Journal of Econometrics 192 (2016) 168–189

Contents lists available at ScienceDirect

Journal of Econometrics

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

A reexamination of stock return predictability[☆]

Yongok Choi [a,1], Stefan Jacewitz [b,2], Joon Y. Park [c,d,*]

[a] Department of Financial Policy, Korea Development Institute, Republic of Korea
[b] Center for Financial Research, Federal Deposit Insurance Corporation, United States
[c] Department of Economics, Indiana University, United States
[d] Department of Economics, Sungkyunkwan University, Republic of Korea

ARTICLE INFO

Article history:
Received 15 January 2014
Received in revised form
24 October 2014
Accepted 9 February 2015
Available online 2 January 2016

JEL classification:
C13
C22

Keywords:
Predictive regression
Time change
Cauchy estimator
Nonstationarity
Stochastic volatility

ABSTRACT

We provide a simple and innovative approach to test for predictability in stock returns. Our approach consists of two methodologies, time change and instrumental variable estimation, which are employed respectively to deal effectively with persistent stochastic volatility in stock returns and endogenous nonstationarity in their predictors. These are prominent characteristics of the data used in predictive regressions, which are known to have a substantial impact on the test of predictability, if not properly taken care of. Our test finds no evidence supporting stock return predictability, at least if we use the common predictive ratios such as dividend–price and earnings–price ratios.

1. Introduction

"STOCK RETURNS ARE PREDICTABLE". So states one of the new facts in finance from Cochrane (1999, 2005). In some ways, it is so widely accepted that stock returns are predictable that it has become a stylized fact. Lettau and Ludvigson (2001) reiterate the broad acceptance that excess returns are predictable by variables such as dividend–price ratios and earnings–price ratios. This conclusion is hardly restricted to academia. Wilcox (2007) states that the clear consensus within the investment industry is that prediction based on these ratios is highly useful. According to Ferson et al. (2003), predictive regressions are used in tactical asset allocation, active portfolio management, conditional performance evaluation, and market timing, among others. As a byproduct of its firm establishment as a stylized fact, there have been many economic models that can support some degree of return predictability in a general equilibrium setting. Theoretical devices used to do so include consumption smoothing in Balvers et al. (1990), habit formation in Campbell and Cochrane (1999), heterogeneous preferences in Chan and Kogan (2002), and time-varying risk preferences in Menzly et al. (2004). Avramov (2004) models the effects of prior beliefs about the extent of predictability explained by asset pricing models.

Despite its broad acceptance as fact, the question of stock return predictability has not been completely settled. It is still considered obvious by many that stock returns simply cannot be predictable. The empirical evidence of the actual out-of-sample predictability of stock returns is indeed mixed and inconclusive. Bossaerts and Hillion (1999) find that even the best prediction models have no out-of-sample predictive power. Welch and Goyal (2008) find

☆ We are grateful to a co-editor, an associate editor and two anonymous referees for useful comments. We would also like to thank Yoosoon Chang, Michael Chung, Hagen Kim, Peter Phillips, Dacheng Xiu and Jun Yu for numerous helpful discussions and suggestions and Willa Chen for making her code available. This paper benefited greatly from discussions at the Center for Financial Research at the FDIC, at the Mid-Atlantic Research Conference in Finance at Villanova University, and at the 2012 International Symposium on Econometric Theory and Applications in China at Shanghai Jiao Tong University. This paper is based in part on dissertations submitted to Indiana University by Yongok Choi and to Texas A&M University by Stefan Jacewitz.

* Corresponding author at: Department of Economics, Indiana University, United States.
E-mail address: joon@indiana.edu (J.Y. Park).

1 The views expressed herein are those of the author and should not be attributed to the Korea Development Institute.

2 The views and opinions expressed here are those of the author and do not necessarily reflect the views of the Federal Deposit Insurance Corporation.

http://dx.doi.org/10.1016/j.jeconom.2015.02.048
0304-4076/

that the standard predictive variables perform poorly in- and out-of-sample and are outperformed by something as simple as the historical average. In contrast, Avramov and Chordia (2006) find evidence of out-of-sample predictability of stock returns by the dividend yield, the term spread, the default spread, and the Treasury bill yield. Moreover, Avramov (2002) and Cremers (2002) each find in-sample and out-of-sample stock return predictability using Bayesian model averaging. Campbell and Thompson (2008) conclude that the out-of-sample predictive ability is small, but economically meaningful. See also Guo (2006, 2009). On the other hand, Ang and Bekaert (2007) find that long horizon predictability to be statistically insignificant and not robust across countries and sample periods. Though, at the same time, they still find that stock return predictability is real, albeit at shorter horizons.[3] See also, among others, Hjalmarsson (2011), Moon et al. (2004), and Valkanov (2003).

