Forecasting the Performance of Hedge Fund Styles*

JOSE OLMO, *Centro Universitario de la Defensa and City University London*
MARCOS SANSO-NAVARRO†, *Universidad de Zaragoza*

March 2012

Abstract

This article predicts the relative performance of hedge fund investment styles one period ahead using time-varying conditional stochastic dominance tests. These tests allow 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 of the Global Macro investment style compared to the other 'Directional Traders' strategies. The Dedicated Short Bias investment style is, on the other hand, stochastically dominated by the other directional styles. These results are confirmed by simple nonparametric tests constructed from the realized excess returns. Further, by exploiting a cross-validation method for optimal bandwidth parameter selection, we find out which factors have predictive power for the density of hedge fund returns. We observe that different factors have forecasting power for different regions of the returns distribution and, more importantly, Fung and Hsieh factors have power not only for describing the risk premium but also, if appropriately exploited, for density forecasting.

Keywords: Conditional density estimation, hedge fund styles, nonparametric methods, portfolio performance, stochastic dominance tests.

JEL codes: C1, C2, G1.

*The authors acknowledge financial support from Ministerio de Ciencia e Innovación (ECO 2009-13675 and ECO2011-22650 projects) and the Regional Government of Aragón (ADETRE Research Group). We are also grateful to seminar participants at Liverpool University and Universidad de Zaragoza and to participants of Foro de Finanzas held in Granada, the Spanish Symposium of Economic Analysis in Málaga, and the CFE conference held in London for useful suggestions and comments. The authors have also benefited from the helpful and constructive comments of an anonymous referee.

†Corresponding author. Address: Departamento de Análisis Económico. Facultad de Economía y Empresa. Gran Vía 2. 50005 Zaragoza (Spain). Phone: +34 876 554 629. Fax: +34 976 761 996. e-mail: marcossn@unizar.es.

1 Introduction

Hedge funds have attracted a great deal of attention during the last fifteen years. These financial instruments are private investment vehicles for wealthy individuals and institutional investors that are less strictly regulated and supervised. Following unconventional trading strategies, these funds have traditionally outperformed other investment strategies partly due to the weak correlation of their returns with those of other financial securities. This stylized fact has recently been disputed: the 2007-08 crisis has revealed the interdependencies of these funds with the rest of the financial industry.

The sequence of papers by Fung and Hsieh (1997, 2001, 2002, 2004, 2011) showed that the risk premium from these funds can be largely explained by a set of financial variables rather different from the standard capital asset pricing formulations widely used in the mutual fund investment literature. These findings are crucial for constructing optimal portfolios. Agarwal and Naik (2004) study the relative performance between hedge funds and also against mutual funds. Related articles include Capocci and Hübner (2004) and Eling and Faust (2010). Patton (2009) also contributed to the study of these investment vehicles by questioning their market neutrality.

Financial return predictability has a long tradition in the empirical finance literature, see Keim and Stambaugh (1986). Return predictability in the hedge fund industry has been investigated by Amenc et al. (2003) and Hamza et al. (2006) and, more recently, by Wegener et al. (2010), Avramov et al. (2011), Bali et al. (2011) and Vrontos (2012), among others. In particular, Wegener et al. (2010) take non-normality, heteroskedasticity and time-varying risk exposures into account to predict the conditional mean of the excess returns on four hedge fund strategies. With the same aim, Bali et al. (2011) exploit 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 these seminal papers are concerned with forecasting the expected excess returns but hardly pay any attention to higher moments of the conditional distribution that are relevant for investment decisions. Interest in density forecasting has recently increased in the empirical finance literature, see Cenesiglou and Timmerman (2008) and Geweke and Amisano (2010), amongst others. In this line, Vrontos (2012) specifies a multivariate GARCH model for the conditional distribution of hedge fund returns.

Efficient investment portfolios are usually the result of an optimization problem subject to some constraints. Optimal portfolios are those that are on the mean-risk efficient frontier or are defined by the combination of risky and riskless assets that maximize a certain expected utility function representing investors' preferences. A powerful statistical method to compare the relative efficiency of investment portfolios is stochastic dominance tests. Fishburn (1977) shows that portfolios that are mean-risk efficient are also stochastically efficient and, hence, a portfolio that stochastically dominates another portfolio is also a better strategy in the mean-risk space. Similarly, this author shows that stochastic dominance implies an ordering of portfolios in terms of investors' expected utility maximization for general forms of the utility function and risk-aversion levels. This methodology has been recently used for comparing investment portfolios. In a seminal paper, 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, make these techniques inappropriate for assessing their relative investment performance.

Recent papers 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, like 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, so 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 one period ahead of hedge fund investment styles. We do this by means of dynamic stochastic dominance tests 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 financial markets movements. These investment styles fall into the broader category of 'Directional Traders', see Agarwal et al. (2009). Our sample period runs between 1994:01 and 2009:12, covering the recent global financial crisis in which these investment vehicles were more exposed to the ups and downs of financial markets than market-neutral strategies. In particular, we study the Dedicated Short Bias (DSB) style, that exhibits exposure to short positions, the Emerging Markets (EM) style, that focuses on investing in the securities of companies from emerging or developing countries, the Global Macro (GM) style, where bets are made on the direction of currency exchange rates or interest rates, and the Managed Futures (MF) style that 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, that we call ALL. Our tests predict a superior performance of the GM investment style compared to the other styles under study. The DSB investment style is, on the other hand, stochastically dominated by the other directional styles. We also find that, whereas the DSB, EM and MF styles do not dominate or are dominated by ALL in the first order, indicating the relative efficiency of these strategies, for the second and third order, we observe that ALL stochastically dominates these directional styles. This can be interpreted as a preference of risk-averse investors for exposure to the whole hedge fund industry over the directional styles. That is, 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 test of stochastic dominance of third order as it shows 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 difference of the observed realized excess returns.

The present study is also related to Li and Kazemi (2007), who estimated conditional density functions for hedge fund indices, and Meligkotsidou et al. (2009), who analyzed hedge fund investment styles using quantile regression methods. Our work is also connected with Billio et al. (2009), who studied hedge fund returns using nonparametric methods, and Giannikis and Vrontos (2011), who dealt with the non-linear relationship between hedge fund returns and risk factors using Bayesian model selection techniques and threshold models. Finally, we join Wong et al. (2008) and Li and Linton (2010) in applying stochastic dominance techniques to study the performance of hedge fund portfolios. Other articles exploring stochastic dominance in related fields[1] are Abhyankar et al. (2008) who compare value versus 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 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. Tables and Figures are collected in an appendix.

[1] Levy (2006) and Sriboonchitta et al. (2010) provide interesting monographs on stochastic dominance and its applications to finance and risk management.

2 Methodology

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

2.1 A Nonparametric Estimator for the Predictive Conditional Density

Let $(Y_t)_{t\in\mathbf{Z}}$ 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*) $F(y)$; let $f_{t-1}(y)$ and $F_{t-1}(y)$ be the corresponding predictive density and predictive distribution functions conditional on the sigma-algebra $\Im_{t-1}$ defined by all the information available up to time t. Our interest is in forecasting these functions. To do this, we consider a $k-$vector of predictive factors, denoted X_t, and a finite information set $I_t = \{(Y_s, X_s), t-m+1 \leq s \leq t\}$ defined on a compact set $\Omega' \in \mathbf{R}^q$, with $q = (k+1)m$. With 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 $f^{I_t}(y)$, and its distributional counterpart, $F^{I_t}(y)$.

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

$$\widehat{f}_x(y) = \frac{n^{-1}\sum_{t=1}^{n} k_{h_Y}(y) W_h(I_{t-1}, x)}{\widehat{f}^{I_1}(x)}, \tag{1}$$

where $W_h(I_{t-1}, x) = \prod_{s=1}^{q} h_s^{-1} w\left(\frac{I_{t-1,s}-x_s}{h_s}\right)$, 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, $1 \leq s \leq q$ and h_Y. The nonparametric estimator of $f^{I_1}(x)$ is $\widehat{f}^{I_1}(x) = n^{-1}\sum_{t=1}^{n} W_h(I_{t-1}, x)$; $I_{t-1,s}$ and x_s denote the $s^{th}-$component of the multivariate random vectors I_{t-1} and x, respectively. Li and Racine (2007) discuss 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 relatively insensitive to the choice of the kernel function. The same cannot be said for bandwidth selection, even more so in our setting given by the search for an appropriate information set I_{t-1} to approximate $f_{t-1}(\cdot)$. Following Hall et al. (2004), we propose a (least squares) cross-validation method to determine the optimal vector of bandwidth parameters. This method allows us to empirically determine I_{t-1}, that is, the vector of conditioning variables that best predicts the density $f_{t-1}(y)$. The cross-validation method automatically determines the irrelevant components of $\Im_{t-1}$ through assigning large smoothing parameters to them and, consequently, shrinking them toward the uniform distribution. The relevant components are precisely those that cross validation has chosen to smooth in the traditional way by assigning them bandwidth parameters of conventional size. A very nice review of the method and properties is given in Li and Racine (2007, Section 5.3).

The choice of the appropriate conditioning information set is very important so as to be able to optimally predict the density of returns and to implement the stochastic dominance tests. Note that one also needs to determine the forecasting scheme: fixed, rolling or recursive. To compare the predictive ability between density forecast competitors, we apply the test developed in Amisano and Giacomini (2007). This method assumes no knowledge of the true predictive density function and simply compares weighted versions of the predictive log-likelihood function of pairwise density forecast competitors over an out-of-sample period.

Let $\hat{f}_x(\cdot)$ and $\hat{g}_x(\cdot)$ be two competing forecasts of $f_{t-1}(y)$ at time $t-1$. The hypotheses of the relative predictive ability test are the following;

$$H_0 : E[WLR_{R,t+1}] = 0, \quad t = 1, 2, \ldots, T \text{ against,} \tag{2}$$

$$H_A : E[\overline{WLR}_{R,P}] \neq 0 \text{ for all } P \text{ sufficiently large,} \tag{3}$$

with $WLR_{R,t+1} = \omega(Y_{t+1}^{st})(log\hat{f}_x(Y_{t+1}) - log\,\hat{g}_x(Y_{t+1}))$ and $\overline{WLR}_{R,P} = P^{-1}\sum_{t=R}^{T-1} WLR_{R,t+1}$;

$Y_{t+1}^{st} = (Y_{t+1} - \hat{\mu}_{R,t})/\hat{\sigma}_{R,t}$ is the realization of the variable at time $t+1$, standardized using estimates of the unconditional mean and standard deviation of Y_{t+1}, $\hat{\mu}_{R,t}, \hat{\sigma}_{R,t}$, computed on the same sample on which the density forecasts are estimated. R corresponds to the in-sample period and $P = T - R$ to the out-of-sample period. The weight function $\omega(Y_{t+1}^{st})$ can be arbitrarily chosen by the forecaster to select the desired region of the distribution of Y_{t+1}. The only requirement imposed on the weight function is that it be positive and bounded. Amisano and Giacomini (2007) propose different alternatives for the centre and the tails of the distribution of the random variable, which will be used in our empirical application.

The relevant test statistic for testing H_0 is

$$t_{R,P} = \frac{\overline{WLR}_{R,P}}{\hat{\sigma}_P/\sqrt{P}}, \tag{4}$$

where $\hat{\sigma}_P^2$ is a heteroskedastic and autocorrelation consistent (HAC) estimator of the a-symptotic long-run variance $\sigma_P^2 = V(\sqrt{P}\ \overline{WLR}_{R,P})$. At a significance level α, this test rejects the null hypothesis of equal performance of forecasts whenever $|t_{R,P}| > z_{\alpha/2}$, where $z_{\alpha/2}$ is the $(1-\alpha/2)$ quantile of a standard normal distribution. In the case of rejection, we would choose $f_x(\cdot)$ if $t_{R,P}$ is positive and $g_x(\cdot)$, otherwise.

We should highlight that this method and more recent improvements (Diks et al., 2011; Gneiting and Ranjan, 2011) do not allow us to determine the correct predictive specification of the model but they do permit us to discriminate between potential forecasting methods. We will use these tests to choose between the fixed and rolling forecasting schemes in an out-of-sample evaluation. The definition of I_t precludes the recursive forecasting scheme in our predictive exercise. Amisano and Giacomini (2007) also discard this method when implementing their predictive ability test for similar reasons.

2.2 Dynamic Stochastic Dominance Tests

Stochastic dominance provides a powerful methodology to compare investment styles. First order stochastic dominance compares the distribution function of returns; the second order compares the expected value of the distributions, and so on. An interesting interpretation of these measures is in terms of expected utilities for different degrees of investors' risk aversion. First order stochastic dominance implies the superiority of an investment strategy for risk-neutral investors. Second order implies the superiority of a strategy for risk-averse investors, that is, investors with preferences that can be modeled by non-decreasing and concave real-valued utility functions. Similarly, third order stochastic dominance implies the superiority of one strategy over another for investors with increasing levels of risk aversion.[2] Seminal contributions to the topic are Stone (1973), Porter (1974) and Fishburn (1977).

From a methodological point of view, Davidson and Duclos (2000) is one of the first articles to introduce tests of stochastic dominance of different orders. These authors, however, do not check the dominance of one distribution function over another for every point of the domain of the corresponding random variables but only for a discrete set of points of these distributions. In this sense, the test may be inconsistent if the stochastic dominance condition is not satisfied for the points not considered in the analysis. On the other hand, the asymptotic theory of the test is standard. Barrett and Donald (2003) extend this test to the entire domain of continuous random variables and use simulation and bootstrap methods to approximate the asymptotic distribution of the test. These authors devised this test to compare income distributions, so it does not make allowance for dependence between the random variables but, even more importantly, it is not valid under serial dependence across time. Linton et al. (2005) incorporate the presence of serial dependence and cross-dependencies between the random variables. The asymptotic theory of the test proposed by these authors is very cumbersome and relies on subsampling techniques. This test

[2] The preferences of investors with increasing levels of risk aversion are characterized by a utility function, $u(x)$, with x denoting wealth, that is non-decreasing, concave and such that $-du/dx$ is concave. It can be shown that investors with a utility function of this type have a preference for a positive skewness of the distribution of wealth. In this line, Harvey and Sidique (2000) show the importance of incorporating skewness preferences into asset pricing models.

compares the existence of stochastic dominance between residuals of parametric time series regression models and, hence, it may be flawed if the parametric regression model proposed to describe the relationship between the response variable and the regressors is inadequate. Scaillet and Topaloglou (2010) also extend Barrett and Donald (2003) and accommodate the presence of serial dependence in unconditional stochastic dominance tests of arbitrary order. Other seminal articles in this literature are Klecan et al. (1991), Anderson (1996) and, more recently, Linton et al. (2010).

Following Linton et al. (2005), we focus on a dynamic setting, characterized by a time-varying information set, and propose a conditional stochastic dominance test that builds on the recent contribution of Gonzalo and Olmo (2011). Our approach is genuinely nonparametric and, therefore, is not affected by misspecification issues. It relies on the nonparametric forecasts of the density functions discussed above. Our testing framework allows for cross-dependence between the returns on the styles and, more importantly, for serial dependence over time. The asymptotic distribution of the tests can be approximated by bootstrap and simple simulation methods. In what follows, we describe our tests for conditional stochastic dominance of arbitrary orders in a dynamic setting.

Let A and B denote two investment portfolios; A stochastically dominates B for order γ conditional on the dynamic information set I_{t-1} if, and only if,

$$\Psi^{A}_{I_{t-1},\gamma}(y) \leq \Psi^{B}_{I_{t-1},\gamma}(y) \text{ for all } y \in \Omega \text{ and } t \in \mathbf{Z} \tag{5}$$

with $\Psi_{I_{t-1},\gamma}(y) = \int_{-\infty}^{y} \Psi_{I_{t-1},\gamma-1}(\tau)d\tau$ and where $\Psi_{I_{t-1},1}(y) = F_{I_{t-1}}(y)$. The integration by parts of these quantities yields the following characterization of the stochastic dominance condition:

$$\int_{-\infty}^{y} (y-\tau)^{\gamma-1} f^{A}_{I_{t-1}}(\tau)d\tau \leq \int_{-\infty}^{y} (y-\tau)^{\gamma-1} f^{B}_{I_{t-1}}(\tau)d\tau \text{ for all } y \in \Omega \subset \mathbf{R} \text{ and } t \in \mathbf{Z}. \tag{6}$$

We are interested in predicting the dynamics of the stochastic dominance relationship between investment styles, that is, our aim is to assess this condition for each period t.

This implies that the conditioning information set I_{t-1} for each t is simply a vector x describing the realization of the variable I_{t-1}. In this case, the characterization of stochastic dominance conditional on $I_{t-1} = x$ is $\Psi^A_{x,\gamma}(y) \leq \Psi^B_{x,\gamma}(y)$ for all $y \in \Omega$ and x fixed, with $\Psi_{x,\gamma}(y) = \int_{-\infty}^{y} \Psi_{x,\gamma-1}(\tau)d\tau$ and $\Psi_{x,1}(y) = P\{Y_t \leq y \mid I_{t-1} = x\}$. The relevant test for predictive stochastic dominance of arbitrary order $\gamma \geq 1$ at time t can be expressed, after some algebra, as the following composite hypothesis:

$$H_{0,\gamma} : E[d_{t,\gamma}(y) \mid I_{t-1} = x] \leq 0 \text{ for all } y \in \Omega \text{ and } x \text{ fixed}, \tag{7}$$

against

$$H_{1,\gamma} : E[d_{t,\gamma}(y) \mid I_{t-1} = x] > 0 \text{ for some } y \in \Omega, \tag{8}$$

with $d_{t,\gamma}(y) = (y - Y_t^A)^{\gamma-1} 1(Y_t^A \leq y) - (y - Y_t^B)^{\gamma-1} 1(Y_t^B \leq y)$.

We follow the extant literature on stochastic dominance tests and obtain the critical values of the test using the least favorable case[3], defined by the equality in (7) for all $y \in \Omega$; that we denote as $\widetilde{H}_{0,\gamma}$. To test for $H_{0,\gamma}$, we propose the supremum of the following process on $y \in \Omega$,

$$\widehat{D}_\gamma(y) = \frac{n^{-1} \sum_{t=1}^{n} d_{t,\gamma}(y) W_h(I_{t-1}, x)}{\widehat{f}^{I_1}(x)}. \tag{9}$$

For $\gamma = 1$, this expression is the difference between the nonparametric kernel estimators of the predictive distribution functions for the returns on portfolio A and B, see Li and Racine (2007, p. 182). For $\gamma > 1$, it compares higher moments of the conditional distribution of both portfolios.

