Detrending Persistent Predictors

Christophe Boucher, Bertrand Maillet

Detrending Persistent Predictors

Christophe BOUCHER, Bertrand MAILLET

2011.19

"Detrending Persistent Predictors"[*]

Christophe M. Boucher[♦]   Bertrand B. Maillet[♦♦]

- March 2011 -

Researchers in finance very often rely on highly persistent – nearly integrated – explanatory variables to predict returns. This paper proposes to stand up to the usual problem of persistent regressor bias, by detrending the highly auto-correlated predictors. We find that the statistical evidence of out-of-sample predictability of stock returns is stronger, once predictors are adjusted for high persistence.

Keywords: Forecasting; Persistence; Detrending; Expected Returns.
JEL Classification: C1, C14, G1.

1. Introduction

This paper explores the forecasting power of detrended persistent predictors for stock returns in standard linear regressions. Rates of return are commonly regressed against the lagged values of some stochastic explanatory variables when forecasting financial returns. Extant research assumes the predictive regression's stochastic explanatory variable is stationary. However, confidence intervals computed for the largest autoregressive root of many explanatory variables

classically used for forecasting stock returns, including the dividend yield, the book-to-market ratio, the short term rate of interest and the term and default spreads, confirm uncertainty surrounding the order of integration of such variables (e.g. Torous et al., 2004).

Stambaugh (1999) point out that persistence leads to biased coefficients in predictive regressions if innovations in the predictor are correlated with returns. Under the same conditions, the standard t-test for predictability has an incorrect size (Cavanagh et al., 1995). These problems are exacerbated if researchers are data mining (Ferson et al., 2003).

An active recent literature discusses alternative econometric methods for correcting the bias and conducting a valid inference (e.g. Torous et al., 2004; Campbell and Yogo, 2006; Jansson and Moreira, 2006; Ang and Bekaert, 2007). Typically, authors conclude that the statistical evidence of forecastability is weaker once tests are adjusted for high persistence.

In this paper, we take an alternative route, by detrending the predictors and then removing their persistence. Contrary to the recent literature on inference in forecasting regressions with persistent regressors, we do not derive asymptotic distributions for OLS regressions under the assumption that the forecasting variable is a close-to unit, yet stationary, root process. Instead, we directly "cut out" the persistence of predictors. We find that the statistical evidence of out-of-sample predictability of stock returns is stronger once predictors are adjusted for high persistence.

The rest of the paper proceeds as follows. Section 2 exposes the empirical framework. Section 3 presents out-of-sample forecasting results. Section 4 concludes.

2. Empirical framework

The data set consists of quarterly observations from 1934:Q4 to 2010:Q2. Stock prices and dividends per share correspond to the Standard and Poor's Composite Index[1]. Real data are deflated by the Consumer Price Index (All Urban

---

[1] The S&P 500 data are available from Robert Shiller's home page at http://www.econ.yale.edu/~shiller and completed from the Standard and Poor's web site (S&P 500 Earnings and Estimate Report).

2

Consumers) published by the BLS. Let r_t denote the real return on the S&P index. The three month T-bill rate is used to construct the real return on the risk free rate, denoted $r_{f,t}$, and the log excess return $r_t - r_{f,t}$. We consider a set of autocorrelated predictive variables. We denote dy_t, the dividend yield, tb_t, the Treasury-bill rate and DEF_t, the default spread defined as the difference between the BAA and AAA corporate bond yields[2].

Figure 1 illustrates our approach based on a basic detrending method such as a backward moving average filter. The black line represents the one-step ahead forecasting ability of the detrended Treasury-bill (measured by the corrected R-square) on quarterly stock returns depending on the smoothing parameter for the low-frequency component. We first observe that the predictive power increases with the filtering parameter and reaches a peak when the low frequency components corresponding to a six-quarter moving average are removed. Secondly, after this peak, the predictive power of the Treasury-bill decreases and converges to the level of the corrected R-square obtained with the original "aggregate" Treasury-bill (the grey line).

