International Journal of Forecasting 32 (2016) 548–558

Contents lists available at ScienceDirect

International Journal of Forecasting

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

Betas and the myth of market neutrality

Nicolas Papageorgiou [a], Jonathan J. Reeves [b,*], Xuan Xie [b]

[a] Department of Finance, HEC Montreal, Canada

[b] School of Banking and Finance, University of New South Wales, Sydney, NSW 2052, Australia

ARTICLE INFO

Keywords:
Evaluating forecasts
Realized betas
Systematic risk
Time series
Zero-beta portfolios

ABSTRACT

Market neutral funds are commonly advertised as alternative investments that offer returns which are uncorrelated with the broad market. Utilizing recent advances in financial econometrics, we demonstrate that using standard forecasting methods to construct market (beta) neutral funds is often very inaccurate. Our findings demonstrate that the econometric methods that are commonly employed for forecasting the beta (systematic) risk typically lack sufficient accuracy to permit the successful construction of market neutral portfolios. The results in this paper also highlight the need for higher frequency returns data to be utilized more commonly. Using daily returns over the past year, we demonstrate an approach that is easy to implement and delivers a substantial improvement, relative to other methods, when attempting to construct a market neutral portfolio.

1. Introduction

Hedge funds are often portrayed as investments which are market (beta) neutral, in that they have little systematic exposure to market risk. However, work by Asness, Krail, and Liew (2001), Bali, Brown, and Caglayan (2011, 2012), and Patton (2009) has shown hedge funds to have substantial market exposure. The recent financial crisis provided further evidence of significant beta exposure by equity market neutral hedge fund managers, as over 70% of funds reporting to Hedge Fund Research (HFR) finished 2008 in the red. Brown, Gregoriou, and Pascalau (2012) further highlight these concerns over the risk characteristics of hedge funds, and in reference to hedge funds during the recent financial crisis, remark "we all fall down together".

The successful construction of a market neutral portfolio depends inherently on the ability of the manager to measure and forecast the beta exposure of his long and short portfolios accurately. The greater the forecast error of the betas, the more likely the fund is to have a significant residual beta exposure, and therefore, the greater the potential exposure to systematic risk factors. This, of course, is precisely what investors hope to avoid when investing in market neutral funds. They do not want to be paying alpha fees (averaging 1.5% in management fees and with 20% in incentive fees) for beta returns (which can be obtained easily through index exchange traded funds for a fee of typically less than 0.2%). In response to these concerns, this study investigates the accuracy of the methods commonly employed for constructing market neutral portfolios.

In general, daily equity returns is the highest frequency that is available reliably for the construction of equity market neutral portfolios, though it is also very common to see beta forecasts generated from monthly equity returns. The overwhelming majority of beta forecasts are generated from a constant beta model with a typical estimation period of between one and five years. This dates back to the work of Fama and MacBeth (1973), who proposed an estimation period of five years of monthly returns, and was further justified by Ghysels (1998), who showed that constant beta models have outperformed more sophisticated models of the time-varying beta. Recently

* Corresponding author. Tel.: +61 2 9385 5874.
E-mail address: reeves@unsw.edu.au (J.J. Reeves).

http://dx.doi.org/10.1016/j.ijforecast.2015.09.005

proposed time-varying realized beta models for quarterly beta forecasting were studied by Andersen, Bollerslev, Diebold, and Wu (2005, 2006), Ghysels and Jacquier (2006), and Hooper, Ng, and Reeves (2008), among others, though Reeves and Wu (2013) showed that these time-varying realized beta models did not outperform the constant beta model estimated on daily returns over the prior year.

In this paper, we evaluate these competing beta forecasting approaches in the setting of the construction of an equity market neutral portfolio. Equity momentum portfolios are constructed, as this is a common portfolio construction technique; see Carhart (1997), Grundy and Martin (2001), Jegadeesh and Titman (1993) and Novy-Marx (2012), among others. However, we are not evaluating the return generating abilities of momentum strategies, but focus instead on evaluating the beta neutrality of portfolios by applying a commonly used and widely studied trading strategy. In addition, as a robustness check for portfolio construction based on momentum, we also construct portfolios by selecting stocks randomly and assessing the bootstrap distribution of statistics. We obtain similar results for both portfolio construction approaches, which provides an indication of the general applicability of the results beyond momentum-based strategies.

