International Journal of Forecasting 32 (2016) 650–668

Contents lists available at ScienceDirect

International Journal of Forecasting

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

Getting the most out of macroeconomic information for predicting excess stock returns

Cem Çakmaklı [a,b,*], Dick van Dijk [c,d,e,1]

[a] Department of Economics, Koç University, Turkey
[b] Department of Quantitative Economics, University of Amsterdam, Netherlands
[c] Econometric Institute, Erasmus University Rotterdam, Netherlands
[d] Tinbergen Institute, Netherlands
[e] Erasmus Research Institute for Management (ERIM), Netherlands

ARTICLE INFO

Keywords:
Return predictability
Model uncertainty
Dynamic factor models
Variable selection

ABSTRACT

This paper documents the fact that the factors extracted from a large set of macroeconomic variables contain information that can be useful for predicting monthly US excess stock returns over the period 1975–2014. Factor-augmented predictive regression models improve upon benchmark models that include only valuation ratios and interest rate related variables, and possibly individual macro variables, as well as the historical average excess return. The improvements in out-of-sample forecast accuracy are significant, both statistically and economically. The factor-augmented predictive regressions have superior market timing abilities, such that a mean–variance investor would be willing to pay an annual performance fee of several hundreds of basis points to switch from the predictions offered by the benchmark models to those of the factor-augmented models. One important reason for the superior performance of the factor-augmented predictive regressions is the stability of their forecast accuracy, whereas the benchmark models suffer from a forecast breakdown during the 1990s.

1. Introduction

Stock return predictability remains an issue of intense debate. On the one hand, substantial numbers of papers have reported positive forecasting results, with the most successful predictor variables being valuation ratios such as the dividend yield, the price-earnings ratio, and the book-to-market ratio, or other financial variables such as the short-term interest rate, the yield spread, or the credit spread, see Ang and Bekaert (2007); Campbell and Shiller (1988); Campbell and Thompson (2008); Fama and French (1988); Ferson and Harvey (1991); Keim and Stambaugh (1986) and Wachter and Warusawitharana (2009), among many others. On the other hand, studies such as that of Goyal and Welch (2008) contest this view, arguing that no single predictor variable outperforms the historical mean return in terms of forecast accuracy over long time spans, either in-sample or out-of-sample.

Interestingly, although stock returns are assumed to be linked to business conditions, macroeconomic variables such as output growth and inflation do not seem to add any predictive power for stock returns beyond that provided by the above-mentioned valuation ratios and interest rate

* Correspondence to: College of Administrative Sciences and Economics, Koç University, Rumelifeneri Yolu, Sarıyer, 34450, Istanbul, Turkey.
E-mail addresses: ccakmakli@ku.edu.tr (C. Çakmaklı), djvandijk@ese.eur.nl (D. van Dijk).
[1] P.O. Box 1738, NL-3000 DR Rotterdam, The Netherlands.

http://dx.doi.org/10.1016/j.ijforecast.2015.10.001
0169-2070/

related variables. The few exceptions to this rule include the consumption-wealth ratio (Lettau & Ludvigson, 2001), the expenditure share on housing (Piazzesi, Schneider, & Tuzel, 2007), expenditures on durables as a fraction of its stock (Gomes, Kogan, & Yogo, 2009), and survey-based measures of expected business conditions (Campbell & Diebold, 2009).

Several complicating issues plague research on the predictive content of macroeconomic information for stock returns. First and foremost, while numerous macroeconomic variables are available, typically only a few of them are considered as possible predictors of returns. Presumably, the information embedded in the selected variables is more limited than that which is actually available to an investor, who might consider a much larger set of macro variables. Obviously, there are sound statistical reasons for limiting the number of (macro) variables in a predictive regression model. Expanding the set of predictors exacerbates the parameter estimation uncertainty, which translates into additional uncertainty in the resulting return forecasts. At the same time, when the number of predictors is restricted, it becomes crucial to select the most informative variables for consideration. This essentially leads to the issue of model uncertainty, as it is not clear *a priori* which macroeconomic variables contain the most relevant information for predicting stock returns.