A suitable test statistic for this test is $T_{n,\gamma} = (nh_1 \cdots h_q)^{1/2} \sup_{y \in \Omega} \widehat{D}_\gamma(y)$, that converges in distribution under $\widetilde{H}_{0,\gamma}$ to the supremum of a Gaussian process with zero mean and a covariance function that depends on the vector x. The proof of this result for a more general setting defined by x varying over the compact set Ω' can be found in Gonzalo and

[3]This technique is standard for composite null hypotheses, that is, those involving an infinite number of conditions, see Barrett and Donald (2003) and Gonzalo and Olmo (2011) in a stochastic dominance context, and Romano and Wolf (2011) for stochastic monotonicity testing.

Olmo (2011). It is well known in the stochastic dominance literature that the asymptotic distribution of these tests cannot be tabulated. Nevertheless, resampling and simulation methods can be implemented to approximate their asymptotic p-values. The following algorithm describes this procedure for $I_{t-1} = x$, with x fixed.

Algorithm:

1. Construct a grid of m points $y_1, ..., y_m$ contained in the compact space Ω and execute the following steps for $j = 1, \ldots, J$.

2. Generate $\{v_t\}_{t=1}^n$ independently and identically distributed (*iid*) $N(0,1)$ random variables.

3. Let $d_{t,\gamma}(y_i) = (y_i - y_t^A)^{\gamma-1} 1(y_t^A \leq y_i) - (y_i - y_t^B)^{\gamma-1} 1(y_t^B \leq y_i)$, with y_t^A, y_t^B being realizations of the random variables Y_t^A and Y_t^B.

4. Set $\widehat{D}_\gamma^*(y_i) = \frac{n^{-1} \sum_{t=1}^n d_{t,\gamma}(y_i) W_h(I_{t-1}, x) v_t}{\widehat{f}^I(x)}$, with x being a realization[4] of I_{t-1} and $h = (h_1, ..., h_q)$ obtained with a cross-validation criterion.

5. Compute $T_{n,\gamma}^{*(j)} = (nh_1 \ldots h_q)^{1/2} \sup_{y_i \in \Omega} \widehat{D}_\gamma^*(y_i)$ for all $y_i \in \Omega$.

This algorithm yields a random sample of J observations from the distribution of the test statistic $T_{n,\gamma}$. The simulated p-value of the stochastic dominance test for a given order γ is

$$\widehat{p}_{n,\gamma}^* = \frac{1}{J} \sum_{j=1}^J 1(T_{n,\gamma}^{*(j)} > T_{n,\gamma}) \tag{10}$$

which, under standard regularity conditions, see Hansen (1996) and Gonzalo and Olmo (2011), converges in probability to the true asymptotic p-value of $\widetilde{H}_{0,\gamma}$ as $J, n \to \infty$.

By repeating the test for each t we can establish a time-varying ranking of portfolios that allows us to construct dynamic trading strategies based on conditional stochastic efficiency.

[4] For simplicity in the exposition, we hereafter consider that $I_{t-1} = I_{t-1}^A \bigcup I_{t-1}^B$ refers to the set that collects the information contained in I_{t-1}^A and I_{t-1}^B, with each of the latter sets containing the information relevant for forecasting f_{t-1}^A and f_{t-1}^B, respectively.

This idea can be extended to analyze portfolio stochastic dominance/efficiency between more than two portfolios by using the test statistics proposed in Barrett and Donald (2003) and Linton et al. (2005). For illustrative purposes, we focus on pairwise comparisons in the next section.

3 Empirical Application

This section is divided into three blocks. The first part discusses the dataset given by the excess returns on the indices of the hedge fund styles above mentioned and the risk factors proposed by Fung and Hsieh (2001, 2004). The second block derives, both in a descriptive and a predictive setting, the optimal set of factors for each style. The final block of this section builds on this analysis to forecast the relative performance of hedge fund styles.

3.1 Data description

The hedge fund returns analyzed have been calculated from the Credit Suisse/Tremont asset-weighted indices expressed in US Dollars.[5] Data are monthly and span the period 1994:01-2009:12. The investment styles considered are Dedicated Short Bias (DSB), Emerging Markets (EM), Global Macro (GM) and Managed Futures (MF). For completeness, we also study an asset-weighted portfolio comprising the whole hedge fund industry (ALL). In what follows, we will refer to the returns in excess over the risk-free asset (3-month Treasury Bill). For simplicity in the implementation of stochastic dominance tests, the returns are defined as the differences between the logarithm prices and not in percentage terms. Some descriptive statistics are reported in Table 1.

[Insert Table 1 about here]

[5] The reader should note that these indices are not investable. The intention of this empirical section is to compare the predictive performance of investment styles in the spirit of Barberis and Schleifer (2003). For this reason, we are more interested in the representativeness of the indices than in the possibility of their being investment vehicles. Another reason for using these styles instead of investable indices is data availability. The Credit Suisse/Tremont (Blue Chip) database on these portfolios begins in August 2003. Nevertheless, Heidorn et al. (2010) report a correlation coefficient around 0.95 for the non-investable and investable indices elaborated by this data provider.

On average, only GM obtains higher excess returns than ALL. The case of DSB, which obtains a negative mean return over the whole sample, is particularly noteworthy. This style obtains, however, the highest maximum return, while EM obtains the lowest minimum one. In terms of skewness, EM, GM and ALL display negative values, more pronounced for EM. The positive skewness of DSB reveals the existence of very large positive returns. In addition, all hedge fund returns have excess kurtosis with the exception of MF, for which the Jarque-Bera test is not able to reject the null hypothesis of normality.[6] In general, these results confirm the non-normality of hedge fund returns reported in the related literature (see Wong et al., 2008; among others). Also in line with previously established evidence, the autocorrelation coefficients and the Ljung-Box statistic p-values suggest that the returns for DSB, GM and MF are serially correlated.

The set of explanatory factors for describing the hedge fund excess returns consists of the seven-factor model of Fung and Hsieh (2004), which has been shown to achieve considerable explanatory power, plus an eighth factor recently proposed by these authors (Fung and Hsieh, 2001) and given by the MSCI Emerging Market index monthly total excess return (MSCIEM). The seven-factor model includes three trend-following risk factors that are the excess returns on portfolios of lookback straddle options on bonds (BTF), currencies (CTF) and commodities (CMTF), constructed to replicate the maximum possible return on trend-following strategies in their respective underlying assets.[7] The two equity-oriented risk factors are the excess monthly total return of the S&P 500 index (EqMkt) and the Russell 2000 index monthly total return minus the S&P 500 monthly total return (SizeSpr). Two bond-oriented factors are the monthly change in the 10-year Treasury constant maturity yield (BMkt) and the monthly change in the Moody's Baa yield minus the 10-year Treasury constant maturity yield (CrdSpr). The corresponding descriptive statistics for these factors

[6] This finding is consistent with empirical studies using the same data (Frydenberg et al., 2008; Switze and Omelchak, 2009). An explanation of this finding may be that the index is asset-weighted and, hence, gives more weight to those funds with large capitalization. The normality of Managed Futures for this period suggests that these large-cap funds exhibit normally distributed returns. To confirm this, we have also constructed the equally-weighted counterpart portfolio of the Managed Futures style and obtained the same highly non-normal behavior observed for the other styles in the Credit Suisse/Tremont database.

[7] Downloadable from http://faculty.fuqua.duke.edu/~dah7/HFRFData.htm.

are shown in Table 2.

[Insert Table 2 about here]

The Jarque-Bera test shows the non-normality of the risk factors. EqMkt, CrdSpr and MSCIEM display negative skewness and all risk factors have excess kurtosis; SizeSpr, CrdSpr and MSCIEM are serially correlated.

3.2 Optimal Descriptive and Predictive Risk Factors

Hall et al. (2004) show that their cross-validation bandwidth selection method not only assigns optimal weights to the different relevant factors for estimating a conditional density, but also automatically determines the factors that are irrelevant. Our interest in this nonparametric estimation procedure is twofold. First, from a descriptive perspective, we use this method to determine the risk factors with power to explain the excess returns observed in hedge funds; and second, from a forecasting perspective, we need to know the set of relevant factors for predicting the conditional density of hedge fund returns. In contrast to the standard linear pricing models popularized by Fung and Hsieh and other authors and to nonlinear refinements, we are interested in finding out which factors have power not only for describing (and predicting) the expected excess return but its complete density.

We consider three different specifications of the excess return both from a descriptive and a predictive point of view. For the former approach, the returns on the hedge fund are regressed on a set of factors measured on the same date and, for the latter, the set of factors is considered one period lagged. We use a simple linear regression model estimated by ordinary least squares (OLS), a quantile regression model for the 25^{th}, 50^{th} and 75^{th} quantiles (QR25, QR50 and QR75, respectively), and the nonparametric conditional density estimation methods (NP). Our analysis covers 1994:01-2006:12, which will be considered later as the in-sample estimation period. The results of the descriptive approach are reported in Table 3.

[Insert Table 3 about here]

For the parametric methods, the risk factors are those variables that are found to be statistically significant at the 5% level; for the nonparametric alternative, the relevant factors are those for which the cross-validation bandwidth selection rule assigns a value lower than one. Our findings can be summarized as follows. First, the set of significant risk factors depends on the investment style. The Fung-Hsieh linear pricing model reveals that the equity-oriented risk factors are sufficient to explain the risk premium of the DSB style. For the EM style, the relevant risk factor is the emerging market index (MSCIEM) and the pricing model resembles a standard CAPM. Second, for a given style and with the exception of DSB, the relevant risk factors depend on the statistical measure under scrutiny. For example, the number of relevant factors for the EM style decreases as we move towards the upper region of the distribution. It is also interesting to observe that, for this style, the equity-oriented factors lose explanatory power beyond the lower tail of the return distribution. The asset-weighted portfolio comprising the whole hedge fund industry (ALL) is explained by the largest number of factors across different statistical measures. Third, the analysis of the whole conditional distribution through the nonparametric approach considers the largest set of explanatory factors. For example, seven of the eight potential factors are considered for explaining the conditional density of the GM style returns. Finally, and unsurprisingly, the nonparametric method achieves the highest log-likelihood values.

The above results dramatically change in the predictive framework. Following Wegener et al. (2010), we also include the lagged hedge fund excess return as a potential predictive factor. Results are reported in Table 4. The standard linear pricing model lacks any predictive power for the DSB style. In addition, except for the GM style, the risk factors are very poor at predicting the quantiles of the hedge fund returns distribution. More specifically, for MF and ALL, bond and equity-related factors have some predictive power for certain quantiles of the distribution; for DSB and EM, however, the predictive ability of the factors is null. Analogous to the descriptive approach, the nonparametric method makes use of the largest number of risk factors for constructing the conditional predictive density. Interestingly, these factors are very similar to those of the descriptive exercise previously discussed. This exercise reveals the importance of obtaining nonparametric estimates of the conditional predictive density of returns. In contrast to standard linear pricing formulations, we have found that the Fung-Hsieh risk factors also have predictive power when properly exploited.

[Insert Table 4 about here]

To assess the persistence of these factors, we compare the forecasting ability of the rolling and fixed forecasting schemes. Whereas, in the former scheme, the bandwidth parameters are recomputed for each rolling window, in the fixed scheme, the bandwidth parameter vector for estimating the conditional predictive density remains constant since it is computed only once. The evolution in the dynamics of the optimal bandwidth parameters obtained through the rolling scheme provide very valuable information on the ability of the factors to predict the conditional density over the out-of-sample period. If the optimal bandwidth parameter corresponding to a potential predictive factor is stable over the out-of-sample evaluation period and takes values of conventional size, there is evidence of the persistence of this factor for predicting the conditional density of returns. On the other hand, an erratic behavior in the dynamics of the optimal bandwidth parameter is evidence of abrupt changes in the predictive ability of the factor. We apply the predictive ability test developed by Amisano and Giacomini (2007) to the density forecast obtained from (1) using a fixed forecasting scheme (the bandwidth parameters are obtained from the period 1994:01-2006:12) and a rolling forecasting scheme in which the bandwidth parameters are re-estimated for each one-month-ahead rolling window. The results of the test in Table 5 reflect a strong persistence in the predictive ability of the factors revealed in Table 4. We observe no significant statistical differences in predictive ability between the fixed and rolling approaches. The Amisano and Giacomini test only finds statistical evidence of a superior predictive ability of the fixed approach for the DSB style when focused on predicting the center of the conditional return distribution. Figure 1 depicts the dynamics of the optimal bandwidth parameters for DSB and the factors BTF, EqMkt, BMkt and MSCIEM. The dynamics of these parameters are reasonably stable over the rolling out-of-sample evaluation period and give support to the results of the predictive ability test.

[Insert Table 5 and Figure 1 about here]

For completeness, Table 6 reports the set of fixed optimal bandwidth parameters corresponding to each potential predictive factor obtained for the in-sample period. The results highlight the differences between styles and factors. The only factor with power for all styles is BTF; SizeSpr and MSCIEM also have an important weight across styles. Nevertheless, the lagged return and the other Fung-Hsieh risk factors are relevant for the conditional density of three out of five of the styles considered.

[Insert Table 6 about here]

3.3 Forecasting the Performance of Hedge Fund Styles

The aim of this study is to predict, one period ahead, the best investment strategy from the set of hedge fund styles involved in directional trading. The null hypothesis of interest is $H_{0,\gamma} : E[d_{t,\gamma}(y) \mid I_{t-1} = x] \leq 0$ for all $y \in \Omega$ and x fixed for a given t. Critical values are obtained under the least favorable case $\widetilde{H}_{0,\gamma}$. The test is one-sided and has power against the hypothesis $E[d_{t,\gamma}(y) \mid I_{t-1} = x] > 0$. This test is implemented for all t in the out-of-sample evaluation period using a rolling scheme to incorporate the information into the test. Expression (9) is estimated using rolling windows of size $R = 160$. For the first out-of-sample observation, we use the sample 1994:01-2006:12 to construct $\widehat{D}_{\gamma}(y)$ for $y \in \Omega$, and simulate the p-value of the test. This exercise is repeated for 1994:02-2007:01 and so on to obtain a time series of 36 p-values over the period 2007:01-2009:12. The optimal bandwidth parameter vector corresponding to I_{t-1} is a fixed vector $(h_{1,A}, \ldots, h_{q,A}, h_{1,B}, \ldots, h_{q',B})$ with $q + q'$ being the number of factors relevant for at least one of the strategies. Following Li and Racine (2007), we consider $h_i = n^{-1/(4+(q+q'))}$, which is the optimal rate of convergence for cross-validation bandwidth selection methods. Another option could be the bandwidth estimates obtained from the nonparametric estimation of each conditional predictive density (1) (see Table 6 for the bandwidths corresponding to the fixed forecasting scheme).

[Insert Figures 2 and 3 about here]

Figure 2 plots the dynamic p-values of stochastic dominance tests of order one for each directional investment style against ALL. The dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of ALL over the individual styles. The solid line represents the p-value of the test defined by the converse null hypothesis, that is, the individual style dominates ALL for order one. Similarly, Figure 3 plots the dynamic p-values of the tests between all the possible pairwise combinations of directional investment styles. For DSB vs. EM, for example, the dashed line corresponds to the p-value of the test whose null hypothesis is given by the dominance of DSB over EM and the solid line reports the p-values of the converse hypothesis.

The test for first order stochastic dominance provides mixed results; pairwise comparisons between ALL and each of the DSB, EM and MF styles reveal no stochastic ordering between these portfolios. These styles are first order stochastically efficient suggesting no strict preference for one over another by risk-neutral investors. These investors should prefer, however, GM over ALL after the third quarter of 2007. This is because, during this period, the GM style strictly dominates the ALL style in the first order. The pairwise comparison between the directional styles also reveals that DSB is dominated in the first order by the other three investment styles during the crisis period; the test also predicts that EM dominates MF. Nevertheless, the dominance of GM with respect to DSB and MF vanishes at the end of the evaluation period.

Investors can use this information to decide, at time t, where to invest at time $t+1$. To shed more light on investors' preferences with respect to these hedge fund portfolios, we also consider higher orders of stochastic dominance reflecting risk aversion. Figures 4 to 7 show the dynamics of the p-values for the tests of second and third order stochastic dominance.

[Insert Figures 4 to 7 about here]

Fishburn (1977) shows that, if portfolio A dominates portfolio B in the first order, it also dominates it for higher orders of stochastic dominance. Our empirical findings are consistent with this theory. DSB is dominated by all the other styles; and ALL dominates DSB, EM and MF. The second order of stochastic dominance also predicts the dominance of GM over ALL; however, for the third order, it can be concluded that GM dominates ALL and that ALL dominates GM. This finding suggests that these portfolios are equally attractive for increasing levels of risk aversion. In contrast, DSB is expected to be the worst-performing style for investors concerned with the risk-return trade-off. The predicted dominance of EM over MF vanishes for higher orders of stochastic dominance, and hence, of risk aversion.

We also contemplate two further experiments as robustness checks. To see if this ordering of styles is robust to the choice and weighting schemes of the individual hedge funds comprising the indices, we have also constructed equally-weighted portfolios from all individual funds in each style, as reported in the Lipper TASS database.[8] The portfolio ALL is now constructed using an equally-weighted combination of all funds comprised in the Lipper TASS database whose styles coincide with those analyzed with Credit Suisse/Tremont. Figures 8 to 13 show the p-values of the test for the same three orders of stochastic dominance.

[8]As of October 2010, this database comprised 12,018 hedge funds and funds of hedge funds, of which 4,577 were actively reporting information. The styles to which each fund is assigned are based on those corresponding to the Credit Suisse/Tremont hedge fund indices. We have considered the monthly returns for the period 1995:01-2009:12 of those funds that report their performance in US Dollars. The number of individual funds that could be included in each portfolio differs among styles (50 funds for DSB, 747 for EM, 593 for GM and 2,095 for MF).

Although the dynamics of the p-values are slightly different for these synthetic indices, the results are comparable to those reported before. However, for the first order, ALL stochastically dominates GM for most of the evaluation period. For stochastic dominance of second and third orders, we observe that the GM style outperforms the other styles except ALL. Finally, DSB is outperformed by the rest of styles.