The findings in this paper show that methods that are designed to deliver beta neutrality (in the sense of low variability from a zero beta) often fail, particularly when the volatility in stock market returns is high. We define high volatility as periods when the CBOE Volatility Index (VXO) is above its median. In addition, we also find that the ex post portfolio betas of the constructed equity market neutral portfolios are generally not correlated with other well known risk premiums, such as the size, value and momentum premiums.

The results in this paper suggest that the inability of equity market neutral funds to exhibit market neutrality in their performances may be due largely to the fact that the commonly utilized approaches lack sufficient accuracy to construct an equity market neutral fund. We also find that, with a portfolio that targets a beta of zero, based on the widely used (Fama & MacBeth, 1973) beta from five years of monthly returns, the ex post beta exceeded one. Even more alarming is the fact that these non-zero ex post betas are amplified by leverage. Heavy leverage from financial institutions to invest, mistakenly, in beta (believing it to be alpha) was a leading factor in the recent financial crisis, see Acharya and Richardson (2009) and Brown (2011). Of the different beta forecasting approaches, we find that the smallest errors occur when beta neutral portfolios are constructed from realized betas computed from daily returns over the past year.

The rest of the paper is organized as follows. Section 2 provides some background on the realized beta, and Section 3 describes the methodology. Sections 4 and 5 present the data and results respectively. Our conclusion is presented in Section 6.

2. Realized beta

In this paper, our forecasting approaches and evaluation rely on realized beta estimates, so we begin by briefly reviewing the measurement of the realized beta. Following the work of Andersen et al. (2006) and Barndorff-Nielsen and Shephard (2004), assume that p_t, which represents the logarithmic $N \times 1$ vector price process, follows a multivariate continuous time stochastic volatility diffusion

$$dp_t = \mu_t dt + \theta_t dW_t \tag{1}$$

where W_t is standard Brownian motion, $\omega_t = \theta_t \theta_t'$ is the instantaneous covariance matrix, and μ_t is the N-dimensional instantaneous drift and is jointly independent of $W(t)$. If h denotes a certain period (e.g., one day, one month, etc.), then the continuously compounded return over period h for stock i can be defined as $r_{i,t+h,h} = p_{i,t+h} - p_{i,t}$.

Under weak regularity conditions,[1] the theory of quadratic variation leads us to the following result for all t as the sampling frequency tends to infinity:

$$\sum_{j=1,\ldots,[h/\Delta]} r_{t+j\cdot\Delta} \cdot r'_{t+j\cdot\Delta} - \int_0^h \omega_{t+s} ds \xrightarrow{p} 0. \tag{2}$$

The realized beta of a security is defined as the realized covariance of the security with the market divided by the realized variance of the market. Following the above discussion, the realized covariance of security i and market M over the period $[t, t+h]$ from the theory of quadratic variation is

$$\widehat{\nu}_{iM,t,t+h} = \sum_{j=1,\ldots,[h/\Delta]} r_{i,t+j\cdot\Delta,\Delta} \cdot r_{M,t+j\cdot\Delta,\Delta}.$$

Similarly, the realized variance of the market over the period $[t, t+h]$ is

$$\widehat{\nu}_{M,t,t+h} = \sum_{j=1,\ldots,[h/\Delta]} r^2_{M,t+j\cdot\Delta,\Delta}.$$

Thus, the realized beta of security i is

$$\widehat{\beta}_{i,t,t+h} = \frac{\widehat{\nu}_{iM,t,t+h}}{\widehat{\nu}_{M,t,t+h}} = \frac{\sum_{j=1,\ldots,[h/\Delta]} r_{i,t+j\cdot\Delta,\Delta} \cdot r_{M,t+j\cdot\Delta,\Delta}}{\sum_{j=1,\ldots,[h/\Delta]} r^2_{M,t+j\cdot\Delta,\Delta}}$$
$$\xrightarrow{p} \frac{\int_0^h \omega_{(iN),t+s} ds}{\int_0^h \omega_{(NN),t+s} ds} = \beta_{i,t,t+h} \tag{3}$$

for all t as $\Delta \to 0$; i.e., the realized beta measure is consistent for the true beta by sampling both the security return and the market return at an ultra high frequency.

2.1. Measurement error

The prior consistency result of the realized beta estimator provides a theoretical justification in the setting of ultra high frequency return measurement; however, in most applications, careful consideration needs to be given to the question of how high the return frequency can be taken without losing accuracy in the return measurement. For most stocks, daily returns is often

[1] See Andersen et al. (2006) and Barndorff-Nielsen and Shephard (2004) for details.