Economics Letters 121 (2013) 19–22

Contents lists available at SciVerse ScienceDirect

Economics Letters

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

Perpetual learning and stock return predictability

Xiaoneng Zhu *

The Chinese Academy of Finance and Development, Central University of Finance and Economics, South Xueyuan Road #39, Beijing, 100081, China

HIGHLIGHTS

- Perpetual learning plays a significant role in excess return forecasts.
- Perpetual learning generates economic gains in portfolio management.
- The results suggest a slow learning process in the stock market.

ARTICLE INFO

Article history:
Received 13 February 2013
Received in revised form
26 June 2013
Accepted 26 June 2013
Available online 2 July 2013

JEL classification:
G10
G17

Keywords:
Excess return
Learning
Forecasts
Stock returns

ABSTRACT

The stock market is evolving, and investors are learning. This paper investigates the role of perpetual learning in excess return forecasts. We find that perpetual learning usually delivers statistically and economically significant out-of-sample gains relative to the historical average.

1. Introduction

Stock market investors face various forms of uncertainty when forecasting excess returns. First, investors do not know predictive variables and the functional form of the true return-generating process, and so face model uncertainty (see, e.g., Avramov (2002)). Second, conditional on a particular forecasting model, investors are uncertain about the true extent of return predictability, as return predictability is time varying (see, e.g., Pettenuzzo and Timmermann (2011)), so investors face model instability. Third, investors do not know the true parameters of a forecasting model, and so face uncertainty about the parameters, also known as estimation risk (see, e.g., Barberis (2000)).

These uncertainties give rise to a highly complex and constantly evolving data-generating process for expected excess returns. In such a case, investors may use estimation methods to learn about the true forecasting model (see, e.g., Branch and Evans (2006), Evans and Honkapohja (2001), and Timmermann (1996)) which matters for the projection of future predictors and excess returns. These uncertainties therefore highlight the importance of perpetual learning (see, e.g., Sargent (1999)) in excess return forecasts. In this paper, we attempt to provide insights into the role of perpetual learning, which is captured by the discounted least squares (DLS) procedure, in excess return forecasts.[1] Indeed, perpetual learning is justified by a theoretical framework along the lines of Merton's (1973) intertemporal capital asset pricing model (ICAPM). Merton (1973) shows that the conditional excess return is related to the conditional covariance between excess returns and the innovations in state variables. If the conditional covariance is time varying, perpetual learning is a potential way of capturing the time-varying conditional covariance. So, perpetual learning is related to the conditional excess return.

Empirically, we use the DLS learning approach and the ordinary least squares (OLS) approach to estimate predictive regressions for excess returns. To evaluate the importance of learning in

* Tel.: +86 10 62288675; fax: +86 10 62288779.
E-mail address: xiaonengz@gmail.com.

[1] We do not claim that perpetual learning can capture the investors' true learning process perfectly. Alternatively, perpetual learning is presumably supposed to be related to the investors' true learning process.

0165-1765/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.econlet.2013.06.035

excess return forecasts, we compare the out-of-sample forecasting performance of the DLS learning approach to that of the OLS non-learning approach.

To anticipate our results, we find that the DLS approach performs considerably better than OLS approach. Indeed, the OLS approach usually underperforms the historical average in the out-of-sample exercise, in line with Welch and Goyal (2008). In contrast, the DLS approach generally outperforms the historical average. In addition, the DLS approach delivers consistent out-of-sample economic gains relative to the historical average. Overall, these results demonstrate the importance of perpetual learning in excess return forecasts.

Our paper is closely related to those of Branch and Evans (2010, 2011, 2013), who suggest the importance of learning in understanding stock market fluctuations. Our paper also relates to that of Branch and Evans (2006), who find that a learning model forecasts inflation and output growth well out of sample. The results of this paper show that a perpetual (or constant-gain) learning model forecasts stock returns well. Furthermore, our work is related to previous studies that take a Bayesian approach to account for uncertainties and to characterize stock return predictability. For example, Barberis (2000) suggests that a long-horizon investor who ignores parameter uncertainty may overallocate to stocks by a sizeable amount.

The remainder of this paper is organized as follows. Section 2 outlines the econometric methodology and the measure for evaluating economic gain. Section 3 discusses predictive variables and presents forecasting results. Section 4 concludes. Section 5 is an appendix.

2. Methodology

A typical forecasting specification regresses excess returns on an independent lagged predictor

$$r_{t+1} = \alpha + \beta x_t + \varepsilon_{t+1}, \tag{1}$$

where r_{t+1} is the excess return and x_t is a variable whose predictive ability is of interest. An interesting empirical benchmark is to use the OLS approach to estimate and predict stock returns. To incorporate perpetual learning into the forecasting procedure, we use the DLS approach to estimate and forecast excess returns. The Appendix provides details of the DLS method. When the model (1) does not include a predictive variable, the OLS forecast is the historical average,

$$\bar{r}_{t+1} = \frac{1}{t}\sum_{j=1}^{t} r_j. \tag{2}$$

Similarly, when the predictive model does not include a predictor, the DLS forecast becomes the weighted historical average. Naturally, the historical average and the weighted historical average are two important empirical benchmarks for evaluating the forecasting performance of predictive variables.

