Journal of Banking & Finance 36 (2012) 2351–2365

Contents lists available at SciVerse ScienceDirect

Journal of Banking & Finance

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

Forecasting the performance of hedge fund styles

José Olmo [a,b], Marcos Sanso-Navarro [c,*]

[a] Centro Universitario de la Defensa, Academia General Militar, Ctra. Huesca s/n, 50090 Zaragoza, Spain
[b] Department of Economics, City University London, Northampton Square, London EC1V 0HB, UK
[c] Departamento de Análisis Económico, Universidad de Zaragoza, Gran Vía 2, 50005 Zaragoza, Spain

ARTICLE INFO

Article history:
Received 3 October 2011
Accepted 20 April 2012
Available online 30 April 2012

JEL classification:
C10
C20
G10

Keywords:
Conditional density estimation
Hedge fund styles
Nonparametric methods
Portfolio performance
Stochastic dominance tests

ABSTRACT

This article predicts the relative performance of hedge fund investment styles using time-varying conditional stochastic dominance tests. These tests allow for the construction of dynamic trading strategies based on nonparametric density forecasts of hedge fund returns. During the recent financial turmoil, our tests predict a superior performance for the Global Macro investment style compared with the other strategies of 'Directional Traders'. The Dedicated Short Bias investment style is stochastically dominated by the other directional styles. These results are confirmed by simple nonparametric tests constructed from realized excess returns. Further, by utilizing a cross-validation method for optimal bandwidth parameter selection, we discover the factors that have predictive power regarding the density of hedge fund returns. We observe that different factors have forecasting power for different regions of the returns distribution and, more importantly, that the Fung and Hsieh factors have power not only for describing the risk premium but also, if appropriately exploited, for density forecasting.

1. Introduction

Hedge funds have attracted a great deal of attention during the last 15 years. These financial instruments are private investment vehicles for wealthy individuals and institutional investors that are less strictly regulated and supervised. These funds follow unconventional trading strategies and have traditionally outperformed other investment strategies, in part because of the weak correlation of their returns with the returns of other financial securities. This apparent fact has recently been disputed, however, as the 2007–2008 crisis has revealed the interdependencies among these funds and the rest of the financial industry.

Return predictability in the hedge fund industry has recently been investigated by Wegener et al. (2010), Avramov et al. (2011), Bali et al. (2011) and Vrontos (2012), among others. In particular, Wegener et al. (2010) account for non-normality, heteroskedasticity and time-varying risk exposures when predicting the conditional mean of the excess returns from four hedge fund strategies. To achieve the same goal, Bali et al. (2011) assess hedge fund exposure to various financial and macroeconomic risk factors. Avramov et al. (2011) find that macroeconomic variables, specifically the default spread and the Chicago Board Options Exchange volatility index (VIX), substantially improve the predictive ability of the benchmark linear pricing models used in the hedge fund industry. All of these seminal papers are concerned with forecasting the expected excess returns of hedge funds but devote little attention to the higher moments of the conditional distribution that are relevant for investment decisions. One exception to this trend is the investigation of Vrontos (2012), which specifies a multivariate GARCH model for the conditional distribution of hedge fund returns.

Efficient investment portfolios are typically the result of an optimization problem subject to certain constraints. Optimal portfolios are either on the mean-risk efficient frontier or defined by the combination of risky and riskless assets that maximize a certain expected utility function representing investors' preferences. Stochastic dominance tests constitute a powerful statistical method to compare the relative efficiency of investment portfolios. Fishburn (1977) demonstrates that portfolios that are mean-risk efficient are also stochastically efficient; thus, a portfolio that stochastically dominates another portfolio is also a better strategy in the mean-risk space. This author also shows that stochastic dominance implies an ordering of portfolios in terms of both risk-aversion levels and investors' expected utility maximization for general forms of the utility function. This methodology has recently been used for comparing investment portfolios. In a seminal paper,

* Corresponding author. Tel.: +34 876 554 629; fax: +34 976 791 996.
E-mail addresses: jolmo@unizar.es (J. Olmo), marcossn@unizar.es (M. Sanso-Navarro).

0378-4266/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jbankfin.2012.04.016

Linton et al. (2005) compare the performance of different worldwide financial indexes. In a similar context, Wong et al. (2008) propose this methodology as an appropriate technique for ranking the performance of Asian hedge funds. These authors also study traditional mean–variance and CAPM approaches for analyzing the performance of these investment instruments and conclude that the nonstandard empirical features of the returns on hedge funds, such as non-normality and option-like behavior, render these techniques inappropriate for assessing their relative investment performances.

Recent studies investigating investor behavior report evidence of the importance of investment styles. According to the style investing hypothesis (Barberis and Shleifer, 2003), investors categorize risky assets into styles and subsequently allocate money to those styles depending on their relative performance. Hedge funds, as with many other investment classes, are grouped into investment styles. Ter Horst and Salganik (2011) find that better-performing and more popular styles are rewarded with higher inflows in subsequent periods; thus, it is important to be able to predict the performance of hedge fund investment styles.

The objective of this paper is to predict the relative performance of hedge fund investment styles one time period in advance. We accomplish this goal through dynamic stochastic dominance tests that are conditional on a time-varying information set. To forecast the conditional density corresponding to each hedge fund investment style, we propose nonparametric kernel methods. The vector of optimal bandwidth parameters is obtained as the solution of the cross-validation method introduced by Hall et al. (2004). This method automatically discards factors with no predictive power to forecast the return on the hedge fund style and hence provides very valuable information on the relevant set of predictive factors.

