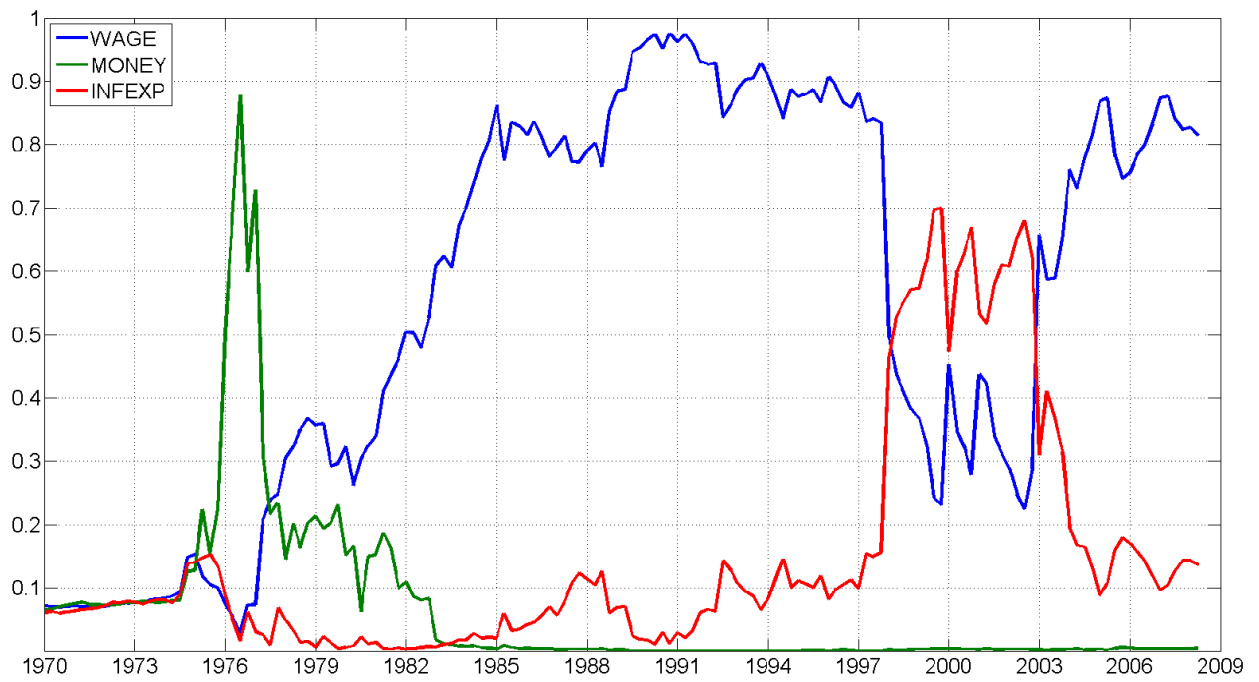


Figure 5: Posterior Probability of Inclusion of Main Predictors (GDP deflator inflation, $h = 1$)

