Journal of Econometrics 186 (2015) 74–93

Contents lists available at ScienceDirect

Journal of Econometrics

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

What is the chance that the equity premium varies over time? Evidence from regressions on the dividend-price ratio[☆]

Jessica A. Wachter[a,b,*], Missaka Warusawitharana[c,1]

[a] Department of Finance, The Wharton School, University of Pennsylvania, 2300 SH-DH, Philadelphia, PA, 19104, United States
[b] NBER, United States
[c] Division of Research and Statistics, Board of Governors of the Federal Reserve System, Mail Stop 97, 20th and Constitution Ave, Washington D.C., 20551, United States

ARTICLE INFO

Article history:
Received 25 January 2012
Received in revised form
22 January 2014
Accepted 2 May 2014
Available online 11 July 2014

JEL classification:
C11
C22
G11
G17

Keywords:
Return predictability
Bayesian statistics
Model uncertainty

ABSTRACT

We examine the evidence on excess stock return predictability in a Bayesian setting in which the investor faces uncertainty about both the existence and strength of predictability. When we apply our methods to the dividend-price ratio, we find that even investors who are quite skeptical about the existence of predictability sharply modify their views in favor of predictability when confronted by the historical time series of returns and predictor variables. Correctly taking into account the stochastic properties of the regressor has a dramatic impact on inference, particularly over the 2000–2005 period.

1. Introduction

In this study, we evaluate the evidence in favor of excess stock return predictability from the perspective of a Bayesian investor. We focus on the case of a single predictor variable to highlight the complex statistical issues that come into play in this deceptively simple problem.

The investor in our model considers the evidence in favor of the following linear model for excess returns:

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

where r_{t+1} denotes the return on a broad stock index in excess of the riskfree rate, x_t denotes a predictor variable, and u_{t+1} the unpredictable component of the return. The investor also places a finite probability on the following model:

$$r_{t+1} = \alpha + u_{t+1}. \tag{2}$$

Namely, the investor assigns a prior probability q to the state of the world in which returns are predictable (because the prior on β will be smooth, the chance of $\beta = 0$ in (1) is infinitesimal), and a probability $1 - q$ to the state of the world in which returns are completely unpredictable. In both cases, the parameters are unknown. Thus our model allows for both parameter uncertainty and "model uncertainty."[2]

☆ This paper previously circulated under the title "What is the chance that the equity premium varies over time? Evidence from predictive regressions." We are grateful to Sean Campbell, Mark Fisher, Michael Johannes, Matthew Pritsker, Robert Stambaugh, Stijn van Nieuwerburgh, Jonathan Wright, Moto Yogo, Hao Zhou and seminar participants at the 2008 meetings of the American Finance Association, the 2007 CIRANO Financial Econometrics Conference, the 2007 Winter Meeting of the Econometric Society, the 2010 Federal Reserve Conference on Financial Markets, the Federal Reserve Board, the 2013 NBER NSF time series conference, the University of California at Berkeley and the Wharton School for helpful comments. We are grateful for financial support from the Aronson+Johnson+Ortiz fellowship through the Rodney L. White Center for Financial Research. This manuscript does not reflect the views of the Board of Governors of the Federal Reserve System or its staff.

* Corresponding author at: Department of Finance, The Wharton School, University of Pennsylvania, 2300 SH-DH, Philadelphia, PA, 19104, United States. Tel.: +1 215 898 7634.

E-mail addresses: jwachter@wharton.upenn.edu (J.A. Wachter), missaka.n.warusawitharana@frb.gov (M. Warusawitharana).

1 Tel.: +1 202 452 3461.

2 However, note that our investor is Bayesian, rather than ambiguity averse (e.g. Chen and Epstein, 2002). Our priors are equivalent to placing a point mass on $\beta = 0$ in (1).

http://dx.doi.org/10.1016/j.jeconom.2014.05.018
0304-4076/

Allowing for a non-zero probability on (2) is one way in which we depart from previous studies. Previous Bayesian studies of return predictability allow for uncertainty in the parameters in (1), but assume uninformative priors (Barberis, 2000; Brandt et al., 2005; Johannes et al., 2002; Skoulakis, 2007; Stambaugh, 1999). As Wachter (2010) shows, flat or nearly-flat priors imply a degree of predictability that is hard to justify economically. Other studies (Kandel and Stambaugh, 1996; Pastor and Stambaugh, 2009; Shanken and Tamayo, 2012; Wachter and Warusawitharana, 2009) investigate the impact of economically informed prior beliefs. These studies nonetheless assume that the investor places a probability of one on the predictability of returns. However, an investor who thinks that (2) represents a compelling null hypothesis will have a prior that places some weight on the possibility that returns are not predictable at all.