[Insert Figures 8 to 13 about here]

The second robustness check consists of studying a period without the tensions produced by the financial markets crisis. With this aim, we repeat the analysis with data prior to the crisis that began in 2007. The in-sample period covers 1994:01 to 2003:12 and the out-of-sample evaluation period is 2004:01-2006:12. The dynamic p-values for the predictive stochastic dominance test for the three orders are those in Figures 14 to 19. The main difference with respect to the crisis period emerges from the analysis of stochastic dominance between the pairs (ALL, GM) and (GM, EM). In this period, ALL and GM are stochastically efficient, in the sense that no portfolio dominates the other, not only for stochastic dominance of first order but also for second order. It is only under increasing levels of risk aversion that our test predicts that ALL and GM are equally attractive. The opposite is observed for this period between GM and EM. In contrast to the crisis period, GM stochastically dominates EM for the first order. The latter observation uncovers an interesting result: the stochastic dominance tests predict that GM is a dominating strategy compared to the other directional styles during tranquil periods. However, during the crisis, this result is only observed for higher orders of stochastic dominance, that is, under investors' risk aversion.[9]

[9] It is worth noting that, for some pairwise comparisons, the dynamics of the p-values are similar to those observed for the crisis period. This phenomenon is due to the fact that the quantity $\widehat{D}_\gamma$ in (9) is constructed using very similar samples for the crisis and noncrisis periods, hence the similarities obtained for those styles for which there are no significant changes in the returns dynamics between 2004 and 2006. The possibility of using windows of the same length, diminishing the extent of overlapping between the samples used for estimating $\widehat{D}_\gamma$ for the crisis and noncrisis exercises, is not feasible since no data on all these directional hedge fund styles is available before 1994.

[Insert Figures 14 to 19 about here]

The results discussed so far are predictions of future performance and, hence, in order to be accepted, they should be compared to ex-post performance. Following the related literature, we propose simple nonparametric proportion tests to assess the difference between the excess returns realized over the out-of-sample period. Let r_t^A and r_t^B denote these realized excess returns for two different investment strategies, A and B, and let $z_t = r_t^A - r_t^B$. We say that strategy A has been better than B for risk-neutral investors if $p_z = P\{z_t > 0\} > 0.5$. Similarly, we say that strategy A is better than B for risk-averse investors if $\widetilde{p}_z = P\{\widetilde{z}_t > 0\} > 0.5$ with $\widetilde{z}_t = r_t^A/\sigma_A - r_t^B/\sigma_B$, where σ_A and σ_B are the unconditional standard deviations of the returns on A and B over the out-of-sample evaluation period. These conditions can be tested as follows[10]:

$$H_0^{(n)} : p_z \leq 0.5 \text{ against } H_1^{(n)} : p_z > 0.5. \tag{11}$$

To test this condition over an evaluation period of length P, we propose the sample version of p_z given by $\widehat{p}_z = \frac{1}{P}\sum_{t=1}^{P} 1(z_t > 0)$. If z_t is serially uncorrelated, it is well known that the test statistic $\sqrt{P}\frac{\widehat{p}_z - p_z}{\sqrt{\widehat{p}_z(1-\widehat{p}_z)}}$ converges, as $P \to \infty$, to a standard normal distribution. Otherwise, we need to correct for the existence of serial correlation between the sequence of indicator functions. One possibility is to estimate the variance of $\widehat{p}_z$ using serially dependent robust estimators. These estimators provide a nice alternative for moderate sample sizes. For small values of P, block bootstrap methods are more suitable to approximate the distribution of the relevant test statistic defined now by $S_P = \sqrt{P}\,(\widehat{p}_z - p_z)$.

These resampling methods are based on blocking arguments in which data are divided into blocks that are resampled. The artificial time series obtained from this resampling procedure are of the same size as the original sample and mimic the dependence structure observed in the data. Let b, l denote integers such that $P = bl$, with b being the block size. There are two different ways of implementing a block bootstrap depending on whether

[10] The test for $\widetilde{p}_z$ is analogous and is omitted to save space.

the blocks are overlapping or non-overlapping. The overlapping rule produces $P - l + 1$ blocks of consecutive observations. We focus on the non-overlapping method that yields a sample of size P from l disjoint blocks $B_1, ..., B_l$ of size b, with $B_j = (1(z_{1+(j-1)b} > 0), \ldots, 1(z_{jb} > 0))$ and $j = 1, \ldots, l$. As in the *iid* bootstrap, the blocks can be repeated when resampling randomly with replacement. The asymptotic distribution of the out-of-sample test S_P can be approximated by the empirical distribution of the test statistic sequence $S_{P,i} = \sqrt{P}\left(\widehat{p}^{*}_{z,i} - \widehat{p}_z\right)$, indexed by $i = 1, ..., M$, with M being the number of Monte Carlo simulations, and $\widehat{p}^{*}_{z,i}$ the bootstrap counterpart of $\widehat{p}_z$ constructed from the simulated block bootstrap sample i. The empirical p-value of the test is obtained as

$$\widehat{p}_{P,b} = \frac{1}{M} \sum_{i=1}^{M} 1(S_{P,i} > S_P) \tag{12}$$

[Insert Table 7 about here]

Table 7 reports the results of the test $H_0^{(n)}$ for p_z and $\widetilde{p}_z$ equal to 0.5. Our choice of block size is based on the optimal data-driven algorithm of Politis et al. (2009). It varies across experiments with an average block size (b) of 2.223 and a standard deviation of 0.677. The number of Monte Carlo simulations is $2,000$. For both p_z and $\widetilde{p}_z$, we observe that the GM style slightly outperforms the other investment styles ex-post, the test being statistically significant at the 5% level against DSB and MF. In addition, the results for ALL confirm that this style has not dominated the other styles over the evaluation period as our tests for $\gamma = 1$ suggested. However, for $\gamma = 2$, only the prediction for DSB is in line with the empirical ex-post test. Finally, the findings for the DSB style reveal its ex-post underperformance against the other strategies.

3.4 Interpretation of results

The interest of this analysis is not only to learn about the performance of 'Directional Traders' hedge funds but also because these directional styles bet on market movements and so their returns are often strongly correlated with the market.

The outperformance of GM and EM over DSB highlights the superiority of 'return enhancer' strategies (Amenc et al., 2003) over 'risk reducer' strategies. The expected increase in overall volatility brought about by the former strategies is compensated by the higher returns obtained. On the other hand, the potential benefit of using DSB strategies concentrating on the short side and thereby sacrificing market-neutrality drops sharply over the crisis period. Hedge fund managers are not capable of properly forecasting how firms are affected by the global economic downturn. Investment on potentially declining funds have not lived up to expectations, indicating that the effect of the crisis has been somewhat unpredictable across firms and economic sectors. There is also the possibility that unexpected regulatory laws forbidding short selling in certain markets put forward to prevent a cascade of short selling orders in some stocks have also prevented these hedge fund strategies from fully capturing profit opportunities.

In contrast, the payoffs of hedge funds specializing in tactical trading strategies that attempt to profit by forecasting major macroeconomic events such as changes in interest rates, currency movements and stock market performance as well as the exact timing of these movements, have exceeded those of the other directional trading strategies. We have observed that the return profile of macro funds is much more volatile than that of other styles. Part of the reason is that macro funds often trade in instruments that are relatively illiquid. In searching for investment opportunities, GM hedge fund managers take into consideration a diverse set of factors such as geopolitical issues, economic indicators, market trends and liquidity flows. Our results suggest that hedge fund managers' expertise in these issues and, more importantly, their knowledge of global markets, has paid off more than taking wrong bets on the performance of firms trading in developed economies.

4 Concluding remarks

The aim of this study is to predict, one period ahead, the best investment strategy from a set of potential candidates. This is done by constructing a predictive test of stochastic dominance of arbitrary order that is applied to directional hedge fund investment styles during the recent global financial crisis. The empirical results provide a clear answer to the question of which style to choose one period ahead.

Under risk neutrality, the test does not provide a clear ranking of portfolio performance during the financial turmoil. It is under risk aversion when the Global Macro style is observed to be superior to the other strategies under analysis, including an asset-weighted portfolio comprising hedge funds from the whole industry. This finding is a little surprising given the diversification properties of the latter. Nevertheless, for increasing levels of risk aversion, we observe that the diversified strategy is equally attractive to the Global Macro style. These results are robust to the composition of the style portfolios. However, our results suggest that the Global Macro style also dominates the other directional styles under risk neutrality when a period previous to the crisis is analyzed.

A byproduct of our analysis has been the study of the optimal set of factors for describing as well as predicting the excess returns on hedge funds. The standard linear pricing formulation for modelling the risk premium on the returns has been extended to analyze the whole predictive density of returns. The cross-validation bandwidth selection method used for estimating these conditional density functions nonparametrically has been instrumental for determining which factors can predict the risk premium. For the linear pricing model, the factors proposed by Fung and Hsieh to explain the risk premium barely have predictive ability one period ahead. Interestingly, these factors are found to be highly significant if we consider the whole predictive density of returns. The nature of these factors depends on the style under consideration.

References

[1] Abhyankar, A., K.-Y. Ho, and H. Zhao, 2008. Value versus growth: Stochastic dominance criteria. Quantitative Finance 8, 693-704.

[2] Agarwal, V., N. Daniel and N. Naik, 2009. Role of Managerial Incentives and Discretion in Hedge Fund Performance. Journal of Finance 64(5), 2221-2256.

[3] Agarwal, V. and N. Naik, 2004. Risks and portfolio decisions involving hedge funds. Review of Financial Studies 17(1), 63-98.

[4] Amenc, N., S. El Bied and L. Martellini, 2003. Predictability in Hedge Fund Returns. Financial Analysts Journal 59(5), 32-46.

[5] Amenc, N. and L. Martinelli, 2002. Portfolio optimization and hedge fund style allocation decisions. Journal of Alternative Investments 5, 7-20.

[6] Amenc, N., L. Martinelli and M. Vaissié, 2003. Benefits and Risks of Alternative Investment Strategies. Journal of Asset Management 4, 96–118.

[7] Amisano, G. and R. Giacomini, 2007. Comparing Density Forecasts via Weighted Likelihood Ratio Tests. Journal of Business and Economic Statistics 25(2), 177-190.

[8] Anderson, G, 1996. Nonparametric Tests of Stochastic Dominance in Income Distributions. Econometrica 64, 1183-1193.

[9] Avramov, D., R. Kosowski, N. Naik and M. Teo, 2011. Hedge funds, managerial skill, and macroeconomic variables. Journal of Financial Economics 99, 672-692.

[10] Bali, T., S. Brown and M. Caglayan, 2011. Do Hedge Funds' Exposures to Risk Factors Predict Their Future Returns? Journal of Financial Economics 101, 36-68.

[11] Barberis, N. and A. Shleifer, 2003. Style Investing. Journal of Financial Economics 68, 161-199.

[12] Barrett, G. and S. Donald, 2003. Consistent tests for stochastic dominance. Econometrica 71, 71-104.

[13] Billio, M., M. Getmansky and L. Pelizzon, 2009. Nonparametric Analysis of Hedge Fund Returns: New Insights from High Frequency Data. Journal of Alternative Investments 12(1), 21-38.

[14] Cappoci, D. and G. Hübner, 2004. Analysis of hedge fund performance. Journal of Empirical Finance 11, 55-89.

[15] Cenesizoglu, T. and A. G. Timmermann, 2008. Is the Distribution of Stock Returns Predictable? Available at SSRN: http://ssrn.com/abstract=1107185.

[16] Davidson, R. and J. Duclos, 2000. Statistical inference for stochastic dominance and for the measurement of povery and inequality. Econometrica 68, 1435-1464.

[17] Diks, C., V. Panchenko and D. van Dijk, 2011. Likelihood-based scoring rules for comparing density forecasts in tails. Journal of Econometrics 163(2), 215-230.

[18] Eling, M. and R. Faust, 2010. The performance of hedge funds and mutual funds in emerging markets. Journal of Banking and Finance 34(8), 1993-2009.

[19] Fishburn, P., 1977. Mean-Risk Analysis with Risk Associated with Below-Target Returns. American Economic Review 67(2), 116-126.

[20] Fong, W.M., W.K. Wong and H.H. Lean, 2005. International momentum strategies: A stochastic dominance approach. Journal of Financial Markets 8, 89-109.

[21] Frydenberg, S., S. Lindset and S. Westgaard, 2008. Hedge Fund Return Statistics, 1994-2005. Journal of Investing 17(1), 7-21.

[22] Fung, W. and D. Hsieh, 1997. Empirical Characteristics of Dynamic Trading Strategies: The Case of Hedge Funds. Review of Financial Studies 10, 275-302.

[23] Fung, W. and D. Hsieh, 2001. The risk in hedge fund strategies: Theory and evidence from trend followers. Review of Financial Studies 14, 313–341.

[24] Fung, W. and D. Hsieh, 2002. The risk in fixed-income hedge fund styles. Journal of Fixed Income 12, 6–27.

[25] Fung, W. and D. Hsieh, 2004. Hedge fund benchmarks: A risk based approach. Financial Analysts Journal 60(5), 65–80.

[26] Fung, W. and D. Hsieh, 2011. The risk in hedge fund strategies: Theory and evidence from long/short equity hedge funds. Forthcoming in Journal of Empirical Finance.

[27] Geweke, J. and G. Amisano, 2010. Comparing and evaluating Bayesian predictive distributions of asset returns. International Journal of Forecasting 26, 216-230.

[28] Giannikis, D. and I. Vrontos, 2011. A Bayesian approach to detecting nonlinear risk exposures in hedge fund strategies. Journal of Banking and Finance 35(6), 1399-1414.

[29] Gneiting, T. and R. Ranjan, 2011. Comparing density forecasts using threshold- and quantile-weighted scoring rules. Journal of Business and Economic Statistics 29(3), 411-422.

[30] Gonzalo, J. and J. Olmo, 2011. Conditional Stochastic Dominance Tests in Dynamic Settings. Mimeo.

[31] Hall, P., J. Racine and Q. Li, 2004. Cross-Validation and the Estimation of Conditional Probability Densities. Journal of the American Statistical Association 99, 1015-1026.

[32] Hamza, O., M. Kooli and M. Roberge, 2006. Further Evidence on Hedge Fund Return Predictability. The Journal of Wealth Management, 9(3), 68-79.

[33] Hansen, B., 1996. Inference when a nuisance parameter is not identified under the null hypothesis. Econometrica 64(2), 413-430.

[34] Harvey, C. R. and A. Siddique, 2000. Conditional Skewness in Asset Pricing Tests. Journal of Finance 55(3), 1263-1295.

[35] Heidorn, T., D. Kaiser and A. Voinea, 2010. The Value-Added of Investable Hedge Fund Indices. Journal of Health Management 13(3), 59-79.

[36] Keim, D. B. and R. F. Stambaugh, 1986. Predicting Returns in the Stock and Bond Markets. Journal of Financial Economics 17(2), 357-390.

[37] Klecan, L., R. McFadden and R. McFadden, 1991. A Robust Test for Stochastic Dominance. Mimeo.

[38] Levy, H., 2006. Stochastic Dominance: Investment Decision Making under Uncertainty. Springer, New York.

[39] Li, Y. and H. Kazemi, 2007. Conditional Properties of Hedge Funds: Evidence from Daily Returns. European Financial Management 13(2), 211–238.

[40] Li, S. and O. Linton, 2010. Evaluating Hedge Fund Performance: A Stochastic Dominance Approach. In Handbook of Portfolio Construction: Contemporary Applications of Markowitz Techniques (J. B. Guerard, ed.), 551-564.

[41] Li, Q. and J. S. Racine, 2007. Nonparametric Econometrics. Princeton University Press.

[42] Linton, O., E. Maasoumi and Y. J. Whang, 2005. Consistent testing for stochastic dominance under general sampling schemes. Review of Economic Studies 72, 735-765.

[43] Linton, O., K. Song, and Y. J. Whang, (2010). An Improved Bootstrap Test of Stochastic Dominance. Journal of Econometrics, 154, 2, 186-202.

[44] Meligkotsidou, L., I. Vrontos and S. Vrontos, 2009. Quantile regression analysis of hedge fund strategies. Journal of Empirical Finance 16, 264-279.

[45] Patton, A., 2009. Are “Market Neutral” Hedge Funds Really Market Neutral? Review of Financial Studies 22(7), 2495-2530.

[46] Politis, D., H. White and A. Patton, 2009. Correction: Automatic Block-Length Selection for the Dependent Bootstrap. Econometric Reviews 28, 373-375.

[47] Porter, R. B., 1974. Semivariance and Stochastic Dominance: A Comparison. American Economic Review 64, 200-204.

[48] Romano, J. P., and M. Wolf, 2011. Alternative Tests for Monotonicity in Expected Asset Returns. Working Paper Series 017. Department of Economics. University of Zurich.

[49] Scaillet, O. and N. Topaglou, 2010. Testing for Stochastic Dominance Efficiency. Journal of Business and Economic Statistics 28(1), 169-180.

[50] Sriboonchitta, S., W-K. Wong, S. Dhompongsa and H. Nguyen, 2010. Stochastic Dominance and Applications to Finance, Risk and Economics. Chapman & Hall/CRC, Boca Raton.

[51] Stone, B. K., 1973. A General Class of Three-Parameter Risk Measures. Journal of Finance 28, 675-685.

[52] Switzer, L. and A. Omelchak, 2009. Time-varying asset allocation across hedge fund indices. Journal of Derivatives and Hedge Funds 15, 70-85.

[53] Ter Horst, J. and G. Salganik, 2011. Style Chasing by Hedge Fund Investors. Available at SSRN: http://ssrn.com/abstract=1362576.

[54] Vrontos, I., 2012. Evidence for Hedge Fund Predictability from a Multivariate Student-t Full-Factor GARCH Model. Forthcoming in Journal of Applied Statistics.

[55] Wegener, C., R. von Nitzsch and C. Cengiz, 2010. An advanced perspective on the predictability in hedge fund returns. Journal of Banking and Finance 34(11), 2694-2708.

[56] Wong, W-K., K. Phoon and H. Lean, 2008. Stochastic dominance analysis of Asian hedge funds. Pacific-Basin Finance Journal 16, 204-223.

Table 1. Descriptive statistics of excess returns. 'Directional Traders' hedge fund investment styles, 1994:01-2009:12.