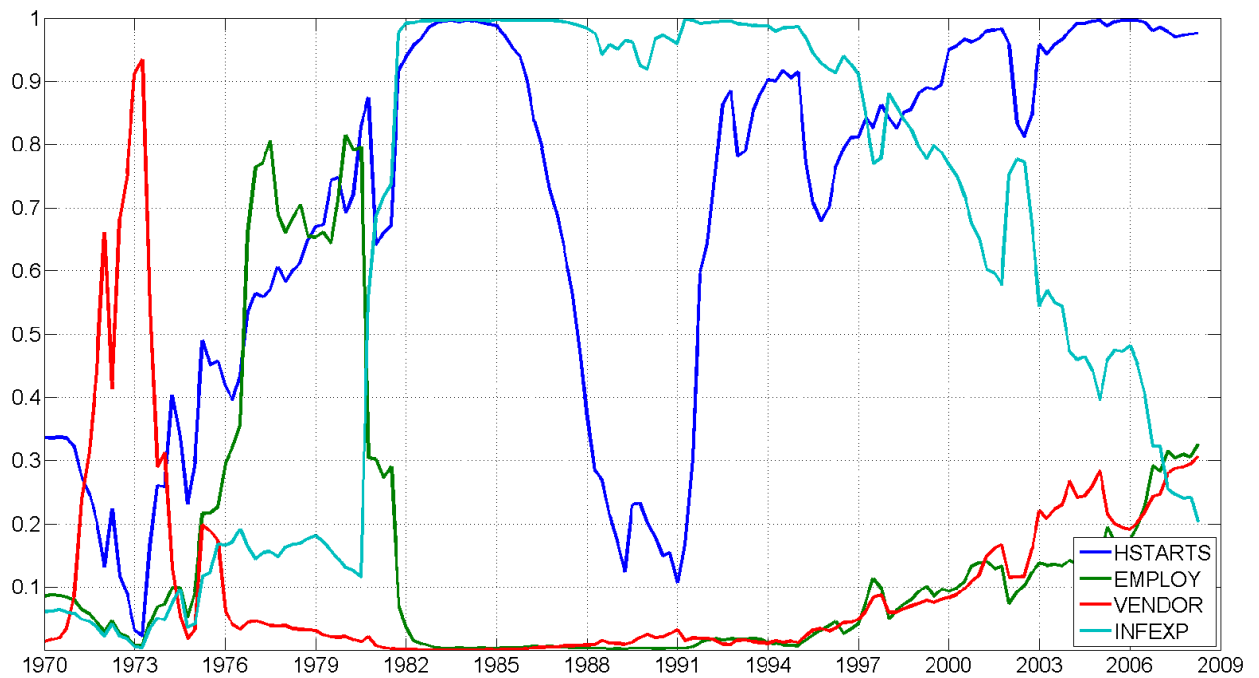


Figure 6: Posterior Probability of Inclusion of Main Predictors (GDP deflator inflation, $h = 4$)

