ARTICLE IN PRESS

Journal of Financial Economics ∎ (∎∎∎∎) ∎∎∎–∎∎∎

Contents lists available at ScienceDirect

Journal of Financial Economics

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

Forecasting stock returns under economic constraints[☆]

Davide Pettenuzzo[a,1], Allan Timmermann[b,c,d,2], Rossen Valkanov[b,*]

[a] Brandeis University, Sachar International Center, 415 South St, Waltham, MA, United States
[b] University of California, San Diego, 9500 Gilman Drive, MC 0553, La Jolla, CA 92093, United States
[c] CEPR, United Kingdom
[d] CREATES, Denmark

ARTICLE INFO

Article history:
Received 9 May 2013
Received in revised form
24 October 2013
Accepted 21 November 2013

JEL classification:
C11
C22
G11
G12

Keywords:
Economic constraints
Sharpe ratio
Equity premium predictions

ABSTRACT

We propose a new approach to imposing economic constraints on time series forecasts of the equity premium. Economic constraints are used to modify the posterior distribution of the parameters of the predictive return regression in a way that better allows the model to learn from the data. We consider two types of constraints: non-negative equity premia and bounds on the conditional Sharpe ratio, the latter of which incorporates time-varying volatility in the predictive regression framework. Empirically, we find that economic constraints systematically reduce uncertainty about model parameters, reduce the risk of selecting a poor forecasting model, and improve both statistical and economic measures of out-of-sample forecast performance.

1. Introduction

Equity premium (EP) forecasts play a central role in areas as diverse as asset pricing, portfolio allocation, and performance evaluation of investment managers.[3] However, more than 25 years of research shows that models allowing for time-varying return predictability often produce worse out-of-sample forecasts than a simple benchmark that assumes a constant risk premium. This finding has led authors such as Bossaerts and Hillion (1999) and Welch and Goyal (2008) to question the economic value of

[☆] We thank an anonymous referee for many constructive comments. We also thank John Campbell, Wayne Ferson, Blake LeBaron, Lubos Pastor, Seth Pruitt, Guofu Zhou, seminar participants at the 2013 Econometric Society Summer Institute and the 2013 European Seminar on Bayesian Econometrics, as well as Anthony Lynch, our discussant at the 2013 Western Finance Association meetings in Lake Tahoe, for helpful comments and suggestions. Xia Meng provided excellent research assistance. We are grateful to John Campbell and David Rapach for making their codes available.

[*] Corresponding author. Tel.: +1 858 534 0898.
E-mail addresses: dpettenu@brandeis.edu (D. Pettenuzzo), atimmerm@ucsd.edu (A. Timmermann), rvalkanov@ucsd.edu (R. Valkanov).
[1] Tel.: +1 781 736 2834.
[2] Tel.: +1 858 534 0894.

[3] Papers on time series predictability of stock returns include Campbell (1987), Campbell and Shiller (1988), Fama and French (1988, 1989), Ferson and Harvey (1991), Keim and Stambaugh (1986) and Pesaran and Timmermann (1995). Examples of asset allocation studies under return predictability include Aït-Sahalia and Brandt (2001), Barberis (2000), Brennan, Schwartz, and Lagnado (1997), Campbell and Viceira (1999), Kandel and Stambaugh (1996), and Xia (2001). Avramov and Wermers (2006) and Ferson and Schadt (1996) consider mutual fund performance under time-varying investment opportunities.

http://dx.doi.org/10.1016/j.jfineco.2014.07.015
0304-405X/© 2014 Elsevier B.V. All rights reserved.

ex ante return forecasts that allow for time-varying expected returns.

Economically motivated constraints offer the potential to sharpen forecasts, particularly when the data are noisy and parameter uncertainty is a concern as in return prediction models. While economic constraints have previously been found to improve forecasts of asset returns, no broad consensus exists on how to impose such constraints. For example, Ang and Piazzesi (2003) impose no-arbitrage restrictions to identify the parameters in a term structure model, Campbell and Thompson (2008) truncate their equity premium forecasts at zero and also constrain the sign of the slope coefficients in return prediction models, and Pastor and Stambaugh (2009, 2012) use informative priors to ensure that the sign of the correlation between shocks to unexpected and expected returns is negative.

