Journal of Empirical Finance 44 (2017) 158–176

Contents lists available at ScienceDirect

Journal of Empirical Finance

journal homepage: www.elsevier.com/locate/jempfin

Using dynamic model averaging in state space representation with dynamic Occam's window and applications to the stock and gold market

Marian Risse [a,*], Ludwig Ohl [b]

[a] Department of Economics, Helmut Schmidt University, Holstenhofweg 85, P.O.B. 700822, 22008 Hamburg, Germany
[b] Center for Medical Biotechnology, University of Duisburg–Essen, Universitätsstraße 2, 45141 Essen, Germany

ARTICLE INFO

JEL classification:
C32
C53
C58
F37
F47
G17

Keywords:
Dynamic model averaging
State space representation
Dynamic occam's window
Forecasting
Trading rule

ABSTRACT

We combine the Onorante and Raftery (2016) dynamic Occam's window approach with the Raftery et al. (2010) DMA/DMS estimator in state space representation to create forecasts using a data-rich forecasting environment. Our approach is mainly related to economic and financial time series that are subject to periods of high volatility, which increases the necessity of a time varying parameter framework. In a forecasting exercise for the stock and gold markets, we highlight the economic value-added of our approach by applying a simple trading rule to the return series. By combining both assets, we show that our approach performs better when compared to alternative forecasting models such as machine learning algorithms and standard DMA/DMS. Results for the complexity of the forecasting models highlight the advantages of high dimensional forecasting approaches in times of economic uncertainty, such as the recent financial crisis. The economic performance of the trading rule weakens when we consider transaction costs.

1. Introduction

Over the last years, model averaging and selection techniques have attracted the attention of many researchers as tools to improve forecasting performance. The techniques discussed in this strand of research are based on the fundamental idea that a forecast should not be derived from a single forecasting model, but should rather be derived by combining or selecting forecasts from a potentially large model space. For example, Pesaran and Timmermann (1995) use ordinary least squares (OLS) regressions to predict out-of-sample excess returns of the S&P 500 index. To account for a large model space, they consider all possible combinations of predictor variables that are expected to influence future returns. Forecast selection is done by using the forecasting model that maximizes an information criterion (in-sample) such as the classical R^2 or the AIC.[1] Averaging these forecasts among the model space is also known as thick modeling (Aiolfi and Favero, 2005).

To overcome the problem of model uncertainty, researches extensively use Bayesian model averaging (BMA) and selection (BMS) to simulate the full predictive density needed to compute individual model probabilities and to weight forecasts (see Raftery et al., 1997; Avramov, 2002; Cremers, 2002; Sala-I-Martin et al., 2004, among others). However, it is well known that BMA/BMS can be very time consuming because, in every period of time, a sufficiently large number of draws from the posterior distribution is necessary to create stable prediction results. More recently, Raftery et al. (2010) propose *dynamic* model averaging (DMA) and selection (DMS) as

* Corresponding author.
E-mail address: marian.risse@hsu-hh.de (M. Risse).

[1] The authors also use an direction of change criterion that is based on the sign of the forecast.

https://doi.org/10.1016/j.jempfin.2017.09.005
Received 29 November 2016; Received in revised form 3 July 2017; Accepted 26 September 2017
Available online 5 October 2017
0927-5398/

a relatively fast computational alternative to BMA/BMS. They use forgetting factors to achieve an analytical solution for the predictive density which avoids using Markov chain Monte Carlo (MCMC) sampling. In a series of publications, Koop and Korobilis (2012), Koop and Tole (2013), Koop and Korobilis (2013, 2014) have introduced DMA/DMS into the economics literature. Meanwhile, the approach is partially extended to, for example, automatically determining the degree of discounting (Bork and Møller, 2015; Risse and Kern, 2016) or to improve the underlying Kalman filter itself (Grassi et al., 2016).

In fact, a huge disadvantage of the forecast selection or thick modeling approach and its Bayesian counterpart is that a researcher has to estimate the full model space to receive a successfully aggregated or selected forecast. Pesaran and Timmermann (1995) use a dataset consisting of nine predictors which ends up in 512 forecasting models to be estimated in every period of time. Koop and Korobilis (2012) forecast inflation using a medium-sized dataset including 14 macroeconomic, monetary, and financial predictors and 16 384 models to be estimated in every period of time.[2] While personal computers do not have problems computing forecasts from less than ten predictors, computational burden drastically increases when using more than 20–25 predictors. For example, a forecasting model with 20 predictors ends up in more than one million models to be computed for the DMA/DMS forecasts in every period of time.