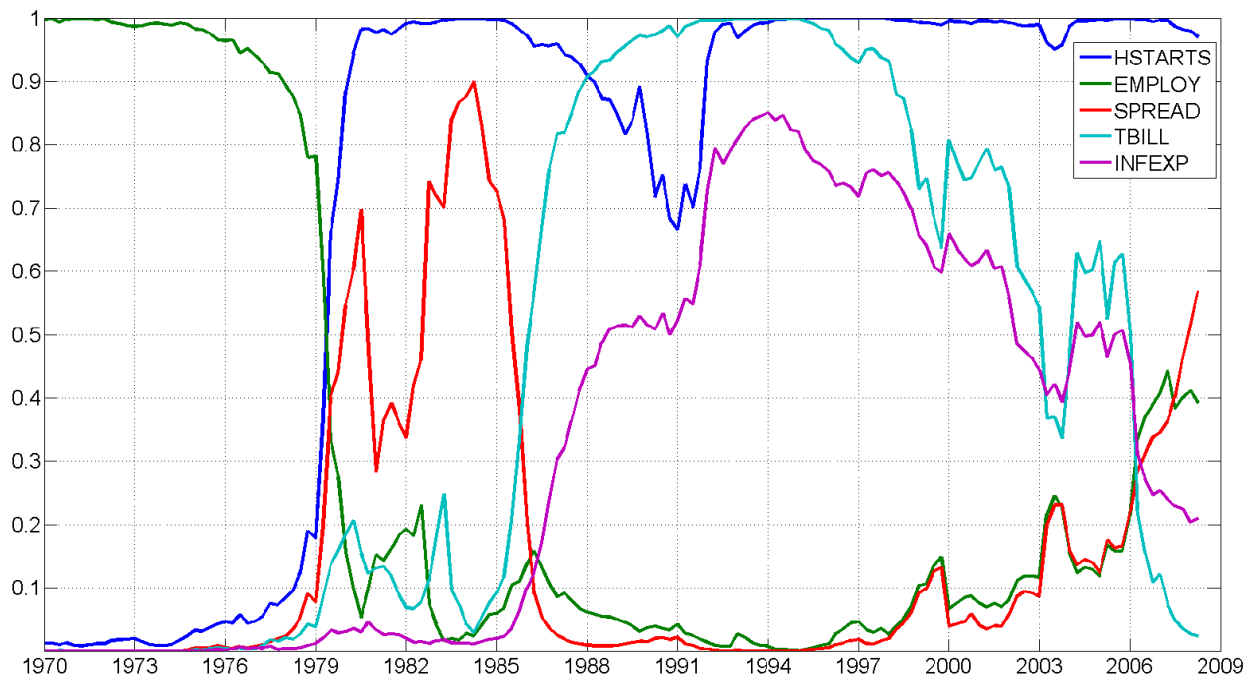


Figure 7: Posterior Probability of Inclusion of Main Predictors (GDP deflator inflation, $h = 8$)

3.2 Forecast Performance: DMA versus Alternative Forecast Procedures

There are many metrics for evaluating forecast performance and many alternative forecasting methodologies that we could compare our DMA and DMS forecasts to. In this paper, we present two forecast comparison metrics involving point forecasts. These are mean squared forecast error (MSFE) and mean absolute forecast error (MAFE). We also present a forecast metric which involves the entire predictive distribution: the sum of log predictive likelihoods. Predictive likelihoods are motivated and described in many places such as Geweke and Amisano (2007). The predictive likelihood is the predictive density for y_t (given data through time $t - 1$) evaluated at the actual outcome. The formula for the one-step ahead predictive density in model l was denoted by $p_l(y_t|y^{t-1})$ above and can be calculated as described in Section 2.3. We use the direct method of forecasting and, hence, the log predictive density for the h -step ahead forecast is the obvious extension of this. We use the sum of log predictive likelihoods for forecast evaluation, where the sum begins in 1970Q1. MSFEs and MAFEs are also calculated beginning in 1970Q1.

In terms of alternative forecasting methods, we present results for:

- Forecasts using DMA with $\alpha = \lambda = 0.99$.
- Forecasts using DMS with $\alpha = \lambda = 0.99$.
- Forecasts using a single model containing all the predictors, but with time varying parameters (i.e. this is a special case of DMA or DMS where 100% of the prior weight is attached to the model with all the predictors, but all other modelling choices are identical including $\lambda = 0.99$). This is labelled TVP in the tables.
- Forecasts using DMA, but where the coefficients do not vary over time in each model (i.e. this is a special case of DMA where $\lambda = 1$).
- Forecasts using BMA (i.e. this is a special case of DMA where $\lambda = \alpha = 1$).
- Recursive OLS forecasts using an AR(2) model.
- Recursive OLS forecasts using all of the predictors.
- Random walk forecasts.

The final three methods are not Bayesian, so no predictive likelihoods are presented for these cases.

Tables 1 and 2 present results for our forecasting exercise for our two different measures of inflation. The big picture story is a clear and strong one: DMA and DMS forecast well. In most cases much better than other forecasting methods and in no case much worse than the best alternative method. We elaborate on these points below.

Consider first the log predictive likelihoods (the preferred method of Bayesian forecast comparison). These always indicate that DMA or DMS forecasts best. One message coming out of Tables 1 and 2 is that simply using a TVP model with all predictors leads to very poor forecasting performance. Of course, we are presenting results for only a single empirical exercise. But TVP models such as TVP-VARs are gaining increasing popularity in macroeconomics and the very poor forecast performance of TVP models found in Tables 1 and 2 should serve as a caution to users of such models (at least in forecasting exercises). Clearly, we are finding that the shrinkage provided by DMA or DMS is of great value in forecasting.

DMA and DMS extend conventional forecasting approaches by allowing for model evolution and parameter evolution. A message provided by the predictive likelihoods is that most of the improvements in forecasting performance found by DMA or DMS are due to their treatment of model evolution rather than parameter evolution. That is, DMA with constant coefficient models typically forecasts fairly well and occasionally even leads

to the best forecast performance (see the results in Tables 1 and 2 for $h = 1$). Recently, macroeconomists have been interested in building models involving parameter change of various sorts. Our results suggest that allowing for the model to change is at least as important. At short horizons, conventional BMA forecasts fairly well, but at longer horizons it tends to forecast poorly.