Our work also relates to the Bayesian model selection methods of Avramov (2002) and Cremers (2002). In these studies, the investor has a prior probability over the full set of possible linear models that make use of a given set of predictor variables. Thus the setting of these papers is more complex than ours in that many predictor variables are considered. However, these papers also make the assumption that the predictor variables are either non-stochastic, or that their shocks are uncorrelated with shocks to returns. These assumptions are frequently satisfied in a standard ordinary least squares regression, but rarely satisfied in a predictive regression. In contrast, we are able to formulate and solve the Bayesian investor's problem when the regressor is stochastic and correlated with returns.

When we apply our methods to the dividend-price ratio, we find that an investor who believes that there is a 50% probability of predictability prior to seeing the data updates to a 86% posterior probability after viewing quarterly postwar data. We find average certainty equivalent returns of 1% per year for an investor whose prior probability in favor of predictability is just 20%. For an investor who believes that there is a 50/50 chance of return predictability, certainty equivalent returns are 1.72%.

We also empirically evaluate the effect of correctly incorporating the initial observation of the dividend-price ratio into the likelihood (the exact likelihood approach) versus the more common conditional likelihood approach. In the conditional likelihood approach, the initial observation of the predictor variable is treated as a known parameter rather than as a draw from the data generating process. We find that the unconditional risk premium is poorly estimated when we condition on the first observation. However, when this is treated as a draw from the data generating process, the expected return is estimated reliably. Surprisingly, the posterior mean of the unconditional risk premium is notably lower than the sample average.

Finally, when we examine the evolution of posterior beliefs over the postwar period, we find substantial differences between the beliefs implied by our approach, which treats the regressor as stochastic and realistically captures the relation between the regressor and returns, and beliefs implied by assuming non-stochastic regressors. In particular, our approach implies that the belief in the predictability of returns rises dramatically over the 2000–2005 period while approaches assuming fixed regressors imply a decline. We also evaluate out-of-sample performance over the postwar period, and show that our method leads to superior performance both when compared with a strategy based on sample averages, and when compared with a strategy implied by OLS regression.

The remainder of the paper is organized as follows. Section 2 describes our statistical method and contrasts it with alternative approaches. Section 3 describes our empirical results. Section 4 concludes.

2. Statistical method

2.1. Data generating processes

Let r_{t+1} denote continuously compounded excess returns on a stock index from time t to $t+1$ and x_t the value of a (scalar) predictor variable. We assume that this predictor variable follows the process

$$x_{t+1} = \theta + \rho x_t + v_{t+1}. \tag{3}$$

Stock returns can be predictable, in which case they follow the process (1), or unpredictable, in which case they follow the process (2).[3] In either case, errors are serially uncorrelated, homoskedastic, and jointly normal:

$$\begin{bmatrix} u_{t+1} \\ v_{t+1} \end{bmatrix} \mid r_t, \ldots, r_1, x_t, \ldots, x_0 \sim N(0, \Sigma), \tag{4}$$

and

$$\Sigma = \begin{bmatrix} \sigma_u^2 & \sigma_{uv} \\ \sigma_{uv} & \sigma_v^2 \end{bmatrix}. \tag{5}$$

As we show below, the correlation between innovations to returns and innovations to the predictor variable implies that (3) affects inference about returns, even when there is no predictability.

When the process (3) is stationary, i.e. ρ is between -1 and 1, the predictor variable has an unconditional mean of

$$\mu_x = \frac{\theta}{1-\rho} \tag{6}$$

and a variance of

$$\sigma_x^2 = \frac{\sigma_v^2}{1-\rho^2}. \tag{7}$$

These follow from taking unconditional means and variances on either side of (3). Note that these are population values conditional on knowing the parameters. Given these, the population R^2 is defined as

$$\text{Population } R^2 = \frac{\beta^2\sigma_x^2}{\beta^2\sigma_x^2 + \sigma_u^2}.$$

2.2. Prior beliefs

The investor faces uncertainty both about the model (i.e. whether returns are predictable or not), and about the parameters of the model. We represent this uncertainty through a hierarchical prior. There is a probability q that investors face the distribution given by (1), (3) and (4). We denote this state of the world H_1. There is a probability $1-q$ that investors face the distribution given by (2)–(4). We denote this state of the world H_0. As we will show, the stochastic properties of x have relevance in both cases.

The prior information on the parameters is conditional on H_i. Let

$$b_0 = [\alpha,\ \theta,\ \rho]^\top$$

[3] The model we adopt for stock return predictability is assumed by Kandel and Stambaugh (1996), Campbell and Viceira (1999), Stambaugh (1999), Barberis (2000) and many subsequent studies. The idea that the price-dividend ratio can predict returns is motivated by present-value models of prices (see Campbell and Shiller, 1988). We have examined the possibility of adding lagged returns on the right hand side of both the return and the predictor variable regression; however the coefficients are insignificant.