	ALL	DSB	EM	GM	MF
Mean	0.005	-0.004	0.005	0.007	0.003
Median	0.005	-0.008	0.011	0.008	0.0002
Maximum	0.081	0.223	0.160	0.101	0.095
Minimum	-0.080	-0.096	-0.234	-0.119	-0.098
Std. Dev.	0.022	0.049	0.045	0.030	0.034
Skewness	-0.268	0.737	-0.799	-0.101	0.033
Kurtosis	5.293	4.545	7.648	6.145	3.064
Jarque-Bera	44.384	36.459	193.241	79.457	0.069
p-value	0.000	0.000	0.000	0.000	0.966
AC(1)	0.202	0.093	0.320	0.084	0.069
p-value	0.005	0.194	0.000	0.240	0.333
AC(4)	-0.037	-0.060	-0.032	-0.072	0.004
p-value	0.036	0.424	0.000	0.465	0.122
AC(12)	-0.010	-0.134	-0.041	0.007	-0.057
p-value	0.132	0.437	0.002	0.031	0.022
Observations	192	192	192	192	192

Note: ALL: Hedge Fund industry, DSB: Dedicated Short Bias, EM: Emerging Markets, GM: Global Macro, and MF: Managed Futures. AC(p) is the autocorrelation coefficient of order p.

Table 2. Descriptive statistics of Fung-Hsieh hedge fund risk factors, 1994:01-2009:12.

	BTF	CTF	CMTF	EqMkt	SizeSpr	BMkt	CrdSpr	MSCIEM
Mean	-0.017	-0.001	-0.006	0.003	0.0008	-0.003	-0.0005	0.003
Median	-0.051	-0.045	-0.032	0.009	0.0003	-0.004	0.000	0.005
Maximum	0.684	0.898	0.644	0.100	0.184	0.275	0.216	0.166
Minimum	-0.256	-0.304	-0.234	-0.168	-0.163	-0.269	-0.253	-0.297
Std. Dev.	0.147	0.198	0.139	0.045	0.036	0.066	0.053	0.072
Skewness	1.459	1.366	1.263	-0.712	0.282	0.469	-0.491	-0.776
Kurtosis	5.995	5.623	5.532	4.096	7.479	7.043	8.276	4.874
Jarque-Bera	139.846	114.796	102.347	25.821	163.069	137.833	230.410	47.378
p-value	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
AC(1)	0.119	0.035	-0.038	0.112	-0.136	0.037	0.199	0.211
p-value	0.096	0.629	0.595	0.117	0.058	0.610	0.005	0.003
AC(4)	-0.060	-0.084	0.005	0.066	-0.035	0.009	-0.001	-0.026
p-value	0.435	0.115	0.988	0.254	0.286	0.001	0.012	0.017
AC(12)	-0.018	-0.074	0.027	0.053	0.039	0.008	0.079	-0.027
p-value	0.406	0.445	0.991	0.614	0.307	0.009	0.002	0.090
Observations	192	192	192	192	192	192	192	192

Note: BTF: Excess returns on portfolios of lookback straddle options on bonds, CTF: on currencies, CMTF: on commodities, EqMkt: Excess monthly total return of the S&P500 index, SizeSpr: Russell 2000 index monthly total return minus the S&P500 monthly total return, BMkt: Monthly change in the 10-year Treasury constant maturity yield, CrdSpr: Monthly change in Moody's Baa yield minus the 10-year Treasury constant maturity yield, MSCIEM: MSCI Emerging Market index monthly total excess return.

Table 3. Relevant risk factors for hedge fund returns, 1994:01-2006:12.

Method	Style	Risk factors	Log-lik.	Style	Risk factors	Log-lik.
OLS	DSB	EqMkt,SizeSpr	357.164	GM	BMkt,CrdSpr	337.860
QR25		EqMkt,SizeSpr	336.851		BTF,BMkt,CrdSpr,MSCIEM	344.080
QR50		EqMkt,SizeSpr	358.758			353.728
QR75		EqMkt,SizeSpr	355.571		BMkt	337.770
NP		EqMkt,SizeSpr	391.211		BTF,CTF,EqMkt,SizeSpr,BMkt,CrdSpr,MSCIEM	480.072
OLS	EM	MSCIEM	339.219	MF	BTF,CTF,CMTF,MSCIEM	336.292
QR25		SizeSpr,MSCIEM	337.975		BTF,CTF,CMTF,CrdSpr,MSCIEM	325.023
QR50		MSCIEM	350.241		BTF,CTF	337.049
QR75		MSCIEM	334.454		BTF,CTF,SizeSpr	326.113
NP		EqMkt,BMkt,MSCIEM	398.105		BTF,CTF,EqMkt,SizeSpr,CrdSpr,MSCIEM	436.241
OLS	ALL	BTF,CMTF,EqMkt,SizeSpr,BMkt,MSCIEM	425.316			
QR25		CTF,SizeSpr,BMkt,MSCIEM	436.739			
QR50		BTF,CTF,EqMkt,SizeSpr,MSCIEM	442.133			
QR75		CMTF,EqMkt,SizeSpr,MSCIEM	421.776			
NP		BTF,CTF,EqMkt,SizeSpr,BMkt,CrdSpr,MSCIEM	562.116			

Note: The returns on hedge funds are regressed on the set of factors measured on the same date using a linear regression model (OLS), a quantile regression model for the 25^{th}, 50^{th} and the 75^{th} quantiles (QR25, QR50 and QR75, respectively) and a nonparametric conditional density estimation (NP). For the parametric methods, the relevant risk factors are the statistically significant variables at the 5% level; for the nonparametric alternative the relevant factors are those for which the bandwidth selection method of Hall et al. (2004) assigns a value lower than one.

Table 4. Relevant predictive risk factors for hedge fund returns, 1994:01-2006:12.

Method	Style	Predictive factors	Log-lik.	Style	Predictive factors	Log-lik.
OLS	DSB		252.392	GM	BTF,BMkt.CrdSpr	330.135
QR25			247.357		BTF,EqMkt,BMkt,CrdSpr	334.424
QR50			247.993		BTF,SizeSpr,BMkt,CrdSpr,MSCIEM	346.940
QR75			233.023		BTF,SizeSpr,CrdSpr	326.792
NP		BTF,CMTF,EqMkt,SizeSpr,BMkt	314.528		AR1,BTF,CTF,EqMkt,SizeSpr,BMkt,CrdSpr,MSCIEM	388.771
OLS	EM	CTF	268.933	MF	CMTF	309.469
QR25			255.660		CrdSpr	301.831
QR50		AR1	278.958		EqMkt	308.049
QR75		AR1	278.252			295.039
NP		AR1,BTF,CTF,EqMkt,SizeSpr,MSCIEM	408.876		AR1,BTF,CMTF,CrdSpr	353.681
OLS	ALL	BTF,BMkt,CrdSpr	379.810			
QR25		BMkt,CrdSpr	380.021			
QR50			390.585			
QR75		BTF	370.117			
NP		BTF,CTF,CMTF,SizeSpr,BMkt,CrdSpr,MSCIEM	456.150			

Note: The returns on hedge funds are regressed on the set of factors lagged one period using a linear regression model (OLS), a quantile regression model for the 25^{th}, 50^{th} and the 75^{th} quantiles (QR25, QR50 and QR75, respectively) and a nonparametric conditional density estimation (NP). AR1 is the lagged excess returns. For the parametric methods, the relevant risk factors are the statistically significant variables at the 5% level; for the nonparametric alternative the relevant factors are those for which the bandwidth selection method of Hall et al. (2004) assigns a value lower than one.

Table 5. Amisano and Giacomini (2007) WLR test. Predictive ability comparison, 2007:01-2009:12 (P=36). NP fixed vs. rolling approaches.

	Distribution	Center	Tails	Right	Left
ALL	-1.440	-1.551	-1.112	-1.402	-0.691
DSB	1.320	1.984**	0.983	1.327	0.362
EM	1.346	0.405	1.627	0.091	1.594
GM	1.258	1.332	0.111	1.465	0.865
MF	-0.434	-0.345	-0.511	-0.167	-0.912

Note: In the fixed forecasting scheme the bandwidh parameter vector for estimating the conditional predictive density is obtained with the method of Hall et al. (2004) for 1994:01-2006:12. In the rolling forecasting scheme the bandwidth parameters are recomputed with rolling windows of 160 observations. A positive value of the test statistic points to a superior predictive ability of the fixed forecasting scheme, and vice versa. ** denotes statistically significant at the 5% level.

Table 6. Conditional predictive density estimation. Optimal bandwidth parameters, 1994:01-2006:12.

	AR1	BTF	CTF	CMTF	EqMkt	SizeSpr	BMkt	CrdSpr	MSCIEM
ALL	35,692	0.388	0.189	0.252	704,557	0.062	0.048	0.010	0.029
DSB	278,464	0,119	230,306	0.160	0.015	0.051	0.024	67,564	542,895
EM	0.018	0.061	0.222	4,767,454	0,024	0,025	437,010	891,992	0.089
GM	0.046	0.083	0.524	1,678,729	0.035	0.034	0.077	0.033	0.509
MF	0.056	0.072	1,593,954	0.019	799,204	1,003,099	622,058	0.027	0.021

Note: Optimal bandwidth parameters determined by the least squares cross-validation method of Hall et al. (2004).

Table 7. Nonparametric proportion tests between realized excess returns, 2007:01-2009:12.

	GM vs.				ALL vs.				DSB vs.			
	ALL	DSB	EM	MF	DSB	EM	GM	MF	ALL	EM	GM	MF
$p_z = 0.50$												
test-stat.	0.000	0.667***	0.000	0.333***	0.333	-0.333	0.000	0.167	-0.333	-0.667	-0.667	-0.500
p-value	0.242	0.000	0.414	0.000	0.254	1.000	0.234	0.248	0.744	1.000	1.000	1.000
$\widetilde{p}_z = 0.50$												
test-stat.	0.000	1.167***	0.333	0.833	0.833***	-0.167	0.000	0.333	-0.833	-0.500	-1.167	-0.667
p-value	0.250	0.000	0.241	0.251	0.000	0.622	0.253	0.271	1.000	1.000	1.000	1.000

Note: For style A vs. style B the null hypothesis of the p_z test is that the probability of the returns of style A being greater than those of style B is less than 50% over the out-of-sample evaluation period. The $\tilde{p}_z$ test compares the standardized excess returns. p-values obtained using a block bootstrap with block size determined by the method in Politis et al. (2009). M=2,000 Monte Carlo simulations. *** denotes rejection of the null hypothesis at the 1% significance level.

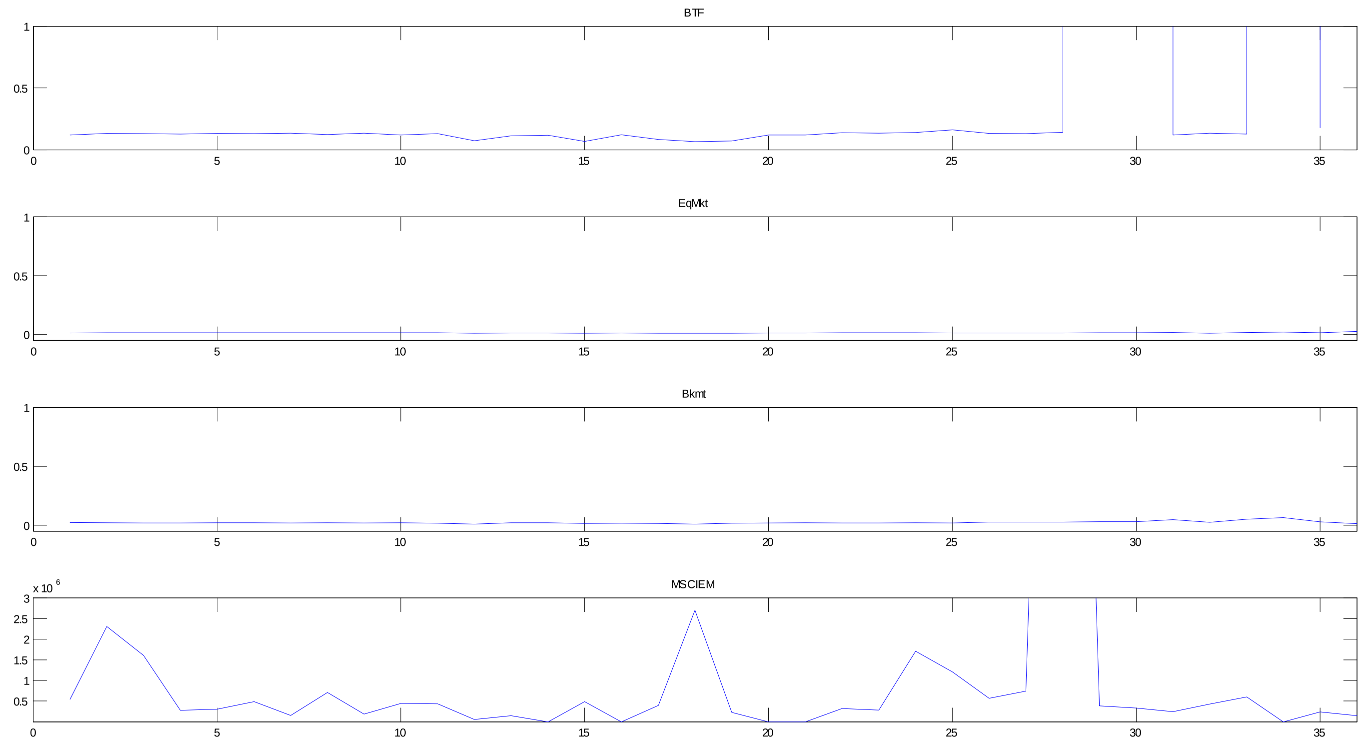

Figure 1: Dynamic optimal bandwidth parameters (Hall et al., 2004) for the DSB style. Rolling out-of-sample evaluation period, 2007:01-2009:12.

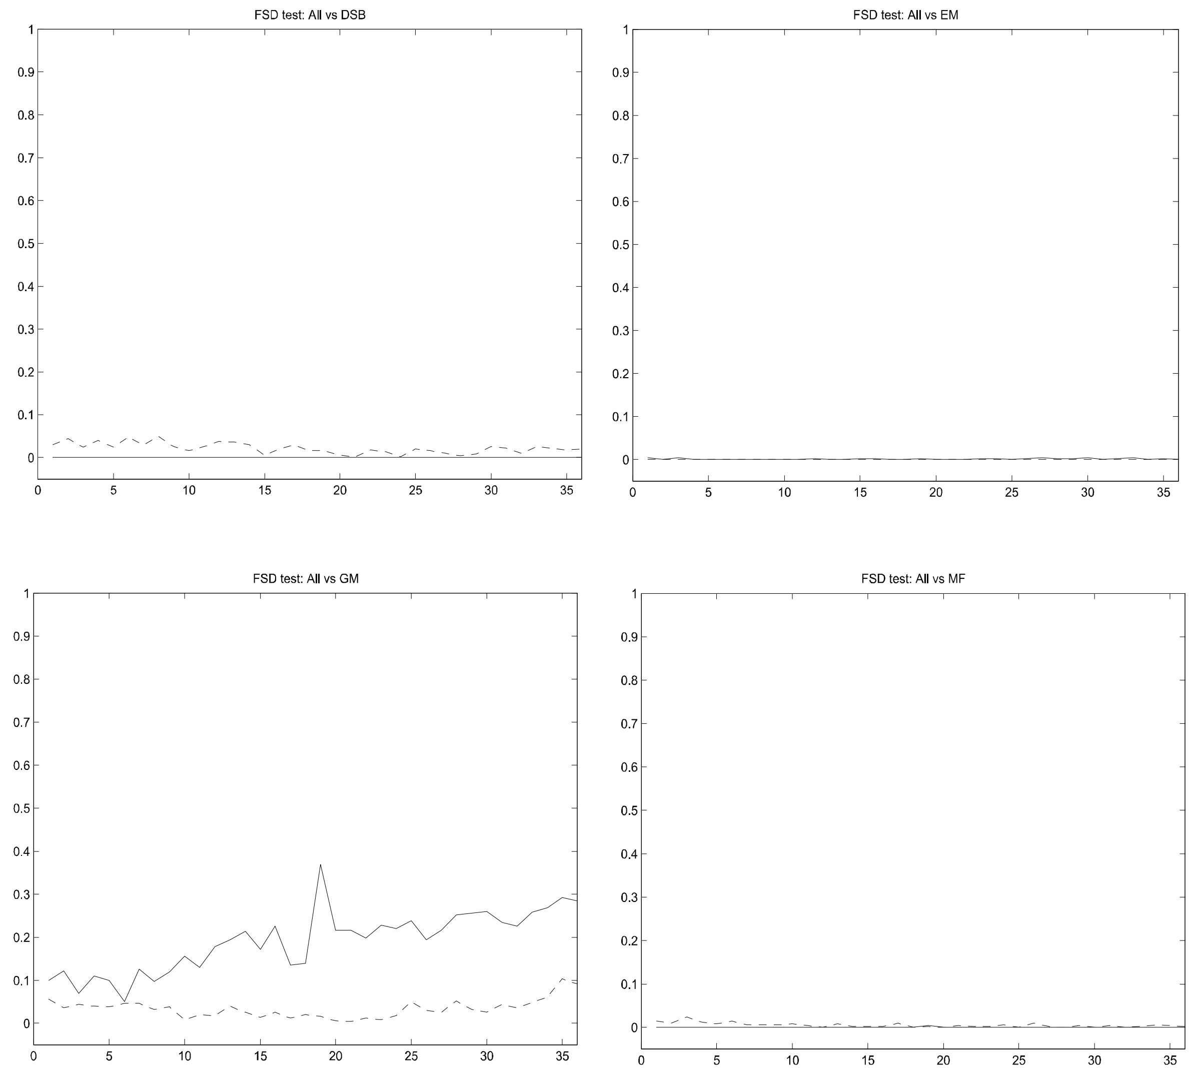

Figure 2: Dynamic p-values of predictive stochastic dominance of order 1. Out-of-sample comparison of the hedge fund industry (ALL) with 'Directional Traders' investment styles, 2007:01-2009:12. The dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of ALL over the individual styles. The solid line represents the p-value of the test defined by the converse null hypothesis.