Second, the relationships between excess stock returns and individual predictors appear to be highly unstable, see Bossaerts and Hillion (1999); Lettau and Van Nieuwerburgh (2008); Paye and Timmermann (2006); Pesaran and Timmermann (2002); Pettenuzzo and Timmermann (2011); Rapach and Wohar (2006); and Ravazzolo, van Dijk, Paap, and Franses (2008); among others. Hence, if a given macroeconomic variable is found to be useful for forecasting stock returns over a certain period, this may very well break down at some point. Put differently, individual variables' predictive abilities fluctuate strongly over time. Note that this aggravates the model uncertainty problem, as the set of variables to be included in a predictive regression model is unstable.

In this paper, we use dynamic factor models to enable us to handle the issues of model uncertainty, parameter estimation uncertainty, and structural instability jointly. Specifically, we use principal component analysis to construct a small number of factors from a large set of macro variables. These factors are then included in a predictive regression model for excess stock returns. The motivation for considering this approach is derived from macroeconomic forecasting, where various successful applications of factor models for predicting variables such as output growth and inflation have appeared in recent years, see Forni, Hallin, Lippi, and Reichlin (2003) and Stock and Watson (2002b), among others, as well as the survey by Stock and Watson (2011). The same approach was also adopted by Ludvigson and Ng (2007) for examining the relationship between expected stock returns and volatility, and by Ludvigson and Ng (2009) for analyzing the effects of macroeconomic conditions on bond risk premia. Our paper is related most closely to the independent concurrent research by Bai (2010), which also demonstrates that excess stock return predictions can be improved by exploiting the factors obtained from a large set of macro variables.[2] Neely, Rapach, Tu, and Zhou (2014) also use factor models to forecast the equity premium, but extract the factors from a limited set of economic variables (14), combined with a number of technical indicators computed from moving average, momentum, and volume-based rules.

When using the factor-based approach for forecasting excess stock returns, several choices need to be made. For one thing, we have to decide which individual variables to use in the factor construction, and how many factors to include in the predictive regression model. While we adopt a standard approach for the latter issue (following previous studies in the macro forecasting literature), we examine the former issue in detail. Specifically, we consider the benefits of several procedures for pre-selecting the macro variables that enter the factor construction stage, relative to the standard approach of simply including all available variables. This relates to the issue that the standard approach assumes the factors to be the first few principal components of the macro variables, which are constructed without taking into account the purpose for which the factors are to be used. If the factors are used in predictive regressions as we do in the present paper, it will not necessarily be the case that these principal components contain the most relevant information for the variable that we aim to forecast, as has been shown by Bai and Ng (2008) and Boivin and Ng (2006), among others. The variable pre-selection procedures that we consider here take into account the predictive ability of the individual macro variables for the excess stock returns explicitly.

We conduct an empirical analysis of the usefulness of the factor-augmented regression models for the out-of-sample prediction of monthly US excess stock returns over the period January 1975–December 2014. In addition to the most commonly applied financial variables, such as the dividend yield and the short-term interest rate, we also construct factors based on a comprehensive set of 118 macroeconomic variables. The value added by the inclusion of these factors in the predictive regressions is assessed in both statistical and economic terms. On the one hand, we consider the directional accuracy of the forecasts to examine the market timing ability of the factor-augmented predictive regressions. For this purpose, we use hit ratios, defined as the proportion of forecasts for which the actual and predicted signs of the excess returns match. On the other hand, we assess the economic value of the return forecasts by using them in active mean–variance investment strategies. In addition to standard performance measures such as the Sharpe ratio, we also use a utility-based metric to evaluate how much an investor would be willing to pay to use the predictions from the factor-augmented models rather than those from a benchmark model.

[2] Several features distinguish our paper from that of Bai (2010). The most fundamental methodological difference is the fact that we examine the possibility of pre-selecting macro variables that are used in the factor construction, while Bai (2010) examines the possibility of selecting the most informative factors after their construction.