Predictive likelihoods also consistently indicates that DMS forecasts a bit better than DMA (although this result does not carry over to MAFEs and MSFEs where DMA tends to do better). DMS and DMA can be interpreted as doing shrinkage in different ways. DMS puts zero weight on all models other than the one best model, thus “shrinking” the contribution of all models except one towards zero. It could be that this additional shrinkage provides some additional forecast benefits over DMA.

If we turn our attention to results using MSFE and MAFE, we can see that the previous picture still holds (although DMA does somewhat better relative to DMS than we found using predictive likelihoods). In addition, we can say that naive forecasting methods such as using an AR(2) or random walk model are clearly inferior to DMA and DMS for both measures of inflation at all forecast horizons. However, with CPI inflation, recursive OLS forecasting using all the predictors does well at the long horizon ($h = 8$). Forecasting at such a long horizon is difficult to do, so it is unclear how much weight to put on this result (and predictive likelihoods for this non-Bayesian method are not calculated). But it is worth noting that the good performance of recursive OLS in this case is not repeated for inflation measured using the GDP deflator nor at shorter horizons.

Table 1: Comparing Different Forecasting Methods: CPI inflation			
Forecast Method	Sum of log pred. like.	MSFE	MAFE
	$h = 1$		
DMA	-85.31	47.48	26.37
DMS	-82.26	48.96	27.82
TVP	-182.36	54.70	32.20
DMA ($\lambda = 1$)	-81.63	45.00	23.02
BMA (DMA with $\alpha = \lambda = 1$)	-84.12	46.07	24.14
Recursive OLS – AR(2)	–	57.52	41.58
Recursive OLS – All Preds.	–	52.76	34.16
Random Walk	–	54.59	35.14
	$h = 4$		
DMA	-108.73	59.34	41.28
DMS	-103.91	59.02	41.02
TVP	-178.30	82.28	72.24
DMA ($\lambda = 1$)	-111.97	60.22	41.94
BMA (DMA with $\alpha = \lambda = 1$)	-123.32	65.43	47.28
Recursive OLS – AR(2)	–	75.00	66.17
Recursive OLS – All Preds	–	58.36	42.08
Random Walk	–	77.30	66.75
	$h = 8$		
DMA	-121.02	67.56	57.31
DMS	-120.29	68.44	63.69
TVP	-154.04	72.82	62.40
DMA ($\lambda = 1$)	-129.76	69.45	60.38
BMA (DMA with $\alpha = \lambda = 1$)	-145.80	79.42	75.77
Recursive OLS – AR(2)	–	83.43	81.11
Recursive OLS – All Preds	–	62.85	46.32
Random Walk	–	99.24	117.35

Forecast Method	Sum of log pred. like.	MSFE	MAFE
$h = 1$			
DMA	-27.10	34.47	12.98
DMS	-24.97	35.61	13.70
TVP	-176.90	38.85	16.99
DMA ($\lambda = 1$)	-21.47	33.17	12.02
BMA (DMA with $\alpha = \lambda = 1$)	-25.00	34.58	13.10
Recursive OLS – AR(2)	–	40.10	17.34
Recursive OLS – All Preds	–	37.34	14.30
Random Walk	–	37.39	15.19
$h = 4$			
DMA	-23.81	35.13	15.23
DMS	-18.20	35.72	15.42
TVP	-180.20	43.82	23.33
DMA ($\lambda = 1$)	-33.38	39.68	19.57
BMA (DMA with $\alpha = \lambda = 1$)	-46.41	42.85	21.64
Recursive OLS – AR(2)	–	48.54	28.97
Recursive OLS – All Preds	–	43.09	19.92
Random Walk	–	44.28	24.57
$h = 8$			
DMA	-59.79	47.73	29.01
DMS	-59.43	48.88	30.38
TVP	-184.60	57.84	33.82
DMA ($\lambda = 1$)	-62.95	47.60	27.78
BMA (DMA with $\alpha = \lambda = 1$)	-81.22	53.37	33.20
Recursive OLS – AR(2)	–	62.90	39.48
Recursive OLS – All Preds	–	53.42	32.33
Random Walk	–	59.16	46.07