This paper proposes a new approach for incorporating economic information via inequality constraints on moments of the predictive distribution of the equity premium. We focus on two types of economic constraints. The first, the equity premium constraint, follows the idea of Campbell and Thompson (2008) and requires the conditional mean of the equity premium to be non-negative.[4] It is difficult to imagine an equilibrium setting in which risk averse investors would hold stocks if their expected compensations were negative, and so this seems like a mild restriction. The second stipulates that the conditional Sharpe ratio (SR) has to lie between zero and a predetermined upper bound. The zero lower bound is identical to the equity premium constraint, and the upper bound rules out that the price of risk becomes too high. The Sharpe ratio of the market portfolio is extensively used in finance and, much like the equity premium, academics and investors can be expected to have strong priors about its magnitude.[5] Yet, SR constraints cast as inequality constraints on the predictive moments of the return distribution have not, to our knowledge, previously been explicitly explored in the return predictability literature.[6]

Other studies consider bounds on the maximum Sharpe ratio in the context of cross-sectional pricing models, which is different from our focus here. MacKinlay (1995) introduces a bound on the maximum squared Sharpe ratio as a way to distinguish between risk and nonrisk explanations of deviations from the Capital Asset Pricing Model (CAPM). MacKinlay and Pastor (2000) provide estimates of factor pricing models that condition on a given value of the Sharpe ratio. In a Bayesian setting, this corresponds to investors having different degrees of confidence in the asset pricing model, with a very large Sharpe ratio corresponding to completely skeptical beliefs about the model.

To incorporate economic information, we develop a Bayesian approach that lets us compute the predictive density of the equity premium subject to economic constraints. Importantly, the approach makes efficient use of the entire sequence of observations in computing the predictive density and also accounts for parameter uncertainty. Our approach builds on the conventional linear prediction model and simplifies to this model if the economic constraints are not binding in a particular sample.

The predictive moments of the return distribution get updated as new data arrive and so the inequality constraints give rise to dynamic learning effects. To see how this works, suppose a new observation arrives that, under the previous parameter estimates, imply a negative conditional equity premium. Because this is ruled out, the economic constraints can force the posterior distribution of the parameter estimates to shift significantly, even in situations in which the estimates of the standard linear model do not change at all. This effect turns out to be empirically important, particularly for large values of the predictor variables. Our empirical analysis finds that the posterior variance of the equity premium distribution—one measure of parameter estimation uncertainty—can be several times bigger for the unconstrained model compared with the constrained models, when evaluated at large values of the predictor variables.

Our approach toward incorporating economic constraints works very differently from that taken by previous studies such as Campbell and Thompson (2008). To highlight these differences, consider the constraint that the equity premium is non-negative. Campbell and Thompson (2008) impose this restriction by truncating the predicted equity premium at zero if the predicted value is negative. While this truncation approach can be viewed as a first approximation toward imposing moment or parameter constraints, it does not make efficient use of the information in the theoretical constraints. In particular, this approach never learns from the information that comes from observing that the estimated model implies negative forecasts of the equity premium and so the underlying model continues to repeat the same mistakes when faced with new data similar to previously observed data. In contrast, our approach constrains the equity premium forecast to be non-negative at each given time. This implies that we have T constraints in a sample of T observations, not just a single constraint. Every time a new pair of observations on the predictor variable and returns becomes available, the non-negativity constraint on the conditional equity premium is used to rule out values of the parameters that are infeasible given the constraint and, hence, to inform the parameter estimates.

In addition to the conditional EP constraint, we explore whether imposing a lower and an upper bound on the Sharpe ratio of the market portfolio provides further

[4] Boudoukh, Richardson, and Smith (1993) develop tests for the restriction that the conditional equity risk premium is non-negative. They find that this restriction is violated empirically for the US stock market.

[5] See Lettau and Wachter (2007, 2011) for recent examples of theoretical asset pricing models that rely on calibrations using the Sharpe ratio. For good treatments of the Sharpe ratio and its theoretical and empirical links to asset pricing models, see Cochrane (2001) and Lettau and Ludvigson (2010).

[6] Ross (2005) and Zhou (2010) consider constraints on the R^2 of the predictive return distribution. In practice, a close relation exists between constraints on the Sharpe ratio and constraints on the R^2. See, e.g., Campbell and Thompson (2008) for investors with mean variance utility. Wachter and Warusawitharana (2009) also consider priors on the slope coefficient in the return equation, which translate into priors about the predictive R^2 of the return equation. Shanken and Tamayo (2012) study return predictability by allowing for time-varying risk and specify a prior on the Sharpe ratio.

Please cite this article as: Pettenuzzo, D., et al., Forecasting stock returns under economic constraints. Journal of Financial Economics (2014), http://dx.doi.org/10.1016/j.jfineco.2014.07.015