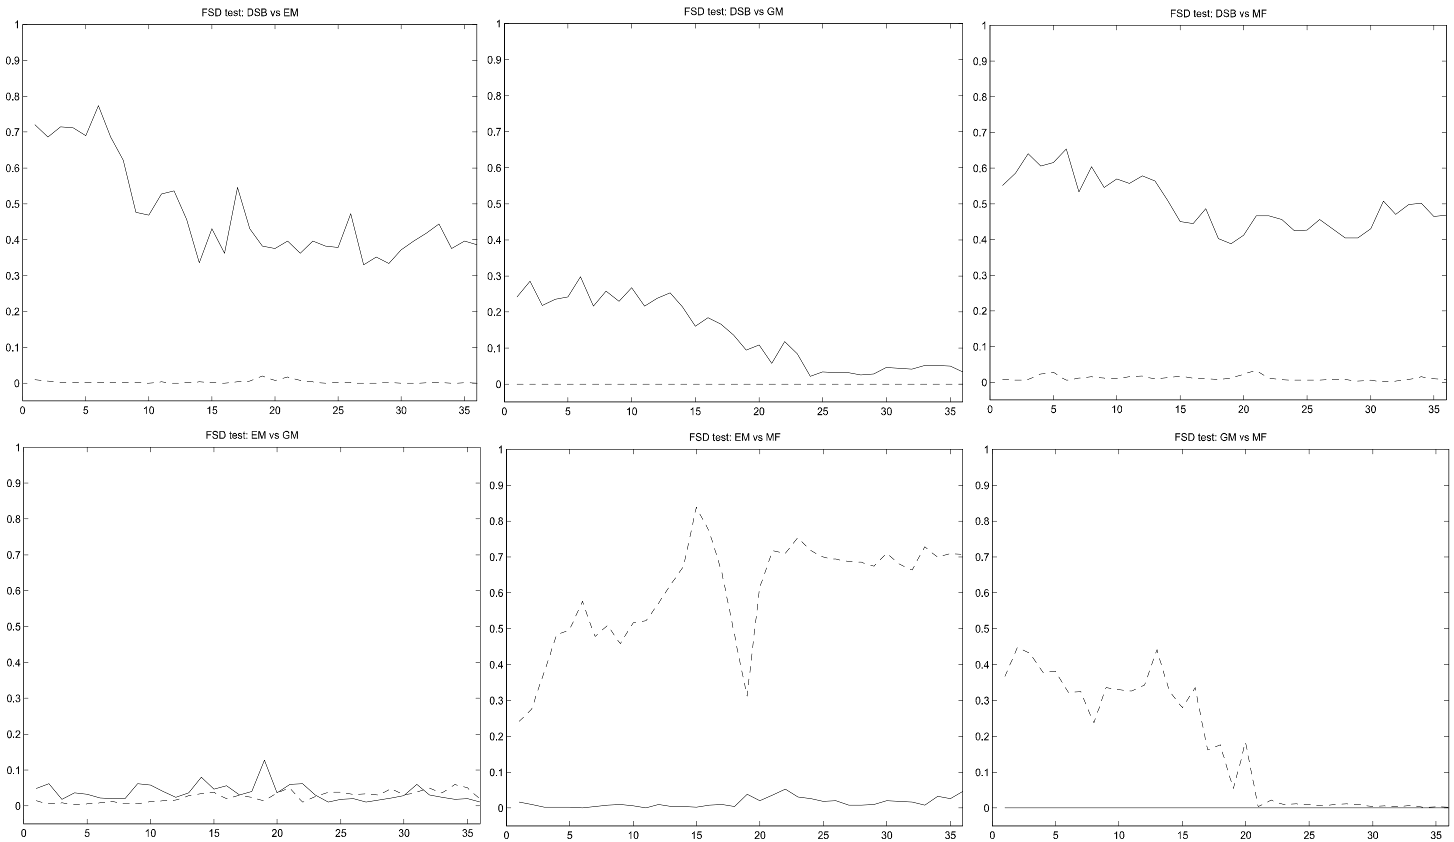

Figure 3: Dynamic p-values of predictive stochastic dominance of order 1. Out-of-sample comparison of 'Directional Traders' hedge fund investment styles, 2007:01-2009:12. For style A vs. style B the dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of style A over style B. The solid line represents the p-value of the test defined by the converse null hypothesis.

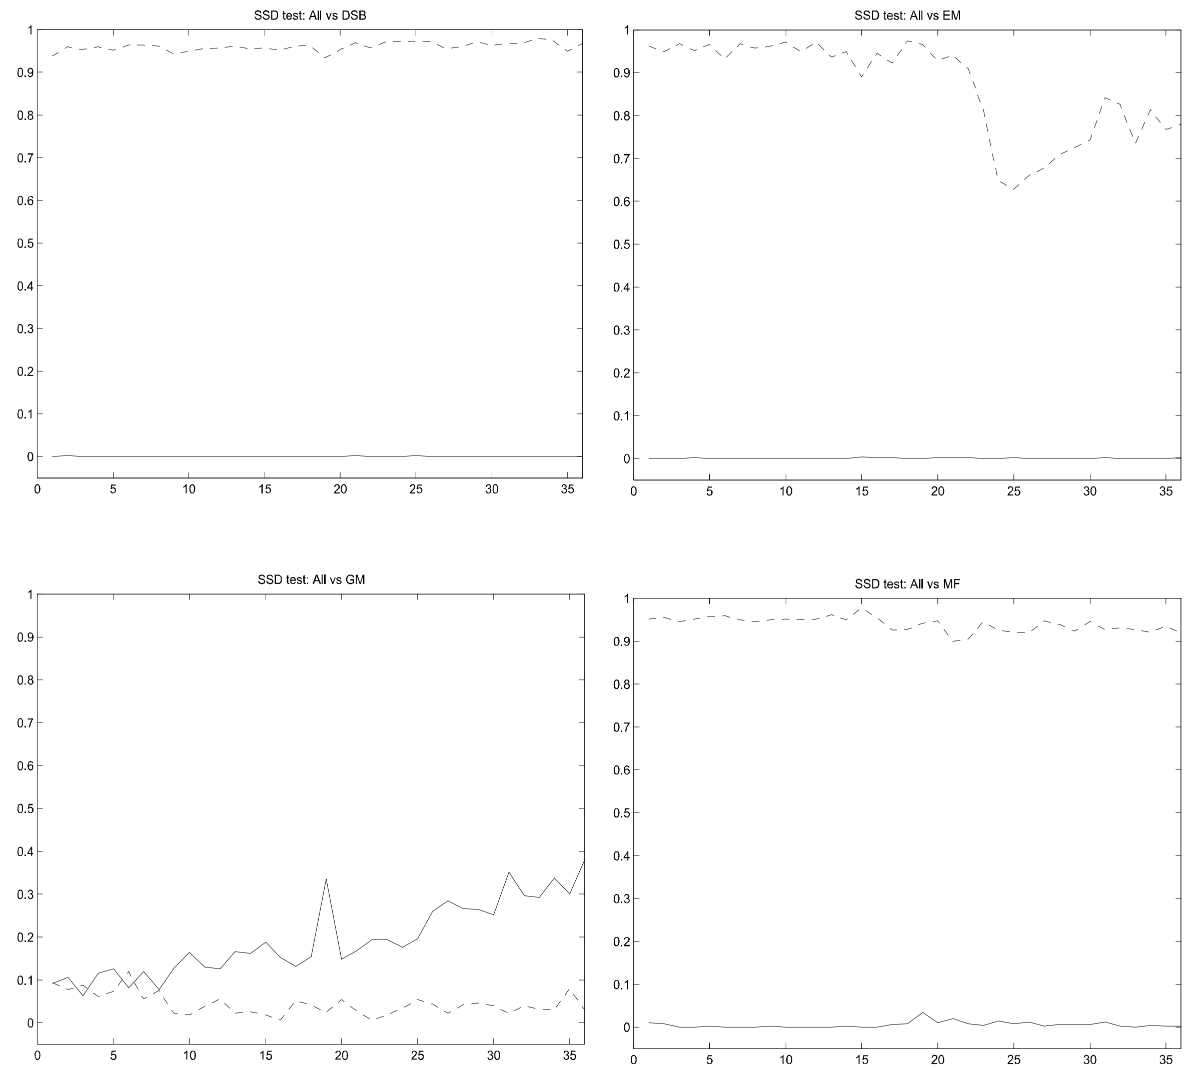

Figure 4: Dynamic p-values of predictive stochastic dominance of order 2. Out-of-sample comparison of the hedge fund industry (ALL) with 'Directional Traders' investment styles, 2007:01-2009:12. The dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of ALL over the individual styles. The solid line represents the p-value of the test defined by the converse null hypothesis.

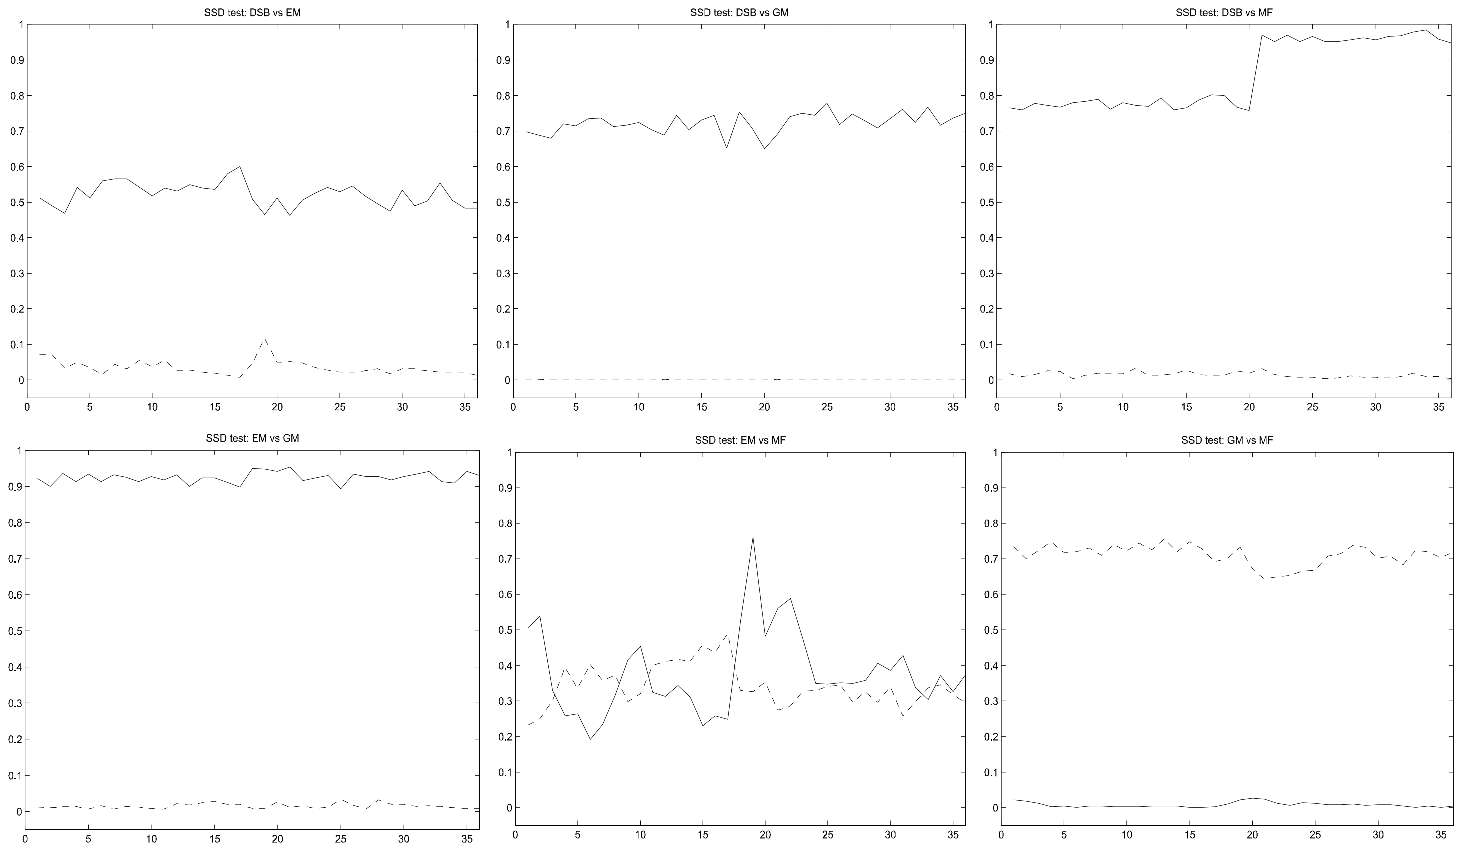

Figure 5: Dynamic p-values of predictive stochastic dominance of order 2. Out-of-sample comparison of 'Directional Traders' hedge fund investment styles, 2007:01-2009:12. For style A vs. style B the dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of style A over style B. The solid line represents the p-value of the test defined by the converse null hypothesis.

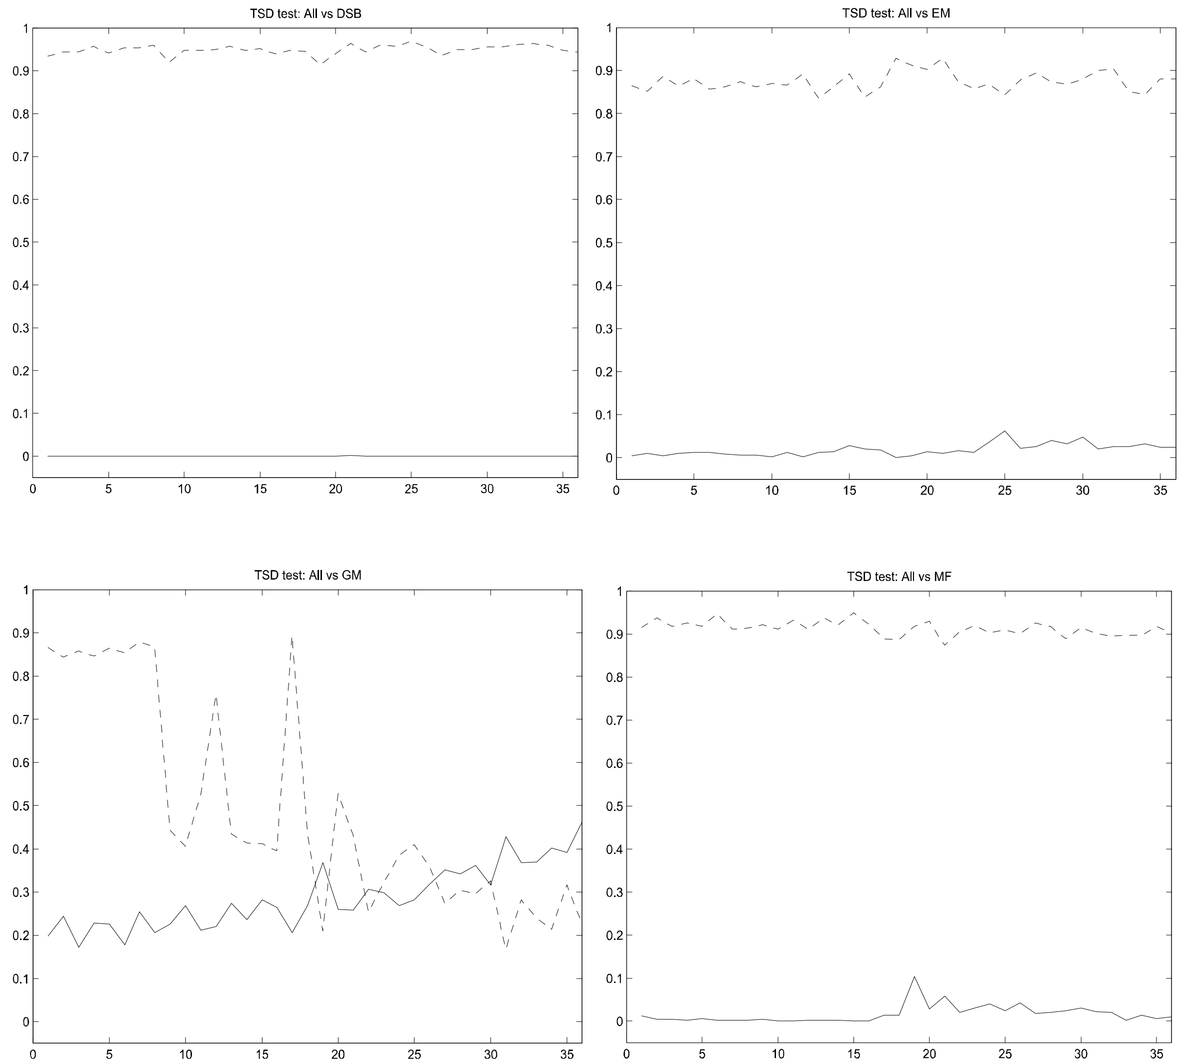

Figure 6: Dynamic p-values of predictive stochastic dominance of order 3. Out-of-sample comparison of the hedge fund industry (ALL) with ‘Directional Traders’ investment styles, 2007:01-2009:12. The dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of ALL over the individual styles. The solid line represents the p-value of the test defined by the converse null hypothesis.