3.3 Sensitivity Analysis

Our previous DMA and DMS results were for our benchmark case where $\alpha = \lambda = 0.99$. As discussed previously, researchers in this field choose pre-selected values for α and λ and the interval (0.95, 0.99) is the empirically sensible one for most empirical applications.

It would be possible to choose α and λ in a data-based fashion, but this is typically not done for computational reasons. For instance, the researcher could select a grid of values for these two forgetting factors and then do DMA at every possible combination of values for α and λ . Some metric (e.g. an information criteria or the sum of log predictive likelihoods through time $t - 1$) could be used to select the preferred combination of α and λ at each point in time. However, this would turn an already computationally demanding exercise to one which was g^2 times as demanding (where g is the number of values in the grid). Accordingly, researchers such as Raftery et al (2007) simply go with $\alpha = \lambda = 0.99$ and argue that results will be robust to reasonable changes in these factors. In order to investigate such robustness claims, Tables 3 and 4 present results for our forecasting exercise using different combinations of the forgetting factors.

Overall, Tables 3 and 4 reveal a high degree of robustness to choice of α and λ . If anything, these tables emphasize the benefits of DMA in that measures of forecast performance are sometimes better than those in Tables 1 and 2 and rarely much worse.

One finding of particular interest is that the combination $\alpha = 0.95$ and $\lambda = 0.99$ tends to forecast very well, for both of our measures of inflation. Note that the value $\alpha = 0.95$ allows for quite rapid change in forecasting model over time. This is consistent with a story we have told before: that it appears that allowing for models to change over time is more important in improving forecast performance than allowing for parameters to change (at least in our data sets).

Table 3: Sensitivity Analysis: CPI inflation			
Forecast Method	Sum of log pred. like.	MSFE	MAFE
	$h = 1$		
DMA, $\alpha = \lambda = 0.95$	-107.19	47.29	24.87
DMS, $\alpha = \lambda = 0.95$	-74.57	43.94	23.16
DMA, $\alpha = 0.99, \lambda = 0.95$	-95.39	58.30	26.61
DMS, $\alpha = 0.99, \lambda = 0.95$	-87.24	48.92	28.81
DMA, $\alpha = 0.95, \lambda = 0.99$	-91.48	45.63	22.57
DMS, $\alpha = 0.95, \lambda = 0.99$	-52.04	40.53	21.65
	$h = 4$		
DMA, $\alpha = \lambda = 0.95$	-106.25	54.26	34.31
DMS, $\alpha = \lambda = 0.95$	-57.34	46.44	26.19
DMA, $\alpha = 0.99, \lambda = 0.95$	-101.91	56.16	35.94
DMS, $\alpha = 0.99, \lambda = 0.95$	-98.64	59.34	42.43
DMA, $\alpha = 0.95, \lambda = 0.99$	-100.38	54.87	34.46
DMS, $\alpha = 0.95, \lambda = 0.99$	-61.37	47.63	26.36
	$h = 8$		
DMA, $\alpha = \lambda = 0.95$	-98.15	56.19	33.84
DMS, $\alpha = \lambda = 0.95$	-51.29	47.28	28.68
DMA, $\alpha = 0.99, \lambda = 0.95$	-111.58	64.04	51.40
DMS, $\alpha = 0.99, \lambda = 0.95$	-114.02	67.02	56.72
DMA, $\alpha = 0.95, \lambda = 0.99$	-92.93	56.48	36.06
DMS, $\alpha = 0.95, \lambda = 0.99$	-66.48	51.30	31.35