While economic time series might be influenced by a broad set of predictors over the full evaluation period, they can be expected to be periodically influenced only by a few predictors (Stock and Watson, 2006). Therefore, using a large set of information for building the forecasting environment is not always advantageous because a researcher might expect an overwhelmingly large number of models to be negligible for forming an optimal forecast. With regard to forecasting the inflation rate, Koop and Korobilis (2012) show that the number of predictors is on average less than approximately one third of the full predictor space and not stable over time. Following this scenario, it is reasonable for researches to shrink the model space to achieve feasible computation without losing too much information and to maintain the advantages of BMA/BMS.

Recently, Onorante and Raftery (2016) introduce a new econometric approach to prevent estimating the full model space and instead to use a subset (a *population*) of models.[3] They call their method *Dynamic Occam's Window* (DOW). DOW is an algorithm that pre-selects, in every period of time, the amount of data that is considered for the underlying econometric model. It does not directly depend on the econometric model itself. Using a database of 30 predictors, Onorante and Raftery (2016) then apply recursive OLS regressions to forecast GDP growth in Europe.

In this paper, we extend the DOW approach by implementing a full-fledged time-varying parameter model in state–space representation. Using such a representation is advantageous, because we control for the speed of adjustment of the individual model coefficients. Our approach is superior to the original recursive OLS estimation since it explicitly assumes parameter instability. Using a state–space representation as in Koop and Korobilis (2012), we can further control for time-varying volatility by implementing an exponentially weighted moving average process. Both aspects become important when predicting macroeconomic or financial time series that are subject to sudden changes (Dangl and Halling, 2012; Johannes et al., 2014).

In a practical application, we use the extended DOW estimator to forecast two time series that have extensively been studied in recent years: the price of gold and the S&P 500 index. Both series are expected to be influenced by a broad range of predictors related to financial markets, monetary economics, and the real economy.[4] We use a combined database of 35 time series for both assets to form our aggregated and selected forecasts.[5] We abstract from a pure statistical forecast evaluation and focus instead on the economic value-added of forecasts. For this reason, we use an extended trading strategy similar to Pesaran and Timmermann (1995, 2000), where an investor permanently switches between an investment in the stock/gold market and an alternative risk-free investment in bonds. Since several authors report that correlation between the price of gold and the stock market is relatively low (or negative, see Hillier et al., 2006; Baur and Lucey, 2010; Baur and McDermott, 2010), this property renders it possible to simulate a trading strategy of a risk-neutral investor that uses the price of gold as a profitable alternative to the S&P 500 index (or vice versa).[6]

We finally compare the DOW trading results with (naive) alternative forecasting techniques such as a least shrinkage regression, a machine learning approach, and subsets of the standard DMA/DMS framework. We additionally introduce a passive buy-and-hold strategy and transaction costs. Our results indicate an extensive increase in performance compared to such alternative forecasting models. We can also show that model complexity increases in turbulent economic times such as the bust of the dot-com bubble or the recent financial crisis demonstrating the necessity of a large database. Finally, we consider several robustness checks.

We organize the remainder of this paper as follows. In Section 2, we lay out the econometric methodology and its evaluation. In Section 3, we present our data. In Section 4, we present our empirical results. Robustness checks are given in Section 5, and we conclude in Section 6.

[2] The exact determination of models follows a simple exponential process. For K predictors in the dataset, 2^K models have to be estimated.

[3] The approach is related to optimal prediction pools, proposed by Geweke and Amisano (2011).

[4] For an overview of recent research on the stock market, see Spiegel (2008). For a recent review of research on the gold market, see O'Connor et al. (2015). A statistically motivated application of the standard DMA/DMS approach on the gold market is published by Aye et al. (2015) and Baur et al. (2016). For a relatively small dataset, Beckmann and Schüssler (2015) apply the standard DMA/DMS approach on the stock market.

[5] To our knowledge, Mittnik et al. (2015) is the only study that uses a similar large dataset to analyze stock prices, but with the emphasis on studying the volatility rather than the returns of the price series. Their econometric methodology is based on a machine learning approach that we also use in this paper for forecast comparison.

[6] Our approach is related for testing the hedging and safe haven properties of both assets, where we focus on an out-of-sample forecasting exercise (for definitions, see Baur and McDermott, 2010).