Filtering and detrending techniques are numerous and applied in a wide range of economic and financial problems. Amongst the various available methods, we consider hereafter a simple backward moving-average filter as well as two different detrending methods applied by economists for measuring the business cycle: a Hodrick-Prescott filter and a wavelet multiscaling filter. On the one hand, the Hodrick-Prescott method is probably the most popular filter indentifying permanent and cyclical components of time series and will serve as a reference in our applications. However, the predictive power of the Hodrick-Prescott detrended variable could suffer from the forward property of this filter in the out-of-sample forecasting exercises presented below. On the other hand, wavelet filtering is an alternative to bandpass filtering and offers better resolution in the time domain (Wang, 1995). Indeed, wavelet basis functions are time-localized (in addition to being scale-localized)[3], which is useful for capturing the date and the magnitude of persistent changes in predictors.

---

[2] These interest rates are available from the FRED II database of the Federal Reserve Bank of St. Louis.

[3] The wavelet analysis consists of a decomposition of a signal into its set of basis functions (wavelets), analogous to the use of sines and cosines in the Fourier analysis. These basis functions are obtained from dilations or contractions (scaling), and translations of a mother wavelet.

3

Since in-sample forecasting results might suffer from a "look-ahead" bias that arises from decomposition using the full sample, we herein only examine the out-of-sample predictability of stock returns.

3. Out-of-sample forecasting results

We use four complementary statistics to compare the out-of-sample performances of our forecasting models: the mean-squared forecasting error (MSE) ratio, the Clark and McCracken's (2001) encompassing test (ENC-NEW), the McCracken's (2007) equal forecast accuracy test (MSE-F) and the modified Diebold-Mariano (MDM) encompassing test proposed by Harvey et al. (1998). We apply the ENC-NEW and MSE-F tests for the nested comparisons and the MDM test for the non-nested comparisons.

The initial estimation period begins with the last quarter of 1934 and ends with the first quarter of 1984 (corresponding to 2/3 of the sample). The detrended predictors and the predictive models are recursively re-estimated until the second quarter of 2010. The smoothing parameters for the three filtering methods used in the out-of-sample forecasts are those corresponding to the higher predictive power, as measured by the corrected R-square, over the initial in-sample period[4].

Table 1 reports results of the out-of-sample one-quarter-ahead nested forecast comparisons of excess returns. We consider a restricted (benchmark) model that includes both a constant and the lagged dependent variables as predictive variables. The nested comparisons are made by enlarging the information set with the one-period lagged values of the predictors as well as the detrended predictors. The three Panels refer to the predictor variable: respectively, the dividend yield, the Treasury-bill and the default spread.

We find that the unrestricted models (which include the predictor itself or the detrending predictor) have smaller MSE than the autoregressive restricted model, except with the Hodrick-Prescott filter. Moreover, MSE ratios show that forecasting power of the detrended predictor is higher for the wavelet filtering than for the backward moving-average method. Both ENC-NEW and MSE-F tests reject the null

---

[4] The persistency of the predictors, as measured by the autocorrelation coefficient, decreases by about 40% after detrending for the various predictors considered.

hypothesis that the detrended predictors with the wavelet or the moving-average filters provide no information about future stock returns at the 5% or 10% significance levels. We also examine whether these differences in forecasting performance between the filtering methods can be statistically discerned using the MDM test (Table 2). The wavelet based forecasting models contain information that produces superior forecasts (lower MSE) to those provided by the competitor models. However, the differences of predictive power between the filtering methods do not appear statistically significant at a 5% significance level.

Finally, detrending predictors based on a simple backward moving average method and the wavelet filtering method deliver superior out-of-sample forecasts relative to a simple random walk model, while "aggregate" (non-detrended) predictor do not.

4. Conclusion

This paper proposes to stand up to the usual problem of persistent regressor bias, by detrending the highly auto-correlated predictors. We find that the statistical evidence of out-of-sample forecastability is stronger once predictors are adjusted for high persistence.

References

Ang, A., Bekaert, G., 2007. Stock return predictability: is it there? Rev. Finan. Stud. 20(3), 651.707.

Campbell, J.Y., Yogo, M., 2006. Efficient tests of stock return predictability, J. Finan. Econ. 81, 27-60.

Cavanagh, C., Graham E., Stock, J., 1995. Inference in models with nearly integrated regressors, Econometric Theory 11, 1131-1147.