Table 4: Sensitivity Analysis: GDP Deflator inflation			
Forecast Method	Sum of log pred. like.	MSFE	MAFE
$h = 1$			
DMA, $\alpha = \lambda = 0.95$	-66.23	36.20	14.03
DMS, $\alpha = \lambda = 0.95$	-46.81	38.12	15.96
DMA, $\alpha = 0.99, \lambda = 0.95$	-48.89	38.04	15.75
DMS, $\alpha = 0.99, \lambda = 0.95$	-48.56	38.77	16.52
DMA, $\alpha = 0.95, \lambda = 0.99$	-34.45	32.17	11.33
DMS, $\alpha = 0.95, \lambda = 0.99$	-0.11	30.76	10.84
$h = 4$			
DMA, $\alpha = \lambda = 0.95$	-26.49	35.52	15.26
DMS, $\alpha = \lambda = 0.95$	-10.48	37.48	18.60
DMA, $\alpha = 0.99, \lambda = 0.95$	-37.50	42.46	22.79
DMS, $\alpha = 0.99, \lambda = 0.95$	-36.97	43.24	24.02
DMA, $\alpha = 0.95, \lambda = 0.99$	-13.04	32.25	13.48
DMS, $\alpha = 0.95, \lambda = 0.99$	25.36	28.87	11.81
$h = 8$			
DMA, $\alpha = \lambda = 0.95$	-42.43	37.90	18.04
DMS, $\alpha = \lambda = 0.95$	-19.91	39.48	22.48
DMA, $\alpha = 0.99, \lambda = 0.95$	-57.95	46.96	29.40
DMS, $\alpha = 0.99, \lambda = 0.95$	-58.65	48.58	30.31
DMA, $\alpha = 0.95, \lambda = 0.99$	-36.93	38.19	18.36
DMS, $\alpha = 0.95, \lambda = 0.99$	-12.26	37.47	20.51

4 Conclusions

This paper has investigated the use of DMA and DMS methods for forecasting US inflation. These extend conventional approaches by allowing for the set of predictors for inflation to change over time. When you have K models and a different one can potentially hold at each of T points in time, then the resulting K^T combinations can lead to serious computational and statistical problems (regardless of whether model averaging or model selection is done). As shown in this paper, DMA and DMS handle these problems in a simple, elegant and sensible manner.

In our empirical work, we present evidence indicating the benefits of DMA and DMS. In particular, it does seem that the best predictors for forecasting inflation are changing

considerably over time. By allowing for this change, DMA and DMS lead to substantial improvements in forecast performance.

References

- Ang, A. Bekaert, G., Wei, M., 2007. Do macro variables, asset markets, or surveys forecast inflation better? *Journal of Monetary Economics* 54, 1163-1212.
- Atkeson, A., Ohanian, L., 2001. Are Phillips curves useful for forecasting inflation? *Federal Reserve Bank of Minneapolis Quarterly Review* 25, 2-11.
- Avramov, D., 2002. Stock return predictability and model uncertainty. *Journal of Financial Economics* 64, 423-458.
- Chen, R., Liu, J., 2000. Mixture Kalman filters. *Journal of the Royal Statistical Society, Series B* 62, 493-508.
- Cogley, T., Sargent, T., 2005. Drifts and volatilities: Monetary policies and outcomes in the post WWII U.S.. *Review of Economic Dynamics* 8, 262-302.
- Cogley, T., Morozov, S., Sargent, T., 2005. Bayesian fan charts for U.K. inflation: Forecasting and sources of uncertainty in an evolving monetary system. *Journal of Economic Dynamics and Control* 29, 1893-1925.
- Cremers, K., 2002. Stock return predictability: A Bayesian model selection perspective. *Review of Financial Studies* 15, 1223-1249.
- Fagin, S., 1964. Recursive linear regression theory, optimal filter theory, and error analyses of optimal systems. *IEEE International Convention Record Part i*, 216-240.
- Geweke, J., Amisano, J., 2007. Hierarchical Markov normal mixture models with applications to financial asset returns.
Manuscript available at <http://www.biz.uiowa.edu/faculty/jgeweke/papers/paperA/paper.pdf>.
- Green, P., 1995. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* 82, 711-732.
- Groen, J., Paap, R., Ravazzolo, F., 2008. Real-time inflation forecasting in a changing world. manuscript.
- Jazwinsky, A., 1970. *Stochastic Processes and Filtering Theory*. New York: Academic Press.
- Koop, G., Leon-Gonzalez, R., Strachan, R., 2009. On the evolution of the monetary policy transmission mechanism. *Journal of Economic Dynamics and Control* 33, 997-1017.
- Koop, G., Potter, S., 2004. Forecasting in dynamic factor models using Bayesian model averaging. *The Econometrics Journal* 7, 550-565.
- Korobilis, D., 2009. Assessing the transmission of monetary policy shocks using dynamic factor models. Discussion Paper 9-14, University of Strathclyde.
- Madigan, D., York, J., 1995. Bayesian graphical models for discrete data. *International Statistical Review* 63, 215-232.
- Primiceri, G., 2005. Time varying structural vector autoregressions and monetary