To measure the forecasting performance, we calculate the mean square forecast error,

$$\mathrm{MSE}(r_t) = \frac{1}{T}\sum_{t=0}^{T-1}(r_{t+1} - \ddot{r}_{t+1})^2, \tag{3}$$

where $\ddot{r}_{t+1}$ is the forecast of the excess return based on t information. In addition, we also use the out-of-sample R^2 statistic, R^2_{OS}, suggested in Campbell and Thompson (2008), to compare the $\ddot{r}_{t+1}$ and $\bar{r}_{t+1}$ forecasts, where $\ddot{r}_{t+1}$ is the forecast based on the OLS/DLS predictive regression model in Eq. (1) and $\bar{r}_{t+1}$ is the (weighted) historical average estimated through period t. Specifically, R^2_{OS} is given by

$$R^2_{OS} = 1 - \frac{\sum_{t=0}^{T-1}(r_{t+1} - \ddot{r}_{t+1})^2}{\sum_{t=0}^{T-1}(r_{t+1} - \bar{r}_{t+1})^2}. \tag{4}$$

The R^2_{OS} statistic measures the reduction in mean square prediction error for the predictive regression model relative to the (weighted) historical average. Thus, when $R^2_{OS} > 0$, the $\ddot{r}_{t+1}$ forecast outperforms the $\bar{r}_{t+1}$ forecast according to the MSE metric.

The R^2_{OS} measure does not explicitly account for the risk borne by an investor over the out-of-sample period. To assess the economic value of a forecasting model, we follow Campbell and Thompson (2008) and Welch and Goyal (2008) to compute the realized utility gains accrued to a mean–variance investor on a real-time basis. Assume that the investor with relative risk-aversion parameter γ allocates his/her portfolio monthly between stocks and risk-free bills using forecasts of excess returns based on the OLS approach. At the end of period t, the investor will decide to allocate the following share of the portfolio to equities in period $t+1$:

$$w_{\mathrm{OLS},t} = \frac{1}{\gamma}\frac{\tilde{r}_{t+1}}{\tilde{\sigma}^2_{t+1}}, \tag{5}$$

where $\tilde{\sigma}^2_{t+1}$ is the rolling-window estimate of the variance of stock returns and $\tilde{r}_{t+1}$ is the OLS forecast of stock returns. Over the out-of-sample period, the investor realizes an average utility level of

$$v_{\mathrm{OLS}} = \tilde{\mu} - \frac{1}{2}\gamma\tilde{\sigma}^2, \tag{6}$$

where $\tilde{\mu}$ and $\tilde{\sigma}^2$ are respectively the sample mean and variance over the out-of-sample period for the return on the benchmark portfolio formed using excess return forecasts based on the OLS approach.

When the investor predicts stock returns using the DLS approach, he/she will choose an equity share of

$$w_{\mathrm{DLS},t} = \frac{1}{\gamma}\frac{\hat{r}_{t+1}}{\hat{\sigma}^2_{t+1}}, \tag{7}$$

where $\hat{\sigma}^2_{t+1}$ is the discounted rolling-window estimate of the variance of stock returns and $\hat{r}_{t+1}$ is the DLS forecast of stock returns. Over the out-of-sample period, the investor realizes an average utility level of

$$v_{\mathrm{DLS}} = \hat{\mu} - \frac{1}{2}\gamma\hat{\sigma}^2, \tag{8}$$

where $\hat{\mu}$ and $\hat{\sigma}^2$ are respectively the sample mean and variance over the out-of-sample period for the return on the benchmark portfolio formed using excess return forecasts based on the DLS approach. In our empirical analysis, the economic gain of learning is the difference between Eqs. (8) and (6),

$$\Delta = 1200(v_{\mathrm{DLS}} - v_{\mathrm{OLS}}). \tag{9}$$

We multiply the difference by 1200 to express it as an average annualized percentage return. The utility gain can be interpreted as the performance fee that the investor would be willing to pay to switch from an OLS trading strategy to a DLS trading strategy. The changing portfolio weights $w_{\mathrm{OLS},t}$ and $w_{\mathrm{DLS},t}$ are driven simultaneously by excess return forecasts and time-varying conditional variances. To isolate the effect of return forecasts from volatility timing, we also calculate portfolio weights by setting $\tilde{\sigma}^2_{t+1}$ and $\hat{\sigma}^2_{t+1}$ to be the unconditional volatility of excess returns. In this case, the OLS (DLS) investor realizes an average utility level of $v_{f,\mathrm{OLS}}$ ($v_{f,\mathrm{DLS}}$). Then, we evaluate the economic gain of learning using

$$\Delta_f = 1200(v_{f,\mathrm{DLS}} - v_{f,\mathrm{OLS}}). \tag{10}$$