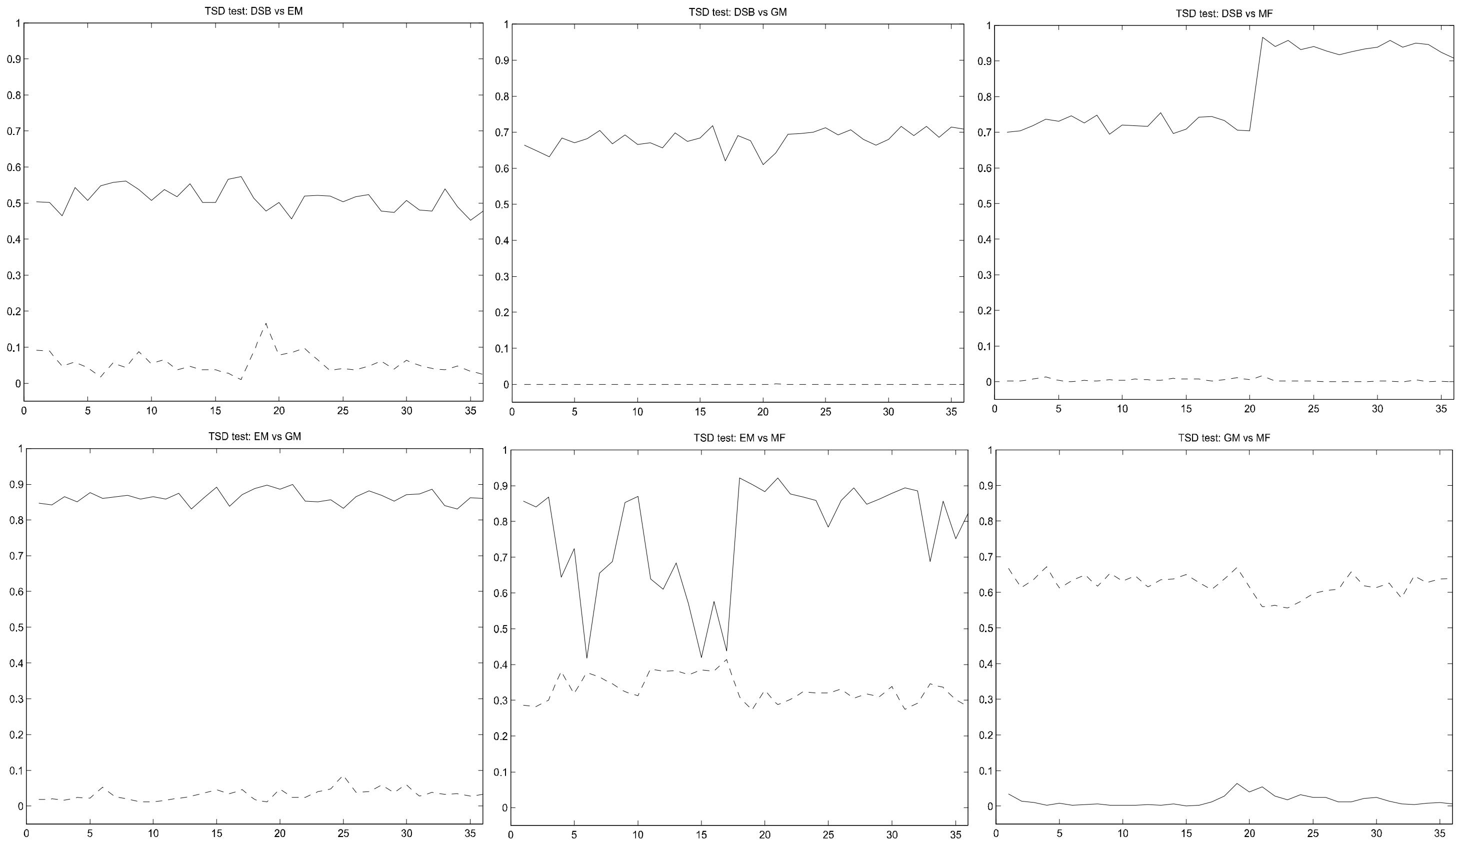

Figure 7: Dynamic p-values of predictive stochastic dominance of order 3. Out-of-sample comparison of 'Directional Traders' hedge fund investment styles, 2007:01-2009:12. For style A vs. style B the dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of style A over style B. The solid line represents the p-value of the test defined by the converse null hypothesis.

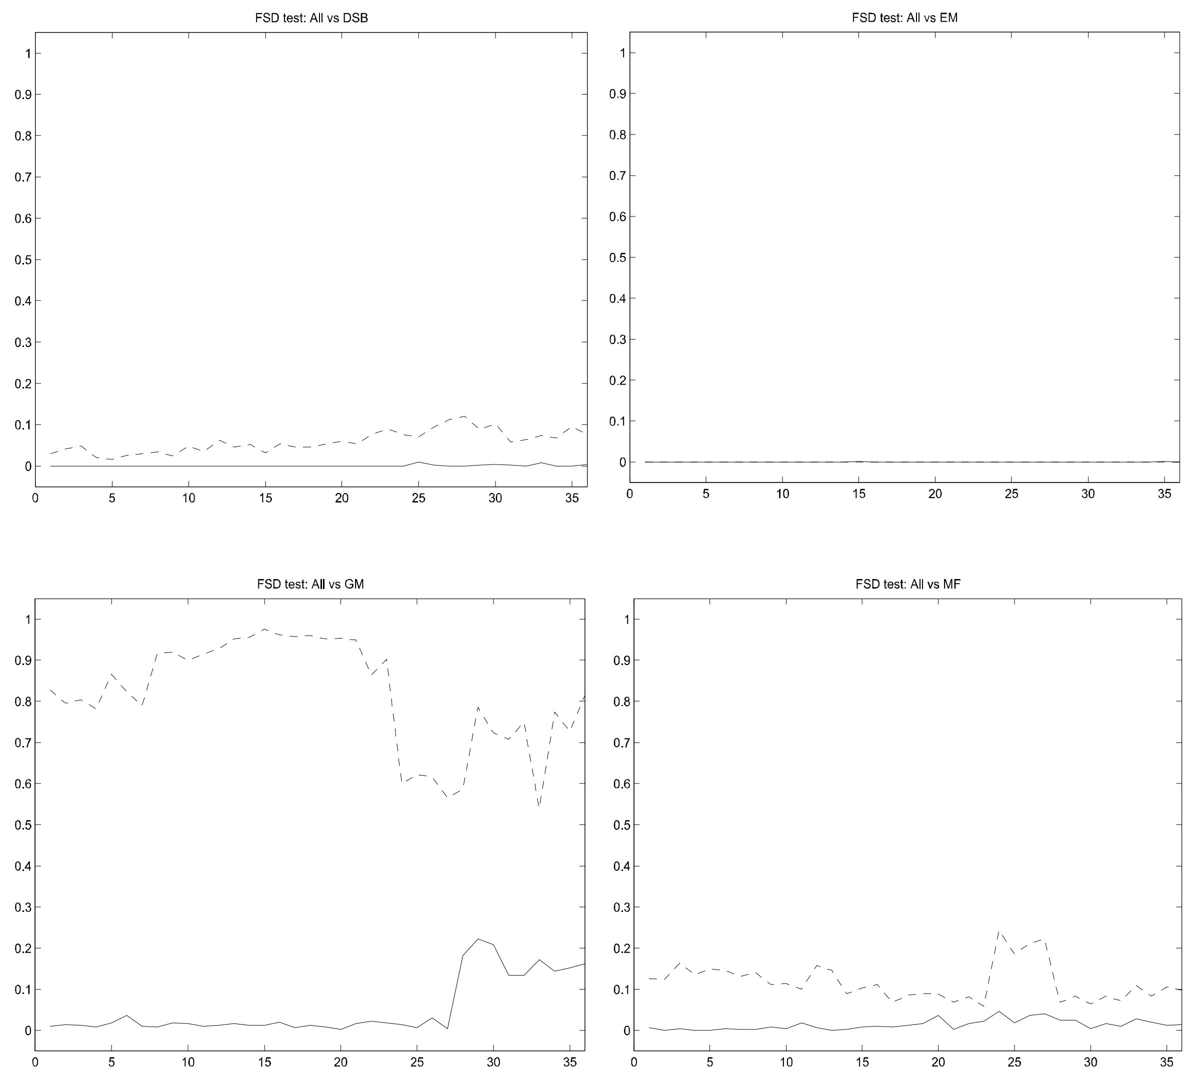

Figure 8: Dynamic p-values of predictive stochastic dominance of order 1 for equally-weighted portfolios comprising the whole universe of individual hedge funds in each directional style as reported in the Lipper TASS database. Out-of-sample comparison of the hedge fund industry (ALL) with 'Directional Traders' investment styles, 2007:01-2009:12. The dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of ALL over the individual styles. The solid line represents the p-value of the test defined by the converse null hypothesis.

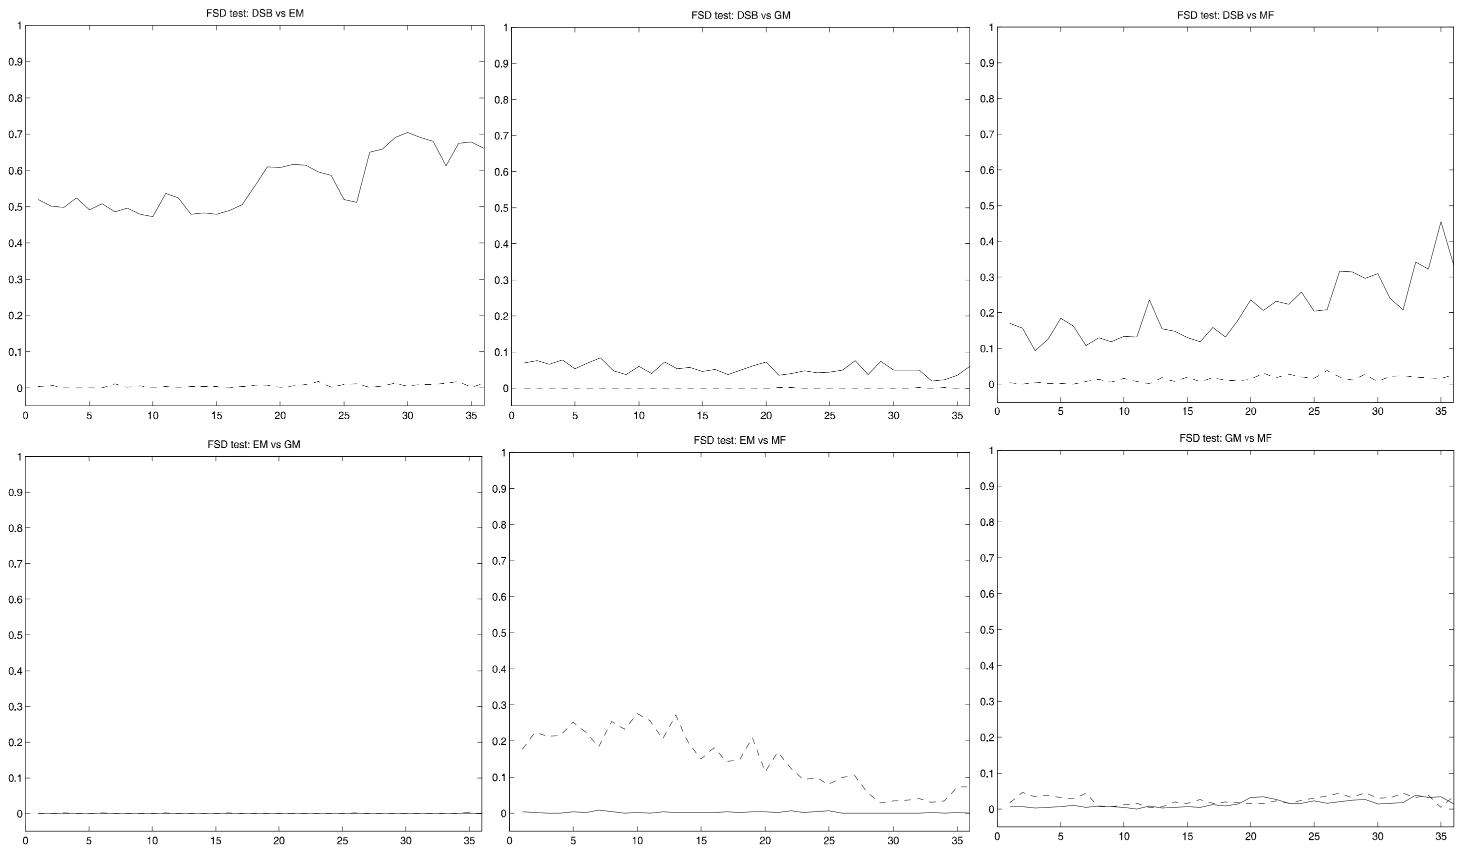

Figure 9: Dynamic p-values of predictive stochastic dominance of order 1 for equally-weighted portfolios comprising the whole universe of individual hedge funds in each directional style as reported in the Lipper TASS database. Out-of-sample comparison of 'Directional Traders' hedge fund investment styles, 2007:01-2009:12. For style A vs. style B the dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of style A over style B. The solid line represents the p-value of the test defined by the converse null hypothesis.

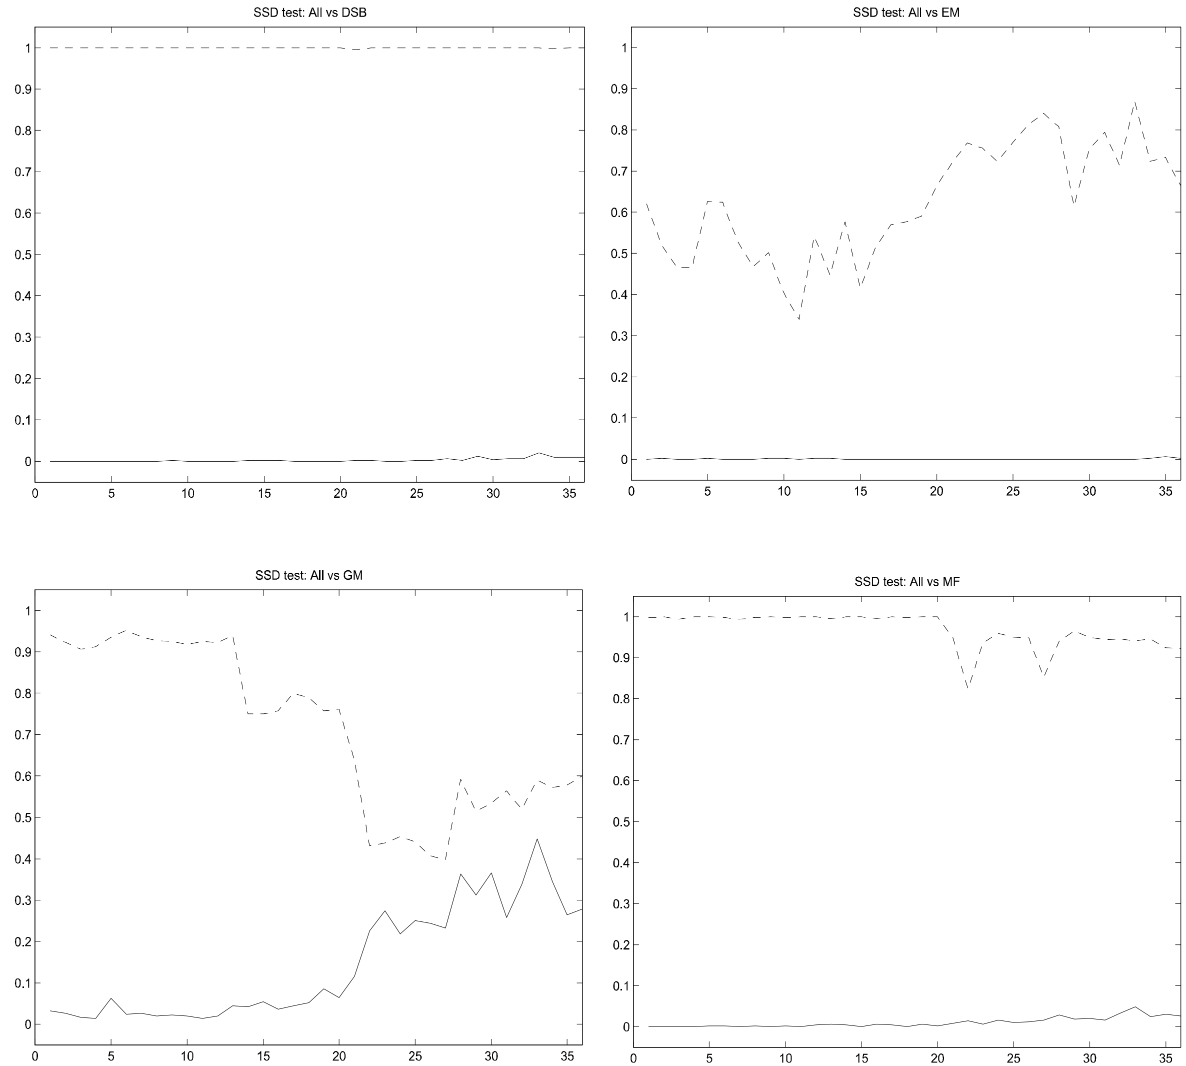

Figure 10: Dynamic p-values of predictive stochastic dominance of order 2 for equally-weighted portfolios comprising the whole universe of individual hedge funds in each directional style as reported in the Lipper TASS database. Out-of-sample comparison of the hedge fund industry (ALL) with 'Directional Traders' investment styles, 2007:01-2009:12. The dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of ALL over the individual styles. The solid line represents the p-value of the test defined by the converse null hypothesis.

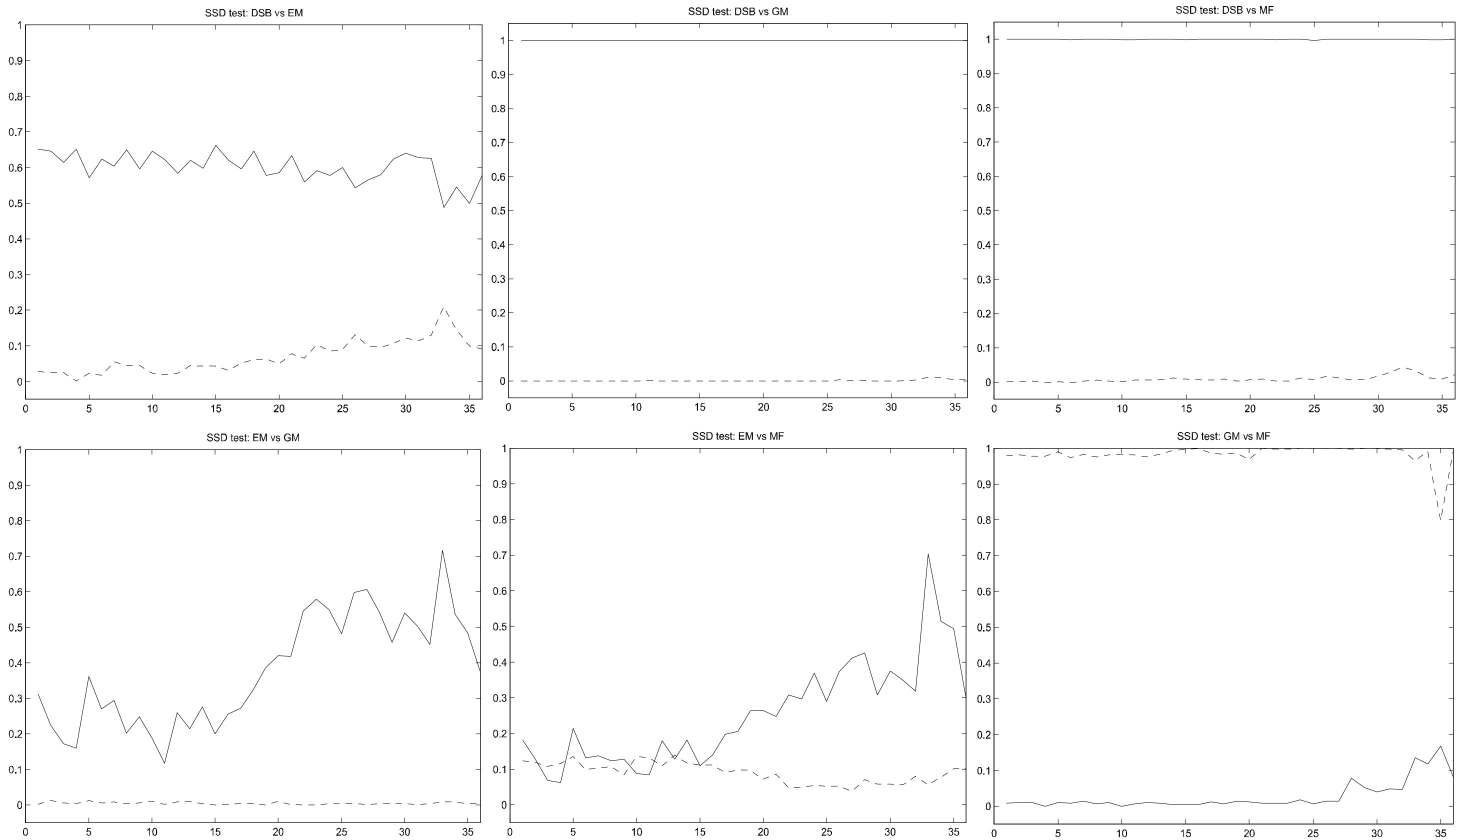

Figure 11: Dynamic p-values of predictive stochastic dominance of order 2 for equally-weighted portfolios comprising the whole universe of individual hedge funds in each directional style as reported in the Lipper TASS database. Out-of-sample comparison of 'Directional Traders' hedge fund investment styles, 2007:01-2009:12. For style A vs. style B the dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of style A over style B. The solid line represents the p-value of the test defined by the converse null hypothesis.