policy. *Review of Economic Studies* 72, 821-852.

Raftery, A., Karny, M., Andrysek, J. and Ettler, P., 2007. Online prediction under model uncertainty via dynamic model averaging: Application to a cold rolling mill. Technical report 525, Department of Statistics, University of Washington.

Smith, J. and Miller, J., 1986. A non-Gaussian state space model and application to prediction records. *Journal of the Royal Statistical Society, Series B* 48, 79-88.

Stock, J. and Watson, M., 1999. Forecasting inflation. *Journal of Monetary Economics* 44, 293-335.

Stock, J. and Watson, M., 1996. Evidence on structural instability in macroeconomic time series relations. *Journal of Business and Economic Statistics* 14, 11-30.

Data Appendix

The variables used in this study were taken from the sources in table below. All series were seasonally adjusted, where applicable, and run from 1959:Q1 to 2008:Q2. Some series in the database were observed on a monthly basis and quarterly values were computed by averaging the monthly values over the quarter. All variables are transformed to be approximately stationary. In particular, if $z_{i,t}$ is the original untransformed series, the transformation codes are (column Tcode below): 1 - no transformation (levels), $x_{i,t} = z_{i,t}$; 2 - first difference, $x_{i,t} = z_{i,t} - z_{i,t-1}$; 4 - logarithm, $x_{i,t} = \log z_{i,t}$; 5 - first difference of logarithm, $x_{i,t} = \log z_{i,t} - \log z_{i,t-1}$.

#	Mnemonic	Tcode	Description	Source
1	GDPDEFL	5	Gross Domestic Product: Implicit Price Deflator	FRED
2	CPI	5	Consumer Price Index For All Urban Consumers: All Items	FRED
3	UNEMP	1	Civilian Unemployment Rate	FRED
4	CONS	5	Real Personal Consumption Expenditures	FRED
5	INV	5	Private Residential Fixed Investment	FRED
6	GDP	5	Real Gross Domestic Product, 3 Decimal	FRED
7	HSTARTS	4	Housing Starts: Total: New Privately Owned Housing Units Started	FRED
8	EMPLOY	5	All Employees: Total Private Industries	FRED
9	PMI	2	ISM Manufacturing: PMI Composite Index	FRED
10	COMPRICE	2	NAPM Commodity Prices Index (Percent)	Bloomberg
11	VENDOR	2	NAPM Vendor Deliveries Index (Percent)	Bloomberg
12	WAGE	5	Average Hourly Earnings: Manufacturing	FRED
13	TBILL	1	3-Month Treasury Bill: Secondary Market Rate	FRED
14	SPREAD	1	Spread 10-year T-Bond yield / 3-month T-Bill (GS10 -TB3MS)	FRED
15	DJIA	5	Dow Jones Industrial Average	Bloomberg
16	MONEY	5	M1 Money Stock	FRED
17	INFEXP	1	University of Michigan Inflation Expectations	Uni. of Mich.