Our empirical application focuses on hedge fund investment styles that bet on the movements of financial markets. These investment styles fall into the broader category of 'Directional Traders', as detailed by Agarwal et al. (2009). Our sample period runs between January 1994 and December 2009 and thus includes the recent global financial crisis in which hedge funds experienced greater exposure to the ups and downs of financial markets than market-neutral strategies. In particular, in this investigation, we study the Dedicated Short Bias (DSB) style, which exhibits exposure to short positions; the Emerging Markets (EMs) style, which focuses on investing in the securities of companies from emerging or developing countries; the Global Macro (GM) style, which bets on the direction of movement of currency exchange rates or interest rates; and the Managed Futures (MFs) style, which exploits short-term patterns in futures markets. The predictive performance of these styles is also compared to an asset-weighted portfolio comprising the whole hedge fund industry, which we call ALL. Our tests predict a superior performance of the GM investment style compared with the other styles under study. By contrast, the DSB investment style is stochastically dominated by the other directional styles. We also find that although the DSB, EM and MF styles neither dominate nor are dominated by ALL in the first-order analysis, indicating the relative efficiency of these strategies, in the second- and third-order analysis, we observe that ALL stochastically dominates these directional styles. This result can be interpreted as a preference of risk-averse investors for exposure to the whole hedge fund industry rather than merely a single directional style. In other words, this result suggests that under risk aversion, investors trade off expected returns for lower risk in the form of more highly diversified portfolios. This finding is reinforced by the third-order test of stochastic dominance, as this test demonstrates that ALL and GM are equally attractive for risk-averse investors with increasing levels of risk aversion. These results are confirmed by simple nonparametric proportion tests on the differences among the observed realized excess returns.

The present study relates to work by Meligkotsidou et al. (2009), who analyzes hedge fund investment styles using quantile regression methods, and Giannikis and Vrontos (2011), who address the nonlinear relationship between hedge fund returns and risk factors using Bayesian model selection techniques and threshold models. We also join Wong et al. (2008) in applying stochastic dominance techniques to study the performance of hedge fund portfolios. Other articles exploring stochastic dominance in related fields include work by Abhyankar et al. (2008) who compare value vs. growth strategies, and Fong et al. (2005) who use stochastic dominance tests to analyze the consistency of general asset-pricing models with the momentum effect.

This article is structured as follows. Section 2 presents the nonparametric techniques used to predict the conditional density of returns of the different hedge fund styles examined and introduces the relevant dynamic tests of stochastic dominance between investment portfolios. Section 3 discusses the data analyzed and the results from the empirical application to the 'Directional Traders' hedge fund styles. Section 4 concludes the manuscript.

2. Methodology

In this section, we first present the nonparametric kernel method used to construct the predictive conditional density function. We then discuss dynamic stochastic dominance tests of arbitrary order.

2.1. A nonparametric estimator for the predictive conditional density

Let $(Y_t)_{t\in Z}$ be a strictly stationary time series process defined on a compact set Ω, with an unconditional density function $f(y)$ and a cumulative distribution function (cdf) $F(y)$; moreover, let $f_{t-1}(y)$ and $F_{t-1}(y)$ be the corresponding predictive density and predictive distribution functions conditional on the sigma-algebra $\mathscr{I}_{t-1}$, which is defined by all the information available up to time $t-1$. Our interest is in forecasting these functions. To accomplish this task, we consider a k-vector of predictive factors, denoted as X_t, and a finite information set $I_t = \{(Y_s, X_s), t-m+1 \leqslant s \leqslant t\}$ defined on a compact set $\Omega' \in R^q$, with $q = (k+1)m$. Using this set, we construct the predictive density function $f_{I_{t-1}}(y)$ that approximates $f_{t-1}(y)$. For completeness, we also introduce the multivariate density function of I_t, denoted as $f^{I_t}(y)$, and its distributional counterpart, $F^{I_t}(y)$.

A natural nonparametric estimator of this conditional density for $I_{t-1} = x$, where x is a multivariate vector that represents a realization of the recent history of the information set, and n represents the number of available observations, is therefore provided by the following expression:

$$\hat{f}_x(y) = \frac{n^{-1}\sum_{t=1}^{n} h_{h_Y}(y) W_h(I_{t-1,x})}{\hat{f}^{I_1}(x)} \quad (1)$$

where $W_h(I_{t-1},x) = \Pi_{s=1}^{q} h_s^{-1} w\left(\frac{I_{t-1,s}-x_s}{h_s}\right)$ and $w(\cdot)$ and $k_{h_Y}(\cdot)$ are univariate kernel functions for the marginal random variables of the vectors I_{t-1} and Y_t, respectively. The corresponding bandwidth parameters are h_s, $1 \leqslant s \leqslant q$ and h_Y. The nonparametric estimator of $f^{I_1}(x)$ is thus $\hat{f}^{I_1}(x) = n^{-1}\sum_{t=1}^{n} W_h(I_{t-1},x)$, and $I_{t-1,s}$ and x_s denote the sth-component of the multivariate random vectors I_{t-1} and x, respectively. Li and Racine (2007) demonstrate the conditions for the uniform consistency of (1) to $f_x(y)$ for all $(x, y) \in \Omega$.

In both theoretical and practical settings, nonparametric kernel estimation has been established as being relatively insensitive to the choice of the kernel function. The same cannot be said for bandwidth selection, particularly in the context of our research, which features the search for an appropriate information set I_{t-1}