In the paper, we reexamine stock return predictability using a new set of econometric methodologies to deal with the problematic characteristics found in predictive regression data. From the econometric perspective, there are several critical issues in predictive regressions for stock returns. First, it is widely recognized that the covariates used to test return predictability have near unit roots and their innovations are strongly correlated with stock returns in the long-run, which causes standard hypothesis tests to substantially over-reject a true null, as shown by Stambaugh (1999). Second, the existence of nearly nonstationary stochastic volatility in returns has also been well documented in the literature. See, e.g., Schaller and Norden (1997) and Jacquier et al. (2004). Potentially, nonstationary stochastic volatility may yield substantial size distortions on standard tests relying on a constant unconditional variance, such as stationary ARCH or GARCH. In fact, Cavaliere (2004) has shown that the standard unit root tests are highly distorted in the presence of nonstationary and stochastic volatility. This immediately implies that predictability tests, as well, will be heavily distorted by nonstationary stochastic volatility in returns. Lastly, other data characteristics of returns and covariates such as the presence of deterministic trends, thick tails, jumps, and structural breaks may also have a serious deleterious effect on the performances of conventional predictability tests.

We provide a new technique for testing predictability that is uniquely suited to address all of the aforementioned econometric problems. We hope that this approach will be appealing not only for its effectiveness, but also for its simplicity. We combine a simple instrumental variable estimator with a simple time change. The estimator, which is called the Cauchy estimator, uses the sign transform of the covariate as an instrument. The instrument directly eliminates the problems caused by the persistent endogeneity and various other aberrant characteristics of the covariates. On the other hand, the time change makes standard hypothesis tests more robust to the presence of a wide variety of nonstationary and nearly nonstationary stochastic volatilities in stock returns, all of which invalidate the use of standard tests in predictive regressions. For the required time change, we use a volatility time in place of the calendar time.[4] Essentially, we wait for volatility to reach a certain threshold before collecting each observation, so that there is a constant level of volatility across all observations in the sample used to test for return predictability. Following the time change, the standard hypothesis tests in predictive regressions become valid regardless of nonstationarity and near nonstationarity in stochastic return volatilities.

The persistent endogeneity of covariates in predictive regressions has been studied by many other authors. In particular, Campbell and Yogo (2006) and Chen and Deo (2009) have recently proposed the tests of return predictability, which are designed specifically to perform well in the presence of persistent and endogenous covariates. However, their tests are not expected to properly deal with nonstationary and nearly nonstationary stochastic volatilities in stock returns and other plausible aberrant characteristics of the covariates. In fact, we find by simulation that they are subject to rather serious size distortions and over-reject the null of no predictability, if any of realistic nonstationary and nearly nonstationary stochastic volatilities considered in the paper are present. In contrast, our approach always yields almost exact sizes in a variety of plausible specifications of stock returns and predictive ratios. The robustness of our approach is quite evident and far reaching. Further, it is achieved without sacrificing powers, i.e., our testing procedure has discriminatory powers at least comparable to, and sometimes better than, the tests by Campbell and Yogo (2006) and Chen and Deo (2009), which are optimal for the prototypical Gaussian model.

If the econometric problems are properly addressed using our new approach combining a simple instrument and a volatility chronometer, we are left with virtually no evidence supporting stock return predictability using some of the most common predictors such as the dividend–price and earnings–price ratios. Our empirical results are strong and unambiguous: For all sampling periods and at all sampling frequencies we consider in the paper, our approach clearly fails to reject the null hypothesis of no predictability. However, it should be emphasized that we do not conclude that stock returns are all together unpredictable. They may well be predictable using alternative predictors, in place of or in addition to the common predictive ratios used here. Instead, we show that several features of the data can cause false predictors to appear valid and, once this is properly accounted for, no evidence of predictability remains. With this in mind, based on the empirical evidence, there seems little reason to believe that stock returns are predictable using only dividend–price or earnings–price ratios.

The remainder of the paper is organized as follows. Section 2 provides the background and summarizes the main issues relating to return predictability. Section 3 introduces a novel approach to effectively deal with the various problems affecting the conventional approach. We introduce a time change to volatility time in order to correct for time-varying stochastic volatilities nonparametrically, and the Cauchy t-ratio to deal with various statistical anomalies in predictive ratios, such as near nonstationarity, structural breaks, and jumps, among many others. In Section 4, we present a new test based on the Cauchy t-ratio for the samples collected after the time change, and show that the limit null distribution of the time-changed Cauchy t-ratio is standard normal under truly mild regularity conditions. Section 5 provides Monte Carlo evidence that compares the performances of traditional predictability tests with our new procedure. In particular, we demonstrate that our procedure in general performs well in terms of both size and power. Next, in Section 6 we apply the technique directly to actual stock return data. The empirical results are clear: We find no evidence for predictability in stock returns. Section 7 concludes the paper, and all the proofs are in the Mathematical Appendix.

[3] More recently, it has also been found that the return predictability becomes stronger if the variance risk premium is used as an additional covariate. See Bollerslev et al. (2009, 2011), Drechsler and Yaron (2011) and Bollerslev et al. (2014).

[4] The time change was used earlier by several authors including Yu and Phillips (2001), Peters and de Vilder (2006), Andersen et al. (2007) and Chang (2012). They all use the realized variance as an estimate of quadratic variation to obtain the required time change. However, we also consider the discrete sample based bi-power variation to allow for the presence of jumps. Moreover, all errors incurred by the use of discrete samples in estimating quadratic variation and bi-power variation are formally and rigorously analyzed in the paper.