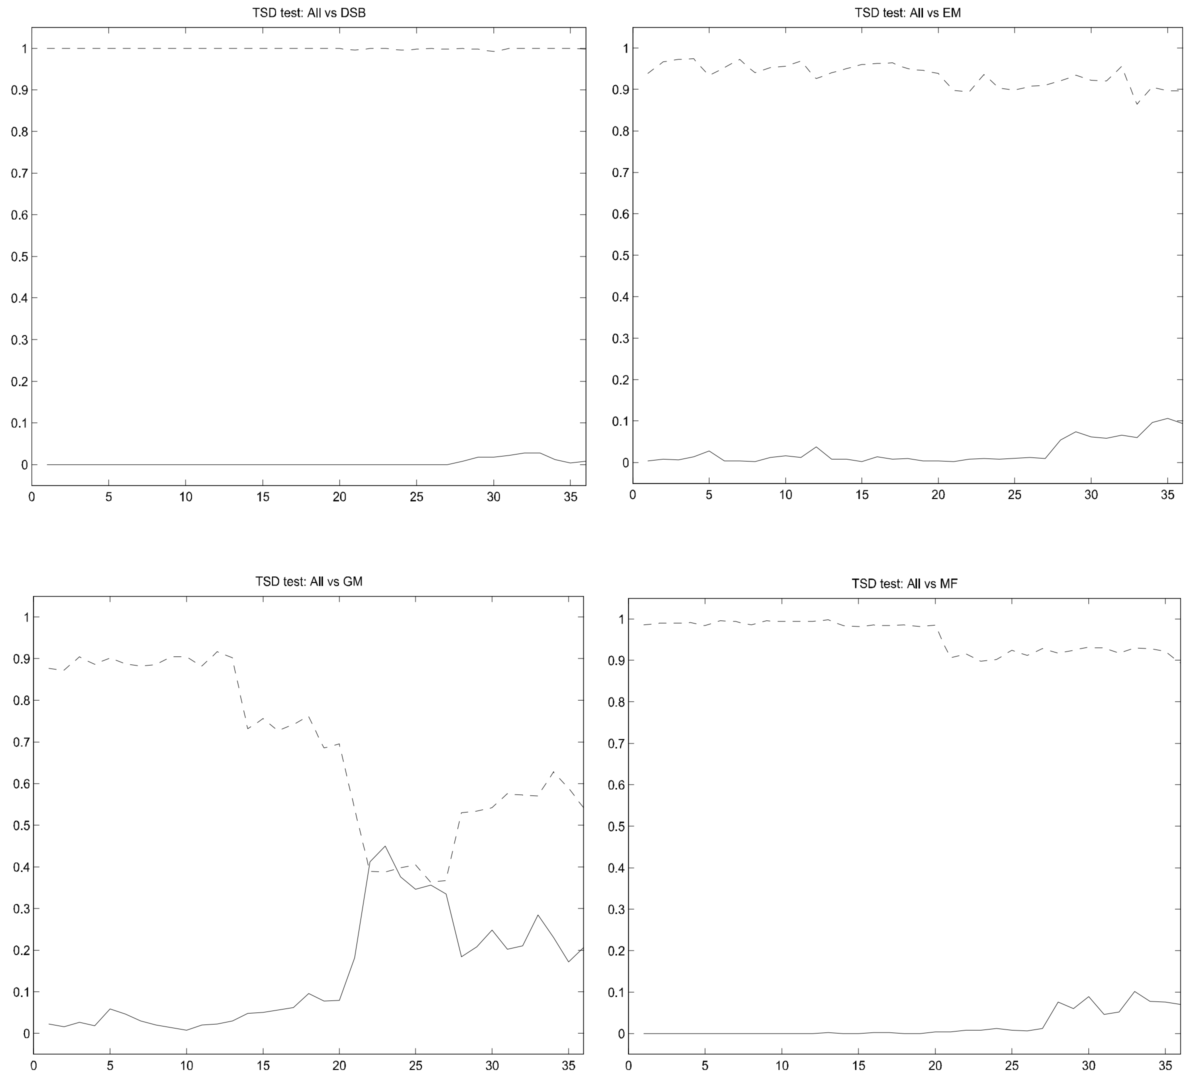

Figure 12: Dynamic p-values of predictive stochastic dominance of order 3 for equally-weighted portfolios comprising the whole universe of individual hedge funds in each directional style as reported in the Lipper TASS database. Out-of-sample comparison of the hedge fund industry (ALL) with 'Directional Traders' investment styles, 2007:01-2009:12. The dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of ALL over the individual styles. The solid line represents the p-value of the test defined by the converse null hypothesis.

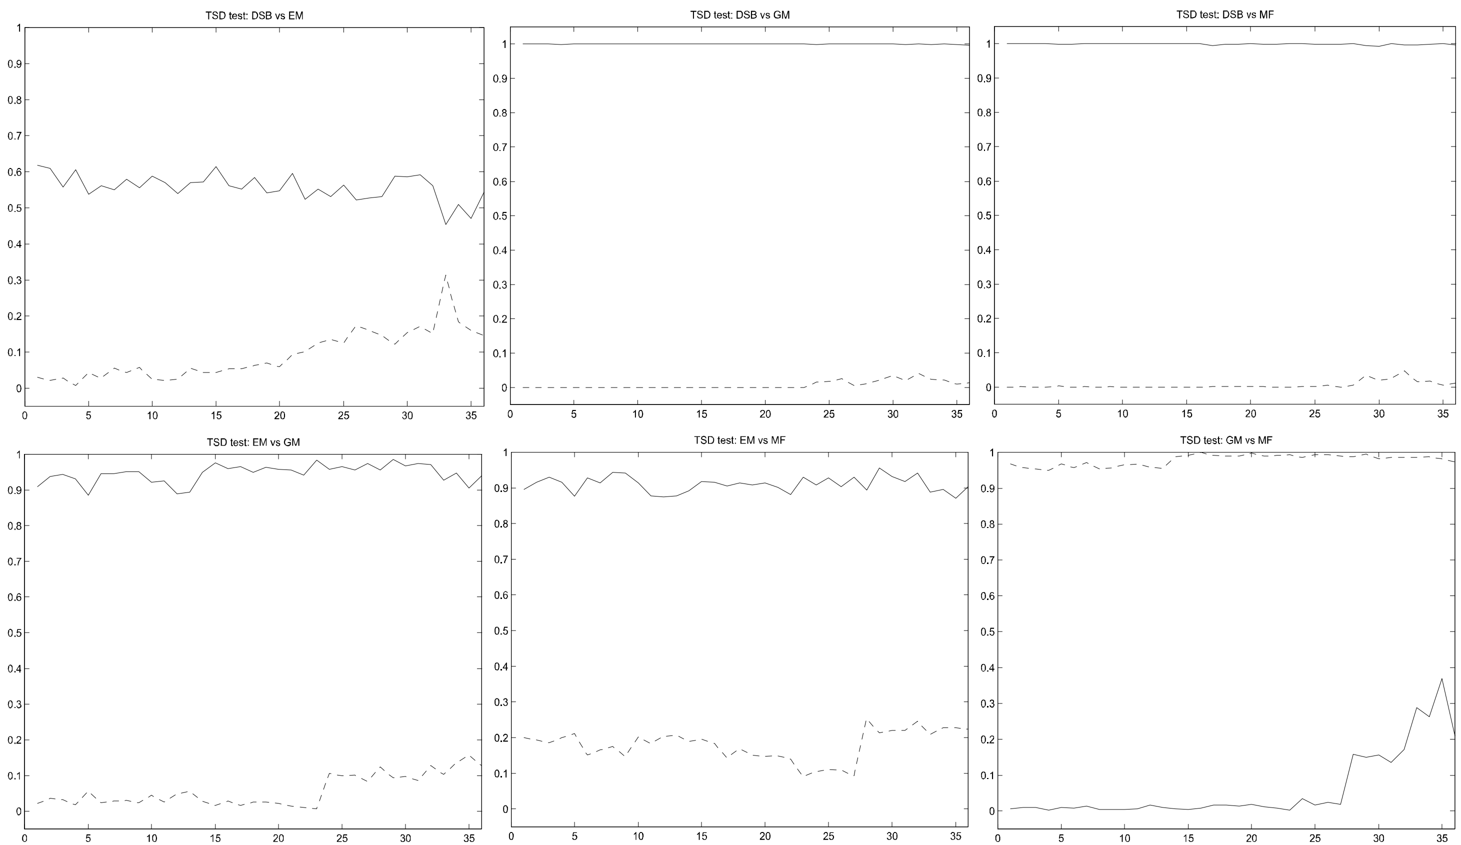

Figure 13: Dynamic p-values of predictive stochastic dominance of order 3 for equally-weighted portfolios comprising the whole universe of individual hedge funds in each directional style as reported in the Lipper TASS database. Out-of-sample comparison of 'Directional Traders' hedge fund investment styles, 2007:01-2009:12. For style A vs. style B the dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of style A over style B. The solid line represents the p-value of the test defined by the converse null hypothesis.

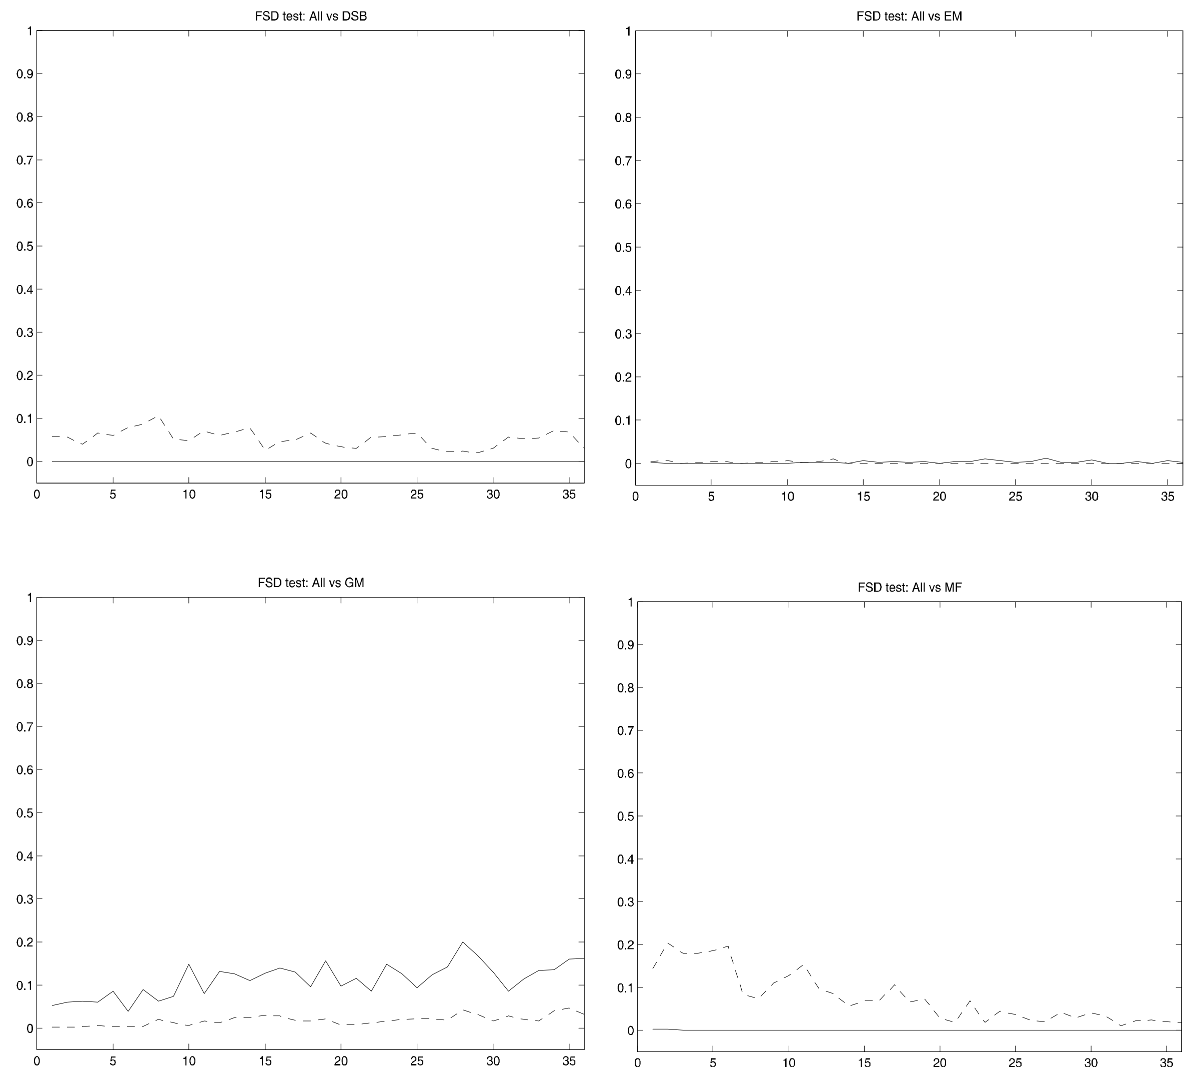

Figure 14: Dynamic p-values of predictive stochastic dominance of order 1. Out-of-sample comparison of the hedge fund industry (ALL) with ‘Directional Traders’ investment styles, 2004:01-2006:12. The dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of ALL over the individual styles. The solid line represents the p-value of the test defined by the converse null hypothesis.

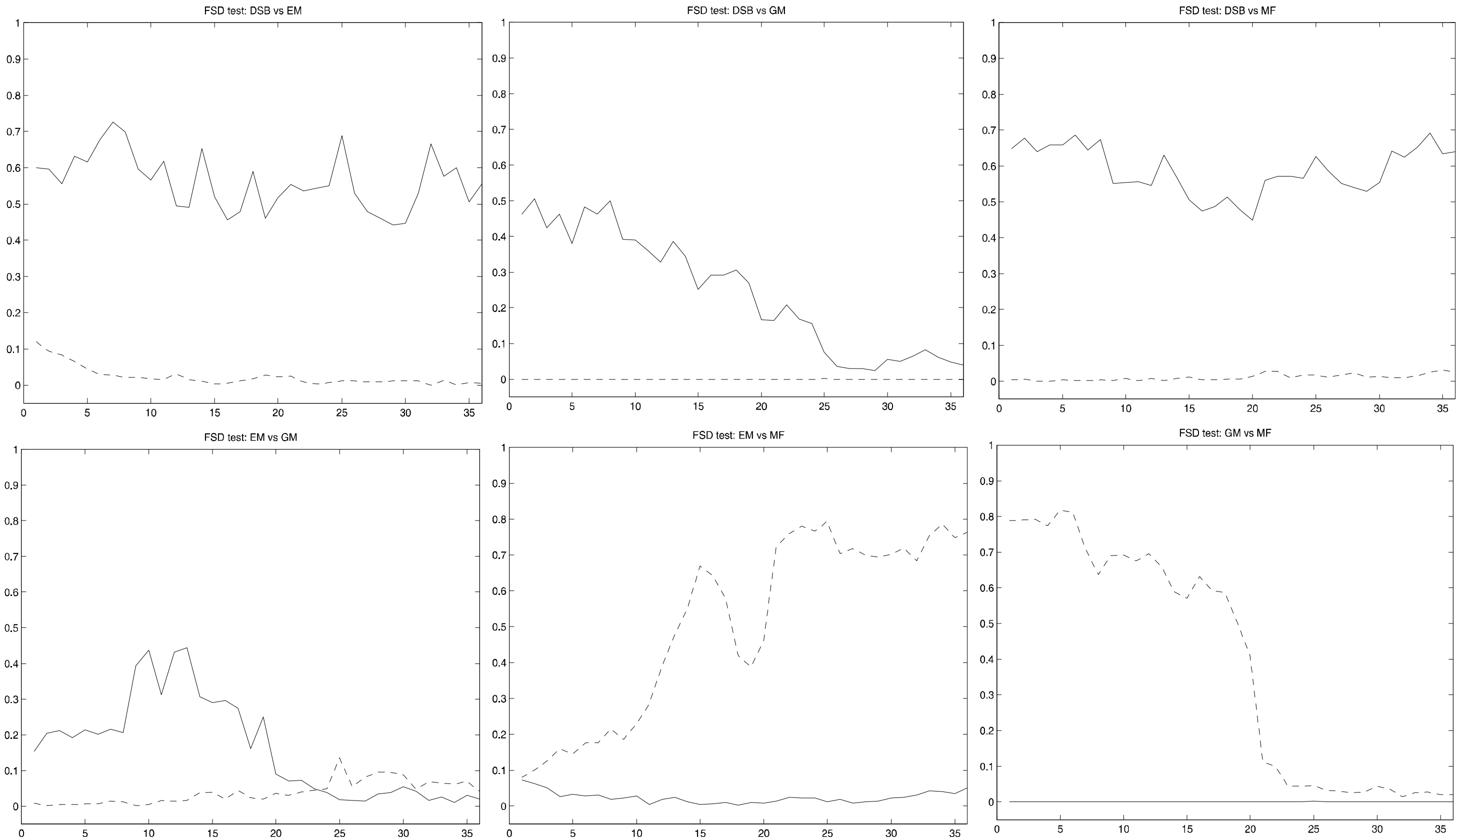

Figure 15: Dynamic p-values of predictive stochastic dominance of order 1. Out-of-sample comparison of 'Directional Traders' hedge fund investment styles, 2004:01-2006:12. For style A vs. style B the dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of style A over style B. The solid line represents the p-value of the test defined by the converse null hypothesis.