Clark, T.E., McCracken, M.W., 2001. Tests of equal forecast accuracy and encompassing for nested models. J. Econometrics 105, 85-110.

Ferson, W.E., Sarkissian, S., Simin, T., 2003. Spurious regressions in financial economics? J. Finance 58, 1393-1413.

Harvey, D., Leybourne, S., Newbold, P., 1997. Testing the equality of prediction mean squared errors. Int. J. Forecasting 13, 281-291.

Jansson, M., Moreira, M.J., 2006. Optimal inference in regression models with nearly integrated regressors, Econometrica 74, 681-715.

McCracken, M.W., 2007. Asymptotics for out-of-sample tests of Granger causality. J. Econometrics 140(2), 719-752.

5

Nelson, C.R., Kim M.J. 1993. Predictable stock returns: the role of small sample bias, J. Finance 48, 641-661.

Stambaugh, R.F., 1999. Predictive regressions. J. Finance Econ. 54(3), 375-421.

Torous W., Valkanov, R., Yan, S., 2004. On predicting stock returns with nearly integrated explanatory variables. J. Bus. 77(4), 937-966.

Wang, Y., 1995. Jump and sharp cusp detection by wavelets. Biometrika 82(2), 385-397.

6

Figure 1
Predictive Power and Backward Detrending

The Figure illustrates the predictive power of the Treasury-bill on US excess returns depending on the low-frequency removing parameter (number of quarters). A backward moving average filter is applied on the Treasury-bill. The black line represents the one-step ahead forecasting ability of the detrended Treasury-bill, measured by the corrected R-square, on quarterly excess returns over the period 1934:Q4-2010:Q2 depending on the filtering parameter. The grey line represents the corrected R-square of the predictive regression obtained with the original "aggregate" Treasury-bill.

7

Table 1

One-quarter-ahead Forecasts of Excess Returns: Nested Comparisons

#	Unrestricted model	MSE_u/MSE_r	ENC-NEW test	MSE-F test
Panel A : Dividend yield				
1	Aggregate	0,9948	1,6790	0,5074
2	Hodrick-Prescott	1,0086	5,5211	-0,8360
3	Moving-average	0,9653	3,0512**	3,5198*
4	Wavelet	0,9327	6,9503**	7,0745**
Panel B : Treasury-bill				
5	Aggregate	0,9892	2,7735	1,0689
6	Hodrick-Prescott	0,9667	3,5402*	3,3732*
7	Moving-average	0,9656	3,3630*	3,4907*
8	Wavelet	0,9646	3,0982*	3,5925*
Panel C : Default spread				
9	Aggregate	0,9685	2,9790	3,1831*
10	Hodrick-Prescott	1,0127	2,9544	-1,2315
11	Moving-average	0,9672	3,1108*	3,3198*
12	Wavelet	0,9661	3,2405*	3,4431*

We consider a restricted (benchmark) model of autoregressive returns (including a constant). MSE_u is the mean-squared forecasting error from the relevant unrestricted model in each row; MSE_r is the mean-squared error from the relevant restricted model. A sign * (**) denotes significance at a five (one) percent confidence level.

Table 2

One-quarter-ahead Forecasts of Excess Returns: Non-nested Comparisons

#	Model 1 *versus* Model 2	MSE_1/MSE_2	MDM Test	p-value
	Panel A : Dividend yield			
1	Wavelet *vs.* Hodrick-Prescott	0.9247	1.3090	0.1936
2	Wavelet *vs.* Moving-average	0.9662	0.8495	0.3977
	Panel B : Treasury-bill			
3	Wavelet *vs.* Hodrick-Prescott	0.9978	0.0850	0.9325
4	Wavelet *vs.* Moving-average	0.9990	0.0602	0.9522
	Panel C : Default spread			
5	Wavelet *vs.* Hodrick-Prescott	0.9539	1.4232	0.1579
6	Wavelet *vs.* Moving-average	0.9988	0.1041	0.9173

All of the models include a constant. The null hypothesis is that the Model 2 encompasses Model 1. The column labelled "*MSE1/MSE2*" reports the ratio of the mean-squared forecasting errors of Model 1 out of those of Model 2.

9