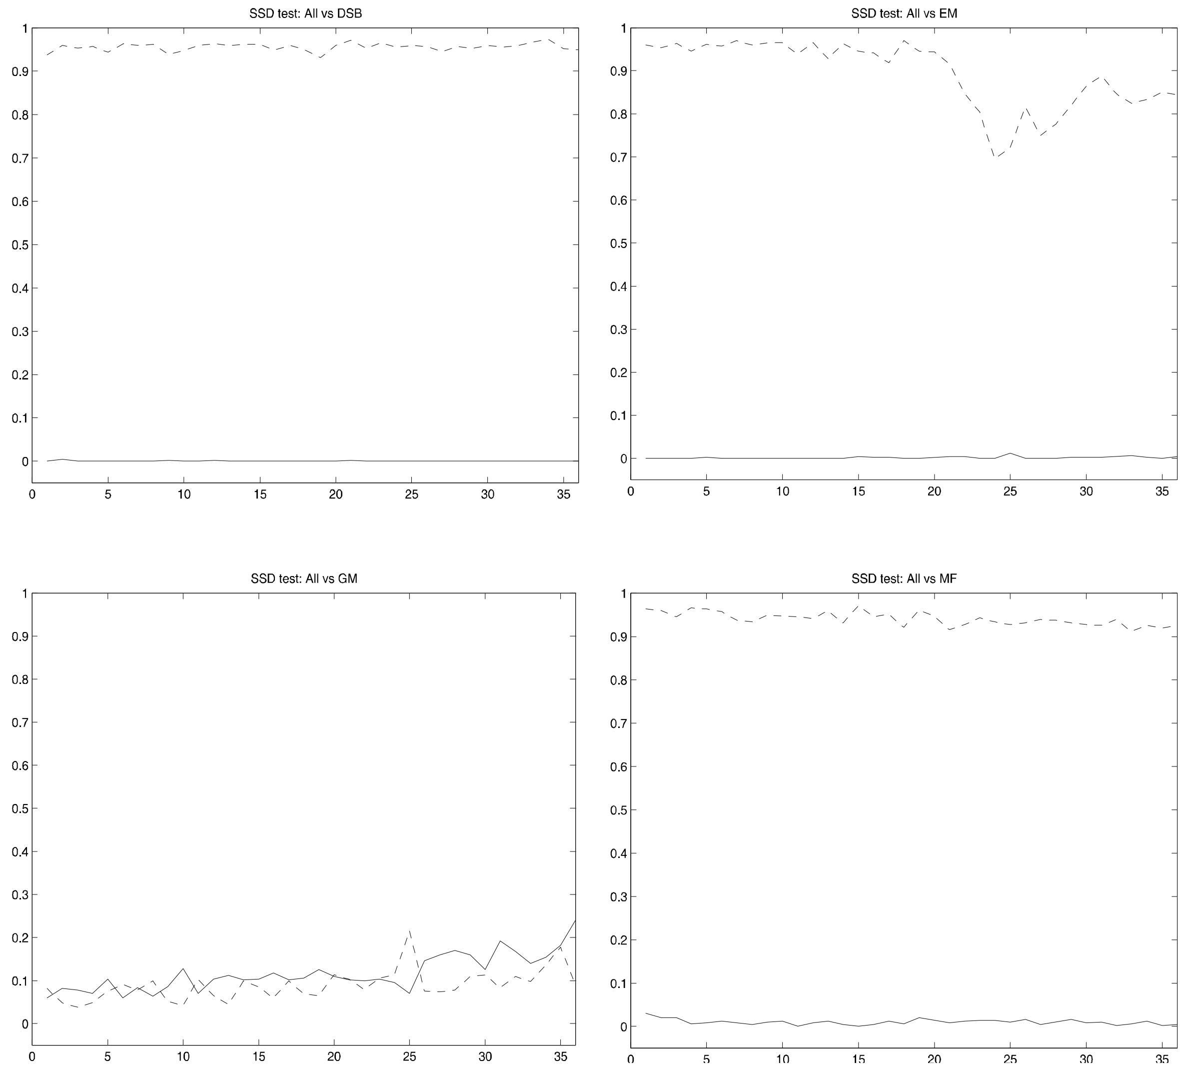

Figure 16: Dynamic p-values of predictive stochastic dominance of order 2. Out-of-sample comparison of the hedge fund industry (ALL) with 'Directional Traders' investment styles, 2004:01-2006:12. The dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of ALL over the individual styles. The solid line represents the p-value of the test defined by the converse null hypothesis.

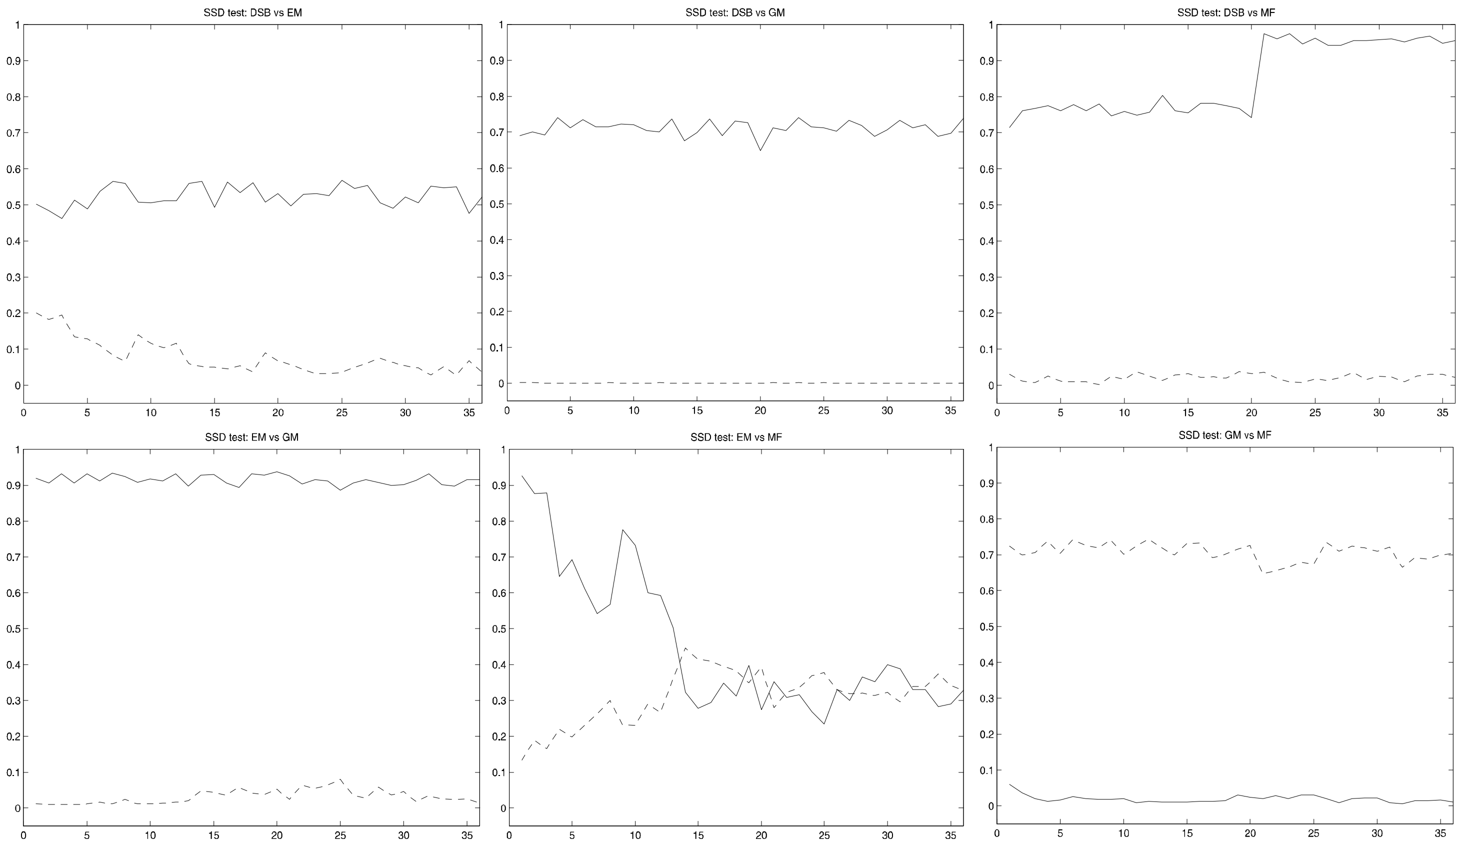

Figure 17: Dynamic p-values of predictive stochastic dominance of order 2. Out-of-sample comparison of ‘Directional Traders’ hedge fund investment styles, 2004:01-2006:12. For style A vs. style B the dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of style A over style B. The solid line represents the p-value of the test defined by the converse null hypothesis.

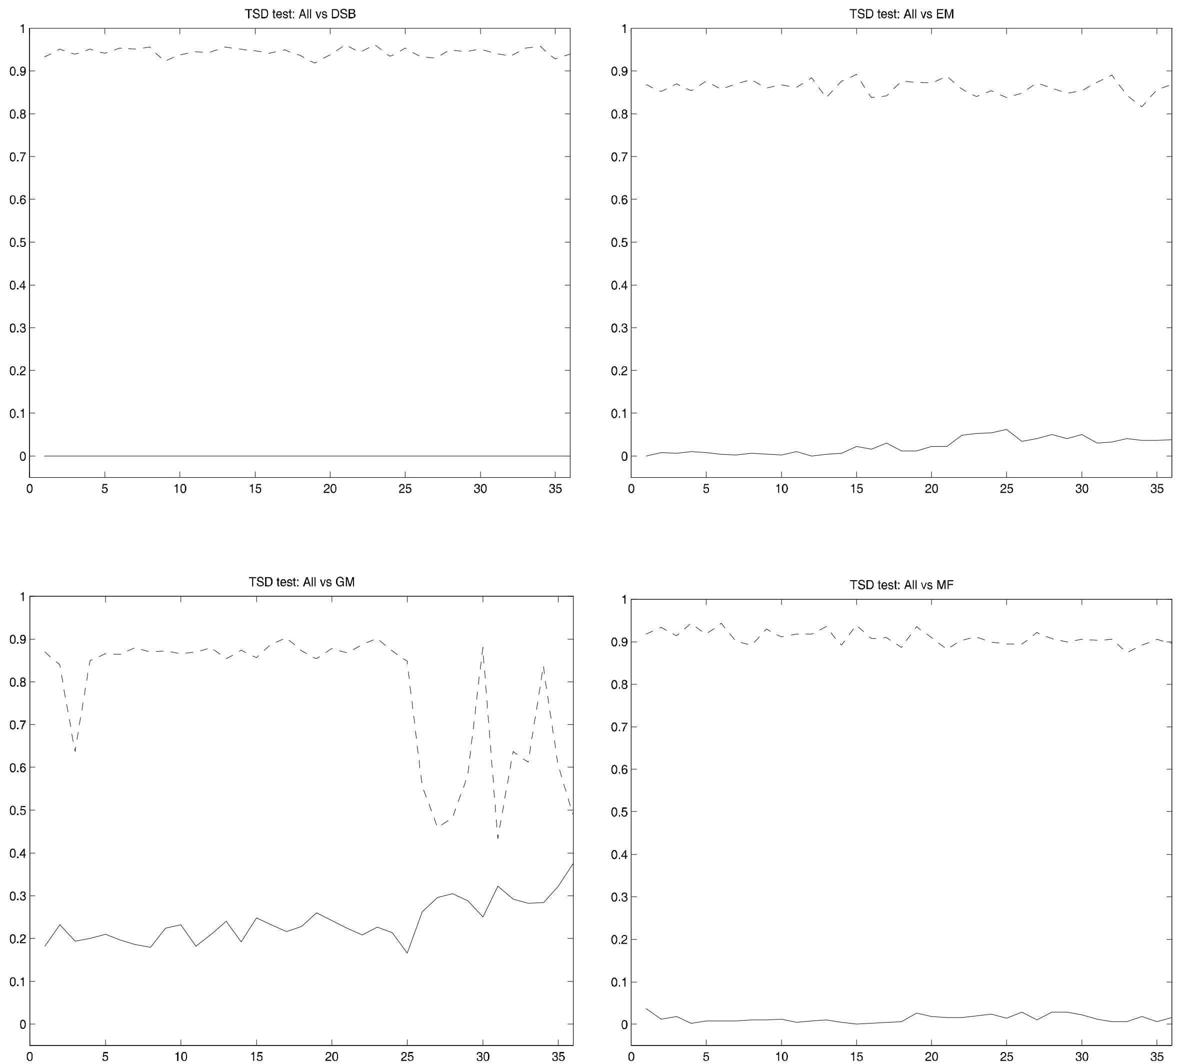

Figure 18: Dynamic p-values of predictive stochastic dominance of order 3. Out-of-sample comparison of the hedge fund industry (ALL) with 'Directional Traders' investment styles, 2004:01-2006:12. The dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of ALL over the individual styles. The solid line represents the p-value of the test defined by the converse null hypothesis.

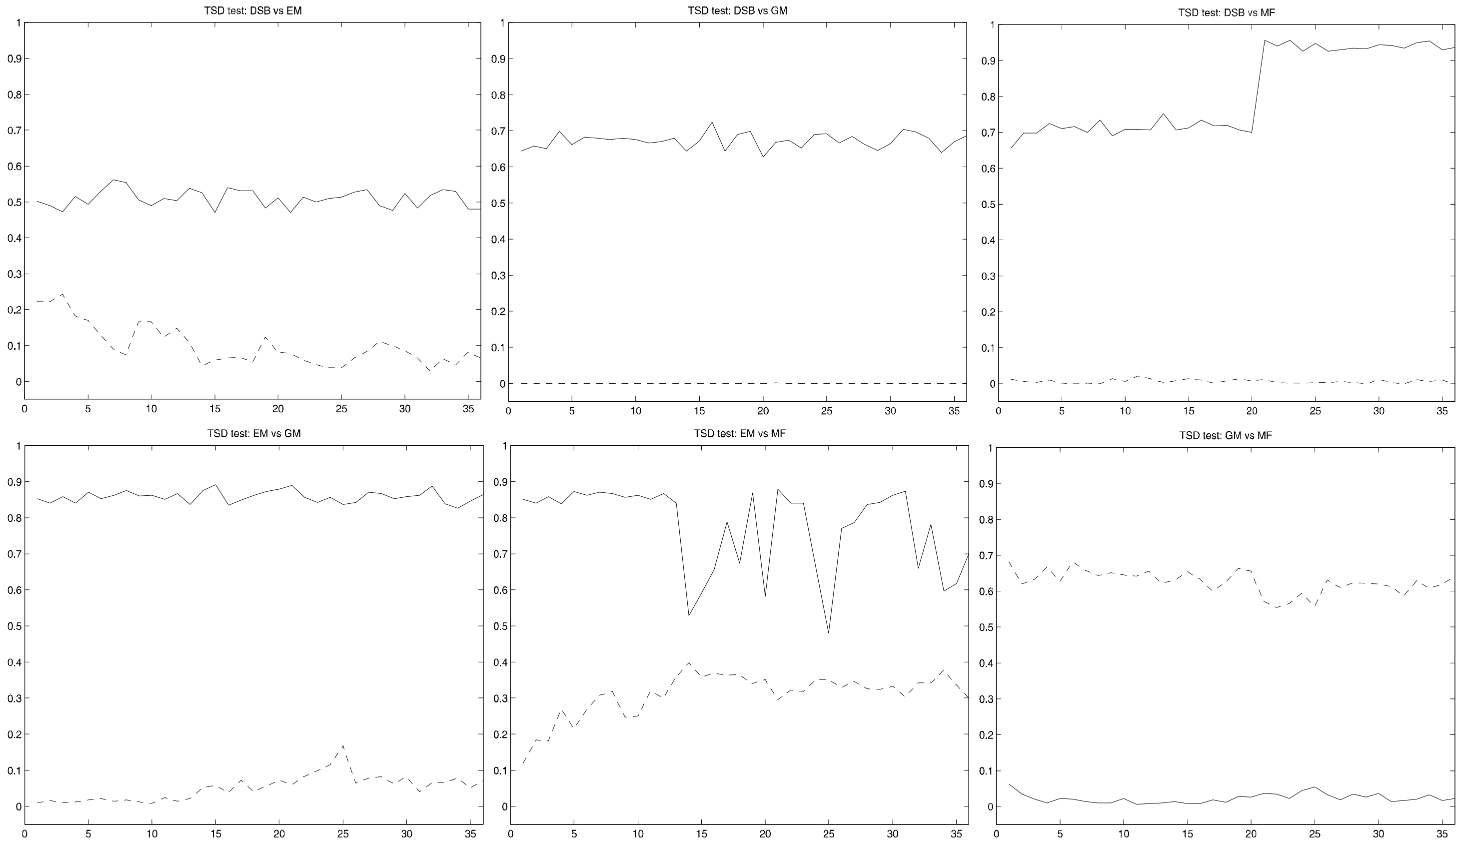

Figure 1: Dynamic p-values of predictive stochastic dominance of order 3. Out-of-sample comparison of 'Directional Traders' hedge fund investment styles, 2004:01-2006:12. For style A vs. style B the dashed line reflects the p-value of the test whose null hypothesis is given by the stochastic dominance of style A over style B. The solid line represents the p-value of the test defined by the